Mapping the Transition of Work in Labor Markets and Entrepreneurial Organizations

by

Theodore Wiley DeWitt III

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in the University of Michigan
2019

Doctoral Committee:

Professor Gerald F. Davis, Chair
Assistant Professor Derek Harmon
Associate Professor Clifford Lampe
Emerita Professor Kathleen Sutcliffe
Professor Amy Wrzesniewski, Yale University

Theodore Dewitt

teddydew@umich.edu

ORCID iD: 0000-0002-4229-1454

DEDICATION

The dissertation is dedicated to my grandmother Martha Coe (1914–1997), my mother Coretta DeWitt, and my daughter Aurora DeWitt. All of the success that I have been able to have in this life is due to the strength and sacrifices of my mother and grandmother. And they taught me how and why I need to do the same for my daughter. The most important thing I want to say here is that, in all of their eyes I saw and I see a deep and probing intelligence. Maybe that is because they are all related. Maybe that is just because I love them all so much. But putting aside those maybes, there are a few things of which I am certain. I am certain that not enough of the world saw just how brilliant my grandmother was, in part because she never really had formal education. I am also certain that not enough people saw my mother's genius because if they had, she would be the one writing this dissertation. (And it would be better!) However, even if the world ignored the unmistakable spark of these black women, I did not. That spark made me everything that I am. And I will cultivate that indelible quality in my daughter, as I can see the strength of those generations past in her imaginative actions and thoughtful contemplations. She will bring such wonder to the world. And somewhere my grandmother's rich, sonorous laugh is echoing as she nods in effusive agreement.

ACKNOWLEDGMENTS

I don't think I will get this right, as there are more people that have helped me on this journey that I can remember in one sitting. So if you do find yourself reading this and you are not mentioned, you should email. And I will promptly tell you the 10 ways in which you rightly assumed that you helped me. And also the 100 ways you helped me that you didn't even know.

To handle the dissertation matters first, I would like to thank my committee: Jerry Davis, Derek Harmon, Cliff Lampe, Kathleen Sutcliffe, and Amy Wrzesniewski. If there is anything brilliant here it is because of them. But if there is anything that is dull and ham-fisted, that was all me. Next, I would like to thank my editors, Kate Epstein and Kelly Clody. Kate has a preternatural knack for conceptual clarity. (She also gets my jokes, which helps.) Kelly has eagle-eyes (and maybe even a third eye!) for technical detail. Together, they took my few decent ideas and scraped away the unnecessary bits and helped me to understand the things that I was attempting to convey. If I learned nothing else from this process, I learned how important it was for me to know these two people. And I am wishin' and hopin' that they will keep working with me in the future. Finally, I would like to think all of my friends that let me interview them for the case study chapter. I count myself very lucky to know such brilliant and accomplished people.

Brandon Bunt, Pat Olson, and the team at Trilogy deserve a paragraph all their own. They are really doing something truly unique and innovative. I truly loved experiencing the richness of their story. But more than anything, I truly appreciate their generosity with their time. They let me pester them with questions for the better part of two years and were always willing to answer

just one more question. I can't believe I was so fortunate to be in Michigan at the time Trilogy was starting out. And I don't know how I was lucky enough to encounter such kindness from Brandon and Pat.

During my six years in Ann Arbor, I experienced more generosity and kindness than I thought possible. First to all, of the PhD students from numerous cohorts including but not limited to: Cassandra Chambers, Ashley Hardin, Lyndon Garrett, Chen Zhang, Lindsey Cameron, Mrudula Nujella, Eun Woo Kim, Laura Sonday, Eun Bit Hwang, Mijeong Kwon, Kat Bae, Hilary Hendricks, Sarah Gordon, and Alaina Segura, your friendship and support has meant the world to me. Next to the faculty. Maxim, thank you for being a driving force and careful scholar in the development of my IERP. Sue, I will always appreciate your encouraging words and deep thoughts. Jim, thanks for always looking into my office and pulling me into yours. I am looking forward to many more two-hour conversations. Wayne, thanks for thinking my research could be a home run. I hope it will be! Dave, thank you for the walks and talks. And to everyone in MO, even though I am leaving, I hope that there are many more great experiences and memories to still be had and formed in the coming years.

To Kathie: I can't thank you enough for that first phone call back in 2013 to let me know that I was coming to Michigan and then for everything that came after. I could not have envisioned a better way to enter academia than under your expert guidance, scholarly brilliance, and indefatigable care.

To Jerry: I thank you for dealing with me for five years. No one has ever given me so many opportunities to pursue success. Thank you for the dry wit, kind words, and deep knowledge. Every day truly was a blessing, comrade.

To Amy: This dissertation and journey has been nearly 12 years in the making and it fittingly started in your Careers class. Thanks for seeing me through to the end of this adventure, even as you embarked on a new one of your own!

To Ali and Rory: Ali, thank you for coming on this epic adventure with me. And Rory, thank you for joining in medias res! I deeply appreciate your collective patience as Daddy vanished into his dissertation hole for hours and days at a time. I am looking forward to our coming adventures in Boston as a family.

And I left this for last because too often I postpone the hard things. Mom, you have been telling me recently that I am your hero. But you have always been mine. All the accomplishments I have, and successes that I experience, are because of everything that you have done for me. I don't know that I can ever properly show how much I appreciate the faith that you have placed in me over the past forty years. No one deserves as magical a person as you to be his mother. How did I get so lucky?

TABLE OF CONTENTS

LIST OF TABLES

TABLE

LIST OF FIGURES

FIGURE

ABSTRACT

Jobs and technology have a complicated history. Advances in technology can enable people to do existing jobs more efficiently, or technology can completely replace the human element of a job. While fear of the impact of job automation is not new, the scope of this concern has expanded due to novel technologies such as robotics, complex neural network architectures, and artificial intelligence. Though wholesale replacement of all human labor may be distant, technology is reconfiguring the jobs landscape, as it has in times past. Understanding how requires a more complete understanding of the structure of jobs and how that structure is impacted by technology than currently exists in the organizations literature, a subject this dissertation explores.

In this work, I advance a theory of jobs and technology that relies on conceptual analogy. This analogy allows me to use what we know about the structure of technology along its combination, recursive, and performative dimensions to infer a similar structure in jobs. I propose that just as a technology is a stack of technological components, a job is a stack of task components. Through this analogy I also propose that jobs and technology can coevolve. To examine the theoretical proposition of a job as a stack of tasks, I take task descriptions of jobs from the U.S. Bureau of Labor Statistics and decompose them into underlying task variables using natural language processing methods. This analysis allows me to represent the underlying theoretical concept of jobs as a stack of tasks that can be represented quantitatively, with jobs consisting of a group of latent task variables. Using paired T-tests and exploratory visualizations

of these variables, I demonstrate that the underlying task structure of jobs varies over time, that jobs are becoming more diverse in terms of their tasks, and that jobs cluster together in sometimes unexpected groups.

Next, to examine the theoretical position that jobs and technology may coevolve, and to understand how technology intersects the creation of new jobs, I performed a case study of a commercial cleaning services company whose workers use computer tablet-enabled, guided workflow software. I supplemented this case with interview data from recent entrepreneurs. Using thematic analysis, I uncovered three key findings across the data set. The first is that in contrast to existing literature on the pursuit of entrepreneurial opportunities, all entrepreneurs seem to experience both risk and uncertainty in the forms of market uncertainty, market risk, execution uncertainty, and execution risk. The second is that these aspects of risk and uncertainty make it very challenging to create fixed and formalized jobs until a company's product and strategy have crystallized. The third is that the technology stack that jobholders use, and the job task stack that jobholders perform, coevolve through a process of on-the-job experimentation.

This work makes three contributions to the literature. The first two are methodological: I demonstrate how natural language processing can help researchers engage in quantitative analysis to address questions past research has only addressed through qualitative methods. This research also provides a new method of examining the underlying task structure of jobs. Finally, this work shows that as new organizations create jobs, that initial task structure is provisional; jobs go through iterative processes of experimentation before they can be formalized.

CHAPTER 1

Introduction

Since the early 1980s, information and communication technology has aided in the completion of tasks people perform in their jobs, and technology's pervasiveness in the workforce has only increased since then. Advancing technology has long posed a threat to the employment prospects of human workers, and stories of the robots in Amazon's fulfillment centers and the self-driving cars of Alphabet's Waymo pepper the news. Yet the impact of the advancement of technology in labor may be more complicated than pundits who warn of massive unemployment allow.

A 2016 McKinsey report predicted that only about 5% of human jobs will be completely automated over the next 20 years, but portions of as much as 60% of occupations will be automated. A review of relevant literature suggests that the evolution of technology and its integration in the workplace also causes the kinds of jobs people perform to evolve. In light of this, I ask, how do we gain a better understanding of the future of work amidst technological change?

No one part of this story is new, per se. Concerns about the automation of work have been around since the dawn of the Industrial Revolution. Research has examined how novel information and computing technology affects work at least since the 1980s (e.g., Zuboff, 1988). And the literature around job and work design is deep and well conducted, with perspectives spanning over 40 years of organizational research (e.g., Grant et al., 2007; Hackman & Oldham, 1976; Parker, Van den Broeck, & Holman, 2017; Wrzesniewski & Dutton, 2001). However, most job design research focuses on the jobholder or how managers can tailor jobs to encourage worker productivity and satisfaction. The job itself has rarely been the object of research, and the researcher's indirect gaze generally views jobs as stable collections of tasks that one person performs within the confines of an organization. This leaves us with an incomplete understanding of how organizations actually bundle organizational tasks into jobs. In turn we are unequipped to predict how the changing technological landscape will reconfigure jobs, especially in new organizations. As emerging technology provides new avenues to organize work and structure jobs, nascent organizations will be at the forefront of job construction.

This dissertation examines the elements of job design from a task perspective and explores how the task components of jobs evolve over time at both a micro and a macro level. At the macro cross-economy level, I use computational methods and natural language processing to construct topic and embedding models that uncover the underlying structure – or task stack – of jobs and how it has changed in recent history, according to the U.S. Bureau of Labor Statistics. Two forms of qualitative data will reveal the micro-level process of job construction and how that intersects with novel technology. The first consists of interviews at a unique field site – a commercial services cleaning company whose workers use tablet computers and augmented reality goggles in their daily work practices. The second consists of interview data from current and former entrepreneurs. The analysis of their experiences will serve as a complementary perspective to what I observe in the field. Thus, entrepreneurial settings will provide a backdrop to my inquiry into how the tasks of jobs have changed and continue to change.

Including this introduction, this dissertation has five chapters. The second chapter presents the theoretical integration between the job and technology literature in organizations studies and asserts that although the literature around jobs and technology have indicated points where they may intersect and have cocrafting effects on one another, these effects – and jobs and technology, generally – can be better understood as part of an analogy. This analogy will review the key sites of technological incursion into the jobs spaces and utilize what we know of technology's structure to make inferences about the structure of jobs.

The third chapter is a computational chapter in which I explore how the underlying task stack of jobs has changed over the past 20 years. In this chapter, I use a combination of topic modeling (Blei, Ng, & Jordan, 2003), thematic analysis (Braun & Clarke, 2012), and paragraph embedding (Grant et al., 2007) to uncover the task structure of jobs. This chapter has four key takeaways. The first is, jobs can be described along their underlying task dimensions. Specifically, using computational methods I move from a qualitative text description of jobs tasks to a quantitative measurement of task variables. The second is that jobs are changing over time. At the individual job level, I demonstrate how each job is composed of a subset of task variables, and that the relative contribution of each task variable to a job changes over time. At the aggregate job level, I show that the average representation of the underlying task variables of jobs changed over the time period of the study. The third is that jobs are getting more diverse in terms of the underlying task structure. I demonstrate this by measuring the entropy of the task variables of which each job is composed. The fourth is that jobs cluster in groupings or neighborhood. And in these neighborhoods, the jobs share similarities in their underlying task structures. Because I can decompose each job task description into measured task variables, I can measure how similar jobs based on their respective task variable compositions these dimensions and render these into data visualizations. These visualizations confirm some of our intuitions about which jobs may be alike in their compositions. However, these visualizations also show some similarities between jobs that outwardly seem very different.

Chapter 4 is a single case study of Trilogy Corporate Services, a start-up that is bringing novel technology to a traditional industry. Start-ups can provide fertile ground for both testing and revising our fundamental organizational theories (Sorenson & Stuart, 2008). Start-ups and small businesses represent the leading edge of job creation in the United States (Haltiwanger, Jarmin, & Miranda, 2013). The particular benefit of this case study is that Trilogy is in the process of job assemblage (L. E. Cohen, 2013, 2016b), which I was able to watch unfold at both the frontline level and the managerial level, while also observing the ways in which the members of the organization navigate the complexity of technological experimentation while attempting to establish consistent work practices. Chapter 4 has three key takeaways. The first two are connected. First, start-ups seem to exhibit similar characteristics in that they all have to navigate market risk, market uncertainty, execution risk, and execution uncertainty in the execution of their entrepreneurial ideas. All of the companies in my sample exhibited these characteristics, which appears to stand in opposition to the entrepreneurial literature regarding the pursuit of entrepreneurial opportunities (Alvarez & Barney, 2010). Second, in navigating these varieties of risk and uncertainty, entrepreneurs find it difficult to create stable and fixed jobs for which to hire. As the case study reveals, prematurely trying to hire into fixed and formalized jobs, before the actual shape of the company's strategy and product offering has crystallized, can lead to problematic outcomes. The third finding is along the specific dimension of execution uncertainty. A company attempting to integrate novel technology with new jobs undertakes an iterative, experimental process that involves not only making changes to both the stack of technologies jobholders use and to the task stack of the job itself. Management can only see the necessity of these changes as the worker carries out the job in real time. Specifically in the case study, these task and technology revisions are made and undertaken in light of evolving technological mishaps as well as internal and external environmental difficulties.

The fifth chapter concludes the work by providing a summary of the entire study and addressing limitations in my theorizing and data choices while providing future directions. In the future directions section, I outline potential paths of development for future research that will mitigate the limitations of the current work.

On two points, I ask for the reader’s indulgence. First, the eventual intention of this work is the creation of three separate papers, one theoretical/conceptual (Chapter 2) and two empirical (Chapters 3 and 4). While they require further revision to achieve this goal, they include some duplicated material because of this intention. Second, this work is provisional. In it I am trying out new ideas, testing new methods, and becoming familiar with a broad array of scholarly literature. I appreciate your patience with my process.

CHAPTER II

Theoretical Lens

Introduction and Context

The glorification of jobs is one of the things that unifies American society. Americans and our leaders, often rightly, see jobs as necessary (and sometimes even sufficient) for economic and psychological well-being. Thus the impact of technology on jobs – through the emergence of new jobs, the transformation of existing jobs, and the relegation of declining jobs to the vicissitudes of economic history – brings both anxiety and joyful anticipation.

Organizations scholars have been broadly examining jobs for 40 years. But a focus on jobs as the primary objects of research, rather than an intermediate construct in a process, has garnered less intention (Parker et al., 2017). As a result, insufficient work has theorized the deep interplay between jobs and technology. For technology to affect jobs, there must be a key site of interaction. This work attempts to provide an understanding of what that key site might be by way of establishing a theoretical analogy between jobs and technology.

In the sections that follow, I will first provide key definitions of terms that will appear throughout this dissertation and provide historical background on key concepts. Then I will review research on jobs and technology, and articulate the specifics of the homological reasoning that connects jobs and technology as well as what gains this reasoning provides from a theoretical perspective. This introduction concludes with a brief example of an evolution of a job under the effects of technology and an explanation through this example of how jobs and technology as homologues have implications for additional research.

Key Terms

The two terms that underpin this study's discussion of the literature are *jobs* and *technology*. Many scholars in both organizations studies and labor economics have sought to define the first of these. Barley and Kunda (2001) note that before the Industrial Revolution, a job was "a discrete task of short duration with a distinct beginning and end," but that by the mid-20th century it became "an ongoing stream of activities attached to a role in a division of labor that was held for an indefinite period of time" (p. 82). Referring to the post–Industrial Revolution meaning that remains in use today, researchers define a job as a stable amalgamation of tasks assembled together and performed under an administrative title (L. E. Cohen, 2013; Miner, 1987). Goos and Manning (2007) state that the particularities of the occupation, often within a specific industry context, primarily define a job. All of these definitions essentially describe jobs as a series of tasks that workers complete.

Technology is a bit more difficult to nail down, as Arthur (2009) notes with apparent frustration in his treatise on technology. Organizations scholars Dewett and Jones (2001) describe technology as "the process of managing the uncertainty and risk surrounding the transactions necessary to convert inputs to outputs" (p. 315). More useful for the current work, Arthur describes technology as an assemblage of practices and components that harnesses the elemental properties of matter for the goal of fulfilling some human purpose.

Historical Background

Jobs

Human beings have always had to work, but Adam Smith's thick descriptions of pin makers in the mid-18th century provide a useful starting point for the contemporary understanding of the design of jobs (Smith & McCulloch, 1838). The gist of Smith's argument is that, given advances in machinery at the time, individuals could divide the tasks involved in making pins and thereby produce more pins than if each individual performed all of the steps to create pins. Dividing tasks into batches would increase cognitive focus, mitigate process interruption, and allow workers to build increased skill and efficiency along one dimension of the pin construction process.[1] While his vision may have fallen short of Henry Ford's assembly line in terms of increased efficiency, it nonetheless changed work profoundly. Frederick Taylor in the early 20th century codifies many of these practices into the notion of scientific management, a perspective in which work practices could be standardized and people almost could be thought of as interchangeable parts, an extension of the machines they operated. A current of thought flowing through this part of job design's history is the idea that managers have the discretion to, in a top-down fashion, impose structure on the workers in order to gain efficient and increased production.[2] The top-down structuring of work by managers remained the dominant theme in the management literature of the time. Beginning with the work of Hawthorne in the 1920s and 1930s, a shift began in thinking about the psychological well-being of workers. The human relation model (described very generally here) advocates that over and

---

[1] Simon (1962) also gets at this in the parable of two watchmakers – Hora and Tempus. They build watches of 1,000 components. Hora builds watches one piece at a time. Tempus builds watches by assembling 10 subcomponents of

[2] One might say this also increases the amount of excess labor capital that the capitalist in question can effectively slough off (Marx [1844], 1972), but that is a matter for another paper.

above the assignments of tasks, one of the primary goals of a manager was to construct work and work context such that it could meet certain unconscious needs (Perrow, 1986).The human relations model, in counterbalance to the grim efficiency of scientific management, is in some ways anticipates the 1970s work of Oldham and Hackman and the Job Characteristics Model. While the locus of control still rests on the assumption of the top-down discretion of management, the aperture has expanded to contemplate the ways in which managers can specifically assign and create work in such a way that the work can induce favorable psychological states in the employee that will improve worker performance. To that end, Oldham and Hackman's (1976) bedrock understanding of job design, Job Characteristics Theory (JCT), identifies five principles managers can use to design jobs that will motivate employees to do their best work: skill variety – the number of skills brought to bear in the work; task identity – allowing workers to complete an entire unit of work, end to end; task significance – the degree to which the job has a positive impact on people's lives; autonomy – the degree to which the job allows for independence or discretion; and feedback – the extent to which the job provides unambiguous information about performance (Hackman & Oldham, 1976; Oldham & Hackman, 2010). Research has shown that focusing on these attributes can in fact induce key psychological states in individuals that will result in meaningful outcomes.

Thus far I have reviewed research in which jobs are an intermediary toward something else, often productivity or worker satisfaction. In spite of recent research calls (e.g., L. E. Cohen, 2016a), many researchers continue to treat jobs in this way, rather than focusing on them as the research object. However, there are exceptions. For example, Miner (1990) examines the structure and the conditions that result in the creation of formal jobs within an organization. She also describes the process and creation of informal or idiosyncratic jobs. Idiosyncratic jobs can accrue through management observing the tasks that an employee does, and thereby intuiting his or her unique set of knowledge and skills, and creating a new role that allows the employee to apply those skills to new tasks. Alternatively, an employee may purposely add to their job portfolio new tasks that are being left undone or that are emergent in company processes. In a subsequent study, Miner (1991) examines how a job's novelty or the way in which it was developed affected the job's persistence over time.

Baron and Bielby (1986) and Strang and Baron (1990) also treat jobs as research objects, more or less. Both of these studies examine the forces resulting in the proliferation of job titles, noting that the specific qualities of the job title reflect the aspects of the institutional, environmental, and internal political struggles over the division of labor. Thus they focus more on context than on jobs specifically, but in problematizing job titles, the researchers had to acquire a deep familiarity with the tasks and compensatory aspects of jobs. Baron and his colleagues understood that a job and its construction have higher-level effects for the structuring of the organizational context itself, a theme that is explored in the evolution of structure and bureaucracy in high tech firms (e.g., Baron, Burton, & Hannan, 1996).

Technology and Work

The inhabitants of the British Isles who built Stonehenge used enclosure ditches combined with elaborate and very strong timber settings to aid in the lifting and setting of large stones removed from quarries with chisels (Michael, 1997). Fast-forward a few thousand years and 35-year-old Theodore W. DeWitt Jr.,[3] an instrument technician in Ohio, worked on a catalytic reactor, or "cat cracker" as he called it, in his hometown refinery. He held this job for the next 20 years. In addition to the job's ability to generate a solid middle-class income, it was

---

[3] That is, my father.

also fuel for intermittent technology lessons – how to turn crude oil into gasoline – for an inquisitive child.[4]

Stories of technology on the job, then and now, largely have the same undercurrent. Technology, when you get to its base, exploits some fundamental, elemental property of matter and harnesses it to a particular end (Arthur, 2009). But the story of technology and its intersection with jobs, as shown in the organizations literature, neither reflects nor often cites any theory of technology. But technology is heavily implicated in the evolution of jobs and work, particularly in the period the current study covers – from the 1980s, as computers and information technology (IT) rose, to the present. Technology in this period has created a fundamental transition in the labor history of the United States: a shift to a service economy. In short, technology changed, again, how work, works.

Attewell and Rule (1984) researched the effects of computing in organizations, including the intersection of technology, control, and information flow. They note that technological change can increase centralization and decentralization of managerial control. Improved communication technologies can allow information to flow more easily and quickly up the chain of command, which can centralize manager control by obviating the need for middle managers. Alternatively, the same systems can result in decentralization as higher-level management allows middle managers to make decisions that top management might have made directly, or because technology gives them better indirect monitoring capacity. Thus they suggest that technology, and the way it changes communication flow and decision-making, plays a crucial role in job

[4] In case you are wondering what that child learned, the "cracking" process, really simplistically, means making crude oil really hot. And then as it cools, various byproducts, such as gasoline, jet fuel, and fuel oil, condense out of it at specific temperatures. The difference between the market price of a barrel of crude oil and its constituent byproducts is often called the crack spread.

design. This viewpoint tracks with the idea common at the time that a crucial part of job design is the overall discretion of managers to assign tasks to workers in a top-down fashion.

Looking at the use of computers and other types of automating technology in the workforce, Zuboff (1988) sensed a shift in the construction of the social reality of the workforce of her time. In a prescient statement, Zuboff said that technology "fundamentally reorganizes the infrastructure of our material world" (1988, p. 5). She agrees with Attewell and Rule that technology upsets the traditional order of the workplace in that it changes the command structure in the workplace, but finds that, rather than just along bases of communications, IT removes restrictions on access to knowledge across the organization, a knock-on effect of this is that underlying connective tissues between employees change. They can connect and collaborate in new ways, with less managerial discretion. According to Zuboff, one of managers' key levers of control had been monopoly of the organizations' knowledge base. Monopolization of knowledge flows allowed managers to more easily impose a command-and-control architecture over workers, an architecture that technology starts to undermine. However, technology also threatens workers, because it both automates – translating information into action – and (in Zuboff's coinage) informates – makes visible processes and activities in the organization that have been previously opaque. The interaction of these forces can reshape the workforce perhaps faster than humans can adapt.

Huber (1990) builds on both of these perspectives, noting technology's ability to change information flow and change access to information, while he also addresses the potential for rapid evolution (Zuboff's work also addressed this third point about technology). Like Attewell and Rule, Huber zeroed in on communication flow and decision-making in creating a theory on the effects of IT. Through his resulting framework, he contemplates how the multidimensional configuration of technology's basic properties and characteristics would affect organizational processes. This framework makes three specific contributions to the literature that have relevance here. First, advanced information technologies lead to increased information accessibility and, by extension, changes in organizational design. Second, many theories of organization are grounded in a period before automation progressed to its current state or advanced information technologies were available, and as such, they may be ill-suited to the task of explaining our informated and automated world. Third, given the rapid evolution of technology, organizational researchers should continue to study its impact, recognizing that a) technology is an intervention or jolt in organizational life, b) technology is a means to enhance decision-making, and c) technology is a means of designing different organizational forms.

Orlikowski, Yates, Okamura, and Fujimoto (1995) use an approach to technological examination that resembles Huber's. In a qualitative study of a research and development project at a large Japanese manufacturing firm, they examine the impact of the introduction of an asynchronous computer conferencing technology, or what were later called *newsgroups*, which resemble comment threads on Reddit or Facebook groups. The system concentrated information and knowledge in one place, allowing it to be accessed widely, thereby improving efficiency and productivity in the firm. But the researchers also identified a meta-structuring process of technology-use mediation in which various actors, often technological managers, essentially helped to guide the way in which technology was structured for employee use. Researchers demonstrated that technology evolves socially around, and interacts with, daily practice.

Powell and Dent-Micallef (1997) approach the exploration of technology by asking whether or not the adoption of a new technology actually improves firm performance. At the time, the view that technology was an unlikely contributor to firm performance because anyone could effectively adopt IT systems was emerging (e.g., Clemons & Row, 1991) in contrast to the understanding among researchers that IT could tighten the connection between strategy and execution. To aid in resolving this debate, Powell and Dent-Micallef examined the adoption of IT in the retail industry, which they found did not in and of itself improve firm performance. However, they found that firms that already had superior human resources (e.g., open organization, great oral and written communication, a company culture stressing experimentation) would realize a substantial benefit through adopting superior IT.

A 2001 review of the organizations literature by Dewett and Jones on the effects of IT on organizational life suggested that understanding the benefits of technology depends on examining technology within the context of an organization's human resources. They point out that IT has information efficiencies that save employees time, save a company money, and provide information synergy benefits in which pooling of human resources produces performance gains. They note the distinctly human perspective with regards to information synergies when they state:

> Knowledge is often inextricably linked to human resources and the way individuals and groups interact, [and] the information synergies possible through IT will only be realized when the firm is able to move further and actually utilize knowledge in its optimal location with the organization. (Dewett & Jones, 2001, p. 326)

In other words, technology's value comes from being interwoven in human interaction. This conclusion reflects two problems in organizational studies at the time, however. The first is that this version of the story of technology gives too much credence to human agency as the dominant actor when considering jobs, rather than acknowledging that technology can have material agency and be deeply involved in the structuring of jobs irrespective of human intentions (Law, 1991). The second problem is a bit subtler. At the time, the field was trying to devise theories around technologies in organizations, but no real theory had emerged about technology as an object, not least because a consensus definition had not emerged. Dewett and Jones (2001) very vaguely refer to IT as “the process of managing the uncertainty and risk surrounding the transactions necessary to convert inputs to outputs” (p. 315). An apparent consequence of this inattention to a theory of technology was that organizations studies largely viewed technology as kind of a black box for converting labor into goods and services, with no attendant ideas about how the evolution of technology would impact the evolution of jobs and, by extension, organizations.

Recent Literature

While all dividing lines are arbitrary, the turn of the current century is coincidentally a useful demarcation line between the historical perspective and current scholarship on technology’s impact on business. The emergence of a bottom-up, relational perspective on job design (Wrzesniewski & Dutton, 2001), a full-throated theory of technology (Arthur, 2009), and a socially based perspective around the use and effects of technology in organizations (e.g., Orlikowski, 2007; Orlikowski & Scott, 2008) marked the years following the turn of the century. These developments highlighted the change in organizations, as they became both more socially focused and more technically based. The reality of the workplace did not reflect the view of the workers as being somewhat isolated within the organization, broadly disconnected from other organizations, and managed on an in-person, face-to-face basis. A new vision of the worker as altruistic, instant-messaging, and telecommuting arose (Oldham & Fried, 2016; Oldham & Hackman, 2010). Both the literature of jobs and of technology reflected the change.

Jobs

Recent research on jobs generally focuses its gaze in one of three ways: (1) on top-down influences (e.g., Parker et al., 2017); (2) on bottom-up influences (e.g., Grant & Parker, 2009; Wrzesniewski & Dutton, 2001); or (3) on the job object itself (e.g., L. E. Cohen, 2013). This review addresses some of the key studies in each of these streams of literature.

Research based on top-down influences draws on an old concept that managers have broad agency with regard to control of the organization's social context in which jobs are situated, as well as the authority to structure jobs outright. The Hawthorne factory studies in the early 20th century reflected management's effect in setting the organizational context, and human relations research essentially did not question this perspective of management as a top-down context setter. And this view of management continues well into the later days of the 20th century. For example, a longitudinal study found that one of the primary goals of effective management is creating an environment of trust and support, which will then engender individual initiative and cooperation among the workers (Ghoshal & Bartlett, 1994). Miner (1987, 1991) finds that management can formally set the specific boundaries around the portfolio of tasks that a worker does, even jobs acquire such tasks in an idiosyncratic fashion.

The job crafting perspective, while it recognizes the significant role of managerial control, asserts that jobholders have significant agency with regards to both their ability to structure their jobs and the social context of the workplace than previous theories have suggested. To turn to this more bottom-up perspective, Wrzesniewski and Dutton (2001) provide the key theoretical model for understanding how employees actively craft their work, acting directly on the task boundaries and the relational boundaries of the job. For example, workers can selectively attend to certain tasks while disregarding others, thereby changing their job and their cognitive sense and perception of their job.[5] Additionally, they can change the relationship boundaries of their job by changing their relationships through the interactions they have with others at work. Finally, individuals' orientation to their jobs can affect managers' decisions to shift the task and relational barriers of the job in the future. Several studies have validated Wrzesniewski and Dutton's theoretical model through empirical work (Leana, Appelbaum, & Shevchuk, 2009; Lyons, 2008). Studies not necessarily focused on how jobs get made addressed the importance of relationships; for example, the performance of workers who relate and connect to the end beneficiaries of performed work perform better (Grant et al., 2007). Likewise, a growing perspective has emerged that work is situated in social relationships that extend throughout the organization and beyond it (Grant & Parker, 2009). It has become an accepted principle in the field that jobholders as well as their managers determine the composition of jobs and that jobs have a relational architecture that affects workers' productivity and sense of meaning (Rosso, Dekas, & Wrzesniewski, 2010)

Other research that focuses on job crafting seeks to factor in a broader range of influences on job creation. Parker et al. (2017) provide a comprehensive review of work design and its influencers. They note that understanding where job design comes from requires understanding all of the factors that might play a role in expanding or limiting the scope of jobs. In addition to intra-organizational, managerial, and employee factors, they point out that international and national factors play a role. For example, as globalization has opened up market access to a global source of suppliers, a large international company in one country can exert influence on its suppliers in another country by demanding adherence to certain safety standards. This can have the effects of altering the jobs at these suppliers. Parker et al. (2017) largely address the

[5] Wrzesniewski's (1997) findings show that the cleaning staff in a hospital's cancer ward understand themselves as doing healthcare work, and indeed craft their approach to their jobs based on improving patient care. This anticipates Wrzesniewski and Dutton's theory, as it presents a clear example of shifting the cognitive boundaries of a job.

mechanism for this broad inter-organizational influence of job design through DiMaggio and Powell's (1983) institutional perspective, describing how coercive (supplier relationships), mimetic (emulating best practices), and normative (professionalized occupations) factors might exert pressure on job design within an organization.

Although her perspective only encompasses the direct ways in which institutional pressures can affect the creation of a new job in her examination of a DNA sequence operator, Cohen's (2013) study of the job as a research object provides the most developed examination of jobs and their assembly in the literature. She describes jobs as essentially a built structure of tasks and probes how these tasks come together. She suggests that, rather than reflecting the will of managers alone, managers, individuals, previous holders of the position, and extra-organizational constraints shape and determine the tasks a job contains. In her multisite qualitative study of the implementation of the then-new technology of DNA sequencers, she noted that university administrators, state insurance regulators, lab managers, as well as job incumbents all played a role in determining the composition of a DNA sequence operator job. Multiple mechanisms led to reconciliation between competing points of view. Beginning to get beyond the understanding of jobs as solely assemblies of tasks, Cohen notes that the boundaries of the job enclose much more than tasks. It extends to the experiences of various actors, bureaucracies, and regulators, all of which affect the contours of task assemblage into a job. In her subsequent work, Cohen (2016b) notes that one of the more unexplored but potentially most influential parts of job design are in fact actors in this inter-organizational space, but her own study continues to focus on actors who have the ability to sanction an organization's work design on institutional grounds of legitimacy, coercion through regulation, or lack of adherence to professional norms.

Cohen's attention to the influence of extra-organizational effects on the task structure of a job at the micro level as well as Parker et al.'s broader approach begin to explore how the economy influences jobs. In this vein, Autor, Levy, and Murnane (2003) examine how technological forces change the task structure of jobs. Specifically, they examine the ways in which computerization determines job tasks. Categorizing jobs in a two-by-two matrix with routine and nonroutine along one dimension, and manual and cognitive (analytical interactive) on the other, they demonstrate that computer technology can serve as a replacement for routine manual and routine cognitive tasks for which the contours of the tasks can be concretely specified.

TABLE I
PREDICTIONS OF TASK MODEL FOR THE IMPACT OF COMPUTERIZATION ON FOUR CATEGORIES OF WORKPLACE TASKS

	Routine tasks	Nonroutine tasks
	Analytic and interactive tasks	
Examples	• Record-keeping • Calculation • Repetitive customer service (e.g., bank teller)	• Forming/testing hypotheses • Medical diagnosis • Legal writing • Persuading/selling • Managing others
Computer impact	• Substantial substitution	• Strong complementarities
	Manual tasks	
Examples	• Picking or sorting • Repetitive assembly	• Janitorial services • Truck driving
Computer impact	• Substantial substitution	• Limited opportunities for substitution or complementarity

Figure II.1 Tasks Table From Autor Levy and Murnane (2003)

On the other hand, technology serves as a complement to nonroutine cognitive tasks. As the capital costs of computing decline, organizations recompose the task composition of jobs and structure more to be nonroutine cognitive tasks.

Other recent studies that focus on the job as a research object include Frey and Osborne, (2017), which introduces a different set of computational methods for examining the underlying task structure of jobs and the extent to which they can be automated. They conclude that 47% of U.S. employment is at some risk of automation in the near future. Alabdulkareem et al (2018) (2018) explore how economic inequality is connected to skills and tasks across multiple industries. Using relatively novel methods of community detection, they examine the tasks, knowledge, skills, and abilities associated with certain jobs. What they find is that essentially there are "neighborhoods" of jobs that are just as economically segregated as actual neighborhoods in America, and exposure to certain sets of tasks and having particular skill sets may persistently relegate people to lower wage jobs – that is, like neighborhoods, some are more accessible than others. Obtaining certain jobs, like neighborhoods, is expensive.

Overall the literature's examination of jobs and what determines their contours reflect a growing understanding of extra-organizational forces that can shape job design. So far scholars have not explicitly described these institutionally based forces empirically. For practical reasons, scholars typically either take a theoretical approach (L. E. Cohen, 2016b; Parker et al., 2017) or begin with the forces that may influence job crafting, which they test and explore. This leaves open the possibility that extra-organizational forces scholars have not yet tested might affect job design, if only we examine them.

Technology

This section addresses two challenges that arise in the organizations studies literature with regards to technology. The first is how to define technology. The second is ontological: What exactly is the nature of technology within an organization?

