NBER WORKING PAPER SERIES

PAYING TO PROGRAM? ENGINEERING BRAND AND HIGH-TECH WAGES

Prasanna Tambe
Xuan Ye
Peter Cappelli

Working Paper 25552
http://www.nber.org/papers/w25552

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2019

We have greatly benefited from feedback from April Franco, Brad Greenwood, Damien Joseph, Boyan Jovanavic, Chris Kemerer, Paul Osterman, Amanda Pallais, and Michael Roach as well as seminar participants at the American Economic Association, the Academy of Management, the American Sociological Association, the Wharton People and Organizations Conference, the University of Minnesota, Temple University, the INFORMS Conference on Information Systems and Technology, Purdue University, the Workshop for Information Systems and Economics, New York University, the University of Washington, the Georgia Tech Sandra Slaughter Software Conference, the University of Illinois-Chicago, the NYU CIRI Seminar series, Arizona State University, Harvard Business School, the Federal Reserve Board, the Wharton School of Business, Michigan State University, and the University of Pittsburgh. We are also grateful to Andrew Chamberlain and Glassdoor for providing access to their data. Prasanna Tambe is grateful to the Alfred P. Sloan foundation for financial assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Paying to Program? Engineering Brand and High-Tech Wages
Prasanna Tambe, Xuan Ye, and Peter Cappelli
NBER Working Paper No. 25552
February 2019
JEL No. J24,J33,J41,L86

ABSTRACT

We test the hypothesis that IT workers accept a compensating differential to work with emerging IT systems, and that employers that invest in these systems can, in turn, capture greater value from the wages they pay. We show that much of the utility IT workers derive from these systems is from skills acquired on the job. This is principally true for younger workers at employers where skill development is encouraged, and the effects are stronger in thicker markets where workers with newer skills have more outside options. An analysis of the text in online employer reviews supports the notion that IT workers value access to interesting IT systems above most other employer attributes. These findings are important because first, they provide evidence of how worker preferences can influence corporate IT investment decisions; second, because they shed light on factors influencing IT skill development; and third, because they point to a potentially important explanation for returns from IT investments.

Prasanna Tambe
Wharton School
University of Pennsylvania
3730 Walnut Street
558 Jon M. Huntsman Hall
Philadelphia, PA 19104
tambe@wharton.upenn.edu

Xuan Ye
Carroll School of Management
Boston College
Fulton Hall 460E
140 Commonwealth Ave
Chestnut Hill, MA 02467
xuan.ye@bc.edu

Peter Cappelli
The Wharton School
Center for Human Resources
University of Pennsylvania
Philadelphia, PA 19104-6358
and NBER
cappelli@wharton.upenn.edu

I. INTRODUCTION

As business process digitization becomes increasingly important for explaining corporate performance, access to IT labor has become a critical differentiator for firms. The supply of trained workers is increasingly seen as a key factor influencing investment in, and the use of, IT systems (e.g., National Academy of Sciences, 2007). The notion that there might be shortfalls in IT labor that should be addressed through immigration has led to public debate. Competition for high-tech workers has also received scrutiny as employers have used aggressive tactics to manage employee hiring and retention—including allegedly collusive measures—drawing attention from the US Department of Justice (Streitfeld, 2014). It has become important therefore, to understand how firms attract and develop high-tech talent. In this study, we test the hypothesis that IT workers accept a compensating differential to work with newer IT systems.[1] Employers, therefore, compete both on the wages and the technologies they offer to prospective workers. These arguments are consistent with evidence showing that some IT investments attract higher quality IT workers (Ross, 1994; Fichman and Kemerer, 1997), and with a growing focus on the importance of a firm's "engineering brand" for recruiting tech talent.[2]

IT workers may prefer working with newer technologies for a variety of different reasons. Human capital theory predicts that workers pay to acquire skills that could be useful later in their careers (Becker, 1964; Ben-Porath, 1967).[3] Additionally, the research on motivation in the management literature has shown that workers prefer jobs that have more task variety (Hackman and Oldham 1980). Because technology changes quickly, it can increase task variety by generating new tasks or creating new ways of performing existing tasks. Other factors, such as the meaningfulness of work, can also influence how workers feel about using

---

[1] Anecdotally, it is not uncommon to hear the following types of suggestions exchanged among IT workers: "I would see working for a cutting-edge company at lower pay worth it just to gain the experience and skills which could be used later on and result in a higher billable rate" (Accessed at http://www.freemoneyfinance.com/2007/10/help-a-reader-a.html on July 24, 2015) or "To me, it may be worth staying (temporarily) at a job where you feel that you are underpaid in money, if you are getting lots of opportunities to learn new skills …. Think of it as similar to the time spent in an MBA or other post-graduate training program." (Accessed at https://news.ycombinator.com/item?id=2309317 on July 24, 2015).

[2] A body of academic work discusses the broader concept of employment brand (Aurand et al 2005; Kim 2012), and the press has more recently discussed the concept of "engineering brand" for high-tech firms. For example, see http://blog.entelo.com/why-an-engineering-employer-brand-is-crucial-to-recruiting-tech-talent. Accessed on Aug 14, 2015.

[3] Popular websites such as Stackshare.io (https://stackshare.io/) exist principally to enable employers and developers to share information about which technologies are used by employers.

new technologies. The goal of this paper is to provide evidence that employees accept a compensating differential to work with newer technologies, and to emphasize the role of acquiring human capital as one key factor shaping these preferences. Parsing out the relative importance of the various factors in motivating workers to prefer one job over another is an important objective but it is a considerable exercise that is a focus of fields like organizational behavior, and is left to specialists in job motivation.

Estimating compensating differentials associated with job characteristics, such as the use of specific technologies on the job, is challenging. When workers pay for some of the costs of acquiring new skills, labor market conditions can affect the extent to which employers and workers share these costs. If labor market conditions are held constant, one could offer workers a schedule of wages and technologies to estimate the partial effect of different technology choices on the worker's wages. However, such an approach would be difficult in a field setting. Few people simultaneously receive offers using multiple, identifiable technologies; obtaining reservation prices for each one would be difficult, and it would lack realism unless job seekers were prepared to say yes to each job. Moreover, offers are not exogenous: candidates are selected based on their human capital.

An alternative approach, similar to the one used in this study, is to elicit the wage premium required to entice workers away from a job in which they are using a new technology. The difference between this figure and the worker's wage can be partially accounted for by the value workers assign to non-wage characteristics of their current job, and as with the notion of consumer surplus, this value can far exceed the wages they are paid. This approach has the advantage that we can clearly measure whether they are using a given IT system (as opposed to whether a potential job would use it if we observed multiple offers). Of course, measuring the size of this premium is challenging, especially in a market setting where the responses have consequences. It is perceptual in nature, only existing inside job seekers' heads, and to our knowledge is not collected by administrative agencies.

To measure this wage premium, we analyze data provided by a leading US jobs board that captures job search behavior for a large cross-section of IT workers, including the wages they earn at their current employers and the "target" wages they seek in the next job. This target measure is similar to the notion of a reservation wage—in this case, the wage at which employed workers can be induced to provide their services to another employer.[4] These figures are generated by actual job search activity, which is noteworthy for two reasons. Prior data sources on reservation wages, when available, have been based on small samples. The data used in our analysis are available for over 50,000 IT workers. Second, the data are reported in a context where the stated target wages have real consequences: potential employers see them. From job seekers' resumes, we text-mine the technologies they use on the job. From Glassdoor employer reviews, we create measures of the value employees place on the learning culture and IT offered by the company. Our final data set includes a cross-section of employed workers' wages and the wages required to leave their employer, the IT systems the worker uses at the firm, measures of the firm's engineering work environment, and other characteristics of the employer and worker.

To understand the advantages of these data for our objectives, it is instructive to consider problems that arise when using approaches that rely on wage observations for IT workers. In principle, we could develop a proxy for the human capital workers derive from IT by assuming they learn more when working with newer IT, but workers are also productive at different levels with these technologies. Moreover, there may be ability bias, i.e., more attractive technology firms may employ higher-quality workers. In both cases, when using cross-sectional data on wages alone, it becomes difficult to separate pay differentials associated with getting to use newer IT from higher wages paid to more productive workers. Examining the wage premium required to leave the firm, on the other hand, separates the amount workers derive from the IT from the value provided to their employers by using this system. If IT were randomly assigned to workers, a hedonic regression of IT on this premium reveals the value of working with newer IT systems.

Of course, random assignment in this context is infeasible, so our analysis is subject to two sources of bias. First, our estimate of the value of IT can be biased upward if use of new IT systems strongly predicts that the amenities of the current job (excluding human capital) are much higher than what workers expect in

---

[4] Reservation wages have been a key parameter in job search models (Lancaster and Chesher, 1983; Feldstein and Poterba, 1984; Shimer and Werning, 2007). Employed workers choose a reservation wage that maximizes their future income, striking a balance between higher wages and the ability to secure job offers.

the next job. Workers cannot specify the non-wage characteristics they desire in future jobs, which limits the number of job seekers looking to change industries or job contexts, but the potential for such bias is still significant. A second source of bias is that unobservable factors, such as risk preferences, can influence both technology choice and how workers set target wages. For example, workers with low risk aversion might seek a higher target wage and be attracted to riskier, unproven technologies.

To assess the robustness of our estimates to these sources of bias, we first demonstrate that using emerging IT systems is associated with higher target wages, after controlling for current wages and education, job tenure, experience, race, and gender. We do not observe a similar relationship for mature or low-growth technologies, and our estimates are robust to including employer and location fixed-effects, which suggests that the differences in search behavior that we document are more closely related to technology than work setting. Second, these effects are stronger for younger workers, and for workers in markets where there are many other employers hiring for new technical skills, both of which are consistent with the argument that some of this preference is due to workers placing value on gaining skills they can use in the future. Indeed, we find that workers who use these technologies leave their jobs more quickly than other IT workers. When considered alongside our wage results, these facts are striking. Among potential omitted variables, few lead workers to set higher target wages *and* exit the firm more quickly, except human capital. Differences in amenities or compensation that can bias estimates of IT use on target wages should induce workers to stay longer with their employers, not leave faster.

Third, using the text from the Glassdoor review data, we demonstrate that a firm's use of emerging technologies is correlated with its technical employees valuing technology and learning as sources of value derived from that employer, and it is negatively correlated with employees valuing compensation at the firm. In contrast, using mature IT systems is positively correlated with technology and negatively with learning. Employers using older, unpopular systems are mostly valued for their compensation practices. In other words, the technical culture of the employer appears to be matched with the technologies it uses. Moreover, direct analysis of the text data provides further evidence that our IT use measure is reflecting the learning benefits of those organizations, not other job characteristics. We also conduct tests to show that potential confounders of our target wage results must be correlated with IT investment and investments in employee learning, but not either factor in isolation. This would suggest that new technologies are not a proxy for other factors influencing whether IT workers value their jobs. Nor is the relationship we observe likely to be due to other attributes correlated with new IT.

Finally, we demonstrate that our results on technology and job preferences are robust to tests of the effects of differences in equity compensation, bonuses, benefits, and visa restrictions, and we use data on employment flows to show that workers tend to move between firms that are similar in terms of non-wage benefits. Along with the worker controls included in our regressions, these latter results provide supporting evidence that unobserved, worker-specific differences related to IT use are not an important source of bias.

The evidence supports the notion that investments in emerging IT can engage more productive labor to work on it at a given cost, in part because it workers have an opportunity to gain valuable skills. Our main effect estimates range from 2% to 4% for workers in our sample, although in the sample of employers that promotes skill development at work, the estimated magnitude is about 5%. This implies that workers who make $100,000 require an additional $5,000 to leave employment where they use new technologies and this figure is likely to be higher in some markets and for some technologies. For comparison, courses on topics such as web development or python programming offered by coding boot camps normally cost about $4,000, and the costs of immersive boot camps that cover a variety of technical topics average slightly less than $12,000.[5] From an employer's perspective, the value that its workers assign to these opportunities is offset by other costs, such as higher churn rates, because employees who learn new skills tend to move on more quickly, and the firm's IT adoption decision must account for these costs as well. In this paper, we do not make a direct comparison of these figures, but rather, demonstrate that, ceteris paribus, firms that invest in these systems can attract higher quality technical labor for a given wage.

Our findings are important for three reasons. First, they document a relationship between labor market competition and IT investment incentives that has not been explored in prior research; second, they

[5]https://www.coursereport.com/blog/coding-bootcamp-cost-comparison-full-stack-immersives. Last accessed on December 13, 2018.

provide evidence of how, in practice, newer IT skills are acquired and how they can be financed: employees trade wages for the chance to use newer technologies. This has implications for explaining stylized facts related to high-tech markets, such as those related to the difficulty in attracting IT employees with the relevant skills in labor markets characterized by early IT adoption. Third, the findings suggest a missing factor in current IT productivity estimates. Sources of returns from new technologies should include the possibility that technical wages are lower, other things being equal, where new technologies are introduced (see Stern 2004 for a similar argument in an R&D context).

