

Reaching for the Stars: Who Pays for Talent in Innovative Industries?

**Fredrik Andersson, Cornell University,
Matthew Freedman, University of Maryland,
John Haltiwanger, University of Maryland and NBER,
Julia Lane, National Opinion Research Center
Kathryn Shaw, Stanford University and CEPR, NBER**

July 2006

Abstract:

Innovation in the U.S. economy is about employing and rewarding highly talented workers to produce new products. Using unique longitudinal matched employer-employee data, this paper makes a key connection between talent and firms in markets with risky product innovations. We show that software firms that operate in product markets with highly skewed returns to innovation, or high variance payoffs, are more likely to attract and pay for star workers. Thus, firms in high variance product markets pay more up-front—in starting salaries—to attract and motivate star employees, because if these star workers produce home-run innovations, the firm's winnings will be huge. However, we also find these same firms pay highly for loyalty: star workers that stay with a firm have much higher earnings in firms with high variance product market payoffs. The large effects on earnings are robust to the inclusion of a wide range of controls for both workers and firm characteristics. One key control is that we also show that in firms that have actually hit home runs, with high revenues, the rewards for star talent are even greater. We also find that the dispersion of earnings is higher within firms with high variance product payoffs.

JEL Codes: J240, J310, L200, L860

Thanks to Tim Bresnahan, Charlie Brown, Ben Campbell, Erica Groshen, Phil Hardiman, Edward Lazear, Alex Mas, Paul Oyer, Julie Wulf and seminar participants at Stanford University, Washington State University, UC Santa Barbara, the Society of Labor Economics meetings, for their helpful comments. This document reports the results of research and analysis undertaken by the U.S. Census Bureau staff. It has undergone a Census Bureau review more limited in scope than that given to official Census Bureau publications. This document is released to inform interested parties of research and to encourage discussion. This research is a part of the U.S. Census Bureau's Longitudinal Employer-Household Dynamics Program (LEHD), which is partially supported by the National Science Foundation Grants SES-9978093 and SES-0427889 to Cornell University (Cornell Institute for Social and Economic Research), the National Institute on Aging Grant R01~AG018854-02, and the Alfred P. Sloan Foundation. The views expressed on statistical issues are those of the author(s) and not necessarily those of the U.S. Census Bureau, its program sponsors or data providers. Some or all of the data used in this paper are confidential data from the LEHD Program. The U.S. Census Bureau supports external researchers' use of these data through the Research Data Centers (see www.ces.census.gov). For other questions regarding the data, please contact Jeremy S. Wu, Program Manager, U.S. Census Bureau, LEHD Program, Demographic Surveys Division, FOB 3, Room 2138, 4700 Silver Hill Rd., Suitland, MD 20233, USA. (Jeremy.S.Wu@census.gov <http://lehd.dsd.census.gov>).

1. Introduction

Innovation occurs because firms employ and reward highly talented workers to produce new products. Yet innovation is inherently risky – particularly in the production of information goods and services. This paper makes a key connection between the riskiness of the product markets in which firms operate and the distribution of earnings for workers. It also finds that firms in high risk markets pay highly for loyalty: employees that stay with a firm have much higher earnings in firms with high variance product market payoffs.

The basic insight of our approach is that we recognize that cost of developing a new information product is very high in both time and worker skill. If the new product is widely adopted, the payoff to the firm can be enormous, but if the new product languishes with little use, the loss can be commensurately large. Software products are the poster child for this phenomenon. Video games are at or near the top of the list for high stakes product development: the latest release of Grand Theft Auto (Vice City) has had more than \$200 million in revenues, while many games make much less. By contrast, in the same industry, business applications software (e.g., database software) is much less risky: once applications have been adopted, software producers have an installed client base that provides a degree of stability for future product development. Our approach examines how these very different potential payoffs in the product market translate into how firms hire and retain “star” software developers.

Our focus is to investigate the link between the differential payoff distribution for different types of products and the earnings distribution of workers in the software industry. We examine how firms in a product line where “home run” products matter, attract, motivate, and retain star employees. We do this using a rich new longitudinal matched employer-employee data source that permits tracking of outcomes for both firms and workers in the software industry over time. On the firm side, we have rich information about the detailed product mix and revenue outcomes for each firm so that we can measure both its actual payoff and its potential payoff distribution based upon its product mix. On the worker side, we measure both earnings levels (including the contribution of exercised stock options and bonuses) and earnings growth due to both within firm pay increases and job-hopping between firms. We use this rich data source to investigate the connection between the payoff to high stakes products and the rewards to stars in the software industry.

The theoretical motivation for the empirical analysis begins with the assumption that all innovative firms want workers who are good at designing or picking new projects. The key insight of the model is that some firms value this talent much more than other firms do (Lazear 2005). If the firm is operating in a product market in which innovation is rarely rewarded, or in which even a great project has little return, then the payoff distribution is low variance and the value of talent is low. On the other hand, if the firm is operating in a product market in which the payoff distribution has a high variance, as it does in the video game example above, then the firm values talent highly because talented workers who pick projects well can win the most in markets where there are huge potential payoffs. The prediction is that those firms operating in high variance payoff markets will hire more talent, and that they will therefore pay higher wages.

Our unique microeconomic data enable us to look inside firms and observe what innovative high potential payoff firms do to attract and reward star talent. First, we show that the high potential payoff firms pay more in starting salaries than other firms. These firms select star software workers who themselves have a history of prior success (i.e., selecting workers who have a personal history of high wage levels and high wage growth rates). Second, we show that high potential payoff firms reward workers for loyalty; that is, they reward workers for staying with the firm. Our findings show that star software workers, who achieve the highest pay levels, do so through loyalty – these workers stay with their firm and receive higher levels of performance pay. There is also an institutional side that produces this loyalty – firms in the software industry tie workers with deferred pay in the form of stock options that fully vest after four years.

It has been rare for researchers to have data that link the product market strategy of the firm to the compensation or human resource management practices of the firm. This has been done largely in the literature on CEO pay, where data are available, and the literature documents CEO pay-for-performance as a function of firm size or underlying strategies. There have been isolated instances in which researchers have done “insider” studies in some firms in which they have documented a link between strategy and performance (Baker and Hubbard 2003; Stern 2004; Wulf 2002, 2005; Garicano and Hubbard 2005), or have used survey data to do so (MacLeod and Parent 1999). There has also been some research on workers in the information industry, as it has become an increasingly dominant part of the U.S. economy (Chevalier and Ellison 1999, Garicano and Hubbard 2005, Fallick et al. 2005, Lerner and Wulf 2005). Typically,

however, there has been a major gap between the theoretical models and the empirical models of incentive contracts and sorting. That is, each theoretical model of some form of incentive pay states the assumptions under which that form of pay is optimally adopted, but empirical researchers may, at best, show that some firms succeed with incentive pay, but do not empirically model its adoption. We show that incentive pay plans and sorting aimed at high talent are optimal when the firm's strategy is to operate in high potential payoff product markets.

In short, our analysis reveals that firms that operate in innovative high payoff product markets will select star workers and will pay stars both higher starting salaries and higher performance pay. These innovating firms put substantial sums of money up front in the form of salaries for talent because they are betting on a high stakes game of producing winning high-payoff products. Of course, these high-stakes firms control the probabilistic outcomes: they pay for performance by sorting workers, or by paying high wages to attract the best talent so the probability of winning is higher for them. They also pay for stars with performance pay with experience and with requirements of loyalty, which increases the probability that they will win the game. As a result, the firms playing in the high-stakes market for innovations cause earnings inequality to rise. When one of their projects hits it big, customers around the world buy that project and it is enormously profitable. That makes the talent at innovative firms enormously valuable, and hence not only increases the variance of pay dramatically for starting salaries but also increases pay *ex post* for experienced workers for winning the innovation competition. The highest skilled stars are much more highly valued and paid than those who are slightly less skilled.

The paper proceeds as follows. In the next section, we provide some background basic facts about the software industry that help motivate our approach and analysis. Our application of the Lazear (2005) model is sketched in Section 3. A detailed description of the data is provided in section 4. The empirical specifications we explore and the results from these specifications are presented in Sections 4 and 5. Concluding remarks are provided in section 6.

2. Background Facts

We present some basic facts in this section that help motivate the approach and analysis that follows. These facts are aimed at describing the revenue payoff distribution for firms and the wage distribution for workers in the software industry.

First, software firms on average pay relatively high salaries, but a small subset of workers in the industry receive particularly high wages. The top half of Table 1 provides summary statistics about the distribution of income from the 2000 Decennial Census Public-Use Microdata Sample (PUMS) for workers in all industries and in the software industry.¹ Workers in the software industry as a whole earn more than twice what workers in all other industries earn (looking at either the mean or median).

The PUMS data suggest that, while the variance of pay in software is relatively large, compensation is not appreciably more skewed to the right for workers in software than in other industries. However, these data do not measure earnings that are important in software – performance bonuses and stock options. Thus, in the bottom half of Table 1, we use data for workers in the software industry from employer-filed Unemployment Insurance (UI) records, which contain data on all earnings including bonuses and stock options.² Because these UI data do not contain hours of work or occupation information, we limit our sample to workers earning at least \$50,000 in the software industry.³ Moreover, we focus on job spells in the software industry that are ongoing in 1997 since this sample of spells is useful for our later analysis exploiting firm level characteristics. Four different measures of earnings related to this spell are reported. The earnings at the start of the job spell, the earnings for experienced workers measured as the earnings at the end of the observed spell in the software industry, the earnings of workers one year prior to the end of the observed spell in the software industry and the earnings in the last quarter of the prior spell (conditional of course on the prior spell being observed). There are left and right censoring issues that we deal with in the standard ways in our econometric analysis below, but even with these limitations there are interesting patterns observed in the lower panel of Table 1. The earnings levels for end of spell are much higher than

¹ We focus on full-time workers between 21 and 44 years of age.

² These data are from the Longitudinal Employer-Household Dynamics (LEHD) Program and are described in detail in our data section below.

³ The \$50,000 cutoff is discussed in more detail below, but note that based on Decennial Census PUMS data, two-thirds of all software workers and four-fifths of software engineers (Census occupation code 102) have total earnings of at least \$50,000. When we replicate the mean of total earnings in Table 1 using only those software engineers, it is little changed, rising from \$90,668 in the table to \$103,881. Note finally that the \$50,000 is the worker's earnings when we last observed him or her in the data – 36 percent of those earning \$50,000 or more when we last observe them have starting salaries less than \$50,000. Fortunately, Table 1 (as well our robustness analysis discussed in more detail below) shows that by using a relatively simple income cutoff, we can identify the software developers and managers in the administrative data. That is, focusing on workers earning more than \$50,000 annually in constant 2001 dollars yields workers that are well identified as software developers and managers.

beginning of spell and are very skewed to the right. The skewness is especially pronounced for the most highly paid workers (the top decile in terms of end-of-spell earnings). The median for the top decile is \$670K and the 90th percentile \$6.7 million at end of spell with a median of \$81K and a 90th percentile of \$174K at the beginning of spell. Clearly, this top group has enormous average within spell wage growth at both the median and especially at the 90th percentile. Much of this high end-of-spell earnings could be bonuses and/or exercised options upon leaving the firm (since it is end-of-spell), so examining the patterns one year prior to the end of spell is also of interest.

Table 1: Summary Information about the Earnings Distribution

	<i>Mean</i>	<i>Median</i>	<i>90th</i>	<i>SD</i>
<i>2000 Decennial Census - Workers 21-44, 35+ Hours/Week & 35+ Weeks/Year</i>				
All Industries				
Total Earnings	40,918	31,891	70,160	183,134
Wage and Salary Income	38,685	31,466	69,097	173,449
Software Publishing Industry (NAICS 5112)				
Total Earnings	80,787	63,782	127,563	334,906
Wage and Salary Income	80,006	63,782	127,563	333,669
Computer Software Engineers (102) in the Software Publishing Industry (NAICS 5112)				
Total Earnings	90,668	70,691	138,193	369,374
Wage and Salary Income	90,496	70,160	138,193	369,777
	<i>Mean</i>	<i>Median*</i>	<i>90th*</i>	<i>SD</i>
<i>Unemployment Insurance (UI) Data - Workers 21-44, Earning \$50,000+ Annualized</i>				
Software Publishing Industry (NAICS 5112)				
Starting Annualized Earnings (Excludes Left-Censored)	69,353	59,665	108,692	82,432
Ending Spell Annualized Earnings (Censored and Uncensored)	344,268	95,508	310,644	2,051,985
One-Year Prior Annualized Earnings (Censored and Uncensored)**	199,172	86,796	220,760	1,101,658
Prior-Spell End-of-Spell Annualized Earnings***	60,951	51,532	100,987	133,153
Top Decile of Workers (by Last Observed Earnings) in Software Publishing Industry (NAICS 5112)				
Starting Annualized Earnings (Excludes Left-Censored)	107,660	80,899	184,951	142,526
Ending Spell Annualized Earnings (Censored and Uncensored)	2,532,500	670,993	6,688,470	6,064,204
One-Year Prior Annualized Earnings (Censored and Uncensored)**	750,551	171,642	1,338,380	2,862,843
Prior-Spell End-of-Spell Annualized Earnings***	98,467	73,434	164,194	150,428


* Average within a 10% band around the true percentile. ** Annualized earnings three quarters prior to last observed full quarter.

