

University of Denver

Digital Commons @ DU

Electronic Theses and Dissertations

Graduate Studies

1-1-2019

Automation and Adaptation: Information Technology, Work Practices, and Labor Demand at Three Firms

Spencer J. Rockwell
University of Denver

Follow this and additional works at: <https://digitalcommons.du.edu/etd>



Part of the [Labor Economics Commons](#)

Recommended Citation

Rockwell, Spencer J., "Automation and Adaptation: Information Technology, Work Practices, and Labor Demand at Three Firms" (2019). *Electronic Theses and Dissertations*. 1546.
<https://digitalcommons.du.edu/etd/1546>

This Thesis is brought to you for free and open access by the Graduate Studies at Digital Commons @ DU. It has been accepted for inclusion in Electronic Theses and Dissertations by an authorized administrator of Digital Commons @ DU. For more information, please contact jennifer.cox@du.edu, dig-commons@du.edu.

Automation and Adaptation: Information Technology, Work Practices, and Labor Demand at Three Firms

Abstract

The use of information and communication technology to automate routine tasks involves two types of innovation: technological and organizational. Together, improvements in technological capabilities and complementary changes made by firms in the way they organize work and implement work practices constitute the conditions under which machines substitute for or complement human workers. Building on the prevailing model of routine-biased technical change and recent insights into organizational complementarities, I conduct three qualitative case studies in health care and real estate to assess the relationship between technology and firm-level labor demand. Unique combinations of technological innovation, organizational complementarity, and decision-making at each firm produce differential impacts for labor demand, with even similar technologies exhibiting quite different patterns of substitution for workers of all skill types. In addition, studying firm-level complementarities illuminates how and why the scope of the routine task may be growing, with particularly important implications for relatively higher skill workers.

Document Type

Thesis

Degree Name

M.A.

Department

Economics

First Advisor

Markus Schneider, Ph.D.

Second Advisor

Kimon Valavanis, Ph.D.

Third Advisor

Christine Ngo, Ph.D.

Keywords

Information technology, Communication technology

Subject Categories

Economics | Labor Economics

Publication Statement

Copyright is held by the author. User is responsible for all copyright compliance.

Automation and Adaptation:
Information Technology, Work Practices, and Labor Demand at Three Firms

A Thesis
Presented to
the Faculty of Arts and Humanities
University of Denver

In Partial Fulfillment
of the Requirements for the Degree
Master of Arts

by
Spencer J Rockwell
March 2019
Advisor: Dr. Markus Schneider

Author: Spencer J Rockwell

Title: Automation and Adaptation: Information Technology, Work Practices, and Labor Demand at Three Firms

Advisor: Dr. Markus Schneider

Degree Date: March 2019

Abstract

The use of information and communication technology to automate routine tasks involves two types of innovation: technological and organizational. Together, improvements in technological capabilities and complementary changes made by firms in the way they organize work and implement work practices constitute the conditions under which machines substitute for or complement human workers. Building on the prevailing model of routine-biased technical change and recent insights into organizational complementarities, I conduct three qualitative case studies in health care and real estate to assess the relationship between technology and firm-level labor demand. Unique combinations of technological innovation, organizational complementarity, and decision-making at each firm produce differential impacts for labor demand, with even similar technologies exhibiting quite different patterns of substitution for workers of all skill types. In addition, studying firm-level complementarities illuminates how and why the scope of the routine task may be growing, with particularly important implications for relatively higher skill workers.

Acknowledgements

I am grateful for the assistance and support of several people, without whom this thesis would not have come together. First, I'd like to thank my wife, Cortney, who provided endless encouragement and moral support throughout my graduate studies, and who put me in touch with the right people to connect with two of my interviewees. I'd also like to thank several members of the Economics faculty for the special support they provided during my journey: Dr. Christine Ngo, for helping get this project off the ground and pushing me to conduct original research; Dr. Peter Ho, for the many casual conversations, both on- and off-topic, which encouraged me to think about the history of economic ideas and economic development differently; and Dr. Markus Schneider, who graciously accepted me as an advisee in a pinch and helped me see this project to the end. Lastly, I'd like to thank my interviewees, Janet Garret, Andy Barkett, and Steve Hess, for offering their time, expertise, and data.

Table of Contents

1. Introduction	1
2. Literature Review	7
2.1 <i>Technological Unemployment</i>	7
2.2 <i>Modeling Trends in Labor Demand and Employment</i>	15
2.3 <i>Technology Adoption and Complementarities at the Firm Level</i>	21
3. Methodology	30
3.1 <i>Theoretical Framework</i>	30
3.2 <i>Data Collection and Analysis</i>	33
3.3 <i>Case Selection</i>	35
4. Case Studies	38
4.1 <i>Background</i>	38
4.2 <i>REX Real Estate Exchange</i>	46
4.3 <i>Sky Ridge Medical Center</i>	51
4.4 <i>UCHealth</i>	57
5. Discussion of Findings	63
5.1 <i>Within-Case Evidence</i>	63
5.2 <i>Cross-Case Comparison</i>	66
6. Conclusion	73
Bibliography	75

1. Introduction

A new wave of “automation anxiety” is currently sweeping the popular discourse on technology, reigniting fears of an imminent employment crisis and resurrecting old debates about redistributive social policies in the world’s advanced economies (Akst, 2013; Avent, 2016; James, 2016). One of the most commonly cited studies estimates that 47% of jobs in the US are at risk of automation in the next two decades (Frey & Osborne, 2017). Another widely cited study of local labor markets in the US finds that the introduction of one industrial robot is associated with the elimination of six jobs, even after controlling for factors like global trade, offshoring, industry structure, and demographics (Acemoglu & Restrepo, 2017). Dramatic statistics such as these make for stunning headlines, especially alongside stories about the progressive encroachment of artificial intelligence into domains once considered uniquely human, such as visual recognition. As our technological capabilities grow, so does the apprehension that human workers are becoming hopelessly redundant.

Economists have long debated the relationship between technological change and employment. Perilous predictions about the end of work, dating from the first industrial revolution to the modern digital era, have so far failed to materialize (Mokyr, Vickers, & Ziebarth, 2015). Nonetheless, the emergence of contemporary innovations in areas such as robotics and machine learning prompts us to ask

whether this time is different. If it is, and given the dizzying pace of technological change today, can the past be any guide to the future?

Even after two hundred years of engagement with the questions invoked by technological change, economists have much to learn about how and for whom automation affects employment. Recent research in labor economics has illuminated an alarming trend of the past quarter century: information and communication technologies (ICT), namely computers, increasingly substituting for many of the routine administrative tasks previously performed by skilled professionals, thereby “hollowing out” the middle of the employment distribution (Acemoglu & Autor, 2010; Autor, 2015). Models of routine-biased technical change help to sketch the broad contours of this dilemma, but in so doing reduce the complexity and diversity of the process of technological adoption across firms to a monolithic phenomenon. At what point does the automation of certain tasks lead to complete substitution for a human worker?

This question points to the broader issue motivating this thesis. Firm-level decisions about why, when, and how to adopt a new technology loom large in determining the specific impact of that technology on outcomes like productivity, business performance, and labor demand (Brynjolfsson, 2010). After all, few technologies are “plug and play.” Empirical evidence demonstrates that successful adopters capitalize on their new technological capabilities by implementing an array of complementary changes to work practices, workplace organization, and even market strategy (Bartel, Ichniowski, & Shaw, 2007; Brynjolfsson & Hitt, 2000). By

taking into account this more encompassing concept of technological change within the firm, economists are better able to identify and investigate the causal mechanisms linking the substitutive (or complementary) properties of technologies with the skill content and magnitude of labor demand (Bresnahan, Brynjolfsson, & Hitt, 2002).

I therefore pose the following research question: How does the adoption of technology and implementation of complementary work practices affect firms' demand for labor? Specifically, I conduct qualitative case studies to investigate the role of ICTs, ranging in technological complexity from DOS-based computer systems to advanced applications of machine learning, in influencing labor demand among three service firms. Though the small sample size limits and methodological approach limits the generalizability of my findings, my research provides a detailed look at the microeconomic factors at work in broader trends related to automation. Indeed, while aggregate analysis remains critical to drawing conclusions about the general direction of technological change and employment, macroeconomic statistics as well as industry-level studies often fail to capture the mechanisms by which new technologies interact with work (Bresnahan et al., 2002; Seamans, 2018).

The case studies examine firms in two previously understudied service industries: health care and real estate. Both play a significant role in the US economy today. In 2017, they combined to generate roughly \$4 trillion in output, just more than one-fifth of US gross domestic product. Real estate constitutes the largest subindustry in the US economy in terms of value-added; since 2008, its output has

exceeded that of the entire manufacturing sector. From 1997 to 2017, value-added output grew more than 180% in health care, fourth most among subindustries. Health care in particular has provided an engine of employment growth for the last two decades. Since 1998, the industry has added 7.1 million jobs, or 35% of all jobs added during that time. Its employment growth rate of 57% is nearly quadruple that of all domestic industries combined. One in eight Americans workers were employed in health care as of 2018; the Bureau of Labor Statistics predicts that 9 of the 15 fastest growing occupations from 2016 to 2026 will be in health care, including jobs like Physician Assistant, Nurse Practitioners, medical assistants, and home health aides (Bureau of Economic Analysis, 2018; Bureau of Labor Statistics, 2018).

Beyond the numbers, both industries potentially represent important proving grounds for the nature and magnitude of automation in post-industrial, service economies. Already, many middle-skill occupations are experiencing disruption as a result of routine-biased technical change. What happens if and when the scope of routinization grows due to innovation in both technology and the organization of work? That is, as additional tasks become subject to machine substitution, what does this mean for industries and jobs previously thought to be relatively immune to automation? Occupations such as real estate agent, doctor, and nurse require, in addition to formal education and training, social and emotional intelligence, professional judgment, and critical thinking; success often seems to be driven by a certain *je ne sais quoi* that seems to be the sole province of the human mind. If the

tasks that make up such jobs become subject to automation, the implications for labor demand and employment may be more significant than previously imagined.

This thesis thus makes an important contribution to economists' understanding of labor demand dynamics. It applies a novel theoretical framework that combines the prevailing model of routine-biased technical change with research on firm-level complementarities to technologies and industries that have received very little attention in the economic literature on automation. Despite many studies on the effects of computers in the workplace, recent ICT innovations, especially machine learning, remain a gap in the literature, if not many a research agenda (Seamans, 2018). Likewise, the vast majority of research on the integration of ICT within work practices and organization focuses on production industries, to the neglect of firms in the service sector (for notable exceptions, see Barley (1986) and Autor, Levy, and Murnane (2001)). This study serves both to ground a more detailed understanding of how adoption plays out inside firms and an approach for applying a set of useful concepts and questions in unique contexts. Most importantly, in calling attention to the usefulness and shortcomings of the theoretical framework itself, it points toward important theoretical problems economists must confront as they seek to understand and predict the effect of automation on firm-level labor demand.

To be clear, the microeconomic research presented here does not allow me to comment directly on the influence of automation in patterns of aggregate employment. The latter are shaped by factors well beyond the scope of this study, especially macro-trends related to effective demand and globalization. Of course,

general equilibrium effects involving these and other factors are, in the final analysis, inseparable even from firm-level labor demand. I therefore focus on highly nontradable industries where unique variables of interest nonetheless provide distinct entry points for analysis so that macroeconomic issues can be effectively bracketed, i.e., as a link from labor demand to employment, rather than a causal mechanism driving labor demand itself.

2. Literature Review

2.1 Technological Unemployment

The debate about the relationship between technology and employment is as old as the dismal science itself. Broadly speaking, the arguments sort into two distinct viewpoints. The first, originally advanced by Classical and then Neoclassical economists, sees in technological progress short-run pain followed by long-run gain. While recognizing the potential for temporary dislocation, it posits an inevitable return to long-run labor market equilibrium along with vast improvements in general welfare that more than compensate for any short-run losses. The second view, held primarily by Keynes and Marx, prophesizes the imminent end of work, which must in turn precipitate significant social change. Yet Marx and Keynes diverge in their appraisal of such disruption, producing starkly different visions of a future without labor. Today, the controversy over technological progress centers on the question of whether or not this time (e.g., the fourth industrial revolution) is different from past moments of technological upheaval and anxiety. At the center of this controversy lie disputes about the nature of contemporary technology and the trajectory of labor productivity.

