

Essays on Information Technology, Intangible Capital, and the Economics of Artificial Intelligence

By

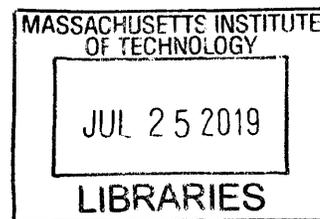
Daniel Rock

B.S. Economics

University of Pennsylvania, Wharton School, 2010

S.M. Management Research

Massachusetts Institute of Technology, Sloan School of Management, 2016



SUBMITTED TO THE SLOAN SCHOOL OF MANAGEMENT IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN MANAGEMENT

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2019

©2019 Massachusetts Institute of Technology. All rights reserved.

Signature redacted

Signature of Author: _____

Department of Management
April 25, 2019

Signature redacted

Certified by: _____

Erik Brynjolfsson
Schussel Family Professor of Management Science
Professor of Information Technology
Thesis Supervisor

Signature redacted

Accepted by: _____

Ezra Zuckerman Sivan
Alvin J. Siteman (1948) Professor of Entrepreneurship and Strategy
Deputy Dean
Professor, Technological Innovation, Entrepreneurship, and Strategic Management and Work
and Organization Studies

Essays on Information Technology, Intangible Capital, and the Economics of Artificial Intelligence

By

Daniel Rock

Submitted to the Department of Management on April 25, 2019 in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Management

ABSTRACT

This dissertation contains four essays concerning the economics of information technology, intangible capital, and artificial intelligence. In the first essay, “Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence” I describe how firms can appropriate some of the value of their employees’ human capital by assigning firm-specific tasks. I then use a database of employment records to document dynamics in the valuation of publicly traded firms as they relate to different types of employment, focusing especially on AI skills.

The second essay, “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies” (coauthored with Erik Brynjolfsson and Chad Syverson) addresses the concern that new technologies with wide applicability throughout the economy can cause both underestimation and overestimation of total factor productivity. As capital is accumulated, intangible investment output, and therefore productivity growth, will be underestimated only to later generate a yield (at which point productivity growth will be overestimated). Presenting a theoretical description of how to use corporate valuations to recover hidden investment value, we discuss how productivity growth and levels can be adjusted to accommodate these changes. Implications for research and development, computer hardware, and computer software investments are considered.

The third essay, “Machine Learning and Occupational Change” (coauthored with Erik Brynjolfsson and Tom Mitchell), develops and implements a method to measure the labor market impact potential of machine learning technologies. Tasks are evaluated for their Suitability for Machine Learning (SML). We find that few occupations can be fully automated with machine learning, but many occupations will potentially be redesigned.

The final essay, “Do Labor Demand Shifts Occur Within Firms or Across Them? Non-Routine-Biased Technological Change 2000-2016” (coauthored with Seth Benzell and Guillermo Lagarda) decomposes labor share shifts of occupational groups into changes between firms, within firms, and due to entry and exit. We find that within-firm compositional shifts are an important component of changes in the overall labor market. We also find that the rate of within-firm shifts has declined in the period from 2000 to 2016. Together, these essays offer insights into how artificial intelligence technologies, particularly machine learning, will impact the U.S. economy.

Thesis Supervisor: Erik Brynjolfsson

Title: Schussel Family Professor of Management Science, MIT Sloan School of Management

Acknowledgements

MIT has been a special place to learn and work these past few years. The student experience at MIT (and at Sloan in particular) is challenging and immensely fulfilling because of the wonderful people who come together to make this institution what it is: open, curious, and humble. I have benefitted greatly from a large community both within and outside of MIT whom I've gotten to know as peers, mentors, collaborators, and lifelong friends.

I could not have asked for a better advisor and mentor than Erik Brynjolfsson. From our first conversation discussing whether an automated laser barber would be economically feasible, to studying the impact of machine learning across the economy in our research, it has been wonderful learning the art (and science) of scholarship from him. He has an exceptionally rare combination of intelligence, integrity, vision, and kindness. It is an amazing privilege to be his student.

I am exceedingly grateful for the tireless support of my dissertation committee. Prasanna Tambe has made me feel welcome as part of the research community, and always has terrific suggestions and exceptional guidance for research (and life in general). Working with Andrew Lo in the past few years always inspired me to aim for big questions, to cross boundaries between fields, and not to abandon the small hours of the morning. Chad Syverson helped me to relate technical details into bigger economic stories. His wisdom and scholarship have forever shaped my perspective on research.

Beyond my committee, I am thankful to the academic community at MIT and beyond. Sinan Aral has been a tremendous influence, introducing me to a broad set of causal inference tools and teaching me what makes a question worth pursuing. Wanda Orlikowski opened my mind to new approaches and frameworks for understanding. Tom Mitchell explained the “black box” of machine learning for me. Neil Thompson showed me how to build and manage large, messy projects and that observational causal inference is “for procrastinators”. My undergraduate advisors, Lorin Hitt and Eric Clemons, encouraged me to pursue this path, providing guidance, insight, and beer along the way. Lynn Wu and Adam Saunders have been amazing “academic older siblings”, aiding with research ideas and answering my endless stream of questions. Meetings with Marshall Van Alstyne always left me excited to understand information’s impact on the world. Bo Cowgill, Dean Eckles, Jeff Furman, John Horton, Kristina McElheran, Frank Nagle, Tavneet Suri, Jon Hersh, and Sean Taylor were always generous with their time and thoughts. Morgan Frank, Guillermo Lagarda, Iyad Rahwan, Izabela Witoszko, and Daniela Rus are fantastic collaborators, and I look forward to our future work on the work of the future. I’m thankful to have the privilege of working with the terrific LinkedIn EGR team, especially Jacqueline Barrett, Nadiya Hayes, Guy Berger, Di Mo, and Robert Stucke. Thank you to Bill Dague – a great friend and constant source of inspiration.

The Initiative on the Digital Economy team was an integral part of my Ph.D. experience. Andy McAfee’s thoughtfulness in analyzing the economy and consistency in doing awesome things are unrivaled. George Westerman gave me ideas to pursue and motivation to engage other disciplines. David Verrill and Christie Ko have built an incredible platform for research of a large group to be successful and integrated. Working with the A-Lab Dream Team of Susan Young and Chuck Gibson was fantastic. Tammy Buzzell, Joanne Batziotegos, Jovi Koene, Carrie Reynolds, and Naomi Stephen have been a pleasure to work with and central to getting projects to run smoothly. I couldn’t be more excited to spend another year with IDE.

I’m indebted to the community of students and post-docs who are now lifelong friends and colleagues. My office mates and fellow PhD students made for a fun and (occasionally) productive

work environment. Thanks to the IT team: Matt Beane, Avi Collis, Dave Holtz, Shan Huang, Arvind Karunakaran, Alex Moehring, Zanele Munyikwa, Sebastian Steffen, and Michael Zhao for engaging conversations and your brilliance in understanding technology. Thanks to IDE post-docs Sagit Bar-Gill, Seth Benzell, Andrey Fradkin, Hossein Ghasemkhani, Xiang Hui, Wang Jin, Meng Liu, Frank MacCrory, Shachar Reichman, Daniela Scur, Ananya Sen, Xiupeng Wang, and Erina Ytsma for lighting the way forward. A special thanks goes to my frère Guillaume Saint-Jacques – “Brother-in-Arms” the whole time through. Thank you to Dan Fehder and Joshua Krieger for tremendous kindness and intellect, and to the 2013 TIES cohort, Ankur Chavda, Danny Kim, and Samantha Zyontz, for welcoming me into the group. Thanks also to Linda Zhao, Mine Kansu, Margo Blank, Maxime Cohen, Manuela Collis, Valere Fourel, Erik Duhaime, Brittany Bond, Carolyn Fu, Wes Greenblatt, Bill Goulding, Heidi Packard, Sarah Bana, Rebecca Grunberg, James Riley, Heather Yang, Madhav Kumar, Minjae Kim, Yonadav Shavit, Becky Karp, and Abhishek Nagaraj for your warmth and friendship. And a huge thank you to Hillary Ross, Davin Schnappauf, and Sarah Massey for help throughout the process.

I have been incredibly fortunate to have the support of wonderful friends and family. I will forever be grateful to Kevin, Emily H., Ted, Emily H., Daniella, Max, Zoey, Meeran, Julia, Yiyi, Marissa, Emily B., Roman, Rob, Jonathan, Arjun, Lando, Matt, Grace, Alanna, Talia, Colleen, Emily W., Lisa, and Becca (my thesis role models), Sam, Inna, Hayley, Kim, Keiko, Maxine, JT, Lawrence, and Rahul. A special thank you to my friend and collaborator Nicholas Fazzari – I have learned so much from you. My siblings Jordan, Elliott, Joey, Rachel, Amy, and Roarke were always there for laughs, encouragement, and counsel. Thank you to my grandparents Doreen and Phil for making me excited about scholarship and always pushing me to do my best. Thank you to my parents, Jean and Michael, and in-laws Karen and Dan for unwavering support and for letting me remain a kid at heart. Above all I owe an enormous debt of gratitude to my wife, Laura. It is the single most significant honor of my life to be your husband and partner. As is the case for all my endeavors, this work is dedicated to you.¹

¹ 8.4.1.1-10.9.2.12-11.1.1.10-12.8.1.2-12.17.3.1-16.6.1.1-17.5.4.5-72.3.2.10-72.8.3.4-86.6.2.1-88.20.1.1-89.22.15.4-92.6.1.6-92.12.3.4-92.6.1.5-92.12.3.4-95.18.16.2-96.1.8.1-129.3.11.5-130.4.5.1-140.3.1.2-141.9.2.5-145.5.1.1-150.9.12.4-155.6.10.2-169.3.3.3-173.3.1.3-178.2.7.1-180.22.17.1-184.6.12.2-//.//.//.-2:

Table of Contents

<i>Acknowledgements</i>	4
<i>Chapter 1 — Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence</i>	8
1. Introduction.....	8
2. Related Research on Human Capital, Technology, and Market Value.....	13
3. Theoretical Framework for Valuing the Firm’s Share of Human Capital Investment.....	19
4. Dataset Construction	29
5. Overall Engineering: Empirical Results and Discussion	40
6. TensorFlow, the Deep Learning Toolkit, and Bottlenecked Talent.....	55
7. Further Robustness Checks.....	64
8. Suitability for Machine Learning Analysis.....	73
9. Conclusion	77
References	79
Appendix: Additional Regression Results on Coverage and Robustness	85
<i>Chapter 2 — The Productivity J-Curve: How Intangibles Complement General Purpose Technologies</i>	92
I. Introduction	92
2. Technology, Investment Theory, and Productivity Growth	95
3. Growth Accounting in the Presence of Unmeasured Intangible Investment.....	98
4. The Productivity J-Curve.....	105
5. Is Hidden AI Capital Investment Already Causing a Productivity Shortfall?	109
6. Deploying the Framework Using R&D, Software, and Computer Hardware Investment.	111
7. Can Intangible Capital Outputs Explain the Productivity Slowdown?.....	123
8. Conclusion	125
References	127
<i>Chapter 3 — Machine Learning and Occupational Change</i>	133
1. Introduction.....	133
2. Automation, Artificial Intelligence, and Labor Demand.....	137

3. Why might this time be different?	140
4. A Model of Machine Learning Automation	141
5. Measuring Suitability for Machine Learning at the Level of Tasks	146
6. Measurability.....	165
7. Conclusion	166
References	167
Appendix A – Machine Learning Paradigms	171
<i>Supervised Learning</i> —	171
Appendix B – The Suitability for Machine Learning Rubric.....	172
Appendix C – A Framework for Predicting the Impact of New Technologies on Labor Demand and Wages.....	177
<i>Chapter 4 — Do Labor Demand Shifts Occur Within Firms or Across Them? Non-Routine Biased Technological Change, 2000-2016</i>	181
1. Introduction.....	181
2. Related Literature and Framework	185
3. Framework	186
4. Total Factor Productivity Changes Across Firms	188
5. Skill-Biased Technological Changes Within Firms	191
6. Data and Methodology	192
7. Net Job Creation Within and Across Firms	194
8. Skill-Biased Technological Change Within and Across Firms	197
9. Firm Employment Mix Dynamism.....	200
10. Skill-Biased Technological Change Within and Across Firms by Industry	204
11. SBTC and Firm Investment.....	207
12. Conclusion	213
References	215

Chapter 1 — Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence

“I was originally supposed to become an engineer but the thought of having to expend my creative energy on things that make practical everyday life even more refined, with a loathsome capital gain as the goal, was unbearable to me.”

– Albert Einstein

1. Introduction

Technological labor is a well-established driver of corporate market value, innovation, and productivity (Hall 1993, 2006; Tambe and Hitt 2012; Tambe 2014). 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, the capital gains in applying highly specialized knowledge can be partially bargained away by employers in markets with competitive labor supply. There is enduring disagreement about how and to what extent firms can appropriate the human capital investments of their workers (Acemoglu and Pischke 1998; Acemoglu and Pischke 1999; Brynjolfsson et al. 2018). There are many ways in which the returns to worker investments in human capital might lead to employer value gains. Engineers, as implementers of technology, are highly complementary to the intangible knowledge assets that firms accumulate.

Firms hire engineers to build. The specifics of what technical workers build are subject to the discretion of their employers (presumably with capital gain as the goal). When workers prefer some tasks over others for non-pecuniary reasons, the firm gains monopsony power. In some cases,

technological workers might sacrifice compensation for other perks or allowances (e.g. the right to publish scientific findings or working with cutting-edge technology) (Stern 2004; Mas and Pallais 2017). In labor market contexts where both supply and demand are competitive, there is little excess surplus to split. Each party earns the marginal product of what it provides. For employees with accumulated firm-specific knowledge and little competition (CEOs and high-level executives, for example), some of the surplus from firm-specific assets can be bargained away by the worker (Brynjolfsson 1994; Hart and Moore 1994). Some kinds of occupations, like managers and tech workers, have a mixture of both kinds of tasks. They do activities that require firm-specific capital that grow productivity and they also implement production processes in tasks that are more competitive.

This paper seeks to address whether technical talent is a source of rents for corporate employers. The answer to this question informs whether strategic technological labor resources are a direct source of sustainable competitive advantage or operate via other factors (Barney 1986; Crook et al. 2011). I approach this question in two ways: firstly by investigating the relationship between market value and aggregated measures of technological talent, and secondly by exploiting the unexpected open-sourcing of Google's TensorFlow, a machine learning software package particularly well-suited for deep learning, at the end of 2015. The overall engineering value estimation aggregates across many different technologies, whereas the TensorFlow analysis illustrates a case where engineering talent is highly scarce. As an emergent technology, I find that the marginal value net of wages of additional AI talent is still well above the breakeven point for the average publicly traded firm.

If technological human capital is highly complementary to the firm's asset base and both cause market value, for example, then failure to invest in worker retention may impair the value of the firm's non-human capital. Technological human capital should cause market value in the case that tech workers are employed building firm-specific assets (even if those assets remain in the workers' heads).

In the first section of the empirical results, I use a firm-level panel of employment by worker type from LinkedIn merged to firm performance and value data to describe the component of market value attributable to technical talent. I estimate the causal effect of additional engineering investment using proximity to land-grant colleges (Moretti 2004; Bloom et al. 2018), changes in state-level covenant-to-not-compete (CNC) policy (Marx, Strumsky, and Fleming 2009; Marx 2011; Ewens and Marx 2017; Starr, Balasubramanian, and Sakakibara 2017; Balasubramanian et al. 2018; Jeffers 2017), and firm and state-level differences in the user cost of R&D capital (Lucking, Bloom, and Van Reenen 2017; Lucking 2018) as instruments for the engineering human capital in firms. I find that while engineering talent expenses and quantities are strongly correlated with market value the causal specifications using instrument variables and correlational estimates including firm fixed effects eliminate statistical significance. Point estimates for the value of engineers increase in specifications using land-grant and CNC instruments, though they are imprecise. Model specifications with the user cost of R&D instruments, to the contrary, indicate that the marginal engineer destroys market value adjusting for the total stock of R&D assets at the firm. This is indicative of the presence of firm-level intangible asset service flows complementing the labor of technological workers, consistent with Brynjolfsson, Hitt, and Yang (2002); Tambe, Hitt, and Brynjolfsson (2012).

Taken holistically, these results suggest that the causal effect of the marginal technological hire seems to add little to the firm's market value, but the average value of these workers is high in equilibrium. I find that each additional engineering worker is correlated with another \$855,000 of market value for the firm, and in wage terms \$1 of engineering wages is correlated with approximately \$11.9 of market value. This suggests that the relationship between market value and technological talent is generated by firm-specific assets that are complementary to or perhaps embodied within generalized engineering labor. The lack of precision on the marginal effect of general engineering indicates the need for more skill-specific studies of labor value. Even aggregate measures of STEM employment fail to suggest easily earned rents.

In the current market environment, there are few, if any, technologies with the transformative potential of artificial intelligence and machine learning (Brynjolfsson, Rock, and Syverson 2017; Brynjolfsson, Mitchell, and Rock 2018b; Cockburn, Henderson, and Stern 2018; Agrawal, Gans, and Goldfarb 2018). One of the primary obstacles to widespread adoption of artificial intelligence is the available labor supply, with top-tier scientists earning more than \$1 million in some cases.² But the market is responding and the supply of machine learning engineers is increasing. Accordingly, to further understand the mechanism generating the large average correlational value of engineering talent and how it might be related to specific labor skills, the second part of the paper exploits the Google's open-source launch of TensorFlow in November 2015. TensorFlow has a Python-based application programming interface (API) which greatly facilitates the ease and efficiency in building (and learning to build) deep learning models. The TensorFlow launch serves as a shock to the fixed costs of learning how to build deep neural nets (DNNs) for software engineers and analysts. Prior to TensorFlow, the ability to train DNNs was rare and highly specialized. The launch of this tool both effectively commodified deep learning as a skill amongst those with Python ability and accelerated *expectations* for how soon deep learning would be easy to learn more generally.

Following the introduction of TensorFlow, I find a rapid increase in the rate and quantity of addition of Artificial Intelligence skills on LinkedIn. Mapping these increases to publicly-traded firms, I find that the TensorFlow shock had differential effects on firm 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 \$3.56 million following the introduction of TensorFlow. The TensorFlow launch provides evidence that talent scarcity can be an important

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

bottleneck to the realization of returns on technological assets. Lowering the barriers to acquiring a formerly rare and valuable skill, as TensorFlow does, makes technological supply competitive. This increased competition of technology suppliers (i.e. engineers) enables their employers to earn returns on firm-specific assets. Additionally, I test whether it is firm-level opportunities to *apply* machine learning causing the increase in market value using the Suitability-for-Machine Learning (SML) measures in (Brynjolfsson, Mitchell, and Rock 2018a, 2018b). If anything, higher average firm SML scores are *negatively* correlated with market value. While it would be premature to assume that the stock market has fully priced in the automation potential of machine learning, this difference-in-differences result suggests AI-related repricing of corporate assets in 2016.

The set of mechanisms by which technology workers might generate market value is generally applicable to all kinds of human capital. However, technological skills can change or depreciate much faster than other kinds of human capital. 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. TensorFlow's introduction is one such example among many. Technological shifts therefore supply outside researchers with a chance to study the outcomes of employment-related relational contracts. Analogous shifts to TensorFlow for managerial workers, for example, might be more challenging to find. Still, studying technological changes can supply insight into how companies and employees divide the gains from business activity. These conclusions, in some cases, can be applied to other kinds of employees. 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.

The paper is organized as follows: Section 2 describes the relevant literature in the economics and strategy of human capital, technology, and market value. Section 3 describes a theoretical model

of how human capital can enter the valuation of firms. Section 4 details the construction of the datasets. Section 5 describes and analyzes the relationship between market value and aggregated engineering talent. Section 6 offers an empirical case study of the TensorFlow launch and discusses the market value effects of making AI talent more abundant. Section 7 concludes.

2. Related Research on Human Capital, Technology, and Market Value

This study fits within a tradition of human capital and technology studies that can be traced back to (Becker 1962). How human capital is accumulated and why firms are motivated to invest in it has long been a puzzle for social scientists. A key question is how easy it is for workers to apply their human capital across different employers. Transferrable skills and knowledge (general human capital) are subject to competitive bidding pressure from firms, while the returns to firm-specific investments are 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 1997; Acemoglu and Pischke 1998; Acemoglu and Pischke 1999). A related literature considers the market power aspect of these frictions, studying monopsony power (Bhaskar, Manning, and To 2002; Ashenfelter, Farber, and Ransom 2010). Firms insulate themselves from competitive pressure in myriad ways, including (but not limited to) regional concentration (Azar et al. 2018), organizational design and technology (Stole and Zwiebel 1996b), and incomplete contracts (Stole and Zwiebel 1996a; Simon 1991).

Match-specific value between employer and employee also leads to productivity-ability sorting, at which point larger surplus values are split between elite matches. Tervio (2008) offers an assignment model approach to understanding the value of CEOs that is instructive for the skilled worker context considered here. If part of the work activities bundled into a given technical job are investments in match-specific value and the remaining portion is assigned to competitive labor tasks, the firm pays wages at the competitive labor margin while appropriating the bundled match-specific value. This will not be the case, as in Tervio (2008), when much of the labor effort is devoted to

match-specific tasks or, as in traditional neo-classical labor supply functions, the labor effort is competitive and commodified.

The context considered here is in-between those extremes: engineers, much like managers, spend part of their time building non-marketable firm-specific assets and part of their time maintaining or implementing production in competitive arenas. This generates an incentive for the firm to use the allocation of tasks (job design) as an instrument of monopsony power. As in Simon (1991), “the combination of uncertainty on the part of the employer (as to what will need to be done in the future) and broad acceptance of the employee (of what he or she will be ordered to do) makes the employment contract a very attractive bargain for both parties”. Of course, firms can modify how workers *perceive* the firm-specificity of their human capital investments, wherein workers might be more willing to learn firm-specific skills if they believe them to be marketable (and conversely, less willing if the skills were perceived as unmarketable) (Coff and Raffiee 2015; Raffiee and Coff 2016).

This paper makes a simplifying assumption that workers and firms correctly understand the firm-specificity of their human capital, but that uncertainty at the time of contracting leads to ex-post rent sharing. Campbell, Coff, and Kryscynski (2012) argue for mobility constraints as a stronger influence than firm-specificity of human capital. Supply-side factors affecting worker bargaining power and the propensity for workers with firm-specific assets to start new firms (as in (Campbell et al. 2012; Eisefeldt and Papanikolaou 2014; Jovanovic 1979)) are an important component of the firm’s share of created value. Additionally there is evidence that “scientists pay to be scientists” (Stern 2004) or, in other words, have preferences favoring academic rewards over monetary ones (Roach and Sauermann 2010). There is evidence this is also the case for technologists, who frequently prefer to work with cutting-edge technology (Tambe, Ye, and Cappelli 2019). More traditionally, monopsony power can be related to employer concentration (Azar et al. 2018; Benmelech, Bergman, and Kim 2018) and financial constraints also affect hiring decisions (Benmelech, Bergman, and Seru 2011).

The mechanism for engineer value in this paper is the interaction of task allocation and firm-specific assets with engineering labor complements.

Fixed costs of capital investment 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 assets in these studies 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) (Hayashi 1982; Tobin et al. 1976; Kaldor 1966).³ 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. These adjustment values for "vanilla" labor have been estimated at low values in the past (Hall 2004; Hall 2017).

Nevertheless, the firm's value share of some types of human capital has been estimated as large and meaningful. Specialized labor is highly sought-after by corporate employers. The H-1B visa program, which expands the talent pool for technically-savvy workers, is typically oversubscribed such that eligible talent from outside the U.S. must file for a lottery. This increased quantity of STEM workers from the H-1B program led to greater productivity (Peri, Shih, and Sparber 2015), within-firm employment (Kerr, Kerr, and Lincoln 2015), and rates of innovation and entrepreneurship (Kerr 2013). Technological talent is deployed not only in implementing the production function, but also in building the knowledge and business process assets of the firm which facilitate growth. 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.

³ This leads to the main market value equation I apply in the empirical section.

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). The estimated average Q-value of R&D assets is nearly an order of magnitude above that of property, plant, and equipment in Compustat firms (Peters and Taylor 2017; Brynjolfsson, Rock, and Syverson 2018). This is suggestive that R&D-intensive firms tend to also accumulate hidden intangible complements – business processes, training, knowledge, and even firm culture – which contribute to market value in ways that are difficult to capitalize on a balance sheet. These factors, even with relatively coarse measures of organizational investment, have predictive power in the cross-section of stock returns (Eisfeldt and Papanikolaou 2013).

Intangible assets are complementary to and correlated with investment in technological human capital (Bresnahan, Brynjolfsson, and Hitt 2002; Brynjolfsson, Hitt, and Yang 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 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 (Syverson 2011; Lustig, Syverson, and Van Nieuwerburgh 2011; Andrews, Criscuolo, and Gal 2015). Intangible assets are inherently hard to measure and constitute an increasingly large component of the U.S. economy's asset stock. One explanation for the high market value of engineers in the empirical results of this paper is that firms with more engineers also tend to build up intangible assets which are left off of the corporate balance sheet. Firms often fail to capitalize software expenses, for example, causing correlational analyses to attribute market value to the observable complements (wage value, in this case). This quantity is the focal object of study in (Tambe et al. 2018). This study begins at the aggregated level of technological labor, and then studies the labor shock of TensorFlow in AI as a means of understanding specifically how scarce labor can serve as a bottleneck to firm value creation.

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. Thompson (2017) studies the economic effects on firm productivity of the switch to multicore processing. Ewens, Nanda, and Rhodes-Kropf (2018) analyze the entrepreneurial effects of the launch of Amazon Web Services (AWS). AWS bundled a number of general-purpose technologies together and made computing infrastructure rentable.⁴ This was a major reduction in the fixed costs of starting a new technology-oriented business. Teodoridis (2017) shows that the hack of the Microsoft Kinect made motion-sensing much cheaper, reducing the need for research teams to collaborate with specialists in motion-sensing technology. The hack democratized the technology, similar to the way in which the launch of TensorFlow has (partially) democratized deep learning. Technological advances need not have such an effect; burden-of-knowledge (Jones 2009) effects might dominate in the case that there is an exogenous increase in knowledge capital that is costly to process (Agrawal, Goldfarb, and Teodoridis 2016). This kind of change might necessitate the use of subject matter specialists in increased proportions. Most recently, (Zyontz 2018) has studied the team expertise structure of cell biology researchers following the advent of CRISPR, a precise gene editing tool.

These studies of how production changes in response to inventive activity have in common a theoretical underpinning in the value of “discovery information” (Hirshleifer 1978). Discovery information refers to “detection of properties of Nature that permit the development of new tools or the utilization of new techniques”. In Hirshleifer’s example, the prices of state-contingent claims adjust when the knowledge of the new state probabilities (e.g. the probability that deep learning will be made easier) will be obtained publicly before the close of trading. As with Eli Whitney’s cotton gin, a “route to profit” other than patent protection for new intellectual property is in speculation on the business prospects of firms technologically exposed industries. In AI, the pecuniary benefits of

⁴ Though not in 2006, AWS now includes AI services.

open-source innovation, by revealed preference, outstripped the benefits of private IP for Google in TensorFlow's case. AI shares characteristics of both types (tool and knowledge) of exogenous changes, though one of Google's stated aims in open-sourcing TensorFlow was to increase usability and accessibility of deep learning for engineers throughout the economy.⁵ A primary pecuniary benefit of making AI models easier to build is an expected subsequent drop in marginal wage rates for AI-intensive human capital.

Like ICT, Artificial Intelligence-related assets are mostly intangible and the returns are mostly in the future at this point (Brynjolfsson, Rock, and Syverson 2017). The recent progress in AI is mostly 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, Hinton, and Williams 1986; LeCun et al. 1998; LeCun, Bengio, and Hinton 2015). AI is an umbrella discipline, including machine learning (which itself includes deep learning) as well as rules-based or expert systems approaches to problem solving. As a new kind of software, however, deep learning and AI more broadly is a general purpose technology (Bresnahan and Trajtenberg 1995; Bresnahan 2010). It is potentially pervasive, improves over time as better and more data arrive, and can spawn complementary innovation. AI is a general purpose prediction technology (Agrawal, Gans, and Goldfarb 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.

Since prediction is pervasive throughout the economy, the promise of AI is that it will lead to business process innovation, job redesign, and new engineering advances across many domains in the

⁵ <https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/>

economy (Furman and Seamans 2018; Brynjolfsson, Mitchell, and Rock 2018b). Critically, deep learning overcomes the obstacle of Polanyi's Paradox where "we know more than we can tell" (Polanyi 1966; Autor 2014). For deep learning models, we need only measure inputs and outputs. The map between them is learned by the algorithm. Artificial General Intelligence (AGI), where machine intelligence equals or surpasses human intelligence in all cognitive tasks, is technologically far away at the moment. But the relatively brittle, bespoke applications of deep learning could feasibly cause large shifts in labor demand and economic value creation processes (Brynjolfsson, Hui, and Liu 2018).⁶

Since the effects 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. Further, as noted by (Raj and Seamans 2018), relatively little data on AI at the firm-level is available. This study builds and measures firm-level proxies for AI investment as of the end of 2017. As a GPT, the effects of machine learning on firms and labor markets will likely be diffuse across many industries (Brynjolfsson, Mitchell, and Rock 2018a; Felten, Raj, and Seamans 2018). Like prior waves of automation, machine learning will differentially impact tasks that are technologically and socially feasible (Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011; Autor and Dorn 2013). Yet for engineering value, we can study now how the decision to make an advanced tool widely available lowered workers' entry costs and facilitated a shift in technological investment.

3. Theoretical Framework for Valuing the Firm's Share of Human Capital Investment

Standard investment theory requires little modification for human capital to enter into market valuations of firms (Tervio 2008). Here I follow a setup common to (Lucas 1967; Hayashi 1982;

⁶ Applications so far include: Self-driving cars, call center automation, insurance claims processing, materials discovery, drug discovery, and language translation

Wildasin 1984; Hayashi and Inoue 1991; Yang and Brynjolfsson 2001; Brynjolfsson, Rock, and Syverson 2018). The firm must choose the right investment and labor quantities to maximize profits:

$$\max_{I,L} V(0) = \int_0^{\infty} \pi(t)u(t)dt; \quad \pi(t) = pf(K, I, L, t) - w^t L - z^t I - r^t K \quad (1)$$

Profits are denoted by $\pi(t)$; $u(t)$ denotes the compound discount factor and p is the price of output. F is assumed nondecreasing and concave in capital (K) and labor (L), and nonincreasing and convex in investment I . W and Z refer to the price vectors at time t for wages and investment, respectively (superscript t is dropped from now on). Assume F is homogeneous in the first degree. We have the following growth constraint on capital stocks of varieties indexed by j :

$$\frac{dK_j}{dt} = I_j - \delta_j K_j \quad \forall j = 1, 2, \dots, J$$

Then the firm's Hamiltonian to maximize is:

$$H(K, I, L, t) = (pf(K, I, L, t) - wL - zI)u(t) + \sum_{j=1}^J \lambda_j (I_j - \delta_j K_j) \quad (2)$$

With the following constraints⁷:

$$\frac{\partial H}{\partial \lambda_j} = \dot{K}_j = I_j - \delta_j K_j \quad \forall j \text{ and } \forall t \in [0, \infty]$$

$$\frac{\partial H}{\partial K_j} = -\dot{\lambda}_j = pF_k u - \lambda_j \delta_j \quad \forall j, t$$

$$\frac{\partial H}{\partial I_j} = 0 = (pF_{I_j} - z_j)u + \lambda_j \quad \forall j, t$$

⁷ $\dot{x} \equiv \frac{dx(t)}{dt} \quad \forall x(t)$

$$\frac{\partial H}{\partial L_i} = 0 = (pF_{L_i} - w_i)u \quad \forall i = 1, 2, 3, \dots, L \text{ and } \forall t$$

$$\lim_{t \rightarrow \infty} \lambda(t)K(t) = 0$$

This implies that the firm's market value is the sum of the quantities of its capital assets multiplied by their replacement value prices added to the installed value (where λ is the installed asset price). The solution for the firm's market value under the condition that marginal and average wages are equivalent in all cases is then⁸:

$$V(0) = \sum_{j=1}^J \lambda_j(0)K_j(0) \quad (3)$$

We can relax the assumption that wages are equal to marginal and average products of labor.

$$\int_0^{\infty} \left[\sum_{j=1}^J (pF_{K_j}K_j + pF_{I_j}I - z_jI_j) + \sum_{i=1}^L (pF_{L_i}L_i - w_i) \right] u(t) dt = V(0) = V_K + V_L \quad (4)$$

And now we are primarily concerned with the value of labor term V_L , and labor varieties are indexed by i .

$$V_L = \int_{t=0}^{\infty} V_L^t u(t) dt = \int_{t=0}^{\infty} \sum_{i=1}^L (pF_{L_i}L_i - w_i) u(t) dt \quad (5)$$

In the case that wages and marginal products of labor are equivalent for all labor types, at all employment quantities, and in all time periods, this term in equation (5) is zero. Since wages are set competitively by the market and the asset holdings of different firms vary, it is unlikely to be the case that all firms individually face the same marginal productivity of labor as the aggregate. Wages, on the

⁸ $V(0) = \sum_{j=1}^J \lambda_j(0)K_j(0) = \sum_{j=1}^J (\lambda_j(0)K_j(0) - \lim_{t \rightarrow \infty} \lambda_j(t)K_j(t)) = \sum_{j=1}^J \int_0^{\infty} (-\dot{\lambda}_j K_j - \lambda_j \dot{K}_j) dt = \int_0^{\infty} (\sum_{j=1}^J (pF_{K_j}K_j + pF_{I_j}I - z_jI_j) + \sum_{i=1}^L (pF_{L_i}L_i - w_i)) u(t) dt = V(0)$

other hand, are more likely to be consistent given ability or skills in the same region and industry, though there is evidence to suggest wage inequality can also be driven by firm characteristics (Song et al. 2015).

Equation (5) describes the potential surplus that an employer receives from the aggregated marginal products of its employees. The worker problem is deliberately simple: workers seek to maximize their wage subject to a constraint that it be above their reservation wage. Workers have one divisible unit of labor to supply. Assume now that the production function can be decomposed as follows as a function of the inputs in (1), but now firms can choose to assign workers to firm-specific labor (H) or general labor (L).

$$\max_{l,H,L} V(0) = \int_0^{\infty} \pi(t)u(t)dt ; \pi(t) = pA^t F(K, I, H, L, t) - w_H^t H - w_L^t L - z^t I - r^t K \quad (6)$$

Subject to additional constraints for workers indexed by l :

$$H_l + L_l = 1 \quad \forall l \in 1,2,3, \dots, N$$

$$H_l \geq 0, L_l \geq 0 \quad \forall l \in 1,2,3, \dots, N$$

$$w_{H_l} \geq w_0, w_{L_l} \geq w_0 \quad \forall l \in 1,2,3, \dots, N \quad \forall t$$

So that all workers have only one unit to supply, that negative labor is not possible, and that the reservation wage for all workers w_0 is met, guaranteeing participation. Of course, every employer has *some* tasks that are firm-specific and that value, rather ironically, will be compensated as part of a general wage because workers can foresee some of what their employers will have them do. The firm-specific value I consider here is “extra” and arises from an incomplete ex-ante contracting problem. The worker does not really know everything their employer might ask of them. Realistically extra firm-specificity might have its roots in job search frictions, worker preferences for more specialized tasks, unique business strategies and production functions, or even programs where workers are

allowed to spend some proportion of their time as they choose (Stern 2004; Tambe, Ye, and Cappelli 2019). Given that this firm-specificity of tasks exists, however, the job design is a channel for the exercise of monopsony power in the labor market.

The firm will maximize the present value of all future discounted profits (productivity term A is included to note that productivity can change over time). $F(K, I, H, L, t)$ serves to transform capital, investment, firm-specific labor, and general labor into production output. In addition to the same set of assumptions on equation (1) to get the solution in equation (4), I assume that F is nondecreasing and concave in H and L . Firms therefore know the ability of the workers they hire, and choose for them whether they work on “H” tasks or on “L” tasks. I also assume that the surplus from firm-specific tasks H are more appropriable to the firm than non-specific tasks, and that there are strictly increasing differences in the marginal value of H and L tasks (this guarantees a single-crossing of the marginal values of firm-specific and general labor task marginal products). Assume the firm’s share of firm-specific task marginal products of labor is $\beta \in [0,1]$. This parameter is assumed to be exogenous, to avoid focusing on a specific bargaining process. Formally, if $Q = H + L$,

$$\exists Q^* = H^* + L^* \text{ s. t.}$$

$$\beta \frac{\partial F}{\partial H}(Q, K, I) > \frac{\partial F}{\partial L}(Q, K, I) \quad \forall Q < Q^*$$

and

$$\beta \frac{\partial F}{\partial H}(Q, K, I) \leq \frac{\partial F}{\partial L}(Q, K, I) \quad \forall Q \geq Q^*$$

$$\therefore \frac{\partial^2 F}{\partial H \partial L}(H, L, K, I) > 0 \quad \forall (H, L, K, I) \quad (7)$$

This assumption characterizes firm-specific and general labor as complements, both required as part of the production function. Though I assume this for tractability purposes, future work could

investigate relaxing this assumption. The firm now observes its hired labor pool and assigns each worker l to some proportion α_l of firm-specific tasks, with the remainder of the worker's time spent on non-specific tasks (or general tasks).⁹ The employers make these decisions perhaps with the workers' skills or abilities in mind. The assignment function for the employer seeks simply to maximize the marginal product of each worker's labor given the other inputs factor vectors. α denotes the $N \times 1$ vector of assignments to firm-specific tasks and each worker has 1 unit of labor to supply:

$$\alpha^* = \operatorname{argmax}_{\alpha} \left\{ \beta \frac{\partial F}{\partial H}(H_{\alpha}, L_{\alpha}, K, I) + \frac{\partial F}{\partial L}(H_{\alpha}, L_{\alpha}, K, I) \right\} \quad (8)$$

$$H_{\alpha} = \int_{l=0}^N \alpha_l dl$$

$$L_{\alpha} = \int_{l=0}^N (1 - \alpha_l) dl = N - H_{\alpha}$$

$$s. t. \alpha_l \in [0,1] \forall l$$

Here the firm only receives β units of marginal product for each unit of H tasks; workers bargain away the remaining proportion. The employer's share of the marginal product of each worker is a sum of the task-specific marginal products. We guarantee participation of each of the workers with the participation constraint in (6). Sorting workers in order of their proportion of work in the firm-specific task (index the highest proportion worker at 0, lowest at N), we get the period-specific employers' and workers' returns to labor in equations (9) and (10) (given t). The workers in general tasks earn a wage w_L for which the firm is a price-taker. While the share is temporarily fixed, workers bargain away surplus proportional to the marginal product of firm-specific tasks at their rank l . The

⁹ Assume, for example, that each worker has an endowed positive productivity level in H and L tasks (ω_H, ω_L) , drawn from a stable distribution, that the firm can observe.

workers' share of firm-specific surplus would be lower if all workers accepted the same share of the aggregate marginal value of firm-specific labor.¹⁰

$$V_L^t = \int_{l=0}^N \left(\alpha_l^* \beta \frac{\partial F}{\partial H}(H_l, L_l, K, I) + (1 - \alpha_l^*) \left(\frac{\partial F}{\partial L}(H_l, L_l, K, I) - w_L \right) \right) dl \quad (9)$$

$$w_l = \alpha_l^* (1 - \beta) \frac{\partial F}{\partial H}(H, L, K, I) + (1 - \alpha_l^*) w_L \quad (10)$$

where

$$X_l = \int_{j=0}^l \frac{\partial X}{\partial j} dj \text{ for } X = H, L$$

And at $l = N$ where N is chosen to maximize profits¹¹:

$$\alpha_N \beta \frac{\partial F}{\partial H}(H_N, L_N, K, I) + (1 - \alpha_N) \left(\frac{\partial F}{\partial L}(H_N, L_N, K, I) - w_L \right) \leq 0$$

The firm's returns to labor investments come from two sources: the share of total surplus they recover from firm-specific labor H and the difference between the marginal product of general-task labor and the general labor wage set by the market. Under the standard neoclassical assumptions that general-task labor is competitive and the demand curve for L is perfectly elastic, this second term under the integral in equation (9) goes to zero. If employers do not invest in firm-specific labor tasks, the market value of labor will be zero. Otherwise, the first term in the integrand in (9) represents the employer share of firm-specific labor. Maintaining the assumption that the firm surplus for general task labor is close to zero, the correlation between human capital measures and market value is

¹⁰ Again for tractability, I assume the rate of change in labor of both types to be differentiable in the index of the workers l . This could easily be discretized and summed instead.

¹¹ If the share of general tasks is 1 at $l=N$, this reduces to difference between the marginal product of general labor and the prevailing general wage in the market (which we expect is close to zero). High frictions are also important, though not directly included in the model (they can be thought of as entering through the firm's bargained share).

affected by 1) changes in the firm's share of surplus β and 2) The quantity of firm-specific labor at the firm, and 3) complementarities (or substitution effects) between capital and firm-specific labor. Market value regressions without firm-specific labor measures may therefore be sensitive to omitted variable bias.

On the labor side, earnings come from two sources: general-purpose labor marginal wages and the workers' bargained shares of *marginal* firm-specific labor value. Differences between average and marginal wages accrue to the firm. Incidentally this model may partly explain the recent separation between labor productivity growth and wage growth (Brynjolfsson and McAfee 2014; Bivens and Mishel 2015; Stansbury and Summers 2017). If general labor wages are falling, perhaps due to outsourcing or increased labor supply, wage growth must come from either increased bargaining for firm-specific surplus or larger overall firm-specific surplus from labor. If employers, perhaps because of specialization or insulation from competition in labor markets, shift their share of work to firm-specific tasks, this could also put downward pressure on wages.¹²¹³ The single-crossing property in (7) suggests too that as employers increase total employment past Q^* , general-task labor will increase as a share of total labor in the firm.¹⁴

The model can be applied to specific occupations as well. Hiring larger quantities of a given occupations will mean that more of the available firm-specific surplus is captured. As bottlenecks of expensive firm-specific labor are alleviated, the firm comes closer to realizing the maximal value of complementary assets. Sample inverse labor demand curves are displayed below in the figure. The

¹² The bargained firm-specific wage must be as least as large as the general wage, or workers can move to a firm that will pay them for entirely general labor.

¹³ Lippman and Rumelt (1982) suggest another possible explanation in business complexity. Competitive pressure is ameliorated when business processes are especially difficult to reverse-engineer. Technology has expanded the combinatorial space of viable business models. In the long-run profits are competed away, but in the short-run a curse of economic dimensionality with scarce productive inputs can make competitive effects slow to act. I leave the exploration of the cardinality of competition to future work.

figure shows wage as a function of quantity demanded (Q_D) for general-task labor (blue), firm-specific task labor less the baseline wage for general labor (red), and the aggregate attainable by varying alpha in a linearly decreasing manner (purple). The single-crossing point is where the red and blue curves intersect.

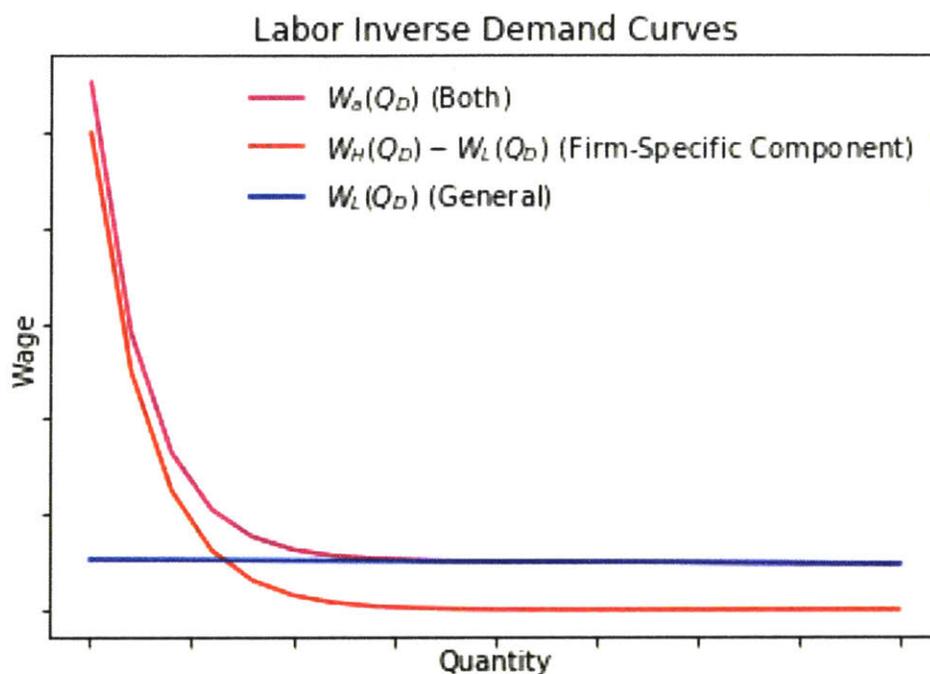


Figure 1 – Labor Inverse Demand Curves with Firm-Specific Tasks

These dynamics provide a motivation for why the labor buyer surplus in (5) might vary as a function of the productivity of firm-specific assets and labor. Since wages are a function of supply of labor conditional on skills and abilities, a technological shock increasing supply and driving wages for general-task labor down (or expected future wages) would change the firm's valuation. This mechanism works both through the direct effect of lower wages, but also through the marginal productivity of firm-specific capital and installed capital assets, i.e. changes in λ , at the new equilibrium. TensorFlow is the expected future labor supply-increasing shock I now turn to for the case of AI. Machine learning talent is an expensive complement to machine learning capital assets like

data and computational power. TensorFlow leads to more abundant machine learning labor, alleviating a bottleneck preventing employers with machine learning assets from realizing larger returns. For the case of machine learning talent, the expectation is that TensorFlow or similar software cause a future increase in the available labor supply. The figure below displays a simplified version of the intended supply shift from the green supply curve to the orange one. This moves the labor market equilibrium from the intersection of the purple inverse demand function (as above) and the green supply curve to the intersection of the inverse demand function and the orange supply curve. Wages are set by the intersection of the supply curves with the purple demand curve, and therefore the shift from the green supply to the yellow supply increases the firm's surplus as the buyer of labor (in red). In the last part of equation (9), the firm sets N such that the marginal surplus of an additional worker is less than or equal to zero. The alleviation of large search frictions and/or inelastic supply (omitted from the model, but a major concern for AI employers) would cause the equilibrium hiring breakeven point to increase. In Section 6 I provide evidence that this is what happens following the launch of TensorFlow.

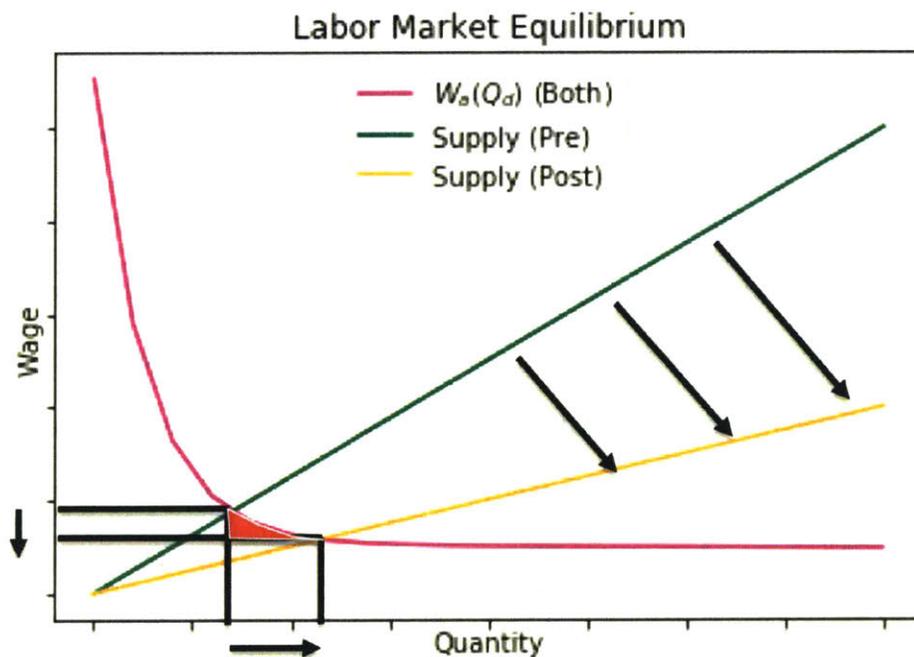


Figure 2 – Labor Market Equilibria with Firm-Specific Tasks and Increasing Skill Supply

4. Dataset Construction

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.¹⁵ The LinkedIn platform has become a standard tool for job seekers in many labor markets. 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 and total wage bills (employee counts multiplied by Bureau of Labor Statistics average wages) of specific varieties of worker. Engineering, information technology, and research worker talent counts and wage bills are some of the aggregated firm variables I construct. The process is similar to the variable construction in (Brynjolfsson et al. 2018; Benzell, Lagarda, and Rock 2018). Figures 3a-c show a representative LinkedIn profile.

¹⁵ Source: The LinkedIn Economic Graph Research team. 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.

Jacqueline Barrett • 1st
Economic Graph at LinkedIn
San Francisco Bay Area

Over 10 years of experience analyzing complex datasets, picking out key insights and communicating to relevant stakeholders. Currently using LinkedIn's rich data set to uncover trends in local economies as well as understand complex labor market phenomena. Economics • Data Science

Experience

- Staff Economic Researcher**
Oct 2015 - Present • 1 yr
Lead a cross-functional team to deliver Skills Gap Analysis, a metric that tracks skills gaps across 100+ cross-industry settings and heavy data in absolute headcount numbers. This metric provides transparency to workers, policymakers, educators and employers so that can begin to close skills gaps locally and nationally. Responsible for hiring and mentoring of Economists, Data Scientists and other employees.
- Economic Graph at LinkedIn**
- Lead Economic Researcher**
Mar 2016 - Sep 2017 • 1 yr 7 mos
Co-founder and Data Lead for the LinkedIn Workforce Report, a monthly labor-market report that tracks changes in hires, migration and skills for US and top 25 metropolitan regions. Insights have been picked up by publications such as the NYT, tweeted by the President and shared.
- Compass Laureate**
4 yrs
- Economist**
2014 - 2016 • 2 yrs
- Senior Analyst**
2012 - 2014 • 2 yrs
- Federal Reserve Bank of Chicago**
4 yrs
- Senior Associate Economist**
2011 - 2012 • 1 yr
- Associate Economist**
2008 - 2011 • 3 yrs

Education

- The University of Chicago - Booth School of Business**
MBA, Finance, Accounting
2014
- Northwestern University**
BA, Economics
2006

Volunteer Experience

- Encores! Board Member**
San Francisco Ballet
- Impact Coach**
TechWomen

Skills & Endorsements

- Data Analysis** 12
Endorsed by Steve Collins, MBA, MS&Z, who is highly skilled at this. Endorsed by Guy Berger, Ph.D. (mutual connection)
- Econometrics** 11
Endorsed by Mariacristina De Haro, who is highly skilled at this. Endorsed by Guy Berger, Ph.D. (mutual connection)
- Economics** 10
Endorsed by Guy Berger, Ph.D. and 1 other who is highly skilled at this. Endorsed by 3 of Jacqueline's colleagues at LinkedIn

Figure 3a, 3b, 3c – LinkedIn Profile, Experience, Education and Skills

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.

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 Brynjolfsson et al. (2018) and Benzell, Lagarda, and Rock (2018), I build a crosswalk between LinkedIn’s internal occupational classification system and the BLS-OES Standard Occupational Classification (SOC) Code by year. For firm-level aggregate employment data, I use the Compustat/Capital IQ North America database value of EMP. 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 form¹⁶:

$$\widehat{EMP}_{it} = \alpha + LI'_{it}\beta + \gamma_{jt} + \epsilon_{it} \quad (10)$$

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, a fixed effect for that industry-year combination, and an error term. 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

¹⁶ Prediction accuracy gains from models with higher complexity (e.g. tree-based models or support vector machines) were relatively small

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 (10) in Table A1. Typically firms have about 1.9 times as many employees as are available on LinkedIn, controlling for the asset base size and industry-year.