*** Includes only individuals for whom we observe a prior spell in the data.

Second, the pay of software workers rises markedly with tenure. Figure 1 compares the earnings distribution for starting-workers salaries to the distribution for experienced-workers earnings (based on the UI data containing options, and exact values are given in the bottom half of Table 1). While 70% of starting earnings are below \$75,000, only 29% of experienced workers earn below \$75,000 (experienced workers have an average tenure of five years). Similarly, only 4% of starting salaries are above \$150,000, but 21% of experienced workers earn

above that amount. Since starting salaries include the salaries paid to new but experienced workers, earnings rise markedly with tenure.⁴

Figure 1



Third, there appears to be a high variance to the gains to innovation in the software industry. We present the distribution of revenues for the top ten video games in Table 2. The distribution is highly skewed; even restricting attention to the top ten games of 2002, the top game earned nearly five times as much as the bottom. We have selected video games as an illustrative example; as mentioned in the introduction, not all software firms have such a skewed payoff distribution for their products. In the consumer video game market, the costs of consumers switching to a new “hot” game is minimal and hence firms in this market have enormous potential gains if the product does ‘hit’ in the market. However, the same is not true for firms that produce enterprise resource software for large mainframe computers; these firms have a lower variance payoff. Since large firms have locked into a software provider and purchase it year after year, the provider is profitable, but software product innovations do not have enormous upside potential gains. A good example of this is the SAS Institute, which produces statistical software for businesses. SAS sells its software through licenses to firms, and have about a 97% renewal rate (Stanford GSB Case 1997).

⁴ By experienced earnings, we are referring to end of spell earnings (see data appendix for details). On average, the experienced earners have more than five years of tenure. For now, we include both censored and uncensored spells for experienced earnings but exclude left-censored spells for starting earnings. In subsequent sections, we control for censoring in all of the empirical specifications relating earnings outcomes.

Table 2: Top Video Games, Ranked by 2002 Sales Revenues

<i>Game</i>	<i>Firm</i>	<i>2002 Revenues (Millions)</i>
Grand Theft Auto Vice City	Take 2	\$218
Grand Theft Auto 3	Take 2	\$120
Madden NFL 2003	Electronic Arts	\$119
Medal of Honor	Electronic Arts	\$73
Kingdom Hearts	Square Enix	\$59
Spider Man	Activision	\$54
Halo	Microsoft	\$51
SOCOM Seals	Sony	\$50
Super Mario Sunshine	Nintendo	\$49
Tony Hawks	Activision	\$46

Source: Merrill Lynch, “Reinstating Coverage of Video Game Industry,” *In-depth Report*. January 21, 2004.

In what follows, we provide a model that links the skewness of firms’ payoff distributions to the skewness and high pay of software workers. Although our empirical analysis encompasses firms from different software product lines, we use the video game product line to illustrate some prima facie evidence.

3. Model of Innovation

We model the process of producing innovative software products, though this process may well apply to innovations undertaken by most knowledge workers. The fundamental characteristic of software production is uncertainty — not knowing whether an innovative product will pay off.⁵ In software innovation, two groups of employees must select projects: software developers must begin working on a new software project not knowing whether they will develop a great product, and software managers must allocate funds to research projects not knowing whether the product will succeed in the market. Thus, a model of project selection pertains to the work of managers and software programmers (or developers).

Given the uncertainty about whether projects will be successful, the key role of an employee seeking to make innovations is to create or pick the best projects. The model by Lazear (2005) demonstrates how employees who are the best at creating or picking projects should be sorted among firms operating in high variance payoff markets.⁶ Assume that projects can have

⁵ There are other related forms of uncertainty about product payoff that have similar predictions. Suppose that a component of the uncertainty is about whether the new idea is implemented well. In this latter case, the star programmers may be those that implement the idea well (e.g., without problematic bugs or other product market features that would have an adverse impact on the returns from the product).

⁶ Similar ideas have been raised before, notably by Prendergast (2000, 2002).

two outcomes, a good outcome with probability P , and a bad outcome with probability $(1-P)$. The uncertainty derives from the fact that software engineers and managers do not know which projects are good and which are bad. As a result, they can make two types of error: “false positive errors” in accepting projects that they believe are good, H' , but that later turn out to be bad; and “false negative errors”, $1-H$, in which they can reject a project that would have turned out to be a good project. More specifically,

$1-H \equiv 1-\text{Probability (accepting a project | project it is actually good)} \equiv \text{false negative}$

$H' \equiv \text{Probability (accepting a project | project it is actually bad)} \equiv \text{false positive}$

If the firm chooses to undertake a good project and it pays off, the firm earns $\$X$. If, on the other hand, the firm chooses to undertake a project that turns out to be a bad one, it costs the firm $-\$Y$. The firm has zero costs and zero revenue if rejects projects early. Given these probabilities and net revenues, the expected payoff for the firm is

$$E(\text{payoff}) = PHX - (1-P)H'Y + P(1-H)*0 + (1-P)(1-H')*0$$

$$E(\text{payoff}) = PHX - (1-P)H'Y$$

Firms that achieve a high payoff are those that have a high value of PHX . Firms that fail, meanwhile, are those that have a high value of the losses $(1-P)H'Y$.

So what is the value of talent or skills in these firms? Lazear (2005) defines a star as one who has a higher probability of accepting good projects when they really are good, and a lower probability of accepting bad projects. This star ability could come from innate talent, or it could be developed as human capital on the job through learning, or it could arise from higher effort in response to incentives. In any event, star managers must pick or accept projects and star programmers must develop (or create) great projects. Both sets of skills are important determinants of success in the software industry. Thus,

$(H + \epsilon) \equiv \text{Star worker's Probability (accepting or developing a project | project is actually good)}$

$(H' - \epsilon) \equiv \text{Star worker's Probability (accepting or developing a project | project is actually bad)}$

where high ε is the measure of the quality of the star – the talent that person has in picking projects relative to non-stars.

Therefore, the value of selecting a star employee relative to a non-star employee is the incremental expected payoff, Δ :

$\Delta \equiv$ value of selecting a star

$$\Delta = [P(H + \varepsilon) X - (1-P) (H' - \varepsilon)Y] - [PHX - (1-P)H'Y]$$

$$\Delta = \varepsilon [PX - (1-P)Y]$$

Hence, firms in high variance payoff markets value star talent the most, since firms that have either high potential payoffs from good project selection (large $\$X$), or large potential losses from bad project selection (large $-\$Y$), gain from having stars with extra talent ε .

This implication is illustrated using a continuous distribution of payoffs in Figure 2. The continuous distribution of payoffs is consistent with the model above; one can think about any given firm having a range of projects with different payoffs but with each project having the type of payoffs and probabilities previously described. The bold line in the Figure 2A shows a high variance payoff distribution and the bold line in Figure 2B shows a low-variance payoff distribution. The dotted (or blue) line is the change in the distribution from star talent. The left tail shifts right due to stars because there are reductions in false positives; that is, for any given project, the star reduces by ε the probability H' of losing $(1-P)Y$ (so the star does not approve or produce projects that later fail because they were “truly” not good projects). The right tail shifts right because the star reduces the number of false negatives; that is, for any given project, the star increases by ε the probability H of accepting a project that is good and has payoff PX . Thus, by shifting the payoff distribution to the right, the mean payoff rises from PA_1 to PA_2 in the payoff distribution of Figure 2A. This is the gain to paying for or hiring stars—and that gain must exceed the cost of the star employee. Figure 2B depicts a narrower underlying project payoff distribution, which would arise when projects are less risky and thus losses and gains are smaller. When a star shifts this low-risk payoff distribution due to their talent or project assessment, the mean gains are smaller; in that case, the gains are $PB_2 - PB_1$. As is evident in the figures, the gains to stars are smaller in low-risk product markets than in high risk: $(PB_2 - PB_1) < (PA_2 - PA_1)$. In sum, stars are more valuable in high-risk product markets of Figure 2A than in

low-risk product markets of Figure 2B, because there are bigger gains (or lower losses) to the assessment or discovery of great projects in high-risk markets.

Figure 2: Shifts in the Payoff Distribution Due to Reductions in False Positive or False Negative Errors

Figure 2A: More Risky Payoff Distribution


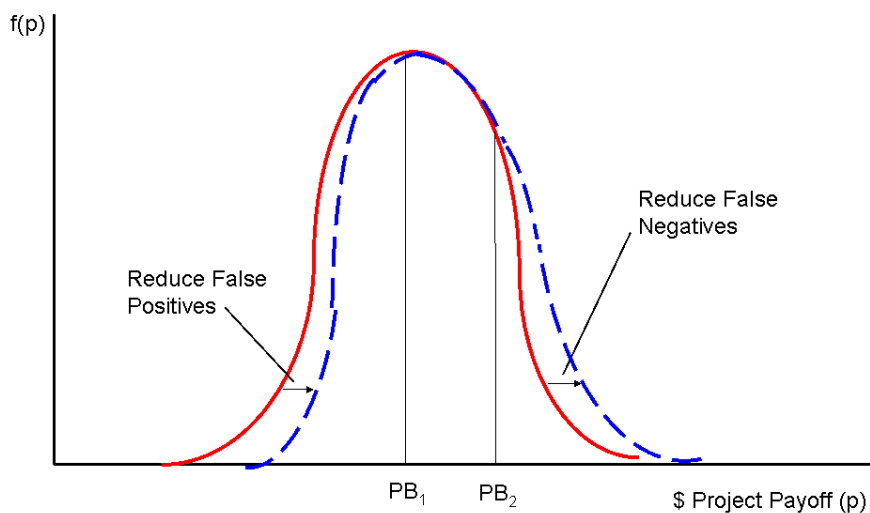


Figure 2B: Less Risky Payoff Distribution



Primary Hypothesis: Firms operating in product markets that have high variance payoffs should pay higher wages, because these firms hire and reward more highly talented software workers.

Underlying this hypothesis is the idea that firms in high variance product markets have human resource practices that select, develop, and reward highly skilled software workers. We do not observe firms' human resource practices, but we do observe in our data all the wages within each software firm, which in turn reflect these practices. We explore the relationship between the variance of product market payoffs and various dimensions of the wage structure.

Given this primary hypothesis, an open question is what mechanisms firms with high variance product markets will use to attract and retain stars. For example, do firms spend a lot of resources selecting star workers very carefully, or they do they spend a lot of resources training workers on the job and providing strong incentives that reward (and sort) star workers over time as they gain experience with the firm?⁷ The simple model above is silent about whether such firms will reward star workers through high initial wages or through sharply rising wage-tenure profiles (potentially via bonuses or stock options). However, as is evident from Table 1 and the discussion in Section 2, our data permit us to make these distinctions, and we examine the differences in our empirical analysis. In our discussion of the results, we provide some interpretation and suggestions on the ways this type of model could be extended to address these issues.

4. Data

In order to study the connection between the structure of firm's product strategies and skill demand, it is necessary to build a data set with detailed information on the earnings and employment histories of workers as well as on the product market characteristics for the firms at which these workers are employed. We take advantage of the existence of a unique employer-employee matched data set constructed and maintained by the U.S. Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) Program, and augment it with highly detailed firm characteristics from the Economic Census and worker characteristics from the 2000 Decennial Census. Appendix 1 provides further details on each of these data sets as well as on our approach to integrating the data.

⁷ Note that an additional way firms can reward star performers is by putting them on the best projects, or giving them time to do their own publishable work. Stern (2004) shows that star scientist 'pay' to be on the best projects by taking lower wages in exchange for preferred work.