Arthur (2009) addresses the definitional part in his work *The Nature of Technology*. In this work, he laments scholars' failure to create a working definition of technology or to create an overall theory of how technology comes into being. He proposes that technology can be defined through purpose, assemblage, and collectivity. The last of these is the simplest: To account for the broad way in which the word is used, Arthur says that technology can be a collective in that it is the entire scope of devices and engineering practices that he defines as technology. His statement, then, that technology is always intended for some human purpose, applies to both individual technologies and the collective technology. He sometimes refers to technology as an executable, intended to carry out some specific set of tasks. He also points out that technology is frequently an assemblage of other technologies, practices, and components, essentially a modular stack of systems, assemblies, subsystems, or subassemblies placed in combination or interconnected to achieve some specific purpose. These systems may be interconnected and cross-linked within the technology at different levels.[6] The modularity aspect of subcomponents can be likened to the discretization of pin making into separate tasks. Technology, except at the most basic level, where basic chemical elements form molecules that comprise the most basic technologies, consists of other technologies. Finally, technology works by harnessing a physical principle in order to perform. For example, a semi-conductor harnesses the fact that silicon, when mixed with certain elements, will only allow current to flow in one direction. Arthur clarifies, then, that a business process is not a technology because it harnesses biological and behavioral processes rather than physical, elemental processes.

Arthur's framework is extraordinarily useful in that it unpacks technology where organizational theory has black boxed it and specifies some ways in which technological objects differ from other organizational objects. What Arthur did not address is, what happens when technology with its purposeful action is placed in a human context? And is it possible to actually treat technology and purpose as completely ontologically separate entities? For example, the

[6] The modularity aspect of subcomponents can be likened to the discretization of pin making into separate tasks.

purpose of a record player in the social context of the Bronx in New York in 1925 was to play recordings of musical instruments. But the purpose of a record player in the social context of the Bronx in New York in 1985, during the rise of rap and hip-hop music, was to be a musical instrument (Faulkner & Runde, 2012). To view the record player as an ontological entity separate from human context is to miss something about the purpose of technology. Because of this, science and technology studies have recognized the importance of the affordance of a technology. Affordances are action potentials – possibilities that emerge from human engagement with a technology (Faraj & Azad, 2012; Gibson, 2014; Olson & Olson, 2000). While a staple in science and technologies studies, affordances do not appear in the organizations literature until Zanmuto et al. (2007), which recognize an affordance perspective as a way to understand how an object might both favor or constrain a specific set of uses. The researchers understood affordance as a bridging concept that unifies IT systems and organization systems, and recognized that the boundary between material agency of technology and the behavioral agency of a human being is not as ontologically separate as we might like.

Orlikowski and Scott (2008) likewise address the question of the boundary between technology and humans. They identify two streams in the literature on technology in organizations studies. The first envisions technology as part of a world of actors and things that are mostly discrete separate ontological entities with mostly stable characteristics linked in unidirectional causal relationships. The second is grounded in interaction, where actors and things are still discrete things, but are part of a web of interactivity leading to interdependent systems. Their main contention is that it is not obvious that either of these perspectives are correct and that there may be no analytical boundary between actors, technology, and the organizations, a view actor-network theory shares (e.g., Latour, 1991).[7] Orlikowski (2015) argues that a preponderance of research falls into the first camp, which has resulted in a failure to understand the ties between technology and organizational life. This has led to a dearth of research on technologies and organizations. Faraj and Azad (2012) suggest that this misspecification has resulted in somewhat of a "ping-ponging" between the poles of complete technological determinism (where technology dictates where human go) and complete human voluntarism (where people dictate where technology goes). Orlikowski and Scott (2008) and Orlikowski (2015) advocate for viewing actors, technology, and organizations as a sociomaterial assemblage in which all parts constantly cocreate one another. Leonardi (2012) is of a similar mind but is slightly restrained from Orlikowski's full sociomaterial assemblage view. Instead he posits that humans have an intentionality, a desire to accomplish something. Technology has a materiality, the arrangement of physical forms in such a way that it endures across time and space. To return to the record player, part of its materiality is that the turntable spins in both 1925 and 1985. As the humans contemplate a way to use the materiality of the object, its functions are activated (or executed, as Arthur [2009] might say) for human use. The interactivity of this process can have an impact on intangible constructs such as jobs, roles, or social structures within systems such as organizations.

Despite their differences, the ontologies of Orlikowski and Leonardi have a unifying thread. This lies in an understanding that, while technology has a suite of specific physical functions that comes from harnessing the materiality of the technology, only human connection and interaction can unlock its specific action potential, and it is very challenging to know before the fact how humans will use technology.

[7] Actor-network theory goes a bit further in this to suggest that not only is there no analytic difference between objects and humans, but there is no a priori reason to give humans preference over objects in an analytical structure. While this point of view has merit, I think it may go a bit too far into organization theory for my purposes.

Technology, then, shares many characteristics with jobs. Some researchers will focus on them to determine how to manipulate them to improve worker and firm performance. They both consist of a subsystem of smaller parts. Finally, scholars recognize that both have relational components. Technology has a set of internal connections, which include connections between technological subsystems or the interdependence of tasks within a job, or a set of cross connections, which are relations between the holders of job, extra-organizational influence on job design, or whole technologies stacked in relation to one another. I will turn to these in short order, but first I want to demonstrate that there is a reason that thinking about jobs and technology is necessary in this particular historical moment.

Critical Occasion

Scholars have been calling on organizations scholars to revisit many of our theories due to the rise of technology at least since Huber (1990). He rightly noted that the technology of the late 1980s was vastly different from the technology in use when key theories in organizations studies were formed, and technology's advancement has continued since 1990. The differences emerge in three key respects: the bases of change, the internal organizational context, and the scale of technology and its potential effect on jobs.

Broadly speaking, organizational theory is grounded in the assumption that workers work in large, publicly traded corporations. Many still do. But the number of publicly traded organizations in the United States has declined sharply, halving in the 15 years between 1997 and 2012 such that there were only about 4,000 (Davis, 2016). It is unlikely the trend reversed in the ensuing 7 years, as there are also fewer IPOs. According to Davis, more companies went public in 1996 than the period spanning from 2008 to 2014. And the companies that do go public employ fewer people. In 2016, Davis estimated the combined workforces of Google, Facebook, Twitter, Dropbox, Zynga, Zillow, LinkedIn, Uber, and Square total 80,000 employees. That is fewer “than Blockbuster had in 2005 or the net number of new employees GM added in 1942 alone” (2016, p. 92).

The internal organizational context has changed. Oldham and Hackman (2010) conceived of people working in fairly static roles, in a fairly static organizational context, with little communication between organizations. Jobs now are in much greater flux than in past decades, with more people working in temporary teams and having greater communication beyond the organizational boundary. In addition, technologies such as Zoom and Slack that allow for distant and asynchronous communication and knowledge sharing had not developed when Oldham and Hackman theorized jobs. These technologies have drastically reduced the in-person, face-to-face communication and interaction that is necessary for the command-and-control organizational structure typical of the time when JCT was formulated (Oldham & Fried, 2016).[8] The flexibility that these new structures offer may affect notions of autonomy as well as how much feedback an employee can expect to receive on the job. Consequently, technology can undercut some of the fundamental assumptions on which the original formulation of JCT is grounded.

Lastly the pace and the scope of technological advancement have changed. Levy and Murnane (2005) contemplated that technology could not replace certain nonroutine cognitive jobs, such as transcriptionist and taxi driver. Their reasoning was that the specifics of these jobs could not be informated, or technologically codified, in such a way that the work could then be automated. But Zuboff not only coined the phrase “informated”; she also noted that the interwoven dynamics of the forces of informating and automating would change the world faster

[8] Companies such as Basecamp have structured their job and organizational structure around minimal face-to-face communications to a quite successful effect (Fried & Heinemeier-Hansson, 2010).

than we could imagine. Today, thanks to advances in computer processing power and the improvement of neural network techniques, there are fully autonomous vehicles (e.g., Bojarski et al., 2016) and fully automated transcription services (e.g., Graves et al.,2014; Manjoo, 2019).[9] Arthur states that "technology is steadily creating the dominant issues and upheavals of our time" (2009, p. 11), and part of this is due to the fact that technology constantly advances by combinatorial evolution – simple technologies can be recombined into ever more complex ones. As some consulting firms predict that 47% of jobs will be automated to some degree in short order (Chui, Manyika, & Miremadi, 2016), how should we try to understand the forthcoming reordering of the labor and organizational landscape?

Conceptual Integration

Types of Analogical Reasoning

One of the dominant modalities we use to advance theorizing in organizational theory is reasoning by analogy, where the researcher takes a model or set of notions from a well-understood source domain and applies it to a target domain that they wish to understand better (Cornelissen & Durand, 2014; Ketokivi, Mantere, & Cornelissen, 2017). Cornelissen and Durand (2014), outline three types of analogical reasoning. The first is a heuristic analogy in which the well-known source domain is used to catalyze thinking and then once new propositions are developed, or hypotheses are constructed for testing, the analogy is largely abandoned. In this first example, the domains don't have to be similar; the analogy simply serves as scaffolding to reach a new theoretical level. The second type is the causal analogy in which a causal template is abstracted from a source domain to a target domain. This is done on the grounds that there are conceptual counterparts between the two domains that behave similarly, or there appear to be

[9] I used Descript, the service Manjoo describes, for all of the transcription work for this dissertation.

some similar dynamics between the two domains. And as such you can use other processes in the source domain to hypothesize about processes in the target domain. The third is a constitutive analogy in which the theory in source domain is brought fully into the target domain such that it essentially serves as the theory and basis for drawing inferences in the target domain. Constitutive analogies differ from casual analogies in that the former imports nearly all of the terminology and knowledge from the source domain, and blends the two domains together in a coherent whole. Vaughan (2014) mainly hews to the causal analogy formulation. She notes that, at least in organizations studies, analogical reasoning can be grounded in the underlying commonalities of organizational objects. All manner of organizations, irrespective of size, share the same building blocks of group life such as hierarchy and division of labor. They also share common processes such as conflict and power. This makes it possible to compare them across cases and generate theory based on analogies. Shepard and Sutcliffe (2015) demonstrate heuristic analogizing in providing a classic perspective on analogical thinking in the form of anthropomorphization, the phenomenon by which humans imbue objects or processes with human characteristics as a starting place for understanding and eventually explanation. In all versions of analogical thinking, the resultant analogy can be the groundwork for further theorizing as we ask what research questions a particular analogy, if valid, should prompt us to explore (Ketokivi et al., 2017; Shepherd & Sutcliffe, 2011, 2015). This dissertation is going to primarily follow Vaughan's notion of a casual analogy, as it will make use of the fact that jobs and technology share similar structures and processes, and because we may potentially import other aspects from the source domain of technology to gain new understanding about jobs.

The Causal Analogy of Jobs and Technology

To ask if I can pursue the notion of whether we can treat a job as if it were technology, we have to ask if there are structural similarities or similar dynamics, or to borrow from Vaughn's version of casual analogy, whether the objects of interest stem from a common origin and shared common processes. Jobs and technology appear to stem from a common origin, to fulfill some specific human purpose. They also share parallel constructs at different levels of analysis. The following is the basis for my analogy of jobs and technology.

Jobs and technology have a combination dimension. Arthur (2009) delineates the combination principle, stating that technologies are essentially stacks of other technologies interrelated and grouped together. Jobs are also essentially stacks of tasks that are grouped together and interrelated, in this case under a common administrative title (Barley & Kunda, 2001; L. E. Cohen, 2013; Miner, 1987).

Jobs and technology have a recursive, self-similarity dimension. Technologies are recursive or self-similar down to the basic components of matter or fundamental natural properties[10] (Arthur, 2009). Jobs are composed of tasks, which are composed of subtasks. Much as chemical elements are the basic building blocks of technology, movements and cognitions of human beings are the basic building blocks of tasks.

Jobs and technology have a performance dimension. All technologies exist as a way to harness some physical property of nature toward executing a specific purpose (Arthur, 2009). Jobs exist to harness a behavioral or biological function of people, such as knowledge or abilities, to a specific purpose.

[10] In Arthur's formulation, a computer programming language is a technology in that it effectively reduces down to information represented in binary math, which is a fundamental property.

If this analogy is plausible, we can use the target domain of technology to infer some characteristics about jobs. Arthur (2009) notes that technology changes through modular substitution. A technology is usually not fixed; its underlying architecture evolves, as new subtechnologies are created and substituted for one another. Televisions still show pictures but light emitting diodes have replaced cathode ray tube architecture. By extension we can perhaps expect the subcomponents of jobs, in this case tasks, to be removed and replaced over time as new organizational tasks become available, or as old tasks are performed by technological artifacts instead of human actors (e.g., Autor et al., 2003). Arthur also notes that technology, over time, tends to move toward greater levels of complexity. This is both because of the first point, that improved subcomponents technologies are made available, but also because the progress of science depends on understanding new ways to capture and harness new phenomena. In this way, new technologies are sources of variations (Aldrich & Ruef, 2006) that are placed in combination with existing ones, resulting in something new. We might imagine jobs evolving in much the same way, with new tasks being introduced into jobs or organizations due to changes in the environment, and as a result the task set of a job has to be reconfigured to manage these environmental changes. As such, much in the ways that we observe technology becoming more complex over time, we might observe jobs becoming more complex over time. What however is somewhat unique about this particular analogy, is that the source domain and the target domain are thought to exist together in a socio-material assemblage (e.g., Leonardi, 2012; Orlikowski et al., 1995), in which each these two items cocreate each other. Therefore we might expect that jobs and technology are constantly changing and cocreating each other. The labor economics literature suggests that, quite specifically, technology can change the job at the specific level of job tasks (.e.g., Frey & Osborne, 2017), and it is not necessarily the entire technology that has an effect, but some subset of its performing capabilities that impacts those job tasks. To further this notion, I present a specific, historical example in which technology seems to be pushing out tasks from a job stack and bringing in new ones over time. As I will describe, the changes in the job at times lead to repurposing of technology to accomplish new ends.

The Evolution of Financial Technology and the Trader Stack

In one of the seminal studies of traders and financial markets, Baker (1984) notes that the social network of floor traders on an options exchange could have an impact on the movement of and volatility of pricing in options.[11] As such, patterns of social ties formed due to economic interactions and the open-outcry, face-to-face market of the floor had an influence on the price of options. When the dominant modality is one of personal interaction, traders will build relationships in the course of shouting to grab people's attention in order to execute a trade, and relational tasks of that sort were one of the primary aspects of the trader job. The Black-Scholes-Merton (BSM) formula (Black & Scholes, 1973; Merton, 1973), an empirical way of pricing options based on the underlying volatility of the stock or commodity on which the option was derived, provides an alternative measurement that is free of the influence of social ties. Traders initially declined to use BSM, which they found cumbersome, until the Chicago Board Options Exchange (CBOE) initiated a system of automatic quotations using BSM in the early 1980s (MacKenzie & Millo, 2003). As BSM became a dominant part of floor life on the CBOE, along with an electronic ordering system, traders no longer had to engage in personal interaction and relationship building to perform their jobs effectively.

As computing power increased, computers became more prevalent, and the Internet became dominant, trading switched from face-to-face interaction to calculations and computer-

[11] Options are a type of financial instrument that, if purchased, give you the right, but not the obligation, to buy (or sell) another financial instrument at a fixed time and a fixed price. For example, if you buy a MSFT June Call @ 200, you have the right to buy 100 shares of Microsoft stock at $200 by the third week of June.

mediated interrelation. Knorr Cetina and Bruegger (Knorr Cetina, 2012; Knorr Cetina & Bruegger, 2002) note that with the rise of computers, financial markets, to a certain degree, no longer exist in a physical place as much as they live in the interconnections between computer terminals. Thus the interconnectivity aspect of computer technology obviated the task of regular interactions with traders at other institutions, replacing them with valuation and order execution by computer. I worked as a trader from 2002 to 2004 and never in that time interacted with a trader outside of my own firm.[12] With the increase in computing power and an increase in computer usages in financial firms, more and more trading operations have transitioned the job of trader from a person who deals in financial instruments, to one who uses computer programming languages such as Python and OCaml to render financial instruments, things of computational abstraction, as fodder for neural networks that can process data faster. One consequence is that as computer scientists and technologists have moved into finance, more technology, not originally intended for financial computation, is repurposed for that use. For example graphics processing units (GPUs) were originally designed to render very complex graphics for gaming and graphic design. However, they do so by performing rapid linear algebra calculations. The financial industry has repurposed this technology from its designed use, and used it instead for deep learning[13] associated with generating new algorithms for order execution in financial markets. However, the constant changes in the job of trader have very real human consequences. A Yale alumnus of my acquaintance who worked for several years as a trader lost his job because the trader task structure almost consisted entirely of complex computational tasks for which he did not have the requisite knowledge or abilities ("Personal Communication,"

[12] In a funny coincidence, the son of the Merton who put the M in BSM interviewed me for my first job as a trader.
[13] Deep learning refers to creating neural networks with many nodes. The underlying math is mostly linear algebra. See Chapter 3's primer on natural language processing for more on neural networks.

personal communication, July 18, 2017).[14] The task structure of trader may well now have more in common with that of computer scientist.

Conclusion

A common term used in software circles is the idea of a "tech stack." This suggests that in order to accomplish a specific technological goal you need an assortment of technologies placed in relation to one another. This aligns with Arthur (2009) but at a higher analytical level. A famous example of a tech stack is the LAMP stack, which stands for Linux (an operating system), Apache (a web server application), MySQL (a database application), and PHP (a scripting language to deliver HTML), and it is used to provide webpages and web applications. Depending on which financial firm you join now, it may well have a technology stack for its trading operations. The tech stack at Jane Street, a proprietary trading firm, includes OCaml (a functional programming language), Merlin (an integrated development environment for software creation), and Iron (a proprietary code repository; "Programming: Jane Street," n.d.). The benefit of establishing this analogy between jobs and established technology is that it calls attention to how jobs and technology share the same fundamental structure. A job is actually a "task stack," a set of tasks put in relation to each other to accomplish a specific end. As new technologies become available to provide a different set of properties for the end user, the technology stacks change. The task stack, then, changes along with it in a process of coevolution.

[14] "None of the banks have offered me an interview during my search over the past year, and though I had more success with hedge funds, it feels like that ship has sailed due to my lack of programming skills (I'm working on it) and my own mistake of not accomplishing enough during this period. Even financial markets have me at arm's length. I've looked into roles as a trader analyst. There has been a long trend of downsizing trading desks, cutting fees, automating, and moving to algorithmic/quantitative trading and away from equity derivatives, my strong suit." (Personal Communication, 2017)

However, this analogy has other implications as well. The composite subassemblies of technologies and the composite tasks of jobs are modular. They can be removed and upgraded or removed and replaced. Additionally, modules can be recombined, and we can see a combinatorial evolutionary process.

The rest of this dissertation will examine both the idea of task stack evolution and the coevolution between a task stack and a technology stack. In the next chapter, I posit that if we had access to myriad job descriptions over many years, it should be possible to see just how the underlying task stack of jobs changes over time. In a later chapter, I will examine the creation of jobs within a start-up and see how the technology stack the company is creating and the task stack of its workers coevolve.

CHAPTER III

Computational Big Data Chapter

Introduction

"What is a job?" is a question that we rarely ask. But having a grounded definition of a job is a precursor for understanding how they might change over time, including in relation to technological change. This chapter explores the key underlying structure of jobs – job tasks – and examines how these tasks have changed and evolved over the past 20 years. Using job descriptions from the U.S. Bureau of Labor Statistics paired with computational methods such as topic modeling and embedding methods, I uncover a latent task structure in jobs across the entire U.S. economy and then use the output from the models to track how jobs have changed and how they have become more or less similar over time. To do so, I establish computational methods as a way to uncover latent structures in temporal, cross-sectional data in organizations studies.

Context

The notion of a job is a common and a familiar one. We can hear the excitement in the youthful voices of the first-time employment seekers as they exclaim, "I got a job!" We can hear the earnestness, both real and affected, in the voices of our political leaders as they speak about "creating good jobs for Americans" (Bush, 1993). I myself have experienced firsthand and secondhand the despair when someone has to say, "I lost my job." Through our own experience, we each know of the economic and psychological reliance that people draw from a job. In addition, organizational scholars have done amazing work in helping us to understand just how that psychological reliance is drawn (Wrzesniewski, Dutton, & Debebe, 2003) and the ways in which certain kinds of jobs and specifically their skill content have economic implications for their holders (e.g., Autor et al., 2003). This chapter seeks to shine further light on a refrain from the press that intersects at all of the above points: the world of work is dramatically changing (Shell, 2018), and a new world of jobs is emerging in our economy (Brynjolfsson & McAfee, 2014; McAfee & Brynjolfsson, 2017)

Most people believe that technology is the main source of change. Jobs that are part of the so-called gig economy – those done on a contract basis or of short-term duration – are technology enabled in that they rely on the ubiquity of smart phones and the concurrent improvements in mobile technology. Because of the market platforms that technology has wrought, you can hire someone to pick up your groceries (Postmates), hail someone to take you to the airport (Uber, Lyft), or contract someone to spend 15 minutes debugging issues in your JavaScript webapp (Uplift). Technology seems to have changed the world of work by creating new avenues of work. The increasing economic impact of the gig economy, and the ways in which it is changing the organization of work in the American economy, demonstrates that it grows ever more important to gain a better understanding of jobs and technology. Some current estimates place gig and temporary jobs at between 10% and 15% of the U.S. workforce (Katz & Krueger, 2016, 2019). Other estimates suggests that over 20% of the workforce is engaged in nontraditional work arrangements (Petriglieri, Ashford, & Wrzesniewski, 2019) with the majority of net employment growth between 2005 and 2015 happening through alternative work arrangements (Katz & Krueger, 2016).

The specter of automation changing or eliminating jobs likewise gets significant attention in the press. For example, the podcast Planet Money (2015) reported that Waymo's self-driving trucks are coming for truck drivers' jobs, and that truck driving is one of the most common jobs in 25 states in the union. Nor are the most prestigious jobs in the country safe; the increased use of algorithmic retirement investing at companies like Betterment and Wealthfront reduces the need for human financial advisors, and surgical practices have begun to integrate robotic tools (Beane, 2018). Consultants at McKinsey predict that 45% of the activities that people are currently paid to perform could be automated, with 60% of all occupations potentially seeing a full third of their tasks automated, and about 5% of the workforce vulnerable to total automation (Chui et al., 2016). Yet it is still vital to define what a job is in order to go more deeply into understanding what it means to suggest that technology will change almost every job in the country.

Theory and Research Questions

The past 20 years have seen some of the most dramatic changes in the background conditions for human jobs. New kinds of jobs have appeared, and jobs that had been around for ages have vanished. Other jobs have gone through a process of transformation. Think, for example, of a plumber. American public restrooms have changed significantly in recent decades. Many have automatic faucets, automatic flush toilets, and toilets with dual settings for different kinds of waste. Some public restrooms have composting features. In addition to this new technology, plumbers face a change in that the variance around the average public restroom is much higher than it used to be. Plumbers need expertise in a far greater range of toilets, pipes, and fixtures. Nor are they alone in experiencing change due to technology. Not only have financial traders become at least part-time computer scientists; campaign managers are now part-time speechwriters, part-time data scientists, part-time behavioral scientists, and part-time marketers.[15]

As I have so far discussed some of the background conditions for a job, it makes sense to address just how "job" is defined. Barley and Kunda note that before the Industrial Revolution a job was "a discrete task of short duration with a distinct beginning and end," but that by the mid-20th century it became "an ongoing stream of activities attached to a role in a division of labor that was held for an indefinite period of time" (Barley & Kunda, 2001, p. 82). More recently scholars have called it a stable amalgamation of tasks assembled together and performed under an administrative title (L. E. Cohen, 2013; Miner, 1987). Goos and Manning (2007) state a job is "a particular occupation or as a particular occupation in an industry." (p. 120) All of these definitions describe jobs, broadly, as series of tasks to be done in order to achieve something specific.

Cohen (2013) examines and specifies just what goes into this notion of job assemblage. In her view, jobs are an assemblage of ideas and tasks that stem from managers, individuals, previous holders of the position, and extra-organizational constraints. In her multisite qualitative study of the implementation of DNA sequencers, she noted the varied sources of ideas about the composition of a DNA sequence operator job: university administrators, state insurance regulators, lab managers, and job incumbents. Various tasks were integrated into jobs through multiple mechanisms, including negotiations between these parties.

The definitions of jobs that emerge at least theoretically square with part of the analogy laid out in the previous chapter regarding the hierarchical self-similarity of and the performing

[15] A former Obama for America campaigner/data scientist and now entrepreneur explained this to me in an interview for this research.

aspects of jobs. The idea of task assemblage is comparable to every job having a task stack. That is, the job comprises interrelated tasks. Given this view of the structure of a job one might plausibly say that when a job changes, the underlying tasks change. The underlying logic of the analogy – how jobs and technology might interact but also how jobs might interact with other jobs – can help us understand how such changes occur.

Jobs and Technology

As Chapter 2 described, jobs and technology link up, with distinct potential implications. Labor economics has examined this in several ways. At a high level, transaction cost economics (Coase, 1946; Williamson, 1975, 1981) suggests that for various reasons, including the inability to direct behavior perfectly through contracts, certain employment relations need to exist within in the bounds of a firm. Technology can generate market platforms that can match buyers and sellers for discrete, specific, and time-delimited tasks more effectively than in years past. The portability and the ubiquity of technologies such as smartphones and the concurrent development of apps has led to the development of platforms that facilitate contracts for very specific tasks. It is possible for more and more tasks to exist outside the confines of a firm or without a firm as an organizational intermediary, and essentially, parts of a jobs task stack can be displaced. Postmates eliminates the need for grocery clerks who deliver, and DoorDash and UberEATS eliminate the need for restaurants to have delivery staff as employees within their organization.

Autor, Levy, and Murnane (2003) look at the potential for technological incursion into the task stack of jobs by examining the ways in which computerization affects certain tasks within work. In their research, they categorize job tasks into a two-by-two matrix with routine and nonroutine along one dimension, and manual and cognitive (analytical interactive) on the other. They essentially demonstrate that computer technology can serve as a replacement for routine manual and routine cognitive tasks for which the contours of the tasks can be concretely specified. On the other hand, technology serves as a complement to nonroutine, cognitive tasks. As the capital costs of computing decline, organizations recompose the task composition of tasks to lean more toward nonroutine cognitive tasks, because it becomes more cost effective to unpack or even eliminate the jobs that have a task stack that a technology stack can perform instead.

Frey and Osborne (2017) assert that, while the general framework Autor, Levy, and Murname introduced was correct, the conditions have changed. Machine learning and artificial intelligence provide a dramatically more performance-capable technology stack, which has the potential to push deeply into the nonroutine cognitive domain. For example, Levy and Murnane (2005) were confident that, given the cognitive complexity of driving, it would never be automated. They turned out to be wrong.[16] Frey and Osborne conclude that Autor, Levy, and Murnane significantly underestimated the extent to which technology will unstack the task stack of jobs. In addition to benefiting from the passage of time, Frey and Osborne also use more rigorous methodology than Autor and colleagues: a combination of interviews with domain experts on automation as well as a host of variables from the U.S. Department of Labor data on jobs (e.g., finger dexterity, originality, social perception) from which they create a training dataset of 70 jobs on the likelihood of automatability. They use this training set to create an algorithm that assigns probabilities of automation to the remaining set of approximately 700 jobs.

The view of jobs as having a task structure is fairly consistent in the labor economics literature, especially when discussing the threat of automation (e.g., Alabdulkareem et al., 2018;

[16] In fact, as I write this, a fully autonomous shuttle is cruising the roads outside of my home on the University of Michigan's North Campus.

Goos & Manning, 2007), even though they do not explicitly recognize that the precise site of technological incursion is the task stack. Both Frey and Osborne and Autor et al. improve our understanding of the likely economic impacts of automation. However, they share a common weakness. One instance appears to be due to the technology available at the time, while the other appears to be due to methodological choice. Both studies rely on human analysts to read task descriptions of jobs to quantify and judge the task structure of the jobs, and to then compress that data into limited and predefined categories. The current study attempts to address these weaknesses through a more fully inductive process.[17]

Jobs and More Jobs

The connection between two jobs may be a site of change. For example, two jobs might have a connection at an input-output interface. One of the implicit and explicit parts of Fordism and to an extent scientific management (Taylor, 1914) is the reduction of the job holder to a kind of machine that mechanistically takes in inputs and renders them outputs. These lines of inputs and outputs connect jobs to one another. While this includes the line chef's output of custard to the pastry chef's éclair, for example, it also includes a financial modeler's output of a sensitivity analysis that is inputted to the final slide deck for the engagement manager's client presentation. But another relationship between jobs that is less often contemplated is the fact that the person moving between jobs can connect them. That person may have the effect of shuttling tasks, routines, or knowledge from the old job to the new job and changing it in the process. The organizations literature has shown that either job holder can restructure the tasks within the job task stack and serve as a vehicle to bridge tasks between stacks. Here the perspectives from job crafting and organizational routines are illuminating.

[17] Even with algorithmic methods and machine learning, human bias can only be mitigated, never eliminated.

Wrzesniewski and Dutton (2001) envisioned employees as active crafters of their work. In this agentic approach to changing the task stack of jobs, job crafters can act upon the task boundaries and the relational boundaries of the job. A job holder can choose to do (or neglect) certain tasks within the scope of the job, which changes their job and, in so doing, their cognitive perception of the job. Additionally, they can change a job's relationship boundaries by changing both their relationships and the interactions that they have with others at work. The particular organizational context in which they find themselves affects the degree to which people can job craft, and these unique circumstances affect the method of crafting if they choose to craft at all. An individual's orientation to his or her job can affect the decision to shift the task and relational barriers of work. While the original piece by Wrzeniewski and Dutton is theoretical, empirical work has substantiated the model (e.g., Leana et al., 2009; Lyons, 2008). One of the most striking assertions that Wrzesniewski makes is that employees are every day altering the tasks that they do and thereby recreating their jobs on a daily basis. This is an agentic framework. As people move from job to job, their job crafting activities may well introduce purposeful changes to the jobs they inhabit, which may result in a drift in job tasks over time.

The philosophy of the job crafting perspective shares the underlying thread of enactment with the literature on organizational routines. Organizational routines are also composed of tasks, and sound similar to portions of the analogy of jobs in that they are "an executable capability for repeated performance in some context that has been learned in response to selective pressures" (M. D. Cohen et al., 1996, p. 683). But organizational routines, though composed of tasks, are not jobs. If jobs are the vertical dimension, organizational routines cut across several jobs horizontally in the performance of an organizational objective. Notably, they involve a repeated performance in response to pressure. Pentland and Rueter (1994) also reference pressure: "An organizational routine is not a single pattern but rather, a set of possible patterns – enabled and constrained by a variety of organizational, social, physical, and cognitive structures from which organizational members enact different performances" (p. 491). Here, too, we see affirmation that organizational routines are enacted as the tasks underlying them are performed over time. Yet as the organizational, social, physical, and cognitive structure changes, so do both the kind and menu of tasks that are enacted. Empirically it has been demonstrated that when personnel move between settings they can bring tacit knowledge with them that would be hard to convey though other means (Kane, Argote, & Levine, 2005). The key idea throughout this is that the underlying tasks of organizational routines are enacted (Feldman, 2000; Weick, 1979), but the very specific cognitive and social qualities of the individual holding the job constrains or enables the enactment of these routines and their underlying tasks. Based on this, in the job-job part of the analogy, human beings serve as a conduit between jobs, across time and organizations, through actively crafting jobs and/or passively changing them through their specific cognitive and social abilities.

Taken together, there are multiple avenues by which technology or the movement of human resources could affect the task stack of a given job. Thus, in spite of the fact that research has often treated jobs as quasi-stable compositions of tasks that change infrequently if at all (e.g., Ilgen & Hollenbeck, 1991), my research question for this study is as follows:

RQ: How has the job task stack of the United States changed over the past 20 years?

But before I examine how jobs have changed, I consider how reliably tasks, the subcomponents of jobs, have been measured.

Jobs Analysis

The idea of jobs analysis has been a substantial part of industrial-organizational psychology since the 1960s. Dunnette, one of the early pioneers of jobs analysis, in fact defined a job in much the same way we have defined it here: “clusters of tasks carried out for some essential purpose” (Dunnette, 1966, p. 69). This notion of a task proved to be important for jobs analysis for two reasons. The first is that because industrial-organizational psychologists posited that once you had the essential tasks of a job defined, then it would be possible to infer the knowledge, skills, and abilities – the behavioral and biological qualities – needed to complete those tasks. For example, McCormick described knowledge, skills, and abilities as “enduring qualities and characteristics of people that a manager infers will bring them success on a job” (McCormick, 1976, p. 654). The second reason that the task proved to be crucial to analyzing a task is that a task is thought to be reliably observable. In the 1990 updated version of the *Handbook of Industrial and Organizational Psychology*, the definition of a job was amended to include the following, “Observable behaviors employed by workers as well as what technologies are employed” (R. Harvey, 1990, p. 74). A meta-analysis of the average levels of reliability that could be expected from jobs analysis data affirmed that tasks are, reliably, the key way of understanding jobs, because intra-rater and inter-rater reliability of the content of jobs was highest when job tasks were the focus of job analysis (Dierdorff & Wilson, 2003). Taken in sum, these points, combined with those already in evidence, provide some confidence that tasks are the key to understanding the structure of jobs and that tasks can generally be measured to understand the content of jobs.

The quantitative and empirical work on jobs analysis in the industrial and organizational psychology literature is quite strong and empirically robust. Further, the empirical work in labor economics looking at the changes in the task stack of work due to technology is equally impressive. However, in contrast to these works, I am simultaneously making both structural and temporal claims that because of various factors the structure of all jobs within the U.S. economy exhibit changes. I am suggesting that because of these dynamics what may well be a latent structure can be uncovered. Another way of saying this is that at the level of the U.S. economy, jobs can be better understood if we take job task descriptions and extrapolate to latent level constructs that capture the essence of the tasks. Doing this allows us to see how jobs change and evolve and find potentially unanticipated connections between jobs. Because I cannot a priori say what that overall structure should look like, I need a set of methods that allows for the discovery of latent structure in data.

Data and Methods

Introduction

One of the theoretical assertions of this work is that each job across the economy has an underlying task stack and that aspects of this task stack can be assessed through analytical methods. One of the assumptions accompanying this is that, before the fact, I don't know and can't make any sort of claims about what that structure would look like. I have, so far, presented a theoretical argument indicating that in every organization there should be sufficient mechanisms at work to generate an overall aggregate, dynamic structure in the task stack of jobs. However, in order to substantiate that this structure exists, I need both an appropriate dataset and analytical methods that correspond to my research question, and that are capable of revealing the underlying dynamics of the phenomena I explore.

From a data standpoint, looking at a comprehensive set of job task descriptions across the economy should be sufficient to address the research questions. Given inter- and intra-rater reliability of analyses of jobs through their tasks (Dierdorff & Wilson, 2003; McCormick, 1976), it is reasonable to assert that descriptions of jobs encode cognitive content about the tasks to be done, that cognitive content can be accessed through a reading of the words (Duriau, Reger, & Pfarrer, 2007), and the changes in the usage of words or changes in the vocabulary say something about the changing ideas behind what those jobs are.

To examine the information content of task descriptions, I will explain some of the potential methods that could be employed, if not all of the linear algebra and multivariable calculus that would be required to understand their intricacies, improve upon them, and actively utilize them in research.[18] Rather, I will focus on the concepts behind the methods, which is sufficient to explain how I implemented them, why they are appropriate methodological choices, and the potential limitations to their applications.

With that said, first I will present the data I used followed by an explanation of the available analytical methods.