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 (10). Re-weighting the BLS-OES occupation-industry-year shares by 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 λ_{jt} 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 θ_{it} , 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 $\overline{\theta}_{jt}$, we 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:

$$LI_{ijt} = \frac{(\theta_{it}\lambda_{jt})}{\theta_{jt}} \text{Compustat}_{it} \quad (11)$$

The end result is Compustat-BLS-OES-consistent firm-year-occupation employment coverage ratios. Occupations like software engineer, unsurprisingly, have high fidelity and near complete coverage 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.¹⁷ 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 (11). Those employee counts are multiplied by their BLS-OES wage in the relevant respective year to construct the wage bill variables.

¹⁷ 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.

I also 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 Brynjolfsson et al. (2018), I aggregate the educational records of the workers according to the years of education required to achieve each listed degree.¹⁸ 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.¹⁹ Descriptive statistics for the LinkedIn human capital measures for 2006 and 2016 can be found in Table 1 below.

2016 Summary Stats	Market Value (MM USD)	Total Assets (MM USD)	IT Employment	Research Employment	Engineering Employment	IT Wage Bill	Research Wage Bill	Engineering Wage	Education Years
Count	7,365	7,365	1,828	1,828	1,828	1,828	1,828	1,828	2,249
Mean	17,081.06	14,133.41	972.40	198.57	1,556.18	96,903,181.60	27,533,898.31	120,046,634.46	35,189.00
Standard Deviation	112,933.43	109,319.50	4,024.43	692.30	4,759.50	405,197,401.23	93,265,923.08	387,763,106.94	115,419.34
0.25	70.24	40.51	35.01	5.21	42.96	3,374,545.19	827,245.57	3,142,992.88	1,053.44
0.50	694.88	436.15	140.77	22.91	239.57	13,708,210.73	3,350,099.03	16,838,521.82	5,495.35
0.75	4,260.37	2,712.24	539.67	109.90	977.20	52,722,407.66	15,822,372.62	70,194,131.19	23,867.57
2006 Summary Stats	Market Value (MM USD)	Total Assets (MM USD)	IT Employment	Research Employment	Engineering Employment	IT Wage Bill	Research Wage Bill	Engineering Wage	Education Years
Count	8,453	8,453	2,211	2,211	2,211	2,211	2,211	2,211	2,994
Mean	11,290.06	9,080.16	517.95	128.16	1,089.52	35,424,886.64	9,797,113.28	63,189,456.84	24,032.11
Standard Deviation	85,117.95	80,507.76	2,033.99	567.22	3,935.39	141,847,510.53	34,689,741.29	242,870,484.33	90,667.33
0.25	77.53	35.89	16.87	2.55	26.54	1,126,353.06	286,148.23	1,584,233.32	600.57
0.50	441.37	246.90	69.90	11.11	128.77	4,666,854.58	1,162,623.08	7,112,400.35	2,647.09
0.75	2,282.32	1,394.58	278.49	51.24	543.20	18,859,136.17	4,921,754.10	29,885,631.95	13,434.37

Table 1: Descriptive Statistics of Employment Measures 2006 and 2016

The final set of LinkedIn-derived values come from the relatively recently constructed panel of detailed skills detail. LinkedIn first rolled out the skills product in 2011, though collection of high-fidelity records of member additions of skills began in 2014. Recently, LinkedIn has categorized and standardized the over 35,000 unique skills on its standard platform into a set of skills clusters using

¹⁸ The normalization of education years to adjust for coverage is an identical process to the count normalization process.

¹⁹ Can be considered as the answer to the question “how many years has the firm gone to school?”

nonlinear embedding spaces.^{20,21} These clusters are seeded by humans and subsequently applied to co-occurrences of skills on profiles across the entire platform. 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.

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 also extract specific skill counts for deep learning, machine learning, R, SPSS, and a handful of other data science skills. All of these measures are then joined to Compustat measures of financial performance by firm and year. Figures 4A-D show the aggregate skill additions for AI-related skills and advertising. There is some seasonality in the data, with more skills getting added in the beginning of the year. The table below shows some example skills for different aggregated categories.

Skills	Artificial Intelligence	Digital Literacy	Data Science	Advertising	Business Management	Data Storage Technologies
1	artificialintelligence	microsoftexcel	forecasting	advertising	management	microsoftsqlserver
2	machinelearning	microsoftword	modeling	campaigns	strategy	mysql
3	classification	microsoftpowerpoint	statistics	collateral	strategicbusiness	sql
4	informationretrieval	microsoftoffice	analytics	sponsorship	smallbusiness	forecasting
5	computervision	microsoftaccess	dataintegrity	directmarketing	strategicplanning	databases
6	neurálnetworks	microsoftoutlook	statisticaltools	searchenginemarketing(sem)	changemanagement	datacenter
7	speechrecognition	email	dataanalysis	branddevelopment	executivemanagement	storage
8	semanticweb	spreadsheets	sas	mediaplanning	serviceproviders	datawarehousing
9	parsing	mac	datamining	emailmarketing	outsourcing	hprproducts
10	patternrecognition	lotusnotes	sampling	media&entertainment	businessplanning	pl/sql

Example Skills Table

Note: Example skills are not exhaustive – some categories have hundreds of skills.

²⁰ 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.

²¹ 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.

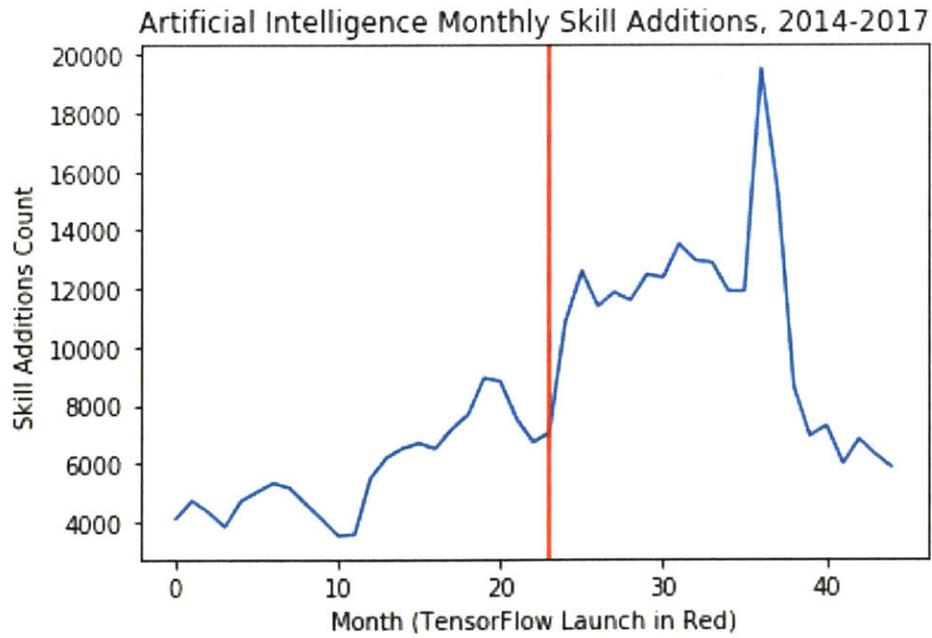


Figure 4A – Artificial Intelligence Skills 2014-2017

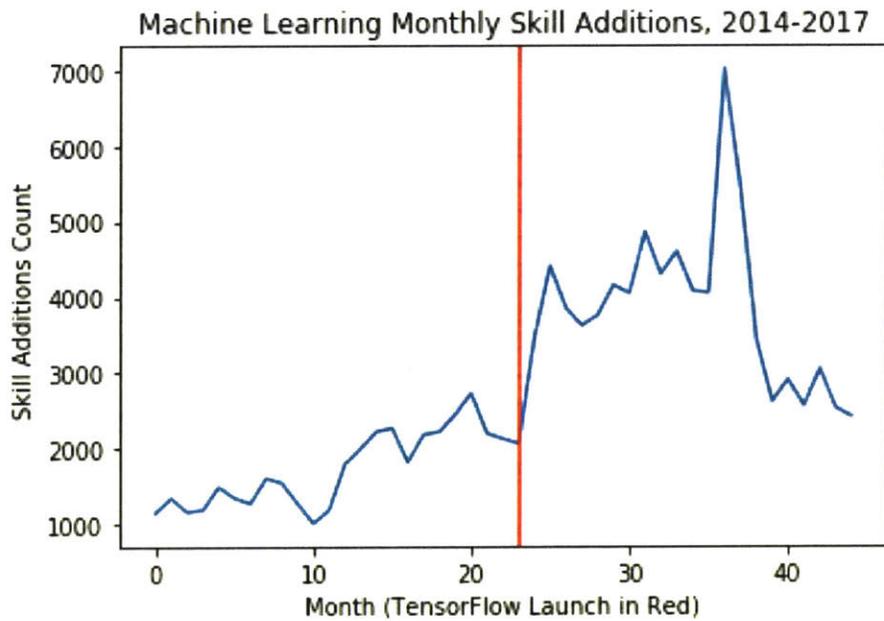


Figure 4B – Machine Learning Skills 2014-2017

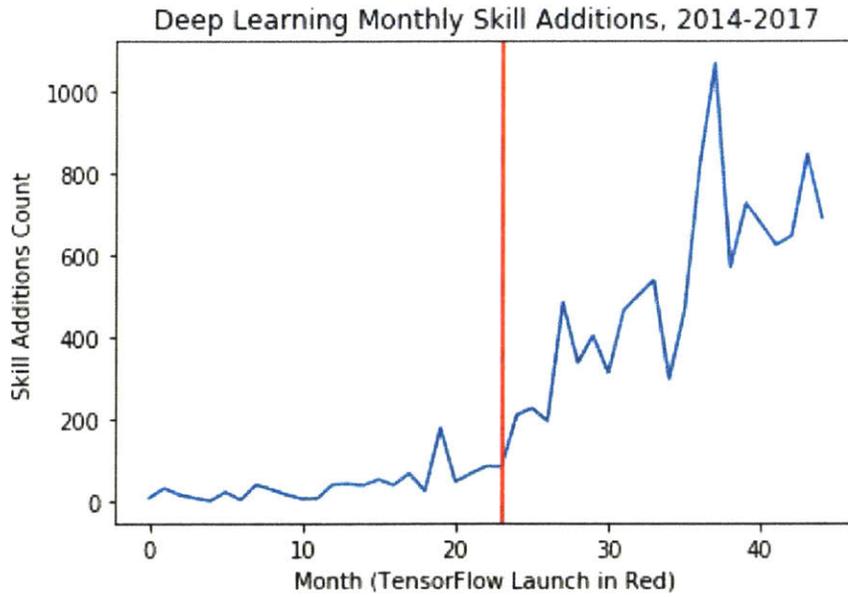


Figure 4C – Deep Learning Skills 2014-2017

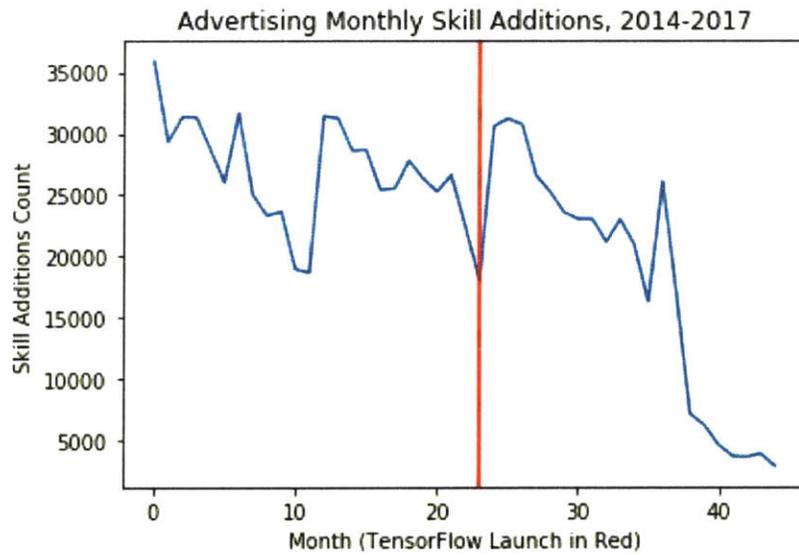


Figure 4D – Advertising Skills 2014-2017

Compustat also indicates the state, zip code, and county for the corporate headquarters. I link these states to changes in covenant-to-not-compete (CNC) policy changes at the state-level using the

set of states in (Ewens and Marx 2017) and (Jeffers 2017) separately. These changes in CNC policy are known to impact knowledge workers with greater propensity than other kinds of employees (Balasubramanian et al. 2018), and they constitute a first-order supply-side shock as mobility restrictions on human capital when the enforceability of these contracts increases. CNC enforceability changes should therefore serve as an instrument for the technical talent hired by the firm. Following Ewens and Marx (2017) and Jeffers (2017), I code increased enforceability by state courts as a 1, unchanged values as a 0, and weakened enforceability as -1.

I also merge the corporate headquarters zip and county codes to the zip and county codes of the land-grant universities. Land-grant universities, established by the Morrill Acts of 1862 and 1890, provided for the creation of colleges in each state following the sale of federal lands. As in Moretti (2004), worker proximity to land-grant institutions predicts higher likelihood of human capital accumulation. To attempt to recover a causal estimate of the market value of technological workers controlling for overall human capital, I include a dummy variable in instrumental variables specifications for proximity to land-grant institutions of the corporate address at the county level.

For the final instrument set, following (Lucking, Bloom, and Van Reenen 2017; Lucking 2018), I match each firm-year to the logged user cost of R&D capital value for both the firm and the best matching state. R&D investment subsidies for each state-year and firm-year lower the cost of corporate R&D investments. Specifications logging these rates and including time fixed effects leave only the state or firm-level variation net of changes in overall capital costs. Since R&D stocks are a known channel for these cost decreases to increase market value, I also include R&D stocks at the firm-level calculated from Peters and Taylor (2017). The conditional effect estimated is therefore adjusted for the innovative activities of the firm, leaving only non-innovative engineering.

NIFA LAND-GRANT COLLEGES AND UNIVERSITIES

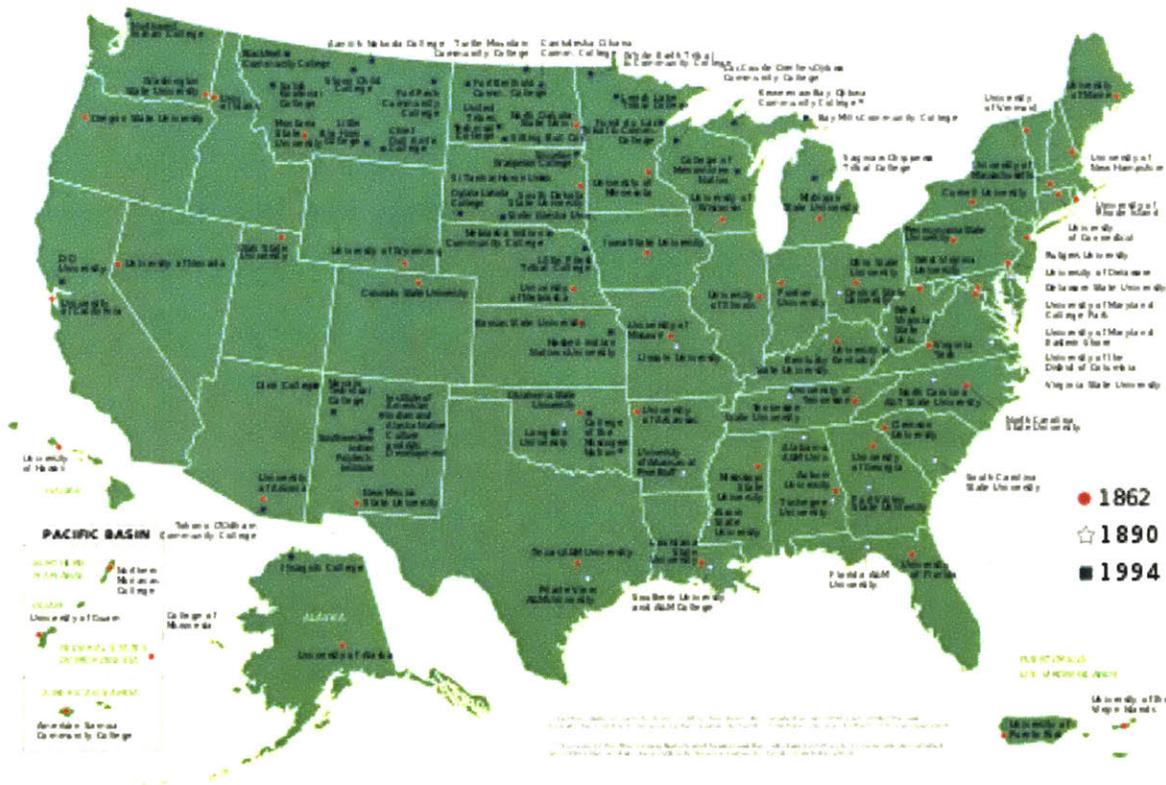


Figure 5 – Land-grant Colleges

The Compustat component of the dataset is mostly taken as-is, with market value (MV) constructed as the total book value of the firm plus the market value of equity at the end of the fiscal year less the book value of common equity. Total Assets (TA) is included in most regression specifications as a control for the capital size of the firm.

I also join in measures of *Suitability for Machine Learning* (SML) at the occupation level. SML measures are from Brynjolfsson, Mitchell, and Rock (2018a), wherein they use the crowdsourced evaluation of thousands of rubric surveys to construct measures of task SML and measurability for each of the detailed work activities (DWAs) supplied in the O*NET database (Brynjolfsson and Mitchell 2017). High relative values of SML indicate an opportunity to use machine

learning (and deep learning in particular) to automate aspects of a task. These scores are aggregated across tasks to the occupational level, and I subsequently aggregate the occupation-level SML scores to the firm-year by generating wage bill-weighted averages of SML scores. If TensorFlow is a shock to the availability of talent, it is also a shock to the opportunities for automation (but not necessarily at the same firm). With measures of which firms have an opportunity to deploy deep learning talent and which firms have an opportunity to use what deep learning engineers build, I can address whether it is the firms who create or the firms who consume technology who capture the rents to technology's effects on labor.

5. Overall Engineering: Empirical Results and Discussion

Technological human capital assets, if they contribute to the market value of the firm, can be priced following an equilibrium relationship that the asset's marginal Q-value above the asset replacement cost must equal the marginal adjustment costs of competitors (Wildasin 1984; Hayashi and Inoue 1991; Hall 2001). Summing over assets within firm and year, we get that the market value of the firm is equal to the sum of the market value of its constituent assets priced at Q. In other words, a regression of market value on measures of the assets of the firm will recover as coefficients the value per unit (dollars) for the asset in equilibrium. To find the value of different types of technological labor, I estimate the coefficient vector for the following regression:

$$MV_{it} = \beta * \text{Total Assets}_{it} + \gamma * HK_{total_{it}} + \sum_{m=1}^M \delta_m \text{TechHK}_{mit} + \mathbf{Z}'_{it} \boldsymbol{\lambda} + \epsilon_{it} \quad (12)$$

In this regression, i indexes the firm, t indexes the year, and m indexes the variety of technological labor. The equation therefore describes the decomposition of the panel of firm market values on the total book value of assets (Total Assets), the total education years in the firm at year t (HK_total), and each type of technological labor TechHK.²² A vector \mathbf{Z} of controls including industry-year fixed

²² Depending on specification, wage bill or counts might be included here

effects and firm fixed effects in some of the specifications are also included. This is a standard market value regression of the sort in (Hall 1993; Brynjolfsson, Hitt, and Yang 2002; Brynjolfsson et al. 2018). However, as suggested by the results in equations (4) and (9), the market value of the firm is also a function of the integrated differences between worker marginal products and worker wages. Equation (12) therefore sets a specification by which we can decompose market value into the value of observable assets and the value of labor-correlated inputs. The coefficient vector for technological talent recovers the “installed” average price of the talent itself and the value of omitted yet talent-correlated assets.

The results for the set of OLS regressions from (12) are in Tables 2 (counts) and 3 (wage bills), revealing strong correlations between market value and engineering talent where firm fixed effects are left out. While IT workers and Research workers appear not to be correlated with market value after controlling for the total education years attained by the firm and the total book value of assets, along with the appropriate fixed effects, an additional engineer is correlated with an increase of approximately \$855,000 of market value (Table 2, column 5), and each dollar of engineering wage bill is correlated with an additional \$11-12 of market value (Table 3, column 5). However, when including firm-level fixed effects, the correlation between all types of technological talent and market value is no longer statistically significant and the point estimates on engineering talent drop. They are not statistically different than the industry-year estimates. Firm fixed effects effectively adjust for the aspects of the firm that allow it to successfully generate value. Including firm fixed effects, in effect, results in an estimate of market value as though the firm were competitive with itself. That is, the variation which generates value is cross-sectional (and over time). The within-firm estimates might be thought of as a lower bound if corporate asset bases are relatively static.

This is consistent with an explanation of firm-specific assets driving the returns to investments in technological labor. That is, there is either a valuable (priced) intangible correlate asset

for technological labor that is firm-specific generating a correlation between market value and engineering labor as an omitted variable, or the component of engineering labor that causes market value is firm-specific and time-invariant.²³ That firm fixed effects drive the coefficient on engineering talent to lose statistical significance is consistent with either omitted firm-specific assets being the primary source of the correlation between market value and engineering talent or the marginal value of firm-specific labor being close to zero. The former explanation suggests off-balance sheet capital is the source of the empirical relationship. This latter explanation would describe a marginal causal effect of engineering talent on market value. Both explanations are consistent with monopsony power coming from the firm-specificity of human capital.

The wage regressions (Table 3) compare the flow of wages to workers of a given type to the market value stock outcome variable, equal to the present value of all future cash flows. Since the wages are what the worker earns, the coefficient on the wage is the present value per \$1 of wages paid in a given year of *all future flows* to the firm. That is, the wage is what workers earn in a year; the table coefficient is the discounted value of what the firm will get in all future years. This inflates the coefficient value because each year's wages must represent all future wages to be comparable to the firm's value. Either representing wages as a stock or representing market value as a one-year flow fixes the comparison problem. Unfortunately, the former approach requires knowledge of the initial stock of firm capital in wages and the depreciation rate (to calculate a perpetual inventory) and the latter requires knowledge of the appropriate discount rate for engineering capital's share of market value. As a back-of-the-envelope calculation, the stock of assets in wages paid to technological talent is 14.3 to 20 times the wage value.²⁴ This would imply that engineer wages as a stock are worth about 59 to 83 cents per dollar net of what the worker is paid if there are no omitted variables. In the case

²³ As a sanity check, the correlations between total assets and market value are close to replacement cost (\$1) and the market value of human capital is positive and statistically significant.

²⁴ Assuming an aggregated human capital depreciation rate of between 5% and 7%, the approximate rate of interest for acquiring human capital: <https://studentaid.ed.gov/sa/types/loans/interest-rates>

that marginal and average products are equal to wages, the 59 to 83 cents estimate is the value of off-balance sheet assets correlated with the presence of engineering talent.²⁵

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Table 2: Market Value – Worker Count Regressions	No Tech	IT Value	Engineering Value	Research Value	All Tech Value	All Tech Value (Firm FE)	Tech w/o HK Value
Total Assets	1.012*** (0.00768)	1.012*** (0.00691)	1.013*** (0.00627)	1.014*** (0.00751)	1.012*** (0.00615)	1.003*** (0.0122)	1.004*** (0.0119)
Total Years of Education	0.0187*** (0.00451)	0.0138** * (0.00429)	0.0125*** (0.00362)	0.0125** (0.00524)	0.00622 (0.00448)	0.00873 (0.00988)	
IT Employees		1.301** (0.498)			0.496 (0.475)	-0.111 (0.302)	0.0737 (0.273)
Engineering Employees			1.129*** (0.289)		0.855*** (0.276)	0.567 (0.422)	0.817* (0.478)
Research Employees				7.164* (4.231)	6.745 (4.428)	0.843 (2.288)	2.208 (2.391)
Observations	50,501	37,813	37,813	37,813	37,813	37,825	37,825
R-squared	0.984	0.984	0.984	0.984	0.984	0.994	0.994
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	No	No
Firm and Year FE	No	No	No	No	No	Yes	Yes

Table Notes: Robust standard errors in parentheses, Standard errors clustered at the industry (3-Digit NAICS) for columns 1-5, firm for columns 6-7. Market value is in millions USD.

*** p<0.01, ** p<0.05, * p<0.1

²⁵ Actual discount rates for this kind of talent vary by firm. Some firms will have much higher (lower) discount rates, in which case the employer share of wage value would be substantially higher (lower).

Table 2 – OLS Market Value Regressions on Worker Counts

	(1)	(2)	(3)	(4)	(5)	(6)
Table 3: Market Value – Wage Bill Regressions	IT Value	Engineering Value	Research Value	All Tech Value	All Tech Value	Tech w/o HK Value
Total Assets	1.012*** (0.00693)	1.013*** (0.00638)	1.011*** (0.00690)	1.010*** (0.00609)	1.003*** (0.0122)	1.004*** (0.0119)
Total Years of Education	0.0148*** (0.00431)	0.0134*** (0.00366)	0.0136*** (0.00424)	0.00985*** (0.00359)	0.00912 (0.00931)	
IT Wage Bill	1.30e-05** (5.10e-06)			1.68e-06 (3.92e-06)	-5.13e-06 (3.38e-06)	-3.63e-06 (2.89e-06)
Engineering Wage Bill		1.57e-05*** (4.75e-06)		1.19e-05** (4.93e-06)	1.02e-05 (6.62e-06)	1.33e-05* (7.57e-06)
Research Wage Bill			7.32e-05** (3.20e-05)	6.04e-05 (4.00e-05)	9.91e-06 (1.17e-05)	1.51e-05 (1.13e-05)
Observations	37,813	37,813	37,813	37,813	37,825	37,825
R-squared	0.984	0.984	0.984	0.984	0.994	0.994
Industry-Year FE	Yes	Yes	Yes	Yes	No	No
Firm and Year FE	No	No	No	No	Yes	Yes

Table Notes: Robust standard errors in parentheses

Standard errors clustered at the industry (3-Digit NAICS) for columns 1-4, firm for columns 5-6. Market value is in millions USD. The wage bill is equal to the prevailing wage for a given occupation-year grouping, where the occupational wage is the employment-weighted average wage of all BLS-OES occupation categories within a given LinkedIn occupational category, matched many-to-one. The wage bill is then a *flow* measure and the market value is a *stock* equal to the present value of all future flows.

*** p<0.01, ** p<0.05, * p<0.1

Table 3 – OLS Market Value Regressions on Worker Wage Bills

In the case that the human capital value is firm-specific, a causal shift in employee supply driven by CNC enforceability tightening, for example, would suggest market values should increase. In this scenario, the policy change forecloses on employment alternatives for employed workers. On the margin, newly lower opportunity costs for employees might make further capital accumulation attractive. If the instrumented engineering labor does not cause market value, it is suggestive evidence that (for compliers), the correlation between market value and engineering talent is driven by hidden intangible assets owned by the firm. If this is instead the case, the valuation of off-balance sheet assets correlated with engineering talent values is less sensitive to employee opportunity cost changes. Monopsonists will hire less than the socially optimal amount if they have to pay all employees the same wage. They receive a mark-down on wages. Supply-increasing shocks will increase both the market value of the firm and the quantity of labor hired if wages for existing workers are renegotiable. Wages might be sticky though. The fact that CNCs also affect existing workers by shifting the optimal balance of firm-specific and general tasks means that even if the firms keep their hiring the same, the market value of firms will respond to CNC policy changes.

As described above, I will use the proximity of the corporate headquarters to land-grant universities and changes in state-level covenant-to-not-compete (CNC) policy as instruments for the firm's stock of human capital (in the land-grant case) and as supply shocks for engineering talent wherein increased (decreased) enforceability reduces (increases) the quality of workers' outside options (in the CNC case). The nature of the instruments does not allow inclusion of firm fixed effects in the specification as the covariate matrix will be of deficient rank. But interacting the CNC policy changes with the land-grant proximity dummies permits an overidentified model. In all specifications, the weak identification test F statistic is well above 20 and the overidentification test (a joint test of local average treatment effect homogeneity and correlation of residuals and instruments) fails to reject the null. This is reassuring, though not dispositive, with respect to endogeneity concerns of some

subset of the instruments. The second stage results for engineering employee counts and wage bills are in Tables 4A and 4B (respectively). The first stage results are reported in Tables 5A and 5B.

	(1)	(2)	(3)	(4)	(5)
Table 4A: IV MV Regressions:	Land-grant	CNC-Jeffers	CNC-Ewens-Marx	LG+CNC-Ewens-Marx	LG+CNC-Ewens-Marx
Engineering Labor					
Engineering Employees	0.965 (1.973)	4.297 (33.29)	0.404 (5.717)	1.067 (1.844)	1.484 (1.692)
Total Assets	1.035*** (0.0113)	1.023*** (0.0339)	1.027*** (0.0117)	1.035*** (0.0110)	1.032*** (0.0104)
Total Years of Education	0.0139 (0.0102)	-0.00137 (0.151)	0.0162 (0.0263)	0.0134 (0.00947)	0.00232 (0.00423)
IT Employees					1.052 (1.684)
Research Employees					6.740 (4.447)
Observations	31,822	32,622	32,622	31,375	31,375
R-squared	0.973	0.976	0.977	0.973	0.974
Industry-Year FE	Yes	Yes	Yes	Yes	Yes

Table Notes: Robust standard errors in parentheses, SEs Clustered by 3-Digit NAICS. Market value is represented in millions USD.

*** p<0.01, ** p<0.05, * p<0.1

Table 4A – IV Market Value Regressions on Engineering Worker Counts

	(1)	(2)	(3)	(4)	(5)
Table 4B: IV MV Regressions:	Land-grant	CNC-Jeffers	CNC-Ewens-Marx	LG+CNC-Ewens-Marx	LG+CNC-Ewens-Marx
Engineering Wage Bill					
Engineering Wage Bill	1.51e-05 (3.10e-05)	-5.60e-05 (0.000524)	7.74e-06 (0.000109)	1.68e-05 (2.90e-05)	2.12e-05 (2.87e-05)
Total Assets	1.035*** (0.0112)	1.032*** (0.0448)	1.027*** (0.0130)	1.035*** (0.0109)	1.029*** (0.00974)
Total Years of Education	0.0141 (0.00972)	0.0332 (0.142)	0.0159 (0.0300)	0.0137 (0.00907)	0.00452 (0.00339)
IT Wage Bill					7.27e-06 (1.98e-05)
Research Wage Bill					9.88e-05* (5.55e-05)
Observations	31,822	32,622	32,622	31,375	31,375
R-squared	0.973	0.963	0.977	0.973	0.975
Industry-Year FE	Yes	Yes	Yes	Yes	Yes

Table Notes: Robust standard errors in parentheses, SEs clustered by Industry (1-3), Firm (4-5). Wage bill is calculated as in Table 3 (see table notes).

*** p<0.01, ** p<0.05, * p<0.1

Table 4B – IV Market Value Regressions on Engineering Wage Bill

Table 4C: R&D User Cost IV (Engineer Counts)	(1)	(2)	(3)	(4)
Log Market Value	User Cost: State	User Cost: All	User Cost: State	User Cost: All
Log Engineers	-1.251	-1.072**	-0.165	-0.579**

	(0.845)	(0.513)	(0.367)	(0.260)
Log(Assets)	0.979***	0.939***	0.932***	0.904***
	(0.0527)	(0.0487)	(0.0231)	(0.0448)
Log(Years of Edu.)	0.676	0.656**	0.119	0.420**
	(0.428)	(0.289)	(0.198)	(0.163)
Log(R&D Capital)	0.109***	0.128*	0.0557***	0.100***
	(0.0397)	(0.0638)	(0.0161)	(0.0291)
Log(IT Workers)	0.467	0.351	0.0279	0.139
	(0.357)	(0.212)	(0.148)	(0.0873)
Log(Researchers)	0.0551	0.0143	0.0361*	0.0313
	(0.0477)	(0.0454)	(0.0203)	(0.0247)
Observations	16,252	12,450	16,246	12,441
Industry-Year FE	Yes	Yes	Yes	Yes
City FE	No	No	Yes	Yes

Standard errors clustered by industry (3-Digit NAICS) in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 4C – IV Market Value Regressions on Engineer Counts Instrumented by R&D User Costs

The causal estimates for the land-grant and CNC local average treatment effects (LATEs) of engineering human capital on market value are imprecise and statistically indistinguishable from zero. While the point estimates are larger (\$7.7 to \$21.2 per dollar of wages for specifications including the Ewens-Marx CNC data), they are also not statistically significantly different than the fixed effect OLS estimates.

Firm fixed effects soak up the variation generating a statistically significant relationship between market value and engineering talent measures. Broadly, this indicates that the average product

of engineering talent in firms is weakly larger than the wages paid over all employed engineers.²⁶ Nevertheless, there is little case that employers can extract freely available rents by hiring more engineers. The imprecisely estimated zero marginal value of an engineer net of wages is consistent with value at or below the opportunity cost of recruiting search frictions, training, and adjustment costs for employers. That engineers are so highly correlated on average with market value is suggestive of complementarities between engineers and firm-specific assets.

The land-grant instrument is designed to approximate (statically) a measure of the available supply of engineers to the firm. Ideally there would be an experiment randomly assigning assets and workers of different types to different companies. Under the assumption that sharing county locations with land-grant universities (as opposed to other universities) is otherwise excluded from market value, the land-grant IV detects contribution of wage changes to firm market value. The first stage suggests that, controlling for worker education and the other inputs, firms near a land-grant college are less likely to hire more engineers. This could be, for example, because human capital-intensive firms far from land-grant colleges hire relatively more engineers. This could be the case if the first high human capital employees are engineers for most firms, and on the extensive margin other types of educated workers are hired.

CNC policy changes, however, shift the outside options of incumbent workers, making it potentially more attractive for firms to invest in firm-specific training for those workers as departure is less likely. At the same time, hiring prospects are somewhat diminished. So CNC changes primarily operate through the price of human capital assets. Since neither of these instruments reveal a conclusive causal effect of engineering employment on market value, even if the coefficients are larger than they are in the OLS case. As a caveat, this is to be expected if equilibria are insufficiently changed by the instruments. Further, the inclusion of firm fixed effects and more precisely estimated

²⁶ And strictly larger for some engineers

lack of statistically significant market value correlation with engineering talent is evidence against a large *within-firm* value of engineering labor investment, but not on *average* in the cross-section. The firm fixed effects control for time-invariant omitted variables. The correlation between market value and engineering talent diminishes within firm. The difference between specifications with firm fixed effects and without them is suggestive then of a possible omitted asset generating rents. A large, statistically significant effect in the IV regressions would suggest the talent itself as the source of rents on the hiring margin. The R&D user cost instrument shows effects in the negative direction, adjusting for the innovative activities of the employers. Table 4C below shows the R&D user cost instrument specifications. These regression specifications have logged market value and logged covariates because of the skewed nature of R&D stocks. Table 5C has the corresponding first stage.

The construction of the R&D user cost instruments is to take the log of the firm and state-level user costs by firm-year.²⁷ The firm-level user cost varies as a result of federal subsidies for R&D at the firm level, but does not apply to all companies as the policy does not apply to newer firms. The state-level subsidy is applicable only to R&D expenditures within a given state. These state-level costs are matched to the firm locations. Additionally included is an interaction of these two variables. Firms that take up the subsidy will have lower costs of capital for R&D investment and ideally will accumulate larger R&D stocks. Adjusting for these R&D stocks (and other types of assets), the conditional IV estimator represents the causal effect of hiring another engineer *not* engaged in R&D activity. Since R&D expenses have a large salary component, this adjusted estimate represents the value of marginal engineering activity. It is unlikely, however, that any engineer is engaged completely in non-innovative activity.

²⁷ Lucking et al.(2018) and Lucking (2019) have greater detail on the construction of these instruments.

The estimate value from the R&D user cost specifications suggest that the marginal engineer engaged in non-innovative activity (somewhat surprisingly) destroys nearly \$580,000 in value with a standard error of \$260,000, adjusting for the corporate city location, innovative activity of the firm, and other assets. This could be due to the conditions under which such an engineer would be hired. Tight labor markets might mean that hiring an engineer under problematic conditions is more an indication that market value is declining. Since location fixed effects do attenuate the coefficient estimate toward zero, the availability of talent is likely an important component of the incentives facing employers and their employees. It may also be the case that more engineers does not always make for easier problem solving and on the margin, removing innovative activity, this is a net drain on the firm's value for the complier companies. Another possibility is that if the value of the engineer is a bundled combination of maintenance activities which have negative value and innovation activities with positive value, then each marginal engineer could have zero value but some activities create a cross-subsidy. Firms facing fixed job designs would therefore have an incentive to redesign the bundle of tasks in engineering jobs, disentangling innovative from non-innovative job tasks.

The CNC, land-grant, and R&D user cost IV regressions fail to provide a conclusive story. The first stage coefficients on CNC policy changes are not statistically significant, but they are for land-grant university proximity. The technological talent markets seem to be in equilibrium with respect to CNCs. The value of engineering talent is high on average, but nearly zero on the margin of the local average treatment effects given by the CNC and land-grant instruments (as suggested by firm fixed effects and both sets of IV regressions). If land-grant proximity does predict hiring, but shows no statistically significant estimated causal effect of hiring on market value, then on the margin we fail to reject the hypothesis that wages are equal to worker product. Finally, the R&D user cost regressions suggest that the marginal engineer may destroy value once the value of innovative activities is priced.

In the case that labor is relatively abundant, we might expect that the marginal hire is devoted principally to general tasks and contributes little to market value net of wages. In the case of engineers, the average value is high and relatively precise but the marginal value is noisy. On balance, these IV specifications rule out little more than large deviations from zero marginal value. There is little evidence of abundant rents for firms seeking to hire more technological talent *in general*. This is possibly because the firm-specific tasks have already reached the point at which general tasks are more valuable to the employer. This suggests the need to look at more granular cases to understand where employers might be able to hire more and generate value at the same time.

Some talent is bottlenecked, i.e. finding and recruiting workers of that type is very costly. For these workers, it is more likely that they work on firm-specific tasks which are higher marginal value. Workers for which there is a bottleneck will have high marginal and average contribution to market value as the frictions to find and/or train them will create a wedge between their wages and the value they create for their employers. AI talent offers a recent technological case study for what can happen when the market expects a previously bottlenecked talent to become much more abundant.

	(1)	(2)	(3)	(4)	(5)
Table 5A: IV Regression First Stage (Worker Counts)	Land-grant	CNC-Jeffers	CNC-Ewens- Marx	LG+CNC-Ewens- Marx	LG+CNC- Ewens-Marx
Total Assets	0.00178 (0.00209)	0.000999 (0.00180)	0.00100 (0.00180)	0.00165 (0.00194)	-0.000599 (0.000848)
Total Years of Education	0.00453*** (0.00138)	0.00452*** (0.00137)	0.00452*** (0.00137)	0.00447*** (0.00135)	0.00114 (0.000931)
CNC (Ewens-Marx)			87.11 (88.37)	64.02 (289.3)	-64.22 (203.1)

Land-grant County Dummy	-485.4***			-460.9***	-324.6**
	(174.5)			(168.6)	(127.5)
CNC (Ewens-Marx) X Land-grant				36.13	163.4
				(271.1)	(196.6)
CNC (Jeffers)	21.59				
	(112.7)				
IT Employees					0.983***
					(0.175)
Research Employees					0.0738
					(0.262)
Observations	31,822	32,622	32,622	31,375	31,375
Industry-Year FE	Yes	Yes	Yes	Yes	Yes
Firm and Year FE	No	No	No	No	No

Robust standard errors in parentheses, SEs Clustered by 3-Digit NAICS

*** p<0.01, ** p<0.05, * p<0.1

Table 5A – Worker Counts First Stage Regression (Companion Table to Table 4A)

	(1)	(2)	(3)	(4)	(5)
Table 5B: IV Regression First Stage (Engineering Wage Bill)	Land-grant	CNC-Jeffers	CNC-Ewens-Marx	LG+CNC-Ewens-Marx	LG+CNC-Ewens-Marx
Total Assets	130.6	82.06	82.34	122.9	-26.08
	(136.6)	(118.0)	(118.1)	(127.6)	(47.90)
Total Years of Education	272.0***	271.4***	271.5***	268.3***	67.41
	(93.71)	(93.16)	(93.12)	(91.67)	(58.16)
CNC (Ewens-Marx)			4.548e+06	909,957	-7.630e+06

		(6.071e+06)	(1.872e+07)	(1.390e+07)
Land-grant County Dummy	-3.092e+07***		-2.938e+07***	-2.010e+07**
	(1.125e+07)		(1.091e+07)	(9.199e+06)
CNC (Ewens-Marx) X Land-grant			4.899e+06	1.397e+07
			(1.739e+07)	(1.340e+07)
CNC (Jeffers)	-1.658e+06			
	(7.834e+06)			
IT Employees				0.712***
				(0.0776)
Research Employees				0.210**
				(0.0891)
Observations	31,822	32,622	32,622	31,375
Industry-Year FE	Yes	Yes	Yes	Yes
Firm and Year FE	No	No	No	No

Table Notes: Robust standard errors in parentheses, SEs Clustered by 3-Digit NAICS. Outcome for first stage is the dollar wage bill (wage times worker counts).

*** p<0.01, ** p<0.05, * p<0.1

Table 5B – Worker Wage Bill First Stage Regression (Companion Table to Table 4B)

	(1)	(2)	(3)	(4)
Table 5C: IV Regression First Stage	User Cost: State	User Cost: All	User Cost: State	User Cost: All
(Engineering Counts)				
Log(Assets)	0.0447**	0.0500*	0.0431**	0.0230
	(0.0215)	(0.0266)	(0.0208)	(0.0184)
Log(Years of Edu.)	0.539***	0.602***	0.539***	0.622***
	(0.0697)	(0.0811)	(0.0678)	(0.0542)

Log(R&D Capital)	0.0378 (0.0272)	0.0186 (0.0743)	0.0331 (0.0202)	0.0351 (0.0624)
Log(IT Workers)	0.388*** (0.0396)	0.359*** (0.0440)	0.369*** (0.0496)	0.310*** (0.0503)
Log(Researchers)	0.0154 (0.0337)	-0.0200 (0.0489)	0.0426 (0.0355)	0.0276 (0.0435)
Log(Firm User Cost)		1.840* (0.952)		1.339* (0.707)
Log(State User Cost)	-0.0314 (0.0247)	-0.0208 (0.0359)	-0.0636 (0.0634)	-0.0751 (0.0861)
Log(Firm*State)		1.723 (1.140)		1.186* (0.654)
Observations	16,252	12,450	16,246	12,441
Industry-Year FE	Yes	Yes	Yes	Yes
City FE	No	No	Yes	Yes

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 5C – Engineering Counts First Stage Regression (Companion Table to Table 4C)

6. TensorFlow, the Deep Learning Toolkit, and Bottlenecked Talent

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.²⁸ 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).²⁹

²⁸ 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...”

²⁹ Usually called “deep” when a standard neural net architecture has 4 or more layers.

TensorFlow was unique 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 training systems running on hundreds of specialized machines with thousands of 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). TensorFlow can 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.

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 abstraction layers like Keras (Chollet and others 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.³⁰

But was the TensorFlow launch 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

³⁰ Keras and TensorFlow are now implemented for R as well. PyTorch and TensorFlow both have another abstraction layer module called fast.ai which is gaining popularity. Its creators, Rachel Thomas and Jeremy Howard, frequently implement new technical advances into the fast.ai module. PyTorch remains a favorite package in the research community.

was trying to “attract developers and new hires to its technology”.³¹ With new technologies, especially open-source software packages, adoption dynamics and value creation can be highly sensitive to network effects (Von Hippel and Krogh 2003). 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, to date TensorFlow’s GitHub project has over 40,000 code commits, 20 branches, 68 releases, and over 1,600 contributors.

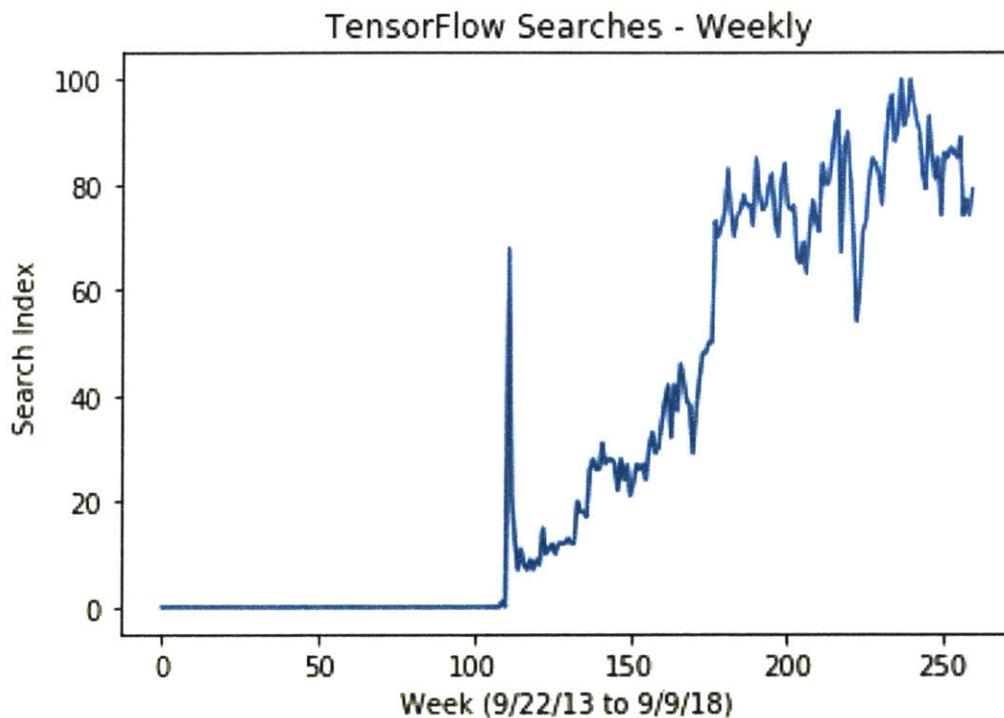


Figure 5 – TensorFlow Searches

Figure 5 shows the Google Trends searches for TensorFlow. Figures 6A and 7 show TensorFlow code and TensorBoard output, respectively. Given that TensorFlow constitutes a surprising reduction in the barriers to learn how to deep learn, it serves as an opportunity to understand

³¹ <https://bits.blogs.nytimes.com/2015/11/09/google-offers-free-software-in-bid-to-gain-an-edge-in-machine-learning/?mtrref=undefined>

between the demand for AI talent and the difficulty in building and implementing deep learning models.

Now I return to the case that a skill goes from being specialized to generalized in expectation. This is the context for TensorFlow, when returns to firm-specific assets are otherwise bottlenecked by a lack of available talent at attractive prices. In spite of the rapid growth in the deep learning-skilled community, it still may be the case that deep learning talent is capacity-constrained. Yet the change that the TensorFlow API, PyTorch, Keras, fast.ai, and similar packages introduce is the expectation that in the future these skills will be generalized instead of specialized. To test the market value effects of the democratization of deep learning via TensorFlow, I use a differences-in-differences continuous treatment specification of the following familiar form:

$$MV_{it} = \beta * \text{Total Assets}_{it} + \gamma * HK_{total_{it}} + \lambda_t + \eta_i + \nu * (POST_{TF} * AISKILL_{it}) + \mathbf{Z}'_{it} \boldsymbol{\delta} + \epsilon_{it} \quad (13)$$

Market value for firm i at time t is a function of total assets, human capital (as total education years), a time fixed effect, a firm fixed effect (which absorbs the firm's average AI skill level), and an interaction between the post-TensorFlow launch dummy and the cumulative AI skill counts for the firm in period t . \mathbf{Z} denotes the vector of additional cumulative skills indices and the AI skills index which might otherwise confound the analysis. I consider a number of related skills, and a handful of seemingly unrelated ones (e.g. advertising) to test the relationship between AI skills and market value following the launch of TensorFlow.³⁴ These regressions are pooled over the course of the year, testing for a structural break related to the acquisition of AI-related skills.

Table 6 has the results of the difference-in-difference analysis. Other than total assets and lagged market value, the index values from LinkedIn are all logged. This allows for an easier interpretation of the coefficients in percentage terms. These difference-in-difference estimates explain

³⁴ Google is excluded from the analysis for obvious reasons.

the Q-value created by firms with different levels of AI skills. Interestingly, the AI skills index on its own negatively predicts market value in the pre-TensorFlow period. AI skills have an economically and statistically significant relationship with market value. Each additional order of magnitude increase in AI skill for the firm is associated with an increase in market value of \$1.9 to \$2.2 billion in the post-TensorFlow period, with standard errors of approximately \$520 million. Including lagged market value (Table 6, Column 7) and fixing the sample to be a balanced panel from 2014-2017 increases the coefficient values by a meaningful factor to the upper end of estimates. Digital Literacy, a category including Microsoft Office and other standard computer skills, and Advertising skills indices are negatively correlated with market value as well. Business Management skills tend to be positively (albeit imprecisely) correlated with market value. Cloud computing skills are negatively correlated with market value, suggesting that lower market-to-book value firms may be investing more in cloud skills at higher rates adjusting for the other skills indices. As the Compustat sample overrepresents manufacturing firms relative to the entire U.S. economy, this is suggestive of smaller market-to-book manufacturing firms investing more in the cloud.

Multiplying the index value by a 1-year lead (Table 7) suggests that the market prices the value of commoditized AI talent at the end of 2015 without earlier anticipation. The 1 year lead coefficient is not statistically significant. In other words, a year prior to the launch of TensorFlow, AI skills increases predicted *no change* in market value. The coefficient is however attenuated, suggesting a need for more detailed analysis of pre-trends.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Table 6: Market Value - AI Skills Difference-in-Difference (MM USD)	AI Skills	+Data Science	+Cloud Computing	+Data Storage	+Digital Literacy	+Bus.Mgmt and Advertising	+Lagged MV	+Balanced Panel

Lagged Market Value							0.149*	
							(0.0761)	
Total Assets	1.099*** (0.0478)	1.099*** (0.0478)	1.099*** (0.0478)	1.099*** (0.0478)	1.099*** (0.0478)	1.099*** (0.0477)	0.957*** (0.0619)	1.122*** (0.0556)
Log(Edu. Years)	932.4 (639.0)	938.7 (633.2)	958.6 (637.2)	964.2 (638.2)	983.4 (629.2)	993.3 (635.7)	828.5 (526.6)	1,224* (738.7)
Log(AI Index)	- 2,058*** (685.1)	- 2,045*** (691.6)	- 1,988*** (677.1)	- 1,985*** (676.4)	- 1,980*** (677.7)	-1,970*** (672.8)	- 2,071*** (643.2)	-2,370*** (777.3)
Log(AI Index x Post TF)	1,900*** (460.7)	1,903*** (459.5)	1,916*** (462.5)	1,917*** (462.7)	1,921*** (462.0)	1,926*** (464.7)	1,998*** (442.9)	2,202*** (520.0)
Log(Data Science Index)		-124.7 (281.7)	-18.81 (290.8)	-6.796 (291.4)	146.7 (259.5)	67.07 (270.7)	-135.6 (278.0)	-105.2 (319.3)
Log(Cloud Computing Index)			-469.5** (231.8)	-458.1** (230.7)	-440.5* (231.2)	-465.5* (243.4)	-519.9** (242.5)	-535.2* (277.6)
Log(Data Storage Technology Index)				-11,102 (10,326)	-9,944 (10,546)	-10,023 (10,438)	-9,750 (9,877)	-20,331 (14,296)
Log(Digital Literacy Index)					-332.7 (293.8)	-435.8* (254.5)	-461.1* (275.1)	-555.8* (319.0)
Log(Bus. Management Index)						559.4 (537.5)	832.2 (676.5)	883.7 (713.5)
Log(Advertising Index)						-160.4 (182.5)	-195.9 (188.3)	-117.3 (218.2)
Observations	6,440	6,440	6,440	6,440	6,440	6,440	6,393	5,284
R-squared	0.998	0.998	0.998	0.998	0.998	0.998	0.999	0.998
Firm and Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 6: Market Value Difference-in-Difference on AI Skills During TensorFlow Launch

	(1)	(2)	(3)
Table 7: AI Difference-in-Difference Leads Robustness Check	AI Cluster	+Data Science	+Other Indices
Total Assets	0.987*** (0.0154)	0.987*** (0.0154)	0.988*** (0.0154)
Log(Edu. Years)	432.0 (554.9)	426.1 (554.4)	485.9 (569.5)
Log(AI Index)	-987.1 (785.8)	-996.9 (781.3)	-963.1 (768.0)
Log(AI Index x Post TF+1 Year Lead)	9.425 (440.5)	6.701 (441.8)	17.49 (447.0)
Log(AI Index x Post TF)	895.5*** (185.4)	893.5*** (186.2)	900.0*** (188.4)
Log(Data Science Index)		120.0 (208.6)	260.5 (212.7)
Log(Cloud Computing Index)			-339.2 (223.7)
Log(Data Storage Technology Index)			-5,548 (9,019)
Log(Digital Literacy Index)			-269.7 (184.0)
Log(Bus. Management Index)			74.10 (428.4)
Log(Advertising Index)			97.97 (161.6)

Observations	4,587	4,587	4,587
Firm and Year FE	Yes	Yes	Yes

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 7: Lead Check for AI Skill Difference-in-Difference

7. Further Robustness Checks

The launch of TensorFlow and other deep learning packages might have coincided with other conditions or endogenous firm decisions which limit any chance to make causal claims. So far, a causal interpretation of the difference-in-differences coefficients in the above tables assumes that market value does not cause adoption of AI, that there are parallel trends in the AI using and non-AI using firms, and the stable unit treatment value assumption (SUTVA) holds.³⁵ The first condition is mitigated by including lagged market value as a control, and the latter condition can be investigated partially with a balanced panel assuming no spillovers between firms. The parallel trends assumption is trickier and requires a more granular time series analysis. I create new group variables “AI Quintile” and “Suitability for Machine Learning (SML) Quintile” for the quintile groups in which each firm falls in AI skill employment and SML ranking (respectively) as of the fourth quarter of 2015. The median firm has no listed AI skills at that time. In fact, up to the 59th percentile of AI use the LinkedIn skill count for AI is zero.³⁶ In these specifications, I calculate the log indices of the input skills to recover a “percentage increase per quantile change” interpretation of the coefficients.³⁷ Using a balanced panel, I estimate the following specification, where total assets and education years are included in the \mathbf{Z} vector:

$$\text{Log}(MV_{it}) = \lambda_t + \eta_i + \nu_t * (D_t * AIQuantile_i) + \mathbf{Z}'_{it}\boldsymbol{\delta} + \epsilon_{it} \quad (14)$$

³⁵ This assumes that spillovers and compositional changes in the sample are not confounding.