The Prepackaged Software Industry

We test the hypotheses of our model by focusing on the prepackaged software industry, which corresponds to the four-digit SIC 7372 (NAICS 5112). This narrow focus has a number of key advantages. The first is the payoff structure: the software industry is characterized by high variance, which is vividly illustrated by the video game example in Table 2. As we show below in our results, there are substantial rewards to producing “hot” products. The second is the link between the firm and the product, which are closely intertwined in this industry. The software firm is an R&D intensive unit with a potentially high variance payoff to innovation; by contrast, many traditional industries, such as automobile manufacturing, while characterized by R&D intensive segments, are not innovative across the board. An additional related advantage of studying the software industry is the ability to trace directly the performance of its primary employees, including software developers and managers, and to link employee performance to the payoff structure of the firm. In other industries, the “knowledge” workers are a smaller component of employment and have a less direct impact on the output of their firms.

A final advantage of studying software is the richness of the available data. In the Economic Census surveys it conducts every five years, the Census Bureau collects a broad array of information on firms that produce software. The data collected for the software industry include detailed product-line information (described below) that we use to construct a measure capturing the variation in the payoff structure by product. These data also provide measures of the size and age of software firms.⁸ These are likely to be important controls that are correlated with product market strategies.

The Dataset on Software Workers and Establishments

The data on software workers is derived from the larger database of individual records within firms created by the Longitudinal Household-Employer Dynamics (LEHD) project housed at the Census Bureau. The LEHD’s longitudinal wage database is derived from the quarterly records of the employment and earnings of individuals from the unemployment insurance (UI) systems data, matched to internal administrative records containing information on workers’ date of birth, race, and sex.

⁸ We thank Ron Jarmin for kindly sharing information on firm and establishment ages with the LEHD Program for this project.

These data have several advantages. First, since the scope of the LEHD data is virtually the full universe of employers and workers, the movements of workers through the earnings distributions within firms as well as across firms can be tracked accurately. Second, the earnings data represent the earnings that workers are actually paid, not workers' memories of their earnings.

A third key advantage of these administrative data is that the earnings measures are quarterly and include bonuses and exercised stock options.⁹ Although the earnings information does not include fringe benefits, bonuses are included. Obviously, valuing stock options is quite difficult: in this case, the options are valued when they are exercised, or when the employee cashes in the options and receives the value. We do not have data on when options are granted to employees. However, our sense is that the exercised options available in our data are the preferred measure of pay for our analysis, rather than options granted to employees. Indeed, as Oyer and Schaefer (2002) point out, it takes four years for stock options to be fully vested, and as Russell (2005) notes, for a typical software company, options are worth nothing for the employee's first two years, and then are vested at a rate of two percent a month for the remaining three years.¹⁰ Thus, the value of options granted depend on two uncertain outcomes – on whether the employee stays with the firm until the options are vested, and on the growth of the stock price of the company.

It is important to emphasize that the LEHD data capture the full universe of covered employers and workers; they are not merely a sample of software workers or firms.¹¹ However, when we look at certain specific features of workers or of firms in the software industry, our data set becomes a smaller sample of the population of workers and firms.

Our basic universe of data workers follows 83,497 employment spells of workers who were employed in the software industry in ten states in the U.S. from approximately 1992 to 2001 (the exact starting years vary by state). This length of time enables us to construct

⁹ No previous studies have included stock options data for a wide range of workers across firms. The nature of our data permit us to exploit the fact that in most employment contracts, employees must exercise all options within 90 days of leaving the firm. We are able to track the earnings of employees for those 90 days and we can thus capture the value of all exercised options.

¹⁰ For very detailed analysis of options granted and other forms of pay within a large software company, see Russell (2005). She shows, for example, that people who have large option grants are also likely to have exercised options (because they are older).

¹¹ There are important exceptions. Most federal employment, and some agricultural and nonprofit employment is not covered. Independent contractors and self employed workers are similarly not covered. See Stevens (2002) for a full discussion of coverage issues.

sufficiently long worker earnings histories to address our research questions. Our primary results are based on two analytical datasets: one consisting of 51,589 employment spells and one of 26,276 spells. These subsets are based on a number of decisions. First, we limit the data to workers between the ages 21 and 44 in order to model the demand for a fairly homogeneous collection of workers in the prime of their careers with similar educational vintages. This reduces the sample from 83,497 to 67,452. Second, we limit our data to those workers making more than \$50,000 at the end of their 1997 spell. The rationale for this is that UI data do not contain information on hours of work or occupation. Therefore, to limit the data to workers who are likely to be full-time and in the more skilled occupations, we choose those making more than \$50,000. Appendix 1 contains an extensive discussion supporting our decision to limit the sample in this way. Together, the age and earnings limitations reduce the sample to 51,589 spells. Finally, while most businesses in our sample of workers could be successfully matched to the Economic Census, a smaller subset had complete information on the establishment, including size, age, sales, and detailed product line information. There are 26,276 spells for which we have complete information regarding firm characteristics as well as worker characteristics. All told, 688 software firms are represented in this sample.¹²

Lastly, we also construct a subset of data of employees in software occupations based on the occupational information in the 2000 Decennial Census. For this sample, we limit our data to those individuals in the software industry who are software engineers, developers, or managers (irrespective of earnings), dropping those who are in other occupations within the software industry. Because the Decennial Census is a 1 in 6 sample of the population in 2000, our sample falls to 2,638 workers in software occupations. This smaller dataset is used to check our results using the larger data set discussed above, and thus the results using these data are referred to in footnotes below.

Measuring Earnings Levels and Earnings Growth

¹² Throughout this paper, when we refer to a firm, we are referring to a firm defined at the State Employer Identification Number (the SEIN, or UI account number), which is the unit of observation in the UI-Wage data. It is an 11-digit number used for reporting taxes at the state level. For single-unit firms, this reflects the entire firm, but for multi-unit firms, the SEIN reflects activity of the firm within a given state. We are able to match the workers to information in to the Economic Censuses since the UI files also include the federal Employer Identification Number (the EIN is on the ES-202 data that is part of the related administrative data system). The EIN is a nine-digit number assigned by Internal Revenue Service (IRS) and used for federal tax purposes by employers, sole proprietors, corporations, partnerships, non-profit organizations, trusts, estates of decedents, government agencies, certain individuals, and other business entities.

A major advantage of the administrative data is that they are longitudinal in both workers and firms. In other words, the data have information about spells of employment with a firm and the associated earnings over long periods of time. These unique data permit us to capture multiple facets of worker earnings profiles. In modeling the link between product markets and compensation, we use information on workers' earnings trajectories within firms, earnings growth associated with transitions between firms, and earnings levels for new and old workers in the firm.

We use four measures of earnings in the empirical analysis. One measure is beginning-of-spell earnings, which corresponds to a given worker's total earnings in the first full quarter of employment with each employer (with the dollar values at annualized 2001 dollars).¹³ The next set of measures is for experienced workers: one includes end-of-spell earnings,¹⁴ which potentially contain exercised stock options, as workers must exercise their options within 90 days of quitting. Another, which does not include exercised stock options, is earnings one year prior to the end of the spell. Finally, for those same workers for whom we observe a prior spell, we measure the level of earnings in the last full quarter of his or her prior job.

We use two measures of earnings growth. Earnings growth within the firm, or within-job earnings growth, is the difference between end-of-spell and beginning-of-spell earnings.¹⁵ Between-job earnings growth is the difference between earnings in the first full quarter of a given worker's new software job and the last full quarter of his or her prior job.¹⁶

Measuring the Product Market Payoff Dispersion for Firms

Investigating the other important component of the hypothesis requires estimates of the variance of the expected payoff for the product market in which each firm operates. For the prepackaged software industry, the 1997 Economic Census delineates 30 detailed product lines, ranging from consumer game and entertainment software to business graphics design and layout

¹³ Beginning-of-spell earnings capture new workers to the firm and censored data of new workers in our data. Sixteen percent of the beginning-of-spell earnings are censored.

¹⁴ We measure this as a worker's last full quarter of annualized earnings in a given spell. End-of-spell earnings captures workers leaving the firm and censored data when our observations end. Forty percent of the end-of-spell earnings are censored.

¹⁵ More specifically, within-job earnings growth is defined as log annualized end-of-spell earnings less log annualized beginning-of-spell earnings, divided by the number of full quarters that a worker was on the job.

¹⁶ More specifically, between-job earnings growth is defined as log annualized beginning-of-spell earnings in the new job less log annualized end-of-spell earnings in the old job, divided by the number of full quarters between jobs. Clearly, between-job earnings growth is only defined for those individuals in the sample for whom we observe them in a job prior to their software job (i.e., those whose software jobs are not left censored and those who are not recent entrants or re-entrants into the labor market).

software to vertical industry banking software to mainframe computer applications. Establishments in the Census are asked to provide data on its revenue for each of the thirty product lines, and we exploit this revenue information to order to construct a measure that reflects the variance of payoffs in each product category.

Each firm's *Product Payoff Dispersion* is created in two steps. First, for each of the thirty product classes, we calculate the 90/50 ratio of the log of revenue per worker. Because firms have multiple products, we treat each product within each firm as though it were a separate product revenue stream, and calculate the 90/50 ratio for the thirty product markets. Second, for each firm, we calculate its payoff dispersion in the product markets in which it is operating by weighting the product-specific 90/50 ratios for the thirty products by the percent of revenue that the establishment has in each product class. More specifically,

$$\text{Product Payoff Dispersion}_j = \sum_{k=1} [\% \text{Revenue}_{jk}] (\text{Product Revenue Dispersion}_k)$$

where $\% \text{Revenue}_{jk}$ is the percent of firm j 's revenue coming from product class k , for product classes $k = 1, 2, \dots, 30$. The variance specific to product k across firms is calculated as a 90-50 ratio of log revenue per worker:

$$\text{Product Revenue Dispersion}_k = \sum_{j=1} [\ln(\text{revenue}_{jk} / \text{worker}_{jk})^{90} - \ln(\text{revenue}_{jk} / \text{worker}_{jk})]^{50}$$

where the *Product Revenue Dispersion* for product k is calculated across all firms producing in product class k , where each product line within each firm is treated as if it is its own independent firm. See Appendix 1 for details.

There are a few key features of the firm-specific *Product Payoff Dispersion* measure. First, this measure reflects each firm's actual product mix, but not its actual revenue. The payoff measure reflects the skewness of revenue per worker in the product classes in which the firm operates. A firm with a high *Product Payoff Dispersion* measure is not necessarily a high or low performing firm, but rather has a product mix with a right skewed distribution of payoffs. Also notably, the measure of revenue dispersion in a given product line is the 90-50 ratio, because the 90-50 ratio is a simple way of capturing the skewness of the upper tail of the revenue distribution. While our model in Section 3 refers to the variance of the entire distribution (thus

the lower tail as well), we focus on the upper tail because we do not observe firm's losses – revenues are truncated at zero.

Table 3: Differences in Payoff Distributions By Product Line*

Detailed Product Line Code	Detailed Product Line Description	90/50 Ratio of Product Line Sales/Worker
<i>Detailed Product Lines in Software Publishing with Greatest Potential Payoffs/Risks</i>		
1122	Game and Entertainment Software	1.31
1183	Networking Software	1.17
1123	Home Productivity Software	1.03
<i>Detailed Product Lines in Software Publishing with Smallest Potential Payoffs/Risks</i>		
1161	Banking and Finance Software	0.66
1142	Distribution Software	0.57
1184	Database Software	0.55

*Based on the national sample of firms.

The means of the *Product Payoff Dispersion* are reported in Table 3 for the highest risk and the lowest risk product lines. The means in Table 3 suggest that one element of our model is borne out: there is variation in the skewness of revenues across product classes.

5. Empirical Model

The model in Section 3 implies that firms operating in product markets with highly dispersed payoffs will hire more highly talented (star) workers. This model generates several testable hypotheses for the empirical analysis.

Product-Specific Payoffs

A standard human capital wage equation takes the following form:

$$(1) \quad \ln(W)_{ijt} = a_0 + a_1 S_{ijt} + u_{ijt}$$

where wages are a function of skills, S_{ijt} (for person i , firm j , and time t), but the key skills in this case (like creativity and programming talent) are unobserved to the econometrician.¹⁷ We hypothesize that highly skilled workers will sort to the high potential payoff firms that value skills the most. Thus, the payoff dispersion will serve as a measure of unobserved skills in the wage regression:

¹⁷ In this paper, we always control for observable skills, and in addition we select only workers earning more than \$50K.

$$(2) \quad \ln(W)_{ijt} = \alpha_0 + \alpha_1(\sigma_j^p) + \varepsilon_{ijt}$$

where σ_{jt}^p is the payoff dispersion or variance that firm faces in its product market(s). The primary hypothesis in Section 3 implies that high payoff dispersion firms pay higher wages because they are selecting highly skilled star workers. It is important to remember that the payoff dispersion measure captures the variance of payoffs *in the firm's product class(es)*, and does not represent the firm's actual historical variance of success.