O*NET

The Occupational Information Network, also known as O*NET, was developed through sponsorship by the U.S. Department of Labor. O*NET's flagship piece of knowledge production is its database. This publicly available tool is released annually, and some updates to the database happen on a biannual basis. It contains occupational descriptions of approximately 1,000 occupations that cover a very broad swath of the U.S. economy. O*NET works under the assumption that every occupation and its tasks require the use of a different combination of knowledge, skills, and abilities. In aggregate, these items are referred to as the O*NET Content Model. The information about these occupations is collected from job incumbents and

[18] Latent semantic analysis relies on linear algebra. Latent Dirichlet Allocation and Embedding Models both rely on some variety of gradient. Neither of these have closed-form solutions, so to "solve" them, researchers rely on minimizing (or maximizing, depending on your viewpoint) some objective function.

occupations experts as part of O*NET's Data Collection Program (O*NET Resource Center Content Model, 2018). Occupational experts identify a statistically random sample of businesses that employ workers in the occupations about which O*NET collects data. Then a random sample of job incumbents within those businesses are contacted to fill out standardized questionnaires, with a portion of the questionnaire customized for the specific occupation of the incumbent

The O*NET project rests on a long history; observation of work and the understanding of jobs has been the bailiwick of the Department of Labor for more than 70 years. In fact, since 1939 the department has conducted surveys to ask domain experts and labor analysts a seemingly simple question: What jobs do people have and what do they do in them? The output of this effort was initially called the Dictionary of Occupational Titles (DOT). By 1991 there were four editions of the DOT, when the O*NET project superseded it. The fourth edition evaluated more than 12,000 occupations along several objective and subjective dimensions. The DOT was more comprehensive, but the O*NET data set goes into far more depth with a smaller collection of occupations by subsuming some occupations listed separately in the DOT into related occupations. Nonetheless, researchers have used both DOT, fourth edition (e.g., Autor et al., 2003) and various years of the O*NET data set (e.g., Frey & Osborne, 2017) to study the effects of the computerization of work. My analysis will follow in these footsteps, using the O*NET datasets released between 2000 and 2018.

Task Measurement

As the task data from O*NET is the main data source used for the analysis, a bit of explanation of that data is set forth here. As mentioned before, the questionnaire that O*NET uses to gather information about occupations economy-wide is standardized, but as certain aspects of work are particular to a given occupation, O*NET has an occupation-specific section that collects task information about that occupation. This occupation-specific questionnaire contains occupational tasks that are assembled by occupational analysts within the O*NET Data Collection effort.[19] The data sheet also leaves room for the respondent to write in tasks associated with the job. Figures III.1 and III.2 present some views of a sample of the O*NET questionnaire for "Nannies."

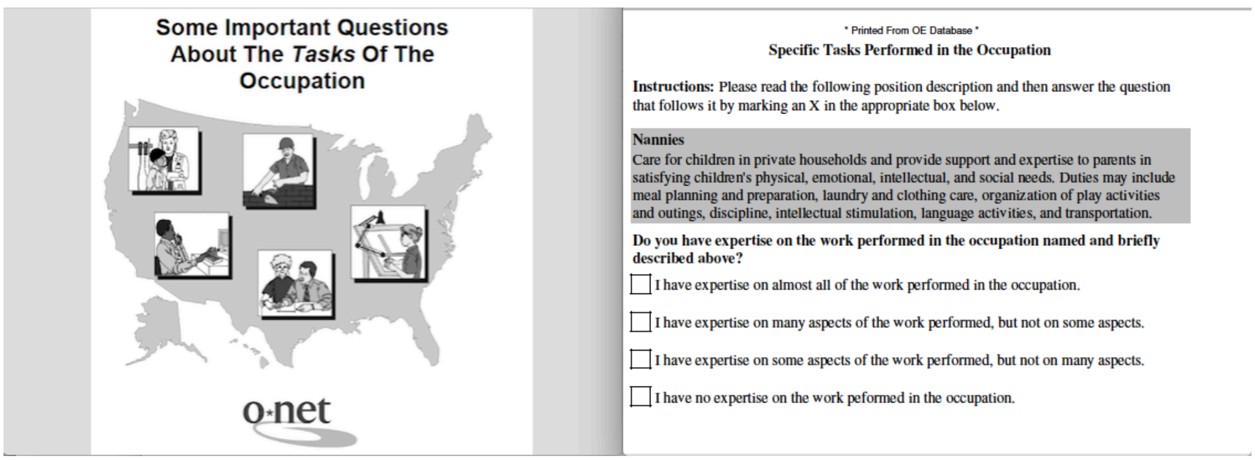

* Printed From OE Database *

Specific Tasks Performed in the Occupation

Instructions: Please read the following position description and then answer the question that follows it by marking an X in the appropriate box below.

Nannies
Care for children in private households and provide support and expertise to parents in satisfying children's physical, emotional, intellectual, and social needs. Duties may include meal planning and preparation, laundry and clothing care, organization of play activities and outings, discipline, intellectual stimulation, language activities, and transportation.

Do you have expertise on the work performed in the occupation named and briefly described above?

☐ I have expertise on almost all of the work performed in the occupation.

☐ I have expertise on many aspects of the work performed, but not on some aspects.

☐ I have expertise on some aspects of the work performed, but not on many aspects.

☐ I have no expertise on the work peformed in the occupation.

Figure III.1 Cover of Survey and Introductory Page With Occupational Description (*Questionnaires: OE - Tasks, 2018 OMB Clearance Package*, 2018)

[19] The Department of Labor actually contracts the data collection out to a nonprofit called the Research Triangle Institute. Sample questionnaires for all occupations as well as other information about the Institutes involvement with O*NET can be found at https://onet.rti.org.

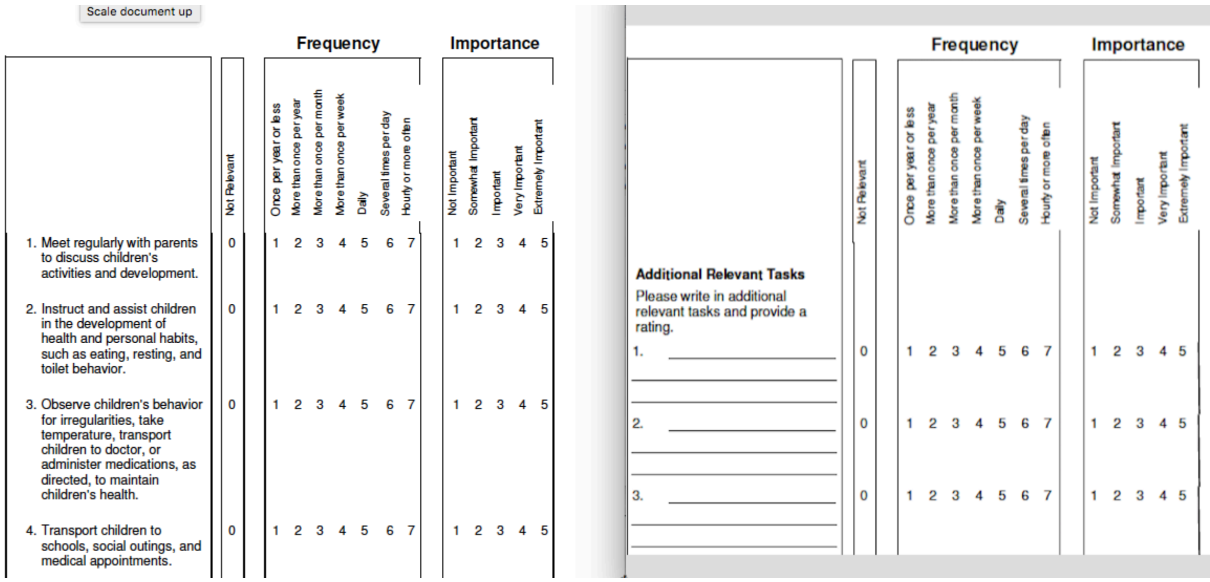

Figure III.2 O*NET Generate Tasks and Respondent Task Write-In Page (*Questionnaires: OE - Tasks, 2018 OMB Clearance Package*, 2018)

As Figure III.2 reflects, respondents indicate the importance of and frequency with which they complete both listed tasks and those they add. Occupational analysts identify the provided list with the mandate of enumerating all tasks 80% of incumbents could be expected to perform (Dierdorff & Morgeson, 2007). O*NET's architects recognize that many occupations have some idiosyncratic tasks in different organizations, but expect significant overlap and seek to identify that overlap. As O*NET is creating the task lists for the questionnaires, it contacts current holders of that occupation to review the lists for accuracy. It then sends questionnaires such as those in the figures to jobholders in the specific occupation randomly selected from those at businesses statistically selected from a random sample ("ETA O*NET Data Collection Program, Employment & Training Administration (ETA) U.S. Department of Labor," n.d.). The number of randomly selected jobholders from each business in the sample is based on the proportion of jobholders in the specific labor market of interest who work for businesses of the same type (Dierdorff & Morgeson, 2007). For example, O*NET samples lawyers from an investment bank based on the percentage of lawyers who work at an investment bank nationwide. O*NET's analysts aggregate the responses to the questionnaires in a report.

O*NET's methodology raises two potential points of concern. The first is whether the incumbent can adequately describe and assess the tasks he or she performs. While there is some disagreement as to whether the socially constructed aspects of jobs render them directly observable and measurable (Sanchez & Levine, 2000), job analysis researchers generally agree that a task is often observable. And people are able to describe precise aspects of the enactment of job tasks because a task is usually a single purposeful action (Dierdorff & Morgeson, 2007; R. J. Harvey & Wilson, 2000). And at the task level, high levels of inter-rater and intra-rater reliability among people evaluating job tasks supports the idea that tasks can be directly observed, codified, and rated (Butler & Harvey, 1988; Dierdorff & Wilson, 2003). A second concern could be that though there is, on average, consensus about job tasks, the level of consensus on those tasks could vary based on the specifics of the work context. Dierdorff and Morgeson (2007) found that for jobs that took place in highly interdependent or highly autonomous work contexts, respondents were more likely to have less consensus on the tasks of the jobs. They therefore tend to describe their tasks somewhat idiosyncratically. Alternatively, work contexts that were described as highly routinized resulted in higher consensus and less idiosyncratic description of tasks.

Ultimately, it seems that some of the O*NET task descriptions for occupations are probably more accurate than others. Given the potential for different levels of consensus about the job tasks in an occupation, the difference in accuracy is inevitable. However, O*NET data collection strategies limit the possibility for *completely* idiosyncratic responses by prepopulating the questionnaires with tasks that job incumbents have verified. To that end, while the O*NET tasks descriptions are unlikely to be the absolute ground truth of what people do in jobs, they are probably sufficiently accurate to uncover broad job trends in the economy over many years.

Natural Language Processing – A Primer

The amount of textual information that human processes create, sometimes unintentionally, has exploded in recent years. This creation has paralleled the growth in processor power and electronic storage capacities. The sheer volume of textual information can make certain kinds of manual content analysis virtually impossible. However, the rise in computing power has brought with it the concurrent development of computational methods to parse meaning out of seas of textual data. In order to discern meaning from these large groups of texts, or corpora, many content analytic strategies have been devised. These include but are not limited to word counts/dictionary methods, topic modeling, and embedding methods.

Word Counts/Dictionary Methods

Word counts are, as the name suggests, totals of the number of words that appear in a given corpus. Dictionary methods and Term Frequency-Inverse Document Frequency (TF-IDF) matrices also use word counts for analysis of a corpus. These are used very commonly across a number of applications, and variations of word counts have been used quite famously and somewhat recently in the analysis of 10-K filings (Hoberg & Phillips, 2010).

Word counts are the simplest of the three methods. This involves using word counts to infer some meaning about a text or a group of texts, depending on the level of analysis for the research question at hand. The unit of analysis could be tweets; it could be academic papers; it could be books. But the overall idea is that counting the number of uses of a particular word in a document will create useful knowledge. This knowledge appears in the form of a word count matrix, typically a spreadsheet with a list of words[20] down the vertical, the unit of analysis (e.g., tweets, books) along the horizontal, with a number at the intersection of word and analytical unit representing the number of times that word appears.

The word counts approach reflects the "bag of words" assumption. These methods do not take word order into account in their methods of analysis (Jurafsky, 2000). While language is complex and some of that complexity is encoded in word order, it is assumed (originally to minimize algorithmic complexity) that the presence of words preserves a lot of the meaning. Thus it is possible to treat words as a bunch of objects in a bag, rather than structured in sentences. It is possible to adjust this assumption, slightly, by looking for what are known as n-grams. Instead of simply counting individual words (1-grams or unigrams), you also count pairs of words (2-grams), or triplets of words (3-grams). But this may not substantially add to the process of inferring meaning (Grimmer & Stewart, 2013). The utility of the word count should not be underestimated, as much can be inferred by knowing how many times a word is used. However, the methodology has some limitations. Notably, word count data, while clearly more compact than the actual text, can still be quite voluminous to extract knowledge from and manage, a weakness that subsequent methods try to assuage.

Dictionary methods address some of the weaknesses of word counts. They rest on the idea that there is a set of words that can all be associated with a given subject, and that counting all of those words will generate useful data. A common example of a dictionary method is sentiment analysis. In this instance, the analyst uses a custom dictionary that has already associated certain words with a broader category. For example the words "great," "wonderful,"

[20] Another common text analysis practice is called lemmatizing (e.g., Grimmer & Stewart, 2013; Hobson et al., 2019). This involves using the word stem of a family of words to represent all the words in the family. Instead of counting "work," "working," "worked," and "works" one time each, you would count "work," the stem, four times. Of course, that poses its own problems. The sentences "I worked on my dissertation" and "This dissertation worked me over" would be ill represented by saying "work" appeared twice.

"positive," "very," "well," "like," "awesome," and "good" could be associated with the subject "positive," while the words "bad," "terrible," "disliked," "hate," and "negative" could be associated with the subject "negative." A software program using this custom dictionary could then return a measure of the positive and negative sentiment of whatever the unit of analysis is. Two common software packages with custom dictionaries that people use for sentiment analysis are the Linguistic Inquiry and Word Count (better known by its acronym, LIWC; Tausczik & Pennebaker, 2010) and the Valence Aware Dictionary and sEntiment Reasoner (better known as Vader; Hutto & Gilbert, 2014).

The TF-IDF Matrix (Sparck Jones, 1972) also builds on basic word counts, in that the words that appear in the text and how often they appear tell you something about the meaning of the text. But TF-IDF allows researchers to compare how different documents use a particular word. It involves calculating terms frequently – how often a given term appears in a document divided by how many words are in that document – and multiplying it by the inverse document frequency, which is the log of the percentage of times it appears in all documents. An example will help clarify this:

Document 1: This world is a dog-eat-dog world.

Document 2: This world is a breeze for cats.

Document 3: Isn't the world a funny place?

In Document 1 the term frequency for dog and world alike is 2/8 = .25. However, the inverse document frequency across the three documents for dog is log(3/1) = 1.09 and for world is log(3/3) = 0. The TF-IDF for dog is .25*1.09 = .2725 and for world is .25*0 = 0. This tells us that dog is a more important word than world in representing the meaning within and between documents. TF-IDF analysis introduces a second important concept to this discussion:

compression. Scaling by the inverse document frequency compresses the data. Many words will now have no significance in the data set, and that allows for a more compact representation of the data.

Topic Modeling Overview

The goal of topic modeling is almost precisely what it sounds like. A researcher takes a group of documents and tries to understand what topics would provide high-level categorizations of the material conveyed by the documents collectively. If a corpus is small, this task is easily accomplished by reading. However, for a large corpus, topic modeling provides understanding that mere reading cannot provide, or cannot provide in a practical way. In addition to saving time, the method can identify common topics between many texts that humans might not identify.

Topic modeling picks up where word counts and dictionary methods leave off by dealing with a representational limitation of these methods. Word counts and dictionary methods show that certain words in documents can help distinguish some meaning within and between documents. These methods were also originally groundbreaking in that they represent documents as mathematical objects. However, these representations cannot fully explain the similarities between documents and convey the extent to which documents have the same meaning. One of the benefits of representing nonmathematical objects as mathematical objects called matrices, which word counts and TF-IDF generate, is that each individual column in a matrix is a vector representation of an observation in multidimensional space. The measured distance between those vectors highlights similarity and difference between those observations. Matrices such as a word count matrix and a TF-IDF matrix are useful data compressions and highlight some meaning. While they are still based on the raw count of individual words, which loses some meaning, they do encode a particular relationship that can be exploited in order to better capture the meaning of documents: which words co-occur together. Topic modeling uses word counts and TF-IDF matrices as inputs and processes them to enumerate word co-occurrences and assign them to topics. Topic modeling comes in several forms, and Latent Semantic Analysis (LSA) and Latent Dirichlet Allocation (LDA) are among the most well known.

Latent Semantic Analysis

The "latent" in LSA stems from an assumption on which LSA operates: documents have latent variables in common. These latent variables are called topics. LSA involves collecting all of the co-occurrences of words in groups of documents. It does this through singular value decomposition (SVD) of a TF-IDF matrix (Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990). SVD is a linear algebraic method by which a higher order matrix is decomposed into three lower order matrices. It can be represented in matrix notation like this: $W_{mxn} => U_{mxp} S_{pxp} V_{pxn}^T$ (Lane, Howard, & Hapke, 2019); decomposition is a factorization. Just as the number 42 can be factored into three numbers (2, 3, and 7) using prime factorization, a large matrix can be factored into smaller matrices using SVD (Hornick, 2016).

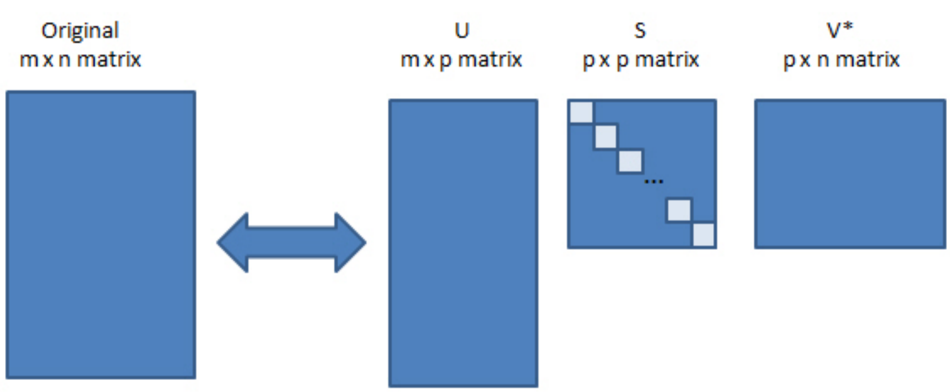

Figure III.3 Singular Value Decomposition (Hornick, 2016)

When you perform an SVD on a TF-IDF (W_{mxn}) the output is a word-topic matrix (U_{mxp}), a topic-topic matrix (S_{pxp}), and a document-document matrix (V_{pxn}^{T}). The values in each cell of the word-topic matrix indicate the importance of each word to a given topic. This is a valuable matrix for two reasons. The first is that it shows which words belong to each topic and that makes it possible to evaluate what latent variable the SVD has uncovered. LSA of a corpus of tweets about the 2016 election might reveal that a topic (here a latent variable, not a topic in the more common meaning of the subject) was composed of the words: "emails," "server," "Comey," "FBI," and "Russia." The latent variable or topic that the matrix identifies would be Secretary Clinton's emails. Word-topic matrices are also valuable because, if you multiply the original TD-IDF term document matrix (W_{mxn}) by the word-topic matrix (U_{mxp}), the output is a matrix of size nxp, generating a topic-document matrix. This topic-document matrix will show that each document is actually a mixture of topics, and this matrix will show the relevance of each topic to a given document. The most precise definition of LSA that I have seen is that LSA "uses a singular value decomposition of the X matrix to identify a linear subspace in the space of TF-IDF features that captures most of the variance in the collection" (Blei et al., 2003, p. 993). A simpler way to think about this is that LSA uses matrix algebra to identify which words co-occur often and then sorts them into groups called topics. LSA can be useful enough, depending on the scope of the text analysis problem. But a challenge with LSA and TF-IDF is that they effectively scale proportionally to the number of documents that are in the corpus under analysis. In LSA specifically, the outcome of SVD is a matrix that captures essentially all the variance in a corpus, and it will create as many topics as it needs to do that. Sometimes maximum variance capture will create topics that are very difficult for humans to interpret. LDA addresses this problem.

Latent Dirichlet Allocation

LDA is more popular than LSA and in certain ways more interpretable. While the kernel of this method was originally positioned for use in analyzing genomics data (Pritchard, Stephens, & Donnelly, 2000), it was revised and popularized for use in natural language processing applications (Blei et al., 2003). Where LSA's attempt to capture variance can result in uninterruptable topics, LDA allocates words to topics probabilistically, which results in a stricter allocation of words to topics. One benefit of this is topics that are more easily interpreted by humans.

To understand how LDA works, start with a mental exercise and, rather than thinking of what topics a corpus might have, think of topics as generative genre machines. For example, the scripts of the movies *The Godfather* (crime/drama), *Analyze This* (crime/comedy), *The Notebook* (romance/drama), and *Notting Hill* (romance/comedy) might represent the corpus. Crime, drama, comedy, and romance genre machines could not *write* the scripts, but they might be able to *build* them. Each might have a bag of Magnetic Poetry™ [21] and pull out words at random and affix them to a refrigerator to create the building blocks of the script. How many genre machines would be necessary, how much would each genre machine contribute to each script, and what would the distribution of words have to be? In this example, one could see that we would probably need four genre machines – a comedy machine, a drama machine, a romance machine, and a crime machine. The crime and the drama machines might contribute equally to the construction of *The Godfather*, and the crime and comedy machines might contribute approximately equally to the construction of *Analyze This*. But in both movies, there is a love interest (or two), so the romance machine would actually contribute to both as well. As for the

[21] Magnetic Poetry™ are sets of tiny magnets with words printed on them, typically related to a particular theme; as of this writing, the website includes sets focused on the moon, bicycles, Emily Dickinson, and grouchy cat, to name a few (www.magneticpoetry.com).

distribution of words in the bag, that is harder to posit. But it seems likely that the crime machine would need to have the words “gun,” “family,” and “don” in fairly high frequency. And the romance robot might need the words “love,” “need,” and “life” in fairly high frequency.

The overall theory of LDA is that all documents being analyzed are a mixture of topics (think genres from above), each topic is a distribution of words, and words can be assigned to multiple topics. The input to a topic model is a TF-IDF matrix or a simple word-document matrix. So, to estimate a topic, LDA is still examining the co-occurrence of words in documents. In order to decide the distribution of words in each topic, a multinomial distribution is used as a prior, and to allow documents to have multiple topics, a Diricihlet prior is assumed, and that gives the name to the process. The process then iterates over the words and the documents in the corpus until it arrives at the topics with the correct word distribution per topic and the correct mixture of topics in the documents, to maximize the likelihood of generating the set of documents. However, some user input is required. And unlike LSA, LDA requires that a user specify a constraint – the number of topics. The output of topic modeling using LDA is a word-topic matrix, in which the value in each cell indicates how important that word is to a given topic, and a document-topic matrix, in which the value in each cell shows what percentage of that document the topic represents. The stricter allocation of words to topics in LDA as compared to LSA provides topics that are more interpretable.

Both LSA and LDA can be used to summarize the topics of large bodies of text. The vector representation of documents in both methods encode more of the meaning of the documents than in word counts or TF-IDF matrices, such that the distance between two vectors in multidimensional space is an indication of the similarities between any two documents. The farther (closer) they are, the more dissimilar (similar) the documents are. This has utility in a wide array of settings and is gaining traction in the management literature (e.g., Croidieu & Kim, 2018). However, embedding methods builds on LDA and addresses some of its limitations.

Embedding Methods

Topic modeling, including LDA and LSA, relies on gathering up the co-occurrences of words across a corpus of documents in order to infer potential topics across the documents. Thus it relies on a "bag of words" assumption, in that all we really need to know to discern meaning about the document is what words are in the document and nothing about the structure. This simplifying assumption is both procedurally and computationally necessary in order to make topic modeling possible. But both advances in scale computational machinery and advances in neural network architecture set the stage for a method that would allow for a more precise representation of the text. These are collectively known as embedding methods.

Understanding how embedding methods work requires an explanation of a neural network. A neural network takes inputs and generates outputs, where the output is essentially a weighted sum of the inputs (Chollet, 2017). It is called a neural network because the inspiration of the concept is the network of neurons in the human brain. Prediction is among the many possible uses of neural networks. To illustrate, let's say for example that we have the following information about a baseball team: average number of toes per player, current win/loss record, and number of fans in millions. From this information, we want to generate a neural network that will predict if the team will win a game. Figure III.3 is a diagram of this task.

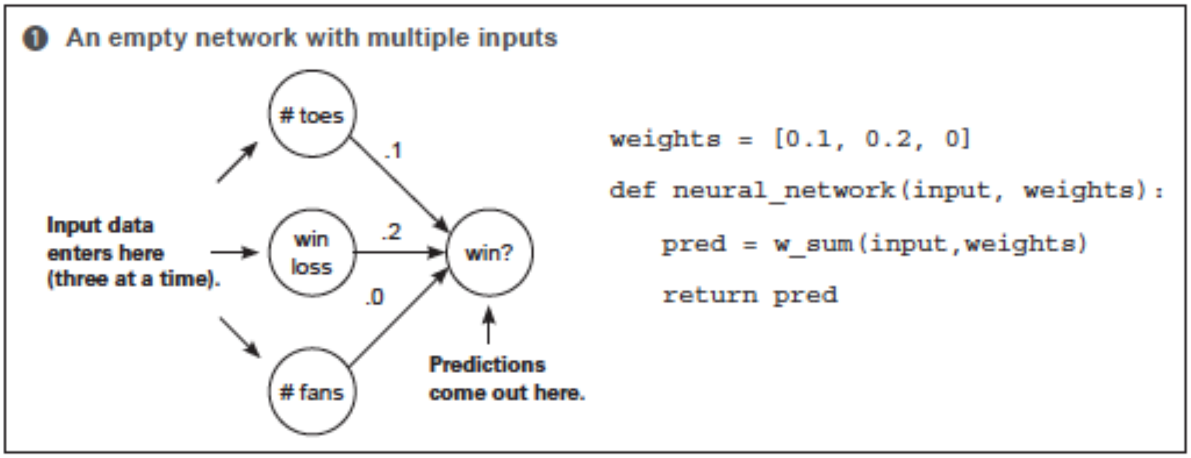

Figure III.4 Simple Neural Network From Trask (2019)

In order to generate a prediction, we have to "train" the neural network. This starts with running training data through the neural network and making adjustments. For example, if you want to predict an unknown performance of the Detroit Tigers in 2019, you could train the model on the known performance of the Detroit Tigers from 2018. This involves inputting data – toes, wins, losses, and millions of fans – into the neural network. The output will be a weighted sum of these values. This output is compared to desired output. The difference between the two values is the error. The error value is used to update the weight of the neural network, and then this is repeated for the next set of data. Again, a diagram is instructive.

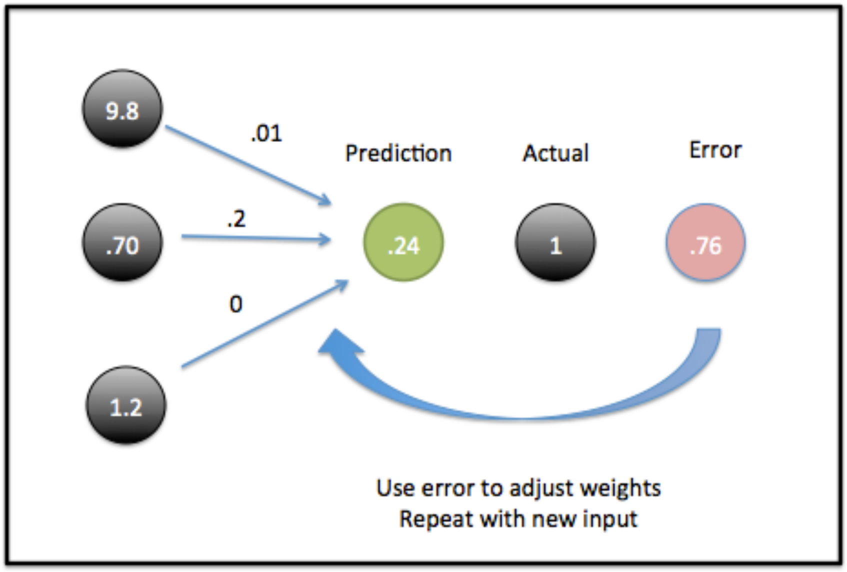

Figure III.5 Neural Network Diagram With Error Correction

The variance in the inputs will generate different levels of error and the weights will fluctuate, but hopefully converge to a fixed set of numbers. The process by which the error term is used to update the weights is called back propagation – that is, the neural network learns from its mistakes. Extending that analogy suggests that these weights are the neural network "storing" knowledge about the relationship between the inputs and the outputs (Trask, 2019). All neural networks at their most fundamental level have this basic architecture; what changes is the type of inputs, the number of neurons, and the numbers of connections between neurons. These are all varied and increased in order to process more types of data and to "store" more complicated representations of that data.

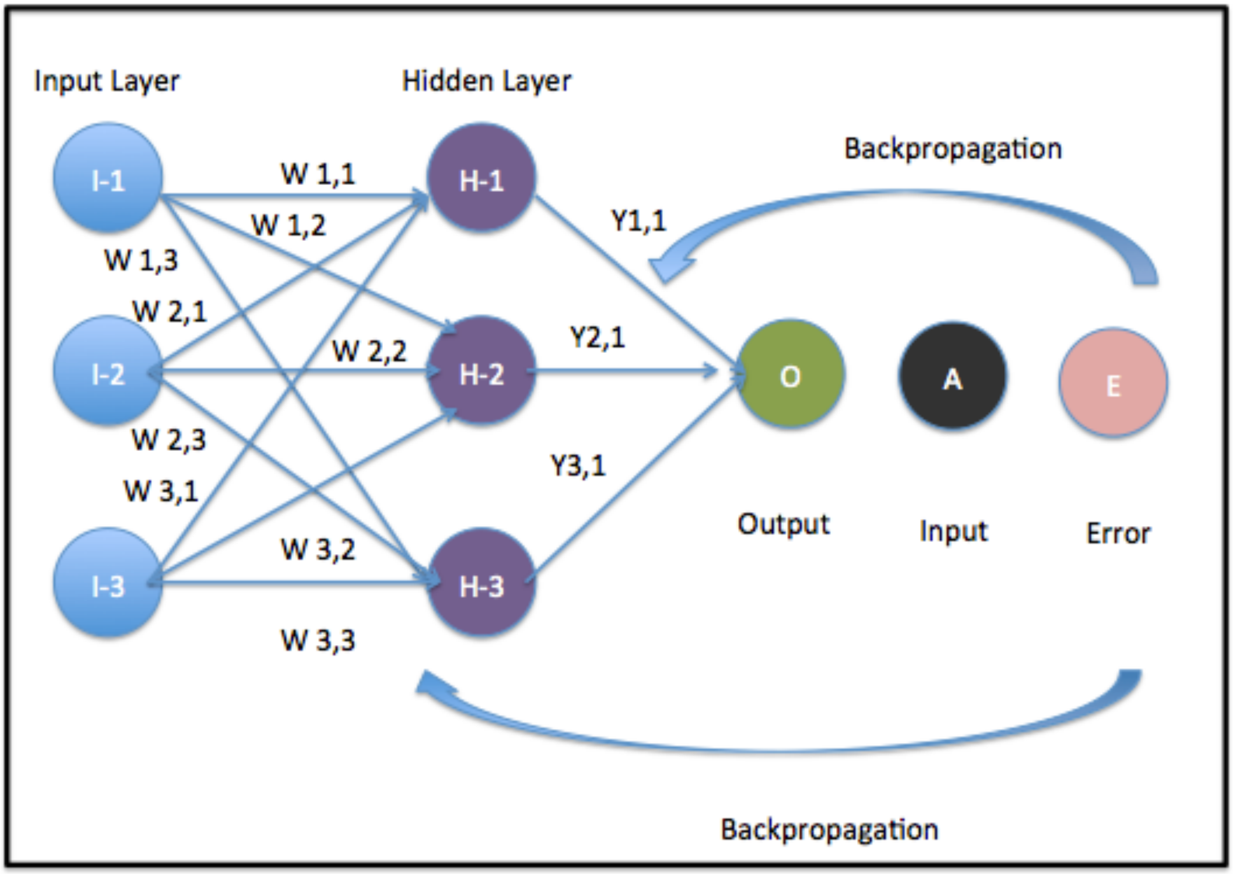

Figure III.6 Neural Network Diagram With Additional Layers

More complicated neural networks have "hidden" layers of neurons. Rather than being a direct connection between input data (I-1, I-2, I-3) and the output (O), there is an intermediary layer of neurons (H-1, H-2, H-3) called the hidden layers, and there are weights between all of the neurons, between the input and the hidden layer, and then between the hidden layer and output layer. The hidden layer "neurons" are equations, the equations vary from context to context, but a commonly used equation is the sigmoid function 1/1+e^-x (Trask, 2019). Much in the way that multiple regression allows for more detailed analysis of certain data sets, neural networks with hidden layers can represent more complicated sets of data. The value passed from the hidden layer includes the weighted sum based on the weights Y1,1, Y2,1, and Y3,1 and the output values of the neurons that are calculated by the equation. But just as with the simple neural network, back propagation helps adjust the weights of the neural network.[22]

Embedding Models Are Neural Networks

Embedding models are neural networks[23] and the weights effectively store knowledge about the words. A word embedding model – a model that uses text as input data and "stores" meanings of those words – is created by training a neural network to use the words in a window around a focal word in a sentence to then predict the probability of that word. This method is known as the continuous bag of words model (Mikolov, Sutskever, Chen, Corrado, & Dean, 2013).

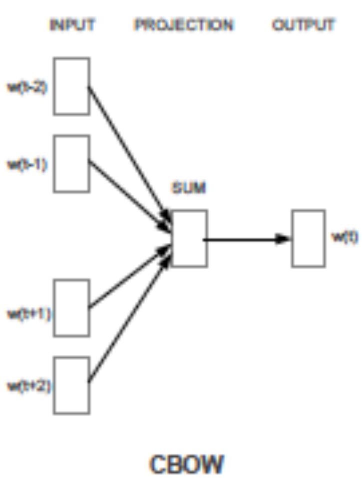

Figure III.7 Continuous Bag of Words Model (Mikolev et al., 2013)

This neural network consists of an input layer, a hidden section with several layers of neurons, and an output layer that is the target word. Take the example, "The cake tastes very

[22] Matrix multiplication makes calculating the values passed between layers trivial. And the key to back propagation involves calculating the gradient of the sigmoid function in this case (Chollet, 2017). This is why multivariable calculus and linear algebra are prerequisites for going deep into this methodology.
[23] Usually they are. But they don't have to be! Moody (2017) established that you can get embeddings by performing an SVD on a matrix of pointwise mutual information between pairs of words.

sweet." The input layer would be the words "the," "cake," "very," and "sweet," and the output layer would be the word "tastes." Neural networks represent word inputs as one-hot vectors (Lane et al., 2019). Imagine a spreadsheet with every word in your corpus in one column in alphabetical order. If you wanted to represent the word "taste," in the column adjacent to "taste" you would place a 1, while putting 0s everywhere else. The column of the 1s and 0s is now a one-hot vector. In most neural network applications, we care about the model's ability to make predictions, but in this case we actually care about the "knowledge" stored in the weights immediately after the first layer. Much like people can interpret the meaning of words by looking at their surrounding context, this neural network does the same thing, and it encodes the meaning of words in the matrix of weights immediately after the first layer.

To determine the validity of this model, Mikolev et al. (2013) asked a simple question: king − man + woman = ? The answer to this is "queen." If their model's "knowledge" of words was accurate, then the model should represent this analogy mathematically. Sure enough, when taking the vector they calculated for "king," subtracting the vector for "man," and adding the vector for "woman," the resultant vector was very close to the vector for "queen." By using neural network architecture, they were able to create a mathematical, or vector, representation of meaning given the context in which words appeared.

An extension of the word embedding model is paragraph embedding (Le & Mikolov, 2014). The key addition is that you include another input to the model beyond the immediate window of the focal word.