³⁶ The appendix has a histogram of AI group percentages. The groups are uneven because of ties in AI skill counts.

³⁷ Technically each index is shifted by 1, so the log of market value variable is log(MV+1).

Each year is represented as a dummy D_t and interacted with the AI Quintile dummy variable. The results are in the Figure 8 below, showing an increase in market value of approximately 4-7% for all AI-using firms relative to non-AI firms (bottom, second, and part of the third quintile) in the quarter of the TensorFlow launch. This is the only quarter for which the change in market value by AI quintile is statistically significant for the second-highest AI-intensive quintile, and one of two for the third quintile. The quarters themselves are not statistically significantly different, but the pooled estimates from Table 8 suggest a structural break. The top AI using firms appear to have a non-parallel trend relative to non-AI firms in the pre-launch period. As seen in green, this high end of AI-using firms has a statistically significantly higher growth rate. This would invalidate the parallel trends assumption *for a quintile group-based* difference-in-difference estimate above and motivates the continuous skill value regression with firm fixed effects. The Appendix has the corresponding table of estimates, including interactions for SML quintile. Interestingly, the SML quintiles interacted with time dummies fail to reflect an effect of Tensorflow. Instead, the market value for those firms decreases later in the same year. This calls into question whether the TensorFlow (and AI talent) shock expectation is driving the downward trend in high SML firms. It is possible that this high SML revaluation occurs with a lag, or that another effect is driving the negative correlation between SML value and market value. Figure 9 shows the continuous treatment version of the effect on deep learning skills in particular. The effect precision narrows as more LinkedIn members post deep learning skills, but here again there is an increase in the market value of firms following the launch of TensorFlow in Q4 2015. This launch is statistically significant from zero AI use (the baseline), while all of the previous quarters are not statistically significantly different from the zero baseline. We fail to rule out parallel trends for the continuous treatment version, though it may be that the statistically significant market value pop due to deep learning skills was by chance. It is convincing, however, that the effects diminish to an imprecisely measured zero in the quarters following Q4 2015. Figures 10A-

D have data science skill, linear regression, management, and advertising skill index specifications (respectively) for comparison. All specifications include the full set of skill index covariates, firm fixed effects, industry-time fixed effects, lagged market value, education years, and total asset. The error bars are the 95% confidence interval using standard errors clustered by firm. The pooled and time-series regressions suggest strong market value effects on AI-using companies from the launch of TensorFlow and related packages. Of course, any shock affecting AI-using companies in the same quarter will also show up in the coefficient estimates. These concurrent unobserved shocks are a threat to any causal interpretation. At a minimum, however, it appears that there was a strong upward repricing of all companies employing AI talent at the same time that Google made the decision to make TensorFlow open-source. The evidence that this launch caused a decrease in the value of companies with high SML tasks is weaker unless the effects occurred with nearly a one-year lag. At the same time, the value of other types of skills do not seem to have responded to the TensorFlow launch.

These results would seem at first glance to stand in stark contrast to the previous section's results on the general value of engineering talent, but instead may indicate that the firm-specific assets already accumulated by companies are strongly complementary to the general deep learning talent pool. The \$1.9 to \$2.2 billion per LinkedIn AI skill order of magnitude increase predicted by the difference-in-difference analysis in Table 6 is unexpectedly large given that even high salaries for AI workers are typically less than \$2 million at the moment. Noticing this pre-trend in the top quintile of AI-using firms, I adjust the difference-in-difference analysis below to exclude the top quintile in Tables 8 and 9.

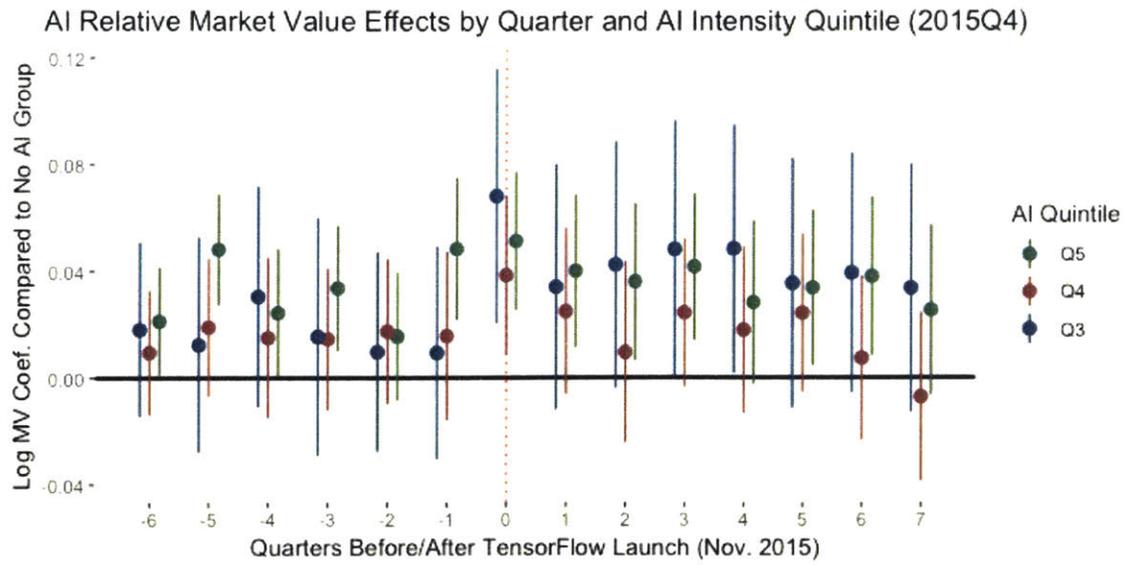


Figure 8 - AI Quintile Effects

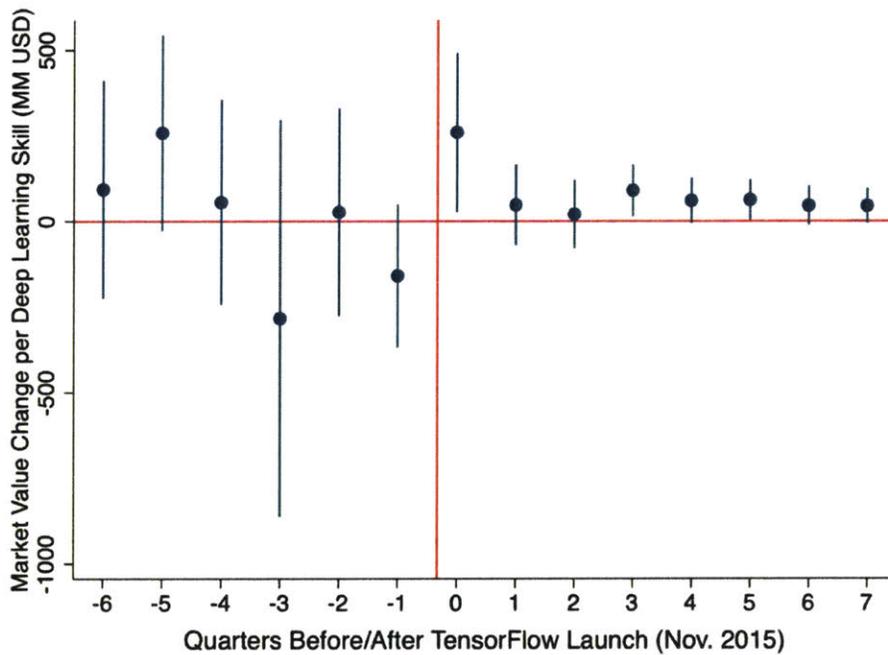


Figure 9 – Market Value Change per Deep Learning Skill

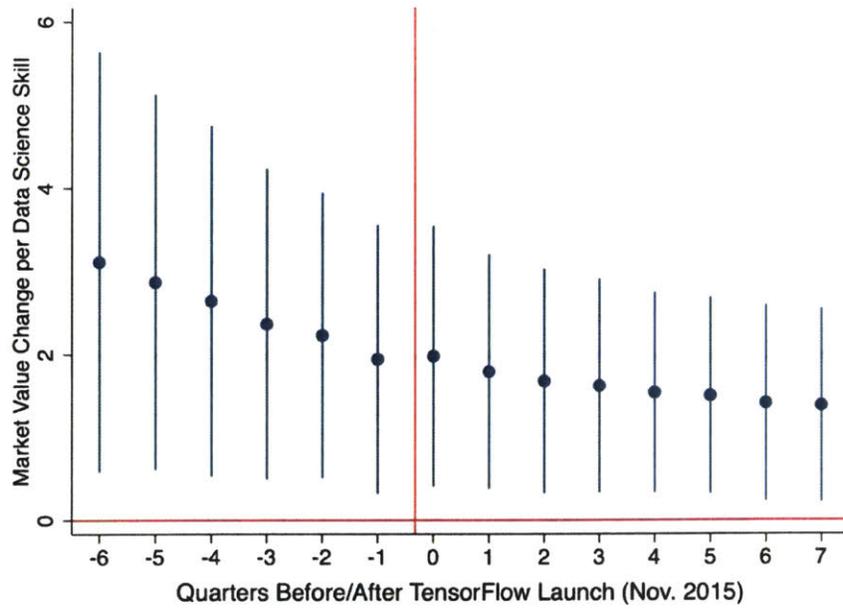


Figure 10A - Data Science Skills

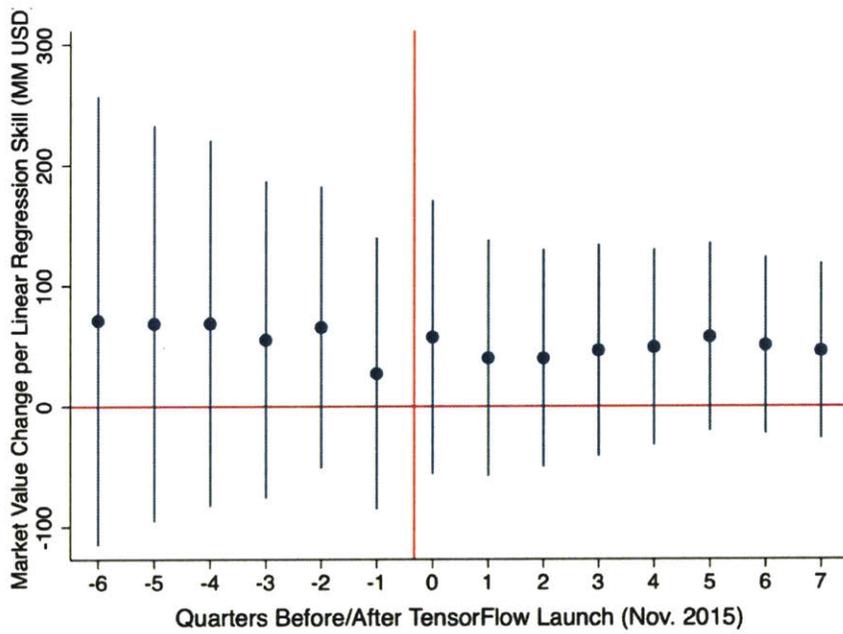


Figure 10B – Linear Regression Skills

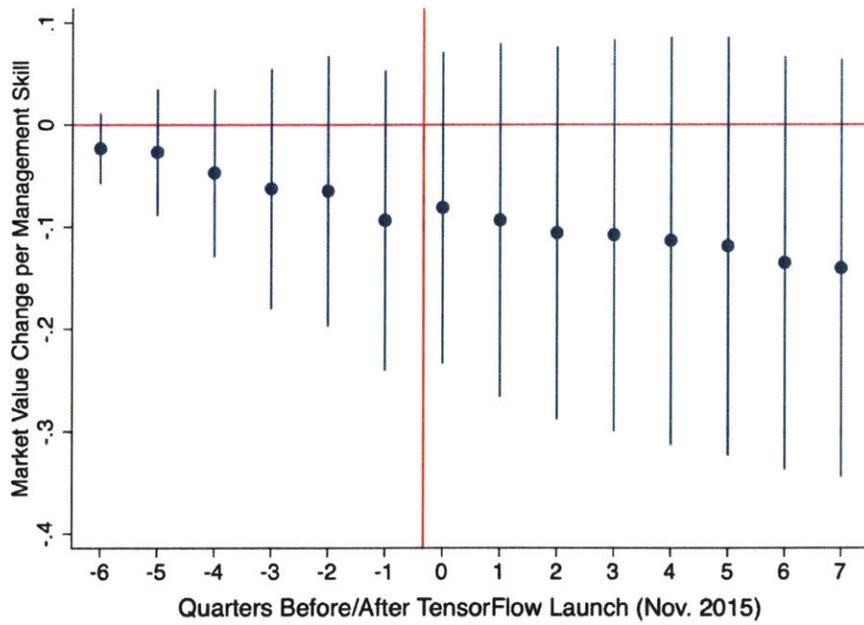


Figure 10C – Management Skills

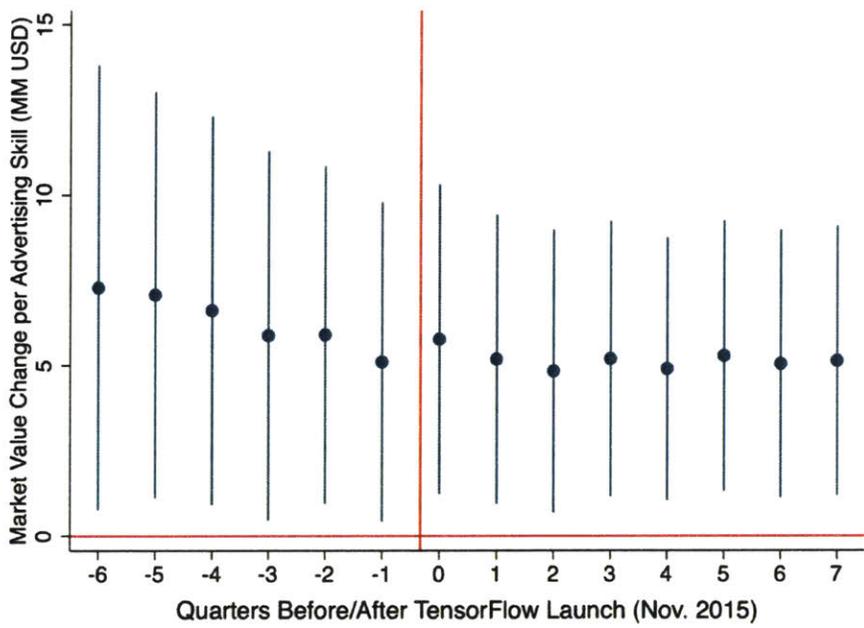


Figure 10D – Advertising Skills

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Table 8: AI Difference-in-Difference w/o Top Quintile	AI Cluster	+Data Science	+Cloud Computing	+Data Storage	+Digital Literacy	+Bus.Mgmt and Advertising	+Lagged MV	+Balanced Panel
Market Value (MM USD)								
Lagged Market Value							0.00473	
							(0.0240)	
Total Assets	0.990*** (0.0180)	0.990*** (0.0180)	0.990*** (0.0180)	0.990*** (0.0180)	0.990*** (0.0180)	0.990*** (0.0180)	0.986*** (0.0209)	1.000*** (0.0164)
Log(Edu. Years)	299.0** (130.0)	306.1** (130.9)	309.1** (131.1)	313.2** (131.0)	322.7** (133.3)	318.1** (132.7)	323.0** (136.3)	360.0** (151.9)
Log(AI Index)	-293.0 (196.4)	-278.5 (197.6)	-272.6 (198.8)	-270.1 (198.9)	-266.9 (198.5)	-269.7 (198.7)	-278.4 (199.4)	-277.7 (222.5)
Log(AI Index x Post TF)	352.3*** (136.2)	357.4*** (136.5)	360.3*** (136.5)	360.3*** (136.5)	361.6*** (136.5)	358.3*** (136.6)	366.1*** (138.6)	356.7** (144.3)
Log(Data Science Index)		-152.9 (130.0)	-139.3 (126.4)	-129.9 (125.4)	-51.85 (119.0)	-8.269 (132.0)	2.779 (139.0)	-99.55 (150.4)
Log(Cloud Computing Index)			-60.82 (137.7)	-51.47 (138.7)	-42.69 (138.0)	-28.63 (140.7)	-39.55 (140.9)	-57.00 (158.1)
Log(Data Storage Technology Index)				-8,939* (5,230)	-8,350 (5,309)	-8,373 (5,245)	-8,481 (5,245)	-14,033* (7,491)

Log(Digital Literacy Index)						-171.4	-117.9	-129.1	-217.3
						(136.7)	(135.4)	(139.0)	(169.4)
Log(Bus. Management Index)							-273.9	-271.5	-224.4
							(207.5)	(223.2)	(255.4)
Log(Ad. Index)							56.82	51.24	149.4
							(126.2)	(126.2)	(137.2)
Observations	5,864	5,864	5,864	5,864	5,864	5,864	5,864	5,819	4,784
Firm and Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 8: Adjusted Difference-in-Differences AI Skill Valuation Post-TensorFlow Ex. Top Quintile of AI-Using Firms

The coefficients excluding the top AI-using quintile are substantially attenuated. Within the remaining AI-using firms as of the end of Q4 2015, the effect of the TensorFlow launch is that a 100% increase in AI skills corresponds to a \$352 to \$367 million dollar increase in market value with standard errors of approximately \$140 million. This is still a substantial amount, and qualitatively similar given that many of the top AI-using firms are the largest firms. The effect of the TensorFlow launch survives excluding the AI superstar firms with pre-trends. Similarly, Table 9 has the leads analysis corresponding to Table 7. The 1 year leads fail to indicate any kind of anticipation, though in this set of specifications the ordinary effects are not detected either.

	(1)	(2)	(3)
Table 9: AI Difference-in-Difference Leads w/o Top AI Quintile	AI Cluster	+Data Science	+Other Indices

Total Assets	0.974***	0.974***	0.974***
	(0.0144)	(0.0144)	(0.0144)
Log(Edu. Years)	-26.97	-28.39	-15.62
	(190.1)	(191.0)	(190.9)
Log(AI Index)	-367.4	-369.5	-366.1
	(269.2)	(268.4)	(269.0)
Log(AI Index x Post TF+1 Year Lead)	102.2	100.9	99.71
	(156.0)	(156.2)	(155.8)
Log(AI Index x Post TF)	51.93	51.20	50.53
	(125.5)	(125.9)	(126.4)
Log(Data Science Index)		33.97	133.3
		(121.6)	(108.4)
Log(Cloud Computing Index)			0.879
			(146.2)
Log(Data Storage Technology Index)			-3,625
			(5,751)
Log(Digital Literacy Index)			-119.5
			(142.9)
Log(Bus. Management Index)			-202.3
			(198.3)
Log(Advertising Index)			81.55
			(138.6)
Observations	4,187	4,187	4,187
Firm and Year FE	Yes	Yes	Yes

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Taken together, these results suggest story like that of Hirshleifer's pecuniary benefits of technological change. Having AI talent in 2015 is a signal that there are other assets at the firm that are

complementary to AI talent. TensorFlow makes a specialist skill into a generalist skill. As a result, more workers can build engineering value using commodity deep learning packages, and firms receive the capital service flow from their assets (for example, large scale databases or automation projects) formerly bottlenecked by the scarcity of available talent. TensorFlow, by making a scarce complement cheaper and more abundant, increases the value of firms positioned to invest in AI. Importantly, these effects are democratizing for at least one quarter. AI-using firms without pre-trends catch up to their AI superstar competitors for the fourth quarter of 2015. Afterward, their market value performance is statistically similar to peer companies without AI talent (as in Figure 9). Further, no similar effects can be observed in other types of skills where we might expect to see similar changes. Data Science, Linear Regression, Advertising, and Management fail to demonstrate any of the effects present for Deep Learning and AI skills more broadly. This suggests that the repricing of AI companies was indeed generated by the launch of TensorFlow.

8. Suitability for Machine Learning Analysis

Independent of AI talent, what happened to companies where the employed workers might be expected to be impacted by machine learning technology? Tables 10 and 11 show the results of the same analysis run for Suitability for Machine Learning (SML) scores aggregated to the firm level (wage bill-weighted averages) and in logged terms (respectively). Notably the launch of TensorFlow has a statistically significant negative association with the value of firms with higher potential to automate tasks with machine learning. A 1% increase in the overall SML of a firm is correlated with a 0.5% decrease in the firm's market value post-2016. This is consistent with the idea that the assets complementary to machine learning engineering are valuable, but potentially productivity-enhancing (not profitability) innovations without a source of rents might force firms to invest more rapidly than they would have otherwise done. Convex ex-ante fixed costs of investment therefore might drive down market value for firms that have to change their business models. Figure 11 shows the correlation between market value and SML for 2016, which is very close to zero.

Asset managers need not have the granular detail of the SML scores for these price effects to be incorporated. Since high SML tasks tend to be clerical and data-intensive routine work, the firms employing lots of these types of workers with intangible assets optimized for these purposes might newly be vulnerable to ML and AI-powered competition. However, these effects are likely not driven by TensorFlow. Interacting the quarter dummies with the SML quintile reveals that the structural change is indeed in the same year, but the repricing does not occur in the same quarter as the TensorFlow launch. Under the assumption that markets efficiently price the effects of TensorFlow at the time of release, it is implausible to conclude that SML-intensive companies lost market value as a result of the open sourcing event.

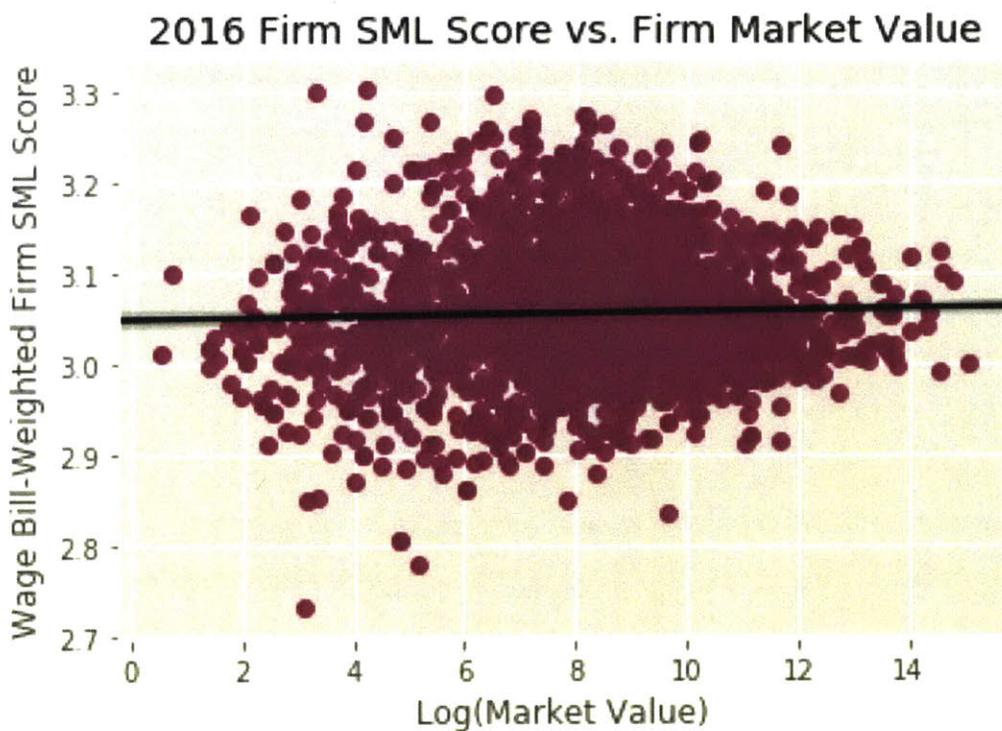


Figure 11: Firm-Level Suitability for Machine Learning (SML) vs. Log Market Value

	(1)	(2)	(3)	(4)	(5)	(6)
Table 10: Market Value (SML) Difference- in-Differences	SML	+Data Science	+Cloud Computing	+Data Storage	+Digital Literacy	+Bus.Mgmt and Advertising
Total Assets	1.062*** (0.0287)	1.099*** (0.0446)	1.100*** (0.0447)	1.100*** (0.0447)	1.078*** (0.0384)	1.043*** (0.0302)
Total Education Years	0.0185 (0.0183)	0.0190 (0.0177)	0.0192 (0.0179)	0.0192 (0.0179)	0.0210 (0.0176)	0.0137 (0.00933)
SML X Post-TensorFlow	-4,336* (2,257)	-3,684 (2,791)	-3,667 (2,803)	-3,661 (2,790)	-6,477** (2,821)	-5,920** (2,476)
Data Science Skill Index		0.503 (0.310)	0.365 (0.402)	0.365 (0.402)	-0.870* (0.483)	1.584** (0.769)
Cloud Computing Skill Index			0.261 (0.761)	0.261 (0.761)	0.213 (0.690)	0.0218 (0.432)
Data Storage Skill Index				1,534 (10,438)	6,843 (10,318)	-1,527 (8,721)
Digital Literacy Skill Index					0.498*** (0.0858)	0.360*** (0.130)
Business Mgmt. Skill Index						-0.717*** (0.198)
Advertising Skill Index						4.523*** (1.586)
Observations	8,764	6,437	6,437	6,437	6,437	6,437
Firm and Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Robust standard errors in parentheses, SEs clustered by Firm

*** p<0.01, ** p<0.05, * p<0.1

Table 10: Suitability for Machine Learning (SML) Difference-in-Differences

(1) (2) (3) (4) (5) (6)

Table 11: SML +Data +Cloud +Data +Digital +Bus.Mgmt and
 Log Market Value – SML Science Computing Storage Literacy Advertising
 Difference-in-Differences

	SML	+Data Science	+Cloud Computing	+Data Storage	+Digital Literacy	+Bus.Mgmt and Advertising
Log(Total Assets)	0.607*** (0.0660)	0.584*** (0.0849)	0.583*** (0.0849)	0.584*** (0.0850)	0.582*** (0.0851)	0.582*** (0.0851)
Log(Education Years)	0.0232 (0.0153)	0.0307 (0.0207)	0.0316 (0.0207)	0.0311 (0.0207)	0.0296 (0.0206)	0.0290 (0.0206)
Log(SMLxPost-TF)	-0.484** (0.235)	-0.520** (0.257)	-0.514** (0.256)	-0.506** (0.256)	-0.525** (0.255)	-0.499* (0.255)
Log(Data Science)		-0.00454 (0.0187)	0.000376 (0.0190)	-0.000913 (0.0193)	-0.0150 (0.0188)	-0.0123 (0.0196)
Log(Cloud Computing)			-0.0154 (0.0109)	-0.0166 (0.0108)	-0.0187* (0.0106)	-0.0183* (0.0105)
Log(Data Storage Tech.)				1.036 (0.871)	0.925 (0.857)	0.945 (0.868)
Log(Digital Literacy)					0.0299* (0.0160)	0.0344** (0.0156)
Log(Business Management)						-0.0292 (0.0237)
Log(Advertising)						0.0136 (0.0165)
Observations	8,764	6,437	6,437	6,437	6,437	6,437
Firm and Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Robust standard errors in parentheses, SEs Clustered by Firm

*** p<0.01, ** p<0.05, * p<0.1

Table 11: Market Value and Log(SML) Difference-in-Difference

9. Conclusion

It remains an open and context-specific question whether the market value of firms is driven by appropriation of the value of technological human capital. Not all varieties of technological human capital are correlated with market value after controlling for generalized education level and the firm's asset base. Engineering talent, however, is. Using a panel of corporate fixed assets and human capital, I have measured the average and marginal returns to investments in technological labor. I find that on average, engineering talent is strongly correlated with market value, but the marginal causal effect of hiring more engineers as estimated by instrumental variables analyses and controlling for time-invariant firm-specific factors is indistinguishable from zero. Nevertheless, exciting growth in market value can occur when formerly specialized skillsets are converted into general ones, as is the case with Google's launch of TensorFlow. The reduced barriers to entry in AI led to growth of roughly \$3.56 million per 1% increase in AI skill excluding top quintile AI firms. This corresponds to a 4-7% contemporaneous increase in market value for firms outside the top quintile of AI skills with the launch of TensorFlow. This is the case even controlling for growth in other kinds of related skillsets, like Cloud Computing or Data Science. Further, Suitability for Machine Learning (SML) is unlikely to be a pathway via which TensorFlow positively impacted market value. If higher SML companies increased in value following the launch of TensorFlow, it would indicate that in expectation companies currently employing lots of high SML labor might appropriate the productivity gains for automating tasks with machine learning in the future. There is little evidence that the SML scores are positively related to market value following the launch of TensorFlow.

This kind of discrete technological change is informative about the processes which limit diffusion of new general-purpose technologies before they become generally-applied technologies. Namely, the tools or training have to be available to the engineers who are to build the assets that generate the technology's value. The choice to make specialized technology with large potential more widely available is sometimes one that can be made by corporate actors in industry. Talent is not

always a bottleneck, but when it is firms may be more likely to designate firm-specific tasks with high marginal value to workers with scarce skills. This makes it more difficult for competitors to bid up the wages of those types of workers, but at the same time more abundant skills are likely to be competitively priced. This combination of firm-specific assets with complementary applications of engineering skills means that firms can appropriate some of their employees' investments in human capital. The paradox that, controlling for non-human capital assets, technological talent can have high value on average but marginally low value is resolved when firms can assign tasks which their competitors do not value (on average) but do value (on the margin).

Managers expecting to pay all incoming workers the same amount as their incumbent staff are faced with a challenge when talent is scarce. Do they want to give everyone a raise just to hire one more person? When competition on the bases of wages is difficult, the assignment of firm-specific tasks is a potential mechanism to bargain away part of the employees' talent value. Managers therefore force workers to compete with each other inside the firm while insulating their employer from outside competition. It is therefore potentially lucrative to expand the available talent pool to capture the full possible value of firm-specific tasks. TensorFlow made expectations of future deep learning talent, a previously scarce skillset, much higher. This suggests that all AI employers expected to find more AI talent in the coming years, permitting assignment of a greater range of firm-specific tasks. This suggests, for example, that open-source production decisions may generate rents for adopting firms via the talent channel. Many previous studies of employer power in the labor market have focused on employer-occupation-market concentration or policy changes changing the bargaining power of employees. Yet if firms are benefitting from exercise of labor market monopsony power, it should show up in their valuations. This paper shows that for a specific type of technological talent – engineers – it can be the case that firm value is in part driven by employer appropriation of employee human capital. Companies can do this by allocating their employees to firm-specific job tasks and finding an employment niche. Meanwhile, workers should carefully consider their contractual

arrangements and how their employers are engineering value. In the case that work tasks are not competitively decided, the distribution of value might not reward the employee for their full contribution.

References

- Abadi, Martin, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. 2016. "Tensorflow: A System for Large-Scale Machine Learning." In *OSDI*, 16:265–83.
- Acemoglu, D., and J.-S. Pischke. 1998. "Why Do Firms Train? Theory and Evidence." *The Quarterly Journal of Economics* 113 (1):79–119.
- Acemoglu, Daron. 1997. "Training and Innovation in an Imperfect Labour Market." *The Review of Economic Studies* 64 (3). Wiley-Blackwell:445–64.
- Acemoglu, Daron, and David Autor. 2011. *Skills, Tasks and Technologies: Implications for Employment and Earnings. Handbook of Labor Economics*. Vol. 4.
- Acemoglu, Daron, and Jorn-Steffen Pischke. 1999. "Beyond Becker: Training in Imperfect Labour Markets." *The Economic Journal* 109 (453):112–42.
- Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb. 2017. "Prediction, Judgment, and Uncertainty." In *Economics of Artificial Intelligence*. University of Chicago Press.
- . 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston: Harvard Business Review Press.
- Agrawal, Ajay, Avi Goldfarb, and Florenta Teodoridis. 2016. "Understanding the Changing Structure of Scientific Inquiry." *American Economic Journal: Applied Economics* 8 (1):100–128.
- Andrews, Dan, Chiara Criscuolo, and Peter N Gal. 2015. "Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence from OECD Countries." *Working Paper*, 39.
- Ashenfelter, Orley C, Henry Farber, and Michael R Ransom. 2010. "Labor Market Monopsony." *Journal of Labor Economics* 28 (2). The University of Chicago Press:203–10.
- Autor, David H. 2014. "Polanyi's Paradox and the Shape of Employment Growth." *Working Paper*, 129–78.
- Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5):1553–97.
- Autor, David, Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4):1279–1333.
- Azar, José A, Ioana Marinescu, Marshall I Steinbaum, and Bledi Taska. 2018. "Concentration in US Labor Markets: Evidence from Online Vacancy Data."
- Balasubramanian, Natarajan, Jin Woo Chang, Mariko Sakakibara, Jagadeesh Sivadasan, and Evan Starr. 2018. "Locked In? The Enforceability of Covenants Not to Compete and the Careers of High-Tech Workers."

- Barney, Jay B. 1986. "Strategic Factor Markets: Expectations, Luck, and Business Strategy." *Management Science* 32 (10). INFORMS:1231–41.
- Becker, Gary S. 1962. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy* 70 (5, Part 2):9–49.
- Benmelech, Efraim, Nittai K Bergman, and Amit Seru. 2011. "Financing Labor."
- Benmelech, Efraim, Nittai Bergman, and Hyunseob Kim. 2018. "Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?"
- Benzell, Seth, Guillermo Lagarda, and Daniel Rock. 2018. "Skill-Biased Technological Change Within and Across Firms: 2000-2016."
- Bhaskar, Venkataraman, Alan Manning, and Ted To. 2002. "Oligopsony and Monopsonistic Competition in Labor Markets." *Journal of Economic Perspectives* 16 (2):155–74.
- Bivens, Josh, and Lawrence Mishel. 2015. "Understanding the Historic Divergence between Productivity and a Typical Worker's Pay: Why It Matters and Why It's Real." *Washington DC: Economic Policy Institute (Instituto de Políticas Económicas)*. [Http://Www. Epi. Org/Publication/Understanding-the-Historic-Divergence-Betweenproductivity-and-a-Typical-Workers-Pay-Why-It-Matters-and-Why-Its-Real](http://www.epi.org/publication/Understanding-the-Historic-Divergence-Between-productivity-and-a-Typical-Workers-Pay-Why-It-Matters-and-Why-Its-Real).
- Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen. 2018. "What Drives Differences in Management Practices?" *American Economic Review*.
- Bresnahan, Timothy. 2010. "General Purpose Technologies." *Handbook of the Economics of Innovation* 2 (1):761–91.
- Bresnahan, Timothy F., and M. Trajtenberg. 1995. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics* 65 (1):83–108.
- Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *The Quarterly Journal of Economics* 117 (1). MIT Press:339–76.
- Brynjolfsson, Erik. 1994. "Information Assets, Technology and Organization." *Management Science* 40 (12). INFORMS:1645–62.
- Brynjolfsson, Erik, Lorin M. (Lorin Moultrie) Hitt, and Shinkyu Yang. 2002. "Intangible Assets: Computers and Organizational Capital." *Brookings Papers on Economic Activity* 2002 (1):137–98.
- Brynjolfsson, Erik, Xiang Hui, and Meng Liu. 2018. "Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform."
- Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company.
- Brynjolfsson, Erik, and Tom Mitchell. 2017. "What Can Machine Learning Do? Workforce Implications." *Science* 358 (6370):1530–34.
- Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018a. "Machine Learning and Occupational Change." *Unpublished Working Paper. MIT*.

- . 2018b. “What Can Machines Learn, and What Does It Mean for Occupations and the Economy?” *AEA Papers and Proceedings*, 43–47.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2017. “Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics.” In *Economics of Artificial Intelligence*. University of Chicago Press.
- . 2018. “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies.”
- Brynjolfsson, Erik, Daniel Rock, Prasanna Tambe, and Jacqueline Barrett. 2018. “People Are Our Greatest Asset: Measuring the Market Value of Human Capital.” MIT Working Paper.
- Campbell, Benjamin A, Russell Coff, and David Kryscynski. 2012. “Rethinking Sustained Competitive Advantage from Human Capital.” *Academy of Management Review* 37 (3). Academy of Management Briarcliff Manor, NY:376–95.
- Campbell, Benjamin A, Martin Ganco, April M Franco, and Rajshree Agarwal. 2012. “Who Leaves, Where to, and Why Worry? Employee Mobility, Entrepreneurship and Effects on Source Firm Performance.” *Strategic Management Journal* 33 (1). Wiley Online Library:65–87.
- Chollet, François, and others. 2015. “Keras: Deep Learning Library for Theano and Tensorflow.” URL: <https://keras.io/>K 7 (8).
- Cockburn, Iain M, Rebecca Henderson, and Scott Stern. 2018. “The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis.” *National Bureau of Economic Research* No. w24449 (September).
- Coff, Russell, and Joseph Raffiee. 2015. “Toward a Theory of Perceived Firm-Specific Human Capital.” *Academy of Management Perspectives* 29 (3). Academy of Management Briarcliff Manor, NY:326–41.
- Crook, T Russell, Samuel Y Todd, James G Combs, David J Woehr, and David J Ketchen Jr. 2011. “Does Human Capital Matter? A Meta-Analysis of the Relationship between Human Capital and Firm Performance.” *Journal of Applied Psychology* 96 (3). American Psychological Association:443.
- Eisfeldt, Andrea L., and Dimitris Papanikolaou. 2013. “Organization Capital and the Cross-Section of Expected Returns.” *Journal of Finance* 68 (4):1365–1406.
- . 2014. “The Value and Ownership of Intangible Capital.” *American Economic Review* 104 (5):189–94.
- Ewens, Michael, and Matt Marx. 2017. “Founder Replacement and Startup Performance.” *The Review of Financial Studies* 31 (4). Oxford University Press:1532–65.
- Ewens, Michael, Ramana Nanda, and Matthew Rhodes-Kropf. 2018. “Cost of Experimentation and the Evolution of Venture Capital.” *Journal of Financial Economics* 128 (3). Elsevier:422–42.
- Felten, Edward W., Manav Raj, and Robert Seamans. 2018. “A Method to Link Advances in Artificial Intelligence to Occupational Abilities.” *AEA Papers and Proceedings* 108:54–57.
- Furman, Jason, and Robert Seamans. 2018. “AI and the Economy,” 1–33. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3186591.

- Hall, Bronwyn H. 1993. "The Stock Market's Valuation of R&D Investment during the 1980's." *The American Economic Review* 83 (2):259–64.
- . 2006. "R&D, Productivity, and Market Value." *IFS Working Papers, Institute for Fiscal Studies (IFS) 06/23*.
- Hall, R. E. 2004. "Measuring Factor Adjustment Costs." *The Quarterly Journal of Economics* 119 (3):899–927.
- Hall, Robert E. 2001. "The Stock Market and Capital Accumulation." *The American Economic Review* 91 (5):1185–1202.
<http://www.jstor.org/stable/2677921>
<http://www.jstor.org/stable/pdfplus/2677921.pdf?acceptTC=true>.
- Hall, Robert E. 2017. "High Discounts and High Unemployment." *American Economic Review* 107 (2):305–30.
- Hart, Oliver, and John Moore. 1994. "A Theory of Debt Based on the Inalienability of Human Capital." *The Quarterly Journal of Economics* 109 (4). MIT Press:841–79.
- Hayashi, Fumio. 1982. "Tobin's Marginal q and Average q: A Neoclassical Interpretation." *Econometrica* 50 (1):213–24.
- Hayashi, Fumio, and Tohru Inoue. 1991. "The Relation Between Firm Growth and Q with Multiple Capital Goods: Theory and Evidence from Panel Data on Japanese Firms." *Econometrica* 59 (3):731–53.
- Hippel, Eric von, and Georg von Krogh. 2003. "Open Source Software and the 'Private-Collective' Innovation Model: Issues for Organization Science." *Organization Science* 14 (2). INFORMS:209–23.
- Hirshleifer, Jack. 1978. "The Private and Social Value of Information and the Reward to Inventive Activity." In *Uncertainty in Economics*, 541–56. Elsevier.
- Jeffers, Jessica S. 2017. "The Impact of Restricting Labor Mobility on Corporate Investment and Entrepreneurship." *Job Market Paper*.
- Jones, Benjamin F. 2009. "The Burden of Knowledge and the 'Death of the Renaissance Man': Is Innovation Getting Harder?" *The Review of Economic Studies* 76 (1). Wiley-Blackwell:283–317.
- Jovanovic, Boyan. 1979. "Firm-Specific Capital and Turnover." *Journal of Political Economy* 87 (6). The University of Chicago Press:1246–60.
- Kaldor, Nicholas. 1966. "Marginal Productivity and the Macro-Economic Theories of Distribution: Comment on Samuelson and Modigliani." *The Review of Economic Studies* 33 (4). JSTOR:309–19.
- Kerr, Sari Pekkala, William R Kerr, and William F Lincoln. 2015. "Skilled Immigration and the Employment Structures of US Firms." *Journal of Labor Economics* 33 (S1). University of Chicago Press Chicago, IL:S147--S186.
- Kerr, William R. 2013. "U.S. High-Skilled Immigration, Innovation, and Entrepreneurship: Empirical Approaches and Evidence." *Nber Working Paper Series* 53:1689–99.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521 (7553):436–

44.

- LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. "Gradient-Based Learning Applied to Document Recognition." *Proceedings of the IEEE* 86 (11):2278–2323.
- Lippman, Steven A, and Richard P Rumelt. 1982. "Uncertain Imitability: An Analysis of Interfirm Differences in Efficiency under Competition." *The Bell Journal of Economics*. JSTOR, 418–38.
- Lucas, Robert E. 1967. "Adjustment Costs and the Theory of Supply." *Journal of Political Economy* 75 (4, Part 1):321–34.
- Lucking, Brian. 2019. "Do R&D Tax Credits Create Jobs ?"
- Lucking, Brian, Nicholas Bloom, and John Van Reenen. 2017. "Have R&D Spillovers Changed?" *SSRN*.
- Lustig, Hanno, Chad Syverson, and Stijn Van Nieuwerburgh. 2011. "Technological Change and the Growing Inequality in Managerial Compensation." *Journal of Financial Economics* 99 (3). Elsevier:601–27.
- Marx, Matt. 2011. "The Firm Strikes Back: Non-Compete Agreements and the Mobility of Technical Professionals." *American Sociological Review* 76 (5). Sage Publications Sage CA: Los Angeles, CA:695–712.
- Marx, Matt, Deborah Strumsky, and Lee Fleming. 2009. "Mobility, Skills, and the Michigan Non-Compete Experiment." *Management Science* 55 (6):875–89.
- Mas, Alexandre, and Amanda Pallais. 2017. "Valuing Alternative Work Arrangements." *American Economic Review* 107 (12):3722–59.
- Moretti, Enrico. 2004. "Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data." *Journal of Econometrics*.
- Paszke, Adam, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. "Automatic Differentiation in Pytorch."
- Peri, Giovanni, Kevin Shih, and Chad Sparber. 2015. "STEM Workers, H-1B Visas, and Productivity in US Cities." *Journal of Labor Economics* 33 (S1). University of Chicago Press Chicago, IL:S225--S255.
- Peters, Ryan H., and Lucian A. Taylor. 2017. "Intangible Capital and the Investment-q Relation." *Journal of Financial Economics* 123 (2):251–72.
- Polanyi, Michael. 1966. "The Logic of Tacit Inference." *Philosophy* 41 (155):1–18.
- Raffiee, Joseph, and Russell Coff. 2016. "Micro-Foundations of Firm-Specific Human Capital: When Do Employees Perceive Their Skills to Be Firm-Specific?" *Academy of Management Journal* 59 (3). Academy of Management Briarcliff Manor, NY:766–90.
- Raj, Manav, and Robert Seamans. 2018. "AI, Labor, Productivity, and the Need for Firm-Level Data." In *Economics of Artificial Intelligence*. University of Chicago Press.
- Roach, Michael, and Henry Sauermann. 2010. "A Taste for Science? PhD Scientists' Academic Orientation and Self-Selection into Research Careers in Industry." *Research Policy*.

- Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. 1986. "Learning Representations by Back-Propagating Errors." *Nature* 323 (6088):533–36.
- Saunders, Adam, and Erik Brynjolfsson. 2016. "Valuing Information Technology Related Intangible Assets." *Mis Quarterly* 40 (1).
- Saunders, Adam, and Prasanna Tambe. 2015. "Data Assets and Industry Competition: Evidence from 10-K Filings," no. September.
- Simon, Herbert A. 1991. "Organizations and Markets." *The Journal of Economic Perspectives* 5 (2):25–44.
- Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2015. "Firming Up Inequality." *NBER Working Paper*.
- Stansbury, Anna M, and Lawrence H Summers. 2017. "Productivity and Pay: Is the Link Broken?"
- Starr, Evan, Natarajan Balasubramanian, and Mariko Sakakibara. 2017. "Screening Spinouts? How Noncompete Enforceability Affects the Creation, Growth, and Survival of New Firms." *Management Science* 64 (2). INFORMS:552–72.
- Stern, Scott. 2004. "Do Scientists Pay to Be Scientists?" *Management Science* 50 (6). INFORMS:835–53.
- Stole, Lars A, and Jeffrey Zwiebel. 1996a. "Intra-Firm Bargaining under Non-Binding Contracts." *The Review of Economic Studies* 63 (3). Wiley-Blackwell:375–410.
- . 1996b. "Organizational Design and Technology Choice under Intrafirm Bargaining." *The American Economic Review*. JSTOR, 195–222.
- Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2):326–65.
- Tambe, Prasanna. 2014. "Big Data Investment, Skills, and Firm Value." *Management Science* 60 (6):1452–69.
- Tambe, Prasanna, Lorin and Hitt, Erik Brynjolfsson, and Daniel Rock. 2018. "The Price and Quantity of IT-Related Intangible Capital."
- Tambe, Prasanna, and Lorin M. Hitt. 2012. "The Productivity of Information Technology Investments: New Evidence from IT Labor Data." *Information Systems Research*.
- Tambe, Prasanna, Lorin M. Hitt, and Erik Brynjolfsson. 2012. "The Extroverted Firm: How External Information Practices Affect Innovation and Productivity." *Management Science* 58 (5):843–59.
- Tambe, Prasanna, Xuan Ye, and Peter Cappelli. 2019. "Paying to Program? Engineering Brand and High-Tech Wages." 25552. National Bureau of Economic Research Working Paper Series. <https://www.nber.org/papers/w25552.pdf>.
- Teodoridis, Florenta. 2017. "Understanding Team Knowledge Production: The Interrelated Roles of Technology and Expertise." *Management Science*. INFORMS.
- Tervio, Marko. 2008. "The Difference That CEOs Make: An Assignment Model Approach." *American Economic Review* 98 (3):642–68.
- Thompson, Neil. 2017. "The Economic Impact of Moore's Law: Evidence from When It Faltered."

Tobin, James, William C Brainard, and others. 1976. "Asset Markets and the Cost of Capital."

White, B. W., and Frank Rosenblatt. 1963. "Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms." *The American Journal of Psychology* 76 (4):705.

Wildasin, David E. 1984. "The q Theory of Investment with Many Capital Goods." *American Economic Review*.

Yang, S, and E Brynjolfsson. 2001. "Intangible Assets and Growth Accounting: Evidence from Computer Investments." *Unpublished Paper. MIT* 85 (61):28.

Zyontz, Samantha. 2018. "Running with (CRISPR) Scissors: Tool Specialists and Innovation with New Technologies."

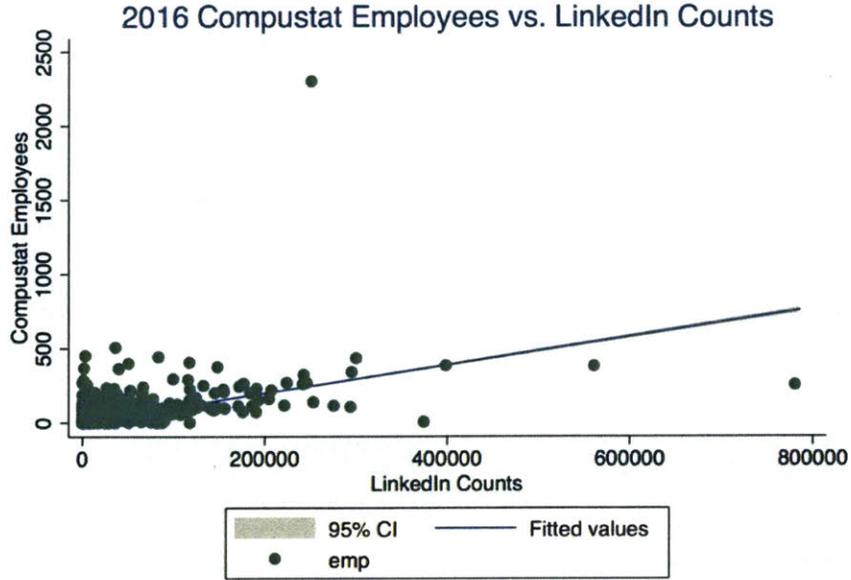
Appendix: Additional Regression Results on Coverage and Robustness

	(1)
LinkedIn Coverage	Compustat Count (Thousands)
LinkedIn Count	0.00190*** (1.59e-05)
Total Assets	4.27e-05*** (1.70e-06)
Observations	52,767
R-squared	0.422
Industry-Time FE	Yes

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note: Coefficients represent the predicted per LinkedIn user employee count (in thousands) for Compustat firms with employee count data populated. This regression is used to predict the employee count in the case that Compustat is missing data.