Two issues are important to bear in mind in the empirical estimation of equation (2). First, σ_j^p will not vary over time; it will be observed in one year and represent a firm-specific effect identifying the product class or classes of the firm. Second, a number of controls, discussed later, are included for other factors that may affect firm wages. Of these, perhaps the most important is the fact that we control for a measure of the actual payoff of the firm using output revenue per worker. It is thus less likely that the estimate of α_1 reflects profit or rent sharing.

We consider a number of refinements of the benchmark specification (2). One is the incorporation of the idea that wages should be more sensitive to the firm's payoff dispersion for more highly skilled workers. In software companies, it is the top talent (or the brilliant programmers) who should be paid the most for their skills in the firms operating in product markets with high payoff dispersion. By contrast, pay for low-wage workers should not be a function of the firm's payoff dispersion (because worker sorting is irrelevant). This suggests estimating a quantile regression of the following form:

$$(2') \quad \ln(W)_{ijt} = \alpha_0 + \alpha_1^{\text{skill}}(\sigma_j^p) + \varepsilon_{ijt}$$

and examining α_1^{skill} to determine whether there is rising sensitivity to payoff potential with skill at each point in the wage distribution.¹⁸

¹⁸ This point is also made in Buchinsky's (1994) model of wages in which he shows that the returns to education are higher at high wage quantiles, though the returns to experience are lower at high wage quantiles. Hallock et al. (2004) show that among CEOs, the sensitivity of wages to firm performance rises as one moves up the earnings distribution.

Another refinement of the model is the investigation of the impact of product variance on different facets of compensation – both in terms of levels and growth. There are a number of reasons why our current model would predict that high product variance firms would offer high base pay that rises sharply with tenure. Such firms might offer higher base pay because they value skills or talent more than do other firms and use this approach to select more highly talented people. Similarly, we would predict that pay should rise with tenure due to sorting; as in all matching models, the return to tenure would be high because stars are both retained and are paid more over time while the losers are fired. Greater human capital investment is another reason to expect wages to grow over time. Firms in high variance product markets may also pay more for effort; that is, firms may have steep incentive pay contracts.¹⁹ Teamwork is another reason we might expect wages to rise over time; if people are working in teams in which their skills are likely to be complementary with the other team members, then it may take time before an employer can identify and reward individual star talent. Lastly, literature on the software industry suggests that firms in the sector want teams to stay together for the product cycle, that they do not want to lose star talent, and that they want to provide incentives for effort.²⁰ Software firms therefore intentionally tie employees to the firm by granting stock options that vest slowly (typically over four years); this further steepens the wage-tenure profile for workers. In sum, it is likely that firms operating high variance product markets pay more for loyalty, compensating their employees for staying with the firm more so than firms in low-variance product markets.

We use the different measures of earnings in our dataset to investigate whether these outcomes occur. In particular, we examine what the rewards for talent are in terms of beginning-of-spell earnings, end-of-spell earnings, one-year prior to end of spell (which is not likely to include exercised stock options), and prior spell earnings. In a closely related fashion, we also explore the relationship between product market payoff variance and earnings growth.²¹

Could a positive coefficient, α_1 on σ_j^p , in (2) and the related specifications simply reflect a compensating differential for risk rather than a firm strategy to attract and retain stars? There are a number of different factors that suggest that this possibility is unlikely. Most importantly, risk differentials are not consistent with the typical structure of earnings in the software industry.

¹⁹ In fact, case study evidence suggests that they do – a larger percent of a workers' pay is performance based as the skill level rises (Russell 2005)

²⁰ See, for example, Russell (2005), Cusumano and Selby (1995), Hoch et al. (2000), and Stross (1997).

²¹ A particular advantage of using the latter as the dependent variable in (2) is that it permits us to abstract from unobserved fixed factors that influence the level of earnings.

Typically, workers in software are rewarded for upside gains, but they are not penalized for losses – base pay typically does not fall when the firm loses money. Therefore, software workers usually choose between two alternative pay packages in job offers: (a) low base pay but high performance-based pay or (b) high base pay but little possibility of rewards if the worker or company does well. Firms operating in high variance product markets who want to induce high performance by offering incentive contracts would therefore offer contract (a), and hence high variance firms will have *lower* starting salaries, or base pay, than firms in low-variance product markets. By contrast, our model states that firms operating in high variance product markets will want to select the highest quality workers, and quality is unobserved to the econometrician, so these firms will have *higher* starting salaries (which are base pay) than firms in low-variance product markets. In sum, a positive value of α_1 resulting from estimating equation (2) with starting salaries as the dependent variable, is likely to reflect a return to skills, not a return to risk taking.²² However, high pay at the end of a job spell could reflect a return to risk taking – if the worker had taken a lower starting salary with the hope of future uncertain gains.²³

²² There is one way in which a firm's losses will translate into lower pay for the worker – the worker will get fired and lose their returns to firm-specific human capital if the firm fails. But this too should produce a compensating wage-risk differential for experienced workers, not for young workers who have yet to invest in firm-specific skills. The lore in this industry is that there are workers who are risk-takers – who seek firms who might produce big hits as in our model – and that these risk-takers accept jobs that offer lower starting salaries for their skills, but that might produce big income gains. Using extensive data for one large software company, Russell (2005) shows that within the firm, pay levels, bonuses, and options are highly correlated across individuals, reflecting the fact that more able workers have higher pay of every kind than the less skilled.

²³ Note also that in jobs with a higher variance of returns for firms, some models would produce the conclusion that there should be less incentive pay in these high variance markets. Aggarwal and Samwick (1999) show that when the variance of stock market returns rises for a company, the percent of pay at risk (or performance pay) falls for top executives—there is a negative relationship between riskiness and pay. But they have a different kind of risk in mind, and their data measure a different kind of risk—their riskiness is the variance in stock returns due to noise or uncontrollable outcomes. It is true that in a tournaments model of incentive pay, increasing the amount of noise or luck reduces the use of incentive pay (Lazear and Rosen 1981). In our model, the variance of payoff outcomes could arise in part from an idiosyncratic shock representing noise or luck, but most importantly arises because some firms hire smarter people who select more successful products and should have pay tied to performance. In the data, we cannot identify whether the variance in the payoff arises from luck or effort, but our model of innovation proposes that it is high skill that produces high payoffs (not luck), so the coefficient α_1 on σ_j^p , should be positive, not negative. Prendergast (2000, 2002) also makes the point that higher risk environments may have more performance based pay, not less, because in those environments, the cost of determining what inputs to monitor is greater than the cost of utilizing output performance based pay. Since the source of the variance in payoffs cannot be identified (and we do not have time-series data product-specific variances or firm-specific variances), we turn to the data to determine the sign. For related empirical models of the risk-pay incentive relationships, for executives see Baker and Hall (2004), Core et al. (2003), Ittner et al. (2003), Murphy (1986), Schaefer (1998), and Wulf (2005); for reviews, see Hallock and Murphy (1999) and Murphy (1999).

In addition to the conceptual argument above, there are two technical reasons why it is unlikely that a positive value of α_1 would reflect a compensating risk differential. For one, in all specifications we include as a control a measure of the worker turnover at the firm, which can be interpreted as a control for job security. For another, we consider alternative specifications that explore the worker's earnings on her prior job spell, reasoning that if firms are paying for talent rather than risk-taking, we should observe higher earnings on prior spell jobs.

5. Empirical Results

The Earnings Levels of Software Workers

The wage regression results in Table 4 explore the relationship between earnings levels and within-job earnings growth rates with the firm's product payoff dispersion (the revenue-weighted 90-50 percentile ratio of revenue per person across the product markets in which the firm operates). All wage regressions control for the person's tenure, censoring of observed tenure, and age. Regressions in every second column also control for firm age, firm size, indicators of the density, education and industrial diversification of the county and in the worker outcome regressions dummies for quarters of accessions and separations to abstract from any macro effects, and the churning rate of jobs in the firm.²⁴

Looking at the regressions for *experienced workers earnings* (columns 1 and 2) and *starting salaries* (columns 3 and 4), it is clear that the *product payoff dispersion* variable has a very significant positive effect on each measure of earnings at all points of the earnings distribution except the very bottom. Thus, even among new hires, firms with high payoff dispersion pay higher wages, suggesting that they select more skilled workers. This is consistent with the extensive industry testimony that describes the software industry's very careful and deliberate hiring practices, all aimed at identifying the right talent and reflecting the high-commitment work environment of the industry (Hoch, et.al., 2000). As discussed above, the

²⁴ The dependent variable in these regressions is the earnings residual for each spell in the sample, where the residual is from a regression of log earnings on quadratics of tenure at job, tenure in industry, and age, fully interacted with each other and with appropriate left and right censoring dummies. See the wage regression results in Appendix 2. Using the residual is equivalent to introducing these controls in the wage regression, but the residual is more intuitively appealing than raw earnings because it reflects only that component of earnings that cannot be accounted for by basic observable worker characteristics. We also appeal to the residual in the star probability regressions used below. When the churning rate is added, it is defined as the accession rate plus the separation rate less the absolute value of the net growth rate of employment at the firm. This churning or excess worker turnover rate is included as a control, since churning may be part of the interaction of product mix and internal labor market strategies that in turn impact the structure of wages. For example, high-risk product strategies may be associated with a high worker turnover. By including this control, we mitigate concerns that the payoff potential variable is capturing a compensating differential for high turnover.

positive coefficient on dispersion for starting salaries is not consistent with the notion that firms pay a compensating wage differential for risk taking: this would have resulted in a negative coefficient.

A comparison of the wage results for *experienced workers earnings* (columns 1 and 2) and *starting salaries* (columns 3 and 4) reveals that earnings are much more sensitive to the product payoff dispersion for experienced workers than for new hires. The earnings for experienced workers can be very large by the end of their spell with the firm: 10 percent of software workers earn more than \$310,000 at the end of their spells (Table 1). The very large wage payoffs for experienced workers could reflect a number of factors: higher marginal products (as in our model), a tournament reward structure, participation in a high-performance team, or improved selection of talented workers over time in the firm. We cannot identify the differences, but Russell (2005) provides very detailed evidence for one software company that suggests that all of these factors enter the earnings of software workers.²⁵ The point is that in each case, the person who can create or pick the best projects will have more skills and more incentive pay in firms with high product payoff dispersion.

The regression results for experienced workers may be heavily influenced by the inclusion of exercised stock options in the earnings measure. To test the sensitivity of our results to the inclusion of options, we run an additional regression in which experienced earnings are redefined as those for people one year before they quit their job or prior to dropping out of our sample due to censoring (since options are most likely to be exercised when departing a firm). The magnitudes of the coefficients in the experienced-earnings regressions decline somewhat, but the basic results remain.

²⁵ All indications are that the firm in Russell's (2005) study looks very much like the typical large firm in our data – the median age is 33 and tenure ranges from 2.7 to 3.1 over 1996 to 1999, and about 65 percent were in research and development and 30 percent in management or administration. Though no previous studies are available describing in detail the forms of incentive pay for software workers, Russell (2005) provides evidence from a case study of one large software firm. She shows that pay is a function three things – base pay and the merit increases in base, bonuses, and stock options. In this firm, individual pay levels and individual deferred pay is very highly positively correlated within the firm (Russell 2005, Figure 23). This example suggests that software companies use a combination of promotion-based pay, as in a tournaments model, of individual-level incentive pay, and of group-based incentive pay. Since these forms of performance pay rise with the grade level and rise the pay level in the firm, they suggest that performance bonuses are higher for star employees, both in their allocation (as in options granted) and in the realization (as in the value of bonuses and exercisable options). As a result, in our wage regressions, the wage growth or levels should contain a portion of wage gains that are attributable to performance-based pay (as will be described in more detail below).