Our analysis sits between several literatures. First, how labor supply affects IT investment has not been widely studied (a recent exception is Branstetter, et al, 2018), but is important for understanding how labor policies influence IT diffusion. This gap is underscored by an existing literature linking R&D investment to scientists' preferences (Stern, 2004; Roach and Sauermann, 2010; Agarwal and Oyhama 2013), and in our focus on skills can be viewed as being nested within a strategy and economics literature that argues that workers exchange wages for acquiring valuable knowledge on the job (Stern, 2004; Franco and Filson, 2006). Second, IT labor markets themselves have been the subject of study (Ang, et al, 2002; Levina and Xin, 2007; Mithas and Krishnan, 2008; Mithas and Lucas, 2010; Tambe and Hitt, 2012; Joseph et al 2012; Bapna, et al, 2013). This work has focused on how human capital, training, and employer characteristics (primarily size and industry) affect IT wages and careers, but has been silent on how IT influences employment choices. Third, our paper is connected to a literature on management practices in IT-intensive firms (Black and Lynch, 2001; Bresnahan, et al, 2002; Bartel, et al, 2007). To the best of our knowledge, it is among the first to discuss management practices related to the firm's technical human capital.

As with all empirical analyses, ours has caveats. Most importantly, there are sampling concerns with these data because participants are actively seeking new jobs. Those who post their resume online are a sizeable proportion of the US labor force—roughly eighteen million people in the United States did so on this job board alone—but our results may not be representative of those who did not post their resumes. However, much of the data used in prior work on IT wages suffer from similar or more severe selection biases. Later in this paper, we report comparisons with administrative data (such as the Current Population

Survey and Occupational Employment Statistics from the Bureau of Labor Statistics) to assess how our sample compares with that of the IT labor force at large.

II. BACKGROUND LITERATURE

A. Mobility, knowledge transfer, and wages

A key question in the academic literature on innovation, dating back to Arrow (1962, p.615), is how labor mobility affects the appropriability of knowledge-based investments. If new knowledge is useful at other firms, workers will pay to acquire it by accepting lower wages (Becker, 1964, Rosen, 1972). This, in turn, can influence how much of their investment firms can appropriate, and what projects they choose. Much of this literature is rooted in R&D investment. These studies demonstrate that the mobility of engineers is a channel for R&D knowledge transmission among firms (Zucker, Darby, and Armstrong, 1998; Almeida and Kogut, 1999). Scholars have shown that scientists accept a compensating differential for R&D experience, such that "spillovers" from the transfer of R&D knowledge are offset by lower scientist wages. Moen demonstrates that technical R&D staff exchange lower wages for gaining R&D experience (2005), and Franco and Filson provide evidence that workers pay for the knowledge acquired from employers using data from the rigid disk drive industry (2006).

A related literature investigates how R&D choices affect the firm's ability to attract scientists. This literature argues that scientists value R&D experience because they enjoy engaging with science and because they acquire knowledge that is valuable for future opportunities. Stern finds that scientists accept compensating differentials when permitted to engage with frontier science (2004). He notes the difficulty in separating wage discounts that scientists accept to work with new science from the bias introduced by the fact that higher-ability workers prefer such jobs. To remove this bias, he collects data on multiple job offers for the scientists in his sample, and uses within-scientist comparisons to estimate wage effects. Recent work has also shown that these preferences have implications for the sorting of R&D workers into industry vs. academic jobs (Roach and Sauermann, 2010; Agarwal and Ohyama, 2013; Roach and Sauermann, 2014), and for the locus of scientific innovation.

B. IT investment and technical careers

Because IT investments are also important for driving productivity, they have also attracted academic attention. IT adoption decisions have been the focus of some of this attention (e.g., see Forman, 2005, or McElheran, 2015). Another strand of this literature uncovers factors influencing IT returns (Brynjolfsson and Hitt, 2000). The most important of these may be the adoption of specific work practices, such as decentralized decision-making and high-powered incentives, which have been shown to complement IT investment (Bresnahan, Brynjolfsson, and Hitt, 2002; Bloom, et al, 2014; Brynjolfsson and McElheran, 2016).

One factor that can influence IT outcomes but that has received less attention is access to labor that can implement and use IT systems. Prior work connects the firm's human capital to IT use (Mata, et al, 1995; Bharadwaj, 2000; Fichman and Kemerer, 1997). Yet, despite the notable success of a literature on IT and labor demand (Bresnahan, Brynjolfsson, and Hitt, 2002; Autor, Levy, and Murnane 2003), how IT investment is influenced by labor supply has only begun to attract attention, and to the best of our knowledge, no prior work examines how IT investment affects which workers firms attract, although our arguments are closely connected to economic theory linking technological change, human capital, and productivity (e.g., Chari and Hopenhayn, 1991; Jovanovic and Nyarko, 1996).

Examining how IT workers' preferences for skills influence their employment choices is also related to the empirical literature on the IT labor force. Existing work in this area focuses on returns to firm-specific and general human capital for IT workers and on institutional determinants of IT wages. These studies analyze data sources such as the Current Population Survey (CPS) or the National Longitudinal Survey of Youth (NLSY) administered by the US government (e.g., Levina and Xin, 2007 or Joseph, et al, 2012), or private survey data such as that collected by *InformationWeek*, a media publication targeted towards the IT industry (e.g., Mithas and Krishnan, 2008). Due in part to the limitations of these data sources, which often do not identify employers, these analyses have been constrained in their ability to examine workers' skills or employers' IT investments.

Nevertheless, understanding employment preferences is particularly important for IT workers because IT skills are often acquired on the job (see Ang, Slaughter, and Ng, 2002; Fong Boh, et al, 2007), so the ability to acquire new skills is a determinant of inter-organizational IT career success (Barley and Kunda, 2006; O'Mahony and Bechky, 2006; Bidwell and Briscoe, 2010; Roberts, et al, 2006). More so than other workers, IT employees may be responsible for maintaining and enhancing their own skills. It is likely, therefore, that IT workers prefer employers that allow them to work with emerging IT systems, which is our key hypothesis.

III. DATA SOURCES AND KEY MEASURES

The key variables and data sources we use are described below and summarized in Table 1.

III.A. Wage and human capital variables

The online job board that provided us with the target wage data is among the largest in the United States. The context makes it clear that prospective employers will see the target wage and interpret it as a statement about what individuals expect from a new employer.

Target wages are similar to reservation wages, which are the wage levels at which workers are indifferent between accepting and rejecting job offers. Reservation wages have important implications for understanding labor market outcomes related to job search, such as labor mobility. The empirical literature on reservation wages has focused on the effects of reservation wage setting on labor outcomes such as unemployment duration, or on factors that affect how workers set their reservation wages, such as unemployment insurance levels. The literature in the latter area typically uses human capital variables such as experience, education, demographic data, and especially employment status to explain reservation wage levels. Although reservation wages are of interest for both employed and unemployed workers, the population of interest in most existing studies has been unemployed workers due to the policy interests around unemployment (e.g., Prasad, 2003). Most empirical work in this area has used surveys asking workers to report their reservation wages (Lancaster and Chesher, 1983; Feldstein and Poterba, 1984; Jones, 1988; Falk, et al, 2006; Hall and Krueger, 2012; Krueger, 2014). These studies have varied in the definitions used. For example, the surveys used by these authors have asked survey respondents for their "desired wage" or for their "asking wage."

As such, the level at which workers set their target wages on the platform we study indicates something important about the pay at which an individual is willing to accept a new job offer, other things equal. These target wages are particularly notable because the consequences associated with answering the question make respondents take the answer seriously. Respondents on this platform also report their current wages, or if unemployed, their most recent wages.[6]

Self-reported wage data of this type have many potential biases, although virtually all labor market studies are based on self-reported data. Individuals may have incentives to inflate their current wages as doing so can affect wage offers from employers.[7] However, there are sanctions against doing so in this context. For instance, inflating one's wages may reduce interest from other employers given that a surprising number of job offers do not appear to allow for wage negotiations (see Hall and Krueger, 2012).

Target and current wages are available for workers who participated on the site through 2007. We keep only employed IT workers who include information about technologies on their resumes,[8] which generates a sample of slightly over 50,000 workers. Workers report their occupation using a drop-down menu. We do not have access to panel data because we do not observe the same workers over multiple years. We can, however, observe how workers differ in their behaviors depending on whether they post before or after the firm adopts a new system. Participants also submit demographic, human capital, and past employment information (i.e. prior jobs and employers) and information about the type of job they seek (e.g., annual or paid hourly, which corresponds to exempt or non-exempt status under the Fair Labor Standards Act). They also identify their MSAs (Metropolitan Statistical Areas).

III.B. Data comparisons

Job board participants are active job seekers who may differ from passive candidates or from those not interested in moving. They may be younger, more mobile, or less likely to be found in firms with developed internal labor markets. On the other hand, job seekers are the principal group of interest for our hypothesis. To characterize these differences, we compare our sample with IT workers in other datasets with known

---

[6] To the best of our knowledge, the only paper that uses similar data is Eriksson and Lagerstrom (2012), which uses Swedish data to examine the wage demands of unemployed workers.

[7] If respondents inflate the value of their current wage, then our dependent variable – the gap between the target and current wage – would be smaller, reducing the range and making it more difficult to find statistically significant effects. This problem is known as range restriction. It is similar to censoring except that what drives the overstated current wage is not easily identified.

[8] Although the data are also available for unemployed workers, we do not include unemployed workers in our sample because the compensating differential is calculated based on how much employees are currently being paid. That being said, ours are reduced form estimates and moving to a tighter labor market may be one reason workers acquire new skills. We cannot sort these explanations out with these data.

sampling properties. For these comparisons, some flexibility was required because the classifications used by government agencies do not fit neatly with the IT classification used on the jobs board. To bridge this gap, we compared IT workers on the job board with workers in the government statistics who fall into "Computer and Mathematical Occupations" (O-Net code 15). Workers in these occupations comprise the majority of the IT workforce, and especially those in technical occupations, such as programmers, network administrators, and database administrators.

Table 2 compares the demographic and human capital statistics of our sample with employed, annually paid IT workers in the 2007 Current Population Survey (CPS), which is administered by the US Census Bureau and designed to be representative of the US workforce. The educational distributions in the two samples are similar. The job board has proportionately fewer white males, but this is difficult to pin down because many respondents did not report their gender in the job board data.

Table 3 compares the wages of job board workers with the wages of workers collected through the Bureau of Labor Statistics Occupational Employment Statistics (OES) program. Like the CPS, data from the OES are statistically representative of the US workforce. We compared wages along two dimensions relevant for our analysis: geography and occupation. When comparing occupations or labor markets, the wages of IT workers in the jobs board sample are not systematically higher or lower than the wages of IT workers collected by the OES. They are close in value in most cases.

These comparisons suggest that the wages of the IT workers in our sample are similar to wages reported in administrative data on IT workers, although they do exhibit important differences. Nevertheless, the results reported below should be interpreted as providing evidence for IT workers who are seeking jobs through online resources. That population appears to be considerable. There are many job boards in addition to the one we access here, and while we do not know how many candidates use multiple job boards, there are millions of individuals who list information on job boards. Kuhn and Skuterud (2000) report that as far back as 1998, fifteen percent of unemployed job seekers, and half of all job seekers who had online access at home, used job boards. A 2014 survey found that one in five individuals in the labor force used job boards (Jobvite, 2014).[9]

III.C. IT use on the job

On their resumes, IT job seekers often include keywords related to programming languages or software packages and the beginning and end dates for projects. Beginning with a list of information technologies that was provided to us by this job board, we algorithmically identify those that appear on a resume and the years in which the user reported working in this job. We also attach the age of each technology, which was gathered by research assistants using online resources.[10] When workers report using multiple technologies in a job, we choose the newest technology. Appendix A discusses some of the details of this process, as well as selection issues that arise from non-reporting of skills.

The framework discussed above suggests that workers place higher value on IT experience that will be of future use (Rosen, 1972), but we face the challenge of identifying the IT systems for which this is true. To the best of our knowledge, no established taxonomy categorizes IT systems. Theoretically, workers derive the greatest value from IT experience that is difficult to acquire through other channels (newer technologies) and when demand is growing. Dying technologies offer less of a future for workers entering the market, and for technologies that are popular and mature, employers expect workers to already possess the relevant skills. For newer technologies, employers expect workers to learn these skills on the job, and for growing technologies, workers are often willing to do so.

We classify technologies by 1) growth rate and 2) whether they are less than ten years old. Ten years is an arbitrary threshold, but technologies are still early in their diffusion path at that point (e.g., even the oldest "big data" technologies are close to within the ten-year window), and our results are not sensitive to shifting this window by five years in either direction. Growth rates are computed from the job search data by measuring the number of firms employing workers who report using that technology from year to year. Further details of how we computed this measure are discussed in Appendix A. Growth rate and age form a

---

[9] http://web.jobvite.com/rs/jobvite/images/2014%20Job%20Seeker%20Survey.pdf
[10] For example, see: https://en.wikipedia.org/wiki/History_of_programming_languages.

2x2 matrix, and Table 4 presents examples of different technologies that, in 2007, fell into each of the quadrants defined by this matrix. We test the hypothesis that workers derive the most utility from working with technologies of the type in the first column--those that are less than ten years of age and with a growth rate is above the median. We refer to these as "emerging" IT systems throughout our analysis.

III. D. Measures of employer attributes

We used online review data, provided by Glassdoor, to create measures of job amenities. Glassdoor is a career intelligence platform where employees anonymously review current and past employers. These reviews provide information about employees' perceptions of their employers (i.e. employer "brand"). In 2014, when we accessed the data, there were over 1.4 million reviews (700,000 for US firms) describing the "pros" and "cons" of employers. To develop measures of what workers value about their employers, we analyzed the text in the "pros" section of the reviews. We used two methods from the literature on text-mining product reviews: 1) unsupervised methods that generate employer features by clustering key phrases, and 2) by manually selecting phrases that indicate learning on the job. Details of these methods are discussed in Appendix A.