	(1)	(2)
Log(Market Value)	SML Quintiles	AI Quintiles
Log(Lagged Market Value)	0.646*** (0.0169)	0.645*** (0.0170)
Log(Total Assets)	0.305*** (0.0181)	0.307*** (0.0181)
Log(Education Years)	-0.00521 (0.00840)	-0.00500 (0.00840)
Log(Business Mgmt.)	-0.0128 (0.00823)	-0.0128 (0.00831)
Log(Cloud Computing)	0.00300 (0.00454)	0.00114 (0.00460)
Log(Data Science)	-0.00590 (0.00755)	-0.00662 (0.00753)
Log(Digital Literacy)	0.00977 (0.00643)	0.00870 (0.00649)
Log(Data Storage)	-0.00268	-0.00266

	(0.00513)	(0.00509)
Log(Big Data)	-0.00111	-0.00276
	(0.00326)	(0.00353)
Quintile 2x6	-0.00201	
	(0.0149)	
Quintile 2x7	-0.00638	
	(0.0143)	
Quintile 2x8	-0.0240	
	(0.0158)	
Quintile 2x9	-0.0224	
	(0.0160)	
Quintile 2x10	-0.0110	
	(0.0209)	
Quintile 2x11	-0.0255	
	(0.0179)	
Quintile 2x12	-0.00180	
	(0.0183)	
Quintile 2x13	-0.00327	
	(0.0207)	
Quintile 2x14	-0.0440**	
	(0.0206)	
Quintile 2x15	-0.0169	
	(0.0166)	
Quintile 2x16	-0.0484**	
	(0.0194)	
Quintile 2x17	-0.0454**	
	(0.0179)	
Quintile 2x18	-0.0187	
	(0.0181)	
Quintile 2x19	-0.0330*	
	(0.0189)	
Quintile 3x6	-0.0156	0.0181
	(0.0174)	(0.0167)

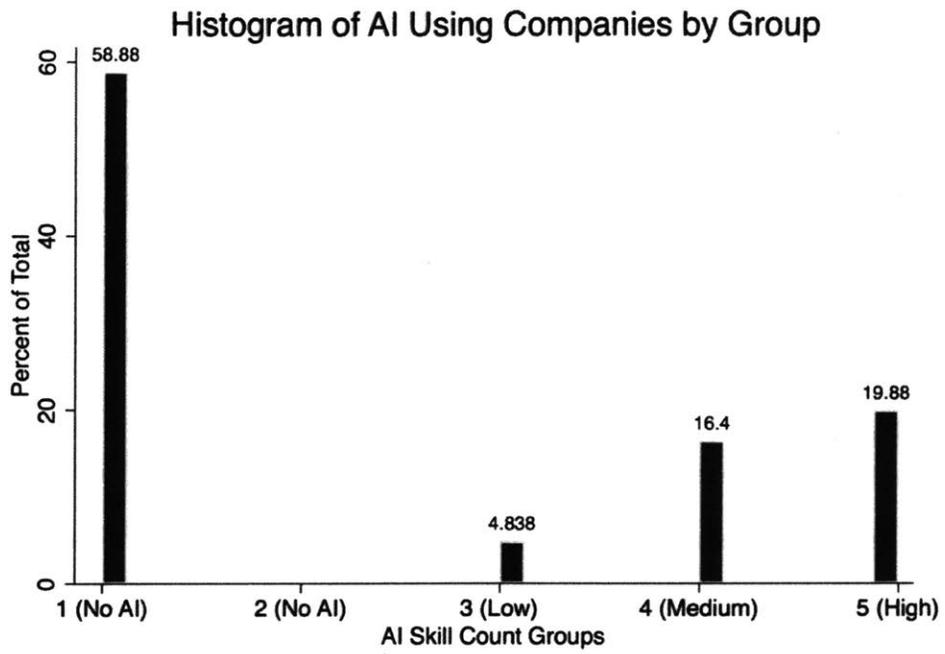
Quintile 3x7	-0.00405 (0.0166)	0.0124 (0.0207)
Quintile 3x8	0.000366 (0.0207)	0.0305 (0.0211)
Quintile 3x9	-0.0340* (0.0198)	0.0155 (0.0228)
Quintile 3x10	-0.00357 (0.0208)	0.00974 (0.0191)
Quintile 3x11	-0.00515 (0.0207)	0.00940 (0.0206)
Quintile 3x12	-0.00215 (0.0209)	0.0682*** (0.0244)
Quintile 3x13	-0.00881 (0.0210)	0.0341 (0.0236)
Quintile 3x14	-0.0146 (0.0211)	0.0424* (0.0236)
Quintile 3x15	-0.0232 (0.0214)	0.0481* (0.0248)
Quintile 3x16	-0.0339 (0.0223)	0.0483** (0.0237)
Quintile 3x17	-0.0249 (0.0210)	0.0354 (0.0239)
Quintile 3x18	0.00932 (0.0197)	0.0393* (0.0230)
Quintile 3x19	-0.0306 (0.0222)	0.0336 (0.0238)
Quintile 4x6	-0.0190 (0.0167)	0.00950 (0.0120)
Quintile 4x7	-0.0159 (0.0159)	0.0190 (0.0132)
Quintile 4x8	-0.00286 (0.0190)	0.0151 (0.0153)
Quintile 4x9	-0.0203	0.0146

	(0.0186)	(0.0137)
Quintile 4x10	-0.00866	0.0175
	(0.0190)	(0.0138)
Quintile 4x11	0.0180	0.0157
	(0.0183)	(0.0162)
Quintile 4x12	-0.00516	0.0385**
	(0.0199)	(0.0153)
Quintile 4x13	-0.00899	0.0249
	(0.0194)	(0.0160)
Quintile 4x14	-0.0311	0.00957
	(0.0200)	(0.0174)
Quintile 4x15	-0.0131	0.0244*
	(0.0204)	(0.0143)
Quintile 4x16	-0.0416*	0.0179
	(0.0231)	(0.0160)
Quintile 4x17	-0.0301	0.0243
	(0.0197)	(0.0152)
Quintile 4x18	-0.00189	0.00740
	(0.0197)	(0.0157)
Quintile 4x19	-0.0357*	-0.00713
	(0.0209)	(0.0162)
Quintile 5x6	-0.0394*	0.0213**
	(0.0209)	(0.0103)
Quintile 5x7	-0.0365	0.0482***
	(0.0223)	(0.0107)
Quintile 5x8	-0.0356	0.0244**
	(0.0264)	(0.0123)
Quintile 5x9	-0.0338	0.0337***
	(0.0229)	(0.0120)
Quintile 5x10	-0.0280	0.0156
	(0.0211)	(0.0122)
Quintile 5x11	-0.0290	0.0484***
	(0.0247)	(0.0136)

Quintile 5x12	-0.0366 (0.0232)	0.0513*** (0.0133)
Quintile 5x13	-0.0502* (0.0264)	0.0401*** (0.0147)
Quintile 5x14	-0.0437* (0.0253)	0.0360** (0.0151)
Quintile 5x15	-0.0272 (0.0229)	0.0416*** (0.0141)
Quintile 5x16	-0.0613** (0.0259)	0.0282* (0.0156)
Quintile 5x17	-0.0388 (0.0252)	0.0337** (0.0149)
Quintile 5x18	-0.0107 (0.0262)	0.0380** (0.0153)
Quintile 5x19	-0.0433* (0.0247)	0.0254 (0.0163)
Observations	20,522	20,526
R-squared	0.998	0.998

Table Note: Robust standard errors in parentheses, SEs Clustered by Firm. TensorFlow Launch corresponds to year-quarter 12. Quintile 5x12 is the coefficient on the fourth quarter of 2015 dummy interacted with the highest quintile for SML (1) and AI (2). None of the second quintile data are available for AI because the second quintile still does not use AI skills. Other values are relative differences to the first quintile.

*** p<0.01, ** p<0.05, * p<0.1



Percentage of Firms by AI Skill Group (Quintiles). Note: Counts are uneven because of ties in the skill counts.

Chapter 2 — The Productivity J-Curve: How Intangibles Complement General Purpose Technologies³⁸

After I left academe in 2014, I joined the technical organization at iRobot. I quickly learned how challenging it is to build deliberative robotic systems exposed to millions of individual homes. In contrast, the research results presented in papers (including mine) were mostly linked to a handful of environments that served as a proof of concept.

-Alexander Kleiner³⁹

I. Introduction

In the late 1980s, Robert Solow (1987) famously pointed out that “a technological revolution, a drastic change in our productive lives” had curiously been accompanied by “a slowing-down of productivity growth, not by a step up.” His famous productivity paradox, that one “can see the computer age everywhere but in the productivity statistics,” named a challenge for economists seeking to reconcile the emergence of exciting technological breakthroughs with tepid productivity growth.

Solow’s Paradox was not unique. In this paper, we argue it was one example of a more general phenomenon resulting from the need for intangible investments in early stages of new general purpose technologies. General purpose technologies (GPTs) are “engines for growth.” They are pervasive, improve over time, and lead to complementary innovation (Bresnahan and Trajtenberg 1995). These are the defining technologies of their times and can radically change the economic environment. They have great potential from the outset, but realizing that potential requires larger intangible and often unmeasured investments and a fundamental rethinking of the organization of production itself. Thus,

³⁸ Joint paper with Erik Brynjolfsson (MIT and NBER) and Chad Syverson (MIT and NBER). We thank Daron Acemoglu, Seth Benzell, John Fernald, Rebecca Henderson, Austan Goolsbee, Adam Saunders, Larry Summers, Manuel Trajtenberg and numerous seminar participants for helpful comments. The MIT Initiative on the Digital Economy provided valuable funding. We dedicate this paper to the memory of Shinkyu Yang, whose pioneering insights on the role of intangibles inspired us.

³⁹ Kleiner, Alexander. 2018. “The Low-Cost Evolution of AI in Domestic Floor Cleaning Robots” *AI Magazine*, (Summer).

the usual measurement of productivity growth as a residual after accounting for input changes in the production function can fall short when the technology changes the production function itself.

The extensive investment required to integrate GPTs into an organization is often forgotten. Along with installing more easily measured items like physical equipment and structures capital, firms must create new business processes, develop managerial experience, train workers, patch software, and build other intangibles. The difficulty for productivity measurement arises because intangible investments are not readily tallied on a balance sheet. The invention of a GPT can lead to the creation of entirely new asset classes and the transformation of existing capital varieties. It also presents abundant opportunity for entrepreneurs to discover new ways to deploy existing capital and labor. Moreover, these transformations of the production process do not occur overnight.

Given all of this, it is easy to see how something like Solow's Productivity Paradox can occur. There is a period during which measurable resources are committed (and measurable output forgone) to build new, unmeasured inputs that complement the new GPT. This period can be of considerable length. For example, the technologies driving the British industrial revolution led to "Engels' Pause," a half-century-long period of capital accumulation, industrial innovation, and wage stagnation (Allen 2009; Acemoglu and Robinson 2013). In the later GPT case of electrification, it took a generation as the nature of factory layouts was re-invented (David 1990). Solow was noting a similar phenomenon roughly two decades into the IT era.

We call this phenomenon the *Productivity J-Curve*. As firms adopt a new GPT, total factor productivity growth will initially be underestimated because capital and labor are spent to accumulate unmeasured intangible capital stocks. Later, measured productivity growth overestimates true productivity growth because the capital service flows from those hidden intangible stocks generates measurable output. The error in measured total factor productivity therefore follows a J-curve shape, initially dipping while the investment rate in unmeasured capital is larger than the investment rate in

other types of capital, then rising as growing intangible stocks begin to affect measured production. As we will explain later, large capital adjustment costs, correlated intangible investments, and high investment shares of income exacerbate the magnitude of the J-curve effect. In the long run, as investment quantities and capital stocks reach their steady-state, the mismeasurement problem disappears even if the intangible investments do not.

We documented the basic idea of the Productivity J-Curve, along with a discussion of the Productivity Paradox in the context of Artificial Intelligence (AI), in Brynjolfsson, Rock, and Syverson (2017), building on earlier work by Yang and Brynjolfsson (2001). The current paper formalizes and expands on these ideas, offering a set of quantitative methods designed to measure the value and productivity effects of intangible investments. Namely, we propose using a set of forward-looking measures derived from stock market valuations as a means of assessing the magnitude of intangible investment value in the style of the traditional growth accounting framework. The basic idea is that these hidden intangibles are still captured by market valuations. As noted by Yang and Brynjolfsson (2001), a combination model of the Q-Theory of investment (Hayashi 1982; Wildasin 1984; Hayashi and Inoue 1991) and neoclassical growth accounting (Solow 1956; Solow 1957; Barro 1998; Corrado, Hulten, and Sichel 2009; Oliner and Sichel 2000; Oliner, Sichel, and Stiroh 2008) can deal simultaneously with the magnitudes of the intangible component of GPT-related investment and lags in implementation. We extend the model in Yang and Brynjolfsson (2001) to adjust the traditional growth accounting methods to include unmeasured intangible capital investments, showing that the J-curve is a consequence of the growth of associated intangible investments. We use market value regressions (following Hall 2001 and 2004, and Hall 2006) to inform estimates of the currently installed value of intangible capital stocks. We show how one can use these estimates to infer more accurate measures of total factor productivity growth on a contemporaneous basis.

2. Technology, Investment Theory, and Productivity Growth

Economic historians have emphasized the transformative effects of GPTs in history. We mentioned the work of David (1990), Allen (2009), and Acemoglu and Robinson (2013) above. Rosenberg and Trajtenberg (2004) identify the Corliss steam engine as an “icon of the Industrial Revolution,” shifting population centers from rural to urban areas as water power was abandoned in favor of steam. Crafts (2004) explores the contribution of steam power to growth in greater detail for the British economy during the Industrial Revolution. Lipsey, Carlaw, and Bekar (2006) offer a list of possible GPTs (including electrification, mass production, and the factory system), also relating those inventions to the presence of the Productivity Paradox. Bresnahan (2010) conducts a wide review of the GPT concept, making the point that the modern era’s information and communication technologies (ICT) broadly constitute a GPT with transformative effects on the economy. Particularly relevant to our analysis work Helpman and Trajtenberg (1994) which notes how general purpose technologies can generate alternating periods of investment and harvesting. Likewise, Jacobs and Nahuis (2002) suggests that GPTs can cause an initial productivity slowdown as high-skilled workers invest in knowledge instead of production.

We focus in this paper on the most recent potential GPT: artificial intelligence. AI, and in particular the subfield of AI called machine learning (ML), is potentially pervasive, improves over time, and can spawn complementary innovation, meeting the Bresnahan and Trajtenberg (1995) criteria for a GPT. We would therefore expect, after an implementation lag period, for AI to significantly impact economic growth as other GPTs have (Brynjolfsson, Rock, and Syverson 2017; Cockburn, Henderson, and Stern 2018; Aghion, Jones, and Jones 2017; Agrawal, McHale, and Oettl 2018; Trajtenberg 2018). Nevertheless, the formal arguments presented here are applicable to other technologies and intangible capital accumulation more generally. The GPT context is useful because this is where we expect firm investment in unmeasured intangible capital goods to be large. Incremental innovations that do not transform productive activity are likely well captured by standard

models. The complementary innovations necessitated by GPTs motivate our approach. If it were not necessary to transform existing business processes via complementary intangible investments, new GPTs would immediately boost output in straightforward and measurable ways. Creating complementary innovation both introduces implementation lags and predisposes the new intangible capital accumulation dynamics to mismeasurement.

Part of the productivity growth slowdown of the past decade may be due to these dynamics.⁴⁰ We argued in earlier work that implementation and restructuring lags are a possible explanation for the juxtaposition of optimism about AI's potential and currently low productivity growth (Brynjolfsson, Rock, and Syverson 2017). An alternative explanation is that the current technological promise is unfounded (Gordon 2015) and we are in a period of secular stagnation (Summers 2015). A third story is that mismeasurement of productivity growth can arise from changes in measures of output quality, consumer surplus, or price indices, particularly for digital goods (Brynjolfsson, Eggers, and Gannamaneni 2018a, 2018b; Goolsbee and Klenow 2018). Syverson (2017) shows that that these types of mismeasurement, while important, are likely to be insufficient to explain the productivity growth decline. We focus instead on a different type of mismeasurement: the forgone output due to investment in unmeasured *capital* goods. Identifying these hidden asset values makes it possible to better measure true productivity growth.

Intangible assets are an increasingly important component of economic activity, especially IT-related intangibles (Brynjolfsson and Hitt, 2000; Hall 2000; Hall 2001; Brynjolfsson, Hitt, and Yang 2002; Tambe, Hitt, and Brynjolfsson 2012, Saunders and Brynjolfsson 2016). This has led to numerous updates to the standard growth accounting frameworks and an emphasis in recent

⁴⁰ Of course, there are many other possible explanations. For instance, Acemoglu (2002) argues that “an acceleration in skill bias could cause a TFP slowdown because it creates an imbalance in the composition of R&D.”

productivity studies on IT's role in productivity dynamics (Jorgenson and Stiroh 2000; Marrano, Haskel, and Wallis 2009; Corrado, Hulten, and Sichel 2009; Byrne, Oliner, and Sichel 2013; Byrne, Fernald, and Reinsdorf 2016), and specifically in the ICT-as-GPT case in (Basu, Fernald, Oulton, and Srinivasan 2003). Haskel and Westlake (2017) argue that intangible capital tends to have high fixed costs, low marginal costs, spillovers, and complementarities with other assets.⁴¹ Further, the existence of significant intangible assets might explain the relatively poor historical performance of Tobin's Q (the ratio of a firm's market-to-book value) in explaining capital investment (Crouzet and Eberly 2018). Accounting for organizational investments, human capital, and business processes can strengthen the link between observed investment and asset prices (Eisfeldt and Papanikolaou 2013, 2014; Peters and Taylor 2017; Kogan et al. 2017; Andrei, Mann, and Moyen 2018).

Our approach applies the Q-theory of investment to recover the portion of productivity growth attributable to unmeasured intangible capital outputs. To adjust aggregate productivity estimates from a growth accounting framework for intangible output, we need to know 1) the growth rates of tangible investment and capital, and 2) the quantity of intangible investment per unit of correlated tangible investment. We empirically pin down this second component from the estimation of market value regressions. Part of Q is treated as intangible capital instead of excess valuation over asset replacement costs. When we observe a firm's market value rise by an amount greater than observed investment, we infer the difference as reflecting the value of intangible capital investments that were correlated with the tangible investment. We call these correlated intangible investments *intangible correlates*. Our framework also handles the case in which intangible capital is used to produce more intangible capital.

The Productivity J-Curve that we describe in this paper is related to, but distinct from, the trade balance J-curve discussed in Rose and Yellen (1989) and Magee (1973).⁴² Their J-curve

⁴¹ They refer to the "4 S's" of intangible capital: sunk costs, scalability, spillovers, and synergies.

⁴² We thank Larry Summers for suggesting how the dynamics we model are similar to the trade J-curve..

describes how trade balances react over time to changes in real exchange rates.⁴³ The similarity between the two J-curves is that there is a change in the sign of derivatives of focal quantities with respect to time as time passes (trade balances in the earlier case, productivity in this one), reflecting the adjustment of production processes in response to an external shock. In Rose and Yellen, the shock comes from a large change in exchange rates. In our paper, it is from a large technological innovation.

3. Growth Accounting in the Presence of Unmeasured Intangible Investment

Our setup builds on the approach of Yang and Brynjolfsson (2001) as follows.

Suppose a competitive firm produces output with a general constant returns to scale production function. Then

$$Y = pF(K, N, A) \quad (1)$$

where Y is the final goods output of the firm, p is the price of final goods output (stable over time), K is the vector of capital goods, N is the vector of variable inputs (e.g., labor), and A represents the level of total factor productivity at time t . With flexible capital and input prices (r , w), we have the following, with g denoting a growth rate:

$$g_Y = \frac{\dot{Y}}{Y} = \frac{p(F_K \dot{K} + F_N \dot{N} + F_A \dot{A})}{Y} = \left(\frac{rK}{Y}\right) g_K + \left(\frac{wN}{Y}\right) g_N + g_A \quad (2)$$

$$g_K = \frac{\dot{K}}{K}; \quad g_N = \frac{\dot{N}}{N}; \quad g_A = \frac{\dot{A}}{A}$$

Values with an upper dot represent the total derivative with respect to time.

In words, one can decompose the growth in output over time into the growth in capital stock multiplied by capital's share of output plus the growth in flexible input quantity multiplied by the

⁴³ Assuming export prices between countries are sticky, one country depreciating its currency makes sticky-priced imports (exports) more (less) attractive, while later prices adjust and foreign import demand increases.

expenditure share of flexible inputs and a final total factor productivity growth term. This last term is the familiar Solow Residual. It represents an improvement in productive efficiency, or more modestly a kind of “measure of our ignorance” in how a firm converts inputs to outputs.⁴⁴ Equation (2) is the basis for traditional growth accounting, which we revisit in equation (9) with an adjustment for unmeasured intangible investments.

To incorporate adjustment costs, we modify (1) following Lucas (1967):

$$Y = pF(K, N, I, A) \quad (3)$$

Now the production function incorporates an investment term I with market price z such that the total cost of investment in one unit of capital goods is $(z - pF_I)$. F is assumed non-increasing and convex in I to represent the idea that adjustment costs grow increasingly costly for larger I . In other words, the firm must forgo an increasing amount of output as its rate of capital investment increases. This helps explain why firms cannot, for example, instantaneously replicate the capital stocks of their competitors without incurring larger costs.

We can relate firm investment behavior to market value using this production function.⁴⁵ For the price-taking firm, market value equals the sum of the capitalized adjustment costs. The firm must solve:

$$\max_{I, N} \left[\int_0^{\infty} \pi(t) u(t) dt \right] = V(0)$$

$$\text{where } \pi(t) = pF(K, N, I, A) - w'N - z'I$$

⁴⁴ Abramovitz, Moses. "Resource and output trends in the United States since 1870." In *Resource and output trends in the United States since 1870*, pp. 1-23. NBER, 1956.

⁴⁵ See, for example, Hayashi (1982), Wildasin (1984), and Hayashi and Inoue (1991).

$$\text{and } \frac{dK_j}{dt} = I_j - \delta_j K_j \quad \forall j = 1, 2, \dots, J. \quad (4)$$

That is, K_j is the capital stock of type j (indexes capital variety), N is a vector of flexible inputs, $u(t)$ denotes the compound discount rate at time t , and δ_j is the depreciation rate of capital of type j . As in Yang and Brynjolfsson (2001), F is assumed non-decreasing and concave in K and N . With homogeneity of degree one for F , we get the solution to the maximization of the Hamiltonian in (5) at time 0:

$$H(K, N, I, A) = (pF(K, N, I, A) - w'N - z'I)u(t) + \sum_{j=1}^J \lambda_j (I_j - \delta_j K_j) \quad (5)$$

with constraints:

$$\frac{\partial H}{\partial \lambda_j} = \dot{K}_j = I_j - \delta_j K_j \quad \forall j \in \{1, 2, \dots, J\}, \forall t \in [0, \infty]$$

$$\frac{\partial H}{\partial K_j} = -\dot{\lambda}_j = pF_{K_j}u - \lambda_j \delta_j \quad \forall j, \forall t$$

$$\frac{\partial H}{\partial I_j} = 0 = (pF_{I_j} - z_j)u + \lambda_j \quad \forall j, \forall t$$

$$\frac{\partial H}{\partial N_i} = 0 = (pF_{N_i} - w_i)u \quad \forall i \in \{1, 2, \dots, L\}, \forall t$$

$$\lim_{t \rightarrow \infty} \lambda(t)K(t) = 0$$

leading to an equation for the value of the firm:

$$V(0) = \sum_{j=1}^J \lambda_j(0)K_j(0) \quad (6)$$

Equation (6) shows that the value of the firm at $t = 0$ is the sum over all varieties of the capital stock quantities multiplied by the “shadow price” of investment of the respective varieties. In our context, this shadow price reflects adjustment costs.⁴⁶

Assuming all asset stocks are measured correctly and market prices correctly represent the value of claims on publicly traded firms, equation (6) suggests that a regression of firm value on dollar quantities of asset varieties will yield a coefficient vector that represents the average present value of one unit of each type of capital. In a frictionless efficient market, that vector would contain values of one for all assets. In the presence of adjustment costs, however, the coefficient will be greater than one.

We can extend this logic to unmeasured intangible GPT investments that are complementary to tangible assets. Suppose a firm adopting a new GPT must invest proportionately in two assets: computer equipment and firm-specific GPT specialist training (e.g., training AI engineers). For a firm with a measurable quantity of tangible computer equipment, the estimated shadow price coefficient for the computer equipment investment will exceed the “true” computer equipment coefficient by the amount necessary to represent the value of the complementary training as well. The specialist training is not capitalized on the firm’s formal balance sheet, yet the financial market must also value the future service flows from training if no arbitrage conditions are to hold. The market value premium

⁴⁶ Following equation (6) in Hall (2000), if λ_j represents the marginal q value (incremental market value created divided by asset replacement cost), then the marginal adjustment cost for the firm (set by the firms’ competitors) at its chosen capital investment rate is set by:

$$c' \left(\frac{k_t - (1 - \delta)k_{t-1}}{k_{t-1}} \right) = q_t - 1 = \lambda_t - 1$$

Where $c'(x)$ is the marginal adjustment cost function and δ is the depreciation rate of capital. In this case, there are no unmeasured intangible correlates, only adjustment costs of investment. Our framework below allows for both adjustment costs and unmeasured intangibles. In that case, the sum of these two elements is our λ value. (One interpretation of this summation is that capitalized convex adjustment costs are, in effect, a nonlinear component of correlated intangible investments.)

over book value implies a Tobin's Q above unity; the value of the firm is higher than the simple replacement cost of its *observed* assets.

When we examine data from financial markets, we find that technology firms have, on average, considerably higher values of Q. This suggests that they have either higher levels of adjustment costs, intangible correlated investments relative to booked assets, or both. This is consistent with the idea that implementing a new GPT requires complementary intangible investment to reorganize production.

In a growth accounting framework, the value of final goods in any given year is divisible into the value of consumption goods and the value of capital goods as follows:

$$p_c C + zI = Y = p_y F(K, N, I) = p_y F_n N + p_y F_k K + p_y F_I I = wN + rK + (z - \lambda)I \quad (7)$$

alternatively, the total payments to capital (including intangible stocks) and labor are:

$$p_c C + zI + (\lambda - z)I = wN + rK$$

This is the growth accounting identity. The value of consumption goods plus the value of capital investment is equal to total output Y . This, in turn, is equal to the total income of flexible inputs, capital rental costs, and investment (both measured and unmeasured).⁴⁷

If the $(\lambda - z)I$ value of capital goods production goes unmeasured, then part of the expenditure on capital goods is missing when the growth decomposition is performed. In the context of a GPT, this means that much of the training, the investment in implementing new decision processes, the reorganization costs, and the incentive designs necessary to generate productive service flows from GPT capital are left out.

⁴⁷ This framework assumes that the firm's Hamiltonian can be aggregated to the economy level, as does the standard growth accounting framework. This is not always the case (Houthakker 1955; Basu and Fernald 1997).

If the economy is accumulating the new GPT-related capital faster than it accumulates measurable capital, TFP will be underestimated. To see why, we can update the growth decomposition equation as follows:

$$g_Y = \frac{\dot{Y}}{Y} = \frac{p(F_K \dot{K} + F_N \dot{N} + F_I \dot{I} + F_A \dot{A})}{Y} \quad (8)$$

Following the first order conditions for the Hamiltonian above, we have

$$\lambda_j(0) = (z_j - pF_{Ij}) \quad \text{and}$$

$$g_Y = \left(\frac{pF_K K}{Y}\right) \left(\frac{\dot{K}}{K}\right) + \left(\frac{pF_N N}{Y}\right) \left(\frac{\dot{N}}{N}\right) + \left(1 - \frac{\lambda}{z}\right) \left(\frac{zI}{Y}\right) \left(\frac{\dot{I}}{I}\right) + \left(\frac{F_A A}{Y}\right) \left(\frac{\dot{A}}{A}\right) \quad (9)$$

There are several differences between this expression and the typical growth accounting equation (2). Because expressions like (2) are the standard method of computing TFP growth, deviations between the true growth decomposition (9) and the decomposition as implemented (2) reflect the sources of TFP mismeasurement.

The first difference is in the first term on the right side of (9). The capital services' share of income needs to be adjusted to account for the capital services from intangibles. Given homogeneity of degree one in the production function, the adjustment is:

$$\frac{pF_K K}{Y} = 1 - \left(\frac{wN}{Y}\right) + \sum_{j=0}^J \left(\frac{\lambda_j}{z_j} - 1\right) \left(\frac{z_j I_j}{Y}\right) \quad (10)$$

Where j indexes different capital varieties. In (2), the summation term was omitted from the value of capital's output share and so in a standard growth decomposition would be rolled into measured productivity growth. The difference between the capital share of income without unmeasured intangibles and the adjusted capital share is equal to the summation term in (10). Labor's share of

income remains the same under either setup. This is why the left side of (10) is stated precisely as the residual income share without labor added to the summation term.

The true growth decomposition (9) also clearly shows in its second-to-last term the investment component missing from (2). The contribution to output growth of this term would also be subsumed into productivity growth when applying decomposition (2).

Denote with g_{TFP}^* the productivity growth measure derived from the true decomposition (9)—i.e., the final term in that expression. Denote the productivity growth measure derived from a standard decomposition (2) as g_{TFP} . Taking the difference gives us an expression for the measurement error in productivity growth:

$$g_{TFP}^* - g_{TFP} = \sum_{j=0}^J \left(\frac{\lambda_j}{z_j} - 1 \right) \left(\frac{z_j l_j}{Y} \right) (g_{I_j} - g_k) \quad (11)$$

where j indexes the capital variety type. This equation states that the measurement error in TFP growth is the sum of the differences between investment growth rates in a given capital variety and the overall growth rate of capital, multiplied by the investment share of observable income and the per-observable–investment-unit value of intangible correlates and adjustment costs. With a fixed multiplier and investment share of income, the growth rate differential between investment and overall capital drive the dynamics of measurement error.

If the shadow price of technological investment is simply the market price of the investment good, so that $\lambda/z = 1$, there is no missing growth in output. However, for GPTs, it is likely that there will be a need for extensive unmeasured investments that correlate with tangible capital goods production, making λ/z greater than one. Because investment share of measurable output (zI/Y) is positive, the sign of productivity mismeasurement will depend on the difference between the growth rates of *investment* in GPT capital and the *installed stock* of GPT capital. Therefore, if GPT diffusion

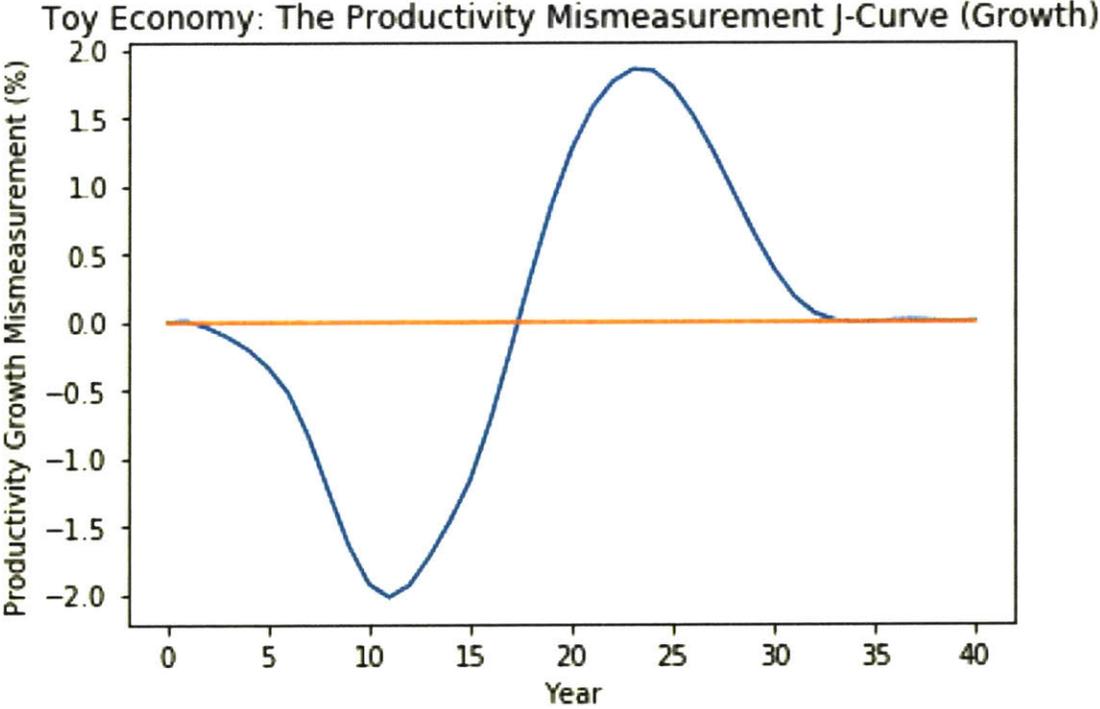
leads to a period of differential growth rates of intangible investments and intangible stocks—and we offer intuition below for why this is likely to be the case—then measured and true productivity growth will diverge during that period. In the long run, if the economy reaches steady state, the two growth rates will converge and productivity will no longer be mismeasured, even in the presence of unmeasured intangibles.

4. The Productivity J-Curve

The most straightforward way to understand the Productivity J-Curve is to consider foregone output used to produce unmeasured capital goods. Suppose a company wants to become more “data-driven” and reorganize its production processes to take advantage of new machine learning prediction technologies (Brynjolfsson and McElheran 2016). This firm might want, for example, to change its labor mix to build more software and to teach its customers to order products online instead of in person. While the company develops online product ordering applications and business processes for that purpose, it will not be able to use those investment resources to produce more final goods inventory. At the same time, though, the capital assets the firm *is* building—institutional software knowledge in the company, hiring practices, organization building, and customer retraining to use digital systems—are left unmeasured on the balance sheet.

On the margin, the (present-discounted and risk-adjusted) value of these unmeasured assets equals the costs incurred to produce them. But during the period in which that output is foregone, the firm’s (traditionally measured) productivity will suffer because it will seem as though the company produces proportionately less output relative to its inputs. Later, when those hidden intangible investments start to generate a yield as inputs, it will seem as though the measured capital stock and employed workers have become much more productive. Therefore, in early investment periods productivity is understated, whereas the opposite is true later when investment levels taper off.

The mismeasurement in this example regards a J-curve in productivity *levels*, and we derive a general expression describing its evolution below. That said, a similar J-curve exists for productivity *growth rates*. The math behind this growth rate J-curve is precisely that from our analysis above and is summarized in equation (11). The underlying intuition is very much like that for the J-curve in the productivity level. Early in the GPT diffusion process, intangible investment growth is likely to be larger than intangible capital stock growth. With missed output growth dominating missed input growth, measured TFP growth is lower than true TFP growth. Later in the GPT diffusion process, investment growth slows below the growth rate of the installed intangible stock. Measured productivity growth then exceeds its true level. Eventually the growth rates equalize in steady state, and productivity mismeasurement disappears.



The mathematics behind the J-curve in level terms follow from equation (7). Differentiating our earlier production function (assuming numeraire output price for simplicity) yields:

$$dY = F_K dK + F_N dN + F_I dI + F_A dA = r dK + w dN + (z - \lambda) dI + F_A dA \quad (12)$$

If we have measured capital service flows, labor service flows, and we know investment prices and their installed shadow values, we can back out the component of output driven by productivity improvements dA . In efficiency units and log terms (for an ordinary Cobb-Douglas production function):

$$\ln(Y + \lambda I) = \ln(A) + \alpha \ln(K) + (1 - \alpha) \ln(N)$$

$$\ln(Y + zI) = \ln\left(\frac{Y + zI}{Y + \lambda I}\right) + \ln(A) + \alpha \ln(K) + (1 - \alpha) \ln(N) \quad (13)$$

That is, with factor shares of α and $1-\alpha$ for capital (including investment) and labor (respectively), we now see that the productivity level decomposition in (13) has an additional term of the ratio of total measured output to total output including intangible correlates and adjustment costs. Fixing a multiplier λ , this ratio is always going to be (weakly) less than one, implying a drag on the measured productivity level. But recall that the measure of capital stock K includes the intangible investment stock as well. We modify (13) to explicitly include the unmeasured intangible capital stock:

$$\ln(Y + zI) = \ln\left(\frac{Y + zI}{Y + \lambda I}\right) + \ln(A) + \alpha \ln(\mathbf{K}) + (1 - \alpha) \ln(N) \quad (14)$$

\mathbf{K} denotes the vector of ordinary capital stocks and unmeasured intangible capital stock. If we have the marginal condition that the present value of investment returns in unmeasured intangibles equals the cost of the investment, in the long run the positive and negative effects of these additional terms on the level of productivity must net out. This is because the total risk-adjusted net present value of capital service flows from unmeasured intangible capital investments must equal the costs to the firm of making those investments.

To see the Productivity J-Curve, consider the case that the stock of intangible capital is zero. In this case, we only have the first negative term creating mismeasurement of productivity. Later, if

net investment is zero instead, we only have the intangible component of the positive adjustment term $\alpha \ln(\mathbf{K})$, causing an overestimate of the productivity level $\ln(A)$.

Recall that the J-curve in productivity *growth* can be seen in the right-hand side of (11). When the growth rate of unmeasured intangible investments exceeds (is lower than) the growth rate of the total capital stock, the true productivity growth rate will exceed (be lower than) the measured productivity growth rate. This effect in either direction is amplified by 1) a large installed-to-purchase price ratio λ/z of investment (or large quantities of unmeasured intangible investment required per unit of measured investment) and 2) a large measured investment share of measured output. This latter effect is part of the explanation for the Productivity Paradox; when the economy is in the early stages of accumulating GPT-related capital, measured investment's share of measured output will be relatively low. The figure below shows the total factor productivity growth J-curve. Because the values are in growth terms, the sum of overestimates and underestimates need not be zero over time. The *level* impact must necessarily be zero in expectation, however, if assets are efficiently priced.

There is a similar description in equation (5) of Brynjolfsson, Rock, and Syverson (2017), where the mismeasured components of investment and capital stock work against each other to generate the difference between measured and actual productivity growth. In this case, r and z , respectively, represent the total prices of capital and investment (adjusted to account for intangible quantities).

$$g_{TFP_{measured}} - g_{TFP_{corrected}} = \left(\frac{r_{intan} K_{intan}}{Y} \right) g_{K_{intan}} - \left(\frac{z_{intan} I_{intan}}{Y} \right) g_{I_{intan}} \quad (15)$$

Note the similarity with equation (11) above. We can extend (15) with (11) to situate this difference in terms of all intangible correlates associated with different capital varieties while separating capital and investment effects. Doing so gives:

$$g_{TFP_{measured}} - g_{TFP_{corrected}} = \left(\sum_{j=0}^J \left(\frac{z_j I_j}{Y} \right) \left(\frac{\lambda_j}{z_j} - 1 \right) g_{K_j} \right) - \left(\sum_{j=0}^J \left(\frac{z_j I_j}{Y} \right) \left(\frac{\lambda_j}{z_j} - 1 \right) g_{I_j} \right) \quad (16)$$

The difference in measurement between the standard approach and ours is just the investment share and intangible value-weighted difference in growth rates between capital stocks and investments.

Note that (16) can be rearranged to form a regression specification, in which the intercept is defined by the corrected measure and the coefficient estimates are defined by the investment shares and lambda values:

$$g_{TFP_{measured}} = g_{TFP_{corrected}} + (g_K - g_I)' \beta + \epsilon \quad (17)$$

In steady state, the growth rates of capital and net investment converge, mitigating the mismeasurement problem. In the short run, the deployment of resources of different types to produce outputs of measured and unmeasured varieties can influence the degree to which productivity growth is mismeasured. These unmeasured intangible capital stocks might be used to produce even more unmeasured intangible assets, in which case the hidden output and hidden input effects can offset each other somewhat. In the case that the rate of intangible capital production accelerates and uses measured capital and labor services in increasingly greater quantities, the J-curve effects are more pronounced.⁴⁸ This will also occur if the quantity of intangible correlates (including adjustment costs) per unit of tangible investment increases.

5. Is Hidden AI Capital Investment Already Causing a Productivity Shortfall?

Gross Domestic Product in the U.S. in 2017 was \$19.5 trillion and in real terms grew at an average annual rate of 2.17% over 2010 to 2017, down from 2.72% per year from 2000 to 2007 (the eight years prior to the Great Recession).⁴⁹ This implies that unmeasured intangible capital investment

⁴⁸ There is also a degenerate scenario in which firms shift toward focusing on intangible output production using intangible assets. In this case, the productivity measurement apparatus starts to lose its value.

⁴⁹ From the Bureau of Economic Analysis GDP statistics.

over 2010 to 2017 would need to average \$107 billion per year ($= 19.5 \text{ trillion} * [2.72\% - 2.17\%]$) in 2017 dollars to explain the entire slowdown in in GDP growth. How much of this slowdown could a Productivity J-Curve for investment in AI and related intangibles explain?

The economy is early in the AI adoption cycle, but recent growth has been impressive. There has been a rapid increase in the use of AI and robotics technology over the past decade (Furman and Seamans 2018). Startup funding for AI has increased from \$500 million in 2010 to \$4.2 billion by 2016 (growing by 40% between 2013 and 2016) (Himel and Seamans 2017). Though concentrated heavily in the information technology sector, overall measurable corporate investment in AI in 2016 was \$26-39 billion, marking 300% growth since 2013 (Bughin et al. 2017). Similarly, international industrial robot shipments since 2004 have nearly doubled overall and almost quadrupled in the consumer electronics industry (Furman and Seamans 2018).

For AI to account for the 0.55% of “lost” output in 2017 GDP, the quantity of correlated intangible investments per unit of tangible investment must be between roughly 2.7 and 4.1 times the observable investment values (using the Bughin et al. (2017) estimate).⁵⁰ This is not implausible. Brynjolfsson, Hitt, and Yang (2002) find that the total market value of measured computer capital investments is as much as \$11.8 per \$1 in measured expenditure, with a standard error of \$4.025.⁵¹ None of the intangible “shadow” value will show up in the productivity statistics. Because the foregone output cannot be explained by growth in labor or observable capital inputs alone, the output shortfall will be attributed to slower productivity growth. Further, this investment (discounted and risk-adjusted) will later generate a capital service flow that produces measurable output.

⁵⁰ The required forgone output in 2017 was \$107 billion. Dividing by the low observed investment figure of \$26 billion implies a required intangible investment that was $107/26 = 4.1$ times the observed investment. Using the larger \$39 billion figure implies intangibles that were $107/39 = 2.7$ times observed investment.

⁵¹ This uses a series of regression specifications motivated by a version of equation (6) in the previous section.

Of course, these numbers are just for 2017, when measured AI investment was several multiples of what it was only a few years prior. Thus analogous pre-2017 values would be notably smaller, and it is unlikely that much of the GDP slowdown gaps in those earlier years would be attributable to AI-related intangibles. However, given that AI investments are likely to continue growing quickly, and the fact that where it exists, AI capital has a high market valuation and as such a considerable shadow value for intangible correlates, we could well be likely entering the period in which AI-as-GPT could have noticeable impacts on estimates of output and productivity growth.

6. Deploying the Framework Using R&D, Software, and Computer Hardware Investment

While the results in the previous section imply AI-related intangibles *per se* have only very recently been large enough to noticeably affect measured output and productivity, other technology-related investments may have had more substantial effects over greater horizons, creating their own J-curve dynamics as a result. We explore this possibility in this section.

Specifically, we estimate the per unit magnitudes of intangible capital investment that coincides with observable R&D, software, and computer hardware capital. We then use those values to adjust total factor productivity estimates and explore if substantial J-curve effects exist for those capital types. To estimate the magnitude of intangible investments, we use the approach described above for obtaining intangible capital shadow values by comparing firms' observable investments to their market capitalization.⁵² We obtain shadow values for R&D, software, and hardware (each) since 1961 and use these to build up time series estimates of intangible stocks from 1961-2017.

⁵² Recall as discussed above that adjustment costs can be thought of as a nonlinear component of intangible investment.

Our productivity baselines, net capital stock, and investment by capital variety estimates all come from Fernald (2014), extended through 2017.⁵³ We take estimates of the total stock of research and development (R&D) capital and the total stock of capitalized selling, general, and administrative (SG&A) expense from the Peters and Taylor (2017) measures available in Wharton Research Data Services (WRDS) (we extend these measures through 2017 as well, following the guidelines in their paper). These estimates are joined to Compustat to construct a panel from 1961-2017 of firm market value, total asset book value, total R&D capital, total “organizational” capital (the capitalized SG&A expenditure), and other firm identifiers of all publicly-held companies in the U.S. In our regressions, we define industry by four-digit NAICS code.

R&D capital provides a useful context for understanding Productivity J-curve dynamics for a few reasons. Corporate research leads to the development of new technologies that diffuse over time, and there has been a steady flow of investment into R&D for decades. Further, the link between R&D investment and market value is well established (Hall 1993, 2006). Because investment in R&D has persisted over the long term, we are more likely to find investment in R&D at nearly steady-state levels. This implies that the intangible-related challenges for productivity estimation coming from R&D are likely to be minimal at present. (Remember from our analysis above that as the growth rates of intangible investment and stocks converge, productivity mismeasurement falls to zero.) Nevertheless, the exercise presented here for R&D is applicable to other capital varieties.

In contrast, heavy software and computer hardware capital investment is a more recent phenomenon in which firm behavior might (still) not have entirely matured, so J-curve dynamics may still be present. We find evidence of this further below after first demonstrating our approach with analysis of R&D-related intangibles.

⁵³ Capital stock estimates for this series are also available from the U.S. Bureau of Economic Analysis (BEA).

The first step in estimating the productivity mismeasurement effect of intangible correlates is estimating how many units of intangible investment correspond to observable investment quantities. We begin with R&D and capitalized SG&A stock measures from Peters and Taylor (2017). For this, we run a market value regression of the style in Hall (1993) and Brynjolfsson, Hitt, and Yang (2002). The specification for firm i in industry j at time t is:

$$Market\ Value_{ijt} = \beta_0 + \beta_1 TotalAssets_{it} + \beta_2 R\&D_{it} + \eta_{jt} + \epsilon_{it} \quad (18)$$

The coefficient on R&D picks up the ratio of dollars of market value created per unit of R&D stock at the firm in a given year. This, which we refer to as the *intangible multiplier*, is the ratio (λ/z) from our analysis above. We estimate specifications both including and excluding capitalized SG&A and industry-year fixed effects. The results are in Table 1 below.

Table 1: Market Value Regressions on R&D and SG&A Stocks

	(1)	(2)	(3)	(4)	(5)	(6)
Market Value Regressions (1962-2017)	Basic R&D	Basic R&D and SG&A	Industry-Year Fixed Effects: R&D	Industry-Year Fixed Effects: R&D and SG&A	Firm and Year Fixed Effects: R&D	Firm and Year Fixed Effects: R&D and SG&A
Total Assets	1.016 (0.00179)	0.998 (0.00232)	1.015 (0.00853)	0.999 (0.0107)	1.013 (0.00725)	0.997 (0.0110)
R&D Stock	2.730 (0.105)	1.753 (0.0970)	2.841 (0.479)	1.953 (0.399)	2.161 (0.297)	1.509 (0.278)
SG&A Stock		1.755 (0.102)		1.657 (0.399)		1.453 (0.374)
Constant	656.8 (14.32)	458.7 (18.06)				
Firm-Year Observations	268,687	268,687	266,795	266,795	267,683	267,683
R-squared	0.987	0.988	0.989	0.989	0.993	0.993

Industry-Year FE	No	No	Yes	Yes	No	No
Firm and Year FE	No	No	No	No	Yes	Yes

Robust standard errors in parentheses

Total Assets are the total assets on the firm's balance sheet, Industry is the four-digit NAICS code. Market Value is the sum of the book value of debt, preferred stock, and the end-of-year equity share price multiplied by common shares outstanding. Specifications (5) and (6) include firm and year fixed effects, but not firm-year fixed effects. Standard errors in parentheses (robust for (1) and (2), clustered by industry in (3) and (4), clustered by firm in (5) and (6)).

The coefficients on Total Assets is very close to 1 as expected, whereas estimates for R&D are 2.73, 2.84, and 2.16 for specifications without capitalized SG&A and with (respectively) no fixed effects, industry-year fixed effects, and firm and year fixed effects). Including capitalized SG&A, the estimates decrease somewhat to 1.75, 1.95, and 1.51 for the respective specifications, with the coefficients on capitalized SG&A picking up much of the difference. Thus, these models suggest that the market value of capitalized research and development expenses is between about two and three dollars per net dollar of investment.

We also estimate a year-by-year regression of market value on total book assets and capitalized R&D with industry fixed effects. Figure 3 presents the time series of R&D coefficient estimates for that specification.⁵⁴ The year-by-year regressions reveal substantial variation in the shadow value of R&D-related intangible assets, consistent with overall valuation dynamics. It is with this set of values that we proceed to adjust productivity growth measurement. Figure 4 shows the same coefficient estimates for Total Assets, which are considerably lower in comparison (note the vertical scale is an order of magnitude smaller than Figure 3).

⁵⁴ The full table of coefficients is available in the Online Appendix. Available: <http://drock.mit.edu/Research>

Figure 3: R&D Market Value Year-by-Year Regression Coefficients 1962-2017

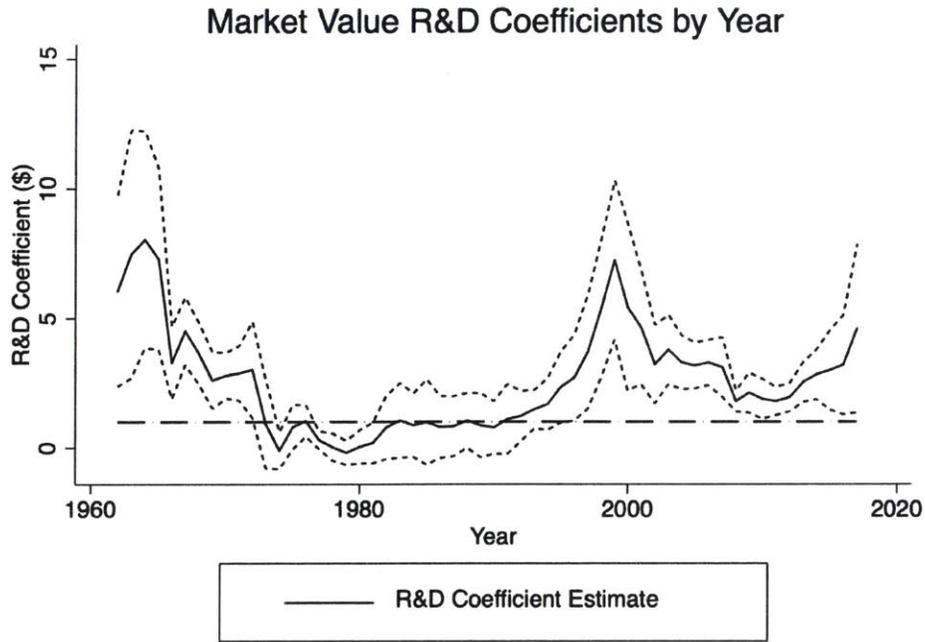
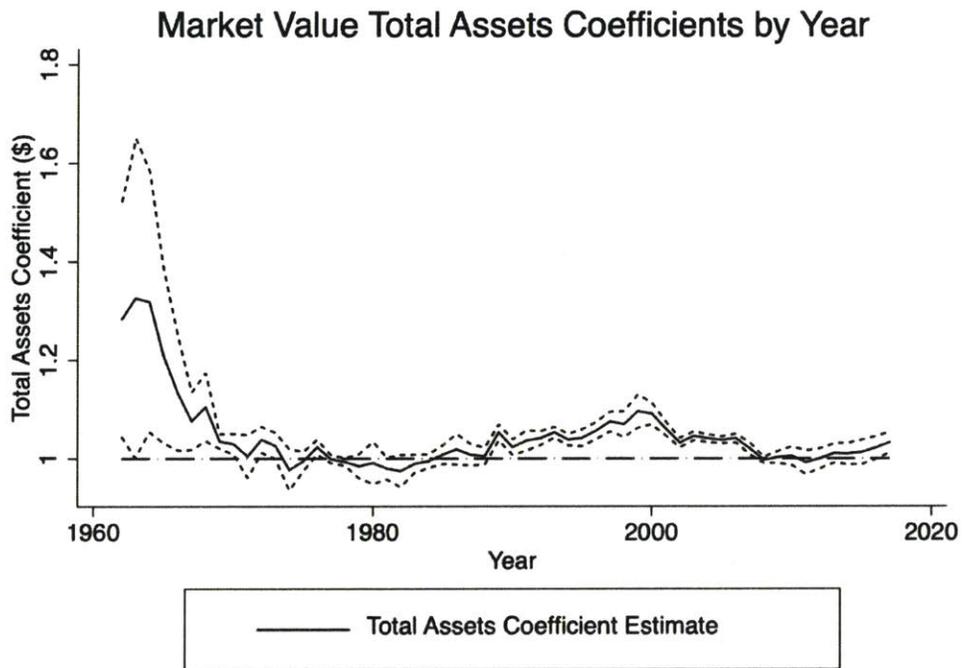


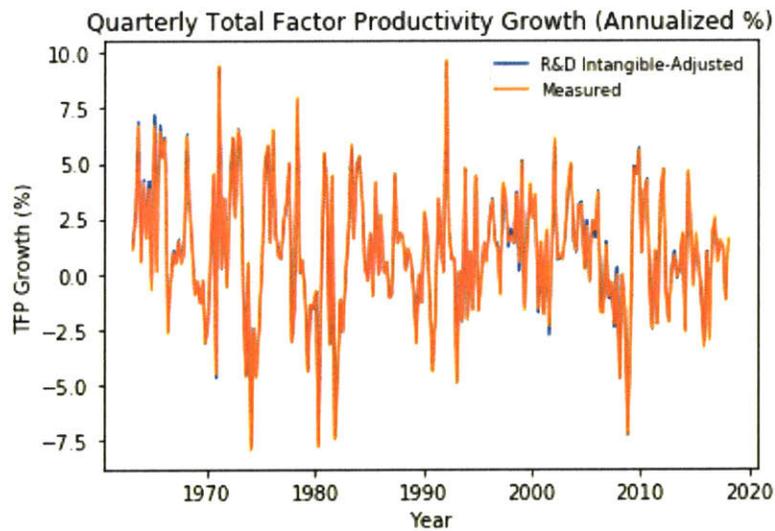
Figure 4: Total Asset Market Value Regression Coefficients 1962-2017

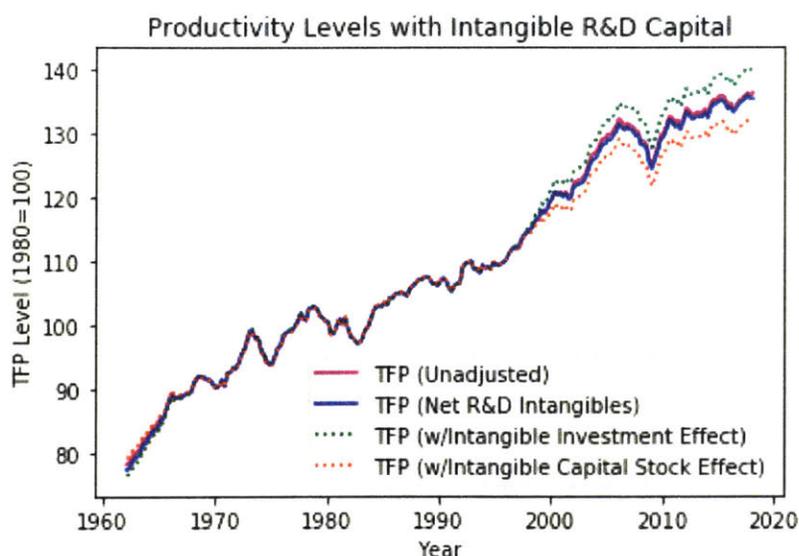


Given an estimate of the total amount of intangible correlates per unit of investment, we proceed to adjust the productivity level and growth estimates to include the missing intangible outputs

and inputs using equation (16). Figure 5 shows the time series of TFP growth, both as measured in Fernald's data and adjusting for unmeasured R&D-related intangibles. Figure 6 shows the effects in level terms, obtained by integrating the growth rates.