As a further check, we develop an entirely different measure of workers' underlying talent, by measuring talent based only on the worker's history prior to the current job. That is we regress each worker's pay from his job prior to 1997 on the characteristics of his job in 1997. The hypothesis is that if high payoff software firms in 1997 are sorting for the highest skilled workers when they hire them, then each worker's pay in his previous job will be a function of his future firm-type as a measure of his underlying talent. Therefore, we estimate wage regressions for the subsample of workers with an observed prior employment spell and examine relationship between current firm product payoff and prior spell earnings. The results here (columns (7) and (8)) are less strong, but at the median we find a positive and statistically significant impact that is quite similar in magnitude to the results for beginning of spell at the median. It makes sense that there is less of a systematic pattern for the low and high percentiles, as the theory is about how the risk-taking firm should especially compensate the most skilled workers. That is, the firm in the prior spell (which may be a firm of any type including a non-software firm) is unlikely to be rewarding especially skilled workers in the same fashion as the current risk-taking software firm.²⁶

The findings in Table 4 show that the impact of the payoff potential rises with skill level; that is, software workers at the upper reaches of the earnings distribution gain the most from working in high variance payoff firms, though workers at the median gain as well.²⁷ The sharply increasing impact of the payoff potential is illustrated in Figure 3 for experienced workers the coefficient at the 90th percentile is more than twice as large as the coefficient at the 10th percentile.

²⁶ There is an econometric concern for these specifications given sample selection (not all workers in our sample have an observed prior spell). We have considered in unreported results selection corrected results (using least squares methods instead of quantile regressions) and found results consistent with those for the median as reported in columns (7) and (8).


²⁷ In Hallock et al. (2004, page 7), they point out that "higher ability managers [would have] higher pay for performance incentives than low ability managers" due to the lower cost of effort for high ability managers.

Table 4: Quantile Regressions for Earnings Level Residuals for Software Spells

	End of Spell Earnings "Experienced Earnings"		Beginning of Spell Earnings "Starting Salaries"		One-Year Prior Earnings "Lagged Earnings"		Prior-Spell Ending Earnings "Prior-Spell Earnings"	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	All Firms	All Firms	All Firms	All Firms	All Firms	All Firms	All Firms	All Firms
<i>Tenth Percentile</i>								
Product Payoff Dispersion +	0.3226 (0.0481)**	0.3068 (0.0460)**	-0.1693 (0.0434)**	-0.0678 (0.0570)	0.1070 (0.0453)*	0.0456 (0.0494)	-0.0666 (0.1304)	-0.0309 (0.1335)
Log Revenue per Worker	0.0622 (0.0045)**	0.0288 (0.0104)**	0.0386 (0.0045)**	0.0385 (0.0066)**	0.0485 (0.0043)**	0.0290 (0.0079)**	0.0358 (0.0140)*	0.0841 (0.0236)**
Firm Average Worker Churn	1.1692 (0.1354)**	1.4403 (0.1534)**	0.9172 (0.1651)**	0.4170 (0.1691)*	0.5065 (0.1293)**	0.8833 (0.1441)**	0.4196 (0.4428)	0.3350 (0.4655)
<i>Fiftieth Percentile</i>								
Product Payoff Dispersion +	0.3750 (0.0450)**	0.3715 (0.0480)**	0.1469 (0.0372)**	0.1631 (0.0346)**	0.2701 (0.0513)**	0.2581 (0.0515)**	0.1487 (0.0566)**	0.1305 (0.0583)*
Log Revenue per Worker	0.1477 (0.0043)**	0.0874 (0.0079)**	0.0332 (0.0026)**	0.0457 (0.0057)**	0.0883 (0.0037)**	0.0634 (0.0076)**	0.0101 (0.0063)	0.0315 (0.0107)**
Firm Average Worker Churn	2.7335 (0.1423)**	3.3058 (0.1555)**	0.8869 (0.0894)**	0.7091 (0.1016)**	2.6445 (0.1490)**	2.8104 (0.1326)**	0.7588 (0.1929)**	0.8171 (0.1981)**
<i>Ninetieth Percentile</i>								
Product Payoff Dispersion +	0.7420 (0.0962)**	0.7218 (0.1180)**	0.2083 (0.0457)**	0.2841 (0.0563)**	0.6391 (0.0930)**	0.5398 (0.1002)**	0.2452 (0.0746)**	0.1252 (0.0811)
Log Revenue per Worker	0.6406 (0.0171)**	0.1469 (0.0168)**	0.0820 (0.0041)**	0.0626 (0.0078)**	0.2847 (0.0116)**	0.1372 (0.0128)**	0.0460 (0.0082)**	0.0473 (0.0149)**
Firm Average Worker Churn	1.3797 (0.2876)**	3.9577 (0.2621)**	0.8751 (0.1460)**	1.7865 (0.1492)**	1.5848 (0.2843)**	3.1629 (0.2583)**	1.3327 (0.2571)**	1.7670 (0.2558)**
<i>Ninety-Fifth Percentile</i>								
Product Payoff Dispersion +	1.0363 (0.2365)**	0.6729 (0.1689)**	0.2803 (0.0677)**	0.2902 (0.0727)**	0.6727 (0.1260)**	0.6290 (0.1425)**	0.1156 (0.1181)	0.0252 (0.1140)
Log Revenue per Worker	0.8205 (0.0166)**	0.1378 (0.0322)**	0.1107 (0.0073)**	0.0748 (0.0153)**	0.4431 (0.0159)**	0.1137 (0.0229)**	0.0311 (0.0131)*	0.0486 (0.0207)*
Firm Average Worker Churn	0.9527 (0.3737)*	4.2284 (0.4450)**	0.6759 (0.2310)**	2.0844 (0.2331)**	1.3174 (0.3657)**	3.4670 (0.3663)**	1.1724 (0.4161)**	1.9094 (0.3594)**
Controls++	No	Yes	No	Yes	No	Yes	No	Yes
Observations	26276	26276	26276	26276	26276	26276	10803	10803

All dependent variables are wage residuals from regression controlling for age and tenure (see Appendix 2). Standard errors in parentheses. * significant at 5%; ** significant at 1%. Marginal Effects are reported.

+ Weighted 90/50 Ratio of Log Revenue per Worker in Firm's Product Lines ++ Controls include (log) firm employment; (log) firm employment squared; Dummies for firm age (<6 years old, 6-10 years old, 11+ years old, with <6 omitted); firm employment growth rate; a dummy for whether the firm is in a high density, high education, and industrially diversified county; and year and quarter job separation dummies.

Figure 3

We also report some of the coefficients on the controls that are of particular interest. First, consider the effect of the firm's actual revenues on the different earnings measures. The earnings regression results show that workers are paid more when the firm does succeed; indeed, pay rises very significantly as a function of the firm's actual log revenue per employee.²⁸ The quantile analysis also suggests that high-wage workers are paid more for the firm's success.²⁹ Note that this should be interpreted as a firm-specific fixed effect: the firm that is highly productive in 1997 (when we measure the firm's revenues) pays more to workers in adjacent years as well. We also find that firm's pay is increasing in the amount of worker churning. This control is another variable that helps capture effects that may be associated with compensating differentials since this is a control for job security. Our finding that the product market payoff dispersion results are robust to the inclusion of this control provides yet more evidence that the impact of

²⁸ In interpreting these results, it is useful again to emphasize that, while the product mix payoff risk measure varies across firms, it is not driven by the realized payoffs of the firm but rather the potential payoff distribution based upon the pool of firms with that product mix. This feature substantially mitigates concerns of contemporaneous endogeneity of the payoff mix measure. This payoff risk measure does reflect a choice by the firm (i.e., the choice of product mix), but this choice is likely made either at the founding of the firm or, at the very least, is made infrequently. After controlling for firm performance, the effects of the product market payoff remain unchanged, which should further reduce concerns about endogeneity.

²⁹ For descriptions of forms of incentive pay for other knowledge workers, and models and empirical results, for the CEO literature, see Hallock and Murphy (1999), Gibbons and Murphy (1992), Jensen and Murphy (1990), and Murphy (1999).

the latter variable is associated with attracting and retaining star workers rather than compensating for risk.

In sum, the quantile regressions show that earnings are higher when firms operate in high variance product markets and that earnings are higher when firms succeed by achieving high revenues. Moreover, those workers at the upper end of the earnings distribution are rewarded disproportionately when firms operate in high variance product markets and when firms succeed by achieving high revenues. These results are robust when we look at different measures of income as well as when we subset the data substantially to only look at individuals who are software programmers (thus excluding all managers and all other well-paid employees).³⁰

The Earnings Growth Rates for Software Workers

The main results on the impact of payoff potential are mimicked when we use within-job earnings growth as the dependent variable instead of the earnings level (columns 1 and 2 of Table 5). Within-job earnings growth rises sharply with the potential payoff of the firm, and this impact is greatest for the highest earnings quantiles. In contrast, between-job earnings growth (columns 3 and 4) is not a function of the potential success of the firm: workers are rewarded for staying with the firm but not for hopping between firms. For the median worker, the effect of payoff potential for between-job earnings growth is actually negative, but at higher earnings quantiles, it is insignificant. Thus, even though starting salaries are higher for individuals working in high variance product classes (columns 3 and 4 of Table 4), firms in these high variance industries do not appear to be stealing stars from other firms by offering high starting salaries. Of course, we cannot observe whether firms are stealing stars by offering high stock option grants. But the point remains that even if they are, the stars typically stay with the firm four or more years to have their options pay out, and the stars must succeed at what they are doing. Job-hopping for higher future earnings may be a common strategy, but such job-hopping is not a viable short-term strategy for wage growth. In this sense, loyalty pays – workers must stay with firms to achieve income growth.

³⁰ This is based on a match of the LEHD data to the long form of the 2000 Decennial Census data. The sample size is reduced substantially, which is why the match is not used in subsequent analysis.

Table 5: Quantile Regressions for Growth Rate Residuals for Software Spells

	<i>Within Job Earnings Growth</i>		<i>Between Job Earnings Growth</i>	
	(1) All Firms	(2) All Firms	(3) All Firms	(4) All Firms
<i>Tenth Percentile</i>				
Product Payoff Dispersion +	-0.0055 (0.0068)	-0.0063 (0.0083)	-0.2828 (0.0676)**	-0.1768 (0.0567)**
Log Revenue per Worker	0.0034 (0.0008)**	-0.0040 (0.0016)*	-0.0216 (0.0080)**	0.0232 (0.0124)
Firm Average Worker Churn	0.0271 (0.0207)	0.1419 (0.0272)**	0.1269 (0.2350)	-0.0886 (0.2488)
<i>Fiftieth Percentile</i>				
Product Payoff Dispersion +	0.0882 (0.0077)**	0.0716 (0.0078)**	-0.2070 (0.0291)**	-0.1636 (0.0356)**
Log Revenue per Worker	0.0162 (0.0008)**	0.0015 (0.0009)	0.0054 (0.0032)	0.0254 (0.0059)**
Firm Average Worker Churn	0.2539 (0.0257)**	0.3678 (0.0166)**	0.2765 (0.1031)**	0.1494 (0.1167)
<i>Ninetieth Percentile</i>				
Product Payoff Dispersion +	0.2543 (0.0264)**	0.2949 (0.0346)**	-0.2194 (0.0939)*	-0.1650 (0.0873)
Log Revenue per Worker	0.0887 (0.0027)**	0.0245 (0.0047)**	0.0057 (0.0097)	-0.0042 (0.0174)
Firm Average Worker Churn	0.6998 (0.0762)**	1.1196 (0.1063)**	0.5118 (0.3064)	0.6335 (0.3640)
<i>Ninety-Fifth Percentile</i>				
Product Payoff Dispersion +	0.2251 (0.0435)**	0.3123 (0.0519)**	-0.1651 (0.2776)	-0.2133 (0.2728)
Log Revenue per Worker	0.1125 (0.0040)**	0.0369 (0.0091)**	0.0232 (0.0269)	0.0228 (0.0473)
Firm Average Worker Churn	0.7821 (0.1194)**	1.4663 (0.1655)**	0.8609 (0.6226)	1.1838 (0.6138)
Controls++	No	Yes	No	Yes
Observations	26276	26276	10803	10803

All dependent variables are wage residuals from regression controlling for age and tenure (see Appendix 2).

Standard errors in parentheses. * significant at 5%; ** significant at 1%.

+ Weighted 90/50 Ratio of Log Revenue per Worker in Firm's Product Lines ++ Controls include (log) firm employment; (log) firm employment squared; Dummies for firm age (<6 years old, 6-10 years old, 11+ years old, with <6 omitted); firm employment growth rate; a dummy for whether the firm is in a high density, high education, and industrially diversified county; and year and quarter job separation dummies.

How Much Do High Payoff Firms Pay for Stars?

The results in Tables 4 and 5 show that firms operating in high variance product classes pay higher wages. How much more do they pay, and to whom?