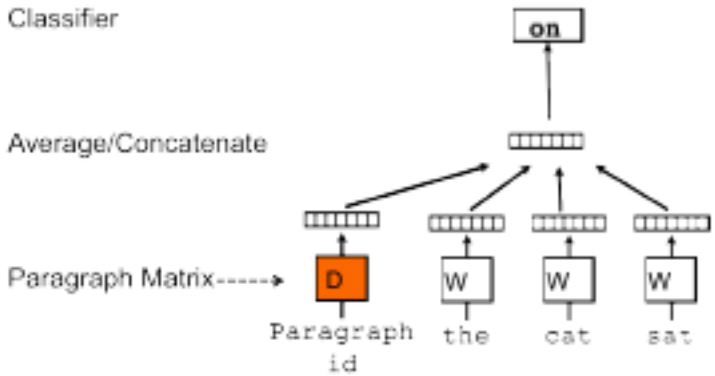

Figure III.8 A Model of Document Embedding (Le and Mikolov, 2014)

In addition to the one-hot vector for each word, you have a one-hot vector for paragraphs. If the word you are examining is in paragraph 1 of your corpus, then you have a vector with a 1 at the paragraph: one location followed by a bunch of 0s. By adding in this paragraph vector, you are also able to encode the meaning of entire paragraphs as vector representations.

Text Analysis Usage in the Literature

Many of the text analytical methods used here can be seen in action in the management literature. They have been used roughly in proportion to how long they have been a part of the natural language processing canon. Thus, word counts and dictionary methods appear the most frequently in the literature. One of the most extensive examples of these methods was the work of Hoberg and Phillips (2010), who were able to discern better industry groupings over and above Standard Industrial Classification codes by examining which words companies had in common in the product descriptions in their annual 10-K flings with the SEC.

Topic modeling is not used as often, but it has increased in the past several years. Kaplan et al. (2015) used topic modeling on a corpus of patents in order to uncover when a patent was introducing new knowledge into a field. Essentially, when a given patent was the first appearance of a new topic, that was indicative of new knowledge. Huang et al. (2017) also used topic modeling to create representations of generation of new knowledge, but the area of focus was securities analysts. They created a topic model based on the text of transcripts of corporate earnings conference calls and security analyst reports within a focal industry. They then compared the topic distribution of the conference calls to the topic distribution of the analyst reports to see the extent to which the analyst report reaffirmed information in the call (had the same topics) or surfaced new information that was not covered on the call (had different topics). The difference in topics was then correlated with abnormal returns in stock prices. The last entry is the work of Croidieu and Kim (2018) in which they examined the legitimization of lay-knowledge in the early history of radio in the United States. They used topic modeling on historical documents to try to identify different periods in the institutional development of the nascent radio industry, while also identifying the various mechanisms associated with processes of gaining legitimacy. Whereas the first two works used the topics results largely as inputs in regression models, Croidieu and Kim actually use the fact that the output of topic models, words and topics, can be reinterpreted as codes and themes, setting a good starting point for grounded theory (Strauss & Corbin, 1997) or thematic analysis (Braun & Clarke, 2012, 2013). To date, I cannot find any instances of embedding models making their way into the management literature, but given the success of other text analytical methods, I believe they will.

The Present Work

I use a combination of LDA topic modeling with thematic analysis and document embedding with the job task data from O*NET as the analytical core of this chapter. The goal is to uncover knowledge about the task stack of jobs in the United States over the past 20 years. Specifically, I seek to quantify the nature of the underlying task stack, to see how elements of that task stack have changed over time, and to see if some jobs have gotten more or less similar over time. Topic modeling is well suited to the first task. Delineating the task stack within and across jobs before the fact would be difficult. Topic modeling will uncover the topics, or in this case, the broad categorizations of tasks both within and across jobs. In order to make these tasks interpretable, I use thematic analysis. This highly flexible tool is very well suited to discerning meaning across documents (Braun & Clarke, 2013). I use iterative rounds of thematic analysis, looking at the output of the topic model, a visual representation of the proximity of the topics to one another, and the actual descriptions of the jobs in order to arrive at a detailed description of the tasks. Then, as the topic model also provides a distribution of topics over documents, or in this case a distribution of tasks over jobs, I can both visualize the evolution of tasks within a job and, using a measure of entropy on the task distribution, quantify just how much jobs have changed and if certain jobs have gotten more or less complex over time.

One of the future goals of this work is to be able to test whether the movement of people between different kinds of occupations has any effect on how similar or dissimilar those occupations will be over time. While topic modeling does encode the task structure of jobs in the economy, embedding models do this with much greater accuracy. So for the second objective of understanding just how similar or dissimilar jobs have become over time, I will create a document embedding model, in which the unit of analysis, the document, is the job task description for every job for every year of data I have from O*NET. As mentioned earlier, using the task descriptions as the input to an embedding model means that one of the results will be a vector representation of each job description. To create a data visualization of the proximity of these jobs, I compress the multidimensional vector representations into two dimensions using a process called t-Stochastic Neighbor Embedding, or t-SNE (Maaten & Hinton, 2008). While these visualizations do not lend themselves as well to print, they are available on my website, www.teddydewitt.com.

Tools

All analysis in this section was performed using the Python packages. Sci-kit learn was used to calculate cosine similarity vectors between measures as well as the LDA analysis. Gensim was used to perform the word embedding analysis. Plotly and Bokeh are the two plotting libraries that were used to create the visualizations that can be seen on my website. Pandas was used to calculate summary statistics. And finally, scipy was used to calculate the paired t-test, t-statistics, p-values, and standard deviations.

Findings

To answer the research question, "How has the task structure of work in the United States changed over the past 20 years," I first had to model the task structure using LDA topic modeling. As I mentioned in the methods sections, LDA requires input from the user in the form of how many topics should be used. This can be both a help and a hindrance. In general, if a researcher has a prior theory about how many topics should be in a data set, he or she should start with that theoretical value. In this case, my theorizing about the task structure of jobs does not suggest potential number of topics to use. To figure out how many topics to use, I relied on what is called a coherence model. Coherence models (Chang, Gerrish, Wang, Boyd-Graber, & Blei, 2009; Mimno, Wallach, Talley, Leenders, & McCallum, 2011) identify pairs of words or the point-wise mutual information between pairs of words within all of the documents and indicate if the top 30 words in each topic adequately reflect these pairs of words. In theory, specifying a number of topics equal to the number of documents in the corpus could yield a perfect coherence measure. However, this would greatly reduce the ability to find commonalities among documents in the corpus. So, following standard practice for coherence models, I performed the analysis with several different numbers of topics and chose the model where the coherence measure started to level off significantly (Mimno & Blei, 2011). This resulted in my choice of 30 topics.

As mentioned in the methods section, I used thematic analysis, going back and forth between the top 30 words in each data set, and examples of the job descriptions in order to describe and then name the 30 task topics generated in the LDA analysis. Tables III.1 through III.3 list a description of each topic and its top 30 words, as well as which jobs display the greatest representation of that topic in selected years of the data set.

Table III.1 Topic Model Results With Thematic Analysis Part 1

Topic Number	Topic Name	Final Topic Theme	Top 30 words for topic	Top 3 jobs for topic (% of job) in 2004	Top 3 jobs for topic (% of job) in 2018
0	Forming Objects	forming, creating,and modifying physcal objects	molds, articles, patterns, mold, metal, garments, needles, fabric, thread, furnaces, clay, stones, smooth, place, sections, garment, edges, rails, seams, plaster, casting, pour, dry, sew, jewelry, wax, sewing, irons, ladles, sand	Foundry Mold and Coremakers (74%), Precision Pattern and Die Casters, Nonferrous Metals (64%),	Foundry Mold and Coremakers (74%) Refractory Materials Repairers, Except Brickmasons (59%) Carpet Installers (58%)
1	Information Analysis and Management	data analysis, planning, coordinating, intepretaton and recommendations associated with business activtities	develop, prepare, review, procedures, direct, activities, analyze, ensure, evaluate, information, work, programs, personnel, coordinate, staff, plan, reports, data, monitor, conduct, management, problems, systems, new, services, policies, maintain, maintenance, clients, design	Private Sector Executive (99%), Government Service Executives (97%),Credit Analysts (96%)	Credit Analysts (96%) Financial Managers, Branch or Department(94%) Compensation and Benefits Managers (93%)
2	Stocking and Cleaning	collecting, restocking supplies, cleaning work areas, with light record keeping	supplies, inform, calculate, aircraft, ensure, work, procedures, information, estimate, examinations, supply, disposal, apply, instructions, wash, distribute, including, operate, keep_records, inventories, operation, arrange, control, records, inmates, supervise, changes, inventory, plan, altitudes	Bartenders(63%), Surgical Technologists(43%), Refuse and Recyclable Material Collectors(43%)	Bartenders (63%) Surgical Technologists (41%) Purchasing Managers (39%)
3	Safety Data	collecting data, gathering information, converying communcations with regards to accidents and emergencies	vehicles, safety, water, systems, ensure, cars, regulations, vehicle, traffic, accidents, compliance, passengers, monitor, damage, loading, direct, procedures, repair, hazards,repairs,maintenance,flight,conditions,report,mechanical,areas,flights,delays,fuel,sites	Insurance Appraisers-Auto Damage(75%), Forest Fire Inspectors and Prevention Specialists 73%,Train Crew Members(64%)	Forest Fire Inspectors and Prevention Specialists (73%) Insurance Appraisers, Auto Damage (63%) Bus Drivers, School or Special Client (55%)
4	Exterior Maintenace	maintenance, repair and keeping of greenspaces and exteriors of buildings	cut, select, tools, work, trees, inspect, materials, areas, operate, spray, clean, remove, trucks, plants, holes, cutting, shovels, saws, set, surfaces, hand, conveyors, brush, tractors, cuts, paint, using_hand, place, load, plant	Tree Trimmers and Pruners(89%), Nursery Workers(84%) ,Insulation Workers, Floor, Ceiling, and Wall (78%)	Landscaping and Groundskeeping Workers (75%) Tree Trimmers and Pruners (71%) Forest and Conservation Workers (68%)
5	Machine Operation	operation and maintenance of machines - reading gauges, using hand tools dealing with fluids	water, oil, repair, valves, gas, pumps, engines, tools, test, maintain, gauges, control, tanks, mechanical, controls, operating,meters, pipe, using_hand, vessels, pipes, air, ensure, maintenance, adjust, systems, machinery, wells, units, drilling	Gas Pumping Station Operators (92%), Auxiliary Equipment Operators-Power (91%),Home Appliance Installers (90%)	Wellhead Pumpers (78%) Rotary Drill Operators, Oil and Gas (72%) Service Unit Operators, Oil, Gas, and Mining (71%0
6	Constructiing Things	construction of buildings, surfacing lands and walkways	surface, materials, surfaces, concrete, stone, finish, sand, material, grout, mix, hand, metal, roofs, chisels, steel, position, trowel, tools, joints, tile, cement, rods, prepare, clean, foundation, posts, required, machines, furnaces, spread	Roof Bolters-Mining (78%), Cement Masons and Concrete Finishers (77%), Fence Erectors (77%)	Cement Masons and Concrete Finishers (77%) Terrazzo Workers and Finishers (72%) Segmental Pavers (70%)
7	Detailed Handcraft	detailed hand craft work and repair often associated with shaping glass, metal, examining instruments	tools, parts, workpieces, specifications, cut, position, templates, patterns, using_hand, metal, shape, glass, measuring_instruments, hand_tools, fit, cutting, blueprints, dimensions, welding, workpiece, remove, according, insert, dies, align, cutting_tools, instruments, set, cuts, adjust	Solderers (99%), Gem and Diamond Workers (81%), Reed or Wind Instrument Repairers and Tuners (80%)	Layout Workers, Metal and Plastic (73%) Drilling and Boring Machine Tool Setters, Operators, and Tenders, Metal and Plastic (71%) Ophthalmic Laboratory Technicians (70%)
8	Healthcare work	adminisrating, peforming and mointoring patient healthcare, health care counseling	patients, medical, treatment, clients, procedures, patient, provide, care, develop, perform, health, information, assist, treatments, physicians, problems, evaluate, administer, assess, monitor, tests, techniques, staff, maintain, health_care, devices, conditions, necessary, record, records	Physician Assistants (90%), Obstetricians and Gynecologists(89%), Family and General Practitioners (88%)	Physician Assistants (90%) Obstetricians and Gynecologists (89%) Family and General Practitioners (88%)
9	Coordinating and Liaising	directs,coordinates organizes, supervises liaises with several populations of people and employees - some financial and general recordkeeping	sales, workers, services, monitor, harvesting, direct, players, crops, products, game, machinery, ensure, farmers, records, arrange, cards, farm, prospective_customers, pay, grading, fields, production, work, advise_customers, supervise, planting, perform, cultivating, irrigation, credit	Agricultural Crop Farm Managers (99%), Pilots-Ship (64%), First-Line Supervisors and Manager/Supervisors - Agricultural Crop Workers (54%)	Farm Labor Contractors(52%) Sales Managers (50%) Gaming Dealers (48%)

Table III.2 Topic Model Results With Thematic Analysis Part 2

Topic Number	Topic Name	Final Topic Theme	Top 30 words for topic	Top 3 jobs for topic (% of job) in 2004	Top 3 jobs for topic (% of job) in 2018
10	Detailed Work	engaged in meticulous detail work, and often repetitive actions	order, maintain, operate, apply, materials, make, necessary, use, check, work, prepare, keep_abreast, provide, areas, used, stock, perform, developments, place, arrange, according, examine, types, registration, items, required, clean, position, set, credit	Proofreaders and Copy Markers (72%) ,Astronomers (46%), Painters and Illustrators (44%)	Proofreaders and Copy Markers (56%) Endoscopy Technicians (39%) Parts Salespersons (35%)
11	Shaping Modification	chemical usage to modify, form, clean shape both human and non-human fixtures	nails, surfaces, clean, water, walls, plates, colors, brush, brushes, remove_excess, negatives, dry, wax, plate, place, strips, rollers, film, customers, ceilings, nail, prepare ,tires, wheels, hair, attach, scissors, plaster, floors, material	Ceiling Tile Installers (78%), Dental Hygienists (77%), Paperhangers (72%)	Dental Hygienists (69%) Hairdressers, Hairstylists, and Cosmetologists (66%) Floor Layers, Except Carpet, Wood, and Hard Tiles (66%)
12	Operation and Maintenance	operating and maitenance of machines and hand tools, usage of hand tools	machines, materials, machine, specifications, products, according, material, production, using_hand, tools, specified, record, adjust, product, processing, work, parts, measure, set, operation, start, ensure, work_orders, data, load, type, examine, records, operate, remove	Coating, Painting, and Spraying Machine Setters and Set-Up Operators (99%),Typesetting and Composing Machine Operators and Tenders(99%), Duplicating Machine(98%) Operators()	Textile Bleaching and Dyeing Machine Operators and Tenders (77%) Textile Winding, Twisting, and Drawing Out Machine Setters, Operators, and Tenders (76%) Adhesive Bonding Machine Operators and Tenders (73%)
13	Human Movers	attends to, monitors, assists and proccesses humans moving between contexts	customers, mail, patrons, information, provide, record, services, train, service, guests, check, areas, ensure, passengers, necessary, answer, assistance, schedules, receive, rooms, supplies, maintain, delivery, keep_records, greet, schedule, contact, messages, facilities, order	Bailiffs (88%),Transportation Attendants, Except Flight Attendants and Baggage Porters (86%), Amusement and Recreation Attendants (76%)	Transportation Attendants, Except Flight Attendants (77%) Amusement and Recreation Attendants (77%) Baggage Porters and Bellhops (76%)
14	Creative Decisioning	evaluation and decision making regarding diverse human activity, coordinating with multiple parties to inform process	film,music,production,energy,create,write,programs,scripts,video,television,performances,effects,media,lighting,clients,design,collaborate,study,sound,research,broadcast,performers,develop,material,ideas,advertising,producers,confer,sets,audio	Cartoonists(90%), Poets and Lyricists(90%0, Insurance Underwriters(84%)	Insurance Underwriters(79%) Film and Video Editors (72%) Sound Engineering Technicians (70%)
15	Information Collection	information collecting, analyzing and assessing;conducting investigatons, creating records, often electronic	information,records,record,data,reports,documents,security,correspondence,medical,maintain,obtain,databases,computers,claims,compile,files,evidence,process,cases,prepare,court,computer,forms,review,obtain_information,identify,interview,perform,regulations,necessary	Insurance Adjusters, Examiners, and Investigators (90%), Government Property Inspectors and Investigators (89%), Fish Hatchery Managers (82%)	Insurance Adjusters, Examiners, and Investigators (85%) Private Detectives and Investigators (78%) Fire Investigators (74%)

Table III.3 Topic Model Results With Thematic Analysis Part 3

Topic Number	Topic Name	Final Topic Theme	Top 30 words for topic	Top 3 jobs for topic (% of job) in 2004	Top 3 jobs for topic (% of job) in 2018
16	Repair and Monitoring	repairing, monitoring, assessing machinery and electical processes	parts,tools,repair,install,test,using_hand,power,components,systems,electrical,specifications,installation,clean,work,maintain,machines,repairs,assemble,instruments,fixtures,measure,replace,hand_tools,machinery,devices,adjust,wiring,remove,maintenance,examine	Station Installers and Repairers, Telephone (97%), Bicycle Repairers (95%), Electric Motor and Switch Assemblers and Repairers (88%)	Helpers--Installation, Maintenance, and Repair Workers (85%) Helpers--Electricians (83%) Mobile Heavy Equipment Mechanics, Except Engines (81%)
17	Post-Secondary Education	tasks broadly assocaited with post-secondary education	research,issues,students,work,colleagues,professional,compile,records,evaluate,undergraduate,participate,prepare,supervise,laboratory,administer,field,knowledge,methods,conferences,journals,government,grades,act,assignments,industry,papers,perform,books,participating,electronic	Atmospheric, Earth, Marine, and Space Sciences Teachers, Postsecondary(69%), Biological Science Teachers, Postsecondary (68%), Environmental Science Teachers, Postsecondary(68%)	Anthropology and Archeology Teachers, Postsecondary (68%) Engineering Teachers, Postsecondary (67%) Physics Teachers, Postsecondary (67%)
18	Engineering Management	engineering related work, monitoring of physical movements, liasing with clients or extra-organizational parties	estimate,processes,temperature,heat,operate,measure,add,processing,test,make_adjustments,supervisors,metals,district,pressure,workers,monitor,transfer,braces,specific_gravity,vats,specified_amounts,based,computer,units,component,knowledge,notify_supervisors,filling,drain,conditions	Chemical Engineers (78%),Civil Engineering Technicians (55%),Tire Builders (51%)	Chemical Engineers (65%) Tire Builders (51%) Civil Engineering Technicians (47%)
19	High Technical Systems	analysis, maintenance, production of computer and other high-tech technical systems	systems,design,develop,test,requirements,computer,procedures,use,users,project,approve,fuel_cell,operating,technicians,web_site,logistics,identify,technologies,installation,supervise,plans,protocols,database,web_sites,related,problems,coordination,programmers,performance,collaborate	Electrolytic Plating and Coating Machine Operators and Tenders, Metal and Plastic (38%), Computer Systems Analysts (37%), Computer and Information Systems Managers (35%)	Radio Frequency Identification Device Specialists (45%) Database Architects (41%) Web Developers (40%)

Table III.4 Topic Model Results With Thematic Analysis Part 4

Topic Number	Topic Name	Final Topic Theme	Top 30 words for topic	Top 3 jobs for topic (% of job) in 2004	Top 3 jobs for topic (% of job) in 2018
20	Lab Work	research, administration,maintenance and management of biologal artificals, plants, insects, animals, and biological systems	animals, research, animal, methods, environmental, study, instruct, diseases, perform, specimens, soil, disease, develop, development, treatment, activities, livestock, control, use, programs, clinical, techniques, examine, conditions, plan, health, laboratory, supervise, feed, plants	Plant Scientists (93%), Range Managers (73%), Pest Control Workers (70%)	Range Managers (73%) Animal Breeders (67%) Pest Control Workers (66%)
21	Public and Legal Interface	data collection and analysis in conjunction with law and regulatory items	data, information, prepare, maps, properties, survey, property, research, public, cases, gis, assessments, evidence, records, conduct, collect, including, surveys, analysis, documents, court, laws, search, regulations, use, land, law, photographs, related, sources	Title Examiners and Abstractors (98%),Title Searchers (93%),Highway Patrol Pilots(92%)	Lawyers (78%) Judges, Magistrate Judges, and Magistrates (75%) Appraisers, Real Estate (74%)
22	Non-Post Secondary Education	tasks for teaching and education broadly assocaited with primary and secondary education, provding and connecting to counselling and resources	students, academic, materials, activities, programs, social, plan, development, needs, progress, student, parents, educational, meet, children, teachers, curricula, required, behavior, use, evaluate, work, school, maintain, funding, organize, physical, policies, matters, objectives	Special Education Teachers, Preschool, Kindergarten, and Elementary School (79%), Special Education Teachers, Middle School (79%), Special Education Teachers, Secondary School (79%)	Special Education Teachers, Kindergarten and Elementary School (79%) Elementary School Teachers, Except Special Education (79%) Preschool Teachers, Except Special Education (79%)
23	Consulting	researches and assess situational contexts; makes recommendations for actions or solutions	inspect, construction, systems, structures, devices, buildings, drawings, installations, represent, operate, design, engineering, materials, fire_prevention, collect, requisition, wind, jobs, plans, codes, required, data, fires, inspections, supervise, damage, observe, ordinances, designs, computer	Fire-Prevention and Protection Engineers (76%), Civil Drafters (64%), Computer Programmers (56%)	Civil Drafters (63%) Computer Programmers (50%) Pile-Driver Operators (47%)
24	Corrective Action	advises and trains parties in corrective actions - connects to appropriate resources	maintain, provide, participate, conduct, train, course_materials, training, observe, topics, prepare, initiate, department, assistance, placement, career, departmental, community_events, handouts, reading_current, publish_findings, teach, services, participants, supervise, performance, plan, evaluate, demonstrations, seminars, document	Probation Officers and Correctional Treatment Specialists (69%) Soil Conservationists (58%) Fitness Trainers and Aerobics Instructors (58%)	Probation Officers and Correctional Treatment Specialists (68%) Fitness Trainers and Aerobics Instructors (57%) Training and Development Specialists (56%)
25	Demonstration and Usage	demonstrates products usage, assess production processes	products, production, design, product, materials, quality, testing, specifications, standards, methods, test, processing, material, develop, analyses, designs, manufacturing, models, process, chemical, conduct, monitor, solutions, engineers, quality_control, new,analyze, procedures, tests, samples	Food Scientists and Technologists (71%) Models (71%), Industrial Engineers (69%)	Industrial Engineers (66%) Food Scientists and Technologists (65%) Technical Writers (62%)
26	Physical Movement	performs intense physical movement, directs physical movement of others	cables, productions, logs, drive, winches, select, hand, loads, trucks, derricks, dance, load, position, pull, pulleys, lengths, cable, control, blocks, scores, direction, according, grade, members, start, trees, animals, trenches, weight, freight	Logging Tractor Operators (60%) Slaughterers and Meat Packers (60%) Fallers(54%)	Slaughterers and Meat Packers (60%) Fallers (56%) Logging Equipment Operators (54%)
27	Customers and Sales	prepares items calculate totals for customers	prepare, food, tables, customers, forms, dishes, orders, documents, order, collect, financial, revision, accounts, meals, preparation, trays, return, kitchen, according, counters, food_preparation, tax, serving, question, bookkeeping, returns, instructions, illnesses, required, salads	Counter Attendants, Cafeteria, Food Concession, and Coffee Shop (71%) Tax Preparers(66%) Food Servers, Nonrestaurant(65%)	Counter Attendants, Cafeteria, Food Concession, and Coffee Shop (67%) Food Servers, Nonrestaurant (65%) Cooks, Short Order (62%)
28	Biological Processing	handles, samples or prepares food and/or human and animal samples	prepare, materials, advise, food, maintain, supplies, clean, serve, shipping, serving, laboratory, media, work_assignments, areas, head, products, variety, record, procedures, store, work, cooking, sort, storage, conduct, containers, records, supervise, weight, weigh	Biochemists (99%) Meat, Poultry, and Fish Cutters and Trimmers(74%) Dental Assistants (73%)	Meat, Poultry, and Fish Cutters and Trimmers (74%) Dental Assistants (69%) Dentists, General (60%)
29	Item Processing	preparing and processing items for further downstream tasks	customers, merchandise, items, containers, inspect, hand, checks, transactions, goods, ships, computers, records, attach, stock, orders, information, sales, shelves, cash, carry, tickets, receive, packages, verify, payment, wrap, price, accounts, determine_priorities, tags	Hunters and Trappers (80%) Statement Clerks (67%) Stock Clerks, Sales Floor (64%)	Statement Clerks (70%) Stock Clerks, Sales Floor (64%) Marking Clerks (62%)

An analysis of the data indicates that the task structure of work has shifted over the study period. Each topic essentially represents a latent task variable, and each job is composed of a mixture of task variables. The topic model also indicates what percentage each task variable contributes to the overall task structure of the job. The charts provide some visual examples of how jobs have changed, as do the task descriptions in various years of plumber and agricultural inspector, in Figures III.11 through III.16.

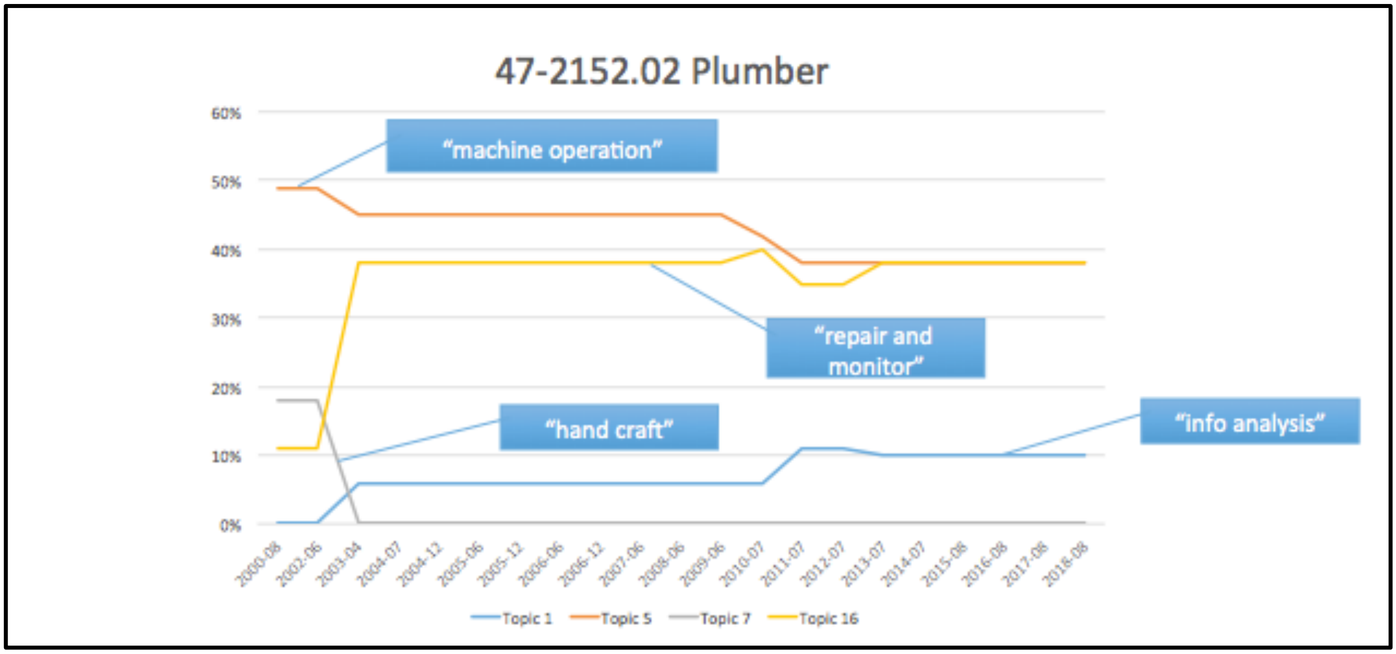

Figure III.9 Plumber Topic Plot 2000–2018

Assemble pipe sections, tubing and fittings, using couplings, clamps, screws, bolts, cement, plastic solvent, caulking, or soldering, brazing and welding equipment. Clear away debris in a renovation. Cut openings in structures to accommodate pipes and pipe fittings, using hand and power tools. Direct workers engaged in pipe cutting and preassembly and installation of plumbing systems and components. Fill pipes or plumbing fixtures with water or air and observe pressure gauges to detect and locate leaks. Hang steel supports from ceiling joists to hold pipes in place. Install oxygen and medical gas in hospitals. Install pipe assemblies, fittings, valves, appliances such as dishwashers and water heaters, and fixtures such as sinks and toilets, using hand and power tools. Install underground storm, sanitary and water piping systems and extend piping to connect fixtures and plumbing to these systems. Keep records of assignments and produce detailed work reports. Locate and mark the position of pipe installations, connections, passage holes, and fixtures in structures, using measuring instruments such as rulers and levels. Measure, cut, thread, and bend pipe to required angle, using hand and power tools or machines such as pipe cutters, pipe-threading machines, and pipe-bending machines. Perform complex calculations and planning for special or very large jobs. Prepare written work cost estimates and negotiate contracts. Repair and maintain plumbing, replacing defective washers, replacing or mending broken pipes, and opening clogged drains. Review blueprints and building codes and specifications to determine work details and procedures. Study building plans and inspect structures to assess material and equipment needs, to establish the sequence of pipe installations, and to plan installation around obstructions such as electrical wiring. Use specialized techniques, equipment, or materials, such as performing computer-assisted welding of small pipes, or working with the special piping used in microchip fabrication.

Figure III.10 Plumber Description From O*NET 2004

Anchor steel supports from ceiling joists to hold pipes in place. Assemble pipe sections, tubing, or fittings, using couplings, clamps, screws, bolts, cement, plastic solvent, caulking, or soldering, brazing, or welding equipment. Calculate costs or savings for water- or energy-efficient appliances or systems. Clear away debris in a renovation. Compile information on governmental incentive programs related to the installation of energy or water saving plumbing systems or devices. Cut openings in structures to accommodate pipes or pipe fittings, using hand or power tools. Determine sizing requirements for solar hot water heating systems, taking into account factors such as site orientation, load calculations, or storage capacity requirements. Direct helpers engaged in pipe cutting, preassembly, or installation of plumbing systems or components. Fill pipes or plumbing fixtures with water or air and observe pressure gauges to detect and locate leaks. Install alternative water sources, such as rainwater harvesting systems or graywater reuse systems. Install green plumbing equipment, such as faucet flow restrictors, dual-flush or pressure-assisted flush toilets, or tankless hot water heaters. Install oxygen and medical gas in hospitals. Install pipe assemblies, fittings, valves, appliances such as dishwashers or water heaters, or fixtures such as sinks or toilets, using hand or power tools. Install underground storm, sanitary, or water piping systems, extending piping as needed to connect fixtures and plumbing. Install, test, or commission solar thermal or solar photovoltaic hot water heating systems. Keep records of assignments and produce detailed work reports. Locate and mark the position of pipe installations, connections, passage holes, or fixtures in structures, using measuring instruments such as rulers or levels. Maintain or repair plumbing by replacing defective washers, replacing or mending broken pipes, or opening clogged drains. Measure, cut, thread, or bend pipe to required angle, using hand or power tools or machines such as pipe cutters, pipe-threading machines, or pipe-bending machines. Perform complex calculations and planning for special or very large jobs. Perform domestic plumbing audits to identify ways in which customers might reduce consumption of water or energy. Prepare written work cost estimates and negotiate contracts. Recommend energy or water saving products, such as low-flow faucets or shower heads, water-saving toilets, or high-efficiency hot water heaters. Review blueprints, building codes, or specifications to determine work details or procedures. Study building plans and inspect structures to assess material and equipment needs, to establish the sequence of pipe installations, and to plan installation around obstructions such as electrical wiring. Weld small pipes or special piping, using specialized techniques, equipment, or materials, such as computer-assisted welding or microchip fabrication.

Figure III.11 Plumber Description From O*NET 2011

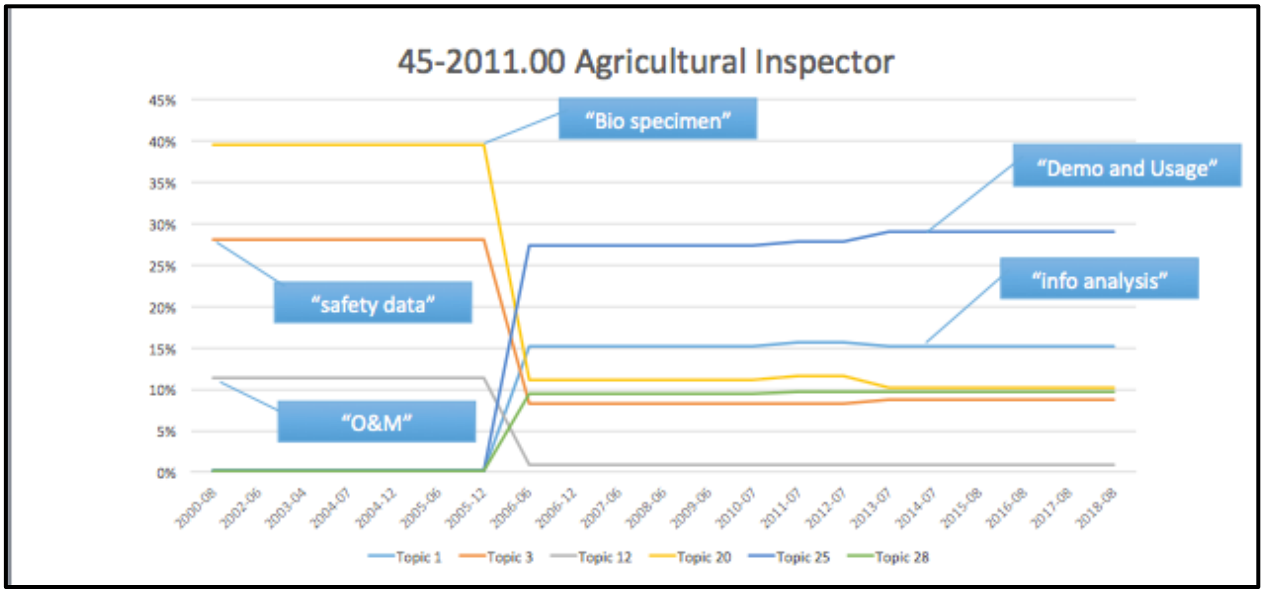

Figure III.12 Agricultural Inspector Topic Plot 2000–2018

Advises farmers and growers of development programs or new equipment and techniques to aid in quality production, applying agricultural knowledge. Collects sample of pests or suspected diseased animals or materials and routes to laboratory for identification and analysis. Examines, weighs, and measures commodities, such as poultry, eggs, meat, and seafood to certify wholesomeness, grade, and weight. Inspects facilities and equipment for adequacy, sanitation, and compliance with regulations. Inspects horticultural products or livestock to detect harmful disease, infestation or growth rate. Inspects livestock to determine effectiveness of medication and feeding programs. Testifies in legal proceedings. Writes reports of findings and recommendations and advises farmer, grower, or processor of corrective action to be taken.

Figure III.13 Agricultural Inspector Description From O*NET 2004

Advise farmers or growers of development programs or new equipment or techniques to aid in quality production. Collect samples from animals, plants, or products and route them to laboratories for microbiological assessment, ingredient verification, or other testing. Compare product recipes with government-approved formulas or recipes to determine acceptability. Direct or monitor the quarantine and treatment or destruction of plants or plant products.

Examine, weigh, and measure commodities, such as poultry, eggs, meat, or seafood to certify qualities, grades, and weights. Inquire about pesticides or chemicals to which animals may have been exposed. Inspect agricultural commodities or related operations, as well as fish or logging operations, for compliance with laws and regulations governing health, quality, and safety. Inspect food products and processing procedures to determine whether products are safe to eat. Inspect livestock to determine effectiveness of medication or feeding programs. Inspect or test horticultural products or livestock to detect harmful diseases, chemical residues, or infestations and to determine the quality of products or animals. Inspect the cleanliness and practices of establishment employees. Interpret and enforce government acts and regulations and explain required standards to agricultural workers. Label and seal graded products and issue official grading certificates. Monitor the grading performed by company employees to verify conformance to standards. Monitor the operations and sanitary conditions of slaughtering or meat processing plants. Provide consultative services in areas such as equipment or product evaluation, plant construction or layout, or food safety systems. Review and monitor foreign product inspection systems in countries of origin to ensure equivalence to the U.S. system. Set labeling standards and approve labels for meat or poultry products. Set standards for the production of meat or poultry products or for food ingredients, additives, or compounds used to prepare or package products. Take emergency actions, such as closing production facilities, if product safety is compromised. Testify in legal proceedings. Verify that transportation and handling procedures meet regulatory requirements. Write reports of findings and recommendations and advise farmers, growers, or processors of corrective action to be taken.