The unadjusted series differs very little from the net adjusted series. The reason is that R&D capital investment rates have been steady over the observation period, roughly canceling out the countervailing influences of intangible outputs and intangible inputs. This is made clear by the dotted green and red lines in Figure 6, which isolate the influence of the two terms in equation (16). The red dotted line shows the downward adjustment to measured productivity due to the failure to measure intangible capital input service flows. The green dotted line reflects the nearly equal-sized upward adjustment to productivity due to uncounted outputs tied to intangible investment.





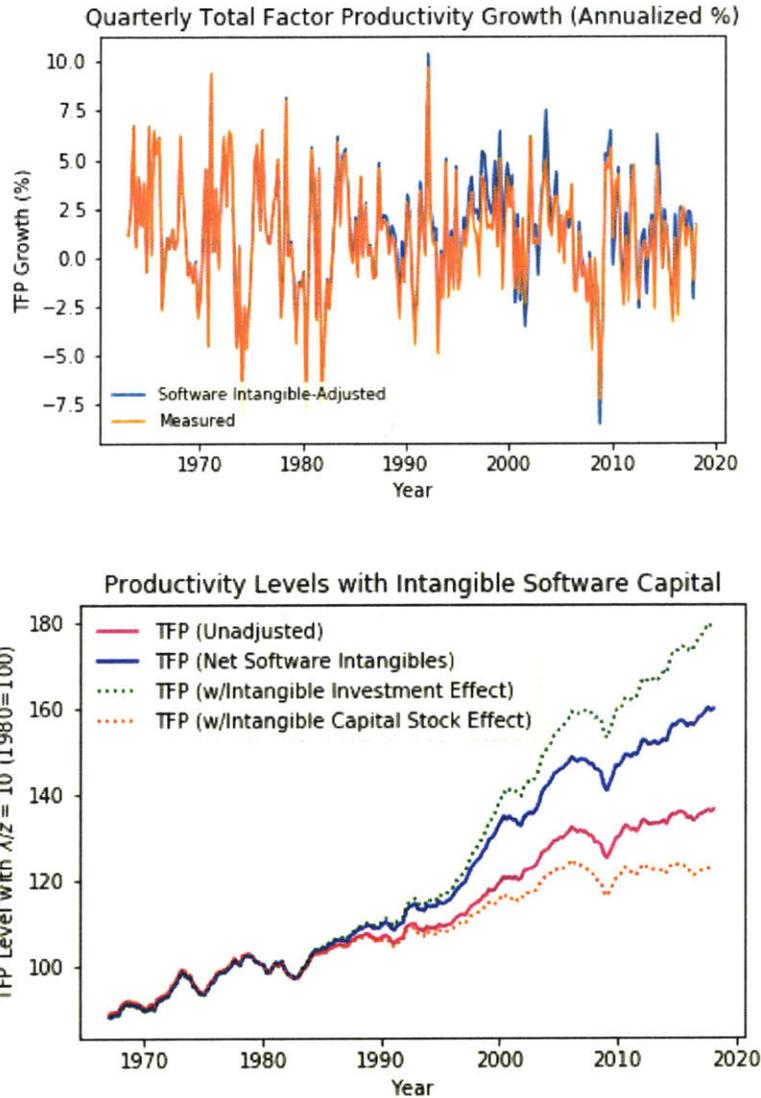
Figures 5 (top) and 6 (bottom): R&D-related Intangible Capital-adjusted Total Factor Productivity

Although the net measurement effects of R&D-related intangibles are negligible, the same is not true for software and computer investments. We do not have similar firm-level data on IT capital stocks and investment to run the market value regression in equation (18) for IT, so we apply the productivity adjustment analysis under a series of plausible values for the intangible multiplier λ/z , guided by the Brynjolfsson, Hitt, and Yang (2002) estimate that each unit of observable software and computer hardware is associated with roughly \$12 (standard error \$4.02) of firm market value.

In contrast to the adjustment for R&D-related intangibles, the Productivity J-curves for both software and computer hardware capital (we separately analyze each) have yet to reach positive territory in terms of levels.

Of the three capital varieties we investigate in this study, software's J-curve is in the least mature stage. Software investment has been and continues to be growing faster than overall capital investment, and its level is sufficiently large to suggest that part of the productivity slowdown might be explained by a compositional shift of investment toward digital assets. Figures 7 and 8, analogously to Figures 5 and 6 for R&D-related intangibles, show the annualized quarterly growth rates and levels

of measured TFP and software-intangible-adjusted TFP. The differences between measured and corrected estimates are starkly larger than those arising from R&D.



Figures 7 (top) and 8 (bottom): Software-related Intangible Capital-adjusted TFP

The J-curve dynamics of software investment began in the 1990s and have not waned since. If we assume an intangible multiplier of \$10, roughly the level but somewhat lower than Brynjolfsson, Hitt, and Yang (2002), the net adjusted TFP level (163.9) is 18.2% higher than measured TFP (138.6) as of the beginning of 2017. Figure 9 shows the productivity level adjustments for more conservative

intangible multipliers. Even for lower levels of the multiplier, the productivity level differences are notable and growing.

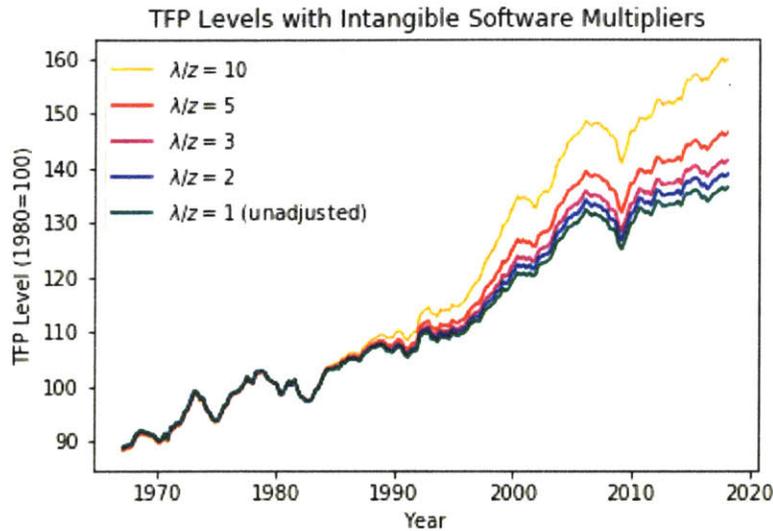
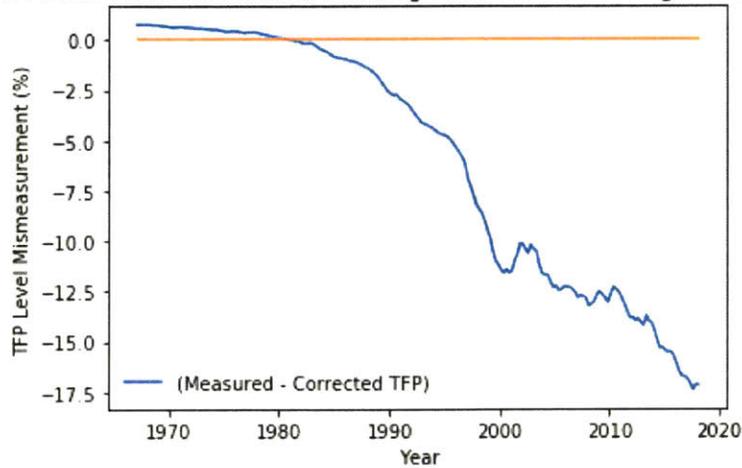


Figure 9: Total Factor Productivity Levels Corrected for Different Software Intangible Multipliers

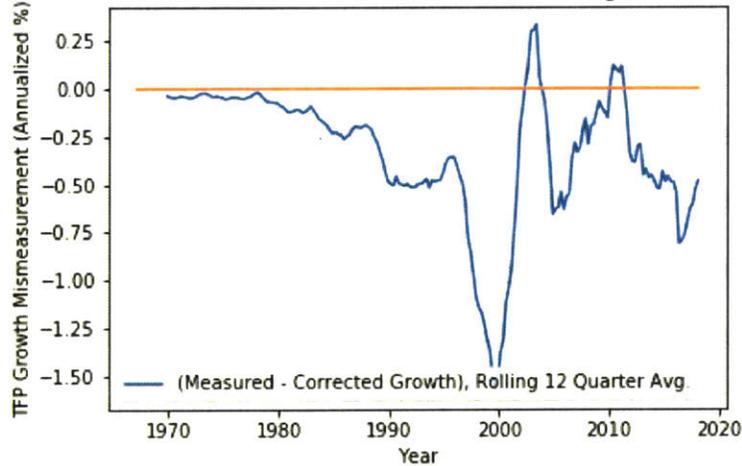
The reason behind the growing understatement of productivity due to software-related intangibles is the growing rate of software investment. Aside from brief periods following the dot-com bust and the financial crisis, investment in software has grown significantly. As a result, software-related intangible investment rates are not yet in steady state. As the analysis above shows, when the investment growth rate exceeds the growth rate of the intangible stock, productivity growth is understated. Since 2010, when the productivity growth mismeasurement effect was very nearly zero, annualized quarterly productivity growth underestimation increased to 0.86% by the end of 2016. The implied understatement was even larger at the end of the 1990s, where measured productivity was 1.6% lower than software-adjusted productivity. Figures 10 and 11 show the respective mismeasurements of TFP levels and growth for software-related intangible capital outputs (i.e., the vertical distances between the adjusted and measured series in Figures 7 and 8) since 1967. At least in level terms, we are still in the capital accumulation phase of a deep Productivity J-curve. Tables in the

Online Appendix show the productivity growth adjustments for R&D, computer software, and computer hardware from 1967-2017.⁵⁵

TFP Level Mismeasurement Percentage with Software Intangibles ($\lambda/z = 10$)



TFP Growth Mismeasurement with Software Intangibles ($\lambda/z = 10$)

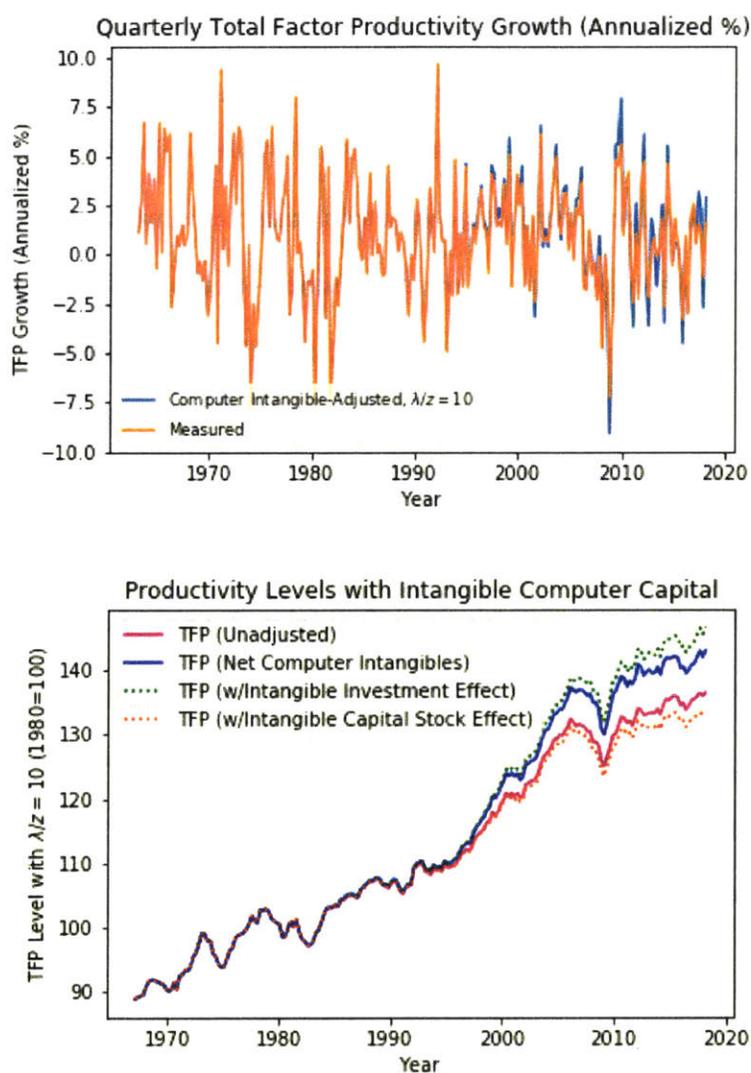


Figures 10 (top) and 11 (bottom): Computer Software-related TFP Mismeasurement in Levels and Growth Terms (respectively)

We extend our analysis to computer *hardware*-related intangible investment. Figures 12 and 13 show adjusted and measured TFP growth and levels, again assuming an intangible multiplier of \$10 for each

⁵⁵ Available at <http://drock.mit.edu/Research>

dollar of hardware investment. Again, the divergence between measured and corrected TFP begins to become noticeable in the 1990s, after Solow's famous quip. Figure 13 also shows where the TFP level would be without adjustment (purple), the net intangible-adjusted series (blue), isolating only the missing intangible inputs effect (dotted red), and isolating only the missing intangible outputs effect (dotted green). Figure 14 compares the adjusted series for an intangibles value of \$10, \$5, \$3, \$2, and \$1 (unadjusted).



Figures 12 (top) and 13 (bottom): Computer-Hardware-Related Intangible Capital-adjusted TFP

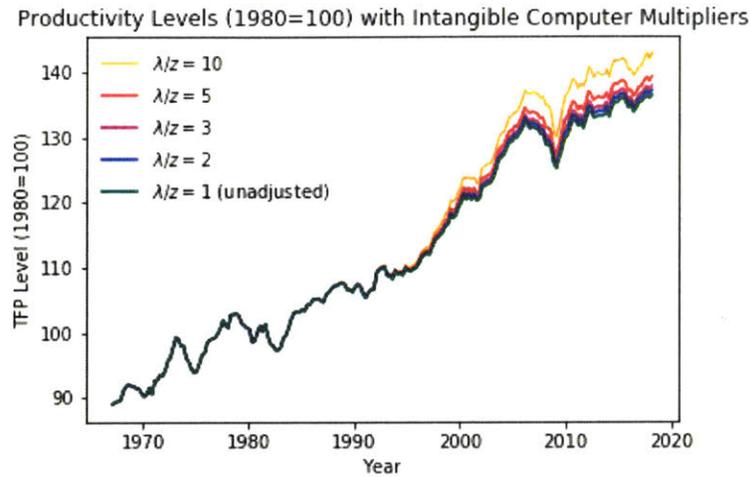
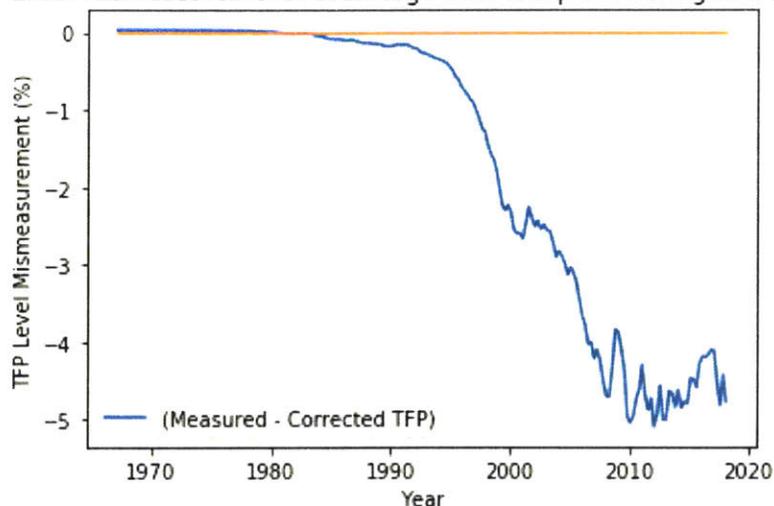


Figure 14: Total Factor Productivity Levels Corrected for Different Computer Intangible Multipliers

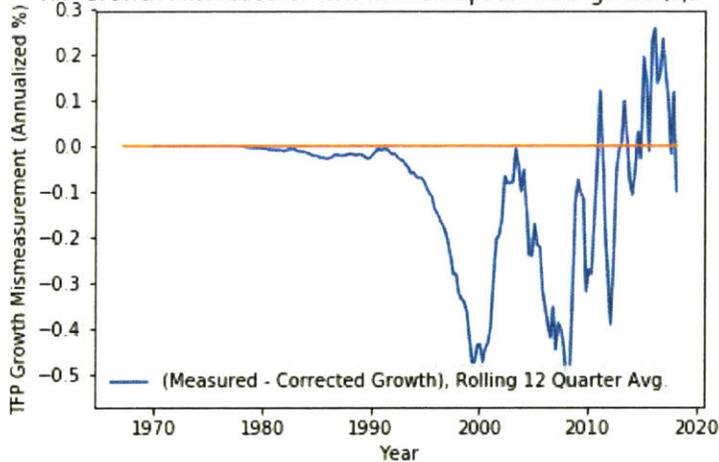
The quantitative patterns for hardware are different than what we found for software. First, the accumulated mismeasurement due to hardware-correlated intangibles is much more modest. Adjusted TFP at the end of 2016 is 4.4% higher than the measured series—a considerably smaller gap than that associated with software-related intangibles. Second, and interestingly, the recent slowdown in the rate of hardware investment has actually caused a small *overstatement* of productivity growth, and as a result, has started to bring the level difference back toward measured TFP. The reversal started following the dot-com bust, reverted as computer hardware investment rebounded in the following years, and then reversed again at the start of the Great Recession. Figures 15 and 16 show the magnitudes of the deviations in TFP growth and levels between the measured and corrected series. In growth terms, the latter (overstatement) part of the Productivity J-curve (at least that component tied to hardware investment) appears to have begun, and in level terms productivity understatement has stabilized.⁵⁶ Assuming an intangible multiplier on hardware capital of 10, the growth overestimate was about 0.2% in 2016.

⁵⁶ Figure 15 is a trailing three-year average of quarterly annualized total factor productivity growth estimates.

TFP Level Mismeasurement Percentage with Computer Intangibles ($\lambda/z = 10$)



TFP Growth Mismeasurement with Computer Intangibles ($\lambda/z = 10$)



Figures 15 (top) and 16 (bottom): Computer Hardware-related TFP Mismeasurement in Levels and Growth Terms (respectively)

7. Can Intangible Capital Outputs Explain the Productivity Slowdown?

We now take the above estimates of the TFP adjustments due to intangible capital related to R&D, software, and computer hardware to ask if the measured productivity slowdown after 2004 (see, e.g., Gordon 2015; Summers 2015; Syverson 2017) can be accounted for by such intangibles. Some role seems plausible; while our calculations above imply intangibles related to software and (to a lesser extent) hardware started having a noticeable influence on true TFP in the 1990s, they also made

contributions in more recent periods. If these are larger than their earlier influence, they would in part explain the measured productivity slowdown.

The slowdown in measured annual TFP growth from 1995-2004 to 2005-2017 was approximately 1.23% per year.⁵⁷ Had measured TFP grown since 2005 at the same rate it did from 1995-2004, and holding labor and tangible capital inputs fixed at their observed levels, U.S. GDP at the end of 2017 would have been \$3.46 trillion higher than it was.⁵⁸

To see if intangible capital accumulation tied to R&D, software, and computer hardware investments can account for any of this shortfall, we use our calculated TFP growth adjustments above to construct an intangible-adjusted TFP series. As discussed above, this series is substantially higher than the measured values in the post-slowdown period. Adjusted annual TFP growth over 2005-2017 was 0.85%, up from the measured value of 0.40%. However, the adjusted series was also larger before the productivity slowdown, averaging 2.53% per year from 1995-2004, higher than the measured value of 1.63%. Thus the productivity slowdown also exists in the adjusted series. Indeed, at 1.68% per year it is larger than the measured slowdown of 1.23%.⁵⁹ Of course, this analysis assumes that the multiplier for intangibles—the amount of intangibles associated with each dollar of tangible investments—is constant throughout the period. If it is higher in recent periods, mismeasurement would be greater in recent periods, and vice-versa.

⁵⁷ We calculate this as the difference between the average quarterly TFP growth values for 1995-2004 and 2005-2017, respectively. We then annualize this average difference.

⁵⁸ At the end of 2017, counterfactual TFP would be 1.235 ($= 1.00407^{52}$) times its level at the end of 2004, where 0.407% was average quarterly TFP growth over 1995-2004. Measured TFP was instead 1.052 times larger in 2017. Thus, assuming observed labor and capital inputs remain as observed, counterfactual GDP at the end of 2017 would be 1.174 ($= 1.235/1.052$) times larger than the observed value of \$19.83 trillion. The difference, \$3.46 trillion, is 17.4% of \$19.83 trillion.

⁵⁹ For the adjusted series, counterfactual TFP is 1.388 ($= 1.00632^{52}$) times its end of 2017 level at the end of 2004, where 0.632% was average quarterly adjusted TFP growth over 1995-2004. Measured TFP was 1.15 times larger in 2017 in adjusted terms.

Note that the fact that intangibles, at least in the simplest formulation with a constant multiplier, do not explain the productivity slowdown (and actually somewhat deepen it) does not imply that intangibles' influence on productivity and GDP is small. Adjusted TFP (again holding observed labor and tangible capital constant) is 15.9% higher than observed at the end of 2004, and 22.8% higher than observed at the end of 2017. To put it in other words, in addition to all the measured assets, including housing, property plant and equipment, and so on that the U.S. economy produced over the past several decades, it also produced trillions of dollars' worth of unmeasured intangible capital. It is just that the long-lived nature of intangibles' effects itself causes these upward adjustments to be differenced out when seeking to explain the slowdown.⁶⁰

8. Conclusion

Our approach has shown how accounting for intangible investments correlated with measurable ones can meaningfully change estimates of productivity growth and dynamics. Both capital inputs and capital outputs are affected by intangibles. Productivity is underestimated in cases where the growth rate of investment (which contributes to output) exceeds the growth rate of capital services (inputs), and overestimated when the investment growth rate is lower. The first of these effects tends to dominate early in the capital accumulation cycle, when firms and organizations devote resources to building unmeasured intangible capital. The second effect dominates later, when these unmeasured assets generate capital services that increase measured output. Finally, when the capital accumulation reaches steady state, there is no longer any mismeasurement. These dynamics generate what we call the Productivity J-curve.

Because technological improvement often leads to the creation of new capital varieties and necessitates investment in intangible complements, the introduction of a new GPT often causes such a

⁶⁰ Our empirical framework can capture firm-specific intangible investments. However, if there are intangibles built at the industry- or economy-wide levels (perhaps by governments or other organizations that can solve free-riding problems), our empirics will miss them even if they create a large J-curve..

J-curve to occur. We show how this has been the case for IT-related capital in recent decades, for which our calculations suggest that trillions of dollars of intangibles output has been produced but not counted in the national income accounts. There is also some evidence that the phenomenon appears to have begun again, very recently, in AI-related intangible investments.

The mere presence of intangible correlate investment is not a guarantee of the existence of the Productivity J-curve. Although R&D investments are large and are associated with large intangibles, we find that mismeasurement related to R&D investments has a negligible effect on the estimation of productivity growth.⁶¹ On the other hand, computer hardware, and to a greater extent software, have a large effect. The difference reflects the interaction of three quantities: the investment share of output of the asset type, the intangible correlate quantity and adjustment costs per unit of observable investment, and the difference between the growth rate of investment in the asset and the growth rate of capital services. In the case of R&D, the investment share is large and the intangible multiplier is historically larger than two. But, as a mature asset type, the difference between the growth rate of R&D investment and the growth rate of capital is not very large. Software, in contrast, has a meaningfully large investment share of output, has intangible multipliers close to 10, and the investment growth rate in software has often exceeded the growth of capital services overall.

By integrating aspects of Q-theory of investment and traditional growth accounting methods, we offer a means of adjusting the productivity statistics such that new, seemingly omnipresent GPTs might show up in the productivity statistics. Assuming that capital markets price corporate securities efficiently, then market value regressions can estimate the value of intangible correlates and adjustment costs per unit of observable capital. The forward-looking nature of market valuation means that lags in capital services would rationally be considered correctly in expectation. Of course, these multipliers reflect a risk-adjusted discounted expected value of the accumulated asset stock which

⁶¹ With a minor deviation present in the late 1990s and early 2000s.

might to come to bear. The mismeasurement issues might accordingly be sensitive to differences in the timing of expected returns. Lower interest rates, for example, could encourage longer duration investments and therefore prolong the effects of the J-Curve. This investment timing component of productivity mismeasurement is left to future research.

The J-Curve method also suggests an indicator of whether or not a new technology is indeed a *general-purpose technology*. If measures of the investment in a given new technology fail to generate economically significant intangibles, that particular technology at that moment in time would not qualify as general-purpose. Equation (17) offers a framework to use firm-level estimates of technological capital (e.g. AI, IT, or robotics capital) to determine if productivity growth estimation is need adjustment in the presence of that new capital type. This framework also might inform whether or not intangible capital accounts for the wide differences between frontier and median productivity firms (Andrews, Criscuolo, and Gal 2015).

The Productivity J-curve explains why a productivity paradox can be both a recurrent and expected phenomenon when important new technologies are diffusing throughout the economy. Adjusting productive processes to take advantage of new types of capital requires the kind of investments the statistics miss. In future, after making appropriate adjustments accounting for the Productivity J-curve, we can see new technologies everywhere *including* the productivity statistics.

References

- Acemoglu, Daron, and James A. Robinson. 2013. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Broadway Business.
- Acemoglu, Daron (2002). Technical change, inequality, and the labor market. *Journal of economic literature*, 40(1), 7-72.
- Aghion, Philippe, Benjamin F Jones, and C Jones. 2017. "Artificial Intelligence and Economic Growth." *National Bureau of Economic Research*, no. w23928.
- Agrawal, Ajay, John McHale, and Alex Oettl. 2018. "Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth." In *Economics of Artificial Intelligence*. University of Chicago Press. <http://www.nber.org/papers/w24541>.

- Allen, Robert C. 2009. "Engels' Pause: Technical Change, Capital Accumulation, and Inequality in the British Industrial Revolution." *Explorations in Economic History* 46 (4):418–35.
- Andrei, Daniel, William Mann, and Nathalie Moyon. 2018. "Why Did the Q Theory of Investment Start Working?" *Working Paper*.
- Andrews, Dan, Chiara Criscuolo, and Peter N Gal. 2015. "Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence from OECD Countries." *Working Paper*, 39.
- Barro, Robert. 1998. "Notes on Growth Accounting." *Journal of Economic Growth* 137 (June):119–37.
- Basu, Susanto, and John G. Fernald. 1997. Returns to Scale in US Production: Estimates and Implications." *Journal of Political Economy* 105 (2): 249-283.
- Basu, Susanto, John G. Fernald, Nicholas Oulton, and Sylaja Srinivasan. 2003. "The case of the missing productivity growth, or does information technology explain why productivity accelerated in the United States but not in the United Kingdom?." *NBER macroeconomics annual* 18: 9-63.
- Bresnahan, Timothy. 2010. "General Purpose Technologies." *Handbook of the Economics of Innovation* 2 (1):761–91.
- Bresnahan, Timothy F., and M. Trajtenberg. 1995. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics* 65 (1):83–108.
- Bresnahan, Timothy F, Shane Greenstein, David Brownstone, and Ken Flamm. 1996. "Technical Progress and Co-Invention in Computing and in the Uses of Computers." *Brookings Papers on Economic Activity: Microeconomics*, 1–83.
- Brynjolfsson, Erik, Felix Eggers, and Avinash Gannamaneni. 2018a. "Measuring Welfare with Massive Online Choice Experiments: A Brief Introduction." *American Economic Association: Papers and Proceedings* 108:473–76.
- . 2018b. "Using Massive Online Choice Experiments to Measure Changes in Well-Being." *NBER Working Paper Series*.
- Brynjolfsson, Erik and Kristina McElheran. 2016. "The Rapid Adoption of Data-Driven Decision-Making." *American Economic Review: Papers and Proceedings* 106(5): 133-39.
- Brynjolfsson, Erik and Lorin M. Hitt. 2000. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *Journal of Economic Perspectives* (4):23-48.
- Brynjolfsson, Erik, Lorin M. Hitt, and Shinkyu Yang. 2002. "Intangible Assets: Computers and Organizational Capital." *Brookings Papers on Economic Activity* (1):137–98.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2017. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." In *Economics of Artificial Intelligence*. University of Chicago Press.

- Bughin, Jacques, Eric Hazan, Sree Ramaswamy, Michael Chui, Tera Allas, Peter Dahlstrom, Nicolas Henke, and Monica Trench. 2017. "Artificial Intelligence – The Next Digital Frontier?" *Artificial Intelligence*, 80.
- Byrne, David M., John G. Fernald, and Marshall B. Reinsdorf. 2016. "Does the United States Have a Productivity Slowdown or a Measurement Problem." *Finance and Economics Discussion Series* 2016 (017):1–74.
- Byrne, David M., Stephen D. Oliner, and Daniel E. Sichel. 2013. "Is the Information Technology Revolution Over?" *SSRN Electronic Journal*.
- Cockburn, Iain M, Rebecca Henderson, and Scott Stern. 2018. "The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis." *National Bureau of Economic Research* No. w24449 (September).
- Corrado, Carol, Charles Hulten, and Daniel Sichel. 2009. "Intangible Capital and U.S. Economic Growth." *Review of Income and Wealth*.
- Crafts, Nicholas. 2004. "Steam as a General Purpose Technology: A Growth Accounting Perspective." *Economic Journal* 114 (495):338–51.
- Crouzet, Nicolas and Janice Eberly. 2018. "Understanding Weak Capital Investment: the Role of Market Concentration and Intangibles." Working Paper.
- David, Paul A. 1990. "The Dynamo And The Computer: An Historical Perspective On T." *American Economic Review* 80 (2):355.
- Eichengreen, Barry. 2015. "Secular Stagnation: The Long View." *American Economic Review* 105 (5):66–70.
- Eisfeldt, Andrea L., and Dimitris Papanikolaou. 2013. "Organization Capital and the Cross-Section of Expected Returns." *Journal of Finance* 68 (4):1365–1406.
- . 2014. "The Value and Ownership of Intangible Capital." *American Economic Review* 104 (5):189–94.
- Fernald, John. 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." *Federal Reserve Bank of San Francisco*.
<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.310.1766&rep=rep1&type=pdf>.
- Furman, Jason, and Robert Seamans. 2018. "AI and the Economy," 1–33.
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3186591.
- Goolsbee, Austan D, and Peter J Klenow. 2018. "Internet Rising , Prices Falling: Measuring Inflation in a World of E-Commerce." *AEA Papers and Proceedings* 108:488–92.
- Gordon, Robert J. 2015. "Secular Stagnation: A Supply-Side View." *American Economic Review* 105 (5):54–59.
- Hall, Bronwyn H. 1993. "The Stock Market's Valuation of R&D Investment during the 1980's." *The American Economic Review* 83 (2):259–64.

———. 2006. “R&D, Productivity, and Market Value.” *IFS Working Papers, Institute for Fiscal Studies (IFS) 06/23*.

Hall, R. E. 2004. “Measuring Factor Adjustment Costs.” *The Quarterly Journal of Economics* 119 (3):899–927.

Hall, Robert E. 2001. “The Stock Market and Capital Accumulation.” *The American Economic Review* 91 (5):1185–1202.

<http://www.jstor.org/stable/2677921>
<http://www.jstor.org/stable/pdfplus/2677921.pdf?acceptTC=true>.

Hall, Robert E. 2000. “E-Capital: The Link between the Stock Market and the Labor Market in the 1990s.” *Brookings Papers on Economic Activity* 2:73–118.

Haskel, Jonathan, and Stian Westlake. 2017. *Capitalism without Capital: The Rise of the Intangible Economy*. Princeton University Press.

Hayashi, Fumio. 1982. “Tobin’s Marginal q and Average q: A Neoclassical Interpretation.” *Econometrica* 50 (1):213–24.

Hayashi, Fumio, and Tohru Inoue. 1991. “The Relation Between Firm Growth and Q with Multiple Capital Goods: Theory and Evidence from Panel Data on Japanese Firms.” *Econometrica* 59 (3):731–53.

Helpman, Elhanan & Trajtenberg, Manuel (1994). A time to sow and a time to reap: Growth based on general purpose technologies (No. w4854). National Bureau of Economic Research.

Henderson, Rebecca. 1993. “Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Alignment Equipment Industry.” *The RAND Journal of Economics* 24 (2):248–70.

Himel, Samuel, and Robert Seamans. 2017. “Artificial Intelligence, Incentives to Innovate, and Competition Policy.” *Antitrust Chronicle* 1 (3).

Houthakker, Hendrik S. "The Pareto distribution and the Cobb-Douglas production function in activity analysis." *The Review of Economic Studies* 23, no. 1 (1955): 27-31.

Jacobs, Bas, and Richard Nahuis. "A general purpose technology explains the Solow paradox and wage inequality." *Economics Letters* 74, no. 2 (2002): 243-250.

Jorgenson, Dale Weldeau, and Kevin J. Stiroh. 2000. “Raising the Speed Limit: U.S. Economic Growth in the Information Age.” *Brookings Papers on Economic Activity* 2000 (1):125–210.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. “Technological Innovation, Resource Allocation, and Growth.” *Quarterly Journal of Economics* 132 (2):665–712.

Lipsey, Richard, Kenneth Carlaw, and Clifford Bekar. 2006. “Economic Transformations: General Purpose Technologies and Long-Term Economic Growth.” *The Economic History Review* 59 (4):881–82.

- Lucas, Robert E. 1967. "Adjustment Costs and the Theory of Supply." *Journal of Political Economy* 75 (4, Part 1):321–34.
- Marrano, Mauro Giorgio, Jonathan Haskel, and Gavin Wallis. 2009. "What Happened to the Knowledge Economy? ICT, Intangible Investment, and Britain's Productivity Record Revisited." *Review of Income and Wealth*.
- Magee, Stephen P. 1973. "Currency Contracts, Pass Through and Devaluation," *Brookings Papers on Economic Activity*, Vol. 1, pp. 303-325.
- Oliner, Stephen D., Daniel E. Sichel, and Kevin J. Stiroh. 2008. "Explaining a Productive Decade." *Journal of Policy Modeling* 30 (4):633–73.
- Oliner, Stephen D, and Daniel E Sichel. 2000. "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives*.
- Peters, Ryan H., and Lucian A. Taylor. 2017. "Intangible Capital and the Investment-q Relation." *Journal of Financial Economics* 123 (2):251–72.
- Rose, Andrew K., and Janet L. Yellen. 1989. "Is There a J-Curve?" *Journal of Monetary Economics* 24 (1):53–68.
- Rosenberg, Nathan, and Manuel Trajtenberg. 2004. "A General-Purpose Technology at Work: The Corliss Steam Engine in the Late-Nineteenth-Century United States." *The Journal of Economic History* 64 (1):61–99.
- Saunders, Adam, and Erik Brynjolfsson. "Valuing Information Technology Related Intangible Assets." *Mis Quarterly* 40, no. 1 (2016).
- Solow, R.M. 1957. "Technical Change and the Aggregate Production Function." *The Review of Economics and Statistics* 39 (3):312–20.
- Solow, Robert. 1987. "We'd Better Watch Out." *New York Times Book Review*, (July 12).
- Solow, Robert M. 1956. "A Contribution to the Theory of Economic Growth." *The Quarterly Journal of Economics* 70 (1):65.
- Summers, Lawrence H. 2015. "Demand Side Secular Stagnation." In *American Economic Review*, 105:60–65.
- Syverson, Chad. 2017. "Challenges To Mismeasurement Explanations for the U.S. Productivity Slowdown." *Journal of Economic Perspectives* 31 (2):165–86.
- Tambe, Prasanna, Lorin M. Hitt, and Erik Brynjolfsson. 2012. "The Extroverted Firm: How External Information Practices Affect Innovation and Productivity." *Management Science* 58 (5):843–59.
- Trajtenberg, Manuel. 2018. "AI as the Next GPT: A Political-Economy Perspective." 24245. Working Paper Series.
- Wildasin, David E. 1984. "The q Theory of Investment with Many Capital Goods." *American Economic Review*.

Yang, S, and E Brynjolfsson. 2001. "Intangible Assets and Growth Accounting: Evidence from Computer Investments." *Unpublished Paper. MIT* 85 (61):28.

Chapter 3 — Machine Learning and Occupational Change⁶²

1. Introduction

Until very recently, the idea of automated prediction technologies guiding motor vehicle traffic without human intervention was little more than science fiction. Now, new developments in machine learning have led researchers to find this and many other AI challenges within reach. The basic neural net models and algorithms thriving today are decades old (White and Rosenblatt 1963; Rumelhart, Hinton, and Williams 1986; LeCun et al. 1998). But now the widespread availability of inexpensive, scalable computing resources and the proliferation of large scale training data to build Deep Neural Nets (DNNs) with viable commercial applications has stirred rapid growth in DNN research and investment.⁶³ According to the 2017 AI Index Annual Report, by 2016 the number of AI papers had increased nine-fold over its 1996 level, the number of AI systems startups had increased by a factor of 14 since 2000, and venture capitalists had invested six times the capital in AI startups as in 2000.⁶⁴ AI is poised to transform occupations across the economy much as rules-driven information technology did in prior years.

AI is a general-purpose technology (GPTs). Like all GPTs, AI is pervasive, improves over time, and spawns complementary innovations (Bresnahan and Trajtenberg 1995; Bresnahan 2010; Brynjolfsson, Rock, and Syverson 2017). In particular, the recent advances in AI depend on improvements in machine learning (ML). ML systems are increasingly adopted in virtually every

⁶² Joint with Erik Brynjolfsson (MIT and NBER) and Tom Mitchell (CMU). Acknowledgments: We thank Zanele Munyikwa and Morgan Frank for their suggestions and help on this project. We are grateful to Betsey Stevenson and Shane Greenstein for their excellent comments as well. We also thank participants in seminars at MIT, Harvard Business School, the American Economic Association annual meeting (2018), The NBER Economics of Artificial Intelligence Conference (2018), and the Workshop for Information Systems and Economics for useful comments on this research, and the MIT Initiative on the Digital Economy for generous funding.

⁶³ Deep learning (or Deep Neural Nets) describes a class of models within the field of Machine Learning (ML), which is itself situated within the broader field of Artificial Intelligence (AI).

⁶⁴ <http://cdn.aiindex.org/2017-report.pdf>

industry. Perhaps more than any other technology, the essence of ML is that it improves with better training data and better algorithms. The practical application of AI and ML almost always requires complementary investments in intangible capital – new business processes, redesign of jobs, training and human capital investments – and enables the development of new goods and services like digital concierges, data center energy optimization, and even fake news filtration.⁶⁵ As noted by Furman and Seamans (2018), recent work on the effects of AI on innovation (Cockburn, Henderson, and Stern 2018), economic growth (Aghion, Jones, and Jones 2017; Agrawal, McHale, and Oettl 2018), productivity (Brynjolfsson, Rock, and Syverson 2018a, 2018b), public policy (Furman and Seamans 2018; Trajtenberg 2018; Goolsbee 2018), and market value (Rock 2019) suggest that the coming impact of AI will be widespread across the economy.⁶⁶

Progress over the past decade has been especially fast using techniques of deep learning. These fall within the machine learning paradigm of supervised learning, where a model is constructed to map observable input features into observable output labels or values (Hinton, Bengio, and LeCun 2015). Deep learning is not a unique solution concept for this kind of task. Linear regression, for example, is commonly used for the same purpose. Similar to linear regression models, a deep learning model will apply different weights and functional transformations to the input values in order to predict or classify the outputs. These models are complicated and it can be difficult to explain why they make a particular prediction.⁶⁷

⁶⁵ <https://ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html>;
<https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/>;
<https://www.technologyreview.com/s/609717/can-ai-win-the-war-against-fake-news/>

⁶⁶ Furman and Seamans (2018) point out in particular that much of the current research on the Economics of AI is in NBER *Economics of Artificial Intelligence* volume (edited by Agrawal, Gans, and Goldfarb).

⁶⁷ For a brief overview, Appendix A describes the types of machine learning paradigms.

The key advantage of this approach over rules-driven software though is that they do not require a pre-determined recipe or series of codified instructions to map input values to outputs. The mapping function is *learned* automatically via an optimization algorithm.⁶⁸ This has at least two major economic implications. First, DNNs help overcome Polanyi’s Paradox that “we know more than we can tell”. This has been a major obstacle to automation by software in the past (Polanyi 1966; Autor 2014). For instance, while humans can easily recognize faces, it is difficult to write formal rules for how to do it. Until recently, that made it difficult for machines to do this task. However, now, using DNNs, machines have learned to do it very well. Furthermore, the machine learning approach has proven surprisingly effective for a wide variety of other tasks, from recommending products to consumers and placing ads on websites, to translating speech and diagnosing cancer. In many cases, machine learn systems now match or even surpass humans doing the same tasks (see e.g. Esteva et al. (2017); Graves, Mohamed, and Hinton (2013)).⁶⁹ One can think of these applications of AI as reducing the cost of various forms of prediction, leaving judgment tasks to humans (Agrawal, Gans, and Goldfarb 2017, 2018). The widespread applicability of “prediction machines” has potentially sizeable implications for changing the nature of economic activity. In particular, because tasks like these are important parts of many occupations, machine learning is likely to have significant effects on content, wages and demand of for many types of workers.

At the same time, there are also many tasks people do that cannot be done by machines at the current frontier of automation. Artificial General Intelligence (or AGI), which seeks to match humans in all tasks, remains a distant goal. Therefore, it is useful to distinguish between the set of tasks mostly like to be affected by machine learning and those which are not likely to be affected, and assess the likely implications for tasks, occupations, and the economy.

⁶⁸ Typically, Stochastic Gradient Descent (SGD) or some variant of SGD.

⁶⁹ Also see the AI Index at <http://www.aiindex.org>

Accordingly, in this paper we seek to assess how and where ML will affect the workforce, extending the ideas of Brynjolfsson and Mitchell (2017) and Brynjolfsson, Mitchell, and Rock (2018). Brynjolfsson and Mitchell (2017) outline a series of factors that relate machine learning and labor demand, offering a 23-question rubric to evaluate each task in the economy for its “Suitability for Machine Learning” or SML. Brynjolfsson, Mitchell, and Rock (2018) provide a first implementation of that rubric using the O*NET database from the U.S. Bureau of Labor Statistics.

This paper expands and improves that implementation in two ways. First, we build a significantly larger data sample using a new crowdsourcing platform. This enables us to we create a variety of measures of SML and SML-related concepts at the level of regions, industries, and firms (in addition to occupations). Second, we build on the summary statistics in Brynjolfsson, Mitchell, and Rock (2018) to provide a more detailed description of where ML is likely to have its greatest impact across the economy, looking at regions, industries, firms, and occupations. The data from this new implementation of the rubric is less noisy and sourced from a more reliable crowdsourced panel. Additionally, we present a model based on Autor, Levy, and Murnane (2003) which describes some of the task-level effects of occupational re-design following ML automation. We supplement the original six factors relating machine learning and labor demand in Brynjolfsson and Mitchell (2017) with the findings from that model.

While building on the ideas in Brynjolfsson, Mitchell, and Rock (2018) and related work by Felten, Raj, and Seamans (2018), these additional measures help contextualize concerns about occupations which are generally exposed to automation technologies. Our results for the SML rubric suggest that very few (if any) occupations can be fully automated by ML, but most occupations have at least one task that is exposed to a possible ML implementation. This suggests that realizing the potential gains to the ML GPT will require substantial investment in the redesign of occupations and business processes. Our rubric is designed to evaluate technological potential in occupations in the

intermediate term (a few years to a decade or so in the future). Nevertheless, some predictions based on which types of tasks are suitable for machine learning are possible.

An important limitation of our approach is that while we assess the effects of ML⁷⁰ on existing tasks, we do not predict the *new* tasks which will be created as technology advances. As stressed by Brynjolfsson and McAfee (2011) and Acemoglu and Restrepo (2017), humans do not simply race against machines, but also work with them to complete tasks they could never have solved before. However, imagining and anticipating these tasks before they exist is far more difficult than analyzing existing tasks. For the purposes of this paper, we leave this challenge to savvy entrepreneurs, creative managers, and perhaps more skilled researchers.

2. Automation, Artificial Intelligence, and Labor Demand

Impressive as they are, the recent growth trends in AI-related investment and research constitute a relatively small fraction of the economy. It is too soon to conclusively measure any major economic effects of AI technologies, much less declare any kind of transformational change. Here we borrow from the economics and management literature on the economic effects of more general trends in information technology and automation. The base unit of analysis in our technological rubrics is the task. Jobs are then bundles of tasks, and those bundles are a design choice by employers. This follows the work of Autor, Levy, and Murnane (2003) (ALM), which explains shifts in labor demand by analyzing the impact of IT-related automation at the task-level. In particular, ALM offers a model whereby some tasks can be substituted for by capital and others cannot. They find that computer capital has reduced the relative demand for routine types of work tasks. Routineness in this context is a measure of how possible it is to describe a given task with rules, indicating a viable automation opportunity for standard types of software. That is, these tasks were not subject to Polanyi's Paradox.

⁷⁰ With rubrics on robotics and rules-based systems under development as well.

Acemoglu and Restrepo (2017) adopt a task-level approach to understanding how the incentives to automate or employ humans are contingent on the available supply of labor for doing tasks (and the prevailing wage), the research costs to develop automation capital for existing tasks, and the research costs to develop new tasks for human workers. These trade-offs might help to understand, for example, the Kaldor (1961) stylized fact that, until recently, the labor and capital shares of income were fairly stable even as more and more tasks were automated. We are concerned principally with the piece describing the research costs of machine learning-based automation and less on the creation of new tasks, though this is obviously a crucial element of the impact of ML technology on labor demand. Much of the impact of IT, another GPT, resulted from the follow-on complementary innovation and reconfiguration of economic activity (Bresnahan et al. 1996; Greenstein and Prince 2005). Our analysis centers on the task components which are most suitable for ML, for rule-based systems, and for robotics.

While ML and robotics can automate a different set of tasks compared to rules-based IT, some of the mechanisms driving shifts in labor demand may be similar. Indeed, earlier waves of automation also had a major impact on productivity, wages, employment, and investment (Acemoglu and Restrepo 2017a; Acemoglu and Autor 2011; Autor and Dorn 2013). While we might not see the same direction of skill-biased technological change and polarization that resulted from earlier waves of automation, it is likely that demand for human tasks easily substituted by ML capital will decline relative to demand for less suitable for machine learning tasks. Furthermore, because ML is differentially suited to different purposes, the distribution of the technology's impact across occupations will likely be uneven (Levy 2017).

Concern about the effect of AI on the “Future of Work” and how to deal with the impending automation of huge swaths of the workforce are growing. Some studies suggest that nearly half of the work activities in the U.S. could be automated with existing technology (Bughin et al. 2017). If

workers could be completely automated, it is possible that the economic effects would immiserate laborers in the long run (Benzell et al. 2016).⁷¹ Yet seemingly in spite of the pace of technological innovation in the past century, the vast majority of prime-age workers are still employed in every economy around the world. This is largely because of the strong complementarities between automation technologies and the tasks that they do not replace (Autor 2015). For example, Bessen (2016) shows that employment of bank tellers initially *increased* after the adoption of the ATM. The tellers were put to work emphasizing tasks that the machines were poorly suited to do. Bessen (2018) argues that the elasticity of labor demand is a major factor too in the overall effects of an automation technology. With elastic demand for the non-automated tasks, we might expect employment in non-automated tasks to increase as the costs of complementary automated tasks drop.

Brynjolfsson and Mitchell (2017) identify demand elasticity as one of six determinants of how ML and labor demand and wages are related, adding as well

2) complementarities,

3) substitution,

4) supply elasticity,

5) income elasticity, and

6) reengineering.

Substitution effects are the primary worry of the neo-Luddites, and the complementarity effects are similar to those in Autor (2015). Brynjolfsson and Mitchell (2017) point out that cheaper production via ML technology has the effect of increasing wealth as well. The effects of greater income change the production function directly through employer demand for labor and indirectly

⁷¹ Acemoglu and Restrepo (2017) refer to this as the “horse equilibrium”.

through consumer demand for products. ML can also affect labor supply elasticity; an inelastic labor supply for complementary non-automatable work could lead to higher wages in that task. Lastly, reengineering the task bundles into new jobs may be an important source of innovation. Firms that take advantage of better job designs that allow machines and humans to specialize will generate output more efficiently. In accordance with the LeChatelier Principle (Milgrom and Roberts 1996), production inputs are typically more elastic in the long-run.

3. Why might this time be different?

DNNs and similar architectures⁷² are appropriate for solving types of problems which other AI and machine learning approaches perform poorly.⁷³ The reason that these structures are so useful in prediction tasks is because the learned parameters in these models can approximate nonlinear functions with high dimensionality. Input features are (in many cases) automatically transformed to represent data in better ways, making it possible to use DNNs in the case where “it is easier to collect the data than to explicitly write the program”.⁷⁴

Models of this type with human-level (and in some cases superhuman) performance have been trained to classify images, translate speech and text, and even to prove mathematical theorems (e.g. Loos et al. 2017). Often DNN algorithms are combined with reinforcement learning frameworks and other types of machine learning to build high performance composite models. Google DeepMind’s AlphaZero algorithm achieved superhuman performance in chess, shogi, and Go within 24 hours of training combining a number of different techniques (Silver et al. 2017). Given an appropriate domain and favorable contextual factors, these algorithms can automate many work activities where the goal is to map inputs to outputs. Recent advances (see e.g. Goodfellow et al. 2016) have implications for

⁷² Including but not limited to convolutional neural nets (CNNs), recurrent neural nets (RNNs), long short-term memory (LSTMs), generative adversarial networks (GANs), and ordinary artificial neural nets (ANNs).