Ricardo (1817 [1973]), reflecting on the accelerating application of machinery during the early stages of capitalist development, acknowledged that machines threatened much of the industrial working class with temporary redundancy. This

result followed not from technological progress, however, but rather from his “wage-fund” theory. According to the theory, investment and wages drew from the same pool of accumulated capital; an increase in the former implied a reduction in the latter *pari passu*. Less funds available for wages meant falling employment, *ceteris paribus* (Berg, 1980). In the long run, Ricardo believed that machinery would increase productivity, accelerating capital accumulation and increasing savings, so that labor demand would eventually rebound.

Later heirs of the Classical tradition maintained Ricardo’s emphasis on the distinction between the short- and long-run consequences of technological progress even as they rejected his wage-fund theory. J.S. Mill (1848 [1965]) accepted that technical improvements in manufacturing might temporarily displace some workers, but doubted the possibility of substantial aggregate or long-term effects during the course of capitalist development. Knut Wicksell (1901 [1934]) made a similar argument in a characteristically Neoclassical way. He posited that technological change could “reduce a number of workers to beggary” by lowering the marginal product of labor and thus wages (p. 164). Unlike Ricardo, Wicksell did not explicitly draw the link from reduced wages to lower employment, likely because he made the strict assumption of flexible, self-equilibrating labor markets. Like Ricardo, he did forecast that rising capital accumulation would eventually raise wages again.

Karl Marx (1868 [1977]) was the first economist to take seriously the prospect of widespread technical substitution resulting in long-term, mass unemployment. He argued that the owners of capital invested in new machines in order to grow profits

through increased productivity and reduced costs. He also recognized that machinery played an important disciplinary role in capitalism, since capitalists could hold the threat of machine replacement over the head of labor to extract more effort (or the same effort for a lower wage). At first, Marx suggested, machines expanded the field of employment by turning workers into mere appendages, destined to work on dull, isolating, repetitive tasks; along with large-scale industry's voracious appetite for expansion, this opened up work to women and children. As the profit-motive drove increasing investment and innovation, however, capitalists would be unable to refuse the opportunity to substitute steady, dependable, and compliant machines for testy and inconstant human laborers. With less workers earning wages, demand for manufactured goods would necessarily fall, causing profits to collapse and precipitating a crisis of capitalism.

Marx also marveled at the innovations of his age and the possibility for human flourishing due in part to technological progress. He recognized that the former grew not only from human ingenuity but also from capitalism's ceaseless drive to revolutionize itself. His philosophy of history helped to reconcile his attitudes toward technology and society. Marx thought that technological progress represented the march of history toward a society of widespread prosperity, finally brought on by social and political revolution, i.e. the overthrow of capitalism. Technological progress, and with it the end of work, could thus be both immiserating, because of the imminent collapse of the political-economic system, and a mere signpost on the road to a more perfect society.

Like Marx, Keynes (1932), who coined the term “technological unemployment,” saw the possibility for widespread technical substitution. Keynes, however, took a notably more sanguine view of the “end” of work. He speculated that technical progress would make society sufficiently wealthy to free people from pressing economic cares, leaving only the problem of how to spend an abundance of leisure time. Recognizing that not all work would simply cease to be valuable or necessary to a well-functioning society, Keynes nonetheless wondered whether people would not settle on a minimal (say, 15-hour) work week, if only to keep themselves busy and useful. In fact, Mill (1848 [1965]) had already made a similar point. He took up the Classical idea of the “stationary state” to imagine a late stage capitalist society whose expansion would be bounded by diminishing returns to capital. While further technical progress would still be possible, he argued, it would serve mostly to reduce humans’ need to work, rather than increase society’s stock of wealth in itself.

Bold prognostications aside, historical evidence suggests that technological unemployment has never materialized (Mokyr et al., 2015). That does not mean that technical progress unambiguously raises welfare always and for everyone. Economists throughout history have also worried that machinery alienates and dehumanizes labor, can have deleterious health effects for workers, and increases income inequality. More to the point of this paper, the “long-run” on which many economists pin their hopes of proclaiming the net benefits of progress may not come

to fruition in a single lifetime; indeed, many estimates suggest that is exactly the case with the first industrial revolution.

The question therefore becomes whether this time is different: can history – and the failed predictions of economists past – be a guide to our future? Or are we reaching a watershed moment when technology drastically reduces the need for human workers across large sections of the economy?

One worrying sign that widespread technical substitution looms on the horizon is the rapidly growing scope of automation. Digitalization, by allowing for infinite and instantaneous copying often at no marginal cost or effort, may have already opened the door to an explosive rate of substitution (Brynjolfsson & McAfee, 2014). Recent exponential growth in the constitutive capabilities underlying robotics – computing power, data storage, and communications – as well as emerging technologies like the cloud and deep learning suggests the possibility of a “Cambrian explosion.” Just as vision is thought to have played a key role in sparking the burst of species diversification during the Cambrian era, the sudden leap in robots’ visual recognition capabilities may portend a rapid proliferation of robot applications (Pratt, 2015). Bullishness on automation extends into the business literature as well: with the technological requirements in place for extensive application and diffusion of artificial intelligence, all that remains is for entrepreneurs to discover and commercialize its many potential uses (Konishi, 2017; Lee, 2018).

Despite recent advancements in robotics and artificial intelligence, humans may yet encounter technical limits to automation. Polanyi’s Paradox offers one

possible reason why: the idea that “we know more than we can tell” suggests that humans inherently struggle to explain, let alone code, what we only know tacitly. Key traits like physical dexterity and social and emotional intelligence, of which our scientific understanding remains relatively limited, therefore prove difficult to impart to machines, which have relied on highly structured and predictable environments and explicit instructions to complete even the most basic tasks, like picking up an object (Autor, 2015). Workarounds like deep learning, by which robots self-learn generalizable knowledge and tasks through open-ended algorithms rather than deterministic code, may make the point moot: the very possibility of a Cambrian explosion speaks to the shifting ground beneath our assumptions about the human mind and the potential for machine intelligence.

The automation of complex knowledge work across industries provides a case in point. At the Memorial Sloan Kettering Cancer Centre, IBM Watson diagnoses diseases by instantaneously matching information about patient symptoms and genetics with more than 600,000 medical evidence reports and two million pages of text from medical journals. Google Translate provides real-time translation in more than 200 languages that improves in accuracy with use; newspapers like the Economist and the LA Times use sophisticated algorithms to craft summary pieces on sports and crime (Berger & Frey, 2016).

Such examples notwithstanding, one might wonder with Autor (2015): if the scope for automation is increasing, why are there still so many jobs? Sluggish productivity growth for much of the last three decades seems to provide *prima facie*

evidence that increasing substitution is a red herring. Low levels of corporate ICT investment since 2000 further suggests that companies are neither highly technologized nor eager to seize on the productivity gains supposedly promised by contemporary innovations (Mishel & Bivens, 2017). How can we account for the productivity slowdown in an age of rapid technological change? Gordon (2016) argues that innovation today isn't what it used to be, i.e., that the internet and smartphones simply do not have the same economic "oomph" of emergent technologies at the turn of the 20th century such as electricity and the automobile. The sheer volume of new, ever-more technologically complex products cannot make up for their inability to generate significant new value or vastly improve the economy's productive capabilities, relative to the breakthrough inventions of the past.

Even if we are not on the eve of a massive wave of automation, such skepticism fails to address several key points. The first is the fundamentally transformative nature of "breakthrough" inventions. For one, such innovations tend to reveal limitations, if not eventually prompt revisions, in productivity statistics. Gordon (2016) himself acknowledges that standard indicators fail to capture important elements of productivity and welfare, such as quality improvements; the rise of free digital services and the sharing economy also cast doubt on the adequacy of those indicators (Brynjolfsson & McAfee, 2014). Moreover, productivity statistics tend to lag actual increases in productivity, sometimes by as many as 10-15 years, opening up the possibility that machine learning and robotics are only now likely to

begin making a noticeable impact in the data. A broader issue concerns the relevance of our current thinking about technology in an era where the latter seems to change faster than the former. The emergence of new general purpose technologies can totally revolutionize the basic conditions of production and with them our assumptions about productivity, i.e. its limits and how they can be overcome (Bresnahan & Trajtenberg, 1995). Should artificial intelligence become the new electricity, as some technologists like to claim, it could be extremely difficult to imagine the future structure of the economy, let alone appropriately measure it in advance.

Human frailties, not least of which our persistent failure to accurately predict the trajectory and outcomes of exponential change, prevent us from achieving certainty about where trends in automation will lead. The debate about technological unemployment is no more likely to be resolved in the near future than it is to go away completely. The key task therefore becomes finding ways of using sound economic theory and reasoning to understand automation how has impacted employment in the past, so as to provide a foundation for assessing the changing relationship between man and machine. While the public debate about the prospects for technological unemployment continues to rage, scholarly work from the past three decades focuses largely on questions of relative labor demand. What types of workers are firms hiring, and why? How are firms really using technology? Through what mechanisms does automation affect employment among different groups? It is to these questions that I now turn.

2.2 Modeling Trends in Labor Demand and Employment

The impact of technological change on labor demand comes into clearer focus when one looks beyond the aggregate employment rate. Simply put, not all workers experience the same effects from automation. Nor is technological change monolithic in its effects on disaggregated groups of workers over time. Trends in relative skill requirements since the first industrial revolution starkly illustrate this fact, providing labor economists with the empirical basis necessary to begin modelling the specific relationship between automation, labor demand, and employment. Such models also help place labor demand in a broader context that includes not only other production inputs but also general equilibrium considerations crucial to establishing definitive claims about automation and employment.

Marx's (1868 [1977]) work again serves as the starting point for modern labor economists' inquiry into automation. He observed how the inventions of the first industrial revolution, namely the steam engine, enabled the automation of much of the work of skilled artisans. The birth of large-scale industry, with its production technology and division of labor that favored large numbers of unskilled workers, sealed the fate of handicraft manufacturers. Skilled labor was not entirely redundant – someone had to invent, improve, and repair the machines, after all – but its value dropped, accompanied by the first notable shift in relative labor demand due to technological change.

The contemporary idea of capital-biased technical change seizes on a kernel of truth in Marx's observation: that technological change inherently favors capital

over labor. The monotonic rise of the capital-output ratio over the last century and a half bears witness to this fact even as the rising aggregate employment rate during this time runs counter to the Marxian intuition that substitution tends to precipitate unemployment (Goldin & Katz, 1998). Could this pattern change? For instance, the falling price of computing power could induce a higher rate of substitution, a possibility the models described below attempt to address. The empirical record regarding inequality, another of Marx's dire predictions, is more mixed. From the end of World War II into the 1970s, income inequality in America moderated even as the country grew more technologically advanced (Gordon, 2016). Since the early 1980s, the labor share of income has fallen in the US and across OECD countries. Meanwhile, the wage share of the least educated workers has declined even as low-skill employment has grown (Berger & Frey, 2016).

Models of skill-biased technical change attempt to address some of the ambiguities that arise from a simple two-input framework. Around the turn of the 19th century, with the shift from steam to electric motive power and the development of new methods of production such as the Fordist assembly line, new technologies began to complement high-skill labor and substitute for many of the tasks previously relegated to low-skill labor. Engineers and electricians, for instance, became more valuable relative to manual factory laborers (Goldin & Katz, 1998). Eventually, this came to imply a "race" between education and technology, to quote the pioneering work of Tinbergen (1974). As the demand for skill *qua* level of education grows, wage inequality between skill levels rises and employment growth among less skilled

workers slows or stops altogether. Assuming continued innovation, only rising education levels can overcome the threat of stagnating labor demand and the relative impoverishment of a large segment of the labor market (Goldin & Katz, 2008).

Empirical evidence from the past two decades, however, suggests a more nuanced picture than that painted by models of skill-biased technical change. In particular, the phenomenon of job polarization calls into question the one-dimensional relationship between technology and skills. Since 2000, the share of total employment in middle-skill positions, including white collar and administrative professional jobs, has decreased substantially: the middle has “hollowed out.” Meanwhile, employment shares in both tails of the skill distribution have grown, with most of the growth concentrated in low-skill employments, most prominently service industry jobs (Autor, 2014). In essence, many skilled workers experienced substitution, while high-skill workers whose education should allow them to outpace technological change faced slowing growth in employment opportunities – two results directly at odds with the skill-biased hypothesis.