There are two caveats. First, participants submitted reviews between 2008 and 2014, but we do not know when they worked at the firms they review, so we build cross-sectional measures of employer attributes. However, the nature of the review process suggests that the reviews describe attributes of a reviewer's employer during jobs ending between 2008 and 2014, so many of these jobs will have spanned 2007, which is year one of our job search data. Ignoring temporal changes in our short window is an approximation. The work practices literature has argued that such practices are costly to adjust and may remain quasi-fixed in the short run relative to changes in IT and labor (Bresnahan et al 2002), but the inability to track changes in work practices is a limitation if employers change attributes quickly.

Another caveat associated with use of these measures is that the review text is constrained by the size of the text entry box (this is similar to answers provided in open-ended survey questions, see Reja, Manfreda, Hlebec, and Vehovar, 2003). Because reviewers have limited space, they list the most favorable employer attributes. Therefore, these measures reflect the ranks of employer attributes, not absolute strengths, so negative coefficient estimates do not indicate that the attributes provide negative utility to the worker, only that they play less of a role than other attributes in target wage setting.

IV. EMPIRICAL FRAMEWORK

We test the idea that the chance to use new technologies in a job moves IT workers along the supply curve for technical work; i.e. at a given price, a firm attracts more workers if they offer more interesting technologies. An assumption is that the demand for IT workers during the period we study is fixed, which is reasonable given that our analysis focuses on a single year. Our measure of interest is within-individual differences between the worker's wage and target wage.

To test the hypothesis that IT workers value working with emerging IT systems, we estimate a model that explains variation in the target wages that employed IT workers post during job search, controlling for current wages and including firm and worker characteristics as explanatory variables. The target wage is equivalent to a reserve price in an online auction with multiple bidders (i.e. the potential employers). For a job seeker *i*, the optimal target wage to post (*r*) is the worker's private value (v_s) for their current job (which includes wages as well as sources of non-monetary value) plus an additional increment based on the distribution of the bidders' (employers') private values (*f(r)*) for the worker's services. It is assumed that job seekers know this distribution. In general, the optimal reserve price solves the following equation, where f(r) is the distribution from which employers' valuations are drawn:

$$(1)\ r = v_s + [(1 - F(r)) / f(r)]$$

The firm's IT investment enters this relationship by increasing the worker's value for the current job.

We assume Cobb-Douglas utility and that agents are risk neutral and have homogenous preferences, which is consistent with much of the theoretical literature on job search that incorporates non-wage job characteristics (e.g., see Blau, 1991). Under these assumptions, consider an empirical (target) wage equation of the form:

$$(2)\ \mathrm{Log}(TW_i) = \beta_0 + \beta_1(EMERGING_i) + \beta_2\mathrm{Log}(CW_i) + \varepsilon_i$$

i indexes the worker. TW is the target wage and CW is worker's current wage, which along with benefits, technology, and other sources of non-monetary value comprise the worker's value for her current job.[11] EMERGING indicates the use of new rather than established technologies on the job. A positive and significant coefficient on β_1 indicates that workers derive value from using these technologies. Because the dependent variable is in logs, the estimates should be interpreted as the percentage change in the target wage associated with a change in the independent variable.

Various market, employer, or worker level attributes can bias β_1. Placing current wages on the right-hand side allows the target wage to be interpreted as a within-person measure: target wage relative to current wage. Current wages absorb many of the factors affecting demand for these workers, except for differences related to the most recent employment spell that can cause the target wage to deviate from the current wage. Differences in the attributes of employers or workers using emerging IT matter only if they affect current or target wages differently from one other.

Our most robust analyses include employer fixed-effects, which eliminates those attributes that are common to all workers at a firm. Our goal is to estimate the partial effect of a worker's choices across different technologies. Across markets, there can be differences that influence how workers set target wages. For instance, in tighter markets, workers may accept less of a pay cut for training (Acemoglu and Pischke 1999), although if firms absorb some of these training costs, our tests should be conservative ones because the effect will disappear in markets for emerging technologies, which tend to be tighter. Nevertheless, to the extent possible, we restrict our comparisons to workers within the same market who are making choices over different technologies. Most of our regressions include MSA and industry fixed effects – to the extent that MSAs and industries can capture the relevant labor market boundaries, these variables should address systematic differences across markets.

[11] There are alternative dependent variables we considered, such as measuring the gap as a percentage of current income, but we use the model in (1) because it is most consistent with the existing literature on determinants of workers' reservation wages. We also tested models using ln(Target wage/Wage) as the dependent variable, which constrains the coefficient on the worker's current wage to be one. To some degree, this imposes a downward bias on the emerging IT variable, which is consistent with what we observe. The IT results using this specification are still significant but smaller in magnitude and less precisely estimated.

We do not have multiple observations for each worker, so cannot remove the effects of worker level unobservables, but we include a number of demographic controls: education, experience, job tenure, race, gender, and ethnicity. The most robust specification we estimate is:

$$(3)\ \mathrm{Log}(TW_i) = \beta_0 + \beta_1(EMERGING_i) + \beta_2\mathrm{Log}(CW_i) + \mu_i + \varphi_i + \delta_i + \varepsilon_i$$

where μ is a market fixed-effect, φ is an employer fixed-effect, and δ is a vector of demographic controls. This specification controls for the worker's current wage, the market in which they are engaging in job search, their employer, and most demographic variables. The key assumption is that conditional on the current wage, market or worker level differences that could otherwise influence target wage setting and are correlated with the technology choice variable are absorbed by the demographic controls and the combination of market, employer, and industry fixed effects. Under these assumptions, this equation estimates the partial effect of using new IT rather than established IT. We assess these assumptions when discussing our results in the next section.

Because we do not observe the job search outcome and cannot measure workers' expectations about future employers, we cannot control for the non-wage value of the current job and expectations about non-wage value at the next one. β_1 is biased upwards if workers systematically expect greater non-wage value at their current jobs than at future ones. This is not very plausible, and we provide evidence of this later in the paper using employer flow patterns and data on employer attributes. Employment transition patterns do not support the notion that workers using emerging IT are systematically moving to employers from which they derive less non-wage value.

V. EMPIRICAL ANALYSIS

V. A. Descriptive statistics and figures

Table 5 presents key statistics. The average wage of IT workers in our sample is about $74,000 and the average target wage is slightly less than $78,000, which is 5% higher than the mean current wage. The fact that the target wage is not far removed from the current wage suggests that the former is not an unrealistic expectation of the respondent's market worth. It also suggests that it would not take much to get a typical IT worker to move, which is consistent with the literature on IT mobility (Saxenian, 1996; Fallick, et al, 2006).

The average IT worker in our sample has about thirteen years of experience and 2.2 years of job tenure, and about 20% of our sample works with emerging IT systems.

Figure 1 plots the percentage increase in wages required to induce IT workers using different IT systems to leave their employers, and it plots illustrative programming languages on the chart. It suggests that as technologies age, workers place progressively less value on hands-on experience with these technologies. This may be because they become less valuable in the market, because they enjoy working with them less, or because channels other than hands-on experience arise through which workers can obtain these skills. Figure 1 suggests that languages that were newer as of 2007, such as Java Servlets, Struts, and Eclipse, were considered more interesting to work with than more mature technologies like Visual Basic. The former languages were dot-com era technologies that were just beginning to spread through the workforce.

V.B. IT use on the job and job search

In Table 6, we report the results of estimating equation (3). In all regressions, the independent variable measures characteristics of the IT systems workers use, and the dependent variable is the target wage. All regressions include current wages and education, gender, experience, job tenure, and race. In addition to IT use and current wages, the estimated coefficients for gender and education are shown in the table. The omitted variable for education is vocational degree and for gender it is female. We first test the hypothesis that IT workers are reluctant to leave employers when working with emerging IT systems. The estimates from column (1) indicate that emerging IT use is associated with a 2% larger premium to leave an employer (t=7.33). There is a higher target wage for males that rises with education levels.[12] This education effect may suggest that more on-the-job learning occurs at higher education levels, or that education is a proxy measure for job positions allowing for greater wage bargaining. The coefficient on emerging IT use remains of similar magnitude in column (2) after including MSA effects (t=10.5), so our results are unlikely to be explained by unobserved market-level factors, at least to the extent that the MSA can capture the relevant market.

---

[12] We replicate this test dropping workers for whom gender is missing, but dropping these workers does not have much of an impact on the magnitude of our key IT estimates.

Column (3) includes variables indicating to which of the four categories from Table 4 the IT systems in use belong, where the omitted category is low growth, older technologies (column 4, Table 4), so the other IT estimates should be interpreted relative to this base group. Of these, only emerging IT is positive and statistically significant (t=7.33). IT workers do not appear less willing to leave their employers when using any of the other IT groups, and these estimates remain similar after adding MSA effects in column (4).

Comparing the job seeker's current wage with their target wage is powerful because this comparison removes the effects of many unobserved worker-level variables, but it does not remove the effects of unobserved characteristics of the current employer that are correlated with the firm's IT choices and make it less desirable for workers to leave. Because we have data from multiple workers for many employers in our sample, we can introduce firm effects. This reduces our sample size because we only retain employers with at least two workers,[13] but the tables in Appendix B show that the sample characteristics of this reduced sample are similar to that of our full sample of workers. After including employer and MSA fixed-effects in column (5), the emerging IT coefficient remains of similar magnitude and statistically significant (t=2.86). With employer fixed-effects, the estimated coefficient on male gender changes sign, suggesting that the positive coefficient on male gender in the earlier analysis was reflecting heterogeneity across firms that was reflected in the gender composition of its workers.

The last three columns test the sensitivity of the estimates to alternate formulations of our IT use measure. The first column raises the threshold used to identify "high-growth" IT systems to those in the top quartile of growth for all IT systems, rather than the top half, and the second uses the standardized value of the growth rate, rather than placing the technologies into quartiles by growth rate. In both cases, the coefficient estimates remain significant and consistent in size with those reported in prior columns.[14] A final test uses a firm-level measure of emerging IT use constructed as the share of the firm's workers using these systems. The coefficient estimate on this measure rises slightly, to about 3.7% (t=7.40), which indicates that it

---

[13] In the reduced sample of employers that have multiple workers, 66% of workers are from employers with between two and five workers in our sample, 17% are from employers with between six and ten workers in our sample, 15% are from workers with between eleven and fifty workers in our sample, and 1% are from employers with more than fifty workers in our sample.

[14] These regression results are also similar when changing the definition of a new technology to being less than five years old or being less than fifteen years old.

is a reasonable proxy measure for the activities of workers at the firm. We use this fact later in our analysis for an analysis in which worker-level IT measures are unavailable.

V. C. Technology preferences, skills, and turnover

One important reason that workers may value working with newer technologies is because of the skills they acquire, and both theoretical and empirical work shows that workers tend to invest in acquiring new skills earlier in their careers (Ben-Porath 1967; Moen 2005).

Column (1) of Table 7 reports results from a model that is similar to the specification in Table 6, column (1) but that only shows the estimates on IT main effects and a dummy variable indicating workers with at least fifteen years of experience.[15] The IT estimates are similar to those reported in prior tables (t=7.00), and the experience estimates suggest that experienced workers require less to leave employers (t=-5.00), which is consistent with prior work on wage growth in the IT sector (Freedman 2008). These estimates are similar after including MSA fixed-effects in column (2). Column (3) adds an interaction term between IT and experience. The IT estimate is of the same magnitude as in prior regressions and is statistically significant (t=9.00), but is offset by the interaction term (t=-4.2). This indicates that it is primarily less experienced workers who accept lower wages to work with emerging IT. These effects change little after including employer effects in column (4). These relationships are illustrated in Figure 2. The fitted lines indicate that IT workers place greater value on using these systems early in their careers, and this value fades with experience. Workers do not place a premium on mature IT, regardless of their experience levels.

A human capital explanation suggests that younger workers take these jobs on their way to higher paying "destination" jobs, and field evidence suggests that high-tech labor markets are characterized by learning and rapid turnover (Saxenian, 1996). The regressions in columns (5) through (9) of Table 7 analyze the relationship between IT systems and the length of time workers stay on the job before exiting. For this analysis, we do not require wage data, so we use archival data on employment histories from workers'

---

[15] Moen (2005) uses twenty years as a threshold when analyzing a similar effect for R&D workers. In other regressions not shown, we have used lower thresholds, such as ten years, with substantively similar results. Thresholds higher than fifteen years limit our sample because older IT workers tend to disproportionately disappear from our sample.

resumes. Time until job exit is measured using each of the prior employment spells reported by workers in our sample, and IT is measured by text-mining the worker's resume for the technology used in each prior employment spell. All regressions control for demographics and we estimate hazard models due to the nature of the dependent variable.[16]

The coefficient estimates in column (5) indicate that at any point during an employment spell, IT workers using emerging technologies are 12% more likely to leave. A source of measurement error in column (5) is that we include the job in which the worker is employed at the time we observe their job search. These observations are right-censored because we observe these workers before their current job has concluded. In column (6) we drop the most recent employment spell but obtain similar estimates.