⁷³ For example, linear regression, support vector machines, or tree-based models, let alone rule-based systems.

⁷⁴ <https://medium.com/@karpathy/software-2-0-a64152b37c35>

prediction and classification tasks in fields ranging from cryptography (Shokri and Shmatikov 2016), to medical imaging (Litjens et al. 2017), and even to archaeology (Guyot, Hubert-Moy, and Lorho 2018). A particular combination of DNNs with reinforcement learning techniques called “inverse” reinforcement learning means that robots can be programmed with videos of the tasks they are meant to perform (Abbeel and Ng 2004; Finn et al. 2016). Machine learning is therefore a method for discovering better programs (or as Tesla Director of AI Andrej Karpathy calls it, “Software 2.0”).⁷⁵ Though particular trained models might not generalize to different tasks, the class of models is potentially useful for a large contingent of oft-repeated but hard to codify tasks which so far have eluded automation.

4. A Model of Machine Learning Automation

The model in Autor, Levy, and Murnane (2003) is a useful starting point for understanding the incentives to re-bundle tasks into new occupations as automation technologies improve. We contrast the ALM model with a constrained version of the same firm problem in the constant elasticity of substitution case. Following ALM, assume production Q (equation 1) is a constant elasticity of substitution function with total factor productivity parameter A , capital inputs K (including intangible capital used in machine learning – like data), labor in efficiency units devoted to machine learnable (SML) tasks L_{ML} , labor in efficiency units devoted to other (non-SML) tasks L_o , share parameter α where $1 \geq \alpha \geq 0$ and elasticity of substitution $1/(1-\sigma)$. Capital and labor suitable for machine learning are assumed to be perfect substitutes. Workers have productivity endowments in ML tasks and other tasks respectively of $E_i = [ML_i, O_i]$ where $1 \geq ML_i, O_i > 0 \forall i$ and the labor supplied is $L_i = [\lambda ML_i, (1-\lambda)O_i]$, $\lambda \in [0,1]$. Our production function Q is then:

$$(1) \quad Q = A(\alpha(K + L_{ML})^\sigma + (1 - \alpha)L_o^\sigma)^{\frac{1}{\sigma}}$$

⁷⁵ In the blogpost from the previous footnote.

where at the margin income-maximizing workers are indifferent between offering ML-able services and other services. If the marginal user cost of capital (Jorgenson 1963) to develop ML applications r sets the wage for ML-able labor W_{ML} , we have wage equations (2) and (3) from first order conditions:

$$(2) \quad W_{ML} = \alpha A^\sigma \left(\frac{Q}{K + L_{ML}} \right)^{1-\sigma} = r$$

$$(3) \quad W_o = (1 - \alpha) A^\sigma \left(\frac{Q}{L_o} \right)^{1-\sigma}$$

and then the ratio of capital and ML labor to other labor is:

$$(4) \quad \theta \equiv \frac{K + L_{ML}}{L_o} = \left(\frac{(1 - \alpha) W_{ML}}{\alpha W_o} \right)^{\frac{1}{\sigma-1}}$$

Letting σ approach negative infinity (i.e. a Leontief production function) with constraints on task groupings illustrates how suboptimal bundling of tasks can block potential productivity gains. Under the conditions presented above, if the rental cost of ML capital went to a marginal value of 0, the workers would prefer to all switch into doing only the tasks that ML cannot. While nonzero in general, marginal costs of building these models need not exceed the costs of computation in specific cases, a cost which is already low and continues to fall very rapidly. The final changes in the ratio of capital and total labor employed, however, depend on the factor shares of different types of inputs and the elasticity of substitution between the different inputs. The Leontief production case demonstrates a version of extreme complementarity, but this need not be the case. Further, while more tractable and useful for this illustration, the elasticity of substitution in multifactor production functions need not be the same across inputs.

Now assume that firms, perhaps because of historical business processes, implementation frictions, and/or adjustment costs, only offer labor contracts that have some preset mixture of SML

and non-SML tasks. This is a pooled hiring model. Formally we have a modification to the first order condition such that it is now constrained (this constrained problem is not in the ALM model). The firm solves:

$$(5) \quad \max_{L_{ML}, L_o, K} Q(L_{ML}, L_o, K) - w_{ML}L_{ML} - w_oL_o - rK$$

$$s. t. L_{ML} - \beta L_o = 0 \text{ and } \beta \geq 0$$

This carries the additional constraint (in addition to feasible production) that the ratio of machine learnable labor and non-SML labor is fixed at a constant β in the short-run. We assume that the costs to reorganize job task bundles in this case might be prohibitively costly, so the firm in the short-run is bound by the constraint. The firm only has two choice variables: K and total labor units L (let this be set equal to the quantity of machine learnable labor). Ignoring the trivial solution (in our case) that no labor of any kind is hired, the new wage conditions are:

$$(6) \quad \frac{\partial Q}{\partial K} - r = 0 = \alpha A^\sigma \left(\frac{Q}{K + L_{ML}} \right)^{1-\sigma} - r$$

$$(7) \quad \frac{\partial Q}{\partial L} = W = (1 - \alpha) \left(\frac{A}{\beta} \right)^\sigma \left(\frac{Q}{L_o} \right)^{1-\sigma} + A^\sigma \left(\frac{Q}{K + L_{ML}} \right)^{1-\sigma} = \frac{(1 - \alpha)(A)^\sigma}{\beta} \left(\frac{Q}{L} \right)^{1-\sigma} + r$$

$$(8) \quad r = \alpha A^\sigma \left(\frac{Q}{K + L_{ML}} \right)^{1-\sigma}$$

Therefore in (8) the marginal product of capital and SML labor is set by the interest rate, but because firms must hire some quantity of SML labor with all quantities of labor, the wages paid to non-SML labor are multiplied by a factor of $\beta^{-\sigma}$ relative to the marginal product of non-SML labor and added to the interest rate (7). That is, labor earns at least the interest rate and then additional wages depending on the quantity Q and the fixed ratio of labor types.

The first term in equation (7) shows that each additional increase in how many units of SML labor must be hired relative to non-SML labor hired drives down the wage at any quantity hired. In the case of full automation and zero interest rates, wages paid to SML labor would be zero and the wage paid to non-SML labor would equal the marginal product of non-SML labor, but scaled by a power of the fixed ratio. In this case the SML labor might be freely disposed (See Benzell and Brynjolfsson (2018) for a related model and some macroeconomic implications). The opportunity cost of necessitating that human workers do what machines could be doing is increasing in the job design parameter β . In the case that the interest rate is less than the difference between the unconstrained non-SML labor wage and the ratio-attenuated wage for labor (the part of the constrained scenario wage separate from the interest rate), workers specializing in the non-SML task are paid less with the pooled equilibrium wage. Figure A1 in the Appendix shows this for a Cobb-Douglas demand function and common linear labor supply curve for both types of labor.⁷⁶ We can consider as well a more flexible adjustment cost function instead of a hard constraint in the ratio of labor units hired. Using the same Q as above, but also penalizing firms for adjusting their labor force in any direction we have:

$$(9) \max_{L_{ML}, L_O, K} Q(L_{ML}, L_O, K) - w_{ML}L_{ML} - w_O L_O - rK - \lambda C(\|L - \bar{L}\|)$$

Where C is a monotonically increasing and convex function in some norm (e.g. the 2-norm) of the vector of labor unit differences across labor types. λ converts units of C into a cost in terms of Q . Then if the net marginal value of capital equals the net marginal value of SML labor it is because interest rates less adjustment costs of SML labor set the wage for SML labor. Depending on the adjustment cost function, this can also depend on the level of non-SML labor. In particular for the 2-norm:

⁷⁶ Fixed parameters include $Q=100$, $\alpha=\beta=0.5$, and $\frac{3}{4}$ of all SML labor being completed by capital for the Figure A1.

$$(10) \quad r - \lambda \frac{\partial C}{\partial L_{ML}} (||L - \bar{L}||) * \frac{(L_{ML} - \bar{L}_{ML})}{||L - \bar{L}||} = W_{ML}$$

A forward-looking firm will have to consider future adjustment costs of labor (and realistically capital as well) when deciding how to maximize the marginal value of each input. Thus, in the simple case of a hard constraint or in the more general case of labor adjustment costs, re-bundling tasks in new occupations will reward clever entrepreneurs with the difference between their new labor demand curve and the demand they would face under the constraint of the old job design. Even with knowledge of which occupations are likely to be affected by a technology, it is challenging to predict changes in labor demand or wages. The Appendix details an approach that, under some assumptions, can serve as a guide for using SML measures to make such predictions.

In the Leontief case all of the labor effort put toward SML tasks has an opportunity cost in efficiency unit output of the share-adjusted non-SML labor. This is because production capability is the minimum of the adjusted input efficiency units. ML could be doing those tasks, and the firm could increase profit if it were able to reorganize their job design. Note that this provides an incentive for lower average O_i productivity workers to take on the non-SML task, leading to cost disease effects similar to those explored in greater detail in Baumol (1967) and specifically for AI in Aghion, Jones, and Jones (2017) and Benzell and Brynjolfsson (2018). Task measurability can have effects on the types of incentives that should ideally be offered as well (Holmstrom and Milgrom 1991). Consider the situation that machine learning can only target measurable tasks. In this case, at every price labor supply would increase for hard-to-measure tasks as humans formerly working on measurable tasks begin to compete in hard-to-measure tasks instead. This is the kind of change we would expect for a reinforcement learning (RL) system, especially an RL system augmented by embedded DNNs. Both humans and machines have better performance when it is easy to map effort and action into rewards. Within the technological frontier for RL, the tasks and incentive schemes which RL can manage well will generate effects of the style presented in the ALM model above. Effectively the elimination of

less noisy work by machines leads to greater competition for noisy, poorly measured work. Even if the costs of effort were identical (i.e. automation costs were similar to labor effort costs on the margin), the risk-averse agents choose a (weakly) lower optimal total effort quantity, and also receive less for their work. On the margin, the risk-neutral machine competitor has a lower reservation wage for effort in tasks with less noisy outcomes. The main effect, however, is that of ALM: for some tasks, the machine capital can do the same task at a lower price.

5. Measuring Suitability for Machine Learning at the Level of Tasks

As in Brynjolfsson, Mitchell, and Rock (2018), we measure the suitability for machine learning (SML) using the SML rubric in Brynjolfsson and Mitchell (2017), although for this paper we improve and extend it.⁷⁷ The degree to which there is an opportunity for machine learning (and DNNs in particular) to automate a task relies on the characteristics of the particular work activity and the context in which it is done. The first iteration of the SML rubric contained 23 distinct evaluative statements designed to measure the ML potential in a given work activity's context and characteristics. The main idea behind these rubrics is to give primacy to the economic relevant features of the technology in question, whether that be ML, IT, or robotics.⁷⁸ Each of these rubrics is customized to reflect the focal "affordances and constraints" of that particular style of automation capital (Pickering 1993).

We applied the SML rubric to 2,069 detailed work activities (DWAs) in the Bureau of Labor Statistics O*NET content model for all of the occupations in the U.S. In Brynjolfsson, Mitchell, and Rock (2018) we used CrowdFlower⁷⁹, a Human Intelligence Task (HIT) platform, to implement a crowdsourced evaluation of the 2,069 DWAs, rating each DWA for each of the 23 statements on a 5-point scale varying from "Strongly Disagree (1)" to "Strongly Agree (5)". To improve the quality of

⁷⁷ With a slight modification to add two statements describing the physical requirements of the activity.

⁷⁸ Ongoing work is building and implementing the robotics and rules-based software rubrics

⁷⁹ Now known as Figure 8: <https://www.figure-eight.com/>

our responses, each DWA was rated by at least 7 distinct respondents. In this iteration, there are fewer questions included so as to focus more on machine learning. The new 21 question version of the rubric is attached in the Appendix. Respondents were recruited via Amazon Mechanical Turk (henceforth AMT). Each DWA is rated at least 10 times in randomized sequential batches. A series of attention checks are implemented and respondents who fail more than 1 attention check are removed from subsequent batches. Respondents for whom there is no variance in responses across all questions are dropped, as are respondents who supplied certain types of feedback in the provided comment box.⁸⁰ For the remaining responses, we take the mean score for each statement as the value for that DWA-statement tuple.

Higher values correspond to more suitability for machine learning, and vice-versa. These DWA scores for each of the statements in our rubric are many-to-many mapped to 18,112 specific occupation-level tasks (with equal weighting within a task). The tasks in O*NET are weighed by the normalized occupational importance scores provided in the O*NET database and aggregated to form our occupation-level measures.

The final measure is constructed via a simple process. We ask about 4 different types of data (image/video, speech, text, and structured/sensor data) for each one of the activities. These values are unlikely to all be rated highly for each activity, but individually a high value might be sufficient to make a task high SML. We therefore create a measure of data intensity as the maximum of the responses to statements 15 through 18 which correspond to these data types. This data score and the scores from the other questions are averaged to generate our primary SML measure. Table 1 below summarizes the primary SML measure for DWAs, tasks, and jobs; Figure 1 below shows a histogram

⁸⁰ We found some respondents who frequently failed attention checks tended to add similar types of comments at the end of the survey, though the precise comments are not reported to avoid invalidating future surveys. For more information, see: <https://www.wired.com/story/amazon-mechanical-turk-bot-panic/>

of the SML scores. The mean SML measure in DWAs, tasks, and jobs is close to 2.97 with a maximal SML score at the DWA-level of 4.0 and a minimal SML score at 2.133. The standard deviation of SML scores for occupations is about 0.183, approximately 60% of the standard deviation of DWA SML. This reflects the diversification effect of bundling multiple activities into the same occupation. The range between the maximal and minimal SML score for the occupational averages is less than that of DWAs and tasks as well. Table 2 shows the top and bottom 15 occupations ranked by SML. About two-thirds of the variance in activity-level scores can be explained by five principal components. The variance explained plot for those components is depicted in Figure 2A. The remaining third of the variance in activity component scores is relatively evenly spread across components.

TABLE 1 — SUITABILITY FOR MACHINE LEARNING: SUMMARY STATISTICS

	Occupations (950)	Tasks (18,112)	Detailed Work Activities (2,059)
Mean Suitability for			
Machine Learning (SML)	2.984	2.976	2.968
Standard Deviation of SML	0.183	0.292	0.303
Minimum SML	2.576	2.133	2.133
25 th Percentile SML	2.844	2.767	2.767
75 th Percentile SML	3.107	3.167	3.167
Max SML	3.552	4.000	4.000

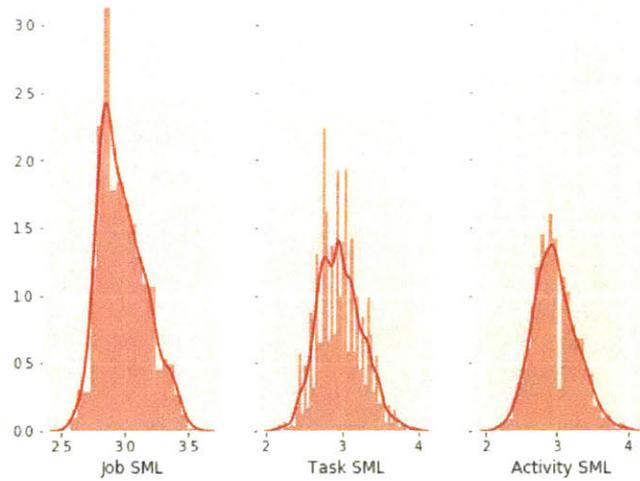


Figure 1- Distribution of Counts of Suitability for Machine Learning (SML) Score for Occupations, Tasks, and Detailed Work Activities

Table 2 - Lowest and Highest 15 SML Score Occupations

Rank	Lowest SML Ranked Occupations	SML	Highest SML Ranked Occupations	SML
1	Clinical Psychologists	2.58	Switchboard Operators, Including Answering Service	3.55
2	Music Composers and Arrangers	2.59	Insurance Claims Clerks	3.50
3	Neuropsychologists and Clinical Neuropsychologists	2.60	Postal Service Mail Carriers	3.50
4	Counseling Psychologists	2.61	Meter Readers, Utilities	3.48
5	Lawyers	2.61	Word Processors and Typists	3.47

6	Product Safety Engineers	2.63	Telemarketers	3.46
7	Industrial-Organizational Psychologists	2.64	Telephone Operators	3.46
8	Coroners	2.64	Police, Fire, and Ambulance Dispatchers	3.44
9	Forensic Science Technicians	2.65	Data Entry Keyers	3.43
10	Fire Investigators	2.65	Couriers and Messengers	3.43
11	Range Managers	2.65	File Clerks	3.43
12	Foresters	2.66	Counter Attendants, Cafeteria, Food Concession, and Coffee Shop	3.43
13	Private Detectives and Investigators	2.66	Payroll and Timekeeping Clerks	3.43
14	Oral and Maxillofacial Surgeons	2.66	Baristas	3.42
15	Biofuels/Biodiesel Technology and Product Development Managers	2.66	Gaming and Sports Book Writers and Runners	3.41

The highest SML occupation in the scores from our first implementation of the rubric is Switchboard Operator, with a SML score of 3.55. They are followed by Insurance Claims Clerks, Mail Carriers, and Meter Readers. Here complementarities may be as important as substitution effects. The typical occupation with a high SML score in the first run tends to be data-intensive and clerical or repetitive in nature, with Word Processors, Telemarketers, Telephone Operators, Dispatchers, Data Entry Specialists, Couriers, File Clerks, Payroll Clerks, and Baristas and Food Service Attendants all showing up in the top 15. In contrast, the lowest SML occupations often have some highly-skilled

interpersonal or otherwise high risk component to their work. Clinical Psychologists are the least SML job, but Composers, Other Psychologists, Coroners, Investigators, and Oral Surgeons also make the low SML list. Economists, for the interested, are relatively close to the lowest of the SML distribution.

An important point, however, is that most occupations do have some component activity or task that is highly rated for SML. This indicates a potential for reorganization. We measure the within-occupation standard deviation of SML as one proxy measure for the “suitability for job re-design” (sdSML). Table 3 below has the results from our first evaluation run of the SML rubric.

Table 3: Lowest and Highest 15 Standard Deviation SML Score Occupations

Rank	Lowest SD SML Occupations	sdSML	Highest SD Occupations	sdSML
1	Fabric Menders, Except Garment	0.35	Ambulance Drivers and Attendants, Except Emergency Medical Technicians	1.03
2	Medical Transcriptionists	0.37	Manufacturing Production Technicians	0.86
3	Title Examiners, Abstractors, and Searchers	0.42	Team Assemblers	0.86
4	Automotive Specialty Technicians	0.43	Childcare Workers	0.84
5	Slaughterers and Meat Packers	0.43	Refuse and Recyclable Material Collectors	0.84
6	Pile-Driver Operators	0.44	Barbers	0.84
7	Agricultural Inspectors	0.44	Counter Attendants, Cafeteria, Food Concession, and Coffee Shop	0.84
8	Dredge Operators	0.44	Bartenders	0.83
9	Electrical Engineering Technologists	0.46	Cleaners of Vehicles and Equipment	0.82
10	Natural Sciences Managers	0.46	Ushers, Lobby Attendants, and Ticket Takers	0.82

			Locker Room, Coatroom, and Dressing Room	
11	Child, Family, and School Social Workers	0.47	Attendants	0.82
12	Animal Scientists	0.47	Nanotechnology Engineering Technicians	0.81
			Merchandise Displayers and Window	
13	Human Resources Managers	0.47	Trimmers	0.80
			Gaming Surveillance Officers and Gaming	
14	Investigators	0.47	Painters, Construction and Maintenance	0.79
15	Park Naturalists	0.47	Cooks, Institution and Cafeteria	0.79

In Table 3 we can see where there might be opportunity. Ambulance Drivers have high variability in the SML values of their constituent tasks, as do Manufacturing Production Technicians and many kinds of food service workers. Also, of note is that Nanotechnology Engineering Technicians have high variability in their task-level SML. These occupations, where it is difficult to imagine full automation by ML, may instead be candidate examples of where the effects of complementarities or incentives to re-bundle tasks will be strong instead. On the other hand, Fabric Menders, Medical Transcriptionists, and Pile Driver Operators all have low variability in their task-level SML values.

Were the constituent tasks of these jobs to be high SML, these would be possible candidates for occupations where we might expect automation substitution effects to dominate. But it seems that few occupations, if any, can be entirely automated away by ML technologies. Figure 2B shows a histogram of the proportion of tasks within an occupation for which SML exceeds the 90th and 50th percentiles. The distribution of occupations that have over half of their tasks above the 50th percentile is fairly flat, but none have over 80% of their constituent tasks exceeding the 90th percentile of task SML.

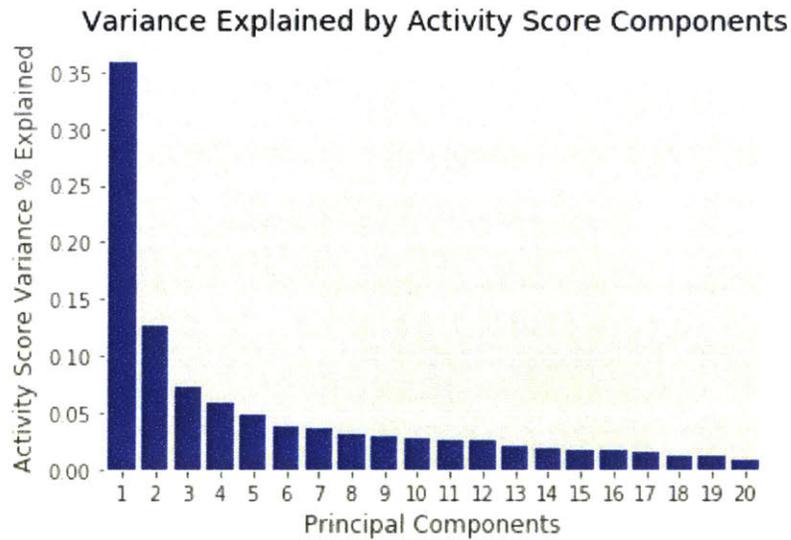


Figure 2A – Principal Component Variance Explained for Activity SML Rubric Scores

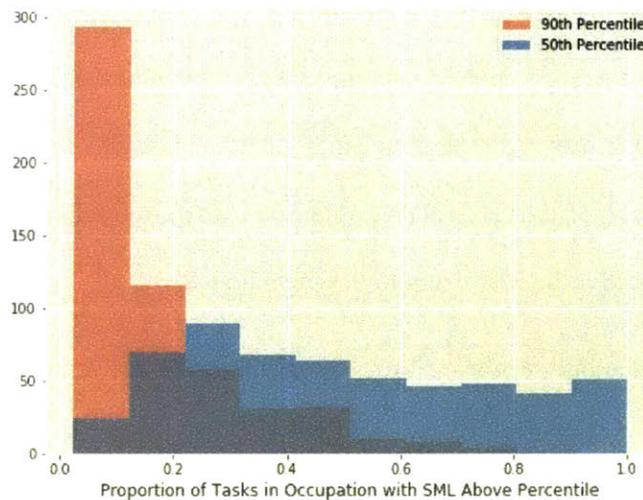


Figure 2B – Histogram of Occupations by Proportion of Tasks with SML Larger than 90th and 50th Percentile Thresholds

Strikingly, our analysis suggests that ML is unlikely to have the same effects on wages and inequality as earlier waves of automation. This can be seen by examining the correlation between SML and wages. Figures 3 and 4 show the scatterplots of wage percentile and 2016 total wage bill

percentile with respect to SML and sdSML scores. SML and sdSML are both negatively correlated with wage percentile, but not with the overall wage bill. This suggests that ML can potentially affect occupations at any skill level. If firms predominantly use wage signals, and not wage bill signals, when determining which tasks to automate first (for the substitution case), it will be the low wage occupations that have tasks automated first. On the contrary, if overall wage bills guide entrepreneurs to automation opportunities, it is occupations like retail salesperson where ML-based automation will take an early foothold.

Of course, these weak correlations we find are the case *before* the technology has been extensively implemented. To the extent technological change is directed, or more precisely, to the extent the adoption and implementation of technology is directed, there will be both wage and market size effects driving incentives to automate (Acemoglu 2002; Acemoglu and Restrepo 2017b). In that case, the occupations near the 100th percentile for total wage bill with high SML scores might be some of the most attractive to automate with ML applications. These jobs, like cashiers and accounting clerks, might later drive job polarization trends similar to those in Autor and Dorn (2013). Such trends would only be observable after the ML technologies are adopted.

This also means that our *ex ante* measures can be useful for testing theories of directed technological change. If technological change and automation are mostly directable by firms, high SML tasks in high wage (or wage bill) occupations could be more attractive automation candidates. In this early stage, sdSML (the standard deviation of task-level SML within a task) is negatively correlated with wages but not wage bills. But we would expect if technological change is highly directable that high sdSML jobs would be attractive to re-design into new task bundles. This could lead to greater year-to-year variation in wages for the occupations due to the variety of aforementioned factors that affect labor demand. We might expect, however, that reorganization could happen for low wage jobs first. Our new method of measuring the technological content of occupations facilitates

testing these kinds of hypotheses. Future studies could use the robotics and software rubrics to test how directable technological change is in technologies that have already diffused, while the SML scores will be more useful for that kind of study once ML investment reaches its steady state.

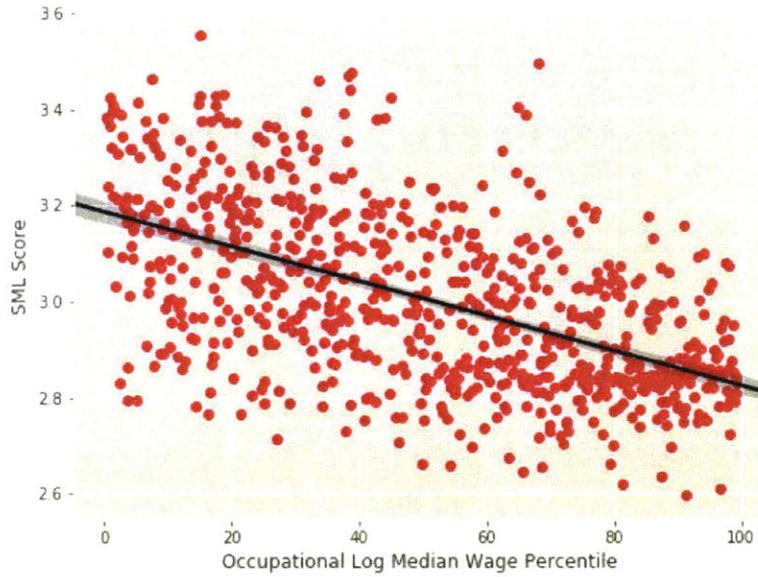


Figure 3A – SML Score vs. 2016 Median Wage Percentile;
Regression Coefficient: -0.0034 (t-stat = 18.5)

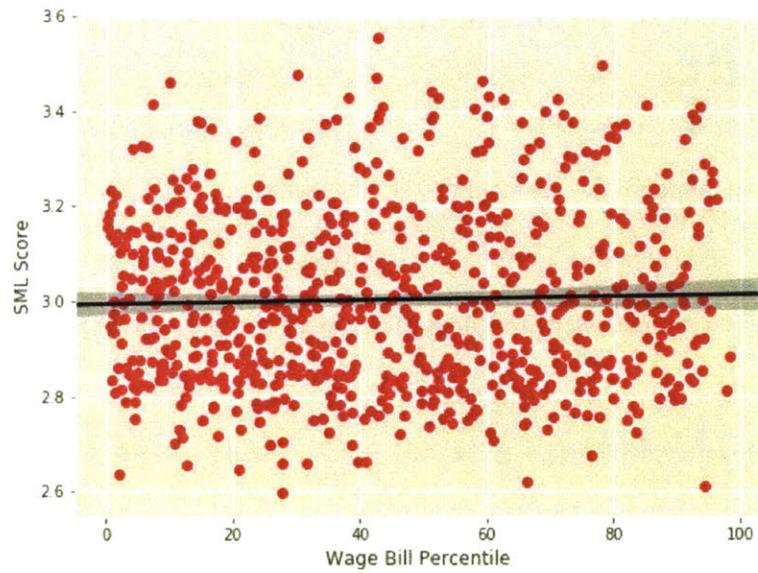


Figure 3B – SML Score vs. 2016 Wage Bill Percentile;

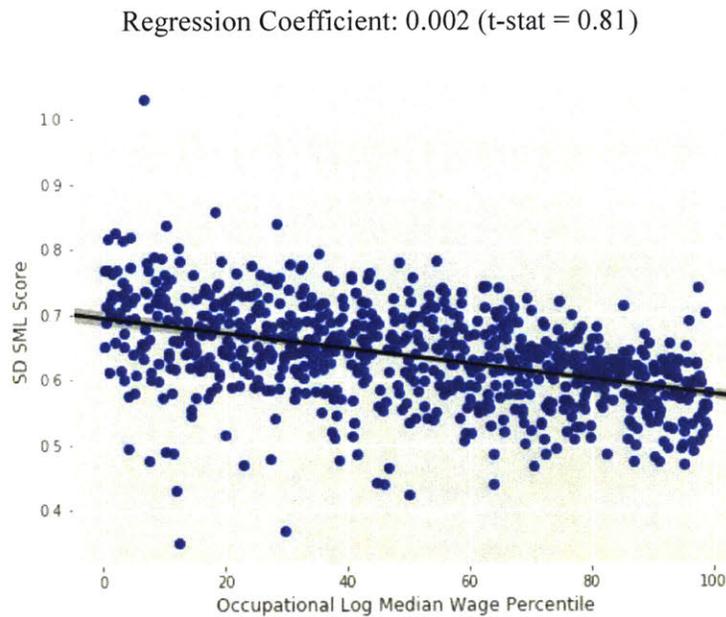


Figure 4 - sdSML Score vs. 2016 Median Wage Percentile

Regression Coefficient: -0.0012 (t-stat = 13.5)

It is further worth noting that the SML measures differ from the occupational automation scores already present in the O*NET database. Figure 5 shows the scatterplot of the O*NET automation measure against the total BLS employment for 2014. Some occupations which O*NET would suggest have low levels of automation, like truck drivers, are some of the occupations for which there are some tasks with high SML scores. Further work with the results for the suitability for rules-based software scores and suitability for robotic systems scores will investigate further how much of the O*NET automation rankings we can explain using the composite scores from our three rubrics.

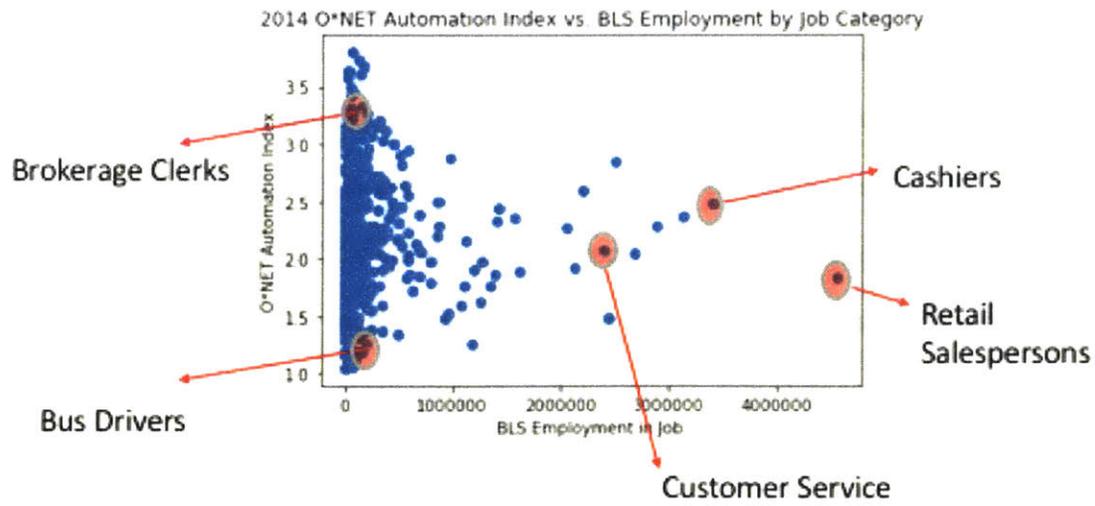
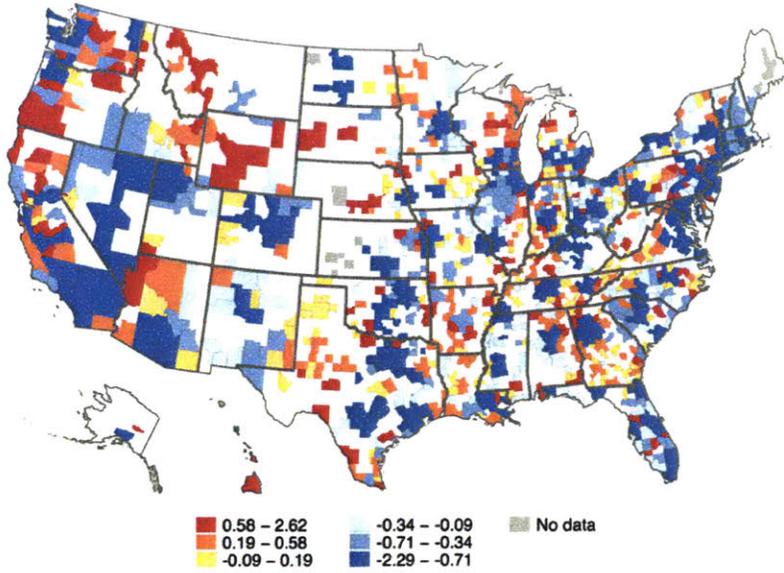


Figure 5 – O*NET Automation Scores

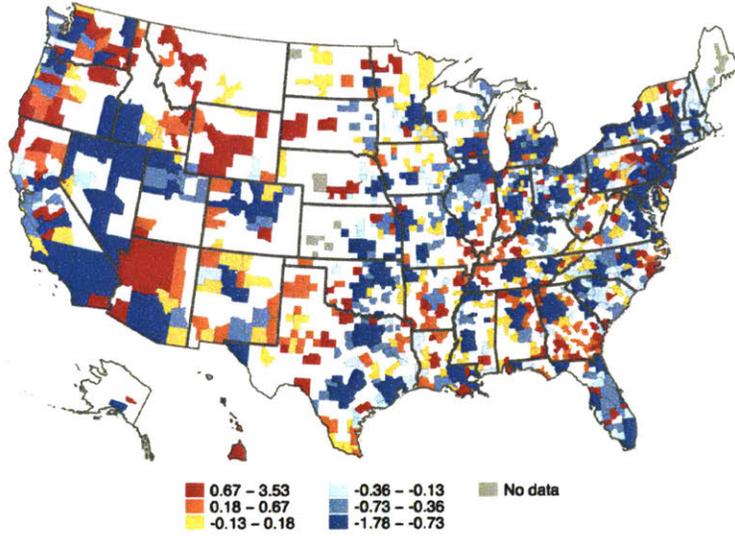
The SML scores vary by region as well, as we might expect from the economic specialization of different regions of the U.S. We join the SML scores to the regional employment data from the BLS Occupational Employment Survey and then build composite SML scores as the employment weighted average by region. Figures 6 and 7 show the standardized SML scores and standardized sdSML scores (respectively) by core-based statistical area (CBSA). There are a few missing CBSA values, which are grey in the map. Large cities tend to have lower average SML scores across all of their occupations, while rural areas are more heterogeneous. CBSAs in the South and parts of away from the Coasts tend to have higher SML scores on average with lower sdSML scores, indicating a relative preponderance of occupations with higher machine learning suitability and less potential for reorganization. In contrast, parts of Texas, Arizona, and the Pacific Northwest have high SML scores on average with high sdSML scores. These trends differ somewhat from the overall automation risks discussed in (Frank et al. 2017).

We can also look into the separate components of the SML scores and see how they are distributed by region. This will be important if various some types of ML flourish sooner than others. Figures 9 and 10 show the regional distribution of the standardized text data and image data scores (respectively) that serve as components in our overall data measure described above. Here we see a starkly different picture to the SML and sdSML regional distributions. The sources of tasks with high intensity in text and image/video data are concentrated in urban centers. The two sets of regional plots are almost mirror images of each other, suggesting that the gains to ML implementation may benefit some regions while affecting labor demand in others. One prediction of our model is that if one type of ML data, e.g. image recognition, starts affecting occupations and regions that score high on our rubric, it will soon start affecting other occupations and regions that also score high, even if they distant geography or in very different industries. Future work will investigate the regional dynamics of occupational SML in greater detail and will incorporate the Suitability for Rules-Based Systems and Suitability for Robotics rubric results.

SML by Region (2013 CBSA)



SDSML by Region (2013 CBSA)



Figures 7 and 8 – SML and sdSML scores by Core-Based Statistical Area (CBSA)

Text Data by Region (2013 CBSA)

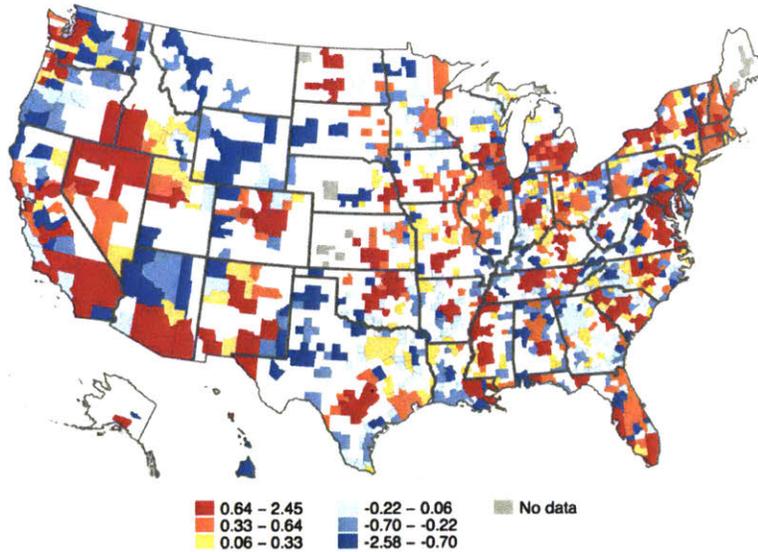
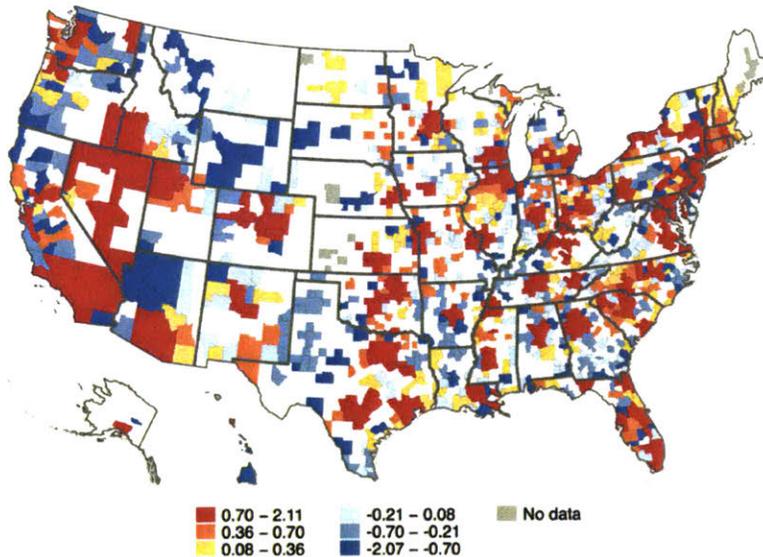


Image Data by Region (2013 CBSA)



Figures 9 and 10 – Text and Image/Video Data Intensity by Core-Based Statistical Area (CBSA)

The incentives to invest in ML, while pervasive throughout the economy, will vary by industry. As in ALM, the factor shares and elasticities of substitution of tasks will influence capital investment and employment decisions. This is the case both within occupations and between occupations in the same firm-level production function. Looking at occupation-level SML averages by

industry, we see uneven potential impact of ML. Figure 11 shows 2016 employment-weighted industry SML scores for 2-digit NAICS Codes. Here we see some early evidence for the kind of potential job polarization that machine learning investments might generate. Service industries with many workers that do complex client or human-facing tasks or critical thinking, and industries with abundant physical labor tend to have lower SML scores on average. Routine services where perception and clerical tasks are common, like retail sales operations, have higher SML scores. At this point, we cannot how labor demand will change by industry. It is too soon, for instance, to tell whether the complementary applications of machine learning will dominate the potential for labor-saving automation in finance. Still, the SML gradient by industry now suggests that the transformative effects of ML will in the long-run will not be evenly spread. The differential returns to ML investment may have important productivity implications in the long-run. Many of the low SML sectors are services industries, suggesting that ML may increase the rate at which less productive industries grow as a proportion of total factor inputs. This is further explored in Baumol (1967) and Aghion, Jones, and Jones (2017).

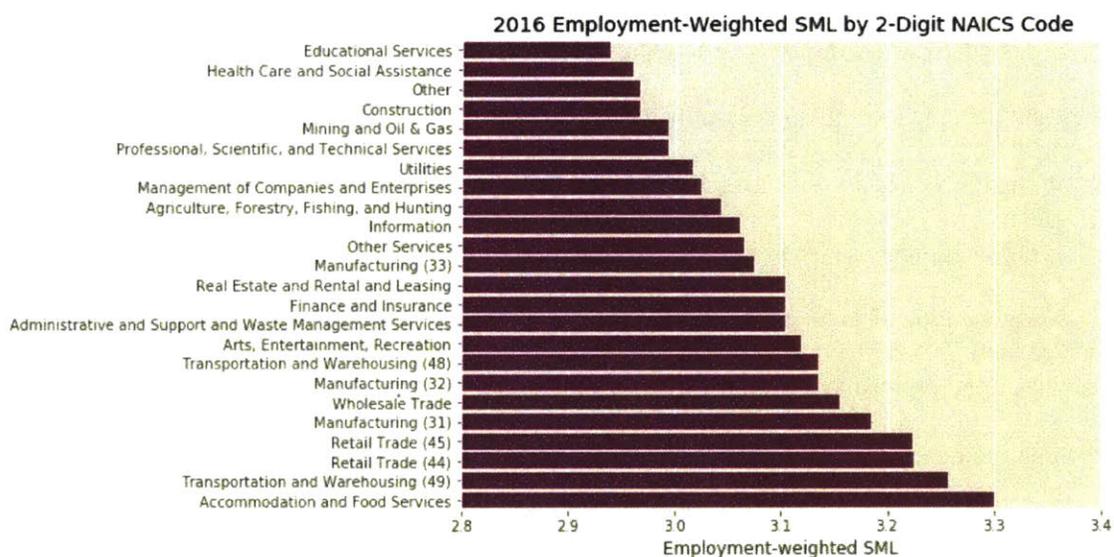
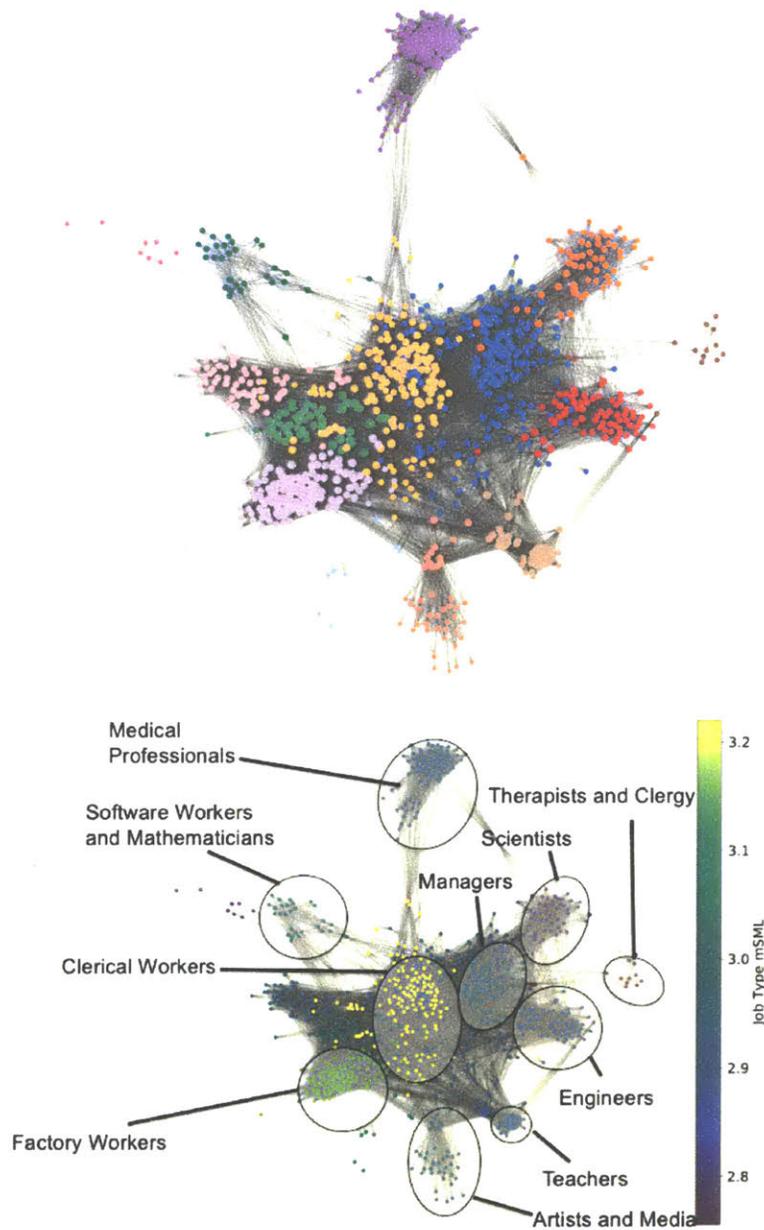


Figure 11 – Employment-Weighted Average SML by 2-Digit NAICS Industry

Augmenting the regional and industrial assessments, the fact that detailed work activities in the O*NET taxonomy are shared between occupations means we can use network analysis to understand the technological content of occupations as well. Each occupation in O*NET is connected to the others via the activities they share. This network is highly modular, meaning that community detection algorithms can easily sort different occupation nodes into clusters based on the activity edges (De Meo et al. 2011). Figure 12 shows the network of occupations grouped by the communities that the Louvain clustering algorithm detected. While these groupings have no official “category”, the light orange cluster, for example, contains many clerical jobs and the dark blue cluster near the top of the figure contains many jobs in the medical field. Applying the SML scores to the clustering, we see some clusters (like the light orange clerical one) have high SML on average. Others, like the dark blue cluster, have relatively lower SML. This SML heatmap on the same Job-Job network as in Figure 12 is represented in Figure 13. In Figure 13, high SML scores are in yellow and low SML scores are in dark purple.

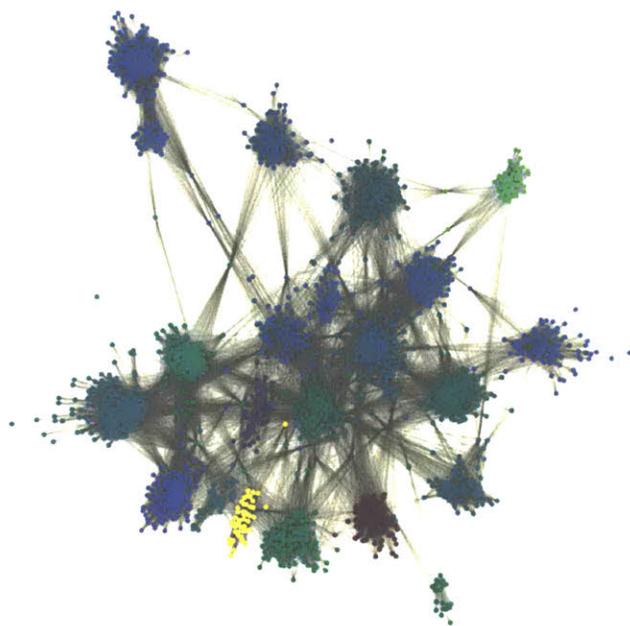
It is perhaps of greater interest to represent this network differently, where nodes are the work activities and the edges are the occupations instead. This is useful to see if the SML characteristics cluster too; that provides a useful lens with which to consider the impact of ML as well as a type of validation of the approach. If the high SML activities are often grouped into the same communities, then we have one type of indication that the dataset is internally consistent. That is, occupations with similar types of activities have similar exposure to machine learning (again, without differentiating between complements and substitutes for human labor). Figure 14 shows the network of activities.

This network is also highly modular. Many of the highest SML activities are in the brown cluster (lower left), which contains clerical tasks. Ongoing work will further investigate occupational connections in activities and how they might relate to the diffusion of technology.



Figures 12 and 13 – The Job-Job Network Colored by Community (Top) and by SML Score (Bottom)
 High (Low) Average SML is pictured in Yellow (Dark Purple)⁸¹

⁸¹ Visualizations are the work of Morgan Frank of the MIT Media Lab’s Scalable Cooperation Group. We are grateful for his assistance on this.



Figures 14 and 15 – The Activity-Activity Network (Connected by Occupations) Colored by Activity Community and SML (respectively)⁸²

6. Measurability

It may be the case that the more measurable an occupation's inputs and outputs are, the more easily a digital system might be able to do some of its tasks. For this reason, we take a subset of the SML questions and separate them out as measures of "measurability". These questions are Q1, Q17, and the Data score (the max of the four data question scores that we have). Averaging these three scores leaves us with the measurability score, another way of considering whether or not a task will be high or low SML. Measurability tends to be positively correlated with wages, as can be seen in Figure 16 below. The measurability range is shown in Figure 17. High wage occupations tend to have lower overall SML, lower sdSML, and higher measurability. In some cases, high wage occupations might be subject to greater disruption from ML technology.

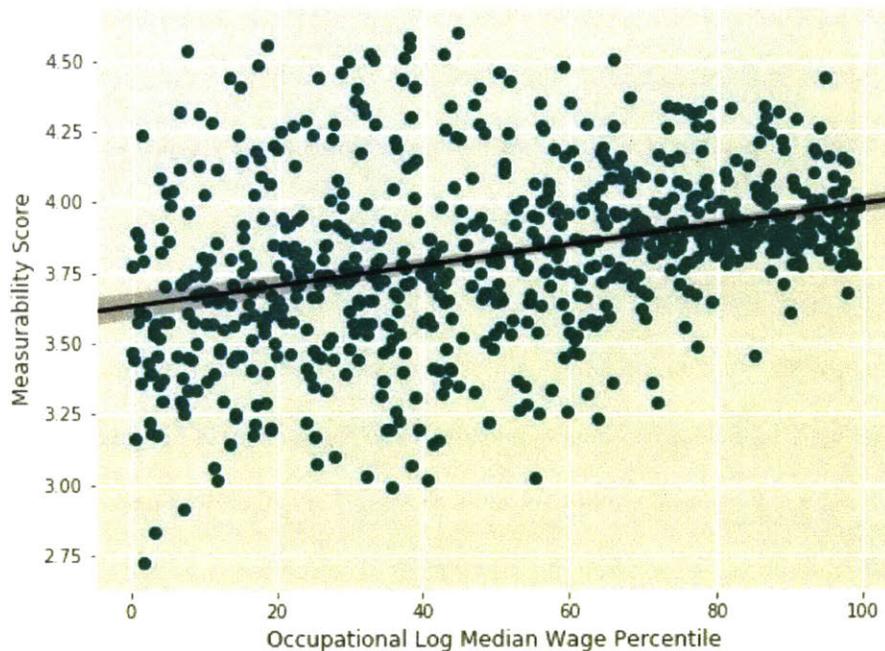


Figure 16 – Scatterplot of Measurability vs. Occupational Wage Percentile;

Regression Coefficient: 0.03 (t-stat = 9.87)

⁸² Visualizations again courtesy of Morgan Frank (MIT Media Lab – Scalable Cooperation)

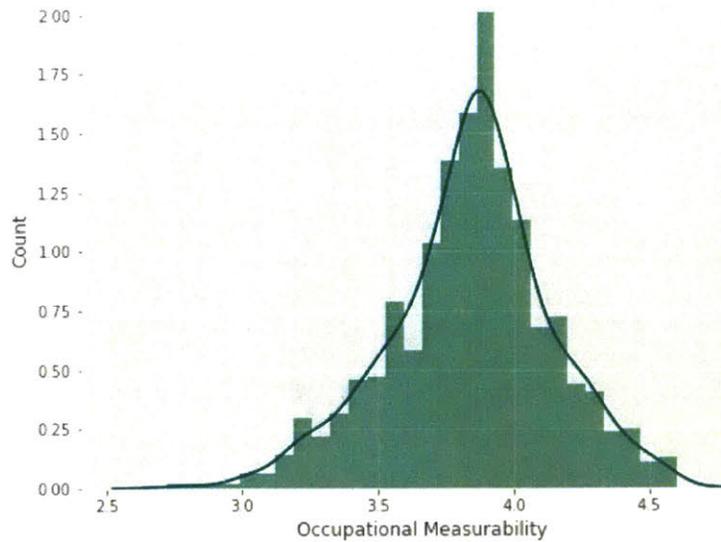


Figure 17 – The Distribution of Occupational Measurability Scores

7. Conclusion

The principal takeaways from all of our empirical analyses are that 1) few occupations can be completely automated using machine learning technology, 2) nearly all occupations have at least one task which is suitable for machine learning, and 3) these circumstances suggest there will be economic gains to re-designing jobs into new bundles of tasks.