To better account for job polarization, researchers have turned to models of routine-biased technical change. In this approach, technology substitutes for human labor in routine, easily codifiable tasks, like computation and “pick and place” assembly line operations, but complements humans performing complex, non-routine tasks that remain difficult to code. Autor (2015) explains how “most work processes draw upon a multifaceted set of inputs...*each* play[ing] essential roles,” so that “improvements in one do not obviate the need for the other. If so, productivity

improvements in one set of tasks almost necessarily increase the economic value of the remaining tasks” (p. 6).

Acemoglu and Autor (2010) provide the theoretical and mathematical foundation for this model. Technology represents an endogenous variable that firms can use to substitute or complement for tasks, with jobs composed of a unique bundle of tasks. In turn, complete substitution *qua* human replacement is possible but not given, since jobs typically involve a variety of more and less codifiable tasks. Labor demand also depends on the elasticity of final product demand and the elasticity of labor supply. If demand for a firm’s good rises with productivity or falling prices, employment may grow even in the face of automation. Jobs with relatively elastic labor supplies, e.g., those with low education requirements, will also tend to absorb more employment than occupations with relatively inelastic supplies due to qualification barriers. As a rising labor supply keeps wages low, firms may find the cost of human labor more favorable than that of capital even where automation is technically feasible.

Figure 1 illustrates the recent trend of job polarization that models of routine-biased technical change help explain. Occupational categories are arrayed along the x-axis from least- to most-skilled. Three categories in the middle of the skill distribution – operators/laborers, production, and office/admin – have experienced declining employment since 1999. Even in the preceding two decades, employment grew more slowly in these occupations than in the two highest and lowest skilled occupations. According to Acemoglu and Autor’s model, the emergence of ICT and

explosion of computing power during this time enabled technical substitution in routine tasks that make up a significant portion of middle-skill jobs, from clerks to production line laborers. Since 1999, employment growth has been strongest among the three least skilled occupations, which include personal care and food/cleaning service, followed by highly skilled professions such as technicians. These jobs tend to involve non-routine tasks, including close personal interaction, situational awareness, and flexibility in adapting to emergent challenges, not easily substituted by computers.

Figure 1. Change in Employment by Major Occupational Category, 1979-2012
(the y-axis plots 100 times log changes in employment, which is nearly equivalent to percentage points for small changes)




Figure and caption adapted from Autor (2015). *Sources:* Author using data from the 1980, 1990, and 2000 Census IPUMS files, American Community Survey combined file 2006–2008, and American Community Survey 2012. The sample includes the working-age (16–64) civilian noninstitutionalized population. Employment is measured as full-time equivalent workers. *Notes:* Figure 2 plots percentage point changes in employment (more precisely, the figure plots 100 times log changes in employment, which is close to equivalent to percentage points for small changes) by decade for the years 1979–2012 for ten major occupational groups encompassing all of US nonagricultural employment. Agricultural occupations comprise no more than 2.2 percent of employment in this time interval, so this omission has a negligible effect.

Understanding the demand for labor and connecting it to aggregate employment also requires taking into account the elasticity of demand for goods and services. If the demand for a firm's output is highly price elastic, productivity improvements may lead to an increase in its demand for labor even as the amount of labor required to produce one unit of output falls. Productivity gains may also increase incomes, boosting sales of income elastic goods and services. Demand also responds to changes in product quality, customization, and speed of delivery brought about by automation. Moreover, demand elasticities change over time, as consumer preferences respond to the interplay of prices and income.

Bessen (2018) demonstrates how these dynamics help explain long-run employment trends in steel, textiles, and automotive manufacturing, which exhibit an inverted "U" shape. Despite rapid productivity growth in these industries at the turn of the 19th century, each experienced substantial employment growth as the price of output fell. In turn, incomes rose and consumers enjoyed new uses for increasingly cheap goods; cotton cloth, for instance, no longer functioned as a luxury good and thus became the fabric of choice in a growing variety of final goods. Eventually, however, demand for steel, textile, and automotive products started to become more inelastic. At a certain point unique to each industry, demand no longer grew sufficiently to offset productivity gains, as diversification in use and the fall in price or rise in income failed to keep pace with the rate of technical substitution.

The literature on relative labor demand opens up but fails to address two important questions. First, what is happening, or what should we expect to happen,

to *absolute* labor demand? In this regard, routine-biased technical change models do not improve much upon their canonical foundation, which explain unemployment only in terms of a mismatch between skills and the demand for labor (Vivarelli & Pianta, 2000). The former model, though superior to the latter in many respects, does not indicate whether technical substitution for routine tasks may have different implications for aggregate employment than did technical substitution for low-skilled labor. Second, and more to the point of the current study, models of routine-biased technical change fail to explain under what specific conditions a technology substitutes or complements for labor (Bresnahan et al., 2002). More precisely, when does substitution for (routine) tasks become substitution for jobs? What's going on "under the hood" when firms demand more or less labor even when a new technology does not necessarily replace every task performed by a particular human being? Specifying the unique, firm-level causal mechanisms involved in technology adoption and the demand for labor may illuminate factors that mitigate or attenuate the substitutive (complementary) properties of technologies themselves.

2.3 Technology Adoption and Complementarities at the Firm Level

Under what specific conditions does technology substitute for or complement labor? Evidence suggests that effective technology adoption at the firm level is not a simple process of "plug and play," but often involves the implementation of complementary organizational changes that enable firms to adapt to new capabilities and competitive environments. Ultimately, the most successful adopters demonstrate

that new technologies change the “fundamental nature of what a firm does and how it does it” (Bartel et al., 2007, p.32). This suggests that how firms adopt technologies – that is, the ways that they adapt work practices, policies and procedures, and organization to new technological capabilities – as well as the extent to which those technologies induce changes in market strategy mediates the effect of technological change on labor demand.

The intra-firm dynamics of technology adoption thus reveal why and to what extent ICT calls for a particular labor mix. “IT is embedded in a cluster of related innovations, notably organizational changes and product innovation,” the three of which together constitute the (skill and/or routine) bias of technical change (Bresnahan et al., 2002, p. 341). The object of analysis therefore shifts from the physical technology itself to the complementarities between the three innovations – ICT, workplace reorganization, and new products and services – underlying adoption. Approaching ICT in this way reflects the fact it is not a traditional capital investment, but a “general purpose technology” whose economic benefits consist largely in facilitating complementary innovations (Bresnahan & Trajtenberg, 1995). In this context, complements represent not only a relation among pairs of inputs, but groups or systems of activities that together generate cumulative effects: raising the level of any one such activity increases (decreases) the marginal return (cost) to any or all of the other activities (Milgrom & Roberts, 1990).

Successfully adopting ICT thus tends to require firms to make several substantial, closely coordinated changes to operations and market strategy. In their

pioneering study of manufacturing firms using Computer-Assisted Design technology and flexible manufacturing processes, Milgrom and Roberts (1990) find pervasive complementarities across functions like marketing, engineering, design, and production. Technologically advanced firms exhibit, for instance, an emphasis on product quality and continuous improvement, integration of product and process engineering, the use of mass data communication and production technologies with low setup times, and short production cycles. As firms adopt some of these characteristics, it becomes profitable to adopt even more of them.

A diverse literature provides further evidence for the important effects of changes in work practices and organization on firm performance. In an empirical study of steel finishing lines across 17 companies, Ichniowski, Shaw, and Prennushi (1997) find that firms implementing clusters of innovative practices – including incentive pay, production teams, flexible job assignments, and training – achieve substantially higher levels of productivity than firms that make changes to individual work practices. The latter see little to no improvement in productivity, underscoring the importance of complementarities among various organizational changes. Black and Lynch (2001) use panel data from 600 manufacturing plants from 1987-1993 to examine the effect of adopting Total Quality Management systems. They conclude that adopting such a system significantly affects productivity only in firms that also use innovative human capital practices, such as profit-sharing programs and employee participation in decision-making, or increase computer use among production workers. A more recent study by Bartels et al. (2007) estimates a

longitudinal model for a narrowly defined industry – valve manufacturing – to show that “the adoption of IT-enhanced machinery involves much more than just the installation of new equipment of the factory floor” (p. 2). Three other mechanisms for increased productivity emerge: (1) a shift in business strategy from long production runs to smaller batch production; (2) improved efficiency at every stage of the production process due to reduced setup, run, and inspection times; and (3) an increase in labor skill requirements and the adoption of new human resource practices.

Applications of the complementarities approach to firms and industries in the service sector has been more limited. One of the earliest studies on the topic examined the adoption of identical computerized tomography scanners in two hospitals in the same metro area (Barley, 1986). The study provides evidence that the new technology had a significant impact on work organization by disrupting the relationship between radiologists and technicians, leading to new roles and forms of interaction between the two occupations. More recently, Autor, Levy, and Murnane (2001) studied the introduction of check imaging and optical character recognition devices in two departments of the same bank branch. In one department, the technology led to computer substitution for high-school educated workers. In the other, where labor skill intensity was already higher pre-adoption, it fostered the integration of tasks with “fewer people doing more work in more interesting tasks” (p. 442). Even within the same firm, therefore, the same technology may exhibit

different effects on workplace reorganization, depending on factors not necessarily related to technology itself, such as human capital.

Unfortunately, few studies of the complementarities between ICT and firm-level organizational changes have addressed their impact on labor demand. Some attempt has been made, however, to incorporate research on complementarities into the canonical skill-biased model of technical change. For instance, two factors may drive up the relative skill intensity of jobs in organizations that adopt both ICT and complementary organizational changes: limited substitution and information overload (Bresnahan et al., 2002). First, the scope of complete worker substitution by, e.g., computers tends to be limited, the more so in jobs consisting of more complex and cognitively demanding work. Thus, firms may restructure work to separate out routine processing tasks from those requiring human skills, which may in turn raise the demand for non-cognitive abilities (“people skills”). Second, increasingly computerized processes produce greater amounts of data, requiring additional skilled labor to perform analytic reasoning and abstract decision-making. But because data volume tends to increase faster than firms can adapt their labor pool, they must also make organizational changes that allow them to better distribute information-processing tasks.

The quickly growing scope for the application of artificial intelligence may soon challenge what we know about ICT adoption, organizational change, and labor demand. Advances in machine learning already make possible the “routinization” of what previously appeared non-routine, from learning (and conquering) Atari video

games to translating foreign languages to recognizing human faces. The expanding field of analytics, propelled by the emergence of big data sets, encroaches on decision-making tasks once considered the sole province of the knowledge worker. For these reasons, adopting machine learning may require even more significant organizational adaptations than computers, reshaping businesses in ways yet to be envisioned, much less studied, by economists (Ransbotham, Kiron, Gerbert, & Reeves, 2017). To be sure, much more firm-level evidence on the adoption of machine learning is necessary in order to make sound claims about the likely impacts of large-scale diffusion and firm transformation (Seamans, 2018). Nonetheless, the nature of new ICTs seems poised to overturn ways of thinking about and dealing with the problems of limited substitution and information overload, and thus produce new types of organizational adaptations with as yet unpredictable consequences for labor demand.

Indeed, previous innovations have posed similar problems of adjustment for business leaders and economists. Reflecting on Robert Solow's famous quip that "Computers show up everywhere but in the productivity statistics," David (1990) compared the introduction of computers to that of the electric dynamo at the turn of the 20th century. He showed that both the physical and social organization of manufacturing plants influenced whether and where firms adopted new motive powers: because the technology required a thorough restructuring of existing plant and work structures, diffusion proceeded slowly and was led almost entirely by new firms that did not need to overcome the challenge or cost of substantial *re-*

organization. Later case studies of computer adoption illustrated how firms might struggle to implement the full slate of complementary changes necessary to improve productivity, and how only partial changes could actually hurt productivity and even put firms out of business (Brynjolfsson & Hitt, 2000).

What do the microeconomics of complementarities mean for our understanding of long-run structural change in the economy? Evolutionary economic theory provides a helpful framework for connecting what happens at the firm level with the broader patterns of change at the industry and global level. What follows draws heavily on Nelson and Winter (1982) and Nelson (2009), though in many ways the seeds of evolutionary approaches had already been sown in Marx's account of large-scale industry, described above. From an evolutionary perspective, firms are behavioral entities made up of building blocks identifiable as organizational routines. This stands in contrast to mainstream, neoclassical definitions of the firm, which tend to treat it as a kind of individual agent made up of a set of additive inputs and outputs. The evolutionary approach thus implies a kind of unique indivisibility pertaining to firms: each one composed of different routines, from production processes to technology to human capital and market strategies, the absence of any one potentially precipitating a major transformation in the firm's output, productivity, and so on.