There may be heterogeneity across markets, due to differences in job hopping rates in high-tech clusters (Fallick, et al, 2006), but our estimates remain stable after adding MSA effects in column (7). After adding employer effects in column (8), the estimate drops slightly, but continues to indicate a greater exit hazard for workers using emerging IT systems. Column (9) also includes employer effects, but like column (6), drops the most recent employment spell. This reduces the sample size, because it limits the sample to prior employment spells for IT workers at firms that employ more than one person in our sample. Yet, with this restriction, the estimate remains positive and significant. When comparing IT workers within the same firm, the estimates indicate that emerging IT use increases exit likelihood by as much as 20%. Across columns, the estimates suggest that workers are 10-20% more likely to exit the firm at any given time if using emerging IT systems.

It is noteworthy that together, Tables 6 and 7 indicate that workers who use these IT systems i) seek a higher wage increase to leave their employers and ii) leave more quickly. These findings are inconsistent with the idea that our IT estimates are biased by other, unobserved job amenities that workers value at these

[16] Average employment spells for workers in each of the four technology classes in Table 4 are 3.4 years (Growth >= Median, Technology Age <= 10 years), 3.6 years (Growth < Median, Technology Age <= 10 years), 3.7 years (Growth >= Median, Technology Age > 10 years), and 3.8 years (Growth < Median, Technology Age > 10 years).

firms. It is difficult to reconcile how job amenities simultaneously make it costlier to remove a worker from their employer but also drive them to leave more quickly.

The results reported in the latter part of this analysis include fixed-effects for the labor market in which the worker is searching for a job, but labor markets can differ significantly in terms of a worker's ability to transfer new technical skills to other employers, so we also investigate how labor market characteristics influence these estimates. The results of these tests are reported and discussed in greater detail in Table 1 of Appendix C. They should be interpreted cautiously because labor markets can differ along a number of dimensions that cannot be easily measured, but can influence job search activity. Nevertheless, the results from these tests are consistent with the argument that the effects are strongest in markets where workers who use emerging IT have a larger number of outside options, in the form of employers who are hiring workers to use these new technologies. These tests provide further support for the argument that workers are reluctant to leave environments where they use newer IT systems because of factors, such as skills, that can be transferred across employers, rather than factors such as non-wage compensation that are not portable.

V.D. Evidence from employer reviews

The analyses reported in the previous section provide circumstantial evidence that patterns of technology use and mobility are consistent with the argument that one reason IT workers value working with new technologies is because they gain new skills. An alternative and more direct approach to exploring this relationship is to simply "ask" IT workers what they value in their employers. In this section, we analyze the text reviews from Glassdoor.com to determine the employer features that IT workers emphasize and how this corresponds with the technologies they use at that employer.

Dimensions of employer value. Using unsupervised methods on the Glassdoor review text, we generated eight dimensions along which IT workers reported valuing employers: *1) LEARNING, 2) TECHNOLOGY, 3) COWORKERS, 4) FLEXIBILITY, 5) COMPENSATION, 6) CAREER, 7) INDUSTRY,* and *8) LOCATION.* Details of the text-mining process and the workplace dimensions that were created can be found in Appendix A. The strongest correlations in the sample are between *LEARNING* and *TECHNOLOGY*, between *LEARNING* and *CAREER,* and between *TECHNOLOGY* and *CAREER*, which supports the argument that IT access is highly valued by workers for acquiring new skills that lead to future opportunities.[17]

In Table 8, we analyze correlations between a) the types of IT used by workers at various employers and b) the workplace dimensions that were mined from the reviews submitted by IT workers from those employers. In this table, each observation corresponds to a worker in our sample. The dependent variables are indicators of whether the worker uses a particular type of technology at work, and the right-hand side variables correspond to whether the worker's employer is above or below the median with respect to the sample in terms of the eight workplace dimensions listed in the paragraph above.

The regression estimates in column (1) suggest that emerging IT use by a worker is correlated with *LEARNING* (t=2.18) and *TECHNOLOGY* (t=3.32) being highly valued by at that worker's employer. Emerging IT use is negatively correlated with *COMPENSATION* (t=-3.35), which is consistent with the argument that workers accept lower pay to work with these systems. In column (2), we find that using high growth but mature technologies is positively correlated with *TECHNOLOGY* (t=3.34) but negatively correlated with *LEARNING* (t=-3.27). This is reasonable because employers expect workers to already have the necessary skills when hired, so it may be less important to provide an environment where workers can build human capital. Columns (3) and (4) analyze low growth technologies. *TECHNOLOGY* is negatively correlated with the use of these systems (t=-3.24), and workers emphasize *COMPENSATION* at these employers (t=3.02), which is consistent with requiring more pay to work with dying or less interesting technologies. Column (5) in this table reports results from a regression of these features against target wages, controlling for the current wage. The estimates are positive and significant for *LEARNING* (t=1.95), *TECHNOLOGY* (t=2.79), and *INDUSTRY* (t=5.08). The coefficient estimate on *LEARNING* indicates that a standard deviation higher value for this measure is associated with a 4% larger difference between the target and current wage, and the *TECHNOLOGY* variable is positive and significant at the same magnitude.

[17] Correlations among the full set of features are reported in Table 2 of Appendix C.

We also performed one additional set of tests intended to probe the robustness of the relationship between skills development and technology investment, the results of which are shown and discussed in detail in Table 3 of Appendix C. Specifically, we tested interactions between IT systems and whether the employer is perceived by reviewers to be a good workplace in which to learn new skills. The statistical power of this analysis is limited by the sample size, but it provides evidence that the measure of employer skills development is positively correlated with using emerging IT and negatively correlated with using mature IT systems, and that both employer skills development and the use of emerging IT must be present together for employees to set a higher target. We do not observe a similar boost in the target wage premium when either measure is high but the other is not, which suggests that sources of bias should act at the confluence of workplace skills development and the use of emerging IT while not being correlated with either factor alone. This narrows the list of omitted factors that can introduce bias into our key estimates.

In sum, the analyses discussed in this section suggest that the opportunity to learn valuable skills is an important component of the value workers derive from using new technology.

V.E. Robustness tests

Finally, we conducted a variety of robustness tests to assess if our results were sensitive to other types of non-wage compensation or to employment patterns that can influence how workers set job search parameters. Extensive use of bonuses, commissions, or equity-based compensation plans can increase a worker's compensation beyond what is reflected in their wages. Employers that use emerging IT systems may make greater use of stock-based pay, which respondents might not be able to easily convert to an annual compensation figure, in which case they may inflate their target wages to adjust. In this section, we analyze a subsample of workers for whom we obtained data on both wage and non-wage compensation. These data were only available for a sample of workers who searched for jobs in 2011. We do not have access to the resume data for this sample, so could not use the methods used earlier to generate worker-level IT measures. Instead, we use firm-level IT measures, which we showed in Table 6 to behave similarly to our individual IT measures.

Column (1) of Table 10 suggests that 2-3% higher pay is required to induce workers to leave their employers (t=3.00), and this estimate changes little in column (2) after including other forms of non-wage compensation (t=2.89). The estimates on the other compensation variables are positive and statistically significant, indicating that workers value these forms of compensation, but the IT estimate remains of similar magnitude to prior estimates. The estimate falls in magnitude but remains significant in column (3), when health and other benefits are added (t=2.44). There is a positive estimate on healthcare, but the coefficients on vacation and retirement benefits are negative, which may be due to selection; older workers, who tend to post lower target wages, may be more attracted to firms with retirement benefits. These estimates are similar in sign and magnitude in subsequent columns after adding MSA effects (t=2.80 in column (4), t=2.40 in column (5)). We cannot add employer effects because our IT variable for these tests is measured at the firm level. This table suggests that the estimated IT coefficient is unlikely to be significantly biased by non-wage compensation levels for workers using emerging IT systems.

Another source of concern is that high-tech markets have large numbers of skilled visa holders with restricted mobility. It is possible they report high targets to offset the legal challenges of changing employers and are more prevalent at companies using emerging IT. To test this explanation, we incorporate citizenship information. Employees in this database report whether they are US citizens. Citizens compose about 78% of the smaller 2011 sample. In column (6), including citizenship does not alter the IT estimates, so mobility restrictions faced by non-US citizens as a group do not appear to be a significant source of estimation bias (t=2.80). The estimates also remain similar in column (7) after limiting the sample to US citizens (t=2.45), who do not face visa-related mobility restrictions.

Finally, if workers who use emerging IT tend to target employers with undesirable work characteristics, it can bias our estimates upwards if they post higher target wages to offset undesirable characteristics. Alternatively, if they target jobs with more desirable characteristics, it will bias our estimates downward. The fact that the job board does not permit workers to set targets for non-wage characteristics should discourage job switchers who are looking for a significant change in work context, but some workers may be looking for such changes. We assessed the importance of this source of bias using the review data along with the resume data to observe employment transitions among firms for which we have Glassdoor data. These results, shown in Table 4 of Appendix C, suggest that i) workers tend to move between firms that are similar in their non-wage attributes and ii) these patterns are similar for workers whether or not they are using emerging IT systems. These patterns are inconsistent with the argument that our results are likely to be sensitive to a large pool of workers who use emerging IT systems and are seeking to move to jobs with undesirable non-wage characteristics.

VI. CONCLUSIONS

This paper provides evidence that IT workers prefer employers that invest in emerging IT systems, and this is in part because they can acquire valuable skills on the job. Although prior studies have examined how IT assets affect performance, we show that those investments can influence firms' competitiveness in the market for high-tech labor. In fact, IT investment may derive some of its performance value because workers are attracted to firms with a superior engineering brand. For a given wage, employers using new technologies can attract higher-productivity technical labor.

These effects are large enough to be interesting. For a firm in which engineers account for a large percent of the workforce, a 2% to 4% reduction in the wage bill is a powerful financial incentive. Results from the productivity literature suggest that the total share of IT investment for Fortune 500 firms is 3% to 5% of total expenditures. Therefore, a reduction of a few percentage points in the wage bill (which can often account for 70% to 80% of total spending) can offset a large fraction of a firm's IT costs. Of course, this does not imply that all firms should adopt emerging IT systems. These investments are associated with higher churn, which is itself costly, and firms differ in the costs they face when first adopting these systems. Our findings suggest only that ceteris paribus, firms with stronger engineering brands, in terms of technology and culture, can attract higher quality technical labor for a given wage.

Our findings have several implications. Competition for IT labor has been a subject of policy discussion for a generation. Firms' IT choices can be an important determinant of retention, and may be one reason why poaching in high-tech markets often requires higher wage premiums than other markets. Where workers feel that their skills are depreciating, they may be harder to retain than in environments where they learn new skills that will be valuable for their future careers. The argument that workers derive value from acquiring transferable skills has additional implications. Because being rewarded for these skills requires switching employers, workers may place greater value on IT experience in markets where mobility is unfettered by non-compete policies. This implies, in turn, that firms have greater incentives to adopt emerging IT in markets without mobility restrictions because they can capture more of the value they provide to workers through this channel. Finally, the heterogeneity of our findings by experience level suggests that IT factors can influence how different age workers sort across firms, which may be important given the many diversity concerns that high-tech firms face.

There are caveats to our study. It would be valuable to know whether the results differ for workers who are not actively seeking jobs. Our data are also from a short panel, so we do not take a position on whether these effects are idiosyncratic to the time period we study or whether they are more broadly generalized to IT-enabled change. A more encompassing test might analyze data from more than one IT diffusion wave. Indeed, in the long run, firms and technology vendors may change the nature of their actions in a way that changes these relationships.

There remains significant scope for work in this area. How firm characteristics, beyond IT investment, influence IT learning and IT wage structure merits further study. Moreover, how IT workers acquire skills, how these skills are rewarded, and how they depreciate all have implications for the IT labor force. Those implications include who chooses to be an IT worker and the wages they earn. These questions have received relatively little attention in the academic literature, but their answers are becoming important for workforce policy and for managing technology in a competitive labor market.

REFERENCES

Acemoglu, D., & Pischke, J. S. (1999). Beyond Becker: training in imperfect labour markets. *The Economic Journal*, *109*(453), 112-142.

Agarwal, Rajshree, and Atsushi Ohyama. 2013. "Industry or Academia, Basic or Applied? Career Choices and Earnings Trajectories of Scientists." *Management Science* 59 (4): 950–70.

Almeida, Paul, and Bruce Kogut. 1999. "Localization of Knowledge and the Mobility of Engineers in Regional Networks." *Management Science* 45 (7): 905–17.

Ang, Soon, Sandra Slaughter, and Kok Yee Ng. 2002. "Human Capital and Institutional Determinants of Information Technology Compensation: Modeling Multilevel and Cross-Level Interactions." *Management Science* 48 (11): 1,427–45.

Archak, N., Ghose, A., & Ipeirotis, P. G. (2011). Deriving the pricing power of product features by mining consumer reviews. *Management Science*, *57*(8), 1,485-1,509.

Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Invention." In The Rate and Direction of Inventive Activity: Economic and Social Factors, 609–26. Princeton University Press.

Aurand, Timothy W., L. Gorchels, T. R. Bishop, (2005) "Human resource management's role in internal branding: an opportunity for cross‐functional brand message synergy", *Journal of Product & Brand Management*, Vol. 14 Issue: 3, pp.163-169.