Table 6 shows the predicted value of earnings for alternative combinations of worker wage classes and firm product market classes. The five columns represent the three wage classes: starting salary for the typical median worker; starting salary for the worker at the 90th wage percentile; experienced-worker salary for the median worker; the experienced-worker salary for the experienced worker in the 90th wage percentile; and in the last column the annual earnings growth rates for median workers. The rows represent product classes across the extremes of variance: the first row is firms in the lowest variance product class (mainframe applications, with a product payoff dispersion of 0.55); the median product class (with a payoff dispersion of 1.00); and the high variance product class (video games, with a product payoff dispersion of 1.33). Thus, the middle row of the table anchors the different earnings measures at the actual median values of our data.

Table 6: Predicted Earnings and Earnings Growth from Table 5 Regressions

	<i>Beginning of Spell Earnings "Starting Salaries"</i>		<i>End of Spell Earnings "Experienced Earnings"</i>		<i>Within-Job Earnings Growth</i>
	Median	90th Percentile	Median	90th Percentile	Median
Low Product Payoff Dispersion Database software: 90/50 Ratio =.55	\$54,394	\$94,845	\$80,285	\$195,047	9.51%
Median Product Payoff Dispersion (90/50 Ratio=1.00)	\$58,000	\$108,000	\$95,000	\$311,000	13.32%
High Product Payoff Dispersion Game/Entertainment Software: 90/50 Ratio 1.31)	\$61,023	\$118,184	\$107,581	\$437,573	15.90%

The predicted values of Table 6 display two pronounced conclusions from the regressions. First, firms operating in high-risk product classes pay more for talent, even in starting salaries. For the median worker, starting salaries rise from \$54,394 to \$61,023 moving from low to high product class dispersion firms. More important, for the high skilled worker in the 90th wage percentile, starting salaries rise from \$94,845 to \$118,184. Second, earnings growth is dramatically higher in firms operating in high dispersion product classes. Experienced workers earn much more in firms operating in higher variance product classes (looking at columns 3 and 4). Higher returns to experience arise because workers who stay with their firm

have strikingly high earnings growth in high variance product classes (looking at column 5). The regressions in Table 5 showed that high variance firms do not reward job hopping, so we don't present simulated wages from the between job earnings growth regressions. The overall conclusion is that loyalty pays, and it pays the most for workers in firms in high variance product classes.

We use the phrase “loyalty pays” to counter graphically the often heard phrase “job hopping pays” to characterize software stars.³¹ A typical perception of the software industry is that stars hop from firm to firm. We find that the typical worker is loyal and is rewarded for that loyalty with higher earnings: in other words, the typical worker stays at least five years, and workers with greater tenure boast much higher potential earnings than workers who hop between jobs. Moreover, the firms that reward such loyalty the most are the very firms that operate in high-risk (and thus high-payoff) product markets. We cannot assess why loyalty pays—it could be teamwork, the firm's protection of its intellectual property, or the long run development duration of projects—but it does pay.

We return finally to the raw data to lend support to our regression results that loyalty pays. Figure 4 divides the source of wage growth for workers into wage growth achieved by moving between jobs (or “job-hopping”) versus wage growth achieved by staying with the same firm and experiencing pay increases (for the within firm pay). For the approximately 4% of the sample who earn over \$1 million in the last period in which we observe them, as we look back over their careers, over 95 percent of their wage growth arose within firms, and less than five percent from movement between firms.³² By contrast, among software workers in the \$50-75K range, the final pay is achieved by a combination of changing jobs and by wage growth when they stay within a firm and experience wage increases.³³

³¹ See Fallick et al. (2005) for results focusing on geographic variation in job hopping among software workers.

³² Our definition of between-firm wage growth is annualized starting compensation minus the ending compensation at the last firm. The starting compensation does not include options granted, so in some sense we could say that we are underestimating the gains to job-hopping if software workers are moving between firms to achieve higher option grants. Nevertheless, our key point is that options granted are not yet compensation – the individual must stay with the firm four years (typically) to vest the options granted and the options must be “in-the-money” as a result of performance. Thus, even if options are granted with job change, the pay is only realized from within firm pay increases – the person must stay and perform.

³³ See Appendix 3 for more analysis of wage growth as a function of the number of different employers.

Figure 4

Thus, the striking result from Figure 4 is that even within the software industry, workers earn more from loyalty to their employer. That is, by far the greatest wage gains come not from hopping between employers, but rather from staying with an employer and earning higher pay over time. The figure corroborates our regression results: high wage growth arises for workers in high payoff product markets when they stay with their current employer—for these workers, the return to loyalty can be very high.

The Variance of Pay Within Firms

Given that high product variance firms pay their stars very well, an auxiliary set of questions emerge along the following lines: is the variance of earnings greater within firms operating in high variance product classes? Because we have earnings data on all individuals within all our firms, we are uniquely able to answer this question with the data.

There is no clear-cut theoretical answer to the question.³⁴ Consider the relationship:

$$(3) \quad \sigma_{jt}^W = \theta_0 + \theta_1(\sigma_{jt}^P) + v_{jt}$$

where σ_{jt}^W is the within firm variance in wages for firm j . In our data (and all other data sets), the earnings variance rises within firms with workers' tenure levels. The question is whether the variance of wages is greater for firms in high variance payoff markets (a positive θ_1) than for firms in low variance payoff markets. The answer depends on the nature of the production

³⁴ For a model of the segregation of jobs into high-skill firms and low-skill firms and implications for wage inequality, see Kremer and Maskin (forthcoming).

function within firms. Do firms that have a high payoff, σ^p_{jt} a) want to employ only stars within the firm, so all are high paying (i.e., $\theta_1 < 0$ or $\theta_1 = 0$), or do firms that a high payoff b) want to employ some stars than other firms have (i.e., $\theta_1 > 0$)? The theoretical answer depends upon the complementarity between stars and non-stars. In addition, if a substantial part of the compensation is group-based pay (as in bonuses), then $\theta_1 < 0$.

Our data provide the answer. For the most part, earnings inequality is greater within firms operating in high variance product classes. In Table 7, we use the individual earnings data, but with a dependent variable that is each worker's experienced earnings minus the median experienced earnings in his firm.³⁵ The product payoff dispersion has a very strong positive effect (θ_1) on earnings inequality, and the effect grows with firm size. The rising effect with firm size is not surprising; it is well known that executive pay rises with firm size, as CEO's in larger firms control more capital. What we show is that large firms operating in high payoff product markets have the highest within firm earnings inequality.

Table 7: Regressions of Individual Wage Levels minus Firm-specific Median Wages

	<i>All Firms</i>		<i>Small Firms</i>		<i>Large Firms</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Product Payoff Dispersion+	0.3738 (0.0385)**	0.2912 (0.0388)**	0.1405 (0.0440)**	0.1495 (0.0430)**	0.8989 (0.0731)**	0.4256 (0.0972)**
Log Revenue per Worker	0.1071 (0.0048)**	0.0357 (0.0068)**	-0.0042 (0.0098)	0.0010 (0.0103)	0.1187 (0.0063)**	0.0634 (0.0111)**
Firm Average Worker Churn	-0.3363 (0.1011)**	-0.6089 (0.1200)**	0.3485 (0.1134)**	-0.2093 (0.1358)	-1.0988 (0.1811)**	-1.3821 (0.2729)**
Controls++	No	Yes	No	Yes	No	Yes
Observations	26276	26276	7840	7840	18436	18436

All dependent variables are wage residuals from regression controlling for age and tenure (see Appendix 2). Standard errors in parentheses. * significant at 5%; ** significant at 1%.

+ Weighted 90/50 Ratio of Log Revenue per Worker in Firm's Product Lines ++ Controls include (log) firm employment; (log) firm employment squared; dummies for firm age (<6 years old, 6-10 years old, 11+ years old, with <6 omitted); firm employment growth rate; a dummy for whether the firm is in a high density, high education, and industrially diversified county; and year and quarter job separation dummies.

These results are largely replicated in Table 8, when the dependent variable is the 90/50 ratio of the earnings within the firm. The disadvantage of this approach is that the sample size decreases substantially when the unit of observation is the firm, for 688 firms, and thus we are losing the information on every person's position in the earnings distribution. The estimated θ_1 is significantly positive for starting salaries. For experienced earnings, the θ_1 is typically twice as

³⁵ In Tables 7 and 8, we continue to restrict the sample to the same earnings (\$50K+) and age categories in all of the prior analysis. Thus, our findings on dispersion of earnings within the firm reflect the dispersion of earnings in the upper tail of the distribution of earnings (i.e., minimum earnings is \$50K).

large, but less precisely estimated. We do not detect firm size differences in θ_1 in this smaller data set. However, most notably, in the lower panel of Table 8, we show that when we take most stock options out of the data, the θ_1 remains positive.

Table 8: Within-Firm Earnings Residual Dispersion Regressions

	<i>Beginning of Spell Earnings "Starting Salaries"</i>		<i>End of Spell Earnings "Experienced Earnings"</i>		<i>One-Year Prior Earnings "Lagged Earnings"</i>	
	(1)	(2)	(3)	(4)	(5)	(6)
Product Payoff Dispersion +	0.3060 (0.1084)**	0.2985 (0.0965)**	0.7685 (0.5585)	0.7178 (0.2705)**	0.5273 (0.3133)	0.4582 (0.1985)*
Log Sales per Worker	0.0707 (0.0235)**	0.0206 (0.0287)	0.6225 (0.1613)**	0.1183 (0.0954)	0.3006 (0.0725)**	0.1069 (0.0579)
Firm Average Worker Churn	0.2521 (0.7097)	0.6981 (0.4461)	-1.5985 (4.1871)	2.8754 (1.7730)	0.9884 (2.1438)	2.8380 (1.2434)*
Constant	-0.1704 (0.1842)	-0.0426 (0.1876)	-2.9961 (1.0964)**	-0.2612 (0.6083)	-1.3989 (0.5538)*	-0.5440 (0.3848)
Controls++	No	Yes	No	Yes	No	Yes
Observations	688	688	688	688	688	688
R-squared	0.26	0.36	0.67	0.89	0.59	0.76

All dependent variables are wage residuals from regression controlling for age and tenure (see Appendix 2).

Standard errors in parentheses. * significant at 5%; ** significant at 1%.

+ Weighted 90/50 Ratio of Log Revenue per Worker in Firm's Product Lines ++ Controls include (log) firm employment; (log) firm employment squared; dummies for firm age (<6 years old, 6-10 years old, 11+ years old, with <6 omitted); firm employment growth rate; a dummy for whether the firm is in a high density, high education, and industrially diversified county; and year and quarter job separation dummies.

Finally, consider how these results fit into the picture portrayed by our earlier earnings regression results. The earnings regression results presented in Tables 4 and 5 provide an explanation for the increasing variance of wages over time across firms in the economy and within sectors and occupations: the increasing movement of the economy towards knowledge workers has increased the value of stars to firms, and thus increased the variance of pay. The variance of pay across all workers is rising due to the nature of the production function (the need for stars), and the sorting and rewarding of stars to high payoff firms. The within-firm regression results of Tables 7 and 8 show that the variance of within-firm earnings is also a function of the variance of the product market payoffs that firms face.

5. Conclusion

The process of innovation in the U.S. economy is fundamentally dependent on employing and rewarding highly talented workers to produce new products. This paper makes a key connection between talent and firms in markets with risky product innovations – we show that

software firms that operate in product markets with highly skewed returns to innovation, or high variance payoffs, are more likely to attract and pay for star workers. Thus, firms in high variance product markets pay more up-front—in starting salaries—to attract and motivate star employees, because if these star workers produce home-run innovations, the firm’s winnings will be huge. These same firms also pay highly for loyalty: employees that stay with a firm have much higher earnings in firms with high variance product market payoffs. These striking effects are robust to the inclusion of a wide range of controls for both workers and firm characteristics. In addition, we show that in firms that have actually hit home runs, with high revenues, the rewards for star talent are even greater.

There may be many factors underlying the finding that loyalty is rewarded more in firms operating in risky product markets. In these firms, workers are likely to have firm-specific human capital, and also worker quality is identified and rewarded with performance bonuses slowly on the job. However, much of their human capital is likely to be product-specific (such as a knowledge of wireless software), not firm specific. Thus, firms tie workers to their firms through deferred compensation – they offer options that become vested only after the employee is with the firm four years. And clearly the deferred compensation is performance pay – stock options pay off only if the firm performs. However, our model emphasizes that firms pay high wages to employees just to be in the game – to play in the market where there are potential big winners – whether the firm actually wins or not.

Though the focus of this paper is on the software industry, our results should generalize to other industries that utilize knowledge workers. There were very pragmatic reasons for focusing on the software industry: the ability to examine one production function as well as use a dataset that matches data on software workers to data on software firms. In addition, firms in the software industry have several attractive features: firms are innovative, and hence have the greatest payoff potential from star talent; the activities of knowledge workers can be directly tied to the performance of the firm; and workers perform the same tasks within and across firms, so we can describe and model the production function for these workers. However, all industries that value project development that pays off should pay high wages to attract the most talented employees.