This also implies that firms must do more than just make marginal decisions about inputs and outputs. Indeed, empirical evidence from the complementarities literature demonstrates that the decision to adopt ICT is marked by important non-

convexities associated with the need to make a set of coordinated choices about technology, workplace organization, product innovation, and supplier and customer relations (Brynjolfsson & Hitt, 2000; Milgrom & Roberts, 1990). In turn, technology adoption is not a simple matter of optimization given a new set of factors. Firms deciding whether to adopt a new technology often face significant costs (including sunk costs), challenges related to identifying and executing appropriate changes in organization, and uncertainty about the outcome of adoption. Investment in new technologies may thus proceed slowly; even when it happens quickly expected productivity gains may not materialize for some time. Across firms, therefore, diffusion tends to occur unevenly. The complementarities literature provides evidence for this: adjustment difficulties and the need for experimentation and coinvention surrounding ICT use at the firm level lead to substantial variation in use, organizational complements to ICT, and performance outcomes (Bresnahan et al., 2002). Complementarities and the indivisibilities they give rise to also account for why firms find it so hard to simply imitate the success of early adopters (Milgrom & Roberts, 1995).

Though the answer is far from automatic, the question of whether to adopt new technology stems naturally from firms' impetus to seek competitive advantage. Schumpeter (1943 [2010]) theorized that innovation often leads to technological, rather than price, competition: firms seek to gain a technical edge on their competitors by developing new products and processes, and by monopolizing the gains from them as long as possible. Additionally, firms are embedded in a particular

set of historical processes and institutional environments, e.g., market structure, government policy, and societal attitudes toward innovation and entrepreneurship, that actively shape the technical and economic conditions in which they make decisions. Successful early adopters encourage imitators and complementary or spinoff innovations by others seeking to wrest a technological advantage from first-movers, in turn generating new technological paradigms and imperatives. Technical change and its consequences, in this view, are not inevitable byproducts simply of the physical properties of innovations, but are also contingent upon social, economic, and political forces that shape adoption and diffusion.

3. Methodology

3.1 Theoretical Framework

This study poses the following research question: How does the adoption of technology and implementation of complementary work practices affect firms' demand for labor? To answer this question, I begin by adopting a theoretical framework grounded in two of the most salient developments from the scholarly literature reviewed above: Acemoglu and Autor's (2010) model of routine-biased technical change, and the insights into the role firm-level complementarities play in determining the effect of technology adoption on firm organization, performance, and labor demand. This framework guides my identification of variables of interest and provides a structure for collecting and analyzing data in the form of three case studies. It also motivates a set of predictions about what kind of evidence case study research will reveal.

Acemoglu and Autor's (2010) characterization of routine-biased technical change motivates the reasoning behind this study: that firms use technology to a) substitute for human labor in tasks which are routine and easily codifiable, and b) complement human labor in tasks which are non-routine and non-codifiable. The work processes a firm utilizes to generate output draw on a complex set of inputs that include both routine and non-routine tasks, with individual jobs typically composed of a unique bundle of both types of tasks. Complete substitution for human labor is

possible but not given: tasks and the ways in which they are (or can be) bundled together into jobs mediates the effect of technology on labor demand.

In turn, identifying firm-level complementarities enhances the model by specifying the conditions under which technology may substitute or complement for any number of tasks. Both complementary investments – for instance, in hardware and training – as well as changes in work practices and the organization of production further mediate the impact of technology on labor demand. Implicitly, the benefits of new technology to a firm, whether in the form of greater efficiency, enhanced product quality, or the development of product innovations, consist not just in the application of a better (faster, more precise, easier to use, etc.) technical apparatus itself but in the ways in which firms adapt their inputs and production processes to leverage new capabilities. Put another way, focusing on the complementarities between three areas of innovation – ICT, workplace reorganization, and new products and services – enables researchers to analyze more closely the causal mechanisms at work in changes to labor demand.

This framework underscores the importance of several variables for analyzing how technology and labor demand interact within a given firm. First, a firm's rationale and strategy for adopting new technology – both why and how it chooses to implement the technology *qua* physical input – establishes the ground for hypothesizing the impact on labor demand. Is adoption intended to reduce costs and generate efficiencies? To improve product quality or consistency? And does the firm approach adoption as “plug and play,” or by seeking to tailor its organization,

production process, and/or market strategy to its emergent capabilities? Second, the primary use of the technology illuminates the particular tasks for which it substitutes and complements, and which aspects of work organization and practice become the focus of complementary changes. In turn, new work practices and complementary investments constitute the third and fourth variables of interests. Fifth, assessing how a firm's demand for labor has changed, both in absolute terms and with respect to the particular labor-skill mix it employs, allows me to assess the theoretical validity of my framework with respect to the particular firms studied as well as compare the consequences of technology adoption across firms.

In addition to establishing the key dimensions of technology adoption I seek to study, my theoretical framework sets up some general expectations about the "story" of adoption and adaptation across firms, and thus what case studies will reveal. At the most basic level, firms will use ICT to substitute for routine tasks and complement non-routine tasks. The effect of doing so on absolute labor demand will be ambiguous; the model is generally agnostic about whether more or less workers are required overall, a result that depends on many features exogenous to the framework. The effect on relative labor demand is more clear: ICT adoption will lead to an increase in demand for labor performing non-routine tasks relative to labor performing substitutable routine tasks, irrespective of worker skill (education) level but mediated by the type and extent of complementarities involved in adoption.

This last prediction gets at the novelty of this study and thus requires some unpacking. It begins with the simple premise that firms will tend to make changes

both simultaneous and subsequent to adoption that complement their new ICT. Specifically, they will implement new work practices and/or redesign workflows, reorganize job responsibilities (tasks), and make related investments, particularly in categories like hardware and employee training, in order to take full advantage of their new technology. Firm choice and strategy also play a critical role in determining the cumulative effect of technology adoption and complementary changes that accompany it. Why firms adopt ICT in the first place matters. Firms aiming to cut costs by generating efficiencies, for instance, likely incorporate technology into their service delivery operations in such a way as to reduce labor demand, while firms adopting ICT in order to improve service quality may not always use technology as a substitute, even where it is technically feasible to do so. Among firms that directly incorporate their use of ICT into their market strategy, the effect of adoption and complementarities should be magnified. Such firms may be driven to discover new complementarities that enhance the effect of technology on performance, productivity, and labor demand. Further, the presence of complementarities, particularly in work practices and organization, may grow the scope of which tasks can be routinized in the first place, thereby extending the application of technology to previously human-performed tasks and jobs.

3.2 Data Collection and Analysis

I collected data by conducting semi-structured interviews with representatives of three firms. Each interview centered on a single technology (e.g., a software

platform) or type of technology (e.g., machine learning). Interview items sought to elicit data on the five variables of interest introduced above:

1. The firm's rationale and strategy for adoption
2. The primary uses to which it puts the identified technology
3. New workflows, organization, and work practices intended to complement the technology
4. Complementary investments made during or subsequent to adoption
5. Changes in labor demand subsequent to adoption.

Firm representatives consisted of director-level and C-suite personnel with direct responsibility for the development, adoption, and ongoing management of the identified technology. Each interview provides the basis for a qualitative firm-level case study, supplemented as appropriate and necessary by additional data retrieved outside of the interviews (e.g., via the web or further personal correspondence).

I conduct both within- and cross-case analyses of the case studies. The former enables me to ascertain the suitability of my theoretical framework for understanding the relationship between technological change, complementary non-technical changes, and labor demand. This is particularly important given that I apply the framework to cases different from those studied elsewhere in the literature, i.e., service firms utilizing recent ICTs, such as machine learning. Insofar as gaps emerge between theory (prediction) and data (finding), I ask whether they are the result of flaws in the framework or its application to cases it is not designed to address. The cross-case analysis illuminates similarities, differences, and potentially emergent

patterns with respect to the variables of interest among the three firms. Though the external validity of the case study findings is limited, the cross-case analysis allows me to tease out possible implications of this study for the current state of knowledge about how technical change and labor demand interact at the firm level, and points to fruitful directions for future research.

3.3 Case Selection

Case selection involved three key criteria. The first was the intrinsic importance of the cases. Intrinsic importance can be broken down along two dimensions: economic and scientific. I define economic importance as the size of an industry, both in terms of its value (as a percentage of national GDP) and aggregate employment. As detailed in the introduction and in Section 4, health care and real estate play outsized roles in the US economy, together contributing more than a fifth of the country's GDP and representing an engine of American employment. Thus, improving our understanding of how technological change affects labor demand in these two industries goes a considerable way toward providing an overall picture of the trajectory of machine substitution and structural employment trends. Surprisingly, the extant literature on technical change and employment pays scant attention to these and other service industries, with just a few exceptions (e.g., Barley (1986) and Autor et al. (2003)): filling this crucial gap is therefore a task of scientific importance in the field of economics. Likewise, this study brings recent theory to

bear on cases of recent ICT innovations, an area where economists are only beginning to understand the labor market consequences of technological change.

The second criteria for case selection addressed the need to minimize the confounding influence of trade and globalization on firms' demand for labor. By selecting firms in nontradeable industries, in addition to focusing on labor demand rather than employment *per se*, I am able to mitigate the role trade plays in determining the microeconomics of the firm. Mutual causality between globalization and the ICT revolution, as well as general equilibrium effects, prevent me from completely eliminating the effect of trade on firm-level outcomes. While accounting for this caveat, however, I can use current economic research to develop a satisfying alternative.

Spence and Hlatishwayo (2012) construct an index of tradability by classifying industries based on the tradable proportion of the value chain of which the goods and services they produce are a part, using value-added as their measure. The tradability of any given slice of the value chain depends on the extent to which production at that point is geographically concentrated. Geographic concentration here serves as a useful if imperfect proxy for the extent to which a good or service can be traded. The assumption is that the extent of the market for a tradable good or service does not depend on the establishment of new facilities for production, since the good can simply be imported. Greater geographic concentration of production therefore implies greater tradability, while the market in nontradable goods grows only as new sites of production emerge. On a 100-point scale, with 100 representing

fully nontradable, Spence and Hlatishwayo (2012) classify both health care (97.8) and real estate (100) as almost or entirely nontradable.

The third criteria for case selection involved the likelihood of generating appropriate evidence given my object of inquiry – the relationship between technology adoption, complementarities, and labor demand – and theoretical framework. Given some preliminary information about a firm, such as industry, use of technology, and the availability of information on the organization and practices of work, I judged the relevance and potential fit of the case for my study. This assures that within-case evidence supports the aim of my research, even if it illuminates problems with my theoretical framework. I also selected firms based on the need to support cross-case analysis. I therefore attempted to select firms using comparable technologies in order to bring to the fore relevant similarities and differences in how and why firms adopted those technologies, and how the choices and complementary changes each firm made influenced labor demand.

4. Case Studies

4.1 Background

The case studies below present firm-level evidence in five domains relevant to my theoretical framework:

1. Rationale and strategy for adoption: why and how did the firm adopt new technology?
2. Primary uses of the technology: what tasks does it perform?
3. New workflows, organization, and work practices related to adoption: how have service delivery operations changed to incorporate the new technology?
4. Complementary investments: did the incorporation of new technology into the firm induce other related investments, such as hardware and training?
5. Changes in labor demand: how has the firm's demands for labor, both skilled and unskilled, changed since adoption?

Table 1. summarizes key features of the three firms. The case studies capture three 'types' of firms. REX Real Estate Exchange is a three-year old startup with fewer than 100 employees operating in 10 metropolitan markets nationwide. Sky Ridge Medical Center (SRMC), with 1,300 employees, is one of 8 hospitals in HealthOne, a for-profit regional healthcare system serving the Denver Metropolitan

Area. UCHealth is a not-for-profit healthcare system, the largest in the State of Colorado, serving the much of the state and portions of neighboring Wyoming and Nebraska. UCHealth employs roughly 19,000 people across 11 hospitals on Colorado’s Front Range.