There are a number of important caveats to our rubric-based analysis. Most importantly, we leave the most important work of inventing new business processes, re-organizing work, and thinking of new occupational tasks to the entrepreneurs, managers, and researchers of the future. This work focuses narrowly on what is already going on in the economy; it is therefore necessarily constrained to silence on change to occur in longer than just a few years. The rubric(s) can and should be updated if the technology changes (for example, if quantum computing becomes viable).

The emphasis on *technological* possibility also belies the myriad *social* facilitators and obstacles for adopting and implementing a GPT. In the case of ML, there are many social, legal, and ethical concerns which must be resolved (Bolukbasi et al. 2016; Bonnefon, Shariff, and Rahwan 2016; Remus

and Levy 2017; Kleinberg, Mullainathan, and Raghavan 2016). Our rubric does not grapple with algorithmic bias, legal challenges, or other policy concerns in ML. So far, the analysis has concentrated on ML as well. Nevertheless, the rubric structure, as a new technologically-focused tool, provides one possible roadmap for social scientists seeking to predict where future GPTs will have the greatest effects.

References

- Abbeel, Pieter, and Andrew Y. Ng. 2004. "Apprenticeship Learning via Inverse Reinforcement Learning." In *Twenty-First International Conference on Machine Learning - ICML '04*, 1.
- Acemoglu, Daron. 2002. "Directed Technical Change." *Review of Economic Studies* 69 (4):781–809.
- Acemoglu, Daron, and David Autor. 2011. *Skills, Tasks and Technologies: Implications for Employment and Earnings. Handbook of Labor Economics*. Vol. 4.
- Acemoglu, Daron, and Pascual Restrepo. 2017a. "Robots and Jobs: Evidence from US Labor Markets." *MIT Working Paper*.
- . 2017b. "The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment." *SSRN Electronic Journal*.
- Aghion, Philippe, Benjamin F Jones, and C Jones. 2017. "Artificial Intelligence and Economic Growth." *National Bureau of Economic Research*, no. w23928.
- Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb. 2017. "Prediction, Judgment, and Uncertainty." In *Economics of Artificial Intelligence*. University of Chicago Press.
- . 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston: Harvard Business Review Press.
- Agrawal, Ajay, John McHale, and Alex Oettl. 2018. "Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth." In *Economics of Artificial Intelligence*. University of Chicago Press. <http://www.nber.org/papers/w24541>.
- Autor, David H. 2014. "Polanyi's Paradox and the Shape of Employment Growth." *Working Paper*, 129–78.
- . 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29 (3):3–30.
- Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5):1553–97.
- Autor, David, Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4):1279–1333.
- Baumol, William J. 1967. "Macroeconomics of Unbalanced Growth: The Anatomy of an Urban Crisis." *American Economic Review* 57 (3):415–26.
- Benzell, Seth G., and Erik Brynjolfsson. 2018. "Digital Abundance and Scarce Genius: Implications

- for Interest Rates and Wages.” *Unpublished Working Paper*. MIT.
- Benzell, Seth G, Laurence J Kotlikoff, Guillermo Lagarda, and Jeffrey D Sachs. 2016. “Robots Are Us: Some Economics of Human Replacement.” *NBER Working Paper Series*, no. Working Paper 20941.
- Bessen, James. 2018. “AI and Jobs: The Role of Demand.” In *Economics of Artificial Intelligence*. University of Chicago Press.
- Bessen, James E. 2016. “How Computer Automation Affects Occupations: Technology, Jobs, and Skills.” *SSRN Electronic Journal* 16 (2):293–321.
- Bolukbasi, Tolga, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. “Man Is to Computer Programmer as Woman Is to Homemaker? Debiasing Word Embeddings,” 1–25.
- Bonnefon, Jean-François, Azim Shariff, and Iyad Rahwan. 2016. “The Social Dilemma of Autonomous Vehicles.” *Science* 352 (6293):1573–76.
- Bresnahan, Timothy. 2010. “General Purpose Technologies.” *Handbook of the Economics of Innovation* 2 (1):761–91.
- Bresnahan, Timothy F., and M. Trajtenberg. 1995. “General Purpose Technologies ‘Engines of Growth’?” *Journal of Econometrics* 65 (1):83–108.
- Bresnahan, Timothy F, Shane Greenstein, David Brownstone, and Ken Flamm. 1996. “Technical Progress and Co-Invention in Computing and in the Uses of Computers.” *Brookings Papers on Economic Activity: Microeconomics*, 1–83.
- Brynjolfsson, Erik, and Tom Mitchell. 2017. “What Can Machine Learning Do? Workforce Implications.” *Science* 358 (6370):1530–34.
- Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. “What Can Machines Learn, and What Does It Mean for Occupations and the Economy?” *AEA Papers and Proceedings*, 43–47.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2018a. “Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics.” In *Economics of Artificial Intelligence*. University of Chicago Press.
- . 2018b. “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies.”
- Bughin, Jacques, James Manyika, Jonathan Woetzel, Frank Michael Mattern, Susan Chui, Anu Lund, Sree Madgavkar, et al. 2017. “A Future That Works: Automation, Employment, and Productivity.” *McKinsey Global Institute*, no. January:1–28.
- Cockburn, Iain M, Rebecca Henderson, and Scott Stern. 2018. “The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis.” *National Bureau of Economic Research* No. w24449 (September).
- Esteva, Andre, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. 2017. “Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks.” *Nature* 542 (7639):115–18.
- Felten, Edward W., Manav Raj, and Robert Seamans. 2018. “A Method to Link Advances in Artificial Intelligence to Occupational Abilities.” *AEA Papers and Proceedings* 108:54–57.

- Finn, Chelsea, Paul Christiano, Pieter Abbeel, and Sergey Levine. 2016. "A Connection between Generative Adversarial Networks, Inverse Reinforcement Learning, and Energy-Based Models."
- Frank, Morgan R., Lijun Sun, Manuel Cebrian, Hyejin Youn, and Iyad Rahwan. 2017. "Small Cities Face Greater Impact from Automation."
- Furman, Jason, and Robert Seamans. 2018. "AI and the Economy," 1–33. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3186591.
- Goodfellow, Ian, Bengio, Yoshua, Courville, Aaron. 2016. "Deep Learning." *MIT Press*.
- Goolsbee, Austan. 2018. "Public Policy in an AI Economy." In *Economics of Artificial Intelligence*. University of Chicago Press.
- Graves, A, A.-R. Mohamed, and G Hinton. 2013. "Speech Recognition with Deep Recurrent Neural Networks." *2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, no. 6:6645–49.
- Greenstein, Shane, and Jeff Prince. 2005. *The Diffusion of the Internet and the Geography of the Digital Divide in the United States. Applied Economics*.
- Guyot, Alexandre, Laurence Hubert-Moy, and Thierry Lorho. 2018. "Detecting Neolithic Burial Mounds from LiDAR-Derived Elevation Data Using a Multi-Scale Approach and Machine Learning Techniques." *Remote Sensing* 10 (2).
- Heckman, James J, Lance J Lochner, and Petra E Todd. 2006. "Earnings Functions, Rates of Return and Treatment Effects: The Mincer Equation and Beyond." *Handbook of the Economics of Education* 1. Elsevier:307–458.
- Hinton, Geoff, Yoshua Bengio, and Yann Lecun. 2015. "Deep Learning Tutorial." *Nips*.
- Holmstrom, Bengt, and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses : Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, & Organization* 7:24–52.
- Jorgenson, Dale W. 1963. "Capital Theory and Investment Behavior." *American Economic Review* 53 (2):247–59.
- Kaldor, Nicholas. 1961. "Capital Accumulation and Economic Growth." In *The Theory of Capital: Proceedings of a Conference Held by the International Economic Association*, 177–222.
- Kleinberg, Jon, Sendhil Mullainathan, and Manish Raghavan. 2016. "Inherent Trade-Offs in the Fair Determination of Risk Scores," 1–23.
- LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. "Gradient-Based Learning Applied to Document Recognition." *Proceedings of the IEEE* 86 (11):2278–2323.
- Levy, Frank. 2017. "Computers and Populism: Artificial Intelligence, Jobs, and Politics in the Near Term." *Oxford Review of Economic Policy*, Forthcoming.
- Litjens, Geert, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen A.W.M. van der Laak, Bram van Ginneken, and Clara I. Sánchez. 2017. "A Survey on Deep Learning in Medical Image Analysis." *Medical Image Analysis*.
- Loos, Sarah, Geoffrey Irving, Christian Szegedy, and Cezary Kaliszyk. 2017. "Deep Network Guided

- Proof Search,” 1–20. <http://arxiv.org/abs/1701.06972>.
- McAfee, Andrew, and Erik Brynjolfsson. 2017. *Machine, Platform, Crowd: Harnessing Our Digital Future*. *Harnessing Our Digital Future*.
- Meo, Pasquale De, Emilio Ferrara, Giacomo Fiumara, and Alessandro Provetti. 2011. “Generalized Louvain Method for Community Detection in Large Networks.” In *International Conference on Intelligent Systems Design and Applications, ISDA*, 88–93.
- Milgrom, Paul, and John Roberts. 1996. “The LeChatelier Principle” 86 (1):173–79.
- Mincer, Jacob. 1974. “Schooling, Experience, and Earnings. Human Behavior & Social Institutions No. 2.” ERIC.
- Pickering, Andrew. 1993. “The Mangle of Practice: Agency and Emergence in the Sociology of Science.” *American Journal of Sociology* 99 (3):559–89.
- Polanyi, Michael. 1966. “The Logic of Tacit Inference.” *Philosophy* 41 (155):1–18.
- Remus, D, and F Levy. 2017. “Can Robots Be Lawyers? Coomputers, Lawyers, and the Practice of Law.” *Geo. J. Legal Ethics* 30:501.
- Rock, Daniel. 2019. “Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence.” *Unpublished Working Paper*.
- Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. 1986. “Learning Representations by Back-Propagating Errors.” *Nature* 323 (6088):533–36.
- Shokri, Reza, and Vitaly Shmatikov. 2016. “Privacy-Preserving Deep Learning.” In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing, Allerton 2015*, 909–10.
- Silver, David, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, et al. 2017. “Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm,” 1–19.
- Trajtenberg, Manuel. 2018. “AI as the Next GPT: A Political-Economy Perspective.” 24245. Working Paper Series.
- White, B. W., and Frank Rosenblatt. 1963. “Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms.” *The American Journal of Psychology* 76 (4):705.

Appendix A – Machine Learning Paradigms

Artificial Intelligence as a field includes approaches to problem solving like symbolic manipulation, knowledge-based systems, and cybernetics in addition to ML. Within ML, there are three main paradigms which are useful for solving different types of problems:

Supervised Learning — This is the area of machine learning that currently has the most widespread commercial use. Algorithms are designed to learn a mapping function between input and output data where both the inputs and the outputs are known (and paired together). Classification and regression are common objectives. Standard deep and convolutional neural nets fall under this category.

Reinforcement Learning — A class of machine learning algorithms which seek to learn an optimal policy function given a set of possible actions and representation of the environment. In training, actions are chosen, the environment changes, and a reward is assigned. The goal of the algorithm is to learn an action policy conditioned on the environmental state that maximizes the (discounted) sum of rewards over time. Reinforcement Learning techniques can often be combined with deep learning methods.

Unsupervised Learning — A class of machine learning algorithms where structure in a set of inputs is computed without having an explicit output target to predict or classify. Examples of applications include clustering and dimension reduction.

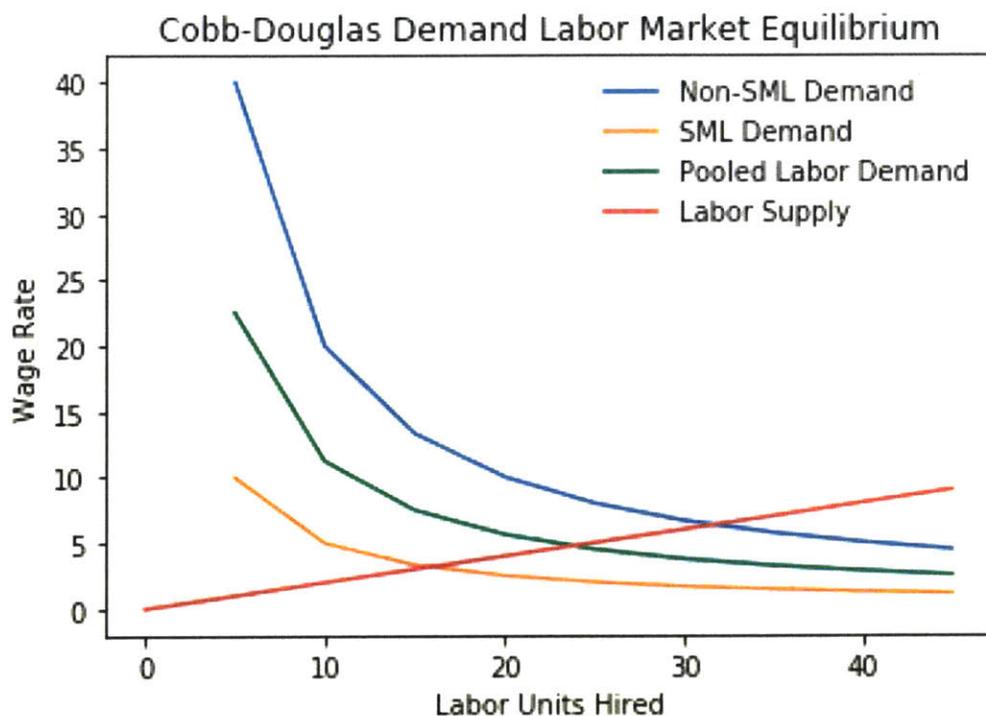


Figure A1 – Cobb-Douglas Labor Market Equilibrium with Fixed Ratio Constraint (Green) and without Constraints (Blue and Orange). Note: depending on the fixed ratio constraint, wages can be higher or lower than either wage in the separate demand case. If capital is weakly more productive than all SML labor in the separate demand case, the orange line is flat at zero as machines replace humans.

Appendix B – The Suitability for Machine Learning Rubric

1. Task information is recorded or recordable by computer

1: It is very difficult or impossible to save particular inputs and outputs in a computerized form (e.g. completely recording ideas and strategies, evaluating relationships with coworkers)

3: It is possible to partially represent inputs and outputs in a digital format (e.g. ranking a series of possible sales leads, recording security footage to evaluate possible dangers)

5: It is easy to store inputs and results on a machine/computer (e.g. calculation of an account balance, generating a record of a transaction, taking a picture or video recording)

2. Task feedback is immediate

1: Feedback is never received or takes a very long time (e.g. making art of different kinds, working for world peace)

3: Feedback is received but response time is inconsistent/unclear and/or unclear on if it is beneficial to progress (e.g. measuring teacher performance using standardized tests)

5: Feedback and results are instantly received when task is completed (e.g. sending a text message, scheduling a meeting, making a reservation)

3. It is okay to make mistakes when completing this task

1: A mistake could lead to serious harm, injury, or death to those involved, or could lead to lasting negative consequences (e.g. mistake during surgery, mistake at a nuclear facility)

3: A mistake will have negative consequences, but can be fixed with some work (e.g. making a clerical error that can be corrected, accidentally writing a bug in software code)

5: A mistake can be easily fixed, and holds few, if any negative consequences (e.g. delivering packages, scheduling a meeting, walking a dog)

4. It is not important that the task is done by a human

1: Task fundamentally requires human connection (e.g. providing psychological therapy services, making a speech)

3: Task could be done by a non-human, but might cause frustration or inefficiency (e.g. customer service)

5: Task requires little human connection, empathy, or emotional intelligence (e.g. preparing taxes, performing calculations, lifting boxes)

5. Task does not require complex reasoning

1: Task requires intuition or highly involved reasoning (e.g. determining hiring needs, coming up with a research proposal/plan, teaching)

3: Task requires some reasoning, but can mostly be broken into well-defined rules (e.g. playing chess, sorting mail)

5: Task is mainly perception and does not require complex reasoning skills (e.g. catching a ball, recognizing an animal)

6. Task matches labels to concepts, predictions, or actions

1: The task does not have clear, consistent categories or labels (e.g. telling a story)

3: The task potentially has well-defined categories or labels, but does not require mapping of the two (e.g. assigning work to direct reports, deciding which products to sell)

5: The task has clear, consistent categories or labels (e.g. translating one language to another, matching an image to words describing the image)

7. Task involves a brief, simple, highly-structured conversation with a customer or someone else

1: Task doesn't require any form of communication/conversation with another person (e.g. solving equations, lifting objects, observing)

3: Task involves conversation with others, but it might not be simple or highly-structured (e.g. providing fashion advice, having a meeting in an office)

5: The task involves simple conversations with people with similar structure (e.g. taking orders and reservations at a restaurant, providing directions to locations of interest)

8. Task is repeated frequently

1: The task is not performed often (e.g. fighting a fire, treating rare and specific issues)

3: Task is performed often, but might be done differently each time (e.g. waiting tables, operating a multi-purpose machine, teaching a class)

5: Task is very repetitive, and is done in the same way each time (e.g. working in an assembly line, delivering things along a route, being a cashier)

9. There is no need to explain decisions when doing the task

1: Decisions highly impact the lives of others and require justification (e.g. persuasion, long term planning, law-making, courtroom decisions)

3: There is some need to explain decisions, particularly when people ask questions (e.g. doctors performing checkups, operating machinery)

5: There is no need to explain decisions. The task is only concerned with having the correct output, and does not depend on the process through which the output is determined. (e.g. correctly predicting the weather, performing the correct calculation, optimizing allocation of resources, determining quickest route)

10. Task is about choosing between multiple predetermined options.

1: Task output does not have to do with choosing one of a few options (e.g. lifting objects, collecting things, making things)

3: Task output typically isn't presented as choosing between one of a handful of preset options, but might be converted into that format (e.g. recommending a plan, choosing a supply company, setting a price on an item for sale)

5: Task is focused on picking between one of multiple options (e.g. grading food quality, diagnosing common conditions, sorting mail)

11. Long term planning is not required to successfully complete the task

1: The task is concerned with planning around a timeline of months or years (e.g. supervising research projects, constructing complex buildings, entrepreneurship, crafting long term cancer treatment plans)

3: The task is concerned with a timeline in the range of weeks or days or an indeterminate amount of time (e.g. managing others' workloads, teaching students a specific set of lessons)

5: The task involves an immediate response and isn't concerned with a future impact (e.g. answer questions in a call center, lifting objects, performing calculations)

12. The task requires working with text data

1: Task does not include working with any text (e.g. making a hamburger, operating machinery)

3: Task may include some light reading and writing (e.g. reading labels, occasionally reading directions)

5: Task includes heavy text processing, reading, or writing (e.g. reading documents, writing a letter)

13. The task requires working with image or video data:

1: Task does not require looking at images or videos, or otherwise using your eyes (e.g. having a phone conversation)

3: Task may occasionally require looking at images and video (e.g. greeting customers, booking a hotel room)

5: Task requires analyzing images and videos (e.g. finding defects in products, looking at surveillance footage, classifying objects in pictures, face recognition)

14. The task requires working with speech data:

1: Task does not require listening to or communicating with speech (e.g. independent tasks such as lifting objects, repetitive assembly work)

3: Task may require occasional listening, talking, or communicating (e.g. construction work, being a cashier, financial analyst)

5: Task requires heavy speech processing, or communicating with speech (e.g. telemarketing, translating between languages, giving a speech, having a conversation)

15. The task requires working with other types of data (other than text, image, video, and speech):

- 1: Task does not require working with data in any form (e.g. making handmade art)
- 3: Task requires working with some types of data at a low frequency (e.g. performing restocking tasks at a grocery store, testing machines for maintenance needs)
- 5: Task requires constant interaction with digital records, sensor data, or other types of data. (e.g. monitoring temperature/weather, analyzing pricing data, pulling and reading medical records)

16. The task can be completed in one second or less

- 1: Task takes a long time to complete (e.g. making a plan, writing a book)
- 3: Task cannot be done instantly, but also does not involve much long-term planning (e.g. performing a surgery, delivering food)
- 5: Task can be done instantly, or can be broken up into smaller choices that can be done instantly (e.g. Identifying a picture)

17. Each instance, completion, or execution of the task is similar to the other instances in how it is done

- 1: Task primarily involves rare or unique situations that cannot be summarized easily with machine-readable data (e.g. making strategic decisions for a company)
- 3: Data can be collected but the data output structure is highly varied (e.g. performing different types of surgery will generate different kinds of feedback)
- 5: Data are already available or can be easily collected (e.g. customer service transcripts, text translation, image classification, stock price movements)

18. Practicing the task to get better is easy

- 1: Task involves many rare-occurring or unique situations that make the task hard to practice (e.g. disaster relief, police detective casework)
- 3: Some parts of the task are possible to practice or learn by repeating (e.g. shipping/receiving in a warehouse, architectural design)
- 5: Sequences can be repeated and tested over and over again, and there are “right” moves that can be used to generate rewards (e.g. video games, learning a language)

19. The task is primarily about predicting something

- 1: The task has little to do with prediction (e.g. writing a novel, washing dishes, installing a solar panel, painting)
- 3: The task has some components which require predicting something or classifying (e.g. driving a vehicle requires guessing what people might do, forecasting financial results as part of business plan)
- 5: The task is entirely about prediction or classification (e.g. predicting the weather, identifying pictures with cats in them)

20. This task is part of this occupation (check 1) (this question is also included in reversed form, i.e. “this task is not part of this occupation”).

- 1: The occupation will almost certainly never do this task as part of their job (e.g. a teacher devising a business plan, a medical doctor underwriting corporate debt, a lawyer walking a dog).
- 3: It is unclear if this occupation will have to perform this task (e.g. a retail employee performing customer base analysis, an athlete preparing office documentation)
- 5: This occupation will have to do this as part of their work (e.g. an accountant has to balance ledgers, a restaurant server will have to wait tables)

21. The task output is not error tolerant (check 2)

- 1: A mistake can be easily fixed, and holds few, if any negative consequences (e.g. a slip up in factory work or mail sorting mistakes could go potentially unnoticed)
- 3: A mistake will have negative consequences, but can be fixed with some work (e.g. a construction mistake, or human resources slip-up will be noticed and reprimanded, but would not result in termination of employment or injury)
- 5: A mistake could lead to serious harm, injury, or death to those involved, or could lead to lasting negative consequences (e.g. mistake during surgery, mistake at a nuclear facility)

Appendix C – A Framework for Predicting the Impact of New Technologies on Labor Demand and Wages

Taking the perspective of a manager at a firm looking to automate tasks, the greatest benefit per task-period unit that the manager could hope to gain by automate is the difference between the wage paid in that period for the task and the cost of capital in the period for the machine replacement. In percentage terms, for any task fully automatable t with wage w we have:

$$(A1) \quad \frac{\Delta w_t}{w_t} = \frac{w_t - r}{w_t}$$

Since capital costs must be paid to finance machine labor, the most a manager could get per task is the difference between the prevailing wage for the task and the cost of capital. Now assume that there is a Suitability for Machine Learning (SML) measure α that represents partial automatability such that we have instead:

$$(A2) \quad \frac{\Delta w_t}{w_t} = \frac{\alpha_t(w_t - r)}{w_t}; \quad \alpha_t \in [0, 1] \forall t$$

The SML measure determines what proportion of the difference between the task wage and the interest rate is appropriable if an automation technology is implemented. SML is constrained to the unit interval. With knowledge of the difference between the task-level wage and the interest rate for some period, the SML measure helps attenuate the total benefit from full automation. Task-level wages, if otherwise unavailable, are possibly recoverable from hedonic regressions to estimate earnings functions, for example (Mincer 1974; Heckman, Lochner, and Todd 2006). Since the incentives to automate may not be continuous along the SML interval, we could choose α^* such that $\alpha \geq \alpha^* = 1$. In either case, what remains is to translate possible changes in wages into changes in quantities of labor demand. Ideally there would be easily accessible data on cross-wage elasticities of labor demand for all of the tasks in the economy. Under the assumption that this information is available, denote the elasticity matrix as follows:

$$(A3) \quad \epsilon_w = \begin{bmatrix} \frac{\partial Q_1}{\partial w_1} \frac{w_1}{Q_1} & \dots & \frac{\partial Q_1}{\partial w_n} \frac{w_n}{Q_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial Q_n}{\partial w_1} \frac{w_1}{Q_n} & \dots & \frac{\partial Q_n}{\partial w_n} \frac{w_n}{Q_n} \end{bmatrix}$$

Where each elasticity on the diagonal is the percentage change in quantity for the task given a change in that task's own wage, and the off-diagonal entries correspond to elasticities of task quantity with

respect to other task wages. These elasticities can of course vary depending on the prevailing wage, and we therefore assume at least locally stationary and constant elasticities in this matrix. The matrix dimension is $n \times n$ to represent all combinations of n tasks at the firm, though the rows and columns could just as easily represent an occupation instead of a firm.

With task weight vector γ ($n \times 1$) designating the proportions of labor in each one of the tasks and the automation potential vector α ($n \times 1$), the resultant equation for the total predicted percentage change in the total quantity of labor demanded (not task-level quantities) is then:

$$(A4) \quad \gamma' \epsilon_w \left(\frac{\Delta w}{w} \circ \alpha \right) \sim \frac{\Delta Q}{Q}$$

This is the task and alpha-weighted total predicted percentage change in labor quantity. Alpha weights are applied using the Hadamard product. The SML weighting scheme can either be defined over the unit interval or can have discrete jumps (or, perhaps, can directly specify which tasks or jobs can be totally automated. Further, known values for changes in quantities, changes in wages, weights, and automation weights can be used to back out possible eigenvalues of the wage elasticity of labor demand matrix. This can be used to assess the plausibility of certain automation scenarios given a prior over the elasticity matrix.

Notably this setup means that evaluating the technological change effects of different technologies is no longer contingent on knowing the entire structure of the wage elasticities of labor demand. It only requires knowledge of the range of possible eigenvalues of the wage elasticity matrix.

$$(A5) \quad \gamma' (\lambda_{\epsilon_w}) \left(\frac{\Delta w}{w} \circ \alpha \right) \sim \frac{\Delta Q}{Q}$$

λ_{ϵ_w} is an eigenvalue of the wage elasticity matrix. What this setup requires is 1) data on interest rates and wages for the tasks or occupations in question, 2) weights of each task or occupation in the entity of student (job, organization, economy, region, etc.), 3) automation scores mapped to the unit interval,

and 4) either predicted quantity changes (to assess automation score realism) or the range of eigenvalues (to evaluate scenarios). Future research will apply this setup towards those ends. Of course there must necessarily be at least one real eigenvalue, guaranteed in this case by the real-valued output in quantity changes.⁸³ We would like to be able to estimate difference in the logged quantities of labor demand as the integral over the wage change of the following where task weights and elasticities are a function of wage vector \mathbf{w} :

$$(A6) \quad d\log Q = (\boldsymbol{\alpha} \circ \boldsymbol{\gamma}(\mathbf{w}))' \boldsymbol{\epsilon}(\mathbf{w}) d \log \mathbf{w}$$

Assuming that elasticities and weights are a constant function of wages permits the approximation in equation (A5).

⁸³ Another assumption is that all entries of the elasticity matrix are real. All quantity changes are intuitively defined over the real space as well, though we fall short of proving that the dimensionality of elasticities is odd.

Chapter 4 — Do Labor Demand Shifts Occur Within Firms or Across Them? Non-Routine Biased Technological Change, 2000-2016⁸⁴

1. Introduction

A large literature has measured wide-scale changes in the composition of US employment. These changes have been biased against occupations that are intensive in routine tasks. Partly for this reason, many have attributed these skill-biased employment shifts to technological change. This paper distinguishes between changes in occupational due to employment changes across and within firms, and measures how skill-biased technological change (SBTC) is related to those changes and firm-level investments in R&D.

An occupation, e.g. business analysts or truck drivers, can increase in relative employment share for several reasons. One way this can happen is if firms that have a disproportionate amount of such workers grow faster than other firms. We call this the *across* firm heterogeneous growth effect. Another possibility is that some workers, like data entry clerks, are being fired and replaced with data analysts. We call this the *within* firm rebalancing effect. A final pair of possibilities is that new firms are created that disproportionately employ data analysts or old firms go out of business that disproportionately employed other kinds of workers.

An occupation, e.g. business analysts or truck drivers, can increase in relative employment share for several reasons. One way this can happen is if firms that have a disproportionate amount of such workers grow faster than other firms. We call this the *across* firm heterogeneous growth effect. Another possibility is that some workers, like data entry clerks, are being fired and replaced with data

⁸⁴ With Seth Benzell (MIT IDE) and Guillermo Lagarda (Boston University and Inter-American Development Bank). We were provided access to de-identified LinkedIn profile information as participants in LinkedIn's Economic Graph Research program. We thank Brian Lucking, Nicholas Bloom, and John Van Reenen for assisting us with their data. We would like to thank David Autor, Jacqueline Barrett, Erik Brynjolfsson, Pascual Restrepo, Sebastian Steffen, the LinkedIn Economic Graph Research team, and the seminar participants at the MIT Sloan Organizational Economics Lunch workshop for their very helpful comments.

analysts. We call this the *within* firm rebalancing effect. A final pair of possibilities is that new firms are created that disproportionately employ data analysts or old firms go out of business that disproportionately employed other kinds of workers.

In this paper we measure the relative importance of these components in job creation and destruction. To generate these new granular measures we rely on several data sets. The most important of these is LinkedIn profile records. These records, when combined with information from the Bureau of Labor Statistics, allow us to estimate occupational employment at the firm-year level. Our dataset covers publicly traded US businesses from 2000 through 2016. Over this interval, US public firms added 10.6 million jobs. The firms in our data had 27.0 million employees in 2000, or about 18 percent of the US labor force (BLS 2018).

Previous research has found faster employment growth for the highest paid and least routine-intensive occupations (Acemoglu and Autor 2011) Therefore in our analyses we group occupations by routine-task intensity and initial wage (in 2000). Consistent with this research, we find that the net new employment created over this interval is highly biased against routine labor. 2.95 million more jobs were created in the bottom third of routineness than in the top third. The most important source of this SBTC was faster growth across firms with high levels of non-routine workers, accounting for 48.9 percent of this bias. The next most important source of this SBTC was within firm rebalancing, which accounts for 30.0 percent of the bias in public firms. New firms also had relatively low levels of routine workers, and accounting for 25.7 percent of the total change. The exit of firms worked slightly against SBTC, explaining -4.6 percent of the trend.

This decomposition is important for better understanding the causes of automation. Leading theories of automation differ in their predictions about whether skill-biased technical change occurs within or across firms. Direct automation (e.g. Benzell et al. 2016) is the simplest story, i.e. complete replacement of human labor with machines. After acquiring new technologies, such as industrial

robots, companies simply fire their automated workers. Other models of interacting industry-level productivity or globalization and aggregate demand suggest that SBTC is caused by a glut of manufactured goods (e.g. Bessen (2018)). Under this theory, routine workers are disproportionately impacted by firings simply because they are concentrated in certain industries. Other theories involving "superstar" firms dominating an industry naturally, such as Autor et al. 2017, lead to a shift toward the occupations and capital varieties most heavily employed in the superstar production function. In these models, within each industry, some new firms or particularly savvy old ones (which are attuned to the new digital economy and have the right mix of non-routine workers to take advantage of it) grow faster and hire more of the workers they always have. Matching the empirical decomposition we document here is a potential desideratum for future models of automation.

In our framework section, we show that firm TFP-driven changes in aggregate routine employment would tend to produce offsetting within-firm employment mix changes in favor of routine workers. Intuitively, this is because if demand for non-routine workers goes up their wages will go up as well all else constant. All types of firms reduce the share of non-routine workers in employment. Therefore, the fact that within-firm rebalancing is a significant cause of the increase in non-routine employment indicates that production function changes within the firm is a main cause of SBTC.

Which, or to what extent each, theory is true is of much more than academic interest. Some government policies have the explicit goal of preventing technological unemployment or softening the blow of automation on routine workers. Whether these policies will be effective will depend in part on the margins where SBTC occurs. Import restrictions on manufactured products or industrial subsidies for manufactured products will fail if routine workers are increasingly not employed by these firms. Discouraging the exit of firms (either through bailouts or bankruptcy law) or encouraging entrepreneurship will help only insofar as the relevant firms employ routine workers. The Tax Cuts

and Jobs Act of 2017, for example, promised to raise the salaries of Rust Belt workers by boosting capital and R&D investment. But if firms that make larger capital investments reduce their demand for routine workers, then this is a dubious solution. Effective policy response is a steep challenge without knowledge of the locus of SBTC or being able to test mechanisms responsible for the shift away from routine workers. This paper provides some description of the different drivers of occupational compositional shifts in the new millennium.

The location of SBTC also has consequences for the long-term future. Theories of directed technological change attempt to understand how economic conditions lead to and are created by innovations. Some theories, such as Peretto and Seater (2013) feature firms which can invest in R&D to change their production functions. In these models, the decision for a firm to develop a skill-biased technology is a function of the prices it faces. The intended consequence of these innovations is to raise demand for workers with low-paid skills. Others, such as Acemoglu (2002) assume that specialized companies create technologies that all other firms implement. In models of that type, technology creators are also sensitive to the relative amounts of workers of each type who might use their tools. The distinction is important because while innovation of the first type will tend to focus on finding new uses for cheap factors, the innovation of the latter type will focus on abundant factors. Under the latter model, higher paid workers may disproportionately benefit from innovation so long as they are sufficiently abundant. We find strong evidence that firms which make additional R&D investments due to tax incentives are more likely to have routine workers.

We also introduce a new measure of within-firm occupational change. Along several dimensions, business dynamism in the United States has been on the decline. From 2000 to 2011, the share of firms less than 5 years old declined about 5 percentage points. Over the same interval, the rates of job creation and destruction decreased from approximately 17 and 14 percent respectively to 13.5 and 12 percent (Decker et al. 2014) We construct a complementary measure, which we call firm

employment mix dynamism. This measure tracks how much firms change their occupational mix. We find that firm employment mix dynamism has decreased steadily since 2000. These results indicate that the decrease in the rate of job creation and destruction corresponds in part to a reduction in the rate at which firms' skill demands change, rather than reflecting only churn in employees within a type. Of course, it may also be that firms have better information about the types of labor they will need to hire. The declining within-firm labor dynamism may reflect a more stable business environment.

2. Related Literature and Framework

What is technology's role in shifting labor demand? Early studies of SBTC focused on the education wage premium. The share of workers with a high school or college education increased dramatically in the latter half of the 20th century. Yet, across nations, the wage premium for the educated stayed constant or increased over this interval (Berman et al. 1998). While other contributing factors have been proposed, such as increased globalization (Acemoglu et al. 2016) and decreased unionization (Western and Rosenfeld 2011) the consensus view is that a leading cause of this phenomena is technological change (Johnson 1997).

More recent papers have diagnosed labor demand *polarization* as the cause. From 1980 to 2005, occupations which were highly compensated in 1980 saw disproportionate growth in both wage and employment. The same is true of occupations compensated poorly. Other occupations saw little employment or wage growth. Autor and Dorn (2013) find that areas that specialized historically in industries which use routine tasks intensively (such as manufacturing) saw larger increases in wage and employment polarization. This finding remains after controlling for the offshorability of jobs. They follow Autor et al. (2003) in attributing this to technological advances, in particular in information technology, which tend to substitute for people in routine jobs. Other papers looking at the role of technological change in wage polarization in developed countries are Acemoglu (1999), Goos and Manning (2007), and Goos et al. (2010, 2014).

There is also some industry and small scale firm-level evidence supporting this hypothesis. Bresnahan et al. (2002) and Berman et al. (1994) provide firm and industry-level evidence respectively of IT and R&D having skill-biased impacts on labor demand. Barth et al. (2017) show that manufacturing establishments with a higher proportion of scientists and engineers pay higher wages.

There is also theoretical and empirical evidence of technology playing a role in the rise of the 1%'s share of income. This trend is driven by labor earnings, not increasing capital income of the wealthy. The increase in top income shares has impacted higher earners of all types Kaplan and Rauh (2013). Rosen (1981) presciently forecasted that economies of scale enabled by new technologies would increase inequality. Innovations in, for example, telecommunications lead more tasks to be winner-take-all where gains might have been more evenly distributed in the past.

Theories of skill-biased employment demand change, whether due to offshoring, technology, changing preferences, or some other mechanism, can be broadly divided into a handful of types. They can be caused by changes in the industrial composition of employment, by the heterogeneous growth of firms with unusual employment mixes, by firms changing their mix of employees, or through the entry or exit of firms with unusual employment mixes.

3. Framework

To understand how different sources of skill-biased technological change can be distinguished, consider the following model of employment.

The economy makes two commodities: manufactured products (M) and services (S). Each is produced in a perfectly competitive industry. The industries contain I_M and I_S firms respectively. Each firm's production is constant elasticity of substitution (CES) in routine and non-routine labor. Within industries, each firm has an identical production function. We have

$$M = \sum_{i=1}^{I_M} A_M^i (z_{R,M}^{\frac{1}{\sigma}} (L_{R,M})^{\frac{\sigma-1}{\sigma}} + (L_{NR,M})^{\frac{\sigma-1}{\sigma}})^{\frac{\sigma}{\sigma-1}} \quad (1)$$

$$S = \sum_{j=1}^{J_S} A_S^j (z_{R,S}^{\frac{1}{\sigma}} (L_{R,S})^{\frac{\sigma-1}{\sigma}} + (L_{NR,S})^{\frac{\sigma-1}{\sigma}})^{\frac{\sigma}{\sigma-1}} \quad (2)$$

The economy has a fixed total endowment of labor \bar{L} . This can be divided into routine labor L_R and non-routine labor L_{NR} , into employment in both sectors E_M and E_S , or by labor and sector. Total time of the representative household is split between the two possible occupations and industries.

$$\bar{L} = L_R + L_{NR} = E_M + E_S = L_{R,M} + L_{NR,M} + L_{R,S} + L_{NR,S} \quad (3)$$

The model has two periods. In the first period, workers are mobile between firms and jobs such that the wage in all occupations is identical. In the second period, in anticipation of a technological change, workers can move between occupations by paying a re-skilling cost C which is increasing and convex in the quantity of labor that switches occupations. Welfare for the representative household in the second period is equal to total wage less these re-skilling costs.

$$U = w_R L_R + w_{NR} L_{NR} - C(|L_R - L'_R|) \quad (4)$$

where L'_R is the wage equalizing share of workers in the routine task before the technological change. Workers will move between occupations so as to maximize their total second period wage less re-skilling costs.

$$\frac{\partial C}{\partial L_R} = |w_R - w_{NR}| \quad (5)$$

which is perfectly mobile between occupations and industries.

There is no specification for the form of aggregate demand. Along that dimension, this framework is very general. We now have enough of a framework to begin analyzing how different types of technological change will influence employment across and within firms.

4. Total Factor Productivity Changes Across Firms

Suppose that firms in industry M get more productive (i.e. A_{Mj} increases for some or all firms). Total employment of routine workers can be written as

$$L_R = E_M \frac{L_{R,M}}{E_M} + E_S \frac{L_{R,S}}{E_S} \quad (6)$$

where E_M and E_S are total employment in the manufacturing and service industry respectively. Total employment is equal to the number of employees working in each type of firm times the share of workers in each type of firm who are routine. The total effect of a technological change can be understood as the sum of a change in the share of workers employed *across* different firms, and a change in the mix in the share of workers who are routine *within* different types of firms.

Consider the effect of a technological change which boosts the TFP of some or all manufacturing firms. This could either be a global effect, such as due to a trade shock, or because a new technology raises the total factor productivity of new (or potential new) firms. This change will have an effect on the total share of workers doing routine work through two mechanisms.

The total effect of an increase in the productivity of manufacturing firms can be written as

$$\frac{\partial L_R}{\partial A_M} = \frac{\frac{\partial E_M}{\partial A_M} \left(\frac{L_{R,M}}{E_M} - \frac{L_{R,S}}{E_S} \right)}{1 - \frac{\partial \frac{w_R}{w_{NR}}}{\partial L_R} \left(E_M \frac{\partial \frac{L_{R,M}}{E_M}}{\partial \frac{w_R}{w_{NR}}} + E_S \frac{\partial \frac{L_{R,S}}{E_S}}{\partial \frac{w_R}{w_{NR}}} \right)} \quad (7)$$

The numerator of (7) summarizes the across firm effect. When manufacturing firms become more productive, this can increase or decrease the share of workers in that industry. If the products of the

two industries are gross complements, the increase in manufacturing firm productivity will decrease the amount of workers needed in those. By assumption manufacturing firms have a higher initial share of workers performing routine tasks. In this case, the across firm effect is to decrease the total amount of routine workers through reducing the share of employment in firms which are initially routine intensive. The denominator summarizes the within firm effect. A decrease in total relative demand for routine workers will decrease the relative wage of routine workers and vice versa. If the relative wage of routine workers increases, then their share of employment in both industries will decrease with the relative importance of both decreases weighed by initial employment in both industries.

Notably, this denominator will always be greater than or equal to one. In other words, any firm level TFP change that leads to a decrease in routine employment will be partially offset by wage decreases for routine workers. This will lead to within-firm rebalancing which favors routine workers.

If workers are perfectly mobile between occupations then the denominator will be 1. In other words, in the special case with perfect labor mobility, (7) reduces to just the numerator. Similarly, if firms are Leontieff in production, and do not change their labor mix as a function of changes in wages then the denominator will also equal one. In either case, the effect of the TFP change on total routine employment will be entirely driven between reallocation in employment across firms. In these cases the equation (7) reduces to just the numerator.

Consider the special case where the technological change decreases employment in industry and production is Leontieff. Intuitively, in that case the technological change will reduce routine employment so long as manufacturing industry is more routine skill intensive. Then

$$M = \sum_{i=1}^{J_M} A_M^i [\min(z_{R,M} L_{R,M,i}, L_{NR,M,i})] \quad (8)$$

$$S = \sum_{j=1}^{J_S} A_S^j [\min(z_{R,S} L_{R,S,j}, L_{NR,S,j})] \quad (9)$$

and the total change in routine employment is only the across firm effect, so the total change in routine employment is

$$\frac{\partial(L_{R,M} + L_{R,S})}{\partial A_M} = \frac{\partial(L_{R,M} + L_{NR,M})}{\partial A_M} \left(\frac{1}{2M+1} - \frac{1}{2S+1} \right) \quad (10)$$

The first effect is through aggregate demand.

Two examples of models with this form of technological change are Bessen (2018) and Autor et al. (2017). In Autor et al. (2017)'s model of superstar firms, increasing industrial concentration explains decreases in labor's share of income. As globalization proceeds, the most productive firms, which tend to have low labor shares, capture an increasing percentage of the market. If this is the correct model of a decrease in labor share, then it could also be a cause of SBTC. If the most productive firms have distinctive employee mixes, then as they disproportionately expand, the occupational composition of the economy will change as well.

In Bessen (2018), reductions in employment for factory workers may be due to TFP changes in the manufacturing industry. That industry disproportionately employs routine manual workers and has seen both increasing import competition and technological progress.

Under either of these hypotheses, the most important source of SBTC will be across-firm growth. Due to the productivity shock, relatively routine-intensive firms will decrease in employment and vice versa. Some of SBTC may also be caused by entry or exit of firms, depending on whether all firms get the TFP shock, or only new ones. However, this form of SBTC is either firm employment mix neutral (if, for example, within firm production is Leontieff), or will be biased *against* non-routine employment. In general the total employment shift towards non-routine employment will be mitigated by within firm employment mix rebalancing.

In the model, there are three factors leading firms to change their employment mix. Shifts in productivity can directly lead to either more or less hiring depending on price elasticity of demand in the product market. If the two industries have different occupational compositions, this will change the overall composition of employment. The second effect is the direct effect of task specific productivity changes. This leads firms to change their occupation mix. Finally there is an indirect effects through wages. If demand for a certain occupation changes, this can lead to changes in relative wages. Shifts in relative wages causes firms to substitute toward cheaper labor.

5. Skill-Biased Technological Changes Within Firms

The other type of technological change possible in this framework is due to a change in Leontieff production functions. This will lead to SBTC within firms in addition to between them.

Suppose that a new technology is invented to automate some routine tasks in the manufacturing industry (i.e. z_M increases). We continue to restrict attention to the special case where production within firms is Leontieff. Then the overall impact on routine employment will be

$$\frac{\partial(L_{R,M} + L_{R,S})}{\partial z_M} = -\frac{L_{R,M} + L_{NR,M}}{(z_M + 1)^2} + \frac{\partial(L_{R,M} + L_{NR,M})}{\partial z_M} \left(\frac{1}{z_M + 1} - \frac{1}{z_S + 1} \right) \quad (11)$$

This equation has two terms. The first term is the direct impact of the automation on routine employment in the manufacturing industry. The second term due to any movement between industries. It will be zero if the two industries have the same initial mix of routine and manual workers (i.e. $z_M = z_S$). The second term will also be zero if the productivity effect of easier production of manufactured goods perfectly cancels the price effect. This would be the case if aggregate demand were Cobb-Douglas.

Examples of models with technological change analogous to this type include Peretto and Seater (2013), Benzell et al. (2016), and Acemoglu and Restrepo (2017). In these models, firms can make investments, either in capital or R&D, to reduce their demand for certain types of workers. If

this is the correct model of SBTC, we should expect to see occupational changes being driven by firms changing their employment mixes. Depending on the strength of the second term, we would still expect to see a role for the other sources. Additionally, given the key role for firm capital and R&D investment in these models, we would expect to see firms with larger stocks of these investments see more dramatic within firm employment mix rebalancing. While the Leontieff case demonstrates the dynamics under a simple assumption of intense complementarity between labor input types, much of the intuition for the change in labor demand remains relevant under less restrictive assumptions.

6. Data and Methodology

Our dataset is created by merging LinkedIn's database of hundreds of millions of position records with Compustat NA firm financial data, occupation-level BLS OES wage and employment data, and O*NET occupational task data. Employment by firm and occupation is constructed from individual LinkedIn resumes and aggregated to 136 occupational groups. BLS occupations are many to one matched to LinkedIn occupation groups. For LinkedIn occupations nesting several BLS occupations, the weighted (by employment) average characteristics (such as wage) of the occupations are attributed to the LinkedIn occupation.

Individuals employed by some firms and with certain occupations are disproportionately likely to have LinkedIn accounts. We normalize the LinkedIn employment measures to account for both sources of bias. As in Brynjolfsson et al. (2018b), we use total firm-level employment from Compustat and industry-level occupational composition from BLS. First, we sum firm-level employment in Compustat to measure total public firm employment by (3-digit NAICS) industry. Starting with the amount of employees we observe within each industry, we use the BLS's estimates of occupational mix by industry to construct the total amount of employees of each occupation in our sample. Aggregating across firms, we then generate an estimate of the total quantities of workers by year in each of the specific occupations in the BLS Occupational Employment Survey. We then estimate LinkedIn coverage by occupation-year. Each firm-occupation-year tuple is deflated (or inflated) for its

coverage relative to our Compustat/BLS estimate. Subsequently, we re-adjust our Compustat/BLS measures such that total firm employment is equal to total firm employment in Compustat. This process maps what we observe in the BLS estimates and the LinkedIn employment estimates by year, occupation, and firm to an estimate that adjusts for occupation and firm-specific differences in LinkedIn coverage. This gives us our measure of the occupational labor force within each firm-year.

Firm and year-level measures of firm capital are taken from Compustat and R&D stocks are taken from Peters and Taylor (2017) available via WRDS. R&D stocks are measured as a perpetual inventory using the BEA's industry specific depreciation rate. We extend the Peters and Taylor R&D stocks to 2016, following Brynjolfsson et al. (2018a). We estimate a firm level wage bill by multiplying our measure of each firm's employment by occupation by the BLS mean annual wage. Our final dataset consists of over 2.5 million firm-year occupation tuples, containing 3662 firms, with an average of 14,687 employees.

We organize occupations and firms into categories for convenient analysis. For occupations, we sort occupations into approximately equal, in terms of year 2000 US publicly traded firm employment, thirds. The first organization is in terms of the routineness of the task. Following Acemoglu and Autor (2011) routineness is measured as the sum of three O*NET questions in 2006 (the earliest year available): (4.C.3.d.3) Pace determined by speed of equipment; (4.A.3.a.3) Controlling machines and processes; (4.C.2.d.1.i) Spend time making repetitive motions. We also organize occupations into equal employment thirds based on year 2000 mean wage.

One way firms are classified is by whether they are in our sample and have positive employees in a given year. If a firm has positive employees in our sample but then leaves the sample before 2016, we say that firm exited. Firms can exit our sample through being bought out, through being de-listed, or through bankruptcy. Firms that have no record of employees, but gained them

before 2016 are labeled as entering. We also classify firms into quintiles by estimated wage bill growth over the interval.

7. Net Job Creation Within and Across Firms

Figures 1 and 2 display total net employment growth, and its components, for selected occupations for firms in our sample from 2000 to 2016. In both figures, the blue bar displays the total net number of jobs created. This is the sum of the next four bars. As figure 1 shows, about one hundred and six thousand data entry clerk jobs were eliminated in our sample over this interval. The green and red bars in each figure distinguish the share of employment changes driven by heterogeneous growth across firms and within firm employment rebalancing respectively. The green bar (across firm growth) indicates how many net jobs would be added if every firm kept their shares of employment constant between 2016 and 2000.

For occupation group j , the across firm growth effect is defined as

$$G_j = \sum_{i=1}^I \left(Emp_{i,j,2016} - \frac{Emp_{i,j,2016}}{\sum_{j=1}^J Emp_{i,j,2016}} \sum_{j=1}^J Emp_{i,j,2000} \right) \quad (12)$$

where emp_{ijt} is the number of employees at firm i , of type j in year t . J is the total amount of occupation groups and I the total amount of firms. The green bars would all be of equal heights if every firm had the same employee mix in 2016, or if all firms added employees at the same rate since 2000. The red bars indicate the share of employment growth due to firms that exist in both 2000 and 2016 changing their employment mixes. We define the within firm rebalancing effect as

$$W_j = \sum_{i=1}^I (Emp_{i,j,2016} - Emp_{i,j,2000}) - G_j \quad (13)$$

A firm can only rebalance towards one occupation if it moves away from another, so rebalancing changes sum to zero. The final source of net job growth is firm entry and exit. The teal bar indicates the number of employees of a type in 2000 at firms that leave the sample before 2016. The orange bar

indicates the number of employees in 2016 of a type at firms which enter the data after 2000. Exit of firms can only eliminate jobs, and entry can only add them.