Figure III.14 Agricultural Inspector Description From O*NET 2011

The changes in individual jobs yield some interesting insights. For plumber, the topics associated with data analysis and with electric and mechanical repair both increased substantially over the time. That definitely aligns with the thought experiment that I proposed earlier in considering the complexity of a plumber's job. Agricultural inspector is surprising in one specific dimension: the topic associated with safety data declined over the period of the study. To the extent that agricultural inspectors play an important role in the safety of our food supply chain, this may be a cause for concern.

To examine changes in the global task structure of all of the jobs in the data set, I used paired t-tests to see how these variables changed over time. I performed this analysis on the task variables, as well as entropy measures, of all of the task variables for each job. This entropy measure can be used as a proxy for the diversity of the underlying task structure. The higher the entropy measure, the more diverse a job's task structure. Entropy as a diversity measure has commonly been in the corporate governance literature (Davis, Diekmann, & Tinsley, 1994), though applications of it in the job analysis literature are new. Entropy is calculated according to the formula for Shannon (1948) entropy $-\Sigma P \log2 (P)$, where P is the % each task topic contributes to the job. For example, a job that consists of one task would have entropy of zero. Alternatively, a job that consists of four tasks, each equally contributing 25% to the job, would have an entropy of 2.

The data set runs from 2000 to 2018. There are 628 jobs with consistent observations between 2000 and 2018, 722 jobs with consistent observations between 2004 and 2018, and 974 between 2011 and 2018. The 2004–2018 and the 2011–2018 ranges were also examined because 2000–2003 marked the beginning of the O*NET, and the data collection practices may have taken time to stabilize.

I calculated paired t-test statistics for each task variable in the 2000–2018, 2004–2018, and 2011–2018 periods. I also calculated t-statistics for the 628 companies in the 2000 cohort in the 2004–2018 period and the 2011–2018 periods. I determined that using paired t-tests on this data set was one of the few appropriate statistical tests available given the particulars of my variables. A paired t-test takes two observations of the same variable at different points in time and calculates the difference between them to see if that difference is significantly different from zero. One reason I used this particular test is that it is fairly robust to nonnormal data and data with some outliers, which are both attributes of this data set.

An explanation of the tables that contain the results is necessary because the usage of paired t-tests here is different from their more common usage. There is also simply a lot to process given the scale of the data. As mentioned above, paired t-tests are often used to compare the difference in a measurement at two points in time, especially after an experimental intervention. In my analysis, there is no experimental intervention, and my usage of the paired t-tests should not be seen as my performing an experiment. Rather, I am only using paired t-tests to see if there are *significant changes* (i.e., statistically different than zero) in measurement of the topic variables and the entropy variable at two different points in time.

In the statistical analysis tables, I have included three measurements for each variable: mean difference, the t-statistic, and the p-value. Mean differences represent the average change in the variable across the population at two different points in time. The t-statistic is the measured t-statistic for the given t-test. The p-value is the measured p-value associated with that t-statistic.

A challenge in calculating these measurements for every job is that not all jobs in the O*NET data set are measured every year. New jobs enter the data set in some years, while others exit the data set. To provide a *consistent basis* for comparison, I have selected the jobs for which I have data points across multiple years. This has resulted in five different "cuts" of the data: 1) the 2000 cohort– the 628 jobs in the O*NET data set for which I have consistent measurements from 2000 to 2018; 2) the 2000 cohort considered from 2004 to 2018; 3) the 2000 cohort considered from 2011 to 2018; 4) the 722 jobs for which I have consistent measurements from 2004 to 2018; 5) the 974 jobs for which I have consistent measurements from 2011 to 2018. The point of taking these different cuts of the data was threefold. First, it allows me to see if the task structure of jobs changes over the entire time of the data set. The second is that it allows me to see if some changes in the task structure of jobs are localized to a shorter time period. The third is that looking at different time periods allows me to account for some stabilization in the data collection methods that are used to collect the O*NET data.

Paired t-tests are comparisons between two different points in time. Each column of data therefore represents comparisons between two different years. In the first column, the comparison is between 2000 and 2018. In the second column, the comparison is between 2004 and 2018 (2000 cohort only). In the third column, the comparison is between 2011 and 2018 (2000 cohort only). In the fourth column, the comparison is between 2004 and 2018. Finally in the fifth column, the comparison is between 2011 and 2018. A guided example will be illuminating. For Topic 1, Information Analysis and Measurement, from 2000 to 2018, the average change in representation of the task variables in all jobs was a 2% increase. In other words, on average, information management and analysis became a greater portion of all jobs, and this change was statistically different from zero.

Because I examined essentially 31 variables (one entropy measure and 30 topic variables) across five different "cuts" of the data, I performed 155 t-tests. I used the Bonferroni adjustment and divided a significance level of .01 by 155, making the new level of significance .000065. The results of the analysis can be seen in Table III.4; three to six significant variables have been starred in each table, and significant p-values have been highlighted and bolded.

Table III.5 Statistical Analysis of Entropy and Topics 0-9

			2000-2018	2004-2018 (2000 Cohort)	2011-2018 (2000 Cohort)	2004-2018	2011-2018
Entropy**		Mean Difference	0.2215	0.1289	0.0298	0.1124	0.0220
		T-statistic	-11.7685	-11.4562	-6.4789	-11.1954	-6.8692
		P-Value	**0.0000**	**0.0000**	**0.0000**	**0.0000**	**0.0000**
Topic 0	Forming Objects	Mean Difference	0.0011	0.0012	-0.0001	0.0010	-0.0001
		T-statistic	-1.0446	-1.4232	0.9730	-1.4078	1.7183
		P-Value	0.2966	0.1552	0.3310	0.1596	0.0861
Topic 1**	Information Analysis and Management	Mean Difference	0.0202	0.0108	0.0020	0.0103	0.0017
		T-statistic	-5.7037	-7.0845	-4.1871	-7.3814	-4.4943
		P-Value	**0.0000**	**0.0000**	**0.0000**	**0.0000**	**0.0000**
Topic 2**	Stocking and Cleaning	Mean Difference	0.0079	0.0008	0.0001	0.0006	0.0000
		T-statistic	-5.1563	-3.1623	-0.9124	-2.7929	-0.6145
		P-Value	**0.0000**	0.0016	0.3619	0.0054	0.5390
Topic 3	Safety Data	Mean Difference	-0.0018	-0.0007	0.0000	-0.0007	-0.0001
		T-statistic	1.2395	0.7218	0.1629	0.8578	0.3136
		P-Value	0.2156	0.4707	0.8707	0.3913	0.7539
Topic 4	Exterior Maintenace	Mean Difference	0.0046	0.0006	0.0001	0.0004	0.0001
		T-statistic	-2.8716	-0.5281	-0.5234	-0.3860	-0.3701
		P-Value	0.0042	0.5976	0.6009	0.6996	0.7114
Topic 5	Machine Operation	Mean Difference	-0.0035	-0.0013	0.0007	-0.0010	0.0006
		T-statistic	2.0236	1.0632	-3.6328	0.9504	-3.8174
		P-Value	0.0434	0.2881	0.0003	0.3422	0.0001
Topic 6	Constructing Things	Mean Difference	-0.0014	-0.0014	-0.0004	-0.0012	-0.0003
		T-statistic	1.3289	1.4868	0.9894	1.4835	0.9290
		P-Value	0.1844	0.1376	0.3228	0.1384	0.3531
Topic 7**	Detailed Handcraft	Mean Difference	-0.0075	-0.0023	0.0002	-0.0020	0.0002
		T-statistic	5.3036	2.1014	-1.1022	2.1271	-1.6105
		P-Value	**0.0000**	0.0360	0.2708	0.0337	0.1076
Topic 8	Healthcare work	Mean Difference	0.0061	0.0009	0.0004	0.0007	0.0002
		T-statistic	-3.0966	-1.1184	-1.3434	-1.0667	-0.9588
		P-Value	0.0020	0.2638	0.1796	0.2865	0.3379
Topic 9	Coordinating and Liaising	Mean Difference	0.0018	-0.0006	-0.0002	-0.0004	-0.0001
		T-statistic	-1.0430	0.7514	0.8822	0.6329	0.8704
		P-Value	0.2973	0.4527	0.3780	0.5270	0.3843

Table III.6 Statistical Analysis of Topics 10-19

			2000-2018	2004-2018 (2000 Cohort)	2011-2018 (2000 Cohort)	2004-2018	2011-2018
Topic 10**	Detailed Work	Mean Difference	0.0179	-0.0070	-0.0032	-0.0076	-0.0025
		T-statistic	-10.5240	5.5720	5.2636	6.6328	6.0605
		P-Value	**0.0000**	**0.0000**	**0.0000**	**0.0000**	**0.0000**
Topic 11	Shaping Modification	Mean Difference	-0.0024	-0.0012	-0.0004	-0.0011	-0.0004
		T-statistic	2.1464	1.5498	2.0218	1.6297	2.3758
		P-Value	0.0322	0.1217	0.0436	0.1036	0.0177
Topic 12**	Operation and Maintenance	Mean Difference	-0.0024	-0.0012	-0.0004	-0.0011	-0.0004
		T-statistic	24.5946	7.0100	0.8779	6.8342	0.7981
		P-Value	**0.0000**	**0.0000**	0.3803	**0.0000**	0.4250
Topic 13	Human Movers	Mean Difference	-0.0015	0.0004	0.0000	0.0006	0.0001
		T-statistic	0.8555	-0.7256	0.1103	-1.0766	-0.5127
		P-Value	0.3926	0.4684	0.9122	0.2820	0.6083
Topic 14	Creative Decisioning	Mean Difference	0.0011	-0.0001	0.0000	0.0000	0.0001
		T-statistic	-0.7798	0.0914	-0.0225	0.0202	-0.5325
		P-Value	0.4358	0.9272	0.9820	0.9839	0.5945
Topic 15	Information Collection	Mean Difference	0.0007	0.0006	0.0004	0.0003	0.0001
		T-statistic	-0.3758	-0.8973	-1.4224	-0.4737	-0.3372
		P-Value	0.7072	0.3699	0.1554	0.6359	0.7360
Topic 16**	Repair and Monitoring	Mean Difference	0.0096	0.0062	0.0001	0.0055	0.0000
		T-statistic	-4.3254	-4.1404	-0.2828	-4.1473	-0.1438
		P-Value	**0.0000**	**0.0000**	0.7774	**0.0000**	0.8857
Topic 17**	Post-Secondary Education	Mean Difference	-0.0122	-0.0008	0.0000	-0.0008	0.0000
		T-statistic	4.9854	1.2989	-0.0964	1.5232	0.2493
		P-Value	**0.0000**	0.1944	0.9232	0.1282	0.8032
Topic 18**	Engineering Management	Mean Difference	0.0055	0.0005	0.0000	0.0005	0.0000
		T-statistic	-4.8938	-1.3838	0.1191	-1.3338	-0.3275
		P-Value	**0.0000**	0.1669	0.9052	0.1827	0.7434
Topic 19**	High Technical Systems	Mean Difference	0.0069	0.0011	0.0002	0.0009	0.0001
		T-statistic	-5.5677	-4.7544	-2.1287	-3.8360	-0.4155
		P-Value	**0.0000**	**0.0000**	0.0337	0.0001	0.6779

Table III.7 Statistical Analysis of Topics 20-29

			2000-2018	2004-2018 (2000 Cohort)	2011-2018 (2000 Cohort)	2004-2018	2011-2018
Topic 20	Lab Work	Mean Difference	-0.0003	-0.0010	-0.0001	-0.0009	-0.0001
		T-statistic	0.2125	1.3806	0.5039	1.3184	0.3194
		P-Value	0.8318	0.1679	0.6145	0.1878	0.7495
Topic 21	Public and Legal Interface	Mean Difference	-0.0057	-0.0011	-0.0003	-0.0009	-0.0001
		T-statistic	3.0306	1.0338	1.2443	1.0389	0.7473
		P-Value	0.0025	0.3016	0.2139	0.2992	0.4550
Topic 22	Non-Post Secondary Education	Mean Difference	0.0001	0.0003	0.0000	0.0004	0.0000
		T-statistic	-0.0735	-0.4690	0.1716	-0.6529	0.1177
		P-Value	0.9415	0.6393	0.8638	0.5140	0.9064
Topic 23**	Consulting	Mean Difference	0.0087	0.0009	-0.0001	0.0008	-0.0002
		T-statistic	-6.8048	-1.3110	0.5696	-1.3557	1.5148
		P-Value	**0.0000**	0.1903	0.5692	0.1756	0.1301
Topic 24**	Corrective Action	Mean Difference	0.0125	0.0009	0.0003	0.0008	0.0003
		T-statistic	-7.1987	-1.4854	-1.5824	-1.5308	-1.8558
		P-Value	**0.0000**	0.1379	0.1141	0.1263	0.0638
Topic 25	Demonstration and Usage	Mean Difference	0.0002	0.0025	0.0004	0.0025	0.0004
		T-statistic	-0.1571	-2.8651	-1.3035	-3.2947	-1.7172
		P-Value	0.8752	0.0043	0.1929	0.0010	0.0863
Topic 26	Physical Movement	Mean Difference	0.0024	0.0015	0.0001	0.0012	0.0000
		T-statistic	-1.8076	-2.2886	-1.0282	-2.0776	-0.4336
		P-Value	0.0712	0.0224	0.3042	0.0381	0.6646
Topic 27	Customers and Sales	Mean Difference	0.0046	0.0002	-0.0003	0.0002	-0.0002
		T-statistic	-2.8960	-0.4124	1.1620	-0.6800	0.8143
		P-Value	0.0039	0.6802	0.2457	0.4967	0.4157
Topic 28**	Biological Processing	Mean Difference	0.0090	0.0016	0.0003	0.0016	0.0001
		T-statistic	-5.1471	-2.7510	-1.3985	-3.0142	-0.7517
		P-Value	**0.0000**	0.0061	0.1625	0.0027	0.4524
Topic 29	Item Processing	Mean Difference	0.0037	0.0003	0.0001	0.0003	0.0001
		T-statistic	-2.9760	-0.3806	-0.3686	-0.4087	-0.7228
		P-Value	0.0030	0.7036	0.7125	0.6829	0.4700

Table III.8 Summary of Statistically Significant Findings

			2000-2018	2004-2018 (2000 Cohort)	2011-2018 (2000 Cohort)	2004-2018	2011-2018
Entropy		Directionality	Increasing	Increasing	Increasing	Increasing	Increasing
		p < .000065	Significant	Significant	Significant	Significant	Significant
Topic 1	Information Analysis and Management	Directionality	Increasing	Increasing	Increasing	Increasing	Increasing
		p < .000065	Significant	Significant	Significant	Significant	Significant
Topic 2	Stocking and Cleaning	Directionality	Increasing				
		p < .000065	Significant				
Topic 7	Detailed Handcraft	Directionality	Decreasing				
		p < .000065	Significant				
Topic 10	Detailed Work	Directionality	Increasing	Decreasing	Decreasing	Decreasing	Decreasing
		p < .000065	Significant	Significant	Significant	Significant	Significant
Topic 12	Operation and Maintenance	Directionality	Decreasing	Decreasing		Decreasing	
		p < .000065	Significant	Significant		Significant	
Topic 16	Repair and Monitoring	Directionality	Increasing	Increasing		Increasing	
		p < .000065	Significant	Significant		Significant	
Topic 17	Post-Secondary Education	Directionality	Decreasing				
		p < .000065	Significant				
Topic 18	Engineering Management	Directionality	Increasing				
		p < .000065	Significant				
Topic 19	High Technical Systems	Directionality	Increasing	Increasing			
		p < .000065	Significant	Significant			
Topic 23	Consulting	Directionality	Increasing				
		p < .000065	Significant				
Topic 24	Corrective Action	Directionality	Increasing				
		p < .000065	Significant				
Topic 28	Biological Processing	Directionality	Increasing				
		p < .000065	Significant				

Explanation of Topic Model Findings

This section describes all of the findings from the statistical analysis based on the LDA constructed variables. Because of the voluminous amount of data associated with the chapter, I have included a summary table, Table III.8, of the statistically significant findings. For each variable, I have included the directionality of the change of the variable, whether it increased or decreased over the period of time. In addition, I have included whether the p-value was significant or nonsignificant, which shows whether the directionality of change is statistically different from zero. Again, an example will be an aide. In the case of Topic 1, informational management and analysis, the table shows that average representation of this topic increased between 2000 and 2018 as indicated by the word increased in the directionality row. In addition, we know that this change is statistically different from zero, as indicated by the word "significant" in the $p < .000066$ column row of the data table. Below, I offer an interpretation of each statistically significant finding.

Entropy. The entropy measure has a significant p-value across all years. This means that the mean difference between entropy of the task variables in 2000 and all years is significantly different from zero. The mean difference between 2000 and all years of the data is positive. This suggests that, on average, the task structure of jobs is getting more diverse over time, in all cohorts of data.

Topic 1: Information management and analysis. Topic 1 has a significant p-value across all years. This means that the mean difference between Topic 1 measured in 2000 and all years is significantly different from zero. The mean difference in Topic 1 between 2000 and all years of the data is positive. This suggests that, on average, information management and analysis is becoming a greater portion of job tasks over time, in all cohorts of data.

Topic 2: Stocking and cleaning. Topic 2 has a significant p-value for the time period 2000–2018. This means that the mean difference between Topic 2 measured in 2000 and 2018 is significantly different from zero. The mean difference in Topic 2 between 2000 and 2018 is positive. This suggests that, on average, stocking and cleaning is becoming a greater portion of job tasks over time for the 628 jobs for which I have data over the 2000–2018 time frame.

Topic 7: Detailed handcraft. Topic 7 has a significant p-value for the time period 2000–2018. This means that the mean difference between Topic 7 measured in 2000 and 2018 is significantly different from zero. The mean difference in Topic 7 between 2000 and 2018 of the data is negative. This suggests that, on average, detailed handcraft is becoming a smaller portion of job tasks over time in the 628 jobs for which I have data over the 2000–2018 time frame.

Topic 10: Detailed work. Topic 10 has a significant p-value across all years. This means that the mean difference between Topic 1 measured in 2000 and all years is significantly different from zero. The mean difference in Topic 10 between 2000 and 2018 is positive. This suggests that, on average, over the 2000–2018 period detailed work has increased. But as the mean difference is negative across all other periods, it also suggests that detailed work became a smaller component of work, on average, over a shorter time frame.

Topic 12: Operation and maintenance. Topic 12 has a significant p-value for the time periods 2000–2018 and 2004–2018. This means that the mean difference between Topic 12 measured in these periods is significantly different from zero. However, the mean difference in Topic 12 appears to be negative over these periods, suggesting that, on average, operation and maintenance as a part of job tasks is declining.

Topic 16: Repairing and monitoring. Topic 16 has a significant p-value for the time periods 2000–2018 and 2004–2018. The mean difference between Topic 16measured in these periods is significantly different from zero. However, the mean difference in Topic 16 is positive over these periods, suggesting that, on average, repairing and monitoring as a part of job tasks is increasing.

Topic 17: Postsecondary education. Topic 17 has a significant p-value for the time period 2000–2018. This means that the mean difference between Topic 17 measured in 2000 and 2018 is significantly different from zero. The mean difference in Topic 17 between 2000 and 2018 is negative, this suggests that, on average, postsecondary education is becoming a smaller portion of job tasks over time.

Topic 18: Engineering management. Topic 18 has a significant p-value for the time period 2000–2018. This means that the mean difference between Topic 18 measured in 2000 and 2018 is significantly different from zero. The mean difference in Topic 18 between 2000 and 2018 is positive. This suggests that, on average, engineering management is becoming a greater portion of jobs tasks over time for the 628 jobs for which I have data over the 2000–2018 time frame.

Topic 19: High technology. Topic 19 has a significant p-value for the time period 2000–2018. This means that the mean difference between Topic 19 measured in 2000 and 2018 is significantly different from zero. The mean difference in Topic 19 between 2000 and 2018 is positive. This suggests that, on average, high technology is becoming a greater portion of job tasks over time for the 628 jobs for which I have data over the 2000–2018 time frame.

Topic 23: Consulting. Topic 23 has a significant p-value for the time period 2000–2018. This means that the mean difference between Topic 23 measured in 2000 and 2018 is significantly different from zero. The mean difference in Topic 23 between 2000 and 2018 is positive, this suggests that, on average, consulting is becoming a greater portion of job tasks over time for the 628 jobs for which I have data over the 2000–2018 time frame.

Topic 24: Corrective action. Topic 24 has a significant p-value for the time period 2000–2018. This means that the mean difference between Topic 24 measured in 2000 and 2018 is significantly different from zero. The mean difference in Topic 24 between 2000 and 2018 is positive. This suggests that, on average, corrective action is becoming a greater portion of job tasks over time for the 628 jobs for which I have data over the 2000–2018 time frame.

Topic 28: Biological processing. Topic 28 has a significant p-value for the time period 2000–2018. This means that the mean difference between Topic 28 measured in 2000 and 2018 is significantly different from zero. The mean difference in Topic 28 between 2000 and 2018 is positive, this suggests that, on average, biological processing is becoming a greater portion of job tasks over time for the 628 jobs for which I have data over the 2000–2018 time frame.

Interpretation of Topic Model Findings

These findings say some unique things about how jobs have changed. Some of the findings suggest trends that are in line with some commonly held expectations, while others do not lend themselves to easy interpretation. The first set of findings that I would group together consist of the findings with the highest level of statistical significance across periods: entropy, information analysis and management, and detailed work. Overall the findings suggest that, over time, jobs have an increasingly diverse task structure. One way to interpret this is that, on average, the number of tasks in the task stack of jobs is increasing. While this does not tell us specifically what jobs require an increasingly diverse suite of tasks, we know that, on average, for the jobs for which O*NET collects data, the aggregate descriptions of these jobs are becoming more diverse in their task representation. This is the one finding that I believe to be very robust as the p-value is so low, and it is consistent across all of the time frames. It is for these same reasons that I also find the results for information management and analysis and detailed work to be very robust as well. The increase in information management and analysis appropriately reinforces the trends that data analysis has become more and more important over time and have led some to call data scientist the sexiest job in the 21st century (Davenport & Patil, 2012). However, my findings suggest that the importance of data analysis has long since bled over into other jobs and is not just the purview of data scientists. The decline in detailed work as part of jobs is interesting and resists direct interpretation. This surprising finding appears to be robust across all time periods. It is possible that while detailed work is important, if it does involve repetitive action, as the thematic analysis suggests, some of this work may have been automated. That might account for a decline in this task in the task stack of various jobs.

The next set of findings that I would put in a similar category is high technology, repair and monitoring, and operations and maintenance. These findings are only appear to be significant in the 2000–2018 and the 2004–2018 time frames. For high technology tasks, this suggests that the largest part of the shift of jobs to high technology – at least for the jobs for which O*NET collected data – happened in the early part of the 21st century. Those jobs do not include those that have been created by new technology. Thus the finding shows that for most jobs that have existed since at least 2004, , most of the high technological tasks have been added. Repair and monitoring and operation and maintenance are unique. They are clearly related, but have slightly different meanings. Repair and monitoring has increased over time whereas operation and maintenance has decreased. This may be simply an artifact of the descriptive language associated with these tasks. However, it may reflect the decrease in the useful life of objects. The tools people use in their jobs are increasingly disposable and replaceable, much like consumer goods.

The last set of findings I would put into a similar category are the findings for stocking and cleaning, detailed handcraft, postsecondary education, engineering management, consulting, corrective action, and biological processing. While the p-values are very small, the changes can only be seen over the 2000–2018 time frame. I recognize that in the early days of O*NET the collection process may not have been as robust or the reporting and descriptive procedures standardized. As such, my findings may simply be an artifact of the peculiarities of the time period. However, they may also reflect real shifts over time. The most interesting is that, for 20 jobs that experience the largest increase in the representation of the corrective action job task, 18 of them experienced decreases in the postsecondary education task, as can be seen in Table III.9. While this is not necessarily a one-for-one increase, there may be a societal impact of jobs once carrying a significant education task now focusing to a greater degree on corrective action.

Table III.9 Postsecondary and Corrective Action Crossover

Title	Topic 17 Differences	Topic 24 Differences
Probation Officers and Correctional Treatment Specialists	-0.00014	0.35405
Health Educators	-0.09720	0.23325
First-Line Supervisors/Managers of Helpers, Laborers, and Material Movers, Hand	-0.00024	0.20511
Musicians, Instrumental	-0.00039	0.18107
English Language and Literature Teachers, Postsecondary	-0.38049	0.17217
Loan Counselors	-0.34302	0.15640
Graduate Teaching Assistants	-0.08562	0.15425
Nursing Instructors and Teachers, Postsecondary	0.55046	0.15010
Clergy	-0.00062	0.14491
Psychiatric Aides	-0.00039	0.13882
Child, Family, and School Social Workers	0.03103	0.13711
Health Specialties Teachers, Postsecondary	-0.27513	0.13353
Economics Teachers, Postsecondary	-0.25409	0.13133
Foreign Language and Literature Teachers, Postsecondary	-0.32014	0.13020
Political Science Teachers, Postsecondary	-0.26599	0.12889
Art, Drama, and Music Teachers, Postsecondary	-0.32897	0.12631
Engineering Teachers, Postsecondary	-0.25177	0.12523
Area, Ethnic, and Cultural Studies Teachers, Postsecondary	-0.28177	0.12235
Agricultural Sciences Teachers, Postsecondary	-0.27148	0.12233
Electrical Engineering Technicians	-0.00034	0.12139

Preliminary Embedding Model Findings

Embedding models create more precise mathematical representations of text than do topic models. This is because embedding models can take word context into account. But while the mathematical representation is arguably more precise, it does not lend itself well to direct human interpretation because the output of an embedding model is a very large matrix, in this case a 100-dimensional matrix. And unlike topic modeling, these matrices do not have any identifying features that would allow me to create particular labels for each of the matrix elements. So while this work does not yet lend itself to statistical analysis, I can present some preliminary findings from the embedding models by visualizing the results.

The embedding model results lend themselves very well to visualizations. The representation of the task structure of jobs as a 100-dimension embedding vector model allows for fairly precise mapping of jobs in a two-dimensional representation, using t-SNE to compress the data such that it is viewable. I will provide links to the visualizations and give brief summaries of what they display. Also, where possible I will provide samples of the visualization that display some key points, and I will provide written descriptions of the imagery.

Visualization 1 - Job Task Similarity [http://bit.ly/31QPzUs]

The visualization referenced above, made using Bokeh, shows a two-dimensional visualization of the job task descriptions for the 2015 version of O*NET. Each point represents a different job. The shape/color of each point represents a different category of job as specified by broad occupational categories. For example, the pink triangles are all jobs that fall into education, training, and library occupations. The black circles are management occupations. All of the various broad occupational categories can be found in the legend on the right side of the diagram. Clicking on the occupational categories will toggle that category of data on and off from view in the diagram.

The way to interpret this diagram is simple. If two points are close to each other, the underlying descriptive language of the job tasks is similar. If the points are far apart, the underlying language of the job tasks is not similar. The axes as currently envisioned resist direct interpretation. They are best understood as the principle components of the data set. These are the two dimensions that capture the most variance of the data set while being orthogonal to one another.

Key Point 1 - Education Jobs are Bifurcated

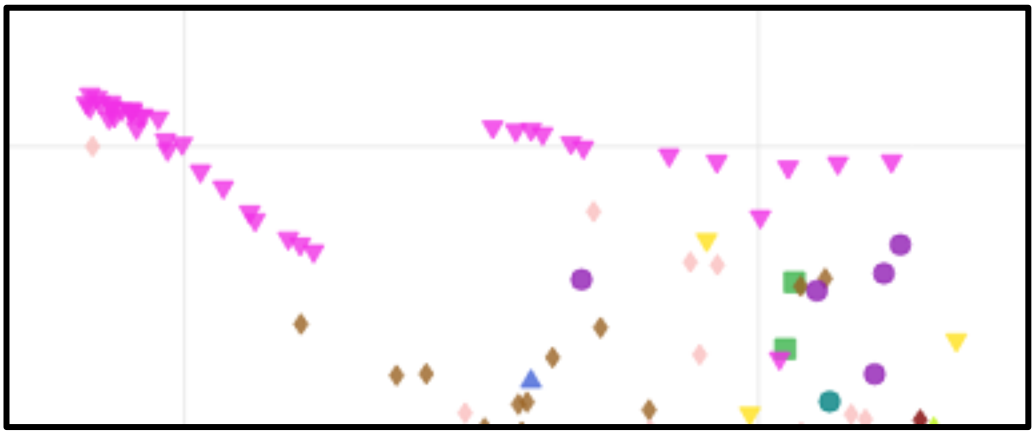

Figure III.15 Educational Occupational Excerpt From Static Web Visualization

Figure III.15 is an excerpt from the broader map described above. On the left side of the figure, the cluster of pink triangles represents postsecondary education jobs or college professors. On the right side of the diagram, the cluster of pink triangles represents elementary and high school instructors. Those two groups of teachers are similar internally, with postsecondary teachers showing extreme similarity, as is evident in the close proximity of the points. But those two groups of teachers differ from each other. An interesting question this raises is, why would this be the case? Some high school students are only a few months younger than the youngest college students, so why the large difference in underlying tasks? Perhaps occupational legitimacy (Abbott, 2014) plays a role in this difference.

Key Point 2 - Is a Supply Chain Manager Actually a Manager?

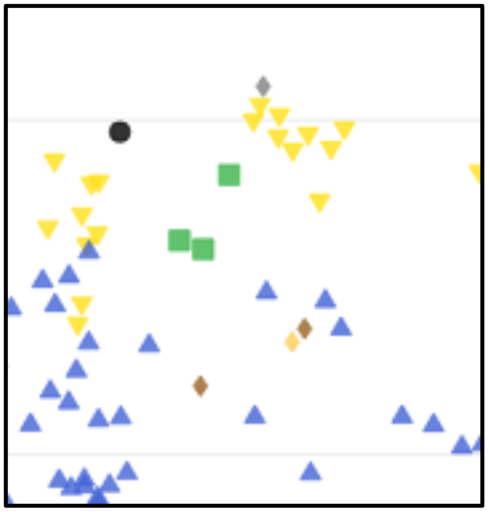

Figure III.16 Supply Chain Manager Excerpt From Static Web Visualization

Figure III.16 is another excerpt from the web visualization. The black circle in the upper left-hand corner is the job of supply chain manager. O*NET classifies it as a management occupation. However, there are no other management occupations nearby. All of the yellow triangles that are close to it are computer and mathematical occupations such as web developers and software developers. And the green squares are business and financial occupations, specifically logistics engineers, logistics analysts, and security management specialists. This suggests that, at least as of 2015, a supply chain manager has less in common with management and does more tasks near the intersection of logistics and technology. A potentially interesting course of study would be to understand how the role of supply chain manager straddles the logics of being ostensibly management-oriented, but practically technologically oriented, as well as how this particular job has evolved over time.

Visualizations are essentially forms of exploratory data analysis. They show which jobs have similar underlying task structures, and then use that data as motivation to probe further into the particularities of those occupations. Job recruiters could use maps such as this to evaluate the applicability of job experience. For example, an organization hiring a supply chain manager might benefit from considering a logistics analysis or security management specialist, given the underlying apparent similarity among these roles.

Visualization 2 - A Gapminder-Inspired Job Viewer [http://bit.ly/2XAZ633]

Hans Rosling (2011) is famous for his TED talk showing his Gapminder visualization that shows how the world's countries have progressed along certain dimensions such as maternal mortality, literacy, and GDP growth. The fame is rightly deserved because it shows the relationship between countries both in terms of these key statistics and, most importantly, how they have moved over time.

Inspired by Gapminder, I have attempted to do something similar with my embedding model data. Using all of the job task data in ONET from 2004 to 2017, I created a Gapminder-like visualization, which shows jobs moving in relation to one another over time. Each job is represented by a circle, the color of the circle represents the broad O*NET occupation classification, and the size of the circle represents the average salary of that job in each year as per the Bureau of Labor Statistics.

Key Point 3 – Difference and Similarities, Within and Between Food Service and Management Occupations

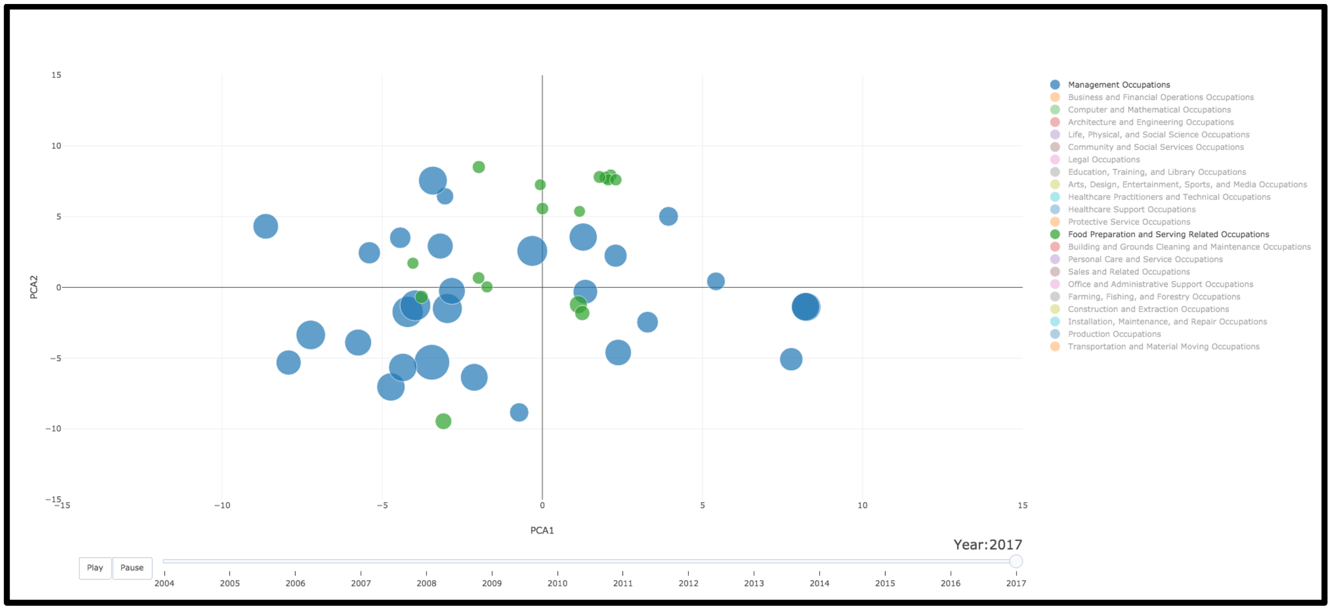

Figure III.17 Excerpt From Animated Web Visualization

Figure III.17 excerpted from the web visualization is a useful example of the concept, illustrative of the visualizations features, while also showing how such a figure can be useful to quickly see the differences between jobs. Much like the other diagram, this one has a legend on the right-hand side that allows users to toggle the visibility of the broad occupational categories, and the proximity of the jobs to one another indicates the similarity of the underlying task structures as a function of the words used in the task descriptions. In this figure, the blue circles are management occupations and the green circles are food preparation and serving related occupations. The data is from the year 2017. The slider bar will show the relative position of jobs in each year, and the slider bar can be set to automatic to show all of the jobs moving around the landscape. The static image above, of course, does not reveal this, but the movement in the full animation shows that the underlying structure of jobs – and which ones are close to which others – changes over time.