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

Bapna, Ravi, Nishtha Langer, Amit Mehra, Ram Gopal, and Alok Gupta. 2013. "Human Capital Investments and Employee Performance: An Analysis of IT Services Industry." *Management Science* 59 (3): 641–58.

Barley, Stephen R., and Gideon Kunda. 2006. Gurus, Hired Guns, and Warm Bodies: Itinerant Experts in a Knowledge Economy. Princeton University Press.

Bartel, Ann, Casey Ichniowski, and Kathryn Shaw. 2007. "How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills." *The Quarterly Journal of Economics* 122 (4): 1,721–58.

Becker, Gary S. 1964. "Human Capital Theory." Columbia, New York.

Ben-Porath, Yoram. 1967. "The Production of Human Capital and the Life Cycle of Earnings." *Journal of Political Economy* 75 (4, Part 1): 352–65.

Bharadwaj, Anandhi S. 2000. "A Resource-Based Perspective on Information Technology Capability and Firm Performance: An Empirical Investigation." *MIS Quarterly*, 169–96.

Bidwell, Matthew, and Forrest Briscoe. 2010. "The Dynamics of Interorganizational Careers." *Organization Science* 21 (5): 1,034–53.

Black, Sandra E., and Lisa M. Lynch. 2005. "Measuring Organizational Capital in the New Economy." In Measuring Capital in the New Economy, 205–36. University of Chicago Press. http://www.nber.org/chapters/c10622.pdf.

Blau, D. M. (1991). Search for nonwage job characteristics: a test of the reservation wage hypothesis. *Journal of Labor Economics*, *9*(2), 186-205.

Bloom, N., Garicano, L., Sadun, R., & Van Reenen, J. 2014. The distinct effects of information technology and communication technology on firm organization. *Management Science*, *60*(12), 2859-2885.

Branstetter, L. G., Drev, M., & Kwon, N. (2018). Get with the program: Software-driven innovation in traditional manufacturing. *Management Science*.

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.

Brynjolfsson, E., and LM. Hitt. 2000. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *The Journal of Economic Perspectives* 14 (4): 23–48.

Brynjolfsson, Erik, and Kristina McElheran. 2016. "Digitization and Innovation the Rapid Adoption of Data-Driven Decision-Making." *The American Economic Review* 106 (5): 133–39.

Cappelli, P., & Chauvin, K. (1991). An interplant test of the efficiency wage hypothesis. *The Quarterly Journal of Economics*, *106*(3), 769-787.

Chari, Varadarajan V., and Hugo Hopenhayn. 1991. “Vintage Human Capital, Growth, and the Diffusion of New Technology.” Journal of Political Economy, 1,142–65.

Falk, Armin, Ernst Fehr, and Christian Zehnder. 2006. “Fairness Perceptions and Reservation Wages: The Behavioral Effects of Minimum Wage Laws.” *The Quarterly Journal of Economics*, 1,347–81.

Fallick, Bruce, Charles A. Fleischman, and James B. Rebitzer. 2006. “Job-Hopping in Silicon Valley: Some Evidence Concerning the Microfoundations of a High-Technology Cluster.” *The Review of Economics and Statistics* 88 (3): 472–81.

Feldstein, Martin, and James Poterba. 1984. “Unemployment Insurance and Reservation Wages.” *Journal of Public Economics* 23 (1): 141–67.

Fichman, Robert G., and Chris F. Kemerer. 1997. “The Assimilation of Software Process Innovations: An Organizational Learning Perspective.” *Management Science* 43 (10): 1,345–63.

Fong Boh, Wai, Sandra A. Slaughter, and J. Alberto Espinosa. 2007. “Learning from Experience in Software Development: A Multilevel Analysis.” *Management Science* 53 (8): 1,315–31.

Forman, Chris. 2005. “The Corporate Digital Divide: Determinants of Internet Adoption.” *Management Science* 51 (4): 641–54. doi:10.1287/mnsc.1040.0343.

Franco, April Mitchell, and Darren Filson. 2006. “Spin-Outs: Knowledge Diffusion through Employee Mobility.” *The RAND Journal of Economics* 37 (4): 841–60.

Freedman, Matthew L. 2008. “Job Hopping, Earnings Dynamics, and Industrial Agglomeration in the Software Publishing Industry.” *Journal of Urban Economics* 64 (3): 590–600.

Ghose, A., P. Ipeirotis, and B. Li. "Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content." *Marketing Science* 31, no. 3 (2012): 493-520.

Hackman, J. Richard. & Oldham, Greg. R. (1980). Work redesign. (Addison-Wesley, Reading, MA).

Hall, Robert E., and Alan B. Krueger. 2012. “Evidence on the Incidence of Wage Posting, Wage Bargaining, and on-the-Job Search.” American Economic Journal: Macroeconomics, 56–67.

Hu, M., & Liu, B. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 168-177). ACM.

Jones, Stephen RG. 1988. “The Relationship between Unemployment Spells and Reservation Wages as a Test of Search Theory.” *The Quarterly Journal of Economics*, 741–65.

Joseph, Damien, Wai Fong Boh, Soon Ang, and Sandra Slaughter. 2012. “The Career Paths Less (or More) Traveled: A Sequence Analysis of IT Career Histories, Mobility Patterns, and Career Success.” *MIS Quarterly* 36 (2): 427–52.

Jovanovic, Boyan, and Yaw Nyarko. 1996. “Learning by Doing and the Choice of Technology.” *Econometrica* 64 (6): 1,299–1,310.

Kim, D. (2012). HR branding: How human resources can learn from product and service branding to improve attraction, selection, and retention [*Electronic article*]. *Cornell Hospitality Report, 12*(14), 6-17.

Krueger, Alan B., and Andreas I. Mueller. 2014. “A Contribution to the Empirics of Reservation Wages.” National Bureau of Economic Research.

Kuhn, Peter, and Mikal Skuterud. 2000. “Job Search Methods: Internet versus Traditional.” Monthly Lab. Rev. 123: 3.

Lancaster, Tony, and Andrew Chesher. 1983. “An Econometric Analysis of Reservation Wages.” Econometrica 51 (6): 1,661–76.

Levina, Natalia, and Mingdi Xin. 2007. “Research Note-Comparing IT Workers’ Compensation across Country Contexts: Demographic, Human Capital, and Institutional Factors.” *Information Systems Research* 18 (2): 193–210.

Marx, Matt. 2011. The Employer Strikes Back: Non-Compete Agreements and the Mobility of Technical Professionals. *American Sociological Review*. 76(5): 695-712.

Mata, Francisco J., William L. Fuerst, and Jay B. Barney. 1995. “Information Technology and Sustained Competitive Advantage: A Resource-Based Analysis.” *MIS Quarterly*, 487–505.

McElheran, Kristina. 2015. “Do Market Leaders Lead in Business Process Innovation? The Case(s) of E-Business Adoption.” *Management Science*, January.

Mithas, Sunil, and Mayuram S. Krishnan. 2008. “Human Capital and Institutional Effects in the Compensation of Information Technology Professionals in the United States.” *Management Science* 54 (3): 415–28.

Mithas, Sunil, and Henry C. Lucas Jr. 2010. “Are Foreign IT Workers Cheaper? US Visa Policies and Compensation of Information Technology Professionals.” *Management Science* 56 (5): 745–65.

Møen, Jarle. 2005. “Is Mobility of Technical Personnel a Source of R&D Spillovers?” *Journal of Labor Economics* 23 (1): 81–114.

O’Mahony, Siobhan, and Beth A. Bechky. 2006. “Stretchwork: Managing the Career Progression Paradox in External Labor Markets.” *Academy of Management Journal* 49 (5): 918–41.

Prasad, Eswar S. 2003. “What Determines the Reservation Wages of Unemployed Workers? New Evidence from German Micro Data.”

Reja, U., Manfreda, K. L., Hlebec, V., & Vehovar, V. 2003. “Open-ended vs. close-ended questions in web questionnaires.” *Developments in applied statistics*, *19*(1), 160-117.

Roach, Michael, and Henry Sauermann. 2010. “A Taste for Science? PhD Scientists’ Academic Orientation and Self-Selection into Research Careers in Industry.” *Research Policy* 39 (3): 422–34.

Roberts, Jeffrey A., Il-Horn Hann, and Sandra A. Slaughter. 2006. “Understanding the Motivations, Participation, and Performance of Open Source Software Developers: A Longitudinal Study of the Apache Projects.” *Management Science* 52 (7): 984–99.

Rosen, Sherwin. 1972. “Learning and Experience in the Labor Market.” *Journal of Human Resources*, 326–42.

Sauermann, Henry, and Michael Roach. 2014. “Not All Scientists Pay to Be Scientists: PhDs’ Preferences for Publishing in Industrial Employment.” *Research Policy* 43 (1): 32–47.

Saxenian, AnnaLee. 1996. Regional Advantage. Harvard University Press.

Shimer, Robert, and Iván Werning. 2007. “Reservation Wages and Unemployment Insurance.” *The Quarterly Journal of Economics* 122 (3): 1,145–85.

Stern, Scott. 2004. “Do Scientists Pay to Be Scientists?” *Management Science* 50 (6): 835–53.

Tambe, Prasanna, and Lorin M. Hitt. 2012. “The Productivity of Information Technology Investments: New Evidence from IT Labor Data.” *Information Systems Research* 23 (3-part-1): 599–617.

Zucker, Lynne G., Michael R. Darby, and Jeff Armstrong. 1998. “Geographically Localized Knowledge: Spillovers or Markets?” *Economic Inquiry* 36 (1): 65–86.

FIGURE 1: AGE OF TECHNOLOGY AND PREDICTED WAGE PREMIUM

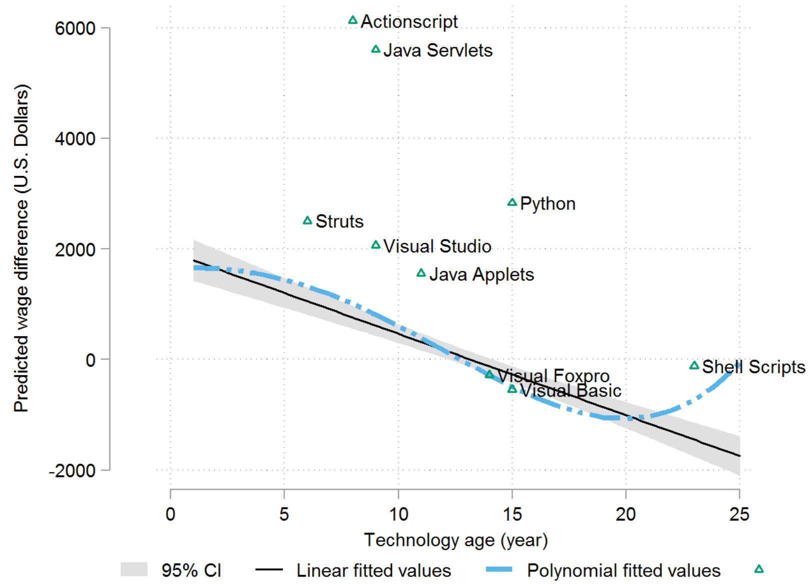

Figure notes: This figure plots a linear fit (the black solid line) and a nonparametric polynomial fit (the blue dashed line) of technology age against the predicted difference between target and current wages (in U.S. dollars). Green triangles denote examples of popular programming languages. The predicted difference is computed as follows: we run a regression using a specification similar to the one used in Table 6, Column (1) except that we drop the "Emerging" variable. We use the residual of this regression, which is interpreted as the percentage change in the current wage. We compute the equivalent dollar value of this residual by evaluating it at the average current wage.

FIGURE 2: JOB EXPERIENCE, IT USE ON THE JOB, AND PREDICTED WAGE PREMIUM

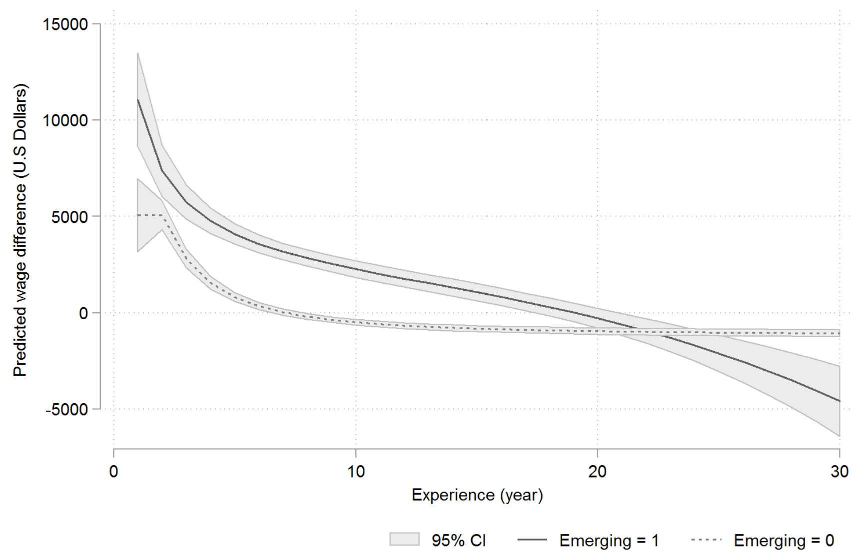

Figure notes: This figure plots fitted lines (nonparametric polynomial fit) of job experience (year) on the predicted difference between the target and current wage (in U.S. dollars) for emerging technologies and all other technologies. The predicted wage difference is computed by running a regression with a specification similar to the one used in Table 6, Column (1) except that we dropped the IT variables. We use the residual of this regression, which is interpreted as the percentage change in the current wage. We compute the equivalent dollar value of this residual by evaluating at the average current wage.