Overall, our results, which document a link between income variance and innovation, complement results in the literature on rising income inequality, changing skill demand, and

economic growth. The income inequality literature has shown that there has been an increasing return to skill both within occupations and industries as well as across occupations and industries. Autor et al. (2005, 2006) and Lemieux (2006) show that rising income inequality reflects an increasing polarization of the job market, noting that income inequality in the upper tail of the income distribution has grown markedly: wages for workers at the 90th percentile of the wage distribution have risen dramatically since 1980 relative to wages of workers at the median, at the same time that there has been no increase in the pay differential between workers at the median and the low wage end at the 10th percentile of the wage distribution.

The rising polarization of the wage distribution suggests that there has been an increase in the demand for highly skilled workers. Those workers in occupations requiring high educational attainment, and in particular those workers in occupations that would be classified as “knowledge” occupations, have seen the greatest growth in demand (Autor et al. 2005, 2006; Autor et al. 2003). Our focus on software therefore provides insights into the wage determination and demand patterns of an industry that may shed light into wage changes occurring in other industries or occupations. The findings are of particular interest since workers in the industry tend to have the same basic education levels and tend to work in firms with similar basic production functions. Therefore, unlike previous studies, we are also able to look very closely at the product demand side of the firm and tie it to skill demand. Our conclusion is that software firms with the highest potential return to innovation activities by skilled workers are the very firms that are increasing their demand for very highly talented workers and are paying these high talent workers much more. The high pay that innovating firms offer top knowledge workers increases the variance of pay in software — both across firms and within firms. But we interpret this increasing income inequality within software as a positive force. The firms paying for high talent among software workers are aiming to increase the skills and effort of software workers. Finally, it is also likely to be a positive force for economic growth. Our data show that firms operating in the most innovative software areas — those firms in product markets that have very high variance value-added payoffs — are those that demand and pay for top talent. Firms such as these, which focus on innovations, have certainly been a source of growth in the U.S. economy in recent years (as in the computer industry), and are likely to be among the biggest drivers of growth in the future.

References

- Abowd, John, John Haltiwanger and Julia Lane (2004). "Integrated Longitudinal Employee-Employer Data for the United States." *American Economic Review*. 224-229.
- Aggarwal, Rajesh K. and Samwick, Andrew A. (1999). "The Other Side of the Trade-off: The Impact of Risk on Executive Compensation." *Journal of Political Economy*. 107(1), 65-105.
- Autor, David, Lawrence Katz, and Melissa Kearney (2005). "Rising Wage Inequality: The Role of Composition and Prices." NBER Working Paper No. 11628.
- Autor, David, Lawrence Katz, and Melissa Kearney (2006). "The Polarization of the U.S. Labor Market." NBER Working Paper No. 11986.
- Autor, David, Frank Levy, and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics*. 118(4). 1279-1333.
- Baker, George P. and Brian J. Hall (2004). "CEO Incentives and Firm Size." *Journal of Labor Economics*. 22(4). 767-798.
- Baker, George P. and Thomas N. Hubbard (2003). "Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information." *American Economic Review*. 23(3). 551-572.
- Buchinsky, Moshe (1994). "Changes in the U.S. Wage Structure 1963-1987: Application of Quantile Regression." *Econometrica*. 62(2). 405-58.
- Chevalier, Judith and Glenn Ellison (1999). "Career Concerns of Mutual Fund Managers." *Quarterly Journal of Economics*. 114(2). 389-432.
- Core, J.E., W. Guay, and David F. Larcker (2003). "Executive Equity Compensation and Incentives: A Survey." *FRBNY Economic Policy Research*. 9.
- Cusumano, Michael A and Richard W. Selby (1995). *Microsoft Secrets: How the World's Most Powerful Software Company Creates Technology, Shapes Markets, and Manages People*. The Free Press: New York, NY.
- Fallick, Bruce, Charles A. Fleischman, and James B. Rebitzer (2005). "Job Hopping in Silicon Valley: The Microfoundations of a High-Technology Cluster." Case Western Reserve University Working Paper.
- Garicano, Luis and Thomas Hubbard (2005). "Managerial Leverage is Limited By the Extent of the Market: Hierarchies, Specialization and the Utilization of Lawyers' Human Capital." CEPR Discussion Paper No. 4924.

- Gibbons, Robert and Kevin J. Murphy (1992). "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence." *Journal of Political Economy*. 100(3). 468-505.
- Hallock, Kevin, Regina Madalozzo, and Clayton Reck (2004). "Uncovering Heterogeneity in Managerial Pay: Firm Performance Relationships Using Quantile Regression." Cornell University ILR School Working Paper.
- Hallock, Kevin and Kevin J. Murphy (1999). *The Economics of Executive Compensation, Volumes I and II*. Edward Elgar: Cheltenham, England.
- Hoch, Detlev J., Cyriac R. Roeding, Gert Purkert, and Sandro K. Lindner (2000). *Secrets of Software Success: Management Insights from 100 Software Firms around the World*. Harvard Business School Press: Boston, MA.
- Ittner, Christopher D., Richard A. Lambert, and David F. Larcker (2003). "The structure and performance consequences of equity grants to employees of new economy firms." *Journal of Accounting and Economics*. 34. 89-127.
- Jensen, Michael C. and Kevin J. Murphy (1990). "Performance Pay and Top-Management Incentives." *Journal of Political Economy*. 98(2). 225-264.
- Kremer, Michael and Eric Maskin. (forthcoming). "Wage Inequality and Segregation by Skill." *Quarterly Journal of Economics*.
- Lazear, Edward P. and Sherwin Rosen (1981). "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy*. 89(5). 841-64.
- Lazear, Edward P. (2005). "Output-Based Pay: Incentives or Sorting?" *Research in Labor Economics*. 23. 1-25.
- Lemieux, Thomas (2006). "Post-Secondary Education and Increasing Wage Inequality." University of British Columbia Working Paper.
- Lerner, Josh and Julie Wulf (2005). "Innovation and Incentives: Evidence from Corporate R&D." Wharton School Working Paper.
- MacLeod, W. Bentley and Daniel Parent (1999). "Job Characteristics and the Form of Compensation." *Research in Labor Economics*. 18. 177-242.
- Murphy, Kevin J. (1986). "Incentives, Learning, and Compensation: A Theoretical and Empirical Investigation of Managerial Labor Contracts." *RAND Journal of Economics*. 17(1). 59-76.
- Murphy, Kevin J. (1999). "Pay, Performance and Executive Compensation." Eds. Orley Ashenfelter and David Card. *Handbook of Labor Economics*. 3. North Holland: New York, NY.

- Oyer, Paul and Scott Schaeffer (2002). "Why Do Some Firms Give Stock Options to All Employees? An Empirical Examination of Alternative Theories." Stanford Research Paper No. 1772.
- Prendergast, Canice (2000). "What Trade-Off of Risk and Incentives?" *American Economic Review*. 90(2). 421-425.
- Prendergast, Canice (2002). "The Tenuous Trade-Off between Risk and Incentives." *Journal of Political Economy*. 110(5). 1071-1102.
- Russell, Kelly Danielle (2005). "The Relationship between Deferred Compensation and Employee Retention." Stanford Research Paper.
- Schaefer, Scott (1998). "The Dependence of Pay-Performance Sensitivity on the Size of the Firm." *Review of Economics and Statistics*. 80(3). 436-443.
- Stern, Scott (2004). "Do Scientists Pay to Be Scientists?" *Management Science*. 50(6). 835-853.
- Stevens, David (2002). "Employment that is not covered by state unemployment insurance laws." LEHD TP 2002-16. U.S. Census Bureau.
- Stross, Randall E. (1997). *The Microsoft Way: The Real Story of How the Company Outsmarts Its Competition*. Perseus Books Group: New York, New York.
- Wulf, Julie (2002). "Internal Capital Markets and Firm-Level Compensation Incentives for Division Managers." *Journal of Labor Economics*. 20(2). S219-S262.
- Wulf, Julie (2005). "Authority, Risk and Incentives: Evidence from Divisional Manager Positions Inside Firms." Wharton School Working Paper.

Appendix 1: Construction of the Data

Sample

The Longitudinal Employer-Household Dynamics (LEHD) Program database links state level data on Unemployment Insurance earnings of all employees within firms to employer data from Census surveys (Abowd et al. 2004).³⁶ Since the scope of the LEHD data is virtually the full universe of employers and workers, movements of workers through the earnings distribution as well as across employers can be tracked accurately.³⁷ Because these data are administrative in nature, both the employment and earnings measures are different from those usually found in surveys. The information in each wage record is simply the total earnings for each individual in a given quarter with a given employer; there is no information on hours or weeks worked, or indeed the duration of employment within the quarter. One key advantage of this administrative data is that the earnings measures are quarterly and include bonuses and exercised stock options.³⁸

For the analysis of compensation over workers' careers, we use a subset of ten states for which we have data for a sufficiently large number of years that we can construct worker earnings histories from which to make meaningful inferences. We also limit our data to workers between the ages 21 and 44, which permits us to focus on the demand for workers in the prime of their careers. We use earnings for individuals who were full-quarter employed in the final quarter of 1997 and whose dominant employer (i.e., the employer at which the worker earned the most in a given quarter) was in the software industry.³⁹ We focus on software-industry spells that span 1997 because software firms in existence that year are most readily matched to the 1997 Economic Census, which contains extensive information about businesses in the industry. We then construct complete employment and earnings histories for individuals in these firms, building backwards from their 1997 software spell and examining prior jobs, including those within the software industry and those in other industries. These distributions helped to inform our decision to set the \$50,000 (in \$2001) threshold for our primary sample of workers.

Using firm identifiers on the UI data, we match LEHD data to the Economic Census. The Economic Census of the services industry is conducted by the U.S. Census Bureau in years ending in "2" and "7." We use the 1997 Economic Census because it characterizes the products

³⁶ As of mid-2006, the LEHD had 40 partner states. This is an ongoing project and additional states are expected to join this program. Because of the sensitivity of these data it is worth noting that the data are anonymized before they are used in any Census Bureau projects; all standard identifiers and names are stripped and replaced by a unique "Protected Identification Key." Only Census Bureau employees or individuals who have Special Sworn Status are permitted to work with the data, and there are serious penalties for disclosing the identity of an individual or business. Any research must be for statistical purposes only, and must be reviewed by the Census Bureau and other data custodians. Under Title 13 of the U.S. code, any breach of confidentiality can result in prosecution in which violators are subject to a \$250,000 fine and/or 5 years in jail.

³⁷ Stevens (2002) describes coverage issues related to the LEHD database.

³⁸ For the laws surrounding the reporting of options, see the example from the California Employment Development Department at <http://www.edd.ca.gov/taxrep/de231sk.pdf>. For analysis of options granted and data available on option values, see Oyer and Schaeffer (2004).

³⁹ Due to the inability to capture hours or weeks in the data, we use a full quarter (FQ) earnings measure. This measure represents earnings for workers who have been employed by the same employer for a full quarter; that is, it represents earnings for a worker who is observed at a firm in quarter t , $t-1$ and $t+1$. While this does not rule out part-time work, it does rule out obviously truncated quarters.

produced by firms in the software industry at a level of detail greater than in any prior year. Unfortunately, products cannot be matched across Census years.

There were 83,497 job spells in software publishing that were ongoing in the fourth quarter of 1997 in the ten states in our sample.⁴⁰ When we condition on workers aged 21-44, the number of spells in the sample falls to 67,452. Finally, because we do not have occupational data in the full LEHD data set but would like to exclude administrative staff and other ancillary workers in the industry, we decide to limit our analysis to only spells that had annualized end-of-spell earnings greater than or equal to \$50,000 (in \$2001). This restriction, which limits the sample to 51,859 spells, is aimed at isolating computer programmers, developers, and managers while weeding out lower administrative and support staff. The \$50,000 threshold was set in light the results of two separate, but related exercises. We first turned to the 5% Public-Use Microdata Sample (PUMS) from the 2000 Decennial Census to get a sense of the distribution of earnings within the relevant set of software occupations (programmers, developers, engineers, and managers in the software industry).⁴¹ ⁴² We also matched our sample of workers with the Long-Form of the 2000 Decennial Census, which provided us with occupation information for slightly less than one-sixth of our sample. We then evaluated the earnings of those workers for whom we obtained occupational information. These exercises suggest that the \$50K threshold is quite reasonable for focusing on software developers, engineers, and managers. We also note that for the core specifications we consider in the empirical analysis, we have estimated the specifications using the restricted matched sample where we directly observe occupation. In unreported results, we show that the main implications of the analysis are robust to the use of this restricted sample.