As the size of the circles represents the average salary of that job, the large difference in salary between management occupations and food service occupations is apparent. The difference between restaurant cooks and private household cooks is also of interest. The latter is represented by the green dot closest to the top of the plot, and the former by the green dot closest to the bottom of the plot, both on the left side of the axis. This suggests that the term "cook" obscures many differences between these roles. Perhaps even more surprising, three intersecting points, two green and one blue, near the x-axis just to the right of the y-axis, represent chef and head cooks, supervisors of food preparation, and funeral service managers. Why these would be close to one another certainty warrants further investigation.[24]

Discussion

This chapter has examined how jobs are structured and how they change over time. The research produces four main takeaways not otherwise present in the literature.

Jobs can be described in terms of their underlying task dimensions. The research supports the claim that jobs have a task structure. By taking the task descriptions of jobs, it is possible, using topic modeling, to decompose them into categories of latent task variables of which all jobs are composed. In this particular analysis, I revealed 30 underlying task variables of jobs. In addition, topic modeling allows me to turn the text-based task description of jobs into 30 task variables that can be measured quantitatively. Each job in the O*NET data set is represented by a subset of the 30 task variables, and topic modeling indicates, in terms of percentages, how much each task variable contributes to a given job. This insight will enable the quantitative measurement of job task variables I use in subsequent chapters of this dissertation.

[24] I *really* hope that these are similar only because they all have management aspects.

Jobs have changed over time. Because I have essentially turned qualitative task descriptions into quantitative task variables, I can use some simple visualization methods and traditional statistical methods to analyze these variables. On the individual job level, I plotted the task variables over time, and demonstrated how certain tasks increased and decreased within a job. This was displayed with plumber and agricultural inspector. Looking at jobs at the aggregate level, I used paired t-tests to see how the average representation of each task variables changed over time among the jobs I analyzed. This showed that certain task variables such as information analysis and management have increased, on average, from 2000 to 2018. Additionally, it displayed that certain task variables such as postsecondary education have decreased over the same period. Because I can show these changes at both the individual job and the aggregate job level, I assert that this demonstrates that jobs are changing over time.

Jobs are becoming more diverse over time. Not only are the underlying task variables of jobs changing over time, but the representation of job tasks is also becoming more diverse. Using entropy, I measured the diversity of the task variables in each job in each year in the data set. Then using the paired t-tests again, I examined whether, at the aggregate level, task diversity changed between 2000 and 2018. What this shows is that, on average, the entropy measure is increasing, which means that jobs are getting more diverse over time, in terms of how many underlying tasks make up each job. It seems possible that automation could have produced this trend in one of two ways. First, workers and/or managers may have engineered increasing job diversity as a response to automation pressures. Incorporating more tasks into any given job is generally going to make it less likely that technology will automate every single one of those tasks away. Secondly and alternatively, the increasing diversity could be a case of survivorship bias. It may be that jobs with a very simple task structure have already been automated away, and hence they have disappeared over time.

Jobs cluster in (unexpected) neighborhoods. By using embedding models to measure the underlying task structure of jobs, I can create a more precise mathematical representation of the underlying task structure of jobs. Specifically, I created a 100-dimensonal-vector model of jobs, which I then compressed along two principal components to allow for the creation of visualizations. These visualizations show us that jobs both exist clustered together in specific neighborhoods, and jobs may cluster together in ways that we would not expect. In the examples above, I showed that teachers are largely similar to one another, but dissimilar to everything else. Restaurant cook and private household cook are quite dissimilar to one another despite both being nominally cooks. Similarly, supply chain managers have very little in common with other managers. A potential next step of this work is to examine the dynamics of certain job neighborhoods over time. For example, do low wage jobs cluster together in such a way that holders of these jobs have difficulty escaping these "job neighborhoods"? And by extension, could this visualization essentially provide a map for traversing the labor landscape? There may well be certain jobs that serve as key intermediaries for eventually progressing to more remunerative work. Though the findings from this section are preliminary, there are many potential avenues for future research.

This work has three main limitations. The first is the relatively circumscribed nature of the data set. At its widest aperture, O*NET has approximately 1,000 jobs in its database. While this might provide an accurate representation of the broad categories of jobs throughout the economy, the number simply feels a little small. Deeper insight may come from a more expansive data set. For example, a collection of actual job postings over the past 20 years might provider a deeper insight into the underlying task structure of jobs. The second limitation is that while I hypothesize about the potential mechanisms for why jobs might change in order to provide the theoretical legitimacy for this study, I did not test any of these mechanisms directly in this analysis. The third limitation is embedded in the natural language processing methods. Coherence model methodology, an established practice, supports the use of the 30 dimensions that I used for the topic analysis. But 30 may not be the ideal number of tasks for understanding the task structure of jobs. For example the tasks variables associated with healthcare and postsecondary education both seem as if they could reduced further. As for the embedding model, the choice of 100 dimensions is standard practice in the world of embedding model scholarship, but this is largely an arbitrary choice.

Future research could address the second of these limitations in depth. The Current Population Survey (CPS) collects data about the U.S. labor force, including which jobs people are moving between. If one of the consequences of people moving between jobs is change to the underlying structure of those jobs, then these transitions should be correlated to the measured changes in jobs over time. This would require harmonizing the O*NET data set with the CPS. A second line of research would be to try to measure how the economy values certain tasks within jobs over time. Given that average salary data is available, it should be possible to measure relationships between changes in salary and the presence of certain job tasks over time. A last line of research would be trying different specifications of these natural language processing models to see how robust these findings are to different model specifications. For example, are the results comparable using a 20 or 60 topic model? And would the embedding visualizations look largely the same using 50 or 300 dimensions to mathematically represent the jobs tasks?

Some sensitivity analysis around those model specifications could be very useful to help these methods gain further traction in organizations and management research.

It is my hope that this work demonstrating the viability of this method for analyzing job task data and for illuminating the underlying structure of jobs will be a foundation for many avenues of research in the coming years.

CHAPTER IV

Case Study

Introduction

This chapter presents a case study of a start-up for two reasons. First, increasingly it is small firms and entrepreneurial start-ups (Haltiwanger et al., 2013) that provide new jobs, not large publicly traded organizations (Davis, 2016). Second, start-ups themselves may be a fertile lab for deeper understanding and potential revision of some of our deeply held theories about organizations studies. Start-ups are ground zero for the transformation wrought by technology (Brynjolfsson & McAfee, 2014).Technology has the potential to reconfigure the landscape of work (Chui et al., 2016), but we do not know much about the process of actually constructing new jobs around novel technology and the potential mechanisms for job construction for a new organization that utilizes novel technology in its daily practices.

The start-up featured here is located in southeastern Michigan and is using proprietary guided workflow software and applying it to the traditional business of commercial cleaning. This company provides a truly unique case for examination (in line with Eisenhardt & Graebner, 2007) because, while there has been qualitative work on entrepreneurial processes (e.g., Baron et al., 1996; Greenwood & Suddaby, 2006), the impact of novel technology in well-established fields (e.g., Beane, 2018; Zuboff, 1988) and “dirty work” such as janitorial services (e.g., Ashforth & Kreiner, 1999), no study has looked at all of these in one setting. In addition to providing new insight into the construction of jobs in relation to technology, this study will illuminate the challenges the company's founders faced in trying to construct a company that resides between the different logics of high-tech organizations and blue-collar traditional organizations. This study sheds light on how entrepreneurs consider and deal with the problems they all face: the uncertainty about how their product will be accepted in the market, and the execution risk associated with creating a new enterprise.

Context

The new world of work is here, if we judge by the plethora of articles, books (Kessler, 2018), and IPOs (Lyft, Uber, and Postmates just to name a few). But I believe that this so-called new world of work actually is emerging from trends that started some time ago. A very quick, data-driven, but rather unscientific use of Google Trends suggests that searches for the phrase "new world of work" actually peaked in 2004. What this may indicate is that the context in which work is situated has been changing for a long time, and that this has implications for how work has continued to develop. It seems that for at least a decade and a half the answer to a basic question about work has been changing: "Who do I work for?"

Very large organizations used to be the answer to the question. Through much of the 20th century, the story of jobs was really a story of dominance of the economic landscape by large organizations. Work primarily took place within the confines of a large, publicly traded corporation, and many people spent their entire working lives at one such corporation. Many of the classic works of organization studies either treat the world of large organizations as a backdrop (e.g., Perrow, 1986) or assume the world and reporting structures of large organizations for subsequent theorizing (e.g., Oldham & Hackman, 2010). However, in the 21st century, we have entered, if not a new era, at least a very different chapter of our economic history, and in this story there are fewer large publicly traded organizations, and work is increasingly not organized in large organizations. Davis (2016) notes that the number of public corporations has declined significantly, with half as many U.S. publicly traded corporations in existence in 2012 (~4,000) as in 1997 (>8,000). This trend is likely to continue as there are also fewer IPOs. Fewer companies went public in the 6 years between 2008 and 2014 than in 1 year in 1996.

Additionally, the large companies at the leading technological edge of society that are publicly traded employ fewer people. Davis estimates the combined workforces of Google, Facebook, Twitter, Dropbox, Zynga, Zillow, LinkedIn, Uber, and Square total 80,000 employees. Blockbuster alone had more than 80,000 in 2005, and GM expanded by more than 80,000 in 1942 (Davis, 2016). Of course, large, publicly traded corporations still exist and still employ people. But they play a diminished role in creating new jobs and a very small role in innovation of job design.

In start-ups, new systems emerge and get constructed without the residue of success. Baron, Hannan, and Burton made this point in their research associated with the Stanford Project on Emerging Companies (e.g., Baron et al., 1996; Baron, Hannan, & Burton, 1999). They wrote that, in contrast to their study of how the employment practices, organizational designs, and business strategies of firms in Silicon Valley developed, most employment systems research suffers from a survivorship bias problem:

> It is not clear whether meaningful theoretical inferences can be drawn from most empirical studies of employment systems in organizations which tend to be either cross sectional comparisons across a sample of survivor organizations or case studies of what has transpired in a particular setting. (1996, p. 6)

Therefore I have tried to settle on a research context that may allow me to examine the forefront of work while mitigating some of the setting issues addressed by Baron et al. Trilogy Corporate Services exists at the intersection of the software-driven space of high technology and the traditionally low-wage–driven commodity space of commercial cleaning. Semi-structured interviews with venture capitalists and founders of other start-ups provide complementary and contrasting viewpoints with the Trilogy story. The overall goal is to describe and understand how work is created in nascent organizations.

Theory and Research Questions

This chapter will draw heavily on the analogy between jobs as task stacks and technology as technology stacks, and how they are interlinked. Technology stacks consist of various kinds of different technologies set in relation to one another, such that their physical properties, or materiality, can be harnessed toward a specific end. In task stacks, the underlying tasks are set in relation to one another such that human biological and behavioral properties can be harnessed toward a specific end. In this chapter, I am specifically looking at entrepreneurial organizations. Because entrepreneurship provides a very specific kind of backdrop for examining jobs and technology, this section begins with some of the viewpoints on entrepreneurship and shows the implications for these viewpoints on jobs and technology.

Entrepreneurship

Entrepreneurship continues to grow as a vital area of study throughout the management disciplines, and classes on entrepreneurship are in constant demand in management schools. Scholars generally consider entrepreneurship to be on the leading edge of economic advancement and our national well-being. One of the most fundamental debates within entrepreneurship research is around the origins of entrepreneurial opportunities. Those who subscribe to the “discovery opportunities” viewpoint argue that exogenous shocks to existing industries form entrepreneurial opportunities and that attentive individuals and firms can discover and exploit these opportunities (Kirzner, 1997; Shane, 2000). The other perspective is the “creation opportunities” viewpoint (Shane & Venkataraman, 2000). This perspective suggests that entrepreneurial opportunities are endogenous and that entrepreneurs construct them through an enactment processes (Aldrich & Ruef, 2006; Weick, 1979). Each viewpoint comes from a different epistemology, and the epistemologies appear to be mutually exclusive (Alvarez & Barney, 2010).

Discovery opportunities are objective and caused by exogenous shocks to already existing markets. These shocks could be changes in demographics or consumer preferences, for example, though I imagine the full list would be quite extensive. The shock causes the potential for economic opportunities that an alert individual can notice and exploit. Scholars who fall into this camp draw on an epistemological framework grounded in critical realism (Bhaskar, 2013; Gorski, 2013). This is evident in three assumptions about discovered opportunities. The first is that they are objectively real and exist independent of human perceptions. The second is that people who exploit these opportunities are different from others, either in personality, prior knowledge, or risk tolerance, before they begin to exploit them. This leads to information asymmetries (Shane, 2000). Lastly, discovery opportunities are marked by risk because, at its core, a discovery opportunity is a market imperfection marked by knowledge asymmetry. Either would-be entrepreneurs already possess knowledge others do not, or they collect this knowledge as they pursue a path toward exploiting the opportunity. Because the opportunity is objective and information can be gathered about whether or not to pursue it, the opportunity is therefore marked by risk. In theory, though the realization of the final outcome of taking an opportunity may be random, the entrepreneur can collect enough information to decide on the probability of potential outcomes (Knight, 2012).

By contrast, evolutionary realism is the ideological forbear of the construction opportunities viewpoint. Evolutionary realism itself emerged out of the social constructionist perspective of Berger and Luckmann (1966) in which all phenomena are in fact constructed through the actions and interactions of individuals. Evolutionary realism attempts to reconcile the strict viewpoint of social constructionism with the classical realist viewpoint in asserting that there can be socially constructed phenomena as well as objective "real" phenomena. The ontological stability of a constructed phenomenon can be tested against objective phenomenon and be discarded if it fails the test. Those retained may have real impact on the world. From the standpoint of entrepreneurial opportunities, the evolutionary realism perspective suggests three things. The first is that opportunities are not objectively real but constructed through processes of enactment. Would-be entrepreneurs test their construction and perceptions of a given opportunity within the context in which they are situated. In other words, they undertake a process of experimentation and adjustment to see if their constructed opportunity has purchase in the market. The second point is that this approach is silent on whether the qualities of entrepreneurs before they take advantage of an opportunity are important. While they may gain knowledge in enacting an opportunity, the ex ante alert entrepreneur of Kirzner (1997) and Shane (2000) is not a factor in this story. Lastly, Knightian uncertainty marks the context in which constructed opportunities take place, which is to say that the outcomes of opportunity are random and the probabilities associated with those potential outcomes are unknowable. This is due to the fact that knowledge needed to pursue creation opportunities is only gained during the enactment

(Alvarez & Barney, 2010). That is to say that it is through the process of pursuing and creating the opportunity and getting feedback from the market that the entrepreneurs get confirmation that the opportunity is in fact “real.”

Two very important implications for both entrepreneurs engaging in entrepreneurship and for researchers examining entrepreneurial contexts emerge. The first is that the epistemology of the opportunity suggests the kinds of strategies the entrepreneur would pursue in enacting an idea, and suggest to the researcher the kinds of elements that they should be looking for in examining the entrepreneurial process. The discovery opportunity perspective suggests that entrepreneurs can get most of the information they need to pursue that opportunity, whereas the construction opportunity perspective suggests that the information will come through a process of enactment and experimentation and that clues from the environment will guide the entrepreneur. These two views serve as baseline logics for examining a potential entrepreneurial opportunity. For example, for examining a creation opportunity, Alvarez and Barney (2010) suggest that process research and examining the entrepreneurial process in action will put the full activity of entrepreneurial enactment on display. I would also offer that the process of enactment has a theoretical home in Weick’s (1979) view of enactment in the environment, and as such, the work on sensemaking may have much to contribute here. In a review of the literature on sensemaking, Maitlis and Christianson (2014) define it as “a process prompted by violated expectation that involves attending to and bracketing cues in the environment creating inter-subjective meaning through cycles of interpretation and action” (p. 67). This parallels the enactment and experimentation process in the markets within which the actor is situated. Weick, Sutcliffe, and Obstfeld (2005) have a bit more specificity in that they say that sensemaking can be seen as a “reciprocal process between actors (Enactment) and their environments (Ecological

Change) that are made meaningful (Selection) and preserved (Retention)" (p. 414). From this viewpoint, entrepreneurial opportunities that are created display a logic consistent with Aldrich and Ruef's (2006) variation, selection, and retention model, where the variation comes from an interactive process with the environment, and other entrepreneurs figure out what sources of variation to keep through a process of experimentation.

The implications of the theoretical connections between sensemaking and the qualities of creation opportunities are particularly exciting. If entrepreneurial opportunities are created, if they are stories formed in the minds of entrepreneurs and then written, then the logic of sensemaking provides the particular grammar that one might expect to see in that story. To that end, one perspective that comes out of the sensemaking tradition is one of contextualized engagement. In my work with Barton, Sutcliffe, and Vogus, we articulate that uncertain and potentially ambiguous conditions require attention, but also discernment and understanding, to see what data and emerging cues signify and how actions and behaviors need to be adjusted in response (Barton, Sutcliffe, Vogus, & DeWitt, 2015). Organizations that understand this balance adjust effectively by engaging different parts of the organizational system, which usually entails taking observations and information from frontline workers and then having managers process that information in light of their experience and expertise. Specifically, contextualized engagement involves processes of anomalizing, in which organizational members take care to notice discrepancies, perturbations, and problems in the environment, while resisting the urge to collapse them into categories. The categorization process normalizes events and makes people treat them as unimportant (Vaughan, 1996). Additionally, leaders have a critical role in facilitating and enabling a process of anomalizing. They can both help to encourage people to engage in this process and help people frame what they see and put it into context. This is called proactive leader sensemaking. Barton et al. found that, in a study of wildland firefighting teams, anomalizing and proactive leader sensemaking were positively associated with performance. While entrepreneurs are not firefighters (usually), the construction opportunity view suggests that they will be in an environment of uncertainty. Taking all of these points in sum, one potential viewpoint is that there is a process of entrepreneurial enactment in which would-be entrepreneurs, as they implement their ideas, are actively perturbing the environment to get information that then informs revisions to their entrepreneurial idea and their implementation strategy. They then select new concepts and use them to revise their idea. They revise it again and, depending on the new information they get from the environment, they may revise again or retain certain concepts. All the while, the leader of the entrepreneurial activity is encouraging people to give them as much detail on these perturbations as possible such that they can be framed to the lens of the leader's experience. The active researcher can look for the aspects of this grammar surrounding the start-up process to understand the presence of a construction discovery.

The second implication behind the dueling epistemologies addressing entrepreneurial opportunities relates to the use of the building blocks of enterprise construction. Technology, for example, is such a building block. Most technology-focused companies that began as start-ups in the past decade – including Uber, AirBnb, Instagram, and Slack – brought a novel idea to market but did not truly invent any new technology per se. They essentially took existing technology and reconfigured it, and then placed it in a technology stack to suit their specific ends. The discovery opportunity epistemology suggests that technology disturbed the economic landscape and created inefficiency. Each of these companies had a sufficiently well-informed view of the potential opportunity, such that they were essentially evaluating the risk of choosing between certain technologies in order to exploit it. The construction opportunity epistemology would suggest that there was a process of experimentation with the technology and reformation of the constructed opportunity that made it possible to bring the product to market.

In sum, opportunities may be discovered or constructed. The building blocks of the organizations and the grammar surrounding their usage can tell us whether they began through discovered or constructed opportunities, especially if we can speak to the particularities of some of these building blocks.

Jobs and Technology

The fundamental challenge with the discovery opportunity/constructed opportunity viewpoint is that it conflicts with the processes of what we know about both jobs and technology. The discovered opportunity viewpoint doesn't conceptually align with the inherently experimental nature of technology's use in the organizational context. As Chapter 2 described, in contemplating the material nature of technology, Leonardi and Orlikowski suggest that the affordances it provides can only be discovered and combined through a process of experimentation with the organizational context. The discovery perspective generally puts forth a concept of risk, and research mitigates risk in the entrepreneurial context (Alvarez & Barney, 2010). Generally, a discovered opportunity is known and understood, and the alert entrepreneur can gather data and use the information gathered to reduce the risk associated with the execution of the opportunity. It is not framed as a process of trial and error. Alternatively, the constructed opportunity perspective is framed very much as a process of experimentation and trial and error. Uncertainty dominates the constructed opportunity modality, and through a process of experimentation, the entrepreneurs reduce uncertainty by essentially perturbing the market, getting feedback from the market, enacting new strategies, retaining ones that work, and discarding ones that don't. The inherent experimentation necessary with new technology seems to be logically connected with the constructed opportunity view.

Alternatively, the risk of execution would also be part of any organization attempting to utilize technology to take advantage of a constructed opportunity. If we take the analogy from Chapter 2 seriously, then the experimentation necessary to discover how technology in a new organization would be used would also necessitate a periodic revision of the task stack in the employees in the organization. The technology stack and the task stack are potentially coevolving structures. It is possible that experimentation with one of these objects will drive fluctuations in the other. The potential instability in the task stack of a job could drive instability in the job itself and the result would be execution risk. The business idea itself may be fundamentally sound, but the team in place and the way the jobs are structured can prevent proper execution of the idea. Execution risk is a concern of venture capitalists, and they typically attempt to lower the risk of investing in a start-up by making management increase the formalization of jobs within the organization as a contingency of their investment (DeWitt Interview Notes, 2019). However, execution risk is supposedly the primary modality of the discovered opportunity epistemology.

One of the possibilities here is that the field's focus on the constructed vs. discovered opportunity is unnecessary. Perhaps all entrepreneurial ventures display some variety of the market uncertainty associated with constructed opportunities and the execution risk associated with discovered opportunities. What is clear through this exercise is that close investigation of a start-up may demonstrate how it is possible to cocreate jobs along with technology. Additionally, a close examination of the start-up process may show us how they think about the concerns of market uncertainty and execution risk. All of these factors motivate my research question:

How do entrepreneurs navigate the organization of work, specifically the organization of tasks into jobs in a new organization, while also taking advantage of opportunities new technology affords?

Trilogy Corporate Services – An Overview

Founded in 2015 in Wixom, Michigan, Trilogy Corporate Services provides commercial cleaning services to a range of clients. While its business area seems very traditional, the way in which it executes this work is quite different from industry norms. The company's cleaning technicians utilize tablet computers and smartphones in their daily work. This gives them access to a fairly deep technology stack that includes, among other things, time reporting software, Slack (chat software), and Co-Pilot, Trilogy's proprietary software. Co-Pilot guides the technicians through the physical layout of an office building while also specifying the steps to be completed in each area of a building, and in what order they should complete those tasks. Cleaners can log problems in the buildings in the software and communicate issues that can be conveyed back to the client. The founders explained that this workflow allows them to deliver a more efficient, high-quality end product to their clients. Further, the software sells both clients and potential employees on the Trilogy experience. Clients love seeing the software in action as it gives them additional confidence that the Trilogy cleaning technicians are doing good work. And the cleaning technicians like the software as it helps with the work but also allows them to recast themselves from a janitor to a cleaning technician.

Data and Methods

In this study, I am largely using an inductive qualitative approach in order to advance my understanding of how entrepreneurs structure jobs around novel technology and navigate the challenges of organizational formation. Qualitative methods are particularly useful for understanding how a complex process unfolds over time (Christianson, Farkas, Sutcliffe, & Weick, 2009) and appropriate for uncovering the process by which an organization discovers how to use new technology in context (Leonardi, 2017). Additionally, the company might serve as a truly unique case (Eisenhardt & Graebner, 2007) in this changing labor landscape. To my knowledge, none of its competitors or similar companies nationwide use technology in the way that Trilogy does. Because of this, the exploration of the application and impact of new technology in this space should offer insights into impending interaction effects of technology and workers across traditional industries. At the same time, other research touches upon my theoretical areas of interests and took place in similar research contexts. These include qualitative work on entrepreneurial processes (e.g., Baron et al., 1996; Greenwood & Suddaby, 2006), the impact of novel technology in well-established fields (.e.g., Beane, 2018; Zuboff, 1988), and the examination of "dirty work" such as janitorial services (Ashforth & Kreiner, 1999; Ashforth, Kreiner, Clark, & Fugate, 2007). However, I have not found a qualitative study that examines all of these aspects in one organization. This makes the opportunity to study Trilogy Corporate Services truly unique.

I have chosen to use a case study of Trilogy supplemented with the analysis of supplemental interviews with other current entrepreneurs, former founders, and venture capitalists. Case study methods have broad applicability whether one is trying to do explanatory or descriptive inquiry (Yin, 2003). My focus is largely descriptive, given Trilogy's distinctiveness. In addition, a case study allows me to construct a deeper case history over time (Ozcan, Han, & Graebner, 2017). Yet I do not seek to advance a theory of one company. Supplementing the case study with an analysis of interviews of other entrepreneurs provides both confirming and disconfirming accounts of the implications of Trilogy's specific context, in keeping with my goal of describing the phenomenon of entrepreneurship pursuits broadly.

Data Collection

The analysis addresses three sources of the data: interviews with Trilogy executives and associates, correspondence with Trilogy associates and company documents, and interviews with entrepreneurs at other start-ups. The interviews with the Trilogy team include semi-structured interviews with the cofounder and CEO Pat Olson[25] as well as semi-structured interviews with the other cofounder and current president Brandon Bunt. I also conducted interviews with cleaning technicians to get their perspective. All Trilogy interviews took place between May 2017 and April 2019. I conducted 21 interviews with Pat, 7 interviews with Brandon, and 1 interview each with three cleaning technicians and one person in the Trilogy office team. The correspondence and company documents corpus consists of emails, slide decks, and organizational charts that I received in my communications with the company from late 2017 through spring of 2019. The interviews with other entrepreneurs consisted of semi-structured interviews with a convenience sample of people from my network of acquaintances who are current or former entrepreneurs and venture capitalists. The entrepreneur interviews occurred in March and April of 2019. I interviewed 13 entrepreneurs once each. A list of questions that formed the basis of the semi-structured interviews can be found in the figure below.

- Questions asked of Pat and Brandon
  - What were some of the most important events from the past week?
  - What are some of the most important things coming up in the next week?
- Questions asked of other Trilogy employees
  - Can you please tell me about your job?
  - What are some of the daily tasks that you do?

[25] Olson has since retired but is still a co-owner of, and advisor to, the business.

    - How do you like using Co-Pilot?
    - Can you tell me about a time that your feedback on a work process was encouraged, if ever?
    - Does the technology make you better at your job?
- Questions asked of entrepreneurs
    - Can you tell me about your current role?
    - Can you tell me the story behind your entrepreneurial venture?
    - How are jobs structured in your organizations?
    - How have you parsed your overall strategic goals into jobs to be done?
    - What is your technology stack?

Figure IV.1 Questions Asked in Semi-structured Interviews

Data Analysis

While the case method is used quite broadly, practitioners of the method have not articulated a preferred way of analyzing case study data. Some case studies have used grounded theory approaches alternating between data and theory to generate findings (e.g., Christianson et al., 2009), and others provide detailed timelines of events, eschewing any particular methodological choice in favor of elaborating in great detail on the key events of the case history (Nelkin, 2004). Yin (2012) notes that though the case study methodology is used often, none follow routine procedures. But he does say that the overall goal of a case study analysis is "to piece together the coded evidence into broader themes and, in essence, create a unique algorithm befitting the particular case study" (Yin, 2012, p. 150). To that end, I have chosen thematic analysis, "a method for systematically identifying, organizing, and offering insight into patterns of meaning (themes) across a data set" (Braun & Clarke, 2012, p. 57). Thematic analysis emphasizes the use of codes and themes to analyze a data set that is reminiscent of grounded theory (Strauss & Corbin, 1997), and indeed Braun and Clarke highlight thematic analysis as being adjacent to that work. But they also note that thematic analysis is just a method for data analysis and does not "prescribe methods of data collection, theoretical positions, epistemological, or ontological frameworks" (Braun & Clarke, 2013, p. 178). The flexibility of this method to reveal meaning across data sets is appropriate for me as I am trying to capture themes into data sets, but I am also interested in describing processes across time to the extent possible. Thematic analysis allows me to do both.

Findings

Overview of Findings Section

The findings here consist of two sections. The first is a summary, description, and partial interpretation of key events in Trilogy's history. One of the values of a case study, especially one in the entrepreneurial space, is to get a sense of the process of how an entrepreneur or set of entrepreneurs actually put their ideas into action. Therefore the description of the story is worthy of being elevated to the level of findings. The description of the story also gives some insight into how a new organization thinks and goes about job creation. This is an important thread to keep in mind, as the point of elaborating on the entrepreneurial process in the Trilogy story is to show that this process is deeply interwoven in the process of job creation. Also, in the presentation of the Trilogy timeline, I include some detail of how the organization adjusted its technology practices as the company saw interactions between its proprietary technology with its workforce and with the environments in which it was operating. The second part of the findings section is a preliminary version of the results of thematic analysis. A more fulsome version of the thematic analysis will be conducted in a later iteration. One of the benefits of thematic analysis that Braun and Clarke (2012) note is that it excels at both inductive and deductive analysis. This was true across the Trilogy case and the examination of other entrepreneurial stories. But for this version of the analysis, I will focus on the themes that emerged from the analysis that speak directly to the entrepreneurial opportunities, jobs, and technology topics that were addressed in the theory section of this chapter.

The Trilogy Story

Early days: Circa 2012. Patrick (Pat) Olson and Brandon Bunt met by chance, and this led to working together by choice. In 2012, Pat was semi-retired. He had founded a publishing company called Hayden McNeil some 20 years before that specialized in the academic textbook market. In 2008 Macmillan Publishing acquired Hayden McNeil and after a period working at Macmillan, Pat left in 2009 and considered what to do next. It was clear to me, observing Pat's enormous energy, that it was inevitable he would reenter entrepreneurship.

Brandon was running a landscaping company that did work in Pat's neighborhood. Pat remembers being impressed by Brandon's fleet of sparkling white trucks moving through the neighborhood. And Pat remembers being very impressed with the reviews that his neighbors were giving of Brandon's company's work. Pat approached Brandon about doing landscaping work and Brandon's company soon began working for Pat.

Brandon had started doing lawn care and landscaping work in high school. And even in those early days he approached this work with an eye toward great quality and efficiency. Brandon told me that his thinking about efficient operations came from his father, who worked on the production floor at one of the "Big Three" automakers for many years. Brandon's father never had the formal degreed education in engineering that was necessary to become a foreman, yet Brandon recalls his father tirelessly studying industrial engineering processes and efficiency, particularly the work of industrial efficiency pioneer W. Edwards Deming. Brandon's father brought that engineer's gaze to teaching Brandon how it was possible to take care of an entire lawn and do things in such a way that you leave your truck with your equipment in hand, and then sequence all of the yard work so you return to your truck only at the end of the job with all of the work complete.

Brandon's company began working on Pat's property – delivering excellent service – and Brandon and Pat got to talking. Short conversations on Pat's lawn turned into 2-hour long conversations on Pat's lawn, which turned into lunches. Pat became an informal consultant to Brandon's landscaping business. Pat explained that he brought the "scar tissue" of entrepreneurial experience to some of the challenges that Brandon faced. When they first said that, given Pat's deep entrepreneurial knowledge and capital and Brandon's youth and efficiency, they could put something together; it was just a joke.

Spring 2015: The formation of Trilogy. Pat and Brandon founded Trilogy Corporate Services in the spring of 2015. In this business they would do the landscape work that Brandon knew so well, but they planned to focus on exclusively corporate clients. Knowing that landscaping is a commoditized business, they figured that if they wanted to pursue the work at scale they would have to differentiate themselves in the market. They decided to use cutting edge technology, solar-powered landscaping equipment, to create such differentiation. They focused on corporate clients because it provided a customer base for which solar-powered equipment would be a positive selling feature. Such equipment makes very little noise, which corporate clients value because the landscapers could avoid disturbing office workers. They also appreciated the opportunity to say they were offsetting carbon emissions. And the younger employees who followed Brandon into this new business love the idea of using environmentally minded equipment. By bringing clean, hi-tech tools, and efficient work processes to what Brandon acknowledged is an "old, dirty, loud and polluting business" Trilogy was able to get traction right out of the gate and quickly fill up their entire landscaping capacity.

Of course, like any nascent business, Trilogy ran into some challenges. Though they filled up their initial capacity, they found it difficult to land additional business in order to grow the organization. While many corporate clients loved the quiet aspects of the equipment, other potential clients preferred lower prices over any kind of quality. The equipment was unreliable, because solar-powered equipment of this magnitude was relatively new. While Trilogy was busy, by it's second year it wasn't profitable. Another thing that Pat noticed was that the landscaping business ended up being very challenging to systematize. He recalled:

> One of the things that I learned that I didn't completely understand about landscape maintenance is, it's hard to systematize around it because [of] its subjective nature in its incredible variations of drought and flooding and insects and pests. You know, like, [it involves] applications of chemicals and variable workflow and seasonality and just everything that you could imagine throwing at an industry to make it hard to do. (2017-05-26 interview notes)

Thus in addition to the market risk associated with customers unwilling to pay a premium for their product, Pat and Brandon encountered the execution risk associated with dealing with the seasonality of landscaping and the challenge of bracketing the external environment in such a way that they could eventually deliver their landscaping product both profitably and at scale. They found a way to pivot into a line of work that solved these problems when one of their landscaping clients asked them: "I don't suppose you guys clean offices, do you?"

2017: The pivot to cleaning. In January of 2017, Trilogy began cleaning office buildings as a way of supporting the landscaping business. They would increase the revenue per client for the company overall without increasing the sales cost associated with marketing to new clients. Most of their corporate clients were not thrilled with their existing cleaning services and were excited at the prospect of consolidating vendors and having Trilogy perform their cleaning work. Pat and Brandon saw an opportunity where they could bring technological innovation to a space – commercial cleaning – that had never seen it, which would allow Trilogy to differentiate itself in the market. They soon saw that the commercial cleaning market was substantially larger than the landscaping market could ever be. In the early spring of 2017, Trilogy informed their clients that they would be winding down the landscaping operations and focusing exclusively on cleaning. That summer I began my discussions about this work with Trilogy.

Pat explained that another factor in the refocus was in that cleaning exhibited much lower variability compared to landscaping. At the same time, every building involves many cleaning tasks. Brandon once mentioned to me casually that I needed to just come see the 75 steps or so needed to clean a restroom. Pat noted this complexity as well:

> We have one client, a relatively small building, it's probably 25,000 square feet. The building has probably 50 to 100 employees. So it's not real big or anything like that. But there are, if you just count them up in that one building, there are 782 discrete actions that our cleaning technicians will make within a given course of work. (Interview notes Summer 2017)

It was recognizing the complexity of cleaning a building that prompted Pat and Brandon to realize the potential benefits of having a technological guidance system that all of their workers could use. This system would guide the workers through the building, and show them all of the cleaning steps needed to deliver a great cleaning product.

March 2017–October 2017: Technological exploration, staffing up. Trilogy explored how they could use technology to improve the commercial cleaning business over the course of eight months in 2017. They originally conceived of using augmented reality headsets as part of the business, thinking that the set of cleaning instructions and the path through the building would be best situated directly in the user's field of vision, with all the data overlaid on the environment they were moving through. Think Pokémon Go™, but with mops and vacuums. After some testing, Trilogy realized that augmented reality technology was several years away from the kind of deployment needed for their business, so the company began to work on a revised plan with guided workflow software on tablets. They called this software Co-Pilot.