TABLE 1: DESCRIPTION OF KEY MEASURES AND DATA SOURCES

Measure	*Description*
CURRENT & TARGET WAGES	Our sample includes 50,628 annually employed IT workers who posted current and target wages on the job board and who identified their wages as being earned on annual basis. We drop hourly workers and those who report earning less than $10,000 per year as their current or target wage and those who report more than $1,000,000 per year as their current or target wage.
EMERGING IT	Our IT measures are derived by text-mining employee resumes to find which technologies the employee was working with at the employer they were with when participating on the job board. For workers that list using a number of different technologies, we use the newest one with which they report working.
EDUCATION	Education is self-reported by workers. Workers report having one of the following seven education levels: 1) High school and below, 2) vocational, 3) two-year degree, 4) four-year degree, 5) graduate degree, 6) doctorate, or 7) other.
RACE AND GENDER	Race and gender are self-reported by workers by choosing values from drop-down boxes. Reporting this information is voluntary.
JOB TENURE AND EXPERIENCE	These measures are computed from the job history (resume) of the employee. Job tenure is the time starting from when they began their most recent job up until the time at which the wage information was posted. Experience is computed beginning from the time they started their first job to when they posted wage information.
YEAR	The year in which the job seeker engaged in job search.
LOCATION	Workers report city, zip code, and state. Using these data, we assigned each worker to a metropolitan statistical area (MSA).
CULTURE	Measures of employer features most valued by their IT workers were constructed by text-mining reviews posted at the Glassdoor.com website, an online labor market intermediary upon which employees review employers. Details of this process are provided in Appendix A.

TABLE 2: DEMOGRAPHIC COMPARISON BETWEEN OUR SAMPLE AND EMPLOYED, NON-HOURLY IT WORKERS IN THE 2007 CURRENT POPULATION SURVEY (CPS)

Year: 2007	**(1) CPS Sample**	**(2) Job Board Sample**
Number of observations	**1,320**	**50,628**
A. By Gender (%)		
Female	29.0	11.1
Male	71.0	47.0
Not Specified	0.0	41.9
B. By Educational Degree (%)		
High School & Below	8.0	2.4
Vocational School	4.7	2.3
Two Year Degree	5.4	11.8
Four Year Degree	45.5	44.7
Graduate Degree	19.1	22.9
Doctorate Degree	1.9	0.9
Other	15.5	15.0
C. By Ethnicity (%)		
African American/Black	6.3	7.2
Asian	13.0	9.5
White	78.9	33.4
Other	1.8	5.7
Not specified	0.0	44.2

Table notes: This table compares a sample of IT workers from the Current Population Survey (CPS), Labor Force Statistics (Column (1)) with our job board sample (Column (2)). We keep workers from the CPS who report a job in a "Computer and Mathematical Science Occupation", their employment status as "Employed", and their hourly status as "Non-Hourly". Comparing column (1) to column (2), the Pearson chi-square statistic for gender is $\chi^2(1) = 77.8$ (p=0.00), for education is $\chi^2(5) = 257.9$ (p=0.00), and for ethnicity is $\chi^2(2) = 114.4$ (p=0.00). Workers with unspecified demographic information are dropped when conducting these chi-square tests.

TABLE 3: COMPARISONS OF IT WAGES WITH WAGES REPORTED FOR IT WORKERS IN THE BUREAU OF LABOR STATISTICS 2007 OCCUPATIONAL EMPLOYMENT SURVEY (OES)

Year: 2007	**(1) 2007 OES Wages**	**(2) Job board Wages**
Number of observations	**3,191,360**	**50,628**
A. By Detailed Occupation (Computer & Mathematical Occupations Only)		
Computer Programmers	72,010	61,726
Software Engineers, Applications	85,660	81,310
Software Engineers, Systems Software	90,780	92,545
Computer Systems Analysts	75,890	67,391
Database Administrators	70,260	74,175
Network and Computer Systems Administrators	67,850	56,878
Computer Specialists, All Other	72,310	74,462
B. By Metropolitan Area (Computer & Mathematical Occupations Only)		
Los Angeles-Long Beach-Glendale, CA	74,510	83,349
San Diego-Carlsbad-San Marcos, CA	73,310	75,698
San Francisco-Oakland-Fremont, CA	87,930	89,285
San Jose-Sunnyvale-Santa Clara, CA	98,160	95,443
Bridgeport-Stamford-Norwalk, CT	83,480	70,937
Washington-Arlington-Alexandria, DC-VA-MD-WV	85,060	82,676
Trenton-Ewing, NJ	76,320	77,000
New York-Northern New Jersey-Long Island, NY-NJ-PA	82,540	78,791
Seattle-Bellevue-Everett, WA	82,120	76,840

Table notes: This table reports wage comparisons between IT workers (SOC major group 15-0000, Computer and Mathematical Occupations) in the 2007 BLS Occupational Employment Statistics (OES) and IT workers in our sample. All figures are expressed in dollars, and are rounded to the nearest dollar. Column (1) reports 2007 averages from the OES; Column (2) reports wages for employed annual IT workers from our sample. Panel A classifies workers by ONET occupation group, identified by manually grouping workers in our sample into ONET categories using their standardized job titles. Panel B shows comparisons between the two data sets for some high-tech regions.

TABLE 4: EXAMPLES OF IT SYSTEMS CLASSIFIED BY AGE AND GROWTH RATE

Growth rate	>= Median	>=Median	< Median	< Median
Technology age	<=10 years	>10 years	<=10 years	>10 years
Examples	Actionscript Agile Development Citrix Metaframe Oracle Clinical Oracle Fixed Assets Paint Shop Pro Sap Basis Struts Unity Google Analytics	Autodesk Citrix Winframe Datastage Informatica Ingenix Java Applets Peopletools Sound Forge Sun Servers Teradata	Avaya Java Servlets Lotus Organizer Macromedia Dreamweaver Macromedia Fireworks Mascot Microsoft Photo Editor Oracle Project Tomcat	Adobe Photoshop Cisco Routers Crystal Reports Delphi Informix Python Samba Shell Scripting Siebel Sybase
Number of workers in sample using each IT type	8,089	5,912	3,920	33,427

Table notes: This table provides examples of information technologies classified by their age and their growth rate. Age is computed in 2007, when the job search data were collected. Growth rates are computed as described in Appendix A. Technologies in the first column are illustrative of a larger set we refer to in our analysis as "emerging" information technologies.

TABLE 5: STATISTICS FOR KEY REGRESSION VARIABLES FOR IT WORKERS IN SAMPLE

Variable	Mean	Std. Dev.	N
Current wage	73,901	34,370	50,628
Target wage	77,775	36,062	50,628
Experience (Years)	13.48	7.14	50,628
Job tenure (Years)	2.23	3.87	50,628
Emerging IT (0/1)	0.22	0.42	50,628

Table notes: This table reports descriptive statistics for key variables used in our regression analysis. Wage statistics are rounded to the nearest dollar.

TABLE 6: USE OF EMERGING IT ON THE JOB AND WAGES

DV: Log(Target)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IT growth measure	>=Median	>=Median	>=Median	>=Median	>=Median	>=25%	Growth	%Emerg
Emerging	0.022***	0.021***	0.022***	0.021***	0.020***	0.020***	0.019***	0.037***
	(0.003)	(0.002)	(0.003)	(0.003)	(0.007)	(0.007)	(0.004)	(0.005)
New=0, Growth=1			0.002	0.001				
			(0.003)	(0.003)				
New=1, Growth=0			-0.000	0.000				
			(0.003)	(0.003)				
Male	0.024***	0.023***	0.024***	0.023***	-0.025***	-0.025***	-0.024***	0.023***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.008)	(0.008)	(0.008)	(0.006)
Gender unreported	0.005	0.006	0.005	0.006	-0.005	-0.005	-0.002	-0.007
	(0.007)	(0.006)	(0.007)	(0.006)	(0.016)	(0.016)	(0.016)	(0.014)
Less than high school	-0.017***	-0.017***	-0.017***	-0.017***	-0.028**	-0.028**	-0.026*	-0.008
	(0.006)	(0.006)	(0.006)	(0.006)	(0.014)	(0.014)	(0.014)	(0.013)
2-year degree	-0.011***	-0.012***	-0.011***	-0.012***	0.010	0.010	0.011	-0.000
	(0.003)	(0.004)	(0.003)	(0.004)	(0.009)	(0.009)	(0.009)	(0.007)
4-Year degree	0.039***	0.037***	0.039***	0.037***	0.033***	0.033***	0.033***	0.035***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.007)	(0.007)	(0.007)	(0.006)
Graduate degree	0.079***	0.075***	0.079***	0.075***	0.071***	0.071***	0.072***	0.080***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.008)	(0.008)	(0.008)	(0.006)
Log (Wage)	0.763***	0.761***	0.763***	0.761***	0.843***	0.843***	0.842***	0.775***
	(0.005)	(0.002)	(0.005)	(0.002)	(0.006)	(0.006)	(0.006)	(0.005)
Employer FE					✓	✓	✓	
Location FE		✓		✓	✓	✓	✓	✓
Robust S.E.	✓		✓					
Demographic controls	✓	✓	✓	✓	✓	✓	✓	✓
R-squared	0.736	0.734	0.736	0.734	0.787	0.787	0.787	0.760
Observations	50,628	50,628	50,628	50,628	11,070	11,070	11,070	11,070

Table notes: This table reports regression estimates of emerging IT use on workers' target wages, controlling for their wages, job tenure, experience, race, gender, and education. The base group for gender is Female, and for educational attainment is vocational degree. Columns (1) through (5) use whether the IT growth rate is greater than or equal to median value as a measure of technology growth to construct the emerging IT variable. Column (6) uses whether the growth rate is in the top quartile as the IT growth measure. Column (7) uses the standardized growth rate. Column (8) uses the fraction of IT workers reporting emerging skills to construct an employer level measure of emerging IT use. The drop in sample size in columns (5) through (8) is due to adding employer fixed-effects. Appendix B compares the sample used in (1)-(4) with the sample in (5)-(9). Standard errors are clustered at the geographic level in Columns (1) and (3); *** p<0.01, ** p<0.05, * p<0.1.

TABLE 7: USE OF EMERGING IT AND CAREER DYNAMICS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
DV	**Log(Target wage)**				**Years until exiting job**				
Model	**OLS**	**FE**	**FE**	**FE**	**Hazard**	**Hazard**	**Hazard**	**Hazard**	**Hazard**
Employment spell sample	**All**	**All**	**All**	**All**	**Current & prior jobs**	**Prior jobs**	**Current & prior jobs**	**Current & prior jobs**	**Prior jobs**
Emerging	0.021***	0.020***	0.027***	0.036***	1.121***	1.124***	1.151***	1.091**	1.210***
	(0.003)	(0.003)	(0.003)	(0.009)	(0.012)	(0.016)	(0.014)	(0.048)	(0.067)
Exp>=15yrs	-0.010***	-0.009***	-0.006***	-0.017***					
	(0.002)	(0.002)	(0.002)	(0.005)					
Emerging * Exp>=15yrs			-0.021***	-0.041***					
			(0.005)	(0.015)					
Log (Wage)	0.770***	0.768***	0.768***	0.847***	0.955***	0.979	0.922***	0.902***	1.001**
	(0.006)	(0.002)	(0.002)	(0.007)	(0.009)	(0.014)	(0.011)	(0.034)	(0.057)
Employer FE				✓				✓	✓
Location FE		✓	✓	✓			✓		
Robust S.E.	✓				✓	✓			
Demographic controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
R-squared	0.742	0.740	0.741	0.685	0.009	0.008	0.006	0.012	0.012
Observations	47,712	47,712	47,712	10,423	47,712	22,713	47,712	10,423	5,456

Table notes: Columns (1) through (4) report regression estimates of how the relationship between emerging IT use and wages varies by workers' experience levels, and columns (5) through (9) report regressions of how emerging IT use is associated with years until the worker either exits their job in the case of prior employment spells, or searches for a new job for the case of the current employment spell. We only retain in our sample workers with less than 40 years of experience and reported job tenure of less than 30 years which slightly lowers the size of the sample from that used in Table 6. The results of hazard models are reported in Columns (5)-(9). Standard errors are clustered at the geographic level in Columns (1), (5), and (6); *** p<0.01, ** p<0.05, * p<0.1.