Table 1. Summary of Case Studies

	REX	SRMC	UCHealth
Industry	Real Estate	Health Care	Health Care
Firm Type	Start up	Single hospital within healthcare system	Healthcare system
Technology	Big data; machine learning	EHR	EHR
Market Size	National	Metropolitan Area	State/Intrastate Region
Employees¹	66	1,300	19,000

The variety in the structural position of the three firms – startup, individual medical center within a healthcare network, corporate healthcare system managing multiple hospitals – poses a challenge and an opportunity for research. On the one hand, cross-case comparisons must be treated with caution, since differences in firm structure or level inevitably affect technological and organizational changes. That case study data arises from the point-in-time perspective, or organizational ‘vantage point,’ of a subjective representative of the firm heightens the positionality of each case. One can draw few if any general conclusions from the cases, even where they suggest distinct patterns or conform to theoretically sound hypotheses.

On the other hand, different firm types allow me to tease out important details about the firm-level complementarities that affect the relationship between

¹ Number of employees reported via interview or other data provided by firm as of November 2018.

technology and labor demand in select service occupations. What I lose in generalizability I gain in detailed exploration of causal mechanisms behind the use of technology to substitute or complement human labor. For instance, the variation provides a useful backdrop against which to understand how each firm makes decisions about innovation and investment, given a particular corporate structure or position. Taking into consideration the latter allows me to study firm-level complementarities in a way that the literature has not yet addressed, and which may have ramifications for workers in a variety of occupations across different types of firms. The difference in position between the two healthcare firms in particular helps illuminate how distinct levels of large corporate entities in the industry view and use technology, both operationally and strategically.

The case studies focus on the firms' adoption and use of information and communication technology (ICT). Major recent developments in ICT have received relatively little attention in the academic literature on technological change at the firm- and industry-level (Seamans, 2018). The spread of big data and advances in machine learning have led to an explosion of new applications of machine learning, trends that emerge in the case studies on REX and UCHealth (Berger & Frey, 2016). SRMC's technology, a DOS-based electronic health records (EHR) system adopted in 2003, reflects much more closely the basic properties of the computer and software systems examined in earlier studies of ICT in service occupations (Autor, Levy, & Murnane, 2001; Bresnahan et al., 2002). In some ways, SRMC can provide a useful

reference point for better understanding whether and how more advanced ICTs change the relationship between technology and human labor.

REX utilizes machine learning to drive customer acquisition and advertising. Algorithms mine vast amounts of microdata on unique features of individual homes, nearby real estate transactions and market conditions, local amenities, and consumer behavior in order to develop marketing strategies targeted at individual consumers.

This enables REX to:

- predict which homes will come up for sale before they are listed, enabling its agents to identify prospective clients before other brokers;
- develop individually-targeted marketing strategies that match highly specific preferences revealed, e.g., through personal consumption and web-browsing patterns; and
- use artificial intelligence to write and continuously modify ad copy as users generate more data about their revealed preferences.

An example comparing REX to a traditional real estate firm illustrates the powerful properties of the former's ICT. A traditional brokerage can use data from a prospective buyer's web browser to observe that they tend to view homes with, say, 3 bedrooms and 3 baths near downtown, allowing them to generate more listings for that consumer with these same features. REX's application of machine learning to granular data reveals that, in fact, the buyer lingers much longer on the page of homes with bay windows and large backyards, and that the desire for downtown proximity is a spurious result of the buyer's preference for living near the city's

highest-performing schools, which happen to be downtown. Moreover, REX can triangulate a single user's data from multiple browsers to create individual-tailored listings across devices and social media platforms.

Case studies of the two hospitals focus on their use of EHR systems. The Healthcare Information and Management Systems Society (HIMSS) provides the following definition of an EHR:

The Electronic Health Record (EHR) is a longitudinal electronic record of patient health information generated by one or more encounters in any care delivery setting. Included in this information are patient demographics, progress notes, problems, medications, vital signs, past medical history, immunizations, laboratory data and radiology reports. The EHR automates and streamlines the clinician's workflow. The EHR has the ability to generate a complete record of a clinical patient encounter - as well as supporting other care-related activities directly or indirectly via interface - including evidence-based decision support, quality management, and outcomes reporting.

Many of these features of EHRs appear in the case studies below. All EHR systems support three basic functions: documenting care providers' actions, issuing orders (e.g., for procedures or medication), and billing for services.

SRMC primarily utilizes its EHR system, Meditech, to perform the first two functions, having made relatively less progress in using the EHR to automate billing. UCHHealth's system, EPIC, incorporates all three EHR functions but goes well beyond these by integrating them with 25 other non-EHR systems into an enterprise-wide platform encompassing human resource management, enterprise planning and analytics, and digital patient portals, among other things. The difference in how the two healthcare firms incorporate EHRs into their operations and strategy suggest

variation in the extensive and intensive use of technology, something I explore in detail in the discussion of findings.

A brief history of EHRs helps to contextualize the two case studies in healthcare. The following draws on Atherton (2011) and Tripathi (2012). Like many ICTs, the development of EHRs corresponds with four major technological changes in the past half-century. First, the development of mainframe computers opened the possibility of digitizing many types of information, including medical records, though only for organizations capable of managing complex IT infrastructures. Second, as computers shrank and the person computer emerged, EHR software became more accessible and affordable for much smaller organizations, such as ambulatory clinics. Third, the internet gradually enabled secure communication and data sharing with patients and other providers; more recently, cloud-based servers have furthered the development of “lightweight” EHRs. Finally, the proliferation of microprocessors across a variety of devices has improved computing and networking capabilities, allowing EHRs to incorporate new types of information for the purposes of improving care and efficiency. Over time, as processing power exploded and digital devices shrank, the customer base for EHRs expanded.

The first EHRs emerged in the 1960s and 1970s as academic and government medical centers sought to improve patient care by developing more consistent and convenient record-keeping methods. Diffusion proceeded slowly until the 1990s, however. In a landmark 1991 report, the Institute of Medicine called for the nationwide implementation of EHRs and identified three major barriers to EHR

adoption: a lack of consistent standards on what EHRs do and how to use them, data privacy and security issues, and the cost of acquiring and implementing EHRs. In particular, the growing heterogeneity of EHR users, each with different informational needs, limited the development and spread of scalable EHRs capable of exploiting the capabilities that computerization offered. Tripathi (2012, p. 27) states the issue succinctly:

As the healthcare industry continued to practice medicine like guilds of independent craftsmen and artisans, they insisted that their tools be custom-crafted as well, which made it impossible for the industry to reap the benefits of economies of scale and scope that have driven high penetration of information technology in other parts of the economy.

EHRs therefore functioned like electronic filing cabinets, mimicking the non-standard record-keeping approaches of narrative- and dictation-based paper systems. EHR developers tried, mostly unsuccessfully, to force-fit computers to paper-oriented workflows, and non-standard contents and formats prevented data sharing between care providers. In turn, healthcare systems systematically underinvested in ICT throughout the 1990s and 2000s, and adoption continued to proceed slowly and unevenly.

The Health Information Technology for Economic and Clinical Health Act (HITECH), a part of the American Recovery and Reinvestment Act of 2009, turned the tide on EHR adoption. It required all public and private healthcare providers to adopt and demonstrate “meaningful use”² of EHRs by January 1, 2014 and

² In HITECH, “meaningful use” is defined as using EHRs for a “meaningful,” i.e., relevant, purpose (such as electronic prescribing); ensuring inter-operability between systems; and submitting quality of care and other measures to the US Department of Health and Human Services.

established a carrot and stick approach to meet this deadline. On the one hand, HITECH provided incentives totaling \$27 billion to encourage care providers to adopt EHRs meeting certain standards. On the other hand, it threatened to penalize non-compliance by reducing Medicare and Medicaid payments. At the same time, HITECH established federal certification for EHR products to ensure they would enable healthcare providers to meet meaningful use requirements, including uniform clinical content across vendor systems and care settings; consistent and robust measurement capabilities; data mining capabilities; public health reporting; and interoperability with other systems.

HITECH accelerated and transformed the use of EHRs. In 2008, just 9.4% of non-federal acute care hospitals had adopted an EHR; by 2015, 83.8% had at least a “basic” EHR and 96% possessed a certified EHR (Henry, Pylypchuk, Searcy, & Patel, 2016). The combination of demand-side incentives and supply-side certification standards facilitated greater commonality in basic EHR functions at a much faster pace than would have occurred in the absence of HITECH. Improvements in dimensions like data sharing and interoperability, combined with increased physician interaction with EHRs, invigorated interest in designing computer-based workflows. Regulatory standards have also helped EHR developers achieve economies of scale and scope that previously allowed large firms to draw value from IT in other services, such as banking, retail, and food services (Tripathi, 2012). With most hospitals having adopted some kind of EHR, a new divide is emerging between hospitals who merely meet standards and “advanced” users, such

as those who use EHR data for performance management or to improve patient engagement (Adler-Milstein et al., 2017).

4.2 REX Real Estate Exchange

REX was founded by a former partner at Goldman Sachs, who sought to capture in real estate the efficiencies generated by computerized trading in the stock market. At its most basic, the firm's value proposition involves using ICT to reduce arbitrage. First, this reduces transaction costs: REX's across-the-board fee of 2% is roughly 80% lower than that of traditional brokerage firms, who charge an average of 5-6% per sale. This can translate into significant savings for both sellers and buyers. Second, REX aims to improve the speed and volume of home sales through targeted, direct-to-consumer advertising that circumvents the Multiple Listing Service (MLS) used by nearly all agents and brokerage firms nationwide.

Strictly speaking, REX did not *adopt* the technology under consideration. After all, its proprietary algorithms are central to its value proposition and thus its attempt to gain a competitive advantage in the real estate market. The company's technology was present at the creation, if you will; better yet, it *is* the "creation," not a later addition meant to upgrade the firm's technological capabilities. Yet REX stands out not only because of the unique properties of its technology, but also – indeed, primarily – how the firm organizes itself around those capabilities. This case therefore underscores the powerful mediating role of organizational complementarities with respect to innovation's effects on a firm's demand for labor.

As I discuss further in the next section, this interaction comes into sharp relief when considered in the context of how most real estate agents and brokerage firms use technology.

REX's algorithms serve as both a substitute and complement for tasks performed by real estate agents. REX agents perform most of the tasks of traditional agents, from recruiting clients to photographing and showing the home to facilitating the closing process. In customer acquisition, however, technology plays a key role. Algorithms rather than agents discover buyers and sellers and thus identify the company's prospective clients, substituting for a task that previously required a significant amount of time and nuanced social interaction. This same function of REX's technology also complements agents' pitches to prospective clients, increasing efficiency and possibly efficacy. First, since agents no longer need to be physically present in a market or have detailed local knowledge of it to target their recruitment efforts, they can pitch services remotely. Second, REX's speed and power of computation allow each of its agents to identify and pitch to roughly 100 clients per month, far exceeding what human agents alone can accomplish.³ With the opportunity to speak to many more potential clients and a wealth of highly specific information about local sales and client preferences, REX agent pitches may also enjoy a higher success volume (if not rate) than non-REX agents, though such a number can only be determined anecdotally.

³ Anecdotally, a "good," i.e., highly motivated and savvy, real estate agent pitches services to about 8-10 potential clients per month.

Machine learning also drives most of REX's marketing efforts. Computers write and update home listings in real time, substituting for copywriting. Direct to buyer advertising may also substitute for the work of buyers' agents, whose value-added derives from their knowledge of and ability to work the MLS. For REX agents, direct to buyer advertising therefore serves as a complement insofar as it improves their ability to reach consumers. Combined with REX's capabilities related to customer acquisition, its advertising technology helps ease market expansion. A staff of two can open a new metropolitan market for REX more or less by "turning on" its algorithms, which immediately begin to crunch heaps of data on local housing sales and target potential clients with advertisements.

In two key ways, REX draws a close analogy with Uber. First, both companies defy long-established wisdom about how to succeed in highly competitive, geographically bounded markets. Just as the value of a "good" cab driver results from their wealth of local knowledge and experience (their ability to navigate complex road systems and avoid traffic jams), so too a "good" real estate agent possessed large personal networks and a keen awareness of local conditions and past sales. Now, Uber leverages advanced GPS technology to substitute for driver familiarity, while REX uses big data and machine learning to circumvent the profession's guild-like restrictions and identify the unique patterns and preferences of consumers in each metropolitan area. Yet while Uber arguably complements the relatively low-skill labor of thousands of drivers, removing barriers to entry (e.g., professional licensing requirements), the impact of REX's technology is less clear.

For the handful of agents it employs as new markets begin to flourish, technology greatly reduces the hurdles they face to both startup and expansion. Though REX's agents must be licensed, they do not have to operate through the MLS, which restricts their access to customers. On the whole, however, the substitution effect almost surely dominates (as REX intends it to), as both the economic and symbolic value of "knowing" a market diminishes significantly in the face of more efficient and equally efficacious technology.