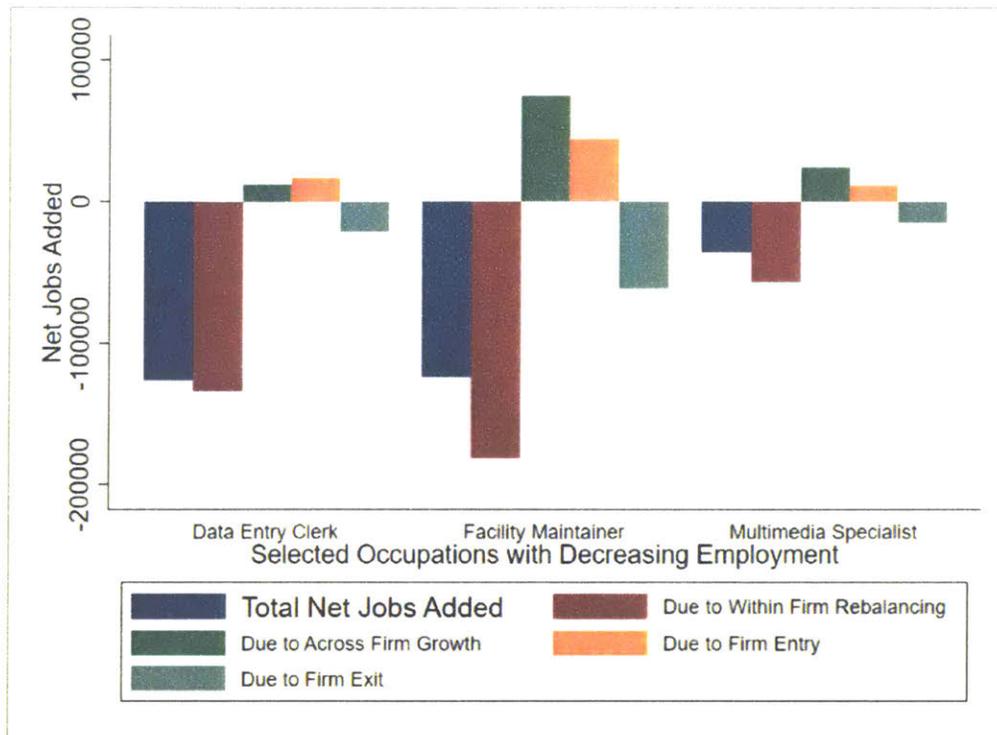


Figure 1: Net Job creation by source for US publicly traded corporations from 2000 to 2016. Selected Occupations.

As shown in figure 1, data entry clerks, facility maintainers, and multimedia specialists lost a net of 126 thousand, 124 thousand, and 36 thousand positions. Data entry clerk is an occupation in the middle third of routineness, and multimedia specialists and facility maintainers are in the top third of routineness. In all of these occupations, the total amount of jobs lost is less than the amount lost due to firings and replacement. For data entry clerks, the amount of jobs lost due to rebalancing is approximately equal to the total amount of jobs lost. This is because firms that employed these types of workers did not grow much, nor did many more firms that hired such workers enter on net. On the other hand, firms that hire facilities maintainers and multimedia specialists increased their employment somewhat more. Therefore, total job loss for these occupations was more limited. Across

all three of these occupations, entry and exit created very few jobs on net. Note that none of these occupations saw net reductions in employment due to firms shrinking in employment.

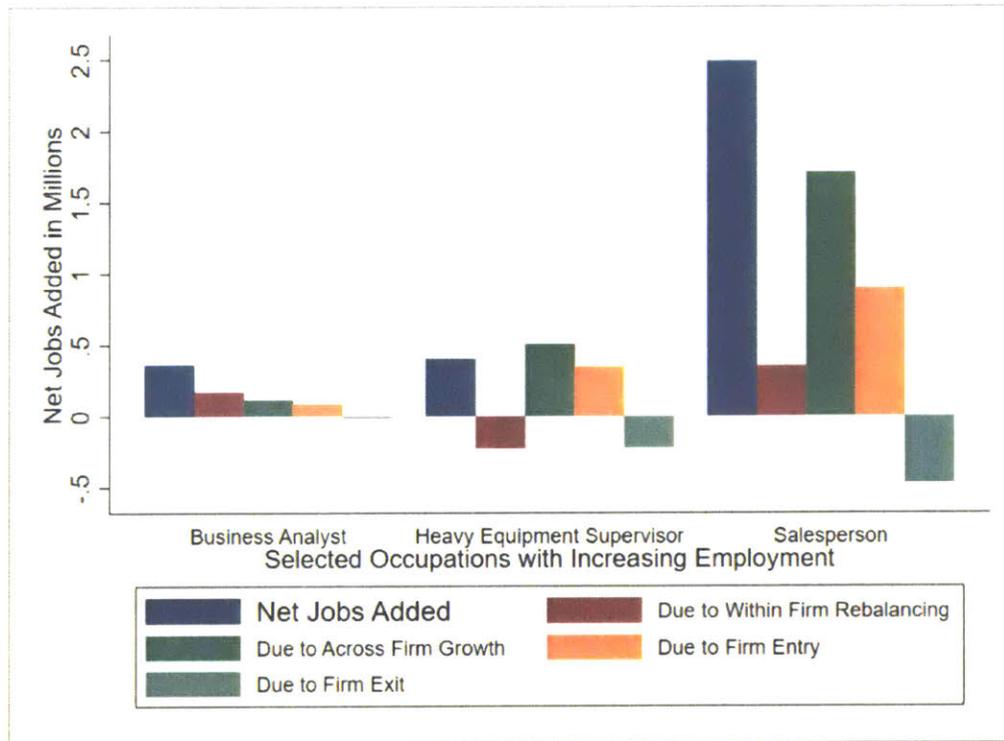


Figure 2: Net job creation by source for US publicly traded corporations from 2000 to 2016. Selected occupations.

Figure 2 repeats this decomposition for three of the fastest growing occupations. Salesperson is the most common occupation in the data, and also the fastest growing. The growth in this occupation is almost entirely due to across firm growth. Heavy equipment supervisors also saw growth in employment despite some within firm rebalancing against them. Business analysts also saw employment increase, in roughly equal measure due to across firm growth, within firm rebalancing, and net entry.

Considering the figures as a pair, it is clear that all components of net job creation need not go in the same direction. However, it does seem to be the case that then number of net jobs created lost by shrinking occupations is closer to and more correlated with within firm rebalancing than for growing

occupation. Figure 3 reflects this relationship. This figure relates an occupation's employment change due to within firm rebalancing against its total change. As can be seen, the top 12 fastest growing occupations are scattered about far from the 45 degree line. The 12 fastest shrinking occupations are right on top of it. Also notable is the green triangle in the upper left. This is Heavy Equipment Supervising. This is the only occupation of those selected with positive employment growth despite firms in the sample replacing them on net.

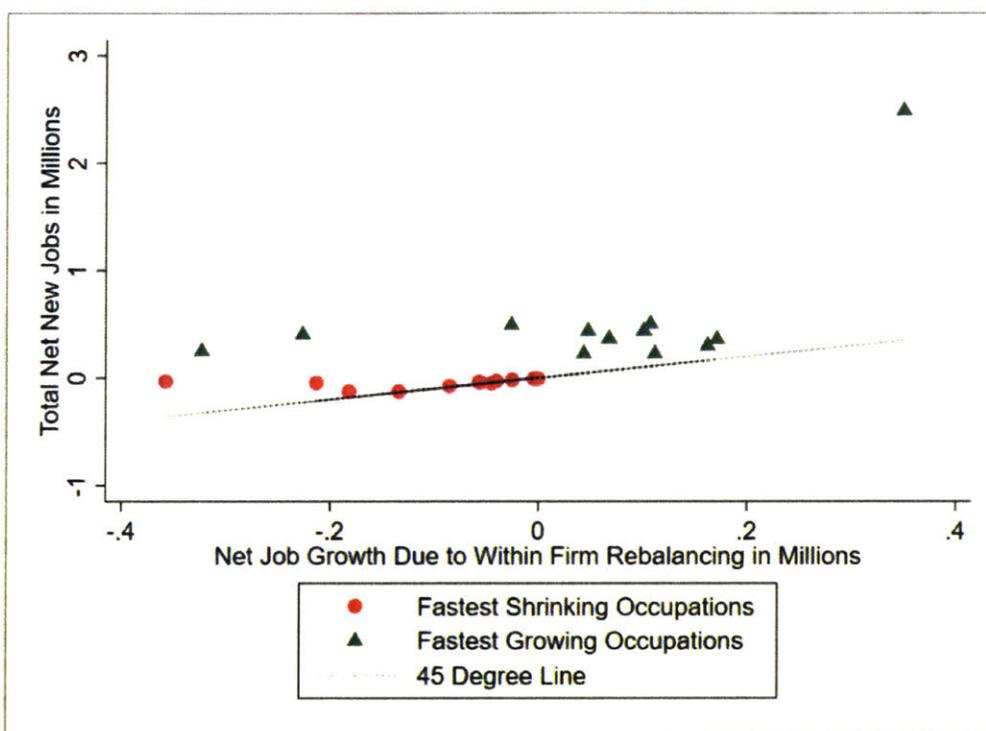


Figure 3: Plot of total net new jobs created and the portion due to within firm rebalancing for the 12 fastest growing and contracting occupations.

8. Skill-Biased Technological Change Within and Across Firms

While it is informative to look at individual occupations, the best way to understand SBTC in the economy as a whole is to group occupations into categories. In figure 4 we present the same decomposition of net employment growth into its components with one key difference. Rather than examining individual occupations, this figure sorts all occupations into one of three categories by their

routine task intensity. Occupations are sorted into three groups with approximately equal employment in 2000.

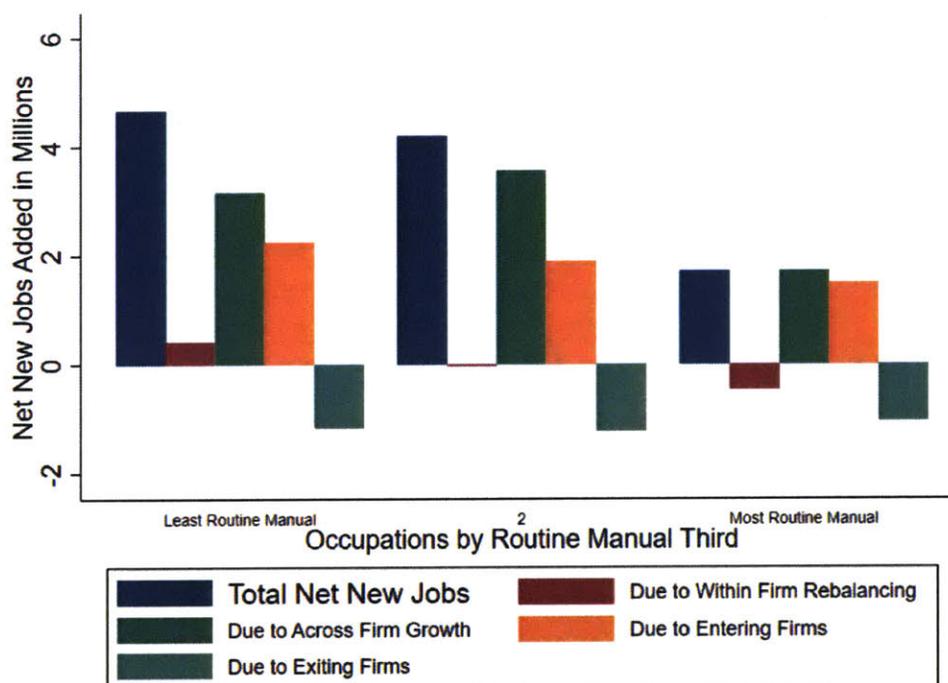


Figure 4: Net job creation by source for US publicly traded corporations from 2000 to 2016. Occupations organized into thirds based on routineness.

SBTC for firms in our sample over this interval was highly biased against manual tasks. As a measure of this, we define the non-routine bias of employment growth as the difference between bottom third routine occupation employment growth less top third routine occupation employment growth. So,

$$GB_R = (E_{R=1,2016} - E_{R=1,2000}) - (E_{R=3,2016} - E_{R=3,2000}) \quad (14)$$

We find that the non-routine bias of employment growth was 2.95 million over this interval. The most important source of this SBTC was faster growth across firms with high levels of non-routine workers, accounting for 48.9 percent of this bias. The next most important source of this SBTC was within firm rebalancing, which accounts for 30.0 percent of the bias in public firms. New firms

also had relatively low levels of routine workers, and accounting for 25.7 percent of the total change. The exit of firms worked slightly against SBTC, explaining -4.6 percent of the trend. These findings are consistent with the hypothesis that non-skill-neutral changes in the production function of firms are critical to explaining aggregate trends in occupational employment.

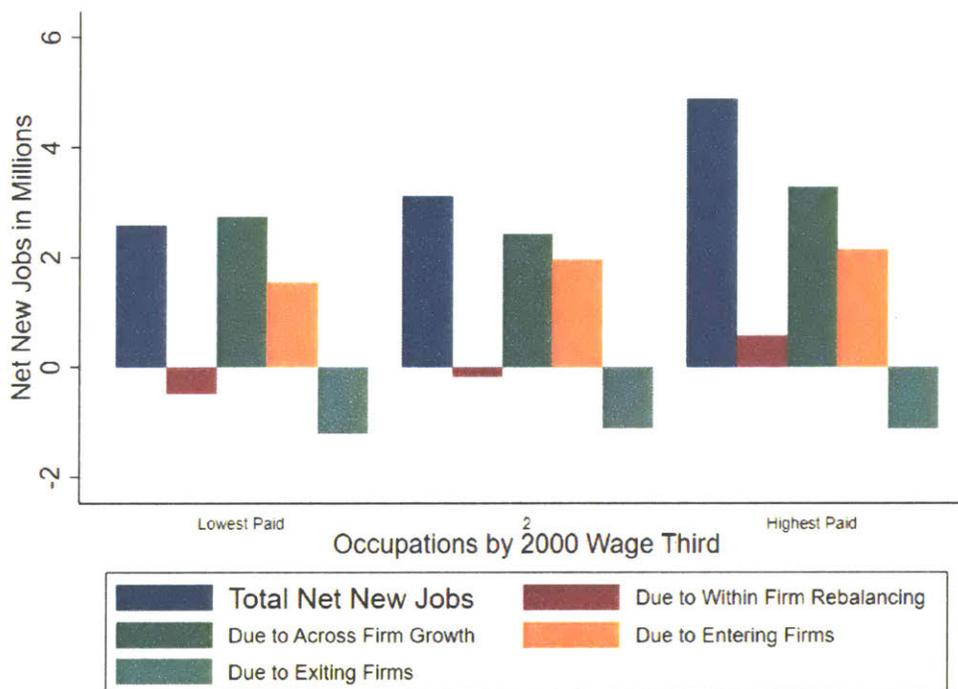


Figure 5: Net job creation by source for US publicly traded corporations from 2000 to 2016. Occupations organized into thirds based on 2000 average wage.

Figure 5 repeats this analysis, this time after dividing occupations into thirds based on their 2000 average wage. We find that total employment was not strongly high wage-biased. 2.02 million more jobs were created in the top third of occupations than the bottom. In contrast with Autor and Dorn (2013) we measure the change over this interval as being strictly skill-biased, rather than polarizing. First, by restricting our attention to publicly traded US firms over the interval 2000 to 2016, we are considering a somewhat different subset of the economy than previous studies. Second, the version of LinkedIn's occupation classification system we use can be somewhat coarse across levels of seniority for occupations with similar routine task intensities. This doesn't impact the

routineness measure, but does cause us to lump together occupations of dissimilar year 2000 wage in some cases. For this reason throughout the paper we focus on organizing occupations by routineness. Still, the results we do have on the sources of SBTC when occupations are organized by wage are consistent with those on routineness. Namely, rebalancing within firms plays an important role.

9. Firm Employment Mix Dynamism

Given that within firm rebalancing is an important component of SBTC, it is natural to ask how the rate at which firms change their workers has changed across firms and time. We define firm employment mix dynamism as the average absolute value of the change in employment share for occupations in a firm. Or,

$$D_t = \frac{1}{136} \sum_{j=1}^{136} \left| \frac{E_{j,t}}{\bar{E}_t} - \frac{E_{j,t-1}}{\bar{E}_{t-1}} \right| \quad (15)$$

where $E_{j,t}$ is employment at a firm of a given occupation, t is the period (here, years) and \bar{E}_t is total firm employment. Since we have 136 principal employment type groups, this give the average absolute value shift in employment share over all categories.

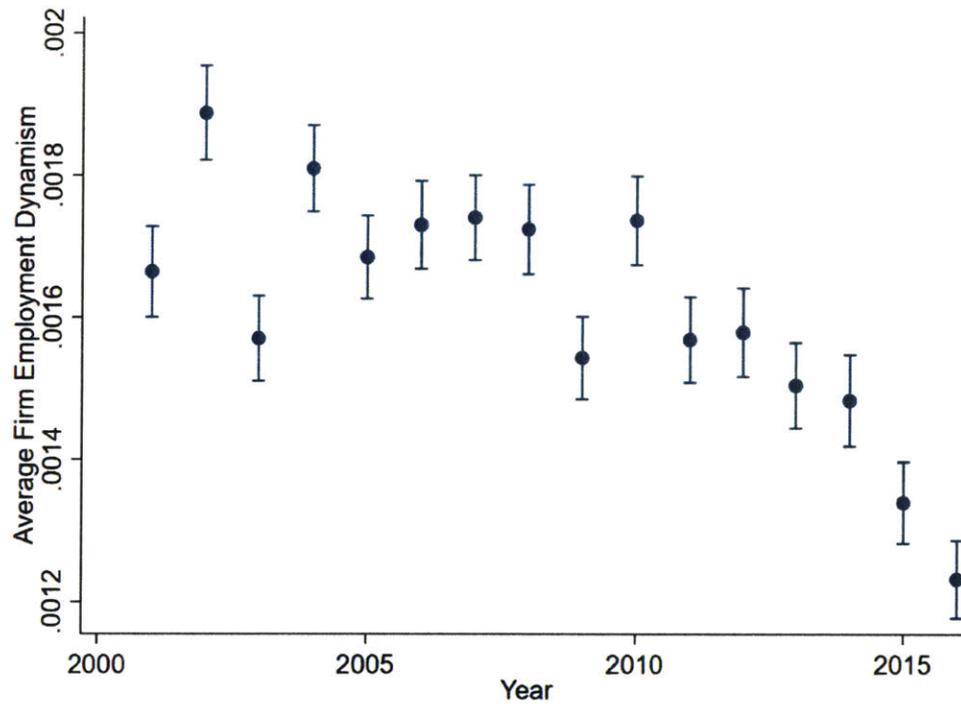


Figure 6: Average employment mix dynamism by year. 95 percent confidence intervals for sample means.

Figures 6 and 7 display trends in average firm dynamism over time. Figure 6 treats all firms equally, while Figure 7 weighs firms by employment and includes total factor productivity growth for the entire US economy. In both figures there is no year 2000 observation, because both the year and the previous year's firm occupation mix need to be observed. For this reason also, firms must be in the data in consecutive years for an observation to appear. Other measures, summarized by Decker et al. (2014) had shown that US firm dynamism has been on the decline. Our measure looking at labor mix shifts is consistent with these other measures.

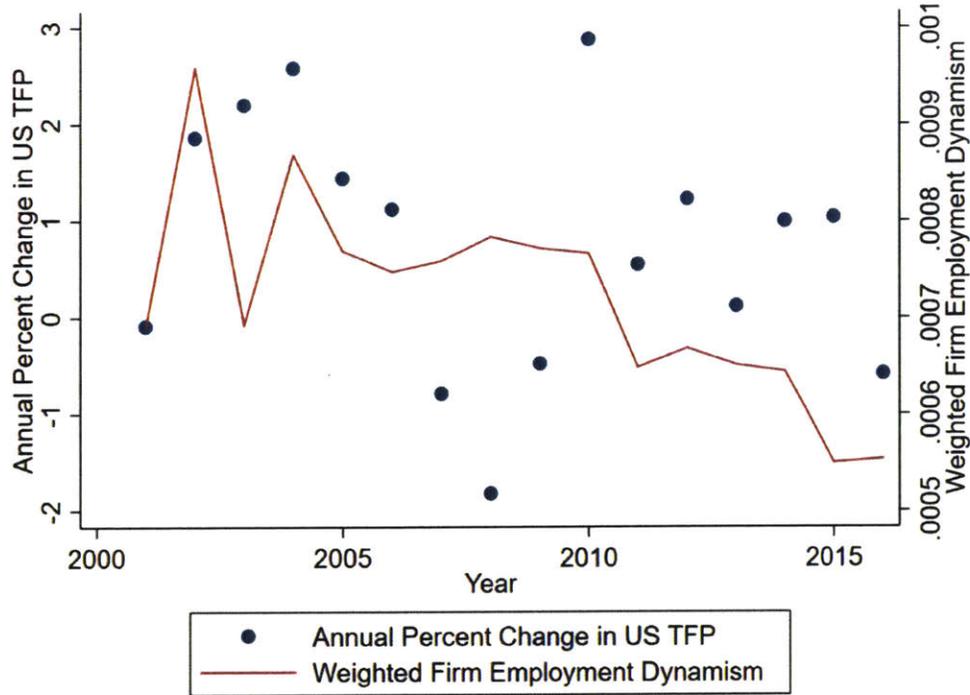


Figure 7: Average employment mix dynamism (weighted by firm employment) and yearly percentage change in U.S. TFP.

We measure a steady decrease in the frequency at which companies change their employees mix. As can also be seen in figure 7 weighing firms by employment when constructing this measure leads to a less dramatic decrease over time. During this period (from 2000 through 2016) the number of firms in the economy overall, as well as publicly traded firms in our sample, declined. The average firm gained more employees. Therefore, it is reasonable to ask whether the decrease in firm employment mix dynamism has been driven or counteracted by an increase in the average size of firms.

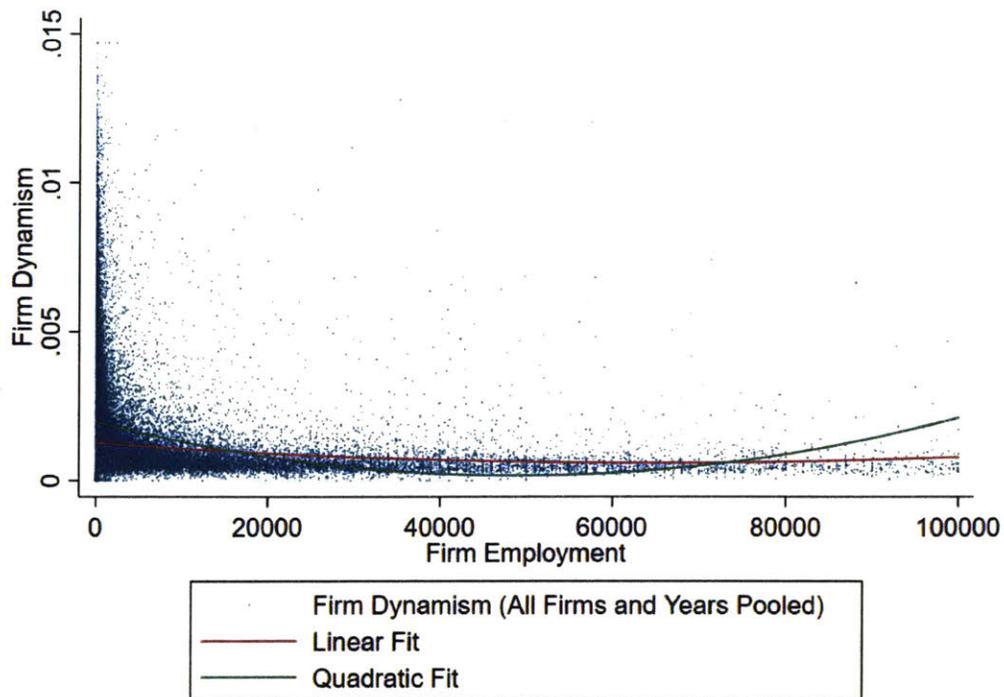


Figure 8: Firm employment mix dynamism and firm employment. All years. Firm-years with less than one hundred thousand employees.

Figure 8 presents the relationship between firm employment mix dynamism and firm employment, pooling all firms and years (with less than one hundred thousand employees). The range of firm variances is highest for very small firms. This is unsurprising because for smaller firms much fewer employees in absolute number need to be replaced (or promoted/demoted to a different occupation in the firm) in order to be measured as having high employment mix dynamism. For these very small firms measurement error also becomes important because each individual LinkedIn observation is weighted to represent several real employees. For large firms, the true and estimated number of employees by occupation is approximately the same. However, for very small firms, discreteness in the number of LinkedIn workers observed becomes important.

Despite this mass of high employment mix dynamism among very small firms, the measure shares no clear relationship overall with firm employment. When weighing by firm employment, a

quadratically fit curve is almost perfectly flat. Future work will explore why some firms across all levels of employment have more dynamic labor mixtures.

10. Skill-Biased Technological Change Within and Across Firms by Industry

To deepen our analysis of within-firm rebalancing, across-firm growth, entry and exit as sources of skill-biased technological change, we repeat our analysis at the industry level. One hypothesis is that industry-level total factor productivity shifts are the most important cause of SBTC. If that is the case we should expect to see non-routine biased employment growth be the most dramatic in industries that already had a relatively low number of routine workers. This would imply that industrial compositional changes are driving SBTC for the aggregate economy.

Figure 9 displays the initial non-routine bias, initial employment level, and non-routine bias of new jobs created for the 17 largest industries in our data. The initial non-routine bias of an industry is the number of employees in the bottom third of routineness less those in the top third of routineness in 2000.

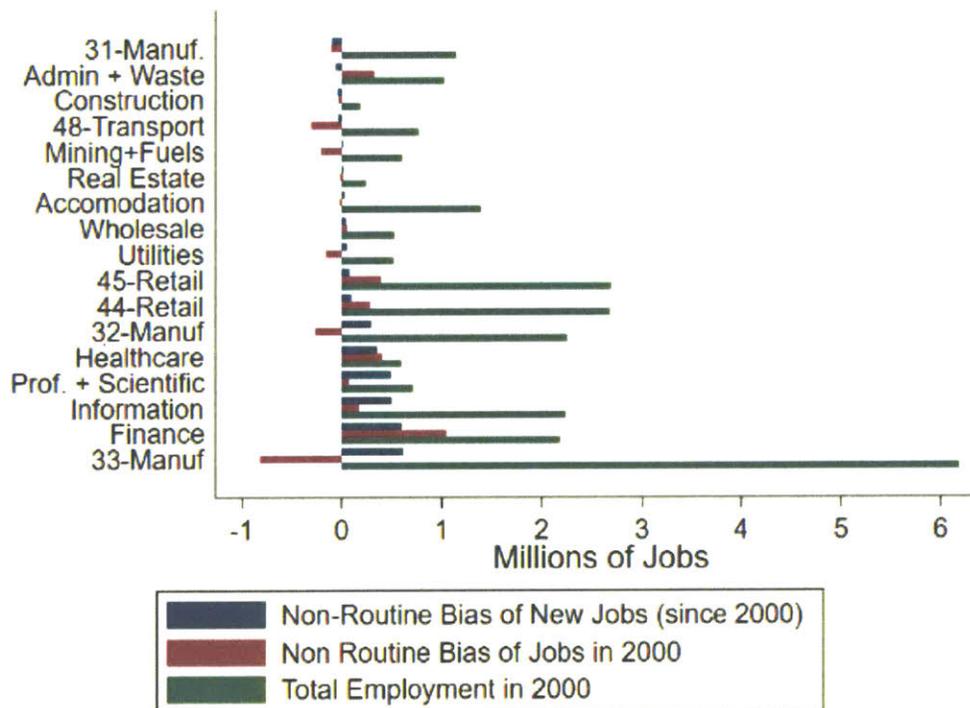


Figure 9: Total employment in 2000, non-routine bias of employment in 2000, and non-routine bias of net new jobs by industry. Largest 17 two-digit industries, sorted by non-routine bias of new job growth. The non-routine bias of employment in a given year is the difference between bottom third routine and top third routine occupation employment. The non-routine bias of employment growth is as defined above.

The industry with the most non-routine-biased employment growth is category 33 manufacturing firms, despite being strongly biased against non-routine labor in 2000. Net employment growth in that industry has been sufficient to switch the industry from employing more bottom-third than top-third routine employees to the opposite. The other industries with strong movement towards non-routine workers are mixed between those that were already strongly non-routine-biased to begin with (including Finance and Healthcare), those who were roughly neutral (including Information, Retail, and Professional and Scientific firms), and those that were biased against (32-Manufacturing). No industry got significantly more routine-biased as a result of new hiring.

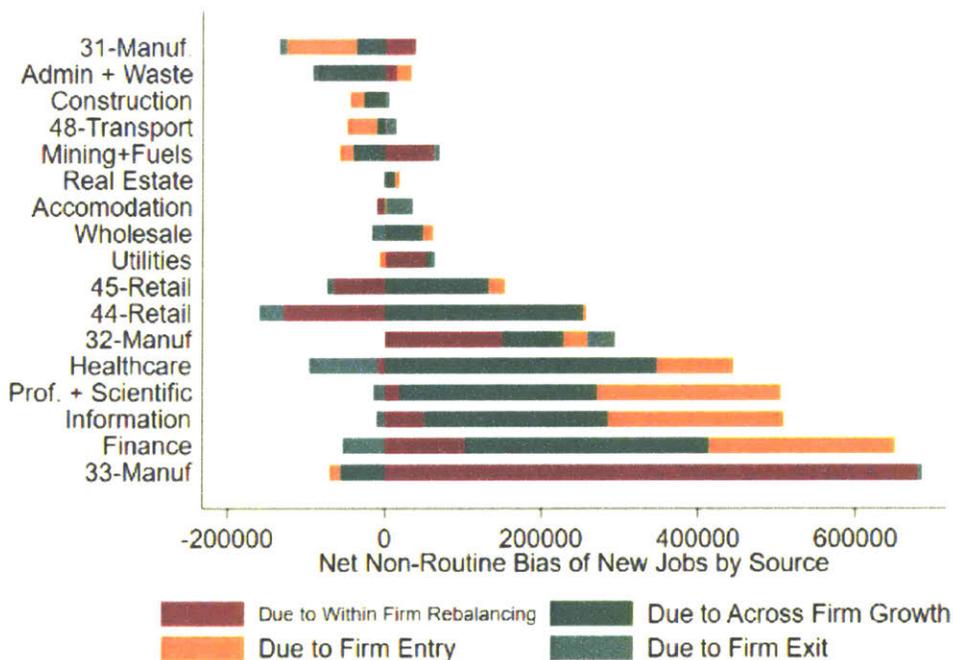


Figure 10: Non-routine bias of new jobs by source. Summing the negative sources to the positive sources yields the blue bar in Figure 9, corresponding to the total net non-routine bias of new employment in the industry.

Figure 10 decomposes the non-routine bias of employment growth into its components by industry. For the manufacturing industries, within-firm rebalancing is by far the most important cause of non-routine-biased employment change. At first this might seem mechanical, given that Figure 9 suggests that employment change was very non-routine-biased, and that manufacturing firms have not had large overall employment growth.⁸⁵ For publicly traded US firms in industry 33-Manufacturing, employment increased by 1.4 million.} Yet even in a shrinking industry, non-routine biased employment growth can be caused by firms with a relatively high amount of routine workers contracting faster than others. This was not the case. Notably, the across firm growth effect is actually slightly negative for 33-Manufacturing firms. Of all industries, only Healthcare seems perfectly consistent with the hypothesis that an industry-wide boom is leading to skill-biased technological change. For industries like Finance, Retail, and Professional and Scientific Services, across firm growth is the most important cause of non-routine- biased employment growth. This is more consistent with the Superstar firm hypothesis. The industry-level dynamics vary considerably. Industry compositional shifts are therefore a primary component of SBTC.

It is also interesting to note the great variety in the contributions of firm entry and exit to non-routine-biased employment growth across industries. For Finance, Information, and Professional and Scientific services, three industries with many successful 'unicorn' entrants in recent years, entry has been an important contributor. In these industries exit has been slightly routine-biased, largely because all firms in the industry are non-routine intensive. The non-routine bias of exit closely inversely follows the routine-bias of firms in the industry in 2000.

A final question we can ask is how employment mix dynamism varies by industry. Figure 11 shows the average, firm employment weighted, firm employment mix dynamism for total economy

⁸⁵ For the country as a whole, employment in all manufacturing industries declined from 17.3 million in 2000 to 12.4 million in 2016 (BLS 2018)

and selected industries. All industries experience an overall downtrend in dynamism. In this figure an industry's employment mix dynamism is unrelated to industry's routine-biased employment growth due to within-firm-rebalancing. It is unclear whether this is because some industries are inherently more dynamic, or a statistical relic of the fact that some industries have a more diverse array of occupations. The range for the measure is quite narrow, and remarkably stable over the observation period.

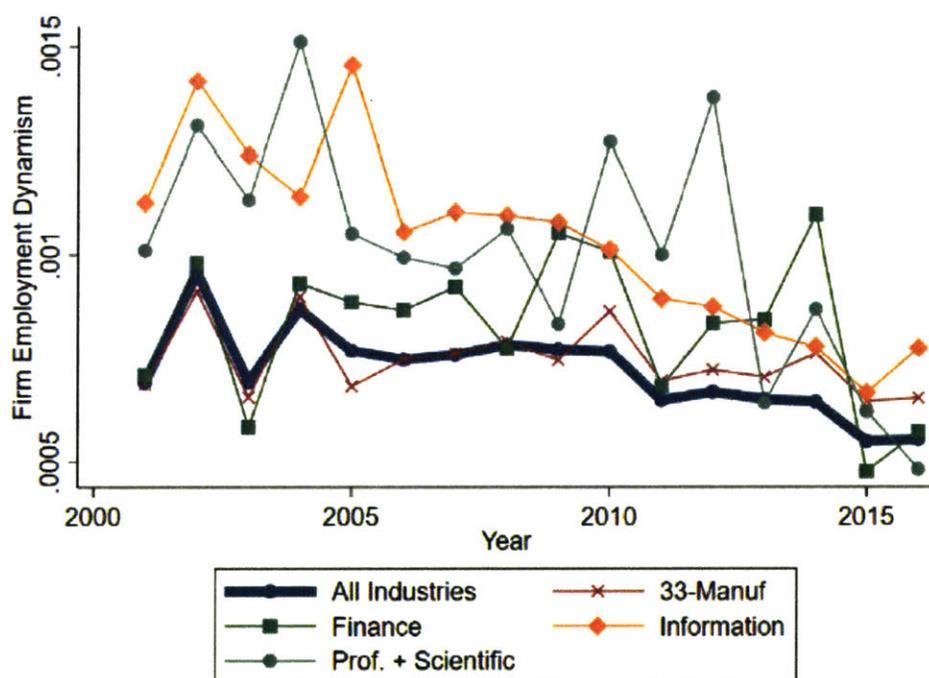


Figure 11: Average (employment-weighted) firm employment mix dynamism for total economy and selected industries

11. SBTC and Firm Investment

As documented above, within-firm rebalancing has played an important role in the relative decline of routine-task intensive occupations. As a final exercise, we document the relationship between within-firm rebalancing and other firm characteristics. This is important because to the extent these relationships are causal, they can drive policies meant to help or hurt the automated. We find that

R&D investment tends to increase firm level demand for routine occupations, suggesting that R&D subsidies may be an effective mechanism for combating occupational displacement.

All theories of SBTC entail some connection between innovation and demand for different tasks. In "factor eliminating" models of technological change, such as in Peretto and Seater (2013) firms can make an investment in R&D to change the factor intensity of their production functions. Over the last 35 years, real interest rates have steadily declined and general automation technologies advanced (Benzell and Brynjolfsson 2018) Therefore, it should be increasingly profitable for firms to develop new products and processes that substitute cheap capital for expensive laborers. Under this hypothesis, we would expect to see firms making greater R&D expenditures experiencing more SBTC. In models of directable technological change at the firm level, such as (Acemoglu and Restrepo 2017), firms can choose between either inventing new tasks for low cost labor, or automating old tasks. If routine tasks are both more easily automated and lower wage, then these models are ambiguous in their prediction of the effect of R&D expenditure on routineness and average firm wage. However, if all tasks are equal in their ability to be automated, then R&D should lead to a lower average firm wage.

However, an OLS regression of R&D on firm routineness or average wage is likely to be misleading. This is because many omitted variables exist that could drive both the mix of workers employed by a firm and R&D decisions. Therefore we require an instrument which can predict R&D investment without being otherwise correlated with routineness.

The instruments for R&D we use are state and federal tax laws as codified in Bloom et al. (2013) and updated through 2015 in Lucking et al. (2018). State tax laws vary the attractiveness of R&D through tax-credits, depreciation allowances, and corporate taxes. These tax incentives differentially apply based on where a firms' R&D is located. To determine which state tax laws a firm is exposed to, Bloom et al. (2013) use the 10 year moving average location of a firms' patent holders.

Federal tax incentives for R&D vary with time and are a function of a firms' historical R&D and revenues. The specific instruments used are the log of the tax price component of R&D user cost.

The first stage results of the state and federal tax component of R&D cost on a firms' R&D stock is reported in table 1. The results are broadly consistent with those reported in Lucking et al. (2018). In each specification forms of tax cost are negatively correlated with R&D stocks. However, we focus on specifications with firm and time fixed effects, as state R&D credits is likely correlated to a states' latent attractiveness for research, and federal R&D likely vary over time as a function of the macroeconomic environment.

Table 2 reports the effect of R&D, after instrumentation with tax policy, on firm routineness. We measure routineness as the average routine-task intensity of occupations at a firm, normalized by year. Across specifications with firm fixed effects, we find a significant positive relationship between a firm's R&D stock and the average routineness of its workers.

This effect is quite large. The most complete specification indicates that a firm that doubles its R&D stock will see the routineness of its average occupation increase .355 standard deviations. To give a sense of how large an effect this is, consider that the average information firms' routineness in 2016 is -.048. In the same year, the construction industry's average routineness is .088. In other words then, a firm increasing its R&D stock by 39% moves it from having the average routineness of an information industry firm to the average routineness of a construction industry firm. It is possible that firms receiving a tax credit to make additional investments in R&D contemporaneously increase the quantity of routine workers hired to implement existing productive processes. It might also be the case that automation activities are infra-marginal, but investment in generating new tasks for routine workers is cross-subsidized by the R&D policy changes. The long-run effects of the R&D policy shifts on routineness are not captured by this analysis. Even if the local average treatment effect for suggests R&D stock increases cause higher average firm routineness, in the long-run firms might devote more

resources to SBTC. However, were this result to generalize, it would suggest R&D tax credits as a possible channel to mitigate the short-term impact of automation technologies on labor. Financial constraints may be another possible confound. If firms respond to R&D subsidies by hiring more flexible sources of labor because of anticipated future difficulties in retaining workers, we would expect to see more routineness in the companies that take up the subsidy. Lucking (2019) provides evidence in support of job creation from innovation for these firms, and fails to find support in favor of financial constraints driving the results.

This result, relating firm R&D stock to routineness, is not driven by manufacturing firms alone. Dividing industries into manufacturing (NAICS codes of 31, 32, or 33) or non-manufacturing, as in table 3, both show a significant negative relationship. However, the F-statistic for non-manufacturing firms is low, indicating these results are potentially biased by a weak instrument. In addition to these results, R&D stocks, instrumented with tax credits, are strongly positively associated with firm size, and not strongly related to the firms' average occupational wage.

There are many extant theories that would connect anticipated wage increases, firm investment, and routineness. For example, if investments -- either in physical capital or R&D have the potential to change a firm's optimal employment mix, then it may make increased sense for a firm to make these investments if it anticipates wages will increase for its current employee mix. For this reason and more, we provide a sample of the rich relationships in the data in table 4.

One other factor that mediates the relationship between, wages, occupational mix and investment is adjustment costs. It stands to reason that if firms face fixed costs in hiring and firing workers, then they are more likely to change their occupational mix when growing or contracting. The extent of within firm rebalancing may also be related to growth through occupation-specific productivity changes at the firm level. Technological shocks of this sort will also change overall firm

productivity, leading the firm to expand or contract in employment as a function of the demand curve for the firms' products.

Figure 12 reports growth in occupations by routineness for firms sorted by their employment growth. The far left figure shows firms in the bottom quintile of total employment growth, and the far right shows firms in the top quintile. Attention is restricted to firms that operate in both 2000 and 2016, and each category of firms had an equal amount of employees in 2000. As can be seen, within firm rebalancing is more common for firms with shrinking employment than those with constant or increasing employment. Over this interval, the growing firms saw about 90 thousand employees moving between occupational routineness categories, constant employment firms saw about 101 thousand, and shrinking firms saw about 258 thousand. This result suggests that the fastest growing firms in our sample maintain relatively constant employment mixtures. This indicates that both causes of SBTC are important. Part of SBTC is driven by the fastest growing firms being routine biased, and part by the fastest contracting disproportionately laying off routine workers.

Table 1: Regressions of Firm Average Occupational Routineness on R&D: First Stage

	(1)	(2)	(3)	(4)
	Log(R&D Stock)	Log(R&D Stock)	Log(R&D Stock)	Log(R&D Stock)
Federal Cost of RD Invest	-9.361*** (0.000)	-1.621*** (0.000)	-1.290*** (0.000)	-1.198*** (0.000)
State Cost of RD Invest	-0.695*** (0.000)	-0.0550 (0.576)	-0.0699 (0.252)	0.00280 (0.924)
Log(Employment)			0.409*** (0.000)	0.148*** (0.000)
Log(Net PP&E)				0.120** (0.003)
Log(Market Value)				0.217*** (0.000)
Firm FE		X	X	X
Year FE	X	X	X	X
Industry FE	X	X	X	X
N	14262	14247	14236	14105

p-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Regressions of Firm Average Occupational Routineness on R&D: IV Results

	(1)	(2)	(3)	(4)
	Norm Firm Routineness	Norm Firm Routineness	Norm Firm Routineness	Norm Firm Routineness
Log(R&D Stock)	-0.0558*** (0.000)	0.378*** (0.000)	0.438*** (0.000)	0.355** (0.001)
Log(Employment)			-0.0525 (0.223)	0.104 (0.075)
Log(Net PP&E)				-0.0002*** (0.000)
Log(Market Value)				-0.102*** (0.000)
Firm FE		X	X	X
Year FE	X	X	X	X
Industry FE	X	X	X	X
N	14262	14247	14236	14105
F-Stat	180.4	37.36	22.08	32.43

p-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Regressions of Firm Average Occupational Routineness on R&D: IV Results

	(1)	(2)	(3)	(4)
	Norm Firm Routineness	Norm Firm Routineness	Norm Firm Routineness	Norm Firm Routineness
Log(R&D Stock)	0.531*** (0.000)	0.435*** (0.000)	0.385* (0.034)	0.493* (0.013)
Log(Employment)		0.0894 (0.085)		0.0887 (0.569)
Log(Net PP&E)		-0.0533** (0.003)		-0.138* (0.022)
Log(Market Value)		-0.122*** (0.000)		-0.122 (0.118)
Firm FE	X	X	X	X
Year FE	X	X	X	X
Manuf or non-Manuf?	Manuf	Manuf	non-Manuf	non-Manuf
N	11377	11247	2870	2858
F-Stat	10762.7	19976.9	3.819	3.362

p-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4: Firm Characteristic Relationships with Firm and Year FEs

	(1)	(2)	(3)	(4)	(5)	(6)
	Norm. Firm Returnness	Avg. firm wage	Log(Employment)	lnMarketValue/lnTotLab	Log(Net PP&E)	Log(R&D Stock)
Avg_firm_wage	0.00000291** (0.001)		0.000000000** (0.000)	0.0000220** (0.000)	-0.000040241 (0.128)	0.00001212** (0.001)
Log(Employment)	0.143** (0.000)	12903.8*** (0.001)		0.167*** (0.000)	0.496*** (0.000)	0.164** (0.000)
Log(Market Value)	0.0137 (0.257)	21471.8*** (0.000)	0.291*** (0.000)		0.439*** (0.000)	0.332*** (0.000)
Log(Net PP&E)	-0.0128 (0.249)	-2115.9 (0.073)	0.206*** (0.000)	0.243*** (0.000)		0.0649 (0.055)
Log(R&D Stock)	0.00756 (0.414)	10719.3*** (0.000)	0.0640** (0.005)	0.171*** (0.000)	0.0480** (0.000)	
Anticipated Wage Incr. 1yr	0.000000376 (0.125)	0.110*** (0.000)	0.000000187** (0.004)	-0.000000554*** (0.000)	-7.54e-108 (0.524)	-0.000000000** (0.001)
Anticipated Wage Incr. 2yr	0.000000011* (0.015)	0.256*** (0.000)	0.000000771*** (0.001)	-0.000000305*** (0.000)	0.000000011** (0.001)	-0.000000103* (0.046)
IT Consultant Share	-1.246* (0.014)	83154.9** (0.004)	-1.641 (0.055)	0.400 (0.078)	0.510 (0.060)	0.158 (0.735)
Norm. Firm Returnness		-23285.2*** (0.000)	0.101*** (0.000)	0.011 (0.233)	-0.0191 (0.355)	0.0321 (0.431)
Firm FE	X	X	X	X	X	X
Year FE	X	X	X	X	X	X
N	23493	23493	23493	23493	23493	23493

p-values in parentheses
* p < 0.05, ** p < 0.01, *** p < 0.001

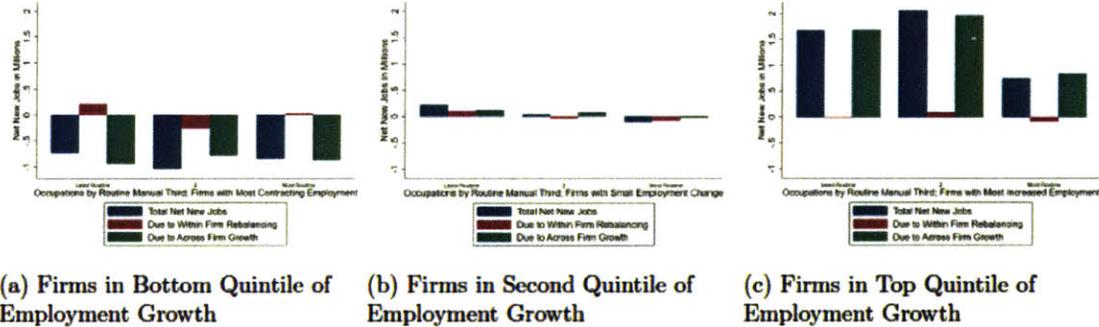


Figure 12: Change in employment by non-routineness of occupation in thirds and firm estimate wage bill growth in quintiles. 2000 through 2016. Net new job growth decomposed into sources. Only firms which do not enter or exit included.

12. Conclusion

In this paper we decompose, for the first time, aggregate skill-biased technological change into its sources within and across firms. There are several caveats to the conclusions in this study. First, much work remains in exploring alternative specifications and measures for the relationship between investment, growth and firm-level skill-biased employment changes. Second, our data source, while unique and powerful, is constructed using self-reported resumes. To deal with this, we introduce a

system of reweighing observations by firm, industry and occupation. What is clear from our analysis is that the shifting of occupational types within firms is an important component of SBTC. Future work will explore the mechanisms behind these shifts in greater detail.

We find that within-firm rebalancing is the second most important source, quantitatively, of non-routine biased employment growth. It is especially important for the increasing non-routineness of labor for firms in initially routine-intensive industries, such as manufacturing. Rebalancing within firms in these industries is the most important sources of skill-biased technological change. Within firm rebalancing explains almost all the decline in employment for the fastest contracting occupations. These observations run counter to theories of SBTC that emphasize firm production function neutral technological changes alone. Still, we find that growth across firms is still important in industries that are already relatively intensive in non-routine tasks. Net firm entry also plays a larger role for growing industries and industries that are initially non-routine intensive.

Contrasting the non-routine bias of within-firm rebalancing across firms with different investment and growth, we find a weak correlation between capital intensity and firms becoming more non-routine intensive. In our regression analysis, we find that firms which do more R&D as a result of a tax break contemporaneously increase their relative share of routine workers. All these trends are consistent with a theory of SBTC driven by technological changes which firms implement infra-marginally.

References

- Acemoglu, Daron, “Changes in unemployment and wage inequality: An alternative theory and some evidence,” *American Economic Review*, 1999, 89 (5), 1259–1278.
- , “Directed Technical Change,” *Review of Economic Studies*, 2002, 69 (4), 781–809.
- and David Autor, “Skills, tasks and technologies: Implications for employment and earnings,” in “Handbook of labor economics,” Vol. 4, Elsevier, 2011, pp. 1043–1171.
- and Pascual Restrepo, “The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment,” 2017.
- , David Autor, David Dorn, Gordon H Hanson, and Brendan Price, “Import competition and the great US employment sag of the 2000s,” *Journal of Labor Economics*, 2016, 34 (S1), S141– S198.
- Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, John Van Reenen et al., *The fall of the labor share and the rise of superstar firms*, National Bureau of Economic Research, 2017.
- Autor, David H. and David Dorn, “The growth of low-skill service job and the polarization of the US labor market,” *The American Economic Review*, 2013, 103 (5), 1553–1597.
- , Frank Levy, and Richard J. Murnane, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, Nov. 2003, 118 (4), 1279–1333.
- Barth, Erling, James C. Davis, Richard B. Freeman, and Andrew J. Wang, “The Effects of Scientists and Engineers on Productivity and Earnings at the Establishment Where They Work,” Working Paper 23484, National Bureau of Economic Research June 2017.
- Benzell, Seth G. and Erik Brynjolfsson, “Digital Abundance and Scarce Genius: Implications for Interest Rates and Wages,” Technical Report, MIT Mimeo 2018.
- Benzell, Seth G, Laurence J Kotlikoff, Guillermo Lagarda, and Jeffrey D Sachs, “Robots Are Us: Some Economics of Human Replacement,” *NBER Working Paper Series*, 2016, (Working Paper 20941).
- Berman, Eli, John Bound, and Stephen Machin, “Implications of Skill-Biased Technological Change: International Evidence,” *Source: The Quarterly Journal of Economics*, 1998, 113 (49), 1245–127961.
- , —, and Zvi Griliches, “Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers,” *Quarterly Journal of Economics*, 1994, (May).
- Bessen, James, “AI and Jobs: the role of demand,” in “Economics of Artificial Intelligence,” University of Chicago Press, 2018.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen, “Identifying technology spillovers and product market rivalry,” *Econometrica*, 2013, 81 (4), 1347–1393.
- BLS, “Labor Force Statistics from the Current Population Survey,” 2018.

- Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt, "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *The Quarterly Journal of Economics*, Feb. 2002, 117 (1), 339–376.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *mimeo*, 2018.
- ___, Prasanna Tambe, and Jacqueline Barrett, "People Are Our Greatest Asset: Measuring the Market Value of Human Capital," in "in" 2018.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, "The role of entrepreneurship in US job creation and economic dynamism," *Journal of Economic Perspectives*, 2014, 28 (3), 3–24.
- Goos, Maarten, Alan Manning, and Anna Salomons, "Explaining Job Polarization in Europe: The Roles of Technology, Globalization and Institutions," *CEP Discussion Papers*, 2010.
- ___, David Autor, Larry Katz, Thomas Lemieux, Stephen Machin, Guy Michaels, John Van Reenen, and Ulrich Zierahn, "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," 2014.
- ___ and ___, "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *The Review of Economics and Statistics*, 2007, 89, 118–133.
- Johnson, George E, "American Economic Association Changes in Earnings Inequality: The Role of Demand Shifts," *Source: The Journal of Economic Perspectives*, 1997, 11 (2), 41–54.
- Kaplan, Steven N. and Joshua Rauh, "It's the Market: The Broad-Based Rise in the Return to Top Talent," *Journal of Economic Perspectives*, 2013, 27 (3), 35–56.
- Lucking, Brian, "Do R&D Tax Credits Create Jobs?," 2019.
- ___, Nicholas Bloom, and John Van Reenen, "Have R&D Spillovers Changed?," 2018.
- Peretto, Pietro F and John J Seater, "Factor-eliminating technical change," *Journal of Monetary Economics*, 2013, 60 (4), 459–473.
- Peters, Ryan H and Lucian A Taylor, "Intangible capital and the investment-q relation," *Journal of Financial Economics*, 2017, 123 (2), 251–272.
- Rosen, Sherwin, "The Economics of Superstars," *The American Economic Review*, 1981, 71 (5), 845–858.
- Western, Bruce and Jake Rosenfeld, "Unions, norms, and the rise in US wage inequality," *American Sociological Review*, 2011, 76 (4), 513–537.

t!