TABLE 8: CORRELATIONS BETWEEN EMPLOYER ATTRIBUTES, EMPLOYEES' IT USE ON-THE-JOB, AND JOB SEARCH

	(1)	(2)	(3)	(4)	(5)
DV:	**Use emerging IT?**	**Use older, High growth IT?**	**Use newer, low growth IT?**	**Use older, low growth IT?**	**Log(Target)**
Workplace attribute	**Logit**	**Logit**	**Logit**	**Logit**	**OLS**
LEARNING	0.124**	-0.232***	-0.007	-0.021	0.037**
	(0.057)	(0.071)	(0.080)	(0.046)	(0.019)
TECHNOLOGY	0.166***	0.177***	0.030	-0.146***	0.039***
	(0.050)	(0.053)	(0.102)	(0.045)	(0.014)
COWORKERS	0.055	-0.112	0.387***	-0.092**	-0.032**
	(0.049)	(0.074)	(0.072)	(0.044)	(0.015)
FLEXIBILITY	-0.073	0.319***	0.018	-0.091**	0.004
	(0.055)	(0.055)	(0.083)	(0.042)	(0.014)
COMPENSATION	-0.184***	0.017	-0.147	0.136***	-0.017
	(0.055)	(0.071)	(0.094)	(0.045)	(0.011)
CAREER	0.047	-0.060	0.096	-0.003	0.015
	(0.051)	(0.065)	(0.082)	(0.052)	(0.020)
INDUSTRY	-0.145**	0.051	0.094	0.036	0.061***
	(0.057)	(0.066)	(0.090)	(0.046)	(0.012)
LOCATION	-0.028	0.024	0.085	-0.012	0.002
	(0.060)	(0.068)	(0.081)	(0.049)	(0.018)
Log(Wage)					0.781***
					(0.018)
Robust S.E.	✓	✓	✓	✓	✓
Demographic controls	✓	✓	✓	✓	✓
R-squared	N/A	N/A	N/A	N/A	0.731
Observations	2,414	2,414	2,414	2,414	2,414

Table notes: This table reports regression estimates of employer features mined from the Glassdoor.com reviews on the worker level IT use and wage measures. Each observation is a worker and each row is an attribute value of that worker's employer. Details about how we produce employer attribute measures are provided in Appendix A. The dependent variable is a binary variable that indicates whether or not the worker uses that type of technology in their job. All regressions control for wages, job tenure, experience, race, gender, and education. Column (1)-(4) estimates correlations of these employer attributes on binary indicators of the technologies that workers use on the job. Column (5) embeds these employer attributes in the target wage specification used in the earlier tables. Robust standard errors are shown in parentheses; Standard errors are clustered at the geographic level; *** p<0.01, ** p<0.05, * p<0.1.

TABLE 9: OTHER COMPENSATION AND BENEFITS AND VISA RESTRICTIONS

DV: Log (Target)	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**
Model	**OLS**	**OLS**	**OLS**	**FE**	**FE**	**FE**	**FE**
Worker sample	**All**	**All**	**All**	**All**	**All**	**All**	**Citizens**
Data year	**2011**	**2011**	**2011**	**2011**	**2011**	**2011**	**2011**
% Emerging	0.027^{***}	0.026^{***}	0.022^{**}	0.028^{***}	0.024^{**}	0.028^{***}	0.027^{**}
	(0.009)	(0.009)	(0.009)	(0.010)	(0.010)	(0.010)	(0.011)
Log (Other comp)		0.010^{***}	0.010^{***}	0.009^{***}	0.009^{***}		
		(0.001)	(0.001)	(0.002)	(0.002)		
Log (Commission)		0.006^{***}	0.006^{***}	0.006^{***}	0.006^{***}		
		(0.001)	(0.001)	(0.001)	(0.001)		
Log (Bonus)		0.003^{***}	0.003^{***}	0.003^{***}	0.003^{***}		
		(0.001)	(0.001)	(0.001)	(0.001)		
DENTAL (Y/N)			-0.013		-0.012		
			(0.016)		(0.014)		
HEALTH (Y/N)			0.060^{***}		0.068^{***}		
			(0.017)		(0.015)		
VACATION (Y/N)			-0.015		-0.026^{**}		
			(0.013)		(0.012)		
RETIREMENT (Y/N)			-0.025^{***}		-0.021^{**}		
			(0.009)		(0.010)		
Non-US Citizen (Y/N)						0.003	
						(0.006)	
Log(Wage)	0.810^{***}	0.802^{***}	0.802^{***}	0.791^{***}	0.792^{***}	0.804^{***}	0.816^{***}
	(0.010)	(0.011)	(0.011)	(0.007)	(0.007)	(0.006)	(0.007)
Location FE				✓	✓	✓	✓
Robust S.E.	✓	✓	✓				
Demographic controls	✓	✓	✓	✓	✓	✓	✓
R-squared	0.778	0.781	0.782	0.768	0.769	0.760	0.804
Observations	6,843	6,843	6,843	6,843	6,843	6,843	5,310

Table notes: This table adds measures of equity compensation, benefits, and citizenship status for a subsample of workers who report this information on the job board. % Emerging is the fraction of workers at the firm who are using emerging IT. This firm level measure, generated using 2007 data, is matched by employer name to a sample of workers from 2011 who reported their compensation information. Job seekers report the value of alternative compensation measures in US dollars. DENTAL, HEALTH, VACATION, RETIREMENT are dummy variables and indicate whether the worker receives these benefits from their employer. Non-US citizen is also a binary indicator. In addition to the variables shown, all regressions control for wages, job tenure, experience, race, gender, and education. Robust standard errors are shown in parentheses; Standard errors are clustered at the geographical level in Columns (1), (2) and (3); *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

APPENDIX A: GENERATION OF TEXT BASED MEASURES OF IT USE AND EMPLOYER CHARACTERISTICS

I. IT MEASURES FROM RESUMES

To develop measures of technologies used on the job, we text mined the resumes that workers submitted to the job board. In their resumes, job seekers provide details about prior employment including information related to programming languages or software packages used on the job. Table A1 provides an illustrative example.

TABLE A1. SAMPLE RESUME WITH PROJECT DESCRIPTIONS

EXPERIENCE
E. xxx Corporation Long Island City. N.Y
11/2005-31/2006 Help Desk Support/Analyst
Tested Proprietary Database System
Gathered User Requirement
Assisted System Design(Design Patterns)
Coded in Java, C#
Sxxx Magazine New York, N.Y
1/2003-6/2003 Web Designer
Designed Front End/ Back End
Assisted System Analyst
Headed 4 Team members in the Project

We wrote software scripts that extracted the following fields:

1) The technologies mentioned by job seekers
2) The beginning and end years of the jobs in which they used the technology

For each year t and skill s, we computed the fraction of IT workers starting a job in that year who reported using that skill, (i.e. the ratio between 1) n_{ts}: IT workers starting a job and reporting the skill and 2) n_t: IT workers starting a job):

$share(s)_t = \frac{n_{st}}{n_t}$.

We measure skill growth as changes in this skill share:

$$Growth(s) = Average_{t \in T}(Share(s)_t - Share(s)_{t-1})$$

We then compute whether that skill has a growth rate that is greater than or equal to the median value for all skills in that year. When workers list multiple skills in a job, we select the newest skill used by the worker in that job.

Measurement error: Because our measures are based on changes in skill counts, whether workers accurately report skills and whether the skills possessed by a worker influence her decision to post a resume can both affect measurement error. Because reporting high-growth skills may increase the likelihood of future employment, workers may over-report these skills. On the other hand, workers using emerging technologies may have less incentive to use job boards if the market is tighter for these skills. In this case, we may underestimate the growth of higher growth skills and overestimate the growth of low growth skills. When we discretize this measure, skills growing at the median rate are most vulnerable to being misclassified, which means that we may identify some new but slower growth skills as emerging skills. This source of error can lead to underestimates of the IT coefficient because it associates a low difference between a target wage and current wage to a skill mislabeled as emerging. For robustness, we used two alternative measures: 1) whether the growth is in the top 25 percentile and 2) the standardized growth value.

II. GLASSDOOR MEASURES OF EMPLOYER ATTRIBUTES

We used the Glassdoor.com data to quantify the value workers assign to employer features by text-mining the online reviews. We focused on the content in the "pros" section of the reviews, which limits the text to those attributes that employees find valuable. We used 1) automated and 2) semi-automated approaches to create different measures of employer features that we used for two separate analyses. The key distinction between these approaches was whether the set of employer features was pre-specified before starting the text mining task.

1. Automated generation of employer features: First, we used a procedure that does not rely on pre-specified employer features. This approach replicates the procedure used to mine product features in Archak et al (2011) and Ghose et al (2012). First, we assigned part-of-speech tags to keywords, "chunked"[18] sentences based on grammar rules, and identified noun phrases as candidates for employer keywords. These were implemented using the Python NLTK package. We dropped infrequent noun phrases by (1) using the algorithm from Liu and Hu (2004) and (2) removing adjectives that did not introduce new concepts. Our final feature candidates included approximately 3,000 unique lemmatized nouns or noun phrases.

Next, we clustered noun phrases based on their context, where context is defined as the five-word window preceding and succeeding each feature. The intuition behind this approach is to cluster phrases that are conceptually similar to one another using association patterns that appeared around these phrases. Each noun phrase context was treated as a separate document. K-means clustering was used to group the noun-phrases. We varied the number of clusters from five to fifteen, and finally settled on eight clusters because each feature set represented a unique employer attribute and the correlations between clusters suggest that they are reasonably orthogonal to one another.

For each cluster, we generated a list of the most frequent phrases and phrases with the top TF-IDF scores (TF-IDF in the feature context documents)[19] and assigned labels to them. In Table B.1, we list the clusters and the top TF-IDF phrases for each cluster. The labels for each topic were assigned using our own judgment, which is standard for many unsupervised machine learning tasks.

Finally, to construct employer values for each cluster measure, we computed whether the fraction of workers referencing each of these measures is greater than or equal to the median value when compared with all organizations in the sample. This transformation absorbs some of the differences in employer size. Moreover, larger companies receive more reviews, but systematic error is introduced into our measures only if workers from larger companies disproportionally favor one feature over another.

2. Manually generated labeling of the SKILLS measure: For another analysis, we manually pre-specified phrases referring to "skill development/training" (SKILL).[20] We analyzed the 2,000 most frequent keywords in the Pros section of the reviews contributed by IT workers. For each keyword, we identified phrases that referred to skills and learning at work. Then, we limited the list to bigrams that appeared at least thirty times in the corpus. In Table B.2, we provide a list of illustrative bigram terms. We coded each review as referring to SKILL if at least one of these bigrams appears in the review. For each employer, we computed the fraction of IT workers referencing SKILL terms in their reviews.

---

[18] Phrase chunking is a natural language process that separates and segments a sentence into its sub constituents, such as noun, verb, and prepositional phrases.

[19] TF-IDF is a value that increases proportionally to the number of times a word appears in a document, but is offset by the frequency of the word in the corpus, which helps to adjust for the fact that some words appear more frequently in general.

[20] This approach is similar to an "Axial coding" paradigm, in which researchers use pre-defined constructs to guide how they analyze interview questions. Because of the size of the review text, we are not able to code each review. We therefore deconstruct the reviews into the smallest meaningful unit of text.

TABLE A2. CLUSTERS OF EMPLOYER FEATURES

LEARNING values, helpful colleague, learn environment, high turnover, tech support, flexibility in work schedule, high salary, pros, learn experience, summer hour, exciting, high tech, family, much stress, technical work, coworker, work schedule, near future, free train, positive
TECHNOLOGY concept, industry leader, interesting product, various project, date technology, supply, project, advance technology, different technology, engineering, software, many location, aerospace industry, automotive industry, agile development, project work, database, employee development, various technology, innovative product
FLEXIBILITY week vacation, hour shift, comp time, flexible work, remote, hours, hour day, short period of time, work time, free time, flexible work time, generous vacation, ample time, flexible work hour, vacation day, time management, amount of time, time-off policy, work hour, other week
COWORKER group of people, team of people, decent people, sense of humor, many people, knowledgeable people, diverse group of people, creative people, netapp, trench, experience people, kind people, bright, work environment, awesome people, helpful people, competent people, quality people, smart people, cool people
COMPENSATION bonus structure, ok salary, meal, medical, free access, solid benefit, federal government, balance between work, care of employee, fitness, excellent benefits, region, pet, package, company car, competitive compensation, life/work balance, company vehicle, stock purchase plan, work life balance
CAREER different opportunity, fast growth, opportunities, job opportunity, growth potential, advancement opportunity, career development, same company, opportunity for career advancement, growth opportunity, company with lot, mobility, promotions, company growth, own career, own career path, educational opportunity, room for growth, carrier, opportunity for career growth
LOCATION drama, cube, perfect place, beautiful location, beautiful place, san diego, easy place, san francisco, beach, nyc, amenity, weather, valley, small office, stable place, private office, proximity, street, other place, seattle
INDUSTRY everything, military, animal, professionalism, red tape, proud, let, paperwork, workload, competitor, face, vast majority, style, age, thats, show, income, recession, personnel, alot
Table notes: We show the top keywords in this table for each employer feature. These keywords were selected based on their TF-IDF score.

TABLE A3. BIGRAMS USED TO CONSTRUCT SKILLS MEASURE

learn lot, cutting edge, great training, great experience, opportunity learn, good training, place learn, training program, learn new, learning experience, good experience, opportunities learn, gain experience, great learning, work experience, training opportunities, learning opportunities, learn grow, experience working, get experience, edge technology, excellent training, new skills, training good, new technologies, training programs, training great, good learning, experience great, new technology, hands experience, job training, skill set, learning environment, experience good, work learn, lot experience, great technology, people learn, learning new, get learn, lots training, sales training, paid training, able learn, leading edge, want learn, learn something
Table notes: This table lists the bigrams used to create the SKILL measure used in our empirical analysis.