Second, REX has "uberized" its employees' work by unbundling and crowdsourcing tasks associated with the job of real estate agent. Rather than match one agent to one client, REX divides up the work related to each client into discrete tasks, such as hosting an open house or delivering a pitch to a prospective client. REX agents use an app to view, select, and receive payment for each task they perform, instead of as a lump sum commission at the conclusion of a sale. What this means for individual agents in terms of total compensation and hours of work remains unclear – as it does in the case of Uber's drivers, too.

The "uberization" of the real estate agent's occupation underscores the technology-driven character of workflows and practices at REX. REX's marketing staff and real estate agents remain key inputs, but their work centers around, and relies upon, the outputs of the company's proprietary algorithms. The computer identifies which customers to target, what price to set, and which features to emphasize in listings. The scope for human manipulation of this data resides mostly on the front end with computer programmers and data scientists; agents and

marketers, on the other hand, use the data to tweak digital campaigns, interact effectively with sellers and buyers, and close deals. The latter's tasks are increasingly bound and determined (and, for agents, even delegated) by technology. Discretion and social and emotional intelligence remain valuable capabilities for REX's staff, and humans still do most of the work that require these capabilities. Overall, though, work is organized to leverage the firm's data-crunching powers, platforms for advertising, and just-in-time contracting.

Beyond operations, technology forms the heart of REX's long-run strategy. In real estate, the use of machine learning both substitutes and complements the work of real estate agents to improve efficiency and likely efficacy. With regard to marketing, technology also allows the company to circumvent the strictures of the MLS. Company leadership envisions an even wider scope of applications for its machine learning, including home insurance, mortgage brokerage, and escrow provision. Should such innovations materialize, one might expect that workflows in those areas would be similarly organized around technological capabilities and utilizing human labor to do what technology can't.

REX's use of technology unambiguously reduces the firm's demand for skilled labor. Far fewer real estate agents are needed to serve a given market, given the wide scope of market discovery and customer acquisition tasks performed by the company's algorithms. A small staff of marketing specialists can design digital advertising campaigns for launch across multiple metropolitan areas. Entering a new market requires just two staff initially. On the other end of the transaction, REX

reduces if not eliminates the need for buyers' agents (and offers such services itself to boot). Human agents and marketing staff remain important and necessary to REX, but in much smaller numbers than most other brokerage firms.

4.3 Sky Ridge Medical Center

SRMC adopted its Meditech EHR system as part of an initiative by their parent company, HCA Healthcare. During the mid-1990s, corporate leadership recognized that EHRs were becoming state of the art among large hospital systems, making adoption critical to maintaining a competitive advantage. Policy also played a role in the company's decision to adopt an EHR: the Medicare Modernization Act of 2003, for instance, required that most hospitals implement systems to support electronic prescribing by 2008 or face reduced Medicare payments. SRMC has used the Meditech system since the hospital opened in 2003.

Adoption took place in stages, with each stage incorporating progressively more complex types of documentation corresponding with the needs of different care providers. Implementation of the Meditech system proceeded in the following order: nursing, medication administration, provider order entry, and provider documentation. Initially, the system primarily supported inputting and retrieving basic patient information (in the form of "charts"), such as medical history and medications taken. The first stage of adoption therefore involved mostly nurses, whose responsibilities include entering and maintaining patient charts, preparing

rooms and patients for care, and communicating their actions and observations with other nurses and care providers in their unit.

Next, SRMC implemented Meditech documentation templates for medication administration. This stage of adoption required further integrating EHR use into the provision of care, as opposed to simply using the system as a repository of information, necessitating the reconfiguration of existing procedures for care providers. When a care provider decides to administer a medication to a patient, Meditech generates the order for that medication only after the provider meets the system's minimum requirements regarding patient information. If the system contains too little information in the patient's chart to validate the medication order, it rejects that order. Before a provider can administer the medication, they must scan the unique barcode on both the medication and the patient's admission sheet to confirm use of the proper dosage and application. Meditech and the procedures designed to use it therefore serve to validate and standardize provider practice across the hospital.

The final stage of adoption involved implementing Meditech as a provider order entry and documentation system. Because of the breadth of this stage in terms of the number and variety of tasks it covered, it proved the most complex and consequential for both workers and the hospital as a whole. With the adoption of Meditech, caring for most patient needs requires that providers use computer-indicated "order sets," or series' of documentable actions for treatment that draw on best practices (when available). In effect, Meditech requires care providers to follow

corporate procedures and, when possible, best practice guidelines in order to execute treatments. When a patient presents with a risk of blood clotting, for example, providers enter this information into Meditech, which returns an order set that includes evidence-based recommendations for clot prevention. The provider must still choose to implement the computer's recommendation, but would not be able to initiate any treatment until reaching this point in the process. This marked a change in provider workflows and further integrated the EHR into the hospital's operational routines.

Likewise, implementing provider documentation capabilities and procedures affected a wide swath of hospital practices and employees. Providers must document every action they take in caring for their patients. For instance, when a provider uses patient restraints, they must enter this information into the patient's digital chart. Two obvious impacts on workflows include the burden of entering information into the computer as well as the ability to access this information later to improve care coordination and efficiency. Another impact relates more broadly to the development of hospital policies and practices: by improving the visibility of provider actions, SRMC used EHR data to tweak and improve the adoption process, identifying areas of weakness in EHR use.

Indeed, Meditech adoption indirectly affects workflows, practices, and organization through its support of certain corporate functions. In addition to improving the ease and efficiency of compliance reporting (no small task for acute care hospitals and healthcare systems), centralized databases pulling from Meditech

allow for the creation of corporate dashboards to monitor hospital performance. In turn, data informs continuous quality improvement, including adjustments to employee training programs, policy development, and workflow design. Notably, SRMC does not use Meditech to perform billing-related functions.

SRMC's adoption of Meditech effectively turned an existing paper-based system of processes and procedures into an electronic one. Workflows and provider routines accommodated this change, but this did not entail a complete overhaul of established work organization. Providers now use Meditech to record patient interactions (e.g., treatments), and the system generates particular actions and workflows accordingly. Providers' lose some autonomy in exchange for standardization and validation of treatment. Some providers struggle to balance direct patient interaction and data entry, which might negatively affect patient perceptions of care and/or require a provider to spend additional time outside of direct interaction recording information about the interaction. In terms of explicit or codified procedures and workflows, however, adoption of Meditech does not appear to have substantially altered so much as formalized established patterns of care provision. The system complements care providers' work – perhaps not always in a way that they prefer – with no apparent substitution effects for tasks performed by nurses and doctors.

Meditech does have a substitution effect on some low-skill positions, such as unit secretary. Unit secretaries are not required to have a medical background or post-secondary education. Prior to Meditech adoption, unit secretaries would

handwrite provider orders, digitally transcribe patient charts, and secure nurse approval for their work. Even more than the technical properties of the system itself, its integration into the flow of care provision substitutes for much of the unit secretary's work: by the end of a provider-patient interaction, providers have already initiated orders and documented their actions digitally, with no need for additional sign-off. In response, SRMC streamlined its chart check procedures and recombined tasks and responsibilities in order to maintain the unit secretary position.

During and after adoption, SRMC made significant complementary investments in hardware and training. Before Meditech, providers used mobile computers that they wheeled from room to room as they made their rounds. Soon after Meditech adoption began, SRMC installed computers in every patient room to improve ease of use for providers. Over time, the hospital has also increased the level of technical support available to hospital staff, investing in equipment and dedicated support positions. Similarly, SRMC's Meditech training program has evolved since initial implementation. Early on, the hospital's clinical informatics staff worked with Meditech consultants to design and deliver four-hour training modules delivered in person to hospital staff. Now, training is offered through 40-minute online tutorials, minimizing the marginal cost of training and providing more flexibility in module design and trainee utilization.

Meditech adoption has had a modest but positive impact on SRMC's demand for labor, partly because of strategic firm-level decisions. Meditech does not substitute for much if any of a care provider's work, and thus leaves the hospital's

demand for skilled labor undisturbed. Unit secretaries faced the greatest risk of technical substitution, and indeed HCA Healthcare recommended that its hospitals eliminate these positions upon adoption in order to cut costs and avoid redundancies. SRMC decided instead to redesign the unit secretary's job description precisely in order to preserve the positions. For example, while secretaries no longer need to digitally transcribe paper charts, they can serve to maintain and improve hospital efficiency by supporting the processing of orders and the coordination of documentation needs within and across hospital units.

In some areas of the hospital, adoption may even have induced additional staffing. One fully staffed department, the hospital pharmacy, experienced an influx of medication orders as a result of Meditech adoption. Many individual provider practices within the hospital, which also use Meditech but are not affiliated with HCA Healthcare, have hired Nurse Practitioners and/or Physician Assistants to support care provision and order entry due to increased patient volumes. What isn't clear from the data collected for this case study is why Meditech adoption has been followed by increased patient volumes. It could be that EHR implementation makes providers more efficient in each patient interaction, thus enabling them to see more patients than prior to adoption in the same amount of time. The ability to easily share and communicate patients' medical information may also improve healthcare access or utilization. Of course, high volumes could also be a response to some exogenous factor, such as increased health needs, perceptions of healthcare quality,

changes in federal policy (such as the individual mandate of the Affordable Care Act), or macroeconomic or demographic trends.

4.4 UCHealth

UCHealth uses EPIC, an enterprise-wide EHR system integrated with 26 unique ICT systems across multiple hospitals, departments, and functions. UCHealth was formed in 2012, when the University of Colorado hospital merged with several other hospitals in the region to form a healthcare system providing acute, ambulatory, and chronic care. Though University of Colorado hospital used EPIC prior to the merger, the other hospitals involved in the initial merger and many since acquired by UCHealth did not. The firm's core ICT strategy was thus to collapse and optimize multiple, disparate systems into a single standardized and fully integrated EHR, in order to provide a common patient and provider experience at every facility and provide the flexibility and scalability to address a variety of access and care needs across its corporate footprint.

Since UCHealth's inception, EPIC has played a critical role in both firm operations and market strategy. The platform supports all three essential EHR functions of order entry, provider documentation, and billing. Four years into the adoption/integration process, UCHealth achieved HIMSS Stage 7, the industry's highest standard for EHR adoption; just 6.4% of US healthcare systems have reached this level.⁴ The "integrated" EHR system combines an impressive array of functions,

⁴ The Healthcare Information and Management Systems Society (HIMSS) uses an eight-stage (0-7) model to measure the "maturity" of an organization's adoption and utilization of EHR functions.

such as enterprise-planning and human resource management, into a single digital platform. As a result, UCHealth's EHR use extends well beyond the technology's standard functions. Such "advanced" functions include:

- *Provide clinical decision support*, a process in which the computer suggests the appropriate actions to take based on the data input. For instance, when a caregiver creates an order for a medication, the system responds by identifying the best medication based on data previously entered into the patient's record, e.g., allergies and genetic makeup.
- *Support cloud-based machine learning*. UCHealth deploys analytics at three levels: descriptive, predictive, and prescriptive. EPIC has the capability of using data to summarize past outcomes, identify probable future outcomes, and prescribe actions to take based on those analyses. For example, operating rooms (OR) constitute the most valuable space in a hospital, with every 1% increase in room utilization generating roughly \$100,000 in annual revenue. Every hospital faces the dilemma of maximizing OR efficiency, a complex task of balancing floor space, equipment, staff, and varying patient needs. EPIC's prescriptive analytics enables corporate leadership to monitor efficiency and patient outcomes across all 80 of its ORs and take immediate action to improve or replace inefficient surgeons.

Features of stage 7 include: a hospital has eliminated the use of paper charts in all departments; uses a variety of data types (e.g., documents, medical images) in its EHR, warehouses and analyzes clinical data to improve quality, safety, and efficiency; shares data via standardized electronic transactions or a health information exchange; and uses data to inform enterprise governance. In its 2017Q4 survey of 5,487 US hospitals, nine of ten (88.9%) were rated between stages three and six, with two-thirds (66.7%) at stage five or six.

- *Create a completely digital patient experience.* Through its secure web portal and mobile app, UCHHealth provides its patients with access to their health information; enables them to schedule appointments, communicate with doctors, view test results, and transfer their records to out-of-system caregivers; and allows for the conduct of virtual visits for many common conditions.