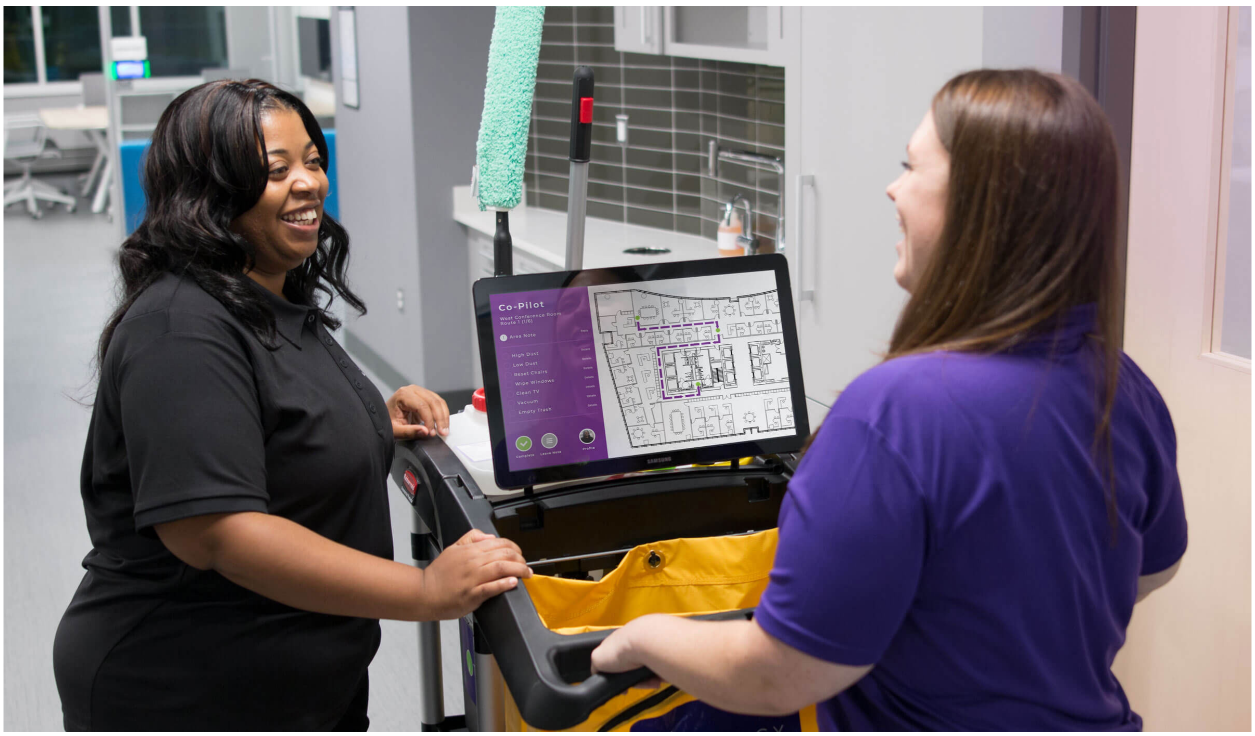

Figure IV.2 Trilogy Cart With Mounted Tablet

Co-Pilot was originally proprietary *networked* software that was on tablet computers mounted on the cleaning carts. Co-Pilot would essentially guide the cleaning technician though the entire floor plan of the building and provide them with all of the hundreds of steps the job would require. The cleaning technician could mark off when they had done each section within the software, a tracking device for themselves as well as Trilogy management. Customer services could use this information when clients had any questions about the cleaning services provided, consulting meticulous records of which cleaning technician cleaned an area of a given building and at what time.

Figure IV.3 View of Co-Pilot Software

Co-Pilot downloaded the floor plans and workflow instructions from Trilogy's servers as a cleaning technician entered a workspace. Creating this sort of software was a substantial upfront investment in both capital and time, the company had to perform and each customer had to participate in an intensive onboarding process for each customer. Pat described the process of integrating a new client into Trilogy's system thus:

> So [we]'ll get floor plans from the client that we convert those floor plans into, they're just PDFs, but we convert them into layered Illustrator files, and then we put those Illustrator files on tablets, and we do walkthroughs, and we mark it up to make sure that the cubicle configurations are correct and that we identify floor types and take

> photographs and attached stuff like are the towel dispensers metal or plastic and on and on and on, and then you have to build the workflow and all the sequences of actions, you know, and this is an astonishing amount of work to do. (Interview notes, 2017)

This seems like – to use Pat's words – an astonishing amount of work for any mostly commoditized business, but Pat and Brandon saw this as part of a cycle that feeds the business. Trilogy wanted to make sure that the cleaning technicians made good wages, an atypical goal in the commercial cleaning business. In order to do that Trilogy has to both charge their clients slightly more than other companies and to minimize employee turnover, which is very high in the commercial cleaning industry. The Co-Pilot software helps cleaning technicians to systematically move through buildings and deliver a higher quality end product, for which Trilogy can command a premium in the market. Beyond this, the outward appearance of the cleaning cart with the tablet attached delivered a little razzle-dazzle to clients. Much in the same way that Brandon's fleet of shining white pickup trucks caught Pat's eye, the distinctive look of a Trilogy Tilt cart (see Figure IV.2) with its mounted tablet display is a visible display of quality. And that matters to workers as well. While better wages can help keep employees longer, employees said in interviews that the software makes them feel they are not just janitors, doing work anyone can do, but cleaning technicians, which has technological trappings not typically associated with cleaning.

Fall 2017: Scaling the organization. In the summer of 2017, the management team consisted of Pat and Brandon, a chief technology officer (CTO), a special projects manager, and a relationship manager. In September of 2017, in anticipation of their plans to scale the business and get more clients, they hired five new managers, including a director of operations and a workflow design manager. Brandon moved from cleaning operations and focused exclusively on bringing in new clients.

The growing of the company required the formalization of positions and of communications styles. Pat noted this communications challenge in September 2017 as he recounted a conversation he had with the new director of operations, who pointed out that the informal style of communications in which any member of the operations team could just call out to another member in the next cubicle to communicate about the daily practice of work was no longer sufficient. The company now had what Pat described as "little companies" or specialized groups inside of it, meaning communication modes would need to change. Pat told me:

> So we have some work to do to try to reorient everybody a little bit to the need to coordinate and communicate more effectively as an organization because this kind of ad hoc verbal sort of thing might have worked when there were three of us. But, but now what we have are three or four little companies inside of our company, and that's not going to work. (Interview notes, September 2017)

October 2017–January 2018: Connection problems. The first iteration of Co-Pilot relied on persistent Internet connectivity to ensure that the correct guided workflow could be downloaded from the company's central servers and so that information about the cleaning technicians' progress throughout the evening could be sent back to Trilogy's systems. However, as Trilogy would discover, persistent Internet connectivity within their clients' buildings was not a guarantee, even in 2017. Clients would only let Trilogy have access to their guest Wi-Fi network, which would not provide coverage in all areas of the buildings. Trilogy bought cellular hotspots across all of the major cellular carriers in an effort to make sure that the tablets would have access to the Internet. However, wireless hotspots would not provide consistent coverage in all parts of the buildings, and cafeterias, a key area of cleaning need, were frequent places of interruption. When the Internet access would disconnect, the cleaning technicians could not log any more progress. Sometimes the software would actually restart entirely, losing track of all progress made in the evening. These Internet connectivity issues sometimes resulted in subpar cleanings, in spite of the diligence of technicians. The management team recognized the scale of the problem when a client complained about a few things being missed by a technician they particularly trusted. When they asked her about it, she said that the software had cut out and she couldn't get the route back up, and therefore she had to rely on her subjective perceptions.

This Internet connectivity problem became an issue in multiple buildings. When the software started to fail, the cleaning technicians stopped using it and instead relied on doing their routes from memory. Realizing what a problem it would be if there was no adherence to the software, Pat went out into the field with the CTO and the workflow design manager and set the goal of going to every building where technicians were experiencing problems and getting the machine working again. Through a process of experimentation, they hoped to figure out a protocol for making sure all the tablets had Internet connectivity. It did not go as he had hoped:

> I said to [the CTO] and [the workflow designer], we're going out tomorrow into the field and I'm going to grab the first one of these tablets and we're going to get it working. And then we're going to go to the next building and get it working and the next building and we're going to establish, like, right protocol right procedures all that stuff so that this is all locked down. (Interview Notes, November 2017)

He concluded, "Honest to God, Teddy, with doing nothing but troubleshooting an entire day, we could not get two machines working properly." It was ultimately clear that the software would need to run locally and that they needed new ways to connect all of the cleaning technicians to management.

November 2017–February 2018: Recruiting challenges. Trilogy experienced the accelerated pace that is common among technology companies, in their early development. The company was on a quick development cycle to resolve the initial challenges with, and iterate on, the technology. At the same time, the company was trying to grow its client base, which required it to grow the organization to accommodate this client growth. Against this backdrop some additional challenges with the company's human capital were exposed. In an interview at the beginning of December, Pat told me that he had concerns about some of the field managers who were responsible for managing the cleaning technicians. He noted that while everyone was good and well intentioned, many of them were not really capable of managing the increasing level of demands that were coming with the scale of the business.

Trilogy was also having trouble recruiting people to be cleaning technicians. In a strong economy with low unemployment, Trilogy needed a lot of really good workers very fast, and they were struggling to find them.

January 2018–March 2018: Tough winter with strategic changes. When Trilogy first started cleaning in January 2017, the company only had approximately three buildings to clean and a relatively mild winter. In the winter of 2018, Trilogy had many more buildings to clean, and it was a brutal winter, marked by heavy snowfalls, temporary thaws, and immediate refreezes. This meant a lot of salt usage on roads and walkways. And inevitably that salt was tracked into buildings. Trilogy estimates that salt-related issues increased their nightly workload by about 40% over the preceding winter. The winter conditions also made the work more dangerous for the cleaning technicians. Several of them slipped and fell and got concussions as they were trying to take garbage out to dumpsters. Others couldn't make it to work. While Trilogy paid its cleaning technicians more than most cleaning companies, its workers generally rely on older cars, and many had to travel as much as an hour in good weather to get to work. Field supervisors had to fill in the gaps as cleaning technicians, which meant there was no active supervision. Trilogy's complaint load from clients went up substantially. During a phone call in February 2018, I could hear Pat's exasperation:

> There have been these compounding factors that have – honestly, I have to be quite candid with you, taken me down a few notches. I will climb back up and get my head right back in the space, but it – but during this winter. It's been like, alright, we just have to find a way to get through today. (Interview notes 2018-02-07)

One of the side effects of this brutal winter was Brandon went back to working on company processes. He had been focusing on new sales full time since September, but the company simply did not have the capacity to take on new clients. The director of operations left the organization, and Brandon took over these duties. The silver lining was that he applied his eye toward process and efficiency to working with the cleaning technicians on the ground and noticing better ways to clean, understanding how to better support technicians, and how to improve both the company's software and its hardware to improve service. One example was his noticing how much time it took for cleaning technicians to find an outlet to plug in a vacuum. This prompted the introduction of cordless vacuums. Another example was seeing that technicians had to go out to the dumpsters to get rid of garbage five or six times on a single job. The company provided longer tilt carts with a larger capacity for holding garbage, reducing the trips out to the dumpsters down to two. Describing these benefits, Pat invoked an old phrase from car production about going to the "gemba," referencing the Japanese term for "the actual place." Going to the gemba, means going to the shop floor and actually observing the production processes in real time. As Pat described:

> What we've really concluded, a deep understanding, is that a bunch of white-collar people in an office separated from the work can't conjure up a thing that is going to work in the field that, you know, like the old Toyota saying you really have to go to the gemba. … We just didn't even understand … until we got out there and watched these things like, [a cleaning technician] was making five trips out to the dumpster at night. (Interview Notes, February 2018)

Not all of the fixes were as simple as bigger carts that cut trips to the dumpster. Brandon's time in the field also led them to recognize that the company could store all of the layouts and the guided workflow on SD cards that could be loaded in and out of the tablets.

Challenges with certain customers also became clear that winter, as being short-handed highlighted just how much time it took to do every part of the business from cleaning to onboarding. Pat noticed how difficult it was to onboard multitenant building clients. Corporate clients that have one massive building provide one set of keys, one alarm code, and one facility manager who sets the guidelines for what needs to be done. All these onboarding factors are multiplied in a multitenant building. Pat also recognized that multi-tenant building managers are generally more cost focused than quality focused, which was not the best fit for Trilogy's business model.

The difficult economics of a technology-focused cleaning business also came to the fore during this difficult time. Pat came to the conclusion that the management structure was very expensive and it was not driving the right results. Brandon's reentry into the field was pushing things in the right direction. He was spending multiple nights at buildings until 3:00 in the morning, carefully watching everything that people were doing and codifying it into best practices. However, the rest of the management staff was not bringing the same level of results and progress. The overhead was simply becoming too much for the business to maintain. In the middle of February, Pat negotiated exit packages for five people in management and redeployed some people who worked in the office into field supervisory roles. Conditions in the field as much as economics drove this decision. During a particularly bad snowstorm in the second week of February 2018, Ann Arbor and the surrounding areas received approximately a foot of snow overnight. While most clients said that the cleaning technicians could clean their buildings over the weekend, one demanded service as usual. Pat recounts this story better than I ever could:

> We reached out to our workers in that building. They live in Southwest Detroit. We figured it'd be a two-and-a-half or three-hour commute each way for that and we call them and said, "Look, it's up to you. If you want to try to do this, we will definitely pay you more than normal. But if you don't we understand." And they said, "We'll do it." As far as we understood, they reported for duty. Then about 9:30 that night Brandon, who is a very caring person, started thinking they're going to be there forever because the [snow] delay, you know and getting there, and he grabbed three or four people and pulled them off their post and put them in a pickup truck, a four-wheel-drive dump truck literally, and drove them out to Ann Arbor to help this team so they can get done at a tolerable hour and get home…. So one of the cleaning techs we brought in to help goes around to the cleaning closet and it hasn't even been opened yet and she got, she's struck with fear like "Oh my God, they must have gone off the road or something. They didn't make it." So she runs up to the security desk and says, "Have you heard from our cleaners?" And he said, "Oh, yeah now, they signed in two and a half hours ago." And

then she hears a little bit of laughter and up in the mezzanine. She walks up the stairs and in the mezzanine, all three of the workers are there, sitting up on the couches feet up bullshitting. Burning hours and hadn't, even in two and a half hours, hadn't even opened the cleaning closet to pull out their carts out yet.

Pat has told me that this was when he realized the very difficult path he is trying to navigate between running a white-collar tech company in a fundamentally blue-collar business. Given the blue-collar nature of the business, he believes that much more supervision than the company had been providing is necessary, in conjunction with the technology, to deliver a high-quality end product. The restructuring was designed to address this need.

March–May 2018: Technology and task stack changes. Trilogy began to roll out a new version of Co-Pilot that was not reliant on Internet connectivity in the spring of 2018. All of the guided workflow and floor layouts were loaded onto SD cards that the application on the tablets would access. They also provided all cleaning technicians Android phones that would operate a custom application the company provided, Slack, and a game they called Money Match. The custom application, which was created with Google Appmakr, allowed technicians to report their progress through an evening's work. Co-Pilot would guide the worker through the space, but the customer software was the means by which the cleaning technicians would mark each section of the building as complete. The company had both building-specific chat lines and general chat lines on Slack that technicians could use. These channels were intended to give technicians an easy way to communicate with each other to ward off loneliness, a way to call for help if they were injured, and a way to report any problems to management. Money Match was designed to incentivize technicians to use Co-Pilot and log their stops. Once technicians had completed various points on the evening's route and logged the progress in the custom application on the phone, they could play Money Match, a simple game, and if they won they might receivea $10 gift card to Starbucks or Target.

Brandon made it clear to me that Trilogy's technology is one piece of a three-part system; the other parts are training and supervision. Technology is what makes the tasks repeatable night after night. This relies on training. Trilogy teaches its technicians to use Co-Pilot and the new app as well as how to have the eye and the precision to complete the cleaning tasks to the desired standard. Supervisors, Brandon explained, not only ensure that people complete the tasks; they also ask questions like, "How are things going?" and "What do you need help with?"

That spring, the benefits of the unconventional approach to hiring technicians were becoming clear. Brandon explained that cleaning experience can actually be a weakness in a job candidate:

> For me, it's no, I want the person with no experience in cleaning, but [who] has some work experience. That's already been in the workplace, has worked for a couple of crummy employers like a McDonald's or a Walmart, or somewhere where they're not really … I know they don't get access to full-time. They don't get access to a good, nurturing environment. And so, because I want people without [cleaning] experience … I think we have a greater chance of getting cleaners not only to stay with us, but we have a greater chance of getting compliance, and getting them to adopt and follow through with our methods. (Interview Notes Spring 2019)

Brandon wants technicians who have not cleaned before because he believes that it is difficult to train people who have done the work for companies that have lower standards. He also indicated that workers who have been ill-treated may appreciate Trilogy's efforts to provide a more supportive environment. He sees this as a way to retain workers and create a superior product.

May 2018–present. After a long and challenging winter in Michigan compounded with some health issues, Pat decided to step away from the business in May 2018. He and his wife closed on a house in Palm Springs, California. Brandon took over full control of the operations of Trilogy. Pat is still an owner, but his role in the day-to-day business is purely advisory. Brandon calls him in as needed and he serves in a quarterly advisory role.

Brandon has continued to grow the business and right sized the client base. The gross margin was 19% and below in the winter of 2018, with improvements driven by the departure of some of the management team, but in the summer of 2018 he achieved a 40% gross margin. Sales increasingly focuses on larger companies that place a premium on the quality of the cleaning product that they receive.

Even with the reduced management overheard, the start-up costs associated with taking on a new client, combined with the ongoing expense of software and process development, effectively imply a minimum efficient client size. In 2017 their average client size was approximately $50,000 in annual revenue. Now clients must pay them at least $300,000 annually. This means a focus on clients that have corporate campuses, and multitenant buildings are no longer in their client list.

Some of the initial pain caused by the company's original client base was masked by Trilogy's substantial access to capital. Because of Pat's success in his previous business, he had lots of capital to deploy in service of growing Trilogy. Brandon notes in hindsight that this was both an advantage and a weakness. The capital allowed them to experiment broadly, develop software faster, and learn from failure at a rate that most companies can't experience. However, the capital also allowed them to perhaps grow too quickly. If he could do it over again he thinks he would spend more time closer to what he called “the dirty core” of the business, with a smaller, more nimble management team working very closely with the cleaning technicians to work out all of the glitches in the cleaning process, before spending the capital and software development time to try to generalize and scale these systems.

Preliminary Findings of Thematic Analysis

On its face, it would seem that Trilogy’s story fits the mode of discovery opportunities. The dominant modality for discovery opportunities is that of risk. For a discovered opportunity, entrepreneurs should be able to get most of the information they need by researching the opportunity objectively. The dominant modality for constructed opportunities is one of uncertainty. In this mode the entrepreneurs don’t even know precisely how their offering will be received, so they go through a process of experimentation and enactment, basically putting their concept out into the world and seeing how the market will react to it and then adjusting the concept as necessary.

Pat and Brandon quite literally discovered the need for cleaning services when a client asked them if they provide it. They then researched this area of business and discovered that many recipients of cleaning services feel insufficiently served. But Trilogy’s story after this discovery reflects the experimentation and enactment associated with a created opportunity, as the company had to experiment and test how their product could serve clients in order to understand how their evolving technology, potential customers, and potential workforce would interact to deliver a profitable product. It seems that the discovered opportunity vs. constructed opportunity dichotomy does not properly encompass Trilogy’s story, as it shows aspects of both risk and uncertainty. Many of the stories from other entrepreneurs in my data set will demonstrate a similar blend. In addition, these entrepreneurs also conveyed information that displayed variation in terms of market concerns and execution concerns.

Based on comparing the data to the theory, four top-level themes emerged that contain many of the common observations across the data set while also addressing the topics presented in the theory and observation section. Given the richness of the data set, I will have to reexamine it in its entirety in future revisions of this work. But for now, many of the events in Trilogy's brief history can be summarized in a two-by-two square around the concepts of risk and uncertainty and market and execution. All aspects of this two-by-two square have implications for job construction in the entrepreneurial setting. In this section, I will mix in examples from the Trilogy case along with other examples from the additional entrepreneurs that I interviewed to explain the points, and I will connect each factor to how jobs are structured in entrepreneurial settings.

Moving from market uncertainty to execution risk. Among the entrepreneurs I interviewed, market uncertainty and execution risk dominated in very different periods in the progression of the organizational evolution. I refer to both uncertainty and risk from both the viewpoint of the entrepreneurs running the businesses and of potential investors in their businesses. Contrary to the entrepreneurial opportunity framework that is dominant in the literature, which frame risk and uncertainty as binary conditions, organizations displayed both risk and uncertainty, just generally at different times.

Market uncertainty can be summed up by a modification of a phrase my dad used to say: "Boy, is anyone buyin' what you're sellin'?"[26]While they do not specifically refer to uncertainty as market uncertainty, Alvarez and Barney's (2010) choice of terminology suggests that

[26] I was not, strictly speaking, selling anything. He was referring to statements I made that he found not entirely plausible.

experimentation and enactment mitigates market uncertainty. It is only through the process of pursuing and enacting the opportunity that a would-be entrepreneur gets the feedback from the market confirming that the opportunity is real. In my interviews, one of the ways entrepreneurs addressed market confirmation was through a concept called product market fit. Product market fit is achieved when you have iterated your product enough that a market segment will actually purchase it. One of the entrepreneurs in my sample originally started an online ordering service in the early 2000s, where individual consumers could use the SMS feature of cell phones to place orders for food ahead of time, and then skip the line once they arrived to the restaurant. While the company gained some modest traction with this product, the founder felt the company really achieved product market fit when they leveraged the technology and infrastructure they built for consumer ordering to instead market directly to restaurants. The company effectively serves as the back-end technological stack provider for any restaurant that wants online ordering. At that point, their revenue grew substantially and a larger company became interested. The entrepreneur explained:

> So we got to profitability. PayPal got very interested in what we were doing and decided to make an investment in the company at the end of 2012, and that was really when we knew we had to scale up, because the market product fit had really clicked, and we brought on a COO, a VP of sales, and a VP of customer success who had all worked together and were much more senior [than the existing staff]. (Interview Notes 2019-04-27)

Through constant interaction with their customer base, this company saw that the technological product that the restaurant industry needed was the creation of a back-end technology stack that would allow them to accept online ordering. By pivoting to this business, the company achieved product market fit. Reaching profitability and catching PayPal's eye was confirmation that they had achieved product market fit. Hiring formal roles in the executive suite was part of this company's scale strategy. The founder and CEO had previously taken what he described as a "stem cell" approach to hiring; product fit had led to a revision of the company's task stack.

The founder of a company that is much younger, one year old, also described switching to a formal hiring process after hitting product market fit. His company had recently emerged from Y Combinator, a tech incubator. He explained that the fact that the product is going to change significantly means that it does not make sense to make formal hires:

> What you're building [in the early days of an entrepreneurial venture] is not actually the thing that people want, like, you've built something, you're going to put it in front of people, and then everyone's going to tell you why it sucks, and then you're going to go and, like, either figure out, like, how to tweak it or how to totally change it. … So it would be a big mistake to hire some super specialized person who, like nails this thing you're doing now, because the thing you are doing now is probably not the thing you're going to be doing, you know, a year from now. (Interview Notes, April 2019)

Execution risk is the other prime modality in the discovery opportunity vs. creation opportunity framework. It is also in some ways on the other side of the resolution of market uncertainty. Alvarez and Barney (2010) do not use the term execution risk, but they largely frame discovered opportunities as being dominated by risk. If the opportunity exists objectively, there are potentially a number of ways to turn the discovered opportunity into a profitable endeavor. Determining which execution strategy is the right one to pursue is accomplished through a process of research and evaluation, and this is the context of risk. Fittingly enough, once a company has hit product market fit, both entrepreneurs and potential providers of capital focus on managing execution risk, especially in the context of scaling the enterprise. One of the interviews that I conducted for my research was with a late-stage venture capitalist. He said that he only invests in a company that has already achieved product market fit, because he prefers to invest in companies that can focus on execution. As a closing condition of their investment, this firm will make both formal and informal requests of the company, along the dimension of execution strategy:

> In post-closing covenants where, you know, we actually require the company to fix certain things within a time period, like, the most up-to-date, like, consumer privacy policy or whatever, you know, if we, from our diligence see that we might build into the covenant that like, “hey, look guys, you need to update this to reflect best practices,” and then you know … through diligence you realize that maybe the CFO is not that good. And you kind of like talk about it and say, “Hey, look,” to the CEO, “maybe this individual, you know, was the right person for the past couple years, but, like, now you’re a bigger company. You need to upgrade this role and that’s the kind of thing, we might not force it, but we might just kind of say it’s our recommendation that you replace this person. (Interview Notes, March 2019)

The idea that suppliers of capital eventually desire greater fidelity in operations was seen in multiple examples in my interviews, with entrepreneurs broadly having a sense of the uncertainty vs. risk tolerance of different classes of investors. For example, a cofounder of an Internet services company mentioned that early-stage investors understand that most of their investments have not yet achieved product market fit. So these investors are placing small quantities of money across many uncertain investments, hoping for one to return 50 times their investment. Later round venture capital investors are not investing in the promise of the idea as much as they are investing in the execution. So they invest larger amounts of capital expecting much smaller returns, but often demand more precise operations. And more precise operations means, at least in part, the formalization of jobs.

The implication for jobs broadly is that in the early days of iteration, of technical product tweaking to try to find product market fit, it may be that highly formalized roles are actually a hindrance to the organization. Entrepreneurs' actions in the early days in the face of market uncertainty are so varied and change so frequently that the structure of highly formalized jobs is counterproductive. Trilogy's example suggests that prematurely hiring into formalized jobs can mean hiring for positions that are actually unnecessary. When Trilogy made those management hires in the fall of 2017, while the company was not strictly trying to find product market fit, it was still in the middle of an iterative process of revising its technology and figuring out work procedures. Trilogy was strategically correct in that to-scale management-level hires were necessary, as was the formalization of roles and lines of communication. However, their example suggests that hiring in *anticipation* of scale may be problematic. With the benefit of hindsight, in a conversation in March of 2019, Pat said he essentially added roles and hires anticipating a scale up and when the scale up came in under that plan, there was a bit of chaos because the specific work he had hired people for didn't materialize. The company then had to pay people to leave at a significant expense.

Another issue that may have weighed on Trilogy on this particular dimension is access to capital. In one respect access to capital was very beneficial. Brandon noted that having lots of capital gave them the ability to experiment broadly and quickly, something that a capital-constrained start-up may have been unable to do. But the constraint caused by limited capital, appears to have the unintended, but beneficial, consequence of restricting the kind of roles for which entrepreneurs can hire. One of the people I interviewed, a serial entrepreneur and a partner at a venture capital firm, noted that first-time entrepreneurs have difficulty making management-level hires even if they want to. She said that with limited capital and a still developing product, a start-up doesn't have the ability to pay or offer attractive equity to attract C-suite–level talent. So instead of hiring top down, they hire bottom up in a very opportunistic way. They hire people who are young and scrappy for very general roles, with not a lot of specification.

Market risk. If market uncertainty concerns whether or not people will buy what an entrepreneur is selling, market risk refers to the sources of variability among customers who are ready to buy. Even if the entrepreneur knows that she has willing buyers, customer buying behavior and interaction with the organization can be highly variable. The goal of some of the entrepreneurs in my sample was to bring down the level of market risk by either training and stabilizing customer behavior or by bringing down the overall variability within the customer population through keeping certain customers while letting other customers go.

An education technology start-up attempted to stabilize customer behavior through education. But the impetus for this customer stabilization came from outcries from their clients interacting with the product. In the early days of the start-up, their clients, primarily middle school teachers, would send them very aggressive emails in which they complained about problems with the start-up's software. Because the founder of the start-up entered into the education technology space in response to a friend's lamentation about the sad state of education technological service, the founder realized that many of their customers had had very poor experiences with software in the past. By coaching their clients about how to productively communicate problems with the software, this company could both constrain the range of potential client behavior, while also positively developing the customer relationship. The founder noted:

> We would have customers in the early days where they would send us like just like flaming emails, you know, like, just like, what is this bullshit, like, it's broken. And like we would have to train them to be like, okay. First of all, not helpful. Second of all, if you want to get this thing fixed, like, here's how you do it. Like, you send me a screenshot of the thing that's broken and you actually describe what you were trying to do and, like, tell me why it's a big problem. Tell me how urgent this thing really is, and now, they've been customers for like six years, you know? But it's like, you really have to train it because otherwise they're like. You know, if you don't tell them what you need from them and they don't know how to behave. (Interview Notes, April 2019)

In this instance, the company gained information about how their technological product offering was being used in practice. As the clients interacted with the product, strategic information to the company was revealed, and it came in three varieties. First, in some cases there was truly a bug in the product, and by training the clients on how to give good feedback, they were able to quickly make needed fixes to the product. Second, the complaints from the client sometimes represented something that they wanted to do, but couldn't figure out how to do. And in many of these cases the functionality that they desired was already a part of the product; the client just needed guidance on how to find it. However, the third kind of information the company received concerned features the client desired but were truly not part of the product's feature set. These items motivated the CEO to revise the overall task stack of the company, and create a new team called Customer Success. Customer Success had the task of proactively going into the schools and training a couple of teachers in-depth, making them comfortable with the company's software. And these teachers could act as a first absorption layer for questions, but also act as an aggregation point for all concerns and potential product requests that the Customer Success team could then filter back to the start-up. This was a more efficient use of the company's time, and it also helped to know their clients' specific pain points. This is an example of how customer interactions with the technology can result in both a change to the company's task stack and, by extension, the features of its product.

Another example of managing market risk is the homogenization of the customer/client base. This is the strategy that Trilogy employed when it introduced the threshold that clients must yield at least $300,000 in revenue. Of course, companies generally benefit from earning more per client. But the delta in revenue actually represents two different things. First, average client revenue of $300,000 represents a proportionally larger cleaning area in terms of square footage as compared to the $50,000 client. So while a $300,000 client brings in more revenue, it is also more work, which means higher costs. But clients in this range are simply larger companies, which are more likely to place a premium on cleaning execution, and therefore appreciate the added value that Trilogy's technological and process innovation solutions bring to commercial cleaning. They also do not bring the variability of the multitenant building client. By focusing on the large corporate campus client, Trilogy brings increased efficiency to the repeated process of onboarding similar clients. It also gains the ability to stabilize its internal work processes, as the company no longer has to continually adjust to the shifting demands of a heterogeneous client base. Much like the educational technology company who developed a process of educating its users, Trilogy was able to figure out the precise job and organizational structure it needs to have in place – lower fixed-management overhead and close-to-the-ground mentoring with building coordinators and field supervisors that circulate among buildings. In the end, I think that what both Trilogy and the education technology start-up found is that it is very difficult to try to systematize work around the idiosyncratic actions and preferences of customers.

Execution uncertainty. Execution uncertainty reflects the interactive complexity among an entrepreneur's knowledge and experience, a company's potential workers, the structuring of jobs and technology, and the environment in which the start-up operates. An entrepreneur may realize that there is a clear market opportunity and market need to be satisfied. However, the entrepreneur may face uncertainty as to whether she can navigate the unanticipated outcomes generated by the complexity of these factors.

As mentioned in the details of the Trilogy's story, the company's first foray into operations was providing quiet carbon-neutral landscaping services to corporate clients, a service for which there was a clear market. However the interactions of a seasonal workforce, unreliable solar-powered equipment, and the complete unpredictability of the actual, outdoor environment generated substantial execution uncertainty.

These aspects made it nearly impossible to execute the landscaping business at scale. And scale was necessary to both achieve profitability and justify the capital and ongoing investment in the equipment that was the key value proposition of Trilogy's service. In contrast, the internal environment of the office was significantly less variable, and it was possible to systematize their processes. However, even in a more stable environment, the business faced many challenges that they could not have anticipated.

The first of these factors was the need for ongoing experimentation with technology to understand both its complications and how to structure it to create a premium cleaning product. Trilogy's first attempt with augmented reality proved unsuccessful because the technology has not yet progressed to a state where it would be implementable. The next iteration was the networked tablet software Co-Pilot, which had promise and worked in testing at the company headquarters. However, interaction with the actual office environments demonstrated that unreliable access to wireless Internet was fatal. This motivated the company to redesign the software to work from SD cards loaded with the building layouts and the guided workflows. But because the persistent connection was lost, the company had to add additional items to their technological stack, including Android phones with additional customized software for tracking progression and communication software in the form of Slack. Leonardi (2012) and Orlikowski (2008) note that you generally can't know the potential benefits of a technology until you see it in repeated contextual use. One specific addendum to this is that you can know neither what particular revision you will need to make to the technology nor what the final form of your technological stack will be until you interact with your work environments repeatedly. However, in addition to revising the technology stack given interactions with the work environment, Trilogy also had to revise its technology based on interactions with the workers. Even after revising the software to work without wireless Internet, the company had trouble driving adherence to Co-Pilot, the guided workflow software. Unlike landscaping, in which it was impossible to do the work without the solar-powered lawn equipment, it was possible for the cleaning technicians to work without Co-Pilot and often they did. Many of the workers felt that they did not need the software in order to clean. This is what prompted the addition of the gamification software, Money Match, to the Trilogy technology stack. This helped to motivate workers to interact with the software on a regular basis to access potential financial rewards. But it also served as a forcing mechanism to help them understand the benefits of the guided workflow software.

The second factor, parallel to this technological exploration process, was that the company had to essentially experiment around the organizational structure in terms of jobs as well as the specifics within the jobs themselves. The organizational structuring of jobs can be most easily summarized with the conflicting frameworks of a white-collar tech start-up and a blue-collar services business. In the early part of Trilogy's foray into cleaning, it was technologically oriented and scale-focused. The start-up made a strategic choice congruent with that particular logic, which is hiring lots of management roles. It turned out that even with technological innovation, commercial cleaning is still a blue-collar services business that requires attendant structures such as direct and focused supervision, however respectful that supervision might be. But these factors were not exposed until the company had a massive shift in the environment with a particularly rough winter. Difficulties serving customers and maintaining staff sent Brandon back into the field where he could observe and interact with the cleaning technicians and watch how they worked (and sometime didn't work) to create and revise the processes for cleaning, add new pieces of equipment to the Trilogy technology stack, create the support and training structures, and figure out the correct potential employees to target. In short, it was only by taking fire and coming up with counterattacks that Trilogy was able to uncover the correct organizational structure, task stack, and technology stack to execute its vision for a different kind of cleaning company.

The theme of execution uncertainty arose in the interviews with other entrepreneurs as well. In talks with a founder of a farmland investment company, I asked him how his start-up figures out which tasks go into roles. He responded that his start-up is still in the process of figuring out just how they are going to do things. As a result, the priorities change frequently. Because of this, they are constantly reorganizing around how to solve new problems. But part of what they seem to be doing is to creating a repertoire of actions that can be brought to bear on all of the situations they are encountering. And these actions can be codified and incorporated into jobs more precisely. He said:

> You know, our priorities are changing like weekly, if not daily, just based on what we have going on and, you know, what fires we're putting out. And so I think we try to do a weekly, like, all hands [meeting] with the team to really review what happens. … We're trying to regularly figure out how to resource against these needs. And you know, in a year's time from now[, we] will be on much more stable ground where we can set like kind of longer-term product road maps, but for now, it's kind of like managing more like 3 to 6 to 9 months out and doing the best we can to put people on those positions, but I think what's unique about, like, the new work [context], if you will. … I think people are more potentially more dynamic these days. (Interview notes, 2019-04-09)

Such comments suggest that it is possible to hire workers who are flexible, such that when they encounter a problem they are able to respond to it and put together a process that can be utilized more in the future. Repeated interactions with an uncertain problem space then make it possible to generate the insights and the processes that will be the foundation of more concretized work processes in the future.

Discussion

These findings have implications for some of the existing theorizing around entrepreneurship as well as my theorizing around jobs and technology. In addition, these findings set the stage for additional work both with this existing data set and for future approaches to researching jobs, entrepreneurship, and technology. In this section, I will discuss some of the limitations of my research. I will also connect back to the theoretical perspectives introduced in the early part of the chapter. Concluding thoughts about my findings will specify my plans for future work building on them.

Limitations

This research has three notable limitations. The first is the use of a convenience sample for interviews. For Trilogy, I was able to interview some of the cleaning technicians, but the vast majority of my interviews were with Pat and Brandon. As a result, I was not able to get as complete of a picture of the company as I would have liked. However, I plan to perform additional interviews in the future as I continue to build out the story of Trilogy.

The second issue is one of insufficient interrogation. For this to be a systematic multiple case study research project, I would have needed more consistent interview questions that I brought across all study subjects, and this would have needed to be paired with multiple interviews with the other entrepreneurs over longer periods of time. This would have allowed me to capture more of the process and variability in the lived experience of the entrepreneurs as they enacted their businesses. While I cannot necessarily reinterview this particular subset of entrepreneurs, for any future work, I now have a better template and a better understanding of how to approach this sort of research.