APPENDIX B: COMPARISONS OF SAMPLES USED IN ANALYSIS

TABLE B1: DESCRIPTIVE STATISTICS FOR SAMPLES USED IN ANALYSIS

	(1)	(2)	(3)
Worker sample	**Full regression sample**	**Employer FE sample**	**Glassdoor sample**
N	**50,628**	**11,070**	**2,414**
Current wage	73,901	72,292	72,843
Target wage	77,775	75,802	76,618
Experience	13.48	13.17	13.63
Job tenure (Years)	2.23	2.22	2.35
Emerging	0.22	0.22	0.20

Table notes: This table compares descriptive statistics of key variables for (1) full regression sample, (2) observations included in employer fixed effects regressions and (3) observations included in HRM practices regressions. Wage statistics are rounded to the nearest dollar.

TABLE B2: DEMOGRAPHIC VARIABLES FOR SAMPLES USED IN ANALYSIS

	(1)	(2)	(3)	(4)	(5)	(6)
Worker sample	**Full regression sample**		**Sample used for Employer FE**		**Glassdoor sample**	
Number of observations	**50,628**		**11,070**		**2,414**	
	N	%	N	%	N	%
A. By Gender						
Female	5,623	11.1	1,245	11.3	250	10.4
Male	23,816	47.0	5,010	45.3	1,086	45.0
Gender not reported	21,189	41.9	4,815	43.5	1,078	44.7
B. By Educational Degree						
Other	7,582	15.0	1,392	12.6	276	11.4
High School & Below	1,217	2.3	268	2.4	38	1.6
Vocational School	1,163	2.4	128	1.2	17	0.7
Two-Year Degree	5,963	11.8	1,388	12.5	310	12.8
Four-Year Degree	22,643	44.7	5,191	46.9	1,124	46.6
Graduate Degree	11,591	22.9	2,662	24.1	649	26.9
Doctorate Degree	469	0.9	41	0.4	0	0.0
C. By Ethnicity						
African American/Black	3,665	7.2	915	8.3	194	8.0
Asian	4,830	9.5	964	8.7	220	9.1
White	25,247	33.4	3,596	32.5	777	32.2
Other	16,886	49.9	5,595	50.5	1,223	50.7

Table notes: This table reports sample comparisons of gender, educational degree and ethnicity groups for (1) full regression sample, (2) observations included in employer fixed effects regressions and (3) observations included in HRM practices regressions.

APPENDIX C: SUPPLEMENTARY TESTS

TABLE C.1: LABOR MARKET FACTORS, EMERGING TECHNOLOGIES, AND WAGES

Table notes. This table analyzes how labor market factors that influence a worker's outside options can affect the worker's willingness to exchange wages for the opportunity to use more interesting IT. These cross-market comparisons should be interpreted with some caution, because the extent to which a worker and employer split the costs of various types of "training" (and therefore the sizes of the effects we observe) are, in part, a function of tightness in the labor market in which the worker is participating (e.g. Acemoglu and Pischke 1999). In this table, we test whether the patterns we observe are consistent with the idea that workers are more likely to exchange wages for skills in markets with better outside options, with the caveat that we cannot directly account for differences in supply and demand conditions across markets that might influence these estimates. The table reports worker-level regression estimates that analyze correlations between the workers' IT use, characteristics of the labor markets in which they are searching, and the difference between the target and current wage. In addition to the variables shown, all regressions control for wages, job tenure, experience, race, gender, and education. Robust standard errors are shown in parentheses; Standard errors are clustered at the geographic level in columns (1), (3), and (5); *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

To characterize labor markets, we measured a) the fraction of local employers hiring workers with emerging IT skills and b) the ratio of the number of employers on the hiring side of the market to the number of workers with these skills. We use these two measures in different sets of regressions. In each, markets are classified according to whether they are above the median value for these measures. We created a 2 x 2 matrix between this market measure and whether or not the focal worker uses emerging IT.

Column (1) uses measures of the fraction of employers in a market hiring workers with emerging IT skills. The effects of workers' emerging IT use on target wages are statistically significant only in markets where more than the median fraction of employers hire for these skills (t=5.75). The effects in smaller markets or where workers are using mature IT systems, are insignificant. Column (2) includes fixed-effects, which compares IT workers who have the same employer and use the same category of IT system, but who are located in different markets. The interaction effect when workers have emerging IT skills and employees are hiring these skills remains positive (t=3.67), but the estimate on emerging IT in small markets turns negative. This may indicate that for smaller market establishments, wages are bid up due to a constrained supply of workers. Columns (3) and (4) use our alternative measure, the ratio of the number of employers using emerging IT systems (i.e. employers with at least one worker using the system) to the number of IT workers using emerging IT systems, which captures relative differences in the worker's bargaining power. The pattern of estimates is similar to those we observed in columns (1) and (2).

DV: Log(Target wage)	**(1)**	**(2)**	**(3)**	**(4)**
Model	**OLS**	**FE**	**OLS**	**FE**
Emerging & Above median fraction of employers use emerging IT	0.023***	0.033***		
	(0.004)	(0.009)		
Emerging & Below median fraction of employers use emerging IT	0.011	-0.092***		
	(0.007)	(0.015)		
Not emerging & Above median fraction of employers use emerging IT	-0.000	-0.003		
	(0.003)	(0.006)		
Emerging & Above median ratio of emerging IT employers to workers			0.020***	0.038***
			(0.004)	(0.009)
Emerging & Below median ratio of emerging IT employers to workers			0.016***	-0.075***
			(0.006)	(0.014)
Not emerging & Above median ratio of emerging IT employers to workers			-0.004	0.001
			(0.002)	(0.006)
Log(Wage)	0.763***	0.763***	0.763***	0.761***
	(0.005)	(0.007)	(0.005)	(0.007)
Employer FE		✓		✓
Robust SE	✓		✓	
Demographic controls	✓	✓	✓	✓
R-squared	0.736	0.686	0.736	0.685
Observations	50,628	11,070	50,628	11,070

TABLE C.2: PAIRWISE CORRELATIONS AMONG EMPLOYER CULTURE ATTRIBUTES

Table notes: This table reports the full set of pairwise correlations between the key workplace dimensions constructed from the Glassdoor review text. Correlations are reported at the employer level for 473 employers that appear both in our sample of workers and in the Glassdoor data. Details about how we produce these workplace dimensions from the online review text are provided in Appendix A, Part II. The strongest correlations in the sample are between *LEARNING* and *TECHNOLOGY*, between *LEARNING* and *CAREER,* and between *TECHNOLOGY* and *CAREER*, which supports the argument that IT access is highly valued by workers for new skills and future opportunities. *** p<0.01, ** p<0.05, * p<0.1.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) LEARNING	1						
(2) TECHNOLOGY	.139***	1					
(3) COWORKERS	-.025	.025	1				
(4) FLEXIBILITY	.085*	.037	.084*	1			
(5) COMPENSATION	.037	-.052	-.047	.100**	1		
(6) CAREER	.153***	.234***	.064	-.013	.032	1	
(7) INDUSTRY	.080*	.122***	.051	.048	.018	.103**	1
(8) LOCATION	.086*	.117**	.081*	.076*	-.079*	.003	.094**

TABLE C.3: WORKPLACE SKILLS DEVELOPMENT, EMERGING IT USE, AND JOB SEARCH

Table notes. This table tests interactions between IT systems and whether the employer is known as a good place to learn new skills, as indicated by the Glassdoor reviews. To develop a measure of whether an employer is known to be a good place to learn skills (*SKILL*), we manually selected phrases that indicate skill development at work.[21] Details of this process are available in Appendix A, Part II. The dependent variable in columns (1) and (2) is a binary variable indicating whether they used emerging IT systems or older IT systems, respectively. The dependent variable in columns (3) through (5) is the logged target wage and the regression is specified similarly to what was used in earlier tables. In addition to the variables shown in the regressions, all regressions control for wages, job tenure, experience, race, gender, and education. Standard errors are shown in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Columns (1) and (2) indicate correlations between emerging IT use and the workplace *SKILL* measure. The estimates suggest that the employer measure of *SKILL* is positively correlated with the use of emerging IT (t=1.97), but negatively correlated with the use of mature IT systems (t=-2.57). Columns (3) through (5) test our target wage specification with the *SKILL* measures. T he estimates in column (3) indicate that more is required to induce IT workers to leave firms that are higher on the *SKILL* measure, which is consistent with the view that opportunities for skill development are valuable to IT workers (t=2.00). The estimates in (4) indicate that this relationship holds only when the use of emerging IT and an emphasis on skill development are present together, although the estimate on the interaction term is significant only at the 10% level. The relationship does not hold when either of these two factors is present in isolation. In fact, the estimated coefficient on emerging IT when present alone is negative, which is consistent with firms paying a premium for workers who are already skilled in that technology. The overall pattern of estimates is similar in column (5) after adding a measure controlling for the extent to which workers note longer hours (*HOURS*) in their reviews.[22] We include this measure because some employers that are known as being good places to learn also require long hours from their workers. The estimate on *HOURS* is negative and significant (t=-2.13). The coefficient estimates on Emerging IT, SKILLS, and the interaction terms in columns (4) and (5) suggest that workers who use emerging IT and are employed at firms that are known as good places to acquire skills set their target wages about 5% higher.

	(1)	**(2)**	**(3)**	**(4)**	**(5)**
DV	**Emerging**	**Old High-growth**	**Log (Target)**	**Log (Target)**	**Log (Target)**
Model	**Logit, FE**	**Logit, FE**	**FE**	**FE**	**FE**
Emerging				-0.059**	-0.066**
				(0.028)	(0.032)
SKILLS	0.371***	-0.388**	0.016**	0.011	0.017**
	(0.119)	(0.151)	(0.008)	(0.008)	(0.008)
HOURS					-0.003
					(0.008)
Emerging x SKILLS				0.095*	0.096*
				(0.053)	(0.050)
Log (Wage)	0.183	-0.440**	0.807***	0.806***	0.831***
	(0.169)	(0.188)	(0.010)	(0.010)	(0.011)
Location FE	✓	✓	✓	✓	✓
Demographic controls	✓	✓	✓	✓	✓
R-squared			0.761	0.761	0.777
Observations	2,414	2,414	2,414	2,414	2,414

[21] We also tested variations that account for phrase negation in the review text with minimal change in coefficients.
[22] We thank an anonymous reviewer for recommending this test.

TABLE C.4: EMPLOYER CHARACTERISTICS OF WORKERS' SOURCE AND DESTINATION FIRMS

Table notes. One potential source of bias is related to the unobserved non-wage characteristics of the job for which the worker is searching. If workers who use emerging IT systems tend to target employers who offer jobs with undesirable characteristics, such as restrictive work practices or the use of old technologies, it can bias our estimates upwards. Alternatively, if they target jobs offering more desirable characteristics, it will exert a downward bias on our estimates. The fact that the job board does not permit workers to set targets for non-wage characteristics should discourage job switchers who are looking for a significant change in work context, but some workers may be looking for such changes. We cannot directly control for the non-wage component of the worker's target job, but we can probe the importance of this source of bias using the resume data to assess patterns of employment transitions between firms for which we have Glassdoor data, which we show in this table.

We use IT workers' career histories through 2011 to identify, for job transitions, characteristics of source and destination firms. For each employer in our Glassdoor sample, we compute the share of workers using emerging IT (*Emerging%)* and the value for each employer work attribute, and we label these as being High if they are greater than or equal to the median value for that feature for all firms. Then, we report the fraction of worker transitions that fall into each of the following four categories: moving from an employer with a high value of a given feature to another employer with a high value for that feature (high to high), and using similar notation, high to low, low to high, and low to low.

Panel A reports these fractions based on workers in the sample who left a *high-emerging IT* employer; Panel B is based on workers who left a *low-emerging IT* employer. For instance, the value in the first row and first column indicates that of works leaving high emerging-IT employers, 71.3% went from an employer where the value of LEARNING was high to another firm where this value was also high. 7.1% of these workers went from an employer where the value of LEARNING was high to a firm where this value was low.

The results suggest that i) workers tend to move between firms that are similar in terms of the non-wage compensation they offer and ii) the patterns are similar for workers whether or not they are using emerging IT systems. The transition patterns in the table are inconsistent with the argument there is a large pool of workers who use emerging IT systems and are seeking to move to jobs with undesirable non-wage characteristics.

	(1)	(2)	(3)	(4)
	High to High	**High to Low**	**Low to High**	**Low to Low**
A. Percentage of workers leaving *High-Emerging IT* Employers, N=86,219				
LEARNING	71.3	7.1	2.5	19.2
TECHNOLOGY	61.9	6.2	4.0	27.9
COWORKERS	52.4	5.6	4.2	37.8
FLEXIBILITY	38.3	4.5	5.5	51.8
COMPENSATION	36.9	4.9	5.1	53.1
CAREER	61.0	6.7	3.4	28.9
INDUSTRY	61.8	13.4	3.7	21.2
LOCATION	45.0	4.9	5.1	45.0
B. Percentage of workers leaving *Low-Emerging IT* Employers, N=80,942				
LEARNING	26.6	5.8	9.9	57.7
TECHNOLOGY	35.2	6.4	8.1	50.3
COWORKERS	43.1	5.6	8.3	43.0
FLEXIBILITY	57.4	8.2	7.3	27.2
COMPENSATION	52.6	6.9	7.6	33.0
CAREER	30.2	5.1	8.1	56.6
INDUSTRY	25.8	5.4	9.4	59.5
LOCATION	48.3	7.7	8.0	35.9