The EPIC integration process involved three basic components: data migration, staff training, and process redesign. First, dedicated IT staff pulled data from legacy EHR systems and mapped it into EPIC. Second, all hospital staff received training geared toward specific job categories, including front desk staff, unit secretaries, nurses, and physicians. Third, IT staff worked with hospital administrators to develop new workflows and procedures appropriate to the technical capabilities of EPIC. During the initial integration process in 2012, and again with each newly acquired hospital, UCHHealth dedicates a significant amount of dedicated IT staff support for a two week “go live” period, after which support is scaled back.

Redesigning workflows constitutes perhaps the most consequential element of the adoption/integration process at UCHHealth. Exploiting the technical capabilities of EPIC, both as an EHR and as an integrated, enterprise-wide ICT platform, has involved reorganizing work around technology, rather than simply inserting digital technology into human- and paper-based procedures. An example of a patient requiring intravenous medication helps to illustrate this shift. A nurse begins by

entering patient information into the system. Based on that information, as well as the patient's medical history and past actions taken to address their issue, EPIC generates an order for an infusion pump. In less technologically advanced systems and hospitals, the nurse would perform the rest of the work: programming the pump, monitoring and adjusting medication levels, and so on. The EHR may provide certain "checkpoints" to ensure the nurse follows proper procedures and best practices, but leaves the human care provider in charge of implementing the procedure. At UCHealth, EPIC itself programs the pump based on the patient's vitals, administers the medication, and monitors and adjusts medication levels, with the nurse simply queuing and validating the computer's actions.

From unit assignments to clinical decision support, EPIC drives the organization of work at UCHealth. The ways in which it complements and substitutes for human labor largely determine the direction that change in work organization and workflows takes. Like other EHR systems, EPIC complements the high-skill work of care providers by supporting information access and sharing, which can improve both efficiency and quality of care (e.g., through better care coordination). Digital interactions with patients may also speed the delivery of care and reduce unnecessary office visits, though confirming this possibility exceeds the scope of the present case study. EPIC also complements the high-skill work of corporate leaders, especially those in human resources and technology-related departments, as tasks like performance monitoring and cross-department coordination become almost seamless.

At the same time, EPIC's substitution effects abound. EPIC now performs some of the work previously done by care providers, such as in the above example of medication administration. Insofar as this frees up providers to perform other, perhaps more complex or nuanced aspects of patient care, the complementary effects may nonetheless dominate. The effect on labor of advanced functionalities like prescriptive analytics is even more ambiguous. If a computer takes on increasingly complex tasks involving higher-level critical thinking, where does this leave doctors? Does it free them up to do even more nuanced, non-routine tasks, or does it simply substitute for a "task" considered central to the science (and art, one might add) of medicine? What might the "even more nuanced, non-routine tasks" include, after all? A similar case may be raised regarding the corporate-level use of EPIC for performance monitoring: it substitutes for data collection, communication, and analysis tasks performed by (generally) white collar professionals; in doing so, it also complements their efforts to improve staffing allocation or technology deployment. Which effect dominates is not clear from this case study.

Further down the skill ladder the consequence of UCHealth's adoption and advanced use of EPIC becomes clearer. Like most EHRs, EPIC substitutes for much of the work performed in the past by unit secretaries, transcriptionists, and other positions filled by individuals that generally lack a medical background and postsecondary degree. The integration of multiple information systems across hospitals and departments appears to reduce even further the need for this type of staff. Though UCHealth does not immediately eliminate redundant positions, it

allows for attrition by not filling vacated positions. Even positions with slightly higher skill requirements, such as medical assistants, face substitution due to EPIC. Basic administrative tasks, such as taking patients to their rooms and taking their vitals, remain the function of medical assistants. As the above example of the infusion pump shows, however, EPIC is increasingly able to perform basic tasks with minimal input and validation from a care provider. In the future, telemedicine – a core component of UCHealth’s market strategy – may significantly reduce demand for medical assistants by improving efficiency: one or two working remotely may be able to room hundreds of patients, with basic procedures completed using digital sensors feeding directly into the EHR.

5. Discussion of Findings

5.1 *Within-Case Evidence*

The initial aim of my analysis is to discern the suitability and usefulness of the theoretical framework presented in Section 3. This is especially important because the literature from which I constructed the framework has focused primarily on the impact of computers in production industries. I therefore begin by posing a series of closely related questions: Does the framework apply to firms in service industries? Does the framework apply in cases where new ICT is involved? Do the predictions generated by the framework bear themselves out in the case study findings? Where there are gaps between the framework and my findings, what do they suggest about the relevance and accuracy of the framework?

Overall, the framework performs well both as a descriptive and an analytical tool. It effectively outlines the general dynamics of technology adoption, complementary changes, and their resulting effects on labor demand, even when extended to industries and technologies that have previously received limited attention in the literature. In the case of REX, a startup who did not *adopt* technology in the strict sense, the framework is a bit more of a stretch as a descriptive tool, though the variables of interest remain important and useful data points for understanding the firm's demand for labor in relation to its use of technology. The framework also grounds sound expectations for the case study findings. All three

firms made important complementary changes during and even after ICT adoption, specifically in the areas of work practices, organization, and flows. Perhaps most compelling, and a helpful confirmation of the theoretical motivation behind this study, are findings in each case that the presence of complementarities contributes to expanding the scope of which tasks ICT can substitute, i.e. by enlarging the range of tasks which can be classified as “routine.” Also notable are the dynamic relationships between firm choice, both operational and strategic, and changes in labor demand.

As discussed below, both the choices and their consequences on labor demand differ in important ways across firms; here, the important point is that the framework identifies this possibility and offers a way for research to illuminate how it manifests in specific examples. Put another way, the theoretical framework offers an effective looking glass for exploring the causal mechanisms that link technology adoption and labor demand. The case studies confirm the validity and significance of the five variables of interest I identify in Section 3; conversely, the case study findings suggest that these variables support a thorough investigation of the subject matter.

Before comparing the case studies, it is worth examining in more detail the still somewhat ambiguous feedback loop between technological substitution for routine tasks and its firm-level complementarities. My theoretical framework poses the likelihood that, as firms adopt ICT and implement changes in work practices and organization, they may find previously unplanned (or even unforeseen) points at which substitution becomes possible. Specifically, a task that at first appears to be

non-routine may, upon the advent of new technical capabilities and in the reordering of tasks and jobs, fall under the domain of routine and thus substitutable tasks.

The problem is that, while contemporary theory identifies this possibility, it has so far made relatively less progress in defining and understanding its consequences. On the one hand, the framework as I present it provides a critical starting point for closely examining the complex intra-firm dynamics involved in substitution and reorganization. My case studies may provide a template for beginning to ask this question in more detail. On the other hand, the framework offers little in the way of specific language or analytical proposition to do more than recognize the problem. Future case study research would do well to incorporate a more robust and appropriate conceptual apparatus for unpacking the processes and relations by which apparently non-routine tasks come to be substituted. What can (or should) researchers designate with the term “routine”? For instance, is machine learning itself a broader process of routinization, or something entirely different? What are the respective roles of technical innovation and firm choice in widening the circle of substitution? Such issues inevitably involve both the technical and economic dimension of innovation and adoption, demanding cross-disciplinary research and dialogue in order to arrive at an appropriate language and set of concepts capable of guiding analysis.

The case of UCHealth illustrates well this dilemma. The hospital system’s EPIC platform not only substitutes for routine tasks such as transcription and order processing, but also provides clinical decision support and prescriptive analytics,

which involve non-routine actions like exercising judgment in light of available data and personal and institutional experience. In other words, EPIC allows UCHHealth to take a non-routine task and make it routine: an act of discretion turned into the mere implementation of procedure, as if completing a checklist. This transformation is only possible insofar as UCHHealth redesigned workflows, however, to incorporate the advanced capabilities of EPIC into its service delivery operations. While the theoretical framework sheds light on this emergent situation, however, it does not offer a way to answer the questions it poses about the (disappearing?) limits to automation and the porous boundary between the technological and social-economic factors involved in mediating the effect of technological change on workers. Eventually, this challenge may call for models of labor demand that reach beyond the current focus on the domain of the routine.

5.2 Cross-Case Comparison

The three firms tend to represent three points on a spectrum in terms of the extent to which they integrate technology and its complements into operations and firm strategy, with varying consequences for labor demand. Broad similarities exist among the three firms, especially insofar as they use technology to substitute for routine tasks and implement new workflows to accommodate and leverage new capabilities, as predicted by the theoretical framework. However, the unique patterns of technology adoption and use, as well as the specific changes to work practices and organization made by each firm, set them apart and demonstrate how technological

change can look very different across firms, with important consequences for firm-level labor demand.

In general terms, all three firms exhibit the same pattern of using ICT to substitute for routine information processing tasks. Both SRMC and UCHealth use EHR systems to perform tasks like transcription, archiving, and order transmission. UCHealth's EPIC system also performs slightly more complex routine tasks, like programming medical devices. Accordingly, the substitution effects of the two EHR systems impact most directly low- and middle-skill positions: transcriptionists, unit secretaries, and nurses. Perhaps unsurprisingly, both firms made subsequent investments in additional hardware, training, and technical support. Both also redesigned workflows and practices, though in quite different ways. The most notable similarity in this regard, which follows mostly from the technical properties of EHRs, involves the use of the EHR by doctors to document their actions and generate orders, often while in the room with the patient.

The case of REX presents a more ambiguous example of substitution for routine tasks. The company's advanced algorithms effectively reduce non-routine information processing tasks into codifiable procedures (routines). In this way, machine learning substitutes for complex knowledge work that typically involves not only brute processing skill but also social and emotional intelligence, critical thinking, and human discretion. Routinization, made possible with technologies that analyze enormous volumes of myriad granular data, paves the way for substitution in a domain previously considered uniquely human. REX's technical capabilities

thus necessitate very different workflows from traditional real estate brokers, so that the firm can leverage both the efficiencies generated through substitution in areas like market research and developing ad copy as well as the powerful complementary effects for REX agents.

Beyond such generalizations, however, the similarities mostly end. In assessing the differences among the three firms, several findings stand out. One notable difference involves the firms' divergent rationales for adoption, a thread that runs through much of the following discussion. Indeed, closely related to the rationale for adoption is the extent to which firms integrate technology into service delivery operations and market strategy, which varies in important ways across the case studies. Together, these two findings provide crucial insights into what new technologies and complementarities mean for labor demand at the firm level.

Each firm expressed a distinct rationale for adopting ICT. For REX, technology is inseparable from market strategy: the efficiencies that arise from using machine learning to substitute for key tasks related to market discovery, customer acquisition, and advertising enable them to undercut the fees of traditional real estate brokers. In contrast, SRMC adopted its Meditech EHR system in order to remain competitive in a changing market, i.e. to keep up with industry and regulatory standards. UCHealth occupies a space in between REX and SRMC. Similar to REX, UCHealth's use of ICT and the way the firm distributes information-processing tasks accordingly is integral to their market strategy, which involves reaping economies of scale and continually innovating to secure a technological lead over other health care

providers. At the same time, UCHealth also integrated their EPIC system in order to improve the quality and consistency of service delivery, much like SRMC, even if for slightly different reasons.

In line with the varying reasons and goals given for ICT adoption, the three firms put their technologies to use in starkly different ways. SRMC makes the most limited use of its EHR, focusing primarily on using Meditech to process, store, and share information directly related to care provision (though not billing). To be sure, even this type of utilization requires modifying workflows and practices to meet new demands and leverage new capabilities. The most obvious way it does so is by imposing formal procedures on care providers: they must meet certain informational requirements and, where applicable, adopt best practice recommendations in order to submit an order, e.g. for medication. In a small but nonetheless important way, this routinizes the work of the provider, reducing the space for professional judgment and altering the workflow involved in treating a patient. Corporate dashboards fed by Meditech data further allow for the firm's leaders to monitor hospital performance and make recommendations to improve care quality and consistency. Yet this latter feature underscores the general theme of SRMC's approach: for the most part, paper-based flows and procedures have become electronically-based, but without a substantial change in the role of care providers within those flows and procedures. Implementation of Meditech necessarily changes the job of lower-skilled workers who primarily performed routine information processing tasks, making some of those tasks redundant; yet as discussed below, such positions still serve to monitor

and validate EHR input as well as facilitate the sharing of information across the hospital.