The last limitation is that this work is preliminary. There are still approximately 15 hours of transcripts that need to be analyzed and that may result in a revision to the thematic analysis as it currently stands. Of all of the limitations, this one is the most repairable. As I try to move this work toward publication, further analysis of my data set can be fixed with nothing more than time and focus. That said, even given the limitation of sample, interrogation, and partial analysis, there were several findings that I was able to uncover in this research.

Review of Findings

One of the dominant debates in entrepreneurial research has been the debate regarding the discovery opportunity and creation opportunity. While Alvarez and Barney (2010) delineated the particulars of this debate and the epistemologies implied by these viewpoints a decade ago, the field seems to have not progressed beyond it. More recent research and research commentary puts this in stark resolution (Braver & Danneels, 2018; Ramoglou & Tsang, 2016, 2018). Ramoglou and Tsang (2016) put forth the concept of entrepreneurial opportunities as propensities. They attempt to sidestep the debate altogether by saying that market demands can be actualized into profits by introducing products and services, and that these products and services come from an entrepreneur's desire to engage in the process of actualization. This approach has received a substantial amount of criticism. Braver and Danneels (2018) have essentially claimed that propensities add nothing to the debate as desires can be both created or discovered. In their rejoinder to Ramoglou and Tsang, Alvarez et al. (2017) note that Ramoglou and Tsang indulge in some of the inherent tautology of the discovered opportunities perspective. They instead advocate for a perspective that I also embrace: even if opportunities are ontologically subjective, they are epistemologically stable and therefore can be studied empirically.

To that end, I find that Alvarez and Barney's (2010) epistemological framework in considering risk and uncertainty is compelling if incomplete. My findings suggest that there is an additional consideration of market concerns and execution concerns to be considered in the actual enactment of entrepreneurial opportunities. To my mind, this is necessary because it will encompass a broader understanding of the specific challenges that entrepreneurs face. That said, my findings support Alvarez and Barney's notion that decisions and research mitigate risk, and enactment and experimentation mitigate uncertainty. Specifically, the experimentation process for mitigating uncertainty seems to reflect a process of enactment, selection, and retention (Aldrich & Ruef, 2006; Weick et al., 2005). However processes of contextualized engagement (Barton et al., 2015) do not seem to be at work here. In appears that in the early stages of entrepreneurial enactment, the founders have to themselves be on the frontlines, rather than gathering information from the frontline via intermediaries. This is in part because the start-ups in my sample all started small (or are still small), with unspecified job tasks in order to quickly react to changing sets of priorities. Trilogy's experience consisted of scaling management too quickly and only finding success when Brandon spent hours directly in the field engaging with what he referred to as the "dirty core" of the business. This however does not mean that contextualized engagement processes are not relevant to entrepreneurship, and it remains possible that further engagement with my data set and future research will highlight evidence of such processes.

In relation to job construction, the overwhelming finding from this data set is that in the early days of the start-up there is too much uncertainty, both market and execution, for highly formalized and structured jobs to be efficient. A start-up is still in a process of iterating to figure out its product offering to achieve product market fit. But it also is still in the process of perturbing the environment in which it is executing and enacting its entrepreneurial venture to get relevant information about how to structure either the jobs within the organization or indeed the organization itself. In some ways the building of a repertoire of job tasks is reminiscent of work that discusses how organizational routines are developed (Feldman, 2000; Pentland & Rueter, 1994). What I would add is that, particularly in entrepreneurial ventures, beyond the notion of product market fit, what entrepreneurs are seeking to resolve is problem solution match. Formalizing and creating highly specialized roles early eliminates the flexibility they need to traverse the landscape of potential solutions quickly. Highly fixed relational and highly specified performative components of jobs do not let you move around the solution landscape fast enough, if at all.

In so far as technology is concerned, in many instances, technology either directly or indirectly incited changes to companies' task stack. In the case of the online ordering company, a company's pivoted use of its technological offering allowed it to reach product market fit. And this necessitated a change from generalist "stem cell" hiring to formalized, specialist hiring to scale the organization. In the case of the educational services start-up, feedback generated by customers' interaction with the technology served as motivation for modifications of the company's task stack and the creation of new jobs. And in Trilogy's case, seeing its technology come into contact with a work environment in flux and with a noncompliant workforce drove multiple revisions to the company's technology.

Future Work

The case study method can clearly illuminate research in jobs, technology, and entrepreneurship. It allows very clear observations of the process of enacting an entrepreneurial opportunity, while also providing an understanding of the experimentation process entrepreneurs undertake while trying to find the right specification of their business idea. In future work, I believe that I can more clearly specify the kinds of questions that would allow for a more thorough interrogation of the jobs and entrepreneurial context. In the future, I will seek exposure to more employees of start-up organizations, beyond the founders and the c-suite. While hearing the stories of founders was vital to gaining an understanding of the entrepreneurial journey, some of the details of the job creation and revision process were almost certainly obscured by my failure to pursue further access to others in these organizations.

Another feature that I hope to implement in future work is a longer time frame over which to observe entrepreneurial ventures. Many new businesses have short lives. Some of the organizations whose founders I interviewed in my study may not be around in another year. Understanding which factors – market uncertainty, market risk, execution uncertainty, or execution risk – play a role in an organization's downfall as well as the specific strategic actions the entrepreneurs in these organizations took over the course of the venture would be useful in furthering our understanding of the peril and promise of entrepreneurial engagement and new job creation.

CHAPTER V

Conclusion

Introduction

The purpose of this dissertation was to broadly explore the underlying structure of jobs. I had the more specific goals of 1) understanding the processes that might change how jobs are structured and 2) how technology and jobs might coevolve. Specifically I asked two questions. The first is how has the task structure of jobs changed over the past 20 years. The second is how do entrepreneurs navigate the organization of work, specifically the organization of tasks into jobs in a new organization, while also taking advantage of opportunities new technology affords? To answer the first question, I used computational methods to extract latent task variables from job task descriptions. To approach the second question, I used qualitative interviews with entrepreneurs in start-ups, as well as people associated with the start-up space, to understand the process of organizing jobs around technology. As a jumping off point for my dissertation work, I theorized about an analogy between jobs and technology, and how they are interrelated.

The idea of the work having a task structure is not new. It is elucidated in both the organizations studies literature (e.g., L. E. Cohen, 2013; Miner, 1987) and in the labor economics literature (e.g., Autor et al., 2003). However this work differs in that I developed a process of examining the task structure of work without making any before-the-fact claims about how work should be structured. Rather I used topic modeling and embedding methods to extract that task structure from the language that people use to describe the jobs they do. My work suggests that this task structure can be mapped and has undergone a process of change over time.

Case studies of entrepreneurs are also not new. One of the largest that I am aware of, conducted by Baron and colleagues (e.g., Baron et al., 1996), was instrumental in providing researchers an understanding of how entrepreneurs structure employment relations in their fledgling companies. My work is not as expansive as theirs. Additionally my work differs in that I sought to understand specifically how entrepreneurs use technology and how that changes the way entrepreneurs structure jobs. To do that I utilized current theory around entrepreneurial opportunities, specifically work on the epistemology of entrepreneurial opportunities put forth by Alvarez and Barney (2010) in order to provide an interpretive framework for the qualitative data that I collected. My analysis showed that the risk/uncertainty duality Alvarez and Barney proposed needs to be extended with considerations of market concerns and execution concerns. Many different entrepreneurial ventures exhibit examples of market uncertainty, market risk, execution uncertainty, and execution risk. And these are all unique concerns that require different considerations and reactions. Uncertainty can be mitigated through a process of enactment, selection, and retention (Aldrich & Ruef, 2006; Weick et al., 2005), but the organizations must resist formalizing specific jobs until they have sufficiently experimented to achieve what I called product solution match. This involves getting close to the core of the business idea to determine what range of repertoires can be converted into task stacks, while also interacting with the environment sufficiently to understand how a technological solution can be brought to bear in the contexts in which entrepreneurs are operating.

Contributions

I believe that one of the largest contributions that I make with this research is methodological. The use of natural language processing is still in its early days in organizations and management studies. In this research, I demonstrate how powerful it can be in examining questions that have long interested scholars but have perhaps resisted answers due to insufficient methodological resources. A second but related contribution is the further explanation of the task structure, or task stack of work. This is a contribution to the broader jobs literature within organizations and management, as it specifies a method for uncovering the task structure of work. In the past, we have relied on job titles (Bielby & Baron, 1986) or purely qualitative methods (L. E. Cohen, 2013) to show this structure. The natural language methods here appear to be an improvement. But this demonstration of the task structure of jobs can also be useful in analyses beyond the jobs and labor literature. There is no boss of the economy, and it is possible that as certain tasks move out of a given job's task stack, other jobs may not absorb them. For example, as agriculture inspectors collect less safety data, where does that task go? And if no other jobs acquire that task, what are the implications? The job task structure mapping can be used to show where there is institutional white space in our labor and economic structure. This study's third contribution is in the examination of the entrepreneurial construction of jobs. The case study demonstrates that not only is it difficult to formalize jobs in a nascent organization but also that early formalization can actually be counterproductive to the company's goals. Until the company's product offering has crystallized, the jobs in the organization are still being revised. To try to formalize jobs before product market fit has been reached may result in conflict, as the jobs people were hired for may not actually match what they are being asked to do. Alongside the flux in task stack, the technology stack workers use is in a state of revision as well. And it is only through the on-the-job, in-the-field experience that the necessity and type of revision can be observed and enacted.

Limitations

One of the limitations of the dissertation, as discussed earlier, is largely one of data collection. The computational paper would be better served by a more expansive data set, whereas the qualitative study would be better served by more access to employees across all levels of the entrepreneurial organizations. A second limitation is that I did not test mechanisms directly. Past research has theorized job crafting as a potential mechanism for why jobs change over time (Wrzesniewski & Dutton, 2001). I accept this without testing for it directly. Further explanations for the mechanisms entrepreneurs use to transition between the different stages of risk and uncertainty would also make the findings in the case study stronger. And in fact one of the ways that Alvarez et al. (2017) propose to get past the entrenched ontological argument of discovered opportunities vs. created opportunities is deeper empirical work on the specific mechanisms behind entrepreneurial enactment. My initial findings suggest that by more specifically tracking the evolution of jobs within a new organization, we can gain better insight into entrepreneurial enactment.

Future Work

Future work will be designed to address the limitations of this dissertation. To fix the issue of potentially limited data for the computational study, I may try to acquire more expansive data sets. A company called Burning Glass maintains one. Alternatively, forging partnerships with LinkedIn and other job hiring sites may be productive in getting a larger data set. And while it would be resource-intensive, data mining of newspaper job listings could potentially generate an great extensive historical data set to see how language around jobs has changed over a longer time frame. To superficially get at potential mechanisms for how jobs change, I will definitely look to using the job transition data from the CPS, which is very high in my research agenda in the coming years.

A multiple case study, over a longer period time, that encompasses a broader swath of the employees in start-up organizations, would be useful in gaining a better understanding of how these organizations specifically structure themselves in order to achieve problem-solution match. The longer time frame would provide more opportunities to see the job structuring process in action. And I can potentially exploit key differences among the start-ups to proceed to more explanatory work of job structuring and entrepreneurial engagement rather than the pure descriptive work of the case study in this dissertation.

Finally, while it does not address a specific limitation of my present research, the visualizations in the computational chapter are potentially very fertile ground for future studies. Knowing the relative similarity between jobs can be an input for policy research. Economic inequality is a pressing issue in our time. To the extent that low-wage jobs cluster together in terms of similarities, my visualizations may show that there may effectively be "disadvantaged job neighborhoods" in the occupational structure of the United States. On a positive note, knowing how far apart jobs are in terms of their underlying task structure may help provide a roadmap for how to use training and education more effectively to transition people to more remunerative kinds of work.

Concluding Remarks

The ways in which work is structured and how it intersects with technological change will continue to be one of the foremost issues of our time. Understanding how jobs and technology intersect and at what levels may be key in ensuring that economic opportunity exists for everyone. This dissertation overall advances our understanding of job structure and change and in so doing also provides a toolkit for continuing this work in the future.

Bibliography

Abbott, A. (2014). *The system of professions: An essay on the division of expert labor*. University of Chicago Press.

Alabdulkareem, A., Frank, M. R., Sun, L., Al Shebli, B., Hidalgo, C., & Rahwan, I. (2018). Unpacking the polarization of workplace skills. *Science Advances*, *4*(7), eaao6030.

Aldrich, H. E., & Ruef, M. (2006). *Organizations Evolving* (2nd edition). Sage.

Alvarez, S. A., & Barney, J. B. (2010). Entrepreneurship and epistemology: The philosophical underpinnings of the study of entrepreneurial opportunities. *Academy of Management Annals*, *4*(1), 557–583.

Alvarez, S. A., Barney, J. B., McBride, R., & Wuebker, R. (2017). On opportunities: Philosophical and empirical implications. *Academy of Management Review*, *42*(4), 726–730.

Arthur, W. B. (2009). *The nature of technology: What it is and how it evolves*. Simon and Schuster.

Ashforth, B. E., & Kreiner, G. E. (1999). "How can you do it?": Dirty work and the challenge of constructing a positive identity. *Academy of Management Review*, *24*(3), 413–434.

Ashforth, B. E., Kreiner, G. E., Clark, M., & Fugate, M. (2007). Normalizing dirty work: Managerial tactics for countering occupational taint. *Academy of Management Journal*, *50*(1), 149–174.

Attewell, P., & Rule, J. (1984). Computing and organizations: What we know and what we don't know. *Communications of the ACM*, *27*(12), 1184–1192.

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

Baker, W. E. (1984). The social structure of a national securities market. *American Journal of Sociology*, *89*(4), 775–811.

Barley, S. R., & Kunda, G. (2001). Bringing work back in. *Organization Science*, *12*(1), 76–95.

Baron, J. N., Burton, M. D., & Hannan, M. T. (1996). The road taken: Origins and evolution of employment systems in emerging companies. *Industrial and Corporate Change*, *5*(2), 239–275.

Baron, J. N., Hannan, M. T., & Burton, M. D. (1999). Building the iron cage: Determinants of managerial intensity in the early years of organizations. *American Sociological Review*, 527–547.

Barton, M. A., Sutcliffe, K. M., Vogus, T. J., & DeWitt, T. (2015). Performing under uncertainty: Contextualized engagement in wildland firefighting. *Journal of Contingencies and Crisis Management*, *23*(2), 74–83.

Beane, M. (2018). Shadow Learning: Building Robotic Surgical Skill When Approved Means Fail. *Administrative Science Quarterly*, 0001839217751692.

Berger, P., & Luckmann, T. (1966). *The Social Construction of Reality*.

Bhaskar, R. (2013). *A realist theory of science*. Routledge.

Bielby, W. T., & Baron, J. N. (1986). Men and women at work: Sex segregation and statistical discrimination. *American Journal of Sociology*, *91*(4), 759–799.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, *3*(Jan), 993–1022.

Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., … Zhang, J. (2016). End to end learning for self-driving cars. *ArXiv Preprint ArXiv:1604.07316*.

Braun, V., & Clarke, V. (2012). Thematic Analysis. In H. Cooper (Ed.), *APA Handbook of Research Methods in Psychology* (Vol. 2, pp. 57–71). American Psychological Association.

Braun, V., & Clarke, V. (2013). *Successful Qualitative Research* (1st ed.). London: SAGE Publications, Ltd.

Braver, L., & Danneels, E. (2018). Propensities return us to the discovery-creation debate about entrepreneurial opportunities. *Academy of Management Review*, *43*(4), 812–815.

Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

Bush, G. (1993). *Public Papers of the Presidents of the United States: George Bush, 1992-1993*. Best Books on.

Butler, S. K., & Harvey, R. J. (1988). A comparison of holistic versus decomposed rating of Position Analysis Questionnaire work dimensions. *Personnel Psychology*, *41*(4), 761–771.

Cetina, K. K. (2012). What is a financial market? Global markets as microinstitutional and post-traditional social forms. *The Oxford Handbook of the Sociology of Finance*, 115.

Chang, J., Gerrish, S., Wang, C., Boyd-Graber, J. L., & Blei, D. M. (2009). Reading tea leaves: How humans interpret topic models. *Advances in Neural Information Processing Systems*, 288–296.

Chollet, F. (2017). *Deep learning with python*. Manning Publications Co.

Christianson, M. K., Farkas, M. T., Sutcliffe, K. M., & Weick, K. E. (2009). Learning through rare events: Significant interruptions at the Baltimore & Ohio Railroad Museum. *Organization Science*, *20*(5), 846–860.

Chui, M., Manyika, J., & Miremadi, M. (2016). Where machines could replace humans—and where they can't (yet). *McKinsey Quarterly*, *7*.

Clemons, E. K., & Row, M. C. (1991). Sustaining IT advantage: The role of structural differences. *MIS Quarterly*, *15*(3).

Coase, R. H. (1946). The marginal cost controversy. *Economica*, *13*(51), 169–182.

Cohen, L. E. (2013). Assembling Jobs: A Model of How Tasks Are Bundled Into and Across Jobs. *Organization Science*, *24*(2), 432–454. https://doi.org/10.1287/orsc.1110.0737

Cohen, L. E. (2016a). Introduction: Brining Jobs Back In: Toward A New Multilevel Approach to the Study of Work ad Organizations. In L. E. Cohen, M. D. Burton, & M. Lounsbury (Eds.), *The structuring of work in organizations* (pp. 1–24). Emerald Group Publishing Limited.

Cohen, L. E. (2016b). Jobs as Gordian Knots: A New Perspective Linking Individuals, Tasks, Organizations, and Institutions. In L. E. Cohen, M. D. Burton, & M. Lounsbury (Eds.), *The structuring of work in organizations* (pp. 25–59). Emerald Group Publishing Limited.

Cohen, M. D., Burkhart, R., Dosi, G., Egidi, M., Marengo, L., Warglien, M., & Winter, S. (1996). Routines and other recurring action patterns of organizations: Contemporary research issues. *Industrial and Corporate Change*, *5*(3), 653–698.

Cornelissen, J. P., & Durand, R. (2014). Moving forward: Developing theoretical contributions in management studies. *Journal of Management Studies*, *51*(6), 995–1022.

Croidieu, G., & Kim, P. H. (2018). Labor of love: Amateurs and lay-expertise legitimation in the early US radio field. *Administrative Science Quarterly*, *63*(1), 1–42.

Davenport, T. H., & Patil, D. J. (2012, October 1). Data Scientist: The Sexiest Job of the 21st Century. *Harvard Business Review*, (October 2012). Retrieved from https://hbr.org/2012/10/data-scientist-the-sexiest-job-of-the-21st-century

Davis, G. F. (2016). *The vanishing American corporation: Navigating the hazards of a new economy*. Berrett-Koehler Publishers.

Davis, G. F., Diekmann, K. A., & Tinsley, C. H. (1994). The decline and fall of the conglomerate firm in the 1980s: The deinstitutionalization of an organizational form. *American Sociological Review*, *59*, 547–547.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, *41*(6), 391–407.

Dewett, T., & Jones, G. R. (2001). The role of information technology in the organization: A review, model, and assessment. *Journal of Management*, *27*(3), 313–346.

Dierdorff, E. C., & Morgeson, F. P. (2007). Consensus in work role requirements: The influence of discrete occupational context on role expectations. *Journal of Applied Psychology*, *92*(5), 1228.

Dierdorff, E. C., & Wilson, M. A. (2003). A meta-analysis of job analysis reliability. *Journal of Applied Psychology*, *88*(4), 635.

DiMaggio, P., & Powell, W. W. (1983). The iron cage revisited: Collective rationality and institutional isomorphism in organizational fields. *American Sociological Review*, *48*(2), 147–160.

Dunnette, M. D. (1966). *Personnel Selection and Placement*. Belmont, CA: Wadsworth Publishing Company.

Duriau, V. J., Reger, R. K., & Pfarrer, M. D. (2007). A content analysis of the content analysis literature in organization studies: Research themes, data sources, and methodological refinements. *Organizational Research Methods*, *10*(1), 5–34.

Eisenhardt, K. M., & Graebner, M. E. (2007). Theory building from cases: Opportunities and challenges. *Academy of Management Journal*, *50*(1), 25–32.

ETA O*NET Data Collection Program, Employment & Training Administration (ETA) U.S. Department of Labor. (n.d.). Retrieved July 21, 2019, from https://www.doleta.gov/programs/onet/datacollection.cfm

Faraj, S., & Azad, B. (2012). The Materiality of Technology: An Affordance Perspective. In P. Leonardi, B. Nardi, & J. Kallinikos (Eds.), *Materiality and Organizing: Social Interaction in a Technological World* (1st ed., pp. 217–236). United Kingdom: Oxford University Press.

Faulkner, P., & Runde, J. (2012). On Sociomateriality. In P. Leonardi, B. Nardi, & J. Kallinikos (Eds.), *Materiality and Organizing: Social Interaction in a Technological World* (1st ed., pp. 49–66). United Kingdom: Oxford University Press.

Feldman, M. S. (2000). Organizational routines as a source of continuous change. *Organization Science*, *11*(6), 611–629.

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, *114*, 254–280.

Fried, J., & Heinemeier-Hansson, D. (2010). *Rework*. Random House.

Ghoshal, S., & Bartlett, C. A. (1994). Linking organizational context and managerial action: The dimensions of quality of management. *Strategic Management Journal*, *15*(S2), 91–112.

Gibson, J. J. (2014). *The ecological approach to visual perception: Classic edition*. Psychology Press.

Goos, M., & Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *The Review of Economics and Statistics*, *89*(1), 118–133.

Gorski, P. S. (2013). *What is critical realism? And why should you care?* SAGE Publications Sage CA: Los Angeles, CA.

Grant, A. M., Campbell, E. M., Chen, G., Cottone, K., Lapedis, D., & Lee, K. (2007). Impact and the art of motivation maintenance: The effects of contact with beneficiaries on persistence behavior. *Organizational Behavior and Human Decision Processes*, *103*(1), 53–67.

Grant, A. M., & Parker, S. K. (2009). 7 redesigning work design theories: The rise of relational and proactive perspectives. *Academy of Management Annals*, *3*(1), 317–375.

Graves, A., & Jaitly, N. (2014). Towards end-to-end speech recognition with recurrent neural networks. *International Conference on Machine Learning*, 1764–1772.

Greenwood, R., & Suddaby, R. (2006). Institutional entrepreneurship in mature fields: The big five accounting firms. *Academy of Management Journal*, *49*(1), 27–48.

Grimmer, J., & Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, *21*(3), 267–297.

Hackman, J. R., & Oldham, G. R. (1976). Motivation through the design of work: Test of a theory. *Organizational Behavior and Human Performance*, *16*(2), 250–279.

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

Harvey, R. (1990). Job Anaysis. In M. D. Dunnette & L. Hough (Eds.), *Handbook of Industrial and Organizational Psychology* (2nd ed., Vol. 2, pp. 71–164). Consulting Psychologists Press, Inc.

Harvey, R. J., & Wilson, M. A. (2000). Yes Virginia, there is an objective reality in job analysis. *Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior*, *21*(7), 829–854.

Hoberg, G., & Phillips, G. (2010). Product market synergies and competition in mergers and acquisitions: A text-based analysis. *The Review of Financial Studies*, *23*(10), 3773–3811.

Hornick, M. (2016, February 5). Using SVD for Dimensionality Reduction. Retrieved June 11, 2019, from https://blogs.oracle.com/r/using-svd-for-dimensionality-reduction

Huang, A. H., Lehavy, R., Zang, A. Y., & Zheng, R. (2017). Analyst information discovery and interpretation roles: A topic modeling approach. *Management Science*, *64*(6), 2833–2855.

Huber, G. P. (1990). A theory of the effects of advanced information technologies on organizational design, intelligence, and decision making. *Academy of Management Review*, *15*(1), 47–71.

Hutto, C. J., & Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. *Eighth International AAAI Conference on Weblogs and Social Media*.

Ilgen, D. R., & Hollenbeck, J. R. (1991). The structure of work: Job design and roles. In *Handbook of industrial and organizational psychology* (2nd ed., Vol. 2, pp. 165–207). Consulting Psychologists Press.

Jurafsky, D. (2000). *Speech & language processing*. Pearson Education India.

Kane, A. A., Argote, L., & Levine, J. M. (2005). Knowledge transfer between groups via personnel rotation: Effects of social identity and knowledge quality. *Organizational Behavior and Human Decision Processes*, *96*(1), 56–71.

Kaplan, S., & Vakili, K. (2015). The double-edged sword of recombination in breakthrough innovation. *Strategic Management Journal*, *36*(10), 1435–1457.

Katz, L. F., & Krueger, A. B. (2016). *The rise and nature of alternative work arrangements in the United States, 1995-2015*. National Bureau of Economic Research.

Katz, L. F., & Krueger, A. B. (2019). *Understanding trends in alternative work arrangements in the United States*. National Bureau of Economic Research.

Kessler, S. (2018). *Gigged: The end of the job and the future of work*. New York: St. Martin's Press.

Ketokivi, M., Mantere, S., & Cornelissen, J. (2017). REASONING BY ANALOGY AND THE PROGRESS OF THEORY. *Academy of Management Review*, *42*(4).

Kirzner, I. M. (1997). Entrepreneurial discovery and the competitive market process: An Austrian approach. *Journal of Economic Literature*, *35*(1), 60–85.

Knight, F. H. (2012). *Risk, uncertainty and profit*. Courier Corporation.

Knorr Cetina, K., & Bruegger, U. (2002). Global microstructures: The virtual societies of financial markets. *American Journal of Sociology*, *107*(4), 905–950.

Lane, H., Howard, C., & Hapke, H. M. (2019). *Natural Language Processing in Action*. Shelter Island, NY: Manning Publications Co.

Latour, B. (1991). Technology is society made durable. In J. Law (Ed.), *A Sociology of Monsters: Essays on Power, Technology and Domination* (pp. 103–131). Routledge.

Law, J. (Ed.). (1991). *A Sociology of Monsters: Essays on Power, Technology and Domination*. Routledge.

Le, Q., & Mikolov, T. (2014). Distributed Representations of Sentences and Documents. *International Conference on Machine Learning*, 1188–1196.

Leana, C., Appelbaum, E., & Shevchuk, I. (2009). Work process and quality of care in early childhood education: The role of job crafting. *Academy of Management Journal*, *52*(6), 1169–1192.

Leonardi, P. (2012). Materiality, Sociomateriality, and Socio-Technical Systems: What Do These Terms Mean? How Are They Different? Do We Need Them? In P. Leonardi, B. Nardi, & J. Kallinikos (Eds.), *Materiality and Organizing: Social Interaction in a Technological World* (1st ed., pp. 3–24). United Kingdom: Oxford University Press.

Leonardi, P. (2017). Methodological Guidelines for the Study of Materiality and Affordances. In R. Mir & S. Jain (Eds.), *The Routledge Companion to Qualitative Research in Organization Studies* (pp. 279–290). Routledge.

Levy, F., & Murnane, R. J. (2005). *The new division of labor: How computers are creating the next job market*. Princeton, NJ: Princeton University Press.

Lyons, P. (2008). The crafting of jobs and individual differences. *Journal of Business and Psychology*, *23*(1–2), 25–36.

Maaten, L. van der, & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, *9*(Nov), 2579–2605.

MacKenzie, D., & Millo, Y. (2003). Constructing a market, performing theory: The historical sociology of a financial derivatives exchange. *American Journal of Sociology*, *109*(1), 107–145.

Maitlis, S., & Christianson, M. (2014). Sensemaking in organizations: Taking stock and moving forward. *Academy of Management Annals*, *8*(1), 57–125.

Manjoo, F. (2019, April 10). I Didn't Write This Column. I Spoke It. *The New York Times*. Retrieved from https://www.nytimes.com/2019/04/10/opinion/screenless-internet.html

McAfee, A., & Brynjolfsson, E. (2017). *Machine, platform, crowd: Harnessing our digital future*. WW Norton & Company.

McCormick, E. J. (1976). Job and task analysis. *Handbook of Industrial and Organizational Psychology*, 651–696.

Michael, J. (1997). Environment and land-use: The economic development of the communities who built Stonehenge (an economy to support the stones). *Proceedings of the British Academy*, *92*, 115–144.

Merton, R. C. (1973). Theory of rational option pricing. *Theory of Valuation*, 229–288.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, 3111–3119.

Mimno, D., & Blei, D. (2011). Bayesian checking for topic models. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 227–237. Association for Computational Linguistics.

Mimno, D., Wallach, H. M., Talley, E., Leenders, M., & McCallum, A. (2011). Optimizing semantic coherence in topic models. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 262–272. Association for Computational Linguistics.

Miner, A. S. (1987). Idiosyncratic jobs in formalized organizations. *Administrative Science Quarterly*, 327–351.

Miner, A. S. (1990). Structural evolution through idiosyncratic jobs: The potential for unplanned learning. *Organization Science*, *1*(2), 195–210.

Miner, A. S. (1991). Organizational evolution and the social ecology of jobs. *American Sociological Review*, 772–785.

Nelkin, D. (2004). The Organization of a Methadone Program. In R. K. Yin (Ed.), *Case Study Anthology* (pp. 205–218). Thousand Oaks, CA: SAGE Publicatons.

Oldham, G. R., & Fried, Y. (2016). Job design research and theory: Past, present and future. *Organizational Behavior and Human Decision Processes*, *136*, 20–35.

Oldham, G. R., & Hackman, J. R. (2010). Not what it was and not what it will be: The future of job design research. *Journal of Organizational Behavior*, *31*(2-3), 463–479.

Olson, G. M., & Olson, J. S. (2000). Distance matters. *Human-Computer Interaction*, *15*(2), 139–178.

O*NET Resource Center Content Model. (2018, May 30). Retrieved May 30, 2018, from https://www.onetcenter.org/content.html

Orlikowski, W. J. (2007). Sociomaterial Practices: Exploring Technology at Work. *Organization Studies*, *28*(9), 1435–1448. https://doi.org/10.1177/0170840607081138

Orlikowski, W. J., & Scott, S. V. (2008). 10 sociomateriality: Challenging the separation of technology, work and organization. *Academy of Management Annals*, *2*(1), 433–474.

Orlikowski, W. J., & Scott, S. V. (2015). Exploring material-discursive practices. *Journal of Management Studies*, *52*(5), 697–705.

Orlikowski, W. J., Yates, J., Okamura, K., & Fujimoto, M. (1995). Shaping electronic communication: The metastructuring of technology in the context of use. *Organization Science*, *6*(4), 423–444.

Ozcan, P., Han, S., & Graebner, M. E. (2017). Single Cases. In R. Mir & S. Jain (Eds.), *The Routledge Companion to Qualitative Research in Organizations*. Routledge.

Parker, S. K., Van den Broeck, A., & Holman, D. (2017). Work design influences: A synthesis of multilevel factors that affect the design of jobs. *Academy of Management Annals*, *11*(1), 267–308.

Pentland, B. T., & Rueter, H. H. (1994). Organizational routines as grammars of action. *Administrative Science Quarterly*, 484–510.

Perrow, C. (1986). *Complex Organizations: A Critical Essay* (3rd ed.). McGraw-Hill, Inc. *Personal Communication*. (2017, July 18).

Petriglieri, G., Ashford, S. J., & Wrzesniewski, A. (2019). Agony and ecstasy in the gig economy: Cultivating holding environments for precarious and personalized work identities. *Administrative Science Quarterly*, *64*(1), 124–170.

Planet Money. (2015). Map: The Most Common* Job In Every State. Retrieved June 11, 2019, from NPR.org website: https://www.npr.org/sections/money/2015/02/05/382664837/map-the-most-common-job-in-every-state

Powell, T. C., & Dent-Micallef, A. (1997). Information technology as competitive advantage: The role of human, business, and technology resources. *Strategic Management Journal*, 375–405.

Pritchard, J. K., Stephens, M., & Donnelly, P. (2000). Inference of population structure using multilocus genotype data. *Genetics*, *155*(2), 945–959.

Programming: Jane Street. (n.d.). Retrieved June 9, 2019, from https://www.janestreet.com/programming/

Questionnaires: OE - Tasks, 2018 OMB Clearance Package. (2018, October 1). Retrieved from https://www.onetcenter.org/ombclearance.html

Ramoglou, S., & Tsang, E. W. (2016). A realist perspective of entrepreneurship: Opportunities as propensities. *Academy of Management Review*, *41*(3), 410–434.

Ramoglou, S., & Tsang, E. W. (2018). Opportunities Lie in the Demand Side: Transcending the Discovery-Creation Debate. *Academy of Management Review*, *43*(4), 815–818.

Rosling, H., & Zhang, Z. (2011). Health advocacy with Gapminder animated statistics. *Journal of Epidemiology and Global Health*, *1*(1), 11–14.

Rosso, B. D., Dekas, K. H., & Wrzesniewski, A. (2010). On the meaning of work: A theoretical integration and review. *Research in Organizational Behavior*, *30*, 91–127.

Sanchez, J. I., & Levine, E. L. (2000). Accuracy or consequential validity: Which is the better standard for job analysis data? *Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior*, *21*(7), 809–818.

Shane, S. (2000). Prior knowledge and the discovery of entrepreneurial opportunities. *Organization Science*, *11*(4), 448–469.

Shane, S., & Venkataraman, S. (2000). The promise of entrepreneurship as a field of research. *Academy of Management Review*, *25*(1), 217–226.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*(3), 379–423.

Shell, E. R. (2018). *The Job: Work and Its Future in a Time of Radical Change*. Random House.

Shepherd, D. A., & Sutcliffe, K. M. (2011). Inductive top-down theorizing: A source of new theories of organization. *Academy of Management Review*, *36*(2), 361–380.

Shepherd, D. A., & Sutcliffe, K. M. (2015). The use of anthropomorphizing as a tool for generating organizational theories. *Academy of Management Annals*, *9*(1), 97–142.

Smith, A., & McCulloch, J. R. (1838). *An Inquiry into the Nature and Causes of the Wealth of Nations*. A. and C. Black and W. Tait.

Sorenson, O., & Stuart, T. E. (2008). 12 Entrepreneurship: A field of dreams? *Academy of Management Annals*, *2*(1), 517–543.

Sparck Jones, K. (1972). A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation*, *28*(1), 11–21.

Strang, D., & Baron, J. N. (1990). Categorical imperatives: The structure of job titles in California state agencies. *American Sociological Review*, 479–495.

Strauss, A., & Corbin, J. M. (1997). *Grounded theory in practice*. Sage.

Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, *29*(1), 24–54.

Taylor, F. W. (1914). *The principles of scientific management*. Harper.

Trask, A. (2019). *Grokking Deep Learning*. Shelter Island, NY: Manning Publications Co.

Vaughan, D. (1996). *The Challenger launch decision: Risky technology, culture, and deviance at NASA*. University of Chicago press.

Vaughn, D. (2014). Analogy, Cases and Comparative Social Organization. In R. Swedberg (Ed.), *Theorizing in Social Science: The Context of Discovery*. Stanford, CA: Stanford University Press.

Weick, K. E. (1979). *The Social Psychology of Organizing*. New York: McGraw-Hill, Inc.

Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (2005). Organizing and the process of sensemaking. *Organization Science*, *16*(4), 409–421.

Williamson, O. E. (1975). Markets and hierarchies. *New York*, *2630*.

Williamson, O. E. (1981). The economics of organization: The transaction cost approach. *American Journal of Sociology*, *87*(3), 548–577.

Wrzesniewski, A., & Dutton, J. E. (2001). Crafting a job: Revisioning employees as active crafters of their work. *Academy of Management Review*, *26*(2), 179–201.

Wrzesniewski, A., Dutton, J. E., & Debebe, G. (2003). Interpersonal sensemaking and the meaning of work. *Research in Organizational Behavior*, *25*, 93–135.

Yin, R. K. (2003). *Applications of Case Study Research* (2nd ed.). Thousand Oaks, CA: SAGE Publications, Inc.

Yin, R. K. (2012). Case Study Methods. In H. Cooper (Ed.), *APA Handbook of Research Methods in Psychology* (Vol. 2, pp. 141–155). American Psychological Association.

Zammuto, R. F., Griffith, T. L., Majchrzak, A., Dougherty, D. J., & Faraj, S. (2007). Information technology and the changing fabric of organization. *Organization Science*, *18*(5), 749–762.

Zuboff, S. (1988). In the age of the smart machine: The future of work and power. *New York: Basic*.