UCHealth's use of EPIC, on the other hand, appears to have done more to transform work in the firm's many hospitals. With EPIC, clinical decision support extends the role of the EHR beyond the best practice recommendations provided by SRMC's Meditech. More to the point, the very organization of the hospital, in terms of both space utilization and personnel, is determined with the significant assistance of the EPIC's prescriptive analytics function. More directly related to the motivations of this study, EPIC changes the role of middle-skilled care providers, i.e. nurses, in subtle but important ways. Rather than performing most of the routine tasks of care provision, from information processing to programming medical devices, they play the role of monitor and validator, ensuring that the EHR system takes the proper actions and providing human input when necessary. This implies a different role from that of care providers in SRMC, even if in a seemingly minute way: now, the EHR carries out many of the tasks of the care provider, with input from, but not always under the direction of, the human worker.

The case study of REX illustrates most clearly what it means for technology and complementary workflows to transform the role of a skilled worker. The very nature of being a real estate agent – the skills it takes to be successful as well as the basic tasks one must perform as an agent – takes on a very different look from a traditional brokerage. Though in some sense the tasks delegated to the agent remain essentially the same – pitch to clients, show homes, and execute transactions – they

are situated around and within a flow of work dictated by the firm's ICT. The latter identifies clients, targets them with customized advertising, adapts that advertising and home listings to capture buyer interest, and even provides a platform for distributing the remaining work to human agents. In this sense, the technology and the workflows that accompany it are inextricable: the only way for REX agents to carry out their job is with the assistance of the technology, given the unique flow of tasks in the firm, while the technology itself obviously necessitates a very different way of organizing and distributing agents' work.

The different purposes and extent of technology adoption and complementarities notwithstanding, the case studies capture how deliberate choice continues to shape firms' demand for labor. At one end of the spectrum, SRMC opted to maintain potentially redundant positions by reorganizing workflows and bundling non-substitutable tasks with emergent tasks, often related to supporting electronic information flows, to repurpose its least-skilled workers. UCHealth does eliminate redundant positions, but only through attrition. REX, motivated by a strategy of keeping brokerage fees low and rapidly expanding into new markets, maintains a minimal staff relative to the other firms studied as well as traditional real estate brokerage firms. By way of "uberization," they also streamline the work of real estate agents, contributing to lower demand for skilled labor. The upshot of this approach may yet be to de-skill the labor force in real estate. Though the data here cannot support that conclusion, the case study of REX points to the possibility that the job of real estate agents may someday be reduced to a handful of non-routine but

relatively straightforward tasks that require few of the characteristics, such as social adeptness and large personal networks, that help traditional agents thrive.

Put another way, under the right conditions, even highly skilled workers may find some of their work taken up by ICTs, complementing (and raising the value of) those tasks which remain under the purview but also potentially limiting the size of the labor pool necessary to carry out those remaining tasks. Without generalizing this observation, it seems likely that the proscription of skilled work in this way may strike some as a new frontier for automation. For the purposes of this study, it is most important to note that such automation is only possible when firms alter the arrangement of jobs and service delivery operations to accommodate and exploit new technological capabilities. Put another way, such automation is not inevitable, either from a technological or economic standpoint. It happens as a result of the confluence of two innovations: one technological, which presents firms with new capabilities to leverage; the other organizational, which involves the clever ways of leveraging those capabilities vis-à-vis new ways of distributing and organizing work, a la UCHealth and especially REX.

6. Conclusion

Though the three case studies do not support general claims about the direction, magnitude, and pace of automation, they can help economists sharpen the analytical tools available for answering such broad and important questions. Current theory regarding routine-biased technical change and firm-level complementarities remains relevant even when applied to cases beyond its original domain. However, in this context the theoretical framework appears to reveal a key limitation surrounding the definition and scope of “routine” tasks. It also suggests that the nature of routine tasks is itself influenced by particular ways of distributing, organizing, and implementing all of the tasks involved in a firm’s service delivery (production) operations. In determining how a technology will substitute or complement for certain kinds of tasks, changes in the properties and capabilities of a given technology matter, but so do the unique and contingent work practices and procedures a company utilizes. The two influence each other, sometimes (as here) to produce very different patterns of technology adoption, workflows, and labor demand. In short, the conditions under which technology adoption results in substitution and complementary effects for workers involves not just the matching of task and technological capability, but also the firm’s strategic and operational decisions related to adoption.

Future research can build on my findings by analyzing a wider range of case studies and pairing detailed qualitative research with aggregate empirical analyses. The institutional and policy setting of technology adoption decisions and their ramifications for labor demand also represents a critical area in need of more investigation. As I describe above, evolutionary economic theory provides a useful framework for placing firm-level data in a broader context, yet some evolutionary theorists themselves acknowledge the school's lack of attention to labor economics (Nelson, 2009). The intersection of labor economics with broader evolutionary theorizing about the social and political dynamics of technological change thus looks like an area ripe for innovative applications to problems of enormous contemporary significance. Such efforts could also support research into both micro- and macro-economic policies that consider employment outcomes related to automation. Though economists have not failed to devote substantial attention to redistributive policies that promise to alleviate unemployment and inequality, they have focused less on how policies can target particular industries or types of firms as they adopt technologies in order to channel the microeconomics of automation.

Bibliography

- Acemoglu, D., & Autor, D. (2010). Skills, Tasks and Technologies: Implications for Employment and Earnings. *NBER Working Paper Series*, 16082.
<https://doi.org/10.3386/w16082>
- Acemoglu, D., & Restrepo, P. (2017). Robots and Jobs: Evidence from US Labor Markets. *NBER Working Paper Series*, 23285. <https://doi.org/10.3386/w23285>
- Adler-Milstein, J., Holmgren, A. J., Kralovec, P., Worzala, C., Searcy, T., & Patel, V. (2017). Electronic health record adoption in US hospitals: the emergence of a digital “advanced use” divide. *Journal of the American Medical Informatics Association: JAMIA*, 24(6), 1142–1148. <https://doi.org/10.1093/jamia/ocx080>
- Akst, D. (2013). Automation Anxiety. *Current*, (558), 3.
- Autor, D. H. (2014). Skills, education, and the rise of earnings inequality among the “other 99 percent.” *Science*, 344(6186), 843–851.
<https://doi.org/10.1126/science.1251868>
- Autor, D. H. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives*, 29(3), 3–30.
<https://doi.org/10.1257/jep.29.3.3>
- Autor, D. H., Levy, F. S., & Murnane, R. J. (2001). *Upstairs, Downstairs: Computers and Skills on Two Floors of a Large Bank* (SSRN Scholarly Paper No. ID

- 289001). Rochester, NY: Social Science Research Network. Retrieved from
<https://papers.ssrn.com/abstract=289001>
- Avent, R. (2016). *The wealth of humans: work, power, and status in the twenty-first century* (First U.S. edition.). New York: StMartin's Press.
- Barley, S. (1986). Technology as an Occasion for Structuring: Evidence from Observations of CT Scanners and the Social Order of Radiology Departments. *Administrative Science Quarterly*, 31(1), 78.
<https://doi.org/10.2307/2392767>
- Bartel, A., Ichniowski, C., & Shaw, K. (2007). How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills. *The Quarterly Journal of Economics*, 122(4), 1721–1758. <https://doi.org/10.1162/qjec.2007.122.4.1721>
- Berg, M. (1980). *The machinery question and the making of political economy, 1815-1848*. Cambridge [Cambridgeshire] ; New York: Cambridge University Press.
- Berger, T., & Frey, C. (2016). Structural Transformation in the OECD: Digitalization, Deindustrialization, and the Future of Work. *OECD Social, Employment, and Migration Working Papers*, (193), 0_1–52.
<https://doi.org/10.1787/5ilr068802f7-en>
- Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence. *The Quarterly Journal of Economics*, 117(1), 339–376.
<https://doi.org/10.1162/003355302753399526>

- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies 'Engines of growth'? *Journal of Econometrics*, 65(1), 83–108.
[https://doi.org/10.1016/0304-4076\(94\)01598-T](https://doi.org/10.1016/0304-4076(94)01598-T)
- Brynjolfsson, E. (2010). *Wired for innovation: how information technology is reshaping the economy*. Cambridge, Massachusetts: MIT Press.
- Brynjolfsson, E., & Hitt, L. M. (2000). Beyond Computation: Information Technology, Organizational Transformation and Business Performance. *Journal of Economic Perspectives*, 14(4), 23–48.
<https://doi.org/10.1257/jep.14.4.23>
- Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: work, progress, and prosperity in a time of brilliant technologies* (First edition.). New York: WWNorton & Company.
- Bureau of Economic Analysis. (2018, November 1). Value Added by Industry.
Retrieved from <https://apps.bea.gov/iTable/iTable.cfm?ReqID=51&step=1>
- Bureau of Labor Statistics. (2018, April 13). Fastest Growing Occupations: 20 occupations with the highest percent change of employment between 2016-26.
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting & Social Change*, 114(C), 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- Goldin, C., & Katz, L. (1998). The Origins of Technology-Skill Complementarity. *The Quarterly Journal of Economics*, 113(3), 693–732.
<https://doi.org/10.1162/003355398555720>

- Goldin, C., & Katz, L. (2008). *The race between education and technology*. Cambridge, Mass.: Belknap Press of Harvard University Press.
- Gordon, R. J. (2016). *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*. Princeton: Princeton University Press.
- Henry, J., Pylypchuk, Y., Searcy, T., & Patel, V. (2016). *Adoption of Electronic Health Record Systems among US non-Federal Acute Care Hospitals: 2008-2015* (No. ONC Data Brief, No. 35). Washington, D.C.: Office of the National Coordinator for Health Information Technology.
- James, H. (2016, September 23). Déclassé: Nothing New Under the Sun. *The American Interest*, 12(2). Retrieved from <https://www.the-american-interest.com/2016/09/23/declasse-nothing-new-under-the-sun/>
- Konishi, Y. (2017). *Toward artificial intelligence technologies as an integral part of service business operations*. London: Center for Economic Policy Research.
- Lee, K.-F. (2018). *AI Superpowers: China, Silicon Valley, and the New World Order*. Boston: Houghton Mifflin Harcourt.
- Milgrom, P., & Roberts, J. (1990). The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *The American Economic Review*, 80(3), 511.
- Milgrom, P., & Roberts, J. (1995). Complementarities and fit strategy, structure, and organizational change in manufacturing. *Journal of Accounting and Economics*, 19(2), 179–208. [https://doi.org/10.1016/0165-4101\(94\)00382-F](https://doi.org/10.1016/0165-4101(94)00382-F)

- Mill, J. S. (1848). *Principles of political economy: with some of their applications to social philosophy*. Toronto: University of Toronto Press.
- Mishel, L., & Bivens, J. (2017). *The zombie robot argument lurches on: There is no evidence that automation leads to joblessness or inequality*. Washington, D.C.: Economic Policy Institute. Retrieved from <http://www.epi.org/publication/the-zombie-robot-argument-lurches-on-there-is-no-evidence-that-automation-leads-to-joblessness-or-inequality/>
- Mokyr, J., Vickers, C., & Ziebarth, N. L. (2015). The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different? *Journal of Economic Perspectives*, 29(3), 31–50. <https://doi.org/10.1257/jep.29.3.31>
- Nelson, R. R. (2009). Technical Change and Industrial Dynamics as Evolutionary Processes. Laboratory of Economics and Management LEM, Sant'Anna School of Advanced Studies, Pisa, Italy. Retrieved from <http://EconPapers.repec.org/RePEc:ssa:lemwps:2009/07>
- Pratt, G. A. (2015). Is a Cambrian Explosion Coming for Robotics? *Journal of Economic Perspectives*, 29(3), 51–60. <https://doi.org/10.1257/jep.29.3.51>
- Ransbotham, S., Kiron, D., Gerbert, P., & Reeves, M. (2017). *Reshaping Business With Artificial Intelligence*. Cambridge, MA: MIT Sloan Management Review and the Boston Consulting Group.
- Seamans, R. (2018). AI, Labor, Productivity and the Need for Firm-Level Data. *NBER Working Paper Series*, 24239. <https://doi.org/10.3386/w24239>

Tripathi, M. (2012). EHR Evolution: Policy and Legislation Forces Changing the EHR. *Journal of AHIMA*, 83(10), 24–29.

Vivarelli, M., & Pianta, M. (2000). *The employment impact of innovation: evidence and policy*. London ; New York: Routledge. Retrieved from <http://ebookcentral.proquest.com/lib/du/detail.action?docID=180324>