While most businesses in our \$50,000+ sample of workers aged 21-44 could be successfully matched to the Economic Census, a smaller subset had complete information, including size, age, revenue, and detailed product line information. There are 26,276 spells for which we have complete information regarding worker characteristics as well as firm characteristics. All told, 688 firms are represented in this sample.

Measurement

The firm-specific potential payoff measure is computed as follows. Beginning with establishment-level Economic Census data, we distribute each establishment's total sales revenue across its product lines according to the reported percentages. We then aggregate to the EIN-product line level, summing sales across establishments within this EIN-product line. Next, we aggregate to the EIN level since this is the common firm identifier in both the Economic Censuses and in the LEHD data. Total sales are summed by EIN, excluding the "other" product line categories. New percentages are then calculated (adjusted so they do not include "other" categories) at the EIN-product line level, and the payoff measure, which we calculate as the natural log of sales per worker, is distributed across all product lines within an EIN (after

⁴⁰ Counting the 1997 software spells and all the previous spells held by workers in these software jobs, we have 143,485 spells in the data.

⁴¹ The primary occupations on which we focused included Census industry occupation codes 100 (Computer and Information Scientists, Research), 101 (Computer Programmers), and 102 (Computer Software Engineers, Applications and Systems Software), as well as 001-043 (managerial occupations).

⁴² We also use information from the Decennial to obtain information regarding basic characteristics of the counties in which people in the sample are employed, including population and employment densities, as well as average educational and income levels. These enter the control variables in the regressions.

adjusting EIN total employment downward by the same percentage that total employer sales were adjusted downward due to the exclusion of the "other" categories). This number is the same across all product lines within an employer (we thus assume that profitability is the same across all product lines within an employer). Then, treating each EIN-product line combination as if it were its own establishment, we obtain the difference between 90th percentile and 50th percentile of log sales per worker by product line. Finally, these 90-50 ratios are merged back onto the EIN-product line file, and the weighted average of the 90-50 ratios associated with all the product lines within an employer is what is used as our measure of potential payoffs.

By using the 90/50 ratio, we focus on the upper tail of the payoff distribution. It is highly correlated with other measures of skewness. However, note that we do not have data on the lower tail of the distribution – we only observe positive revenue for the firms, not negative revenue – so we focus on the returns to high payoffs.

Appendix 2: Calculation of Residuals

Using the worker-spells data, we wish to characterize what it means to be a star worker – a worker whose wage level and/or wage growth rate is abnormally high. We do so by deriving residuals from regressions of earnings levels and growth rates; the purpose of using residuals is to find those individuals who have abnormally high earnings levels and/or earnings growth after controlling for their age and experience.

We begin by computing earnings level residuals for individuals in the sample. We compute the end-of-software-spell earnings level residual for each spell in the sample as the residual from a regression of log end-of-spell earnings on quadratics of tenure at job, tenure in industry, and age, fully interacted with each other and with appropriate left and right censoring dummies. In a similar fashion, we compute beginning-of-software-spell earnings level residuals and lagged end-of-software-spell earnings level residuals.

We then compute within job earnings growth residuals and between job earnings growth residuals for individuals in the sample. The former is the residual from a regression of within job earnings growth for each spell in the sample on quadratics of tenure at job and age, fully interacted with each other and with appropriate left and right censoring dummies. The latter is the residual from a regression of between job earnings growth for each transition on tenure at (previous) job and age, fully interacted with each other and with appropriate left and right censoring dummies as well as dummies for a switch within industry and a switch between industries. We then form an earnings growth residual for each individual that is simply the sum of within job and between job residuals over their work histories.

Table A1: Log End-of-Spell Annualized Earnings

Dependent Variable: Log End-of-Spell Annualized Earnings				
Observations	143,485			
Parameter	Estimate	Standard Error	t Value	Pr> t
Intercept	6.4046	0.0347	184.4000	<.0001
Tenure in Industry	0.0545	0.0039	14.1300	<.0001
Tenure in Industry Squared	-0.0010	0.0001	-12.0500	<.0001
Tenure in Job	0.0992	0.0040	25.0300	<.0001
Tenure in Job Squared	-0.0003	0.0001	-3.8800	0.0001
Age	0.1694	0.0017	97.5500	<.0001
Age Squared	-0.0016	0.0000	-75.0000	<.0001
Right Censored Spell	0.8766	0.0687	12.7600	<.0001
Left Censored Spell	-0.1425	0.0133	-10.7300	<.0001
Tenure in Industry x Left Censored Industry Spell	-0.0144	0.0039	-3.6800	0.0002
Tenure in Industry Squared x Left Censored Industry Spell	0.0002	0.0001	3.0300	0.0025
Tenure in Job x Left Censored Job Spell	-0.0187	0.0044	-4.2200	<.0001
Tenure in Job Squared x Left Censored Job Spell	0.0003	0.0001	3.1900	0.0014
Tenure in Industry x Right Censored Industry Spell	-0.0269	0.0219	-1.2200	0.2207
Tenure in Industry Squared x Right Censored Industry Spell	0.0003	0.0003	0.9000	0.3691
Tenure in Job x Right Censored Job Spell	-0.0418	0.0219	-1.9100	0.0564
Tenure in Job Squared x Right Censored Job Spell	0.0011	0.0003	3.2000	0.0014
Tenure in Industry x Age	-0.0004	0.0001	-4.1300	<.0001
Tenure in Job x Age	-0.0016	0.0001	-17.6900	<.0001
Tenure in Industry x Age x Left Censored Industry Spell	0.0004	0.0001	4.4600	<.0001
Tenure in Job x Age x Left Censored Job Spell	0.0002	0.0001	2.1300	0.0336
Tenure in Industry x Age x Right Censored Industry Spell	0.0001	0.0003	0.3600	0.7208
Tenure in Job x Age x Right Censored Job Spell	-0.0003	0.0003	-0.8600	0.3913
Separated Q1	0.0378	0.0072	5.2900	<.0001
Separated Q3	0.0595	0.0070	8.4700	<.0001
Separated Q4	0.0436	0.0074	5.8700	<.0001
		Sum of		
Source	DF	Squares	Mean Square	F Value
Model	25	43,074.80	1,722.99	2,309.20
Error	143,459	107,040.92	0.7461	
Corrected Total	143,484	150,115.72		
	R-Square	Coeff Var	Root MSE	LHS Variable Mean
	0.286944	7.844303	0.863796	11.01176

Table A2: Log Beginning-of-Spell Annualized Earnings

Dependent Variable: Log Beginning-of-Spell Annualized Earnings				
Observations	143,485			
Parameter	Estimate	Standard Error	t Value	Pr> t
Intercept	6.8097	0.0257	265.3500	<.0001
Age	0.1829	0.0013	138.2700	<.0001
Age Squared	-0.0019	0.0000	-117.1200	<.0001
Left Censored Spell	-0.1313	0.0052	-25.0700	<.0001
Accessed Q1	-0.0862	0.0057	-15.1600	<.0001
Accessed Q3	0.0134	0.0054	2.4900	0.0128
Accessed Q4	-0.0296	0.0058	-5.0800	<.0001
		Sum of		
Source	DF	Squares	Mean Square	F Value
Model	6	16,191.69	2,698.62	5,946.98
Error	143,478	65,107.30	0.45	
Corrected Total	143,484	81,298.99		
	R-Square	Coeff Var	Root MSE	LHS Variable Mean
	0.199162	6.300298	0.673631	10.69205

Table A3: Lagged Log End-of-Spell Annualized Earnings

Dependent Variable: Lagged Log End-of-Spell Annualized Earnings				
Observations	143,485			
Parameter	Estimate	Standard Error	t Value	Pr> t
Intercept	6.4128	0.0328	195.5500	<.0001
Tenure (Adjusted) in Industry	0.0313	0.0041	7.7200	<.0001
Tenure (Adjusted) in Industry Squared	-0.0005	0.0001	-4.7700	<.0001
Tenure (Adjusted) in Job	0.1602	0.0042	38.2400	<.0001
Tenure (Adjusted) in Job Squared	-0.0023	0.0001	-23.6400	<.0001
Age	0.1683	0.0017	99.9600	<.0001
Age Squared	-0.0017	0.0000	-77.8300	<.0001
Right Censored Spell	0.5539	0.0363	15.2400	<.0001
Left Censored Spell	0.0352	0.0085	4.1500	<.0001
Tenure (Adjusted) in Industry x Left Censored Industry Spell	0.0006	0.0045	0.1300	0.8958
Tenure (Adjusted) in Industry Squared x Left Censored Industry Spell	-0.0003	0.0001	-2.4100	0.0161
Tenure (Adjusted) in Job x Left Censored Job Spell	-0.0777	0.0048	-16.0300	<.0001
Tenure (Adjusted) in Job Squared x Left Censored Job Spell	0.0019	0.0001	17.0800	<.0001
Tenure (Adjusted) in Industry x Right Censored Industry Spell	-0.0367	0.0196	-1.8800	0.0606
Tenure (Adjusted) in Industry Squared x Right Censored Industry Spell	0.0009	0.0003	2.6300	0.0085
Tenure (Adjusted) in Job x Right Censored Job Spell	-0.0271	0.0197	-1.3800	0.1689
Tenure (Adjusted) in Job Squared x Right Censored Job Spell	0.0007	0.0003	2.1500	0.0314
Tenure (Adjusted) in Industry x Age	-0.0003	0.0001	-2.9200	0.0035
Tenure (Adjusted) in Job x Age	-0.0017	0.0001	-18.0700	<.0001
Tenure (Adjusted) in Industry x Age x Left Censored Industry Spell	0.0004	0.0001	3.8300	0.0001
Tenure (Adjusted) in Job x Age x Left Censored Job Spell	0.0005	0.0001	4.4300	<.0001
Tenure (Adjusted) in Industry x Age x Right Censored Industry Spell	-0.0003	0.0003	-1.0500	0.2947
Tenure (Adjusted) in Job x Age x Right Censored Job Spell	0.0003	0.0003	0.8100	0.4184
Quarter 1 Dummy	-0.0231	0.0070	-3.3100	0.0009
Quarter 2 Dummy	0.1038	0.0068	15.1600	<.0001
Quarter 4 Dummy	0.1318	0.0072	18.2100	<.0001
		Sum of		
Source	DF	Squares	Mean Square	F Value
Model	25	53929.872	2157.1949	3041.4
Error	143459	101752.14	0.7093	
Corrected Total	143484	155682.01		
	R-Square	Coeff Var	Root MSE	LHS Variable Mean
	0.34641	7.752834	0.842186	10.86294

Table A4: Annualized Within-Job Earnings Growth

Dependent Variable: Annualized Within-Job Earnings Growth				
Observations	143,485			
Parameter	Estimate	Standard Error	t Value	Pr> t
Intercept	0.0877	0.0122	7.2200	<.0001
Tenure in Job	0.0143	0.0006	23.6300	<.0001
Tenure in Job Squared	-0.0002	0.0000	-12.6600	<.0001
Age	-0.0037	0.0006	-6.2100	<.0001
Age Squared	0.0000	0.0000	4.1400	<.0001
Right Censored Spell	0.3153	0.0234	13.4600	<.0001
Left Censored Spell	0.0097	0.0048	2.0100	0.0449
Tenure in Job x Left Censored Job Spell	-0.0039	0.0008	-4.7900	<.0001
Tenure in Job Squared x Left Censored Job Spell	0.0001	0.0000	4.5900	<.0001
Tenure in Job x Right Censored Job Spell	-0.0283	0.0019	-14.7000	<.0001
Tenure in Job Squared x Right Censored Job Spell	0.0005	0.0000	14.3900	<.0001
Tenure in Job x Age	-0.0001	0.0000	-9.1400	<.0001
Tenure in Job x Age x Left Censored Job Spell	0.0000	0.0000	0.4800	0.6339
Tenure in Job x Age x Right Censored Job Spell	0.0000	0.0000	0.9100	0.3627
Separated Q1	-0.0021	0.0025	-0.8300	0.4091
Separated Q3	0.0242	0.0025	9.8400	<.0001
Separated Q4	0.0263	0.0026	10.1200	<.0001
Accessed Q1	0.0153	0.0026	5.9800	<.0001
Accessed Q3	-0.0049	0.0024	-2.0200	0.0433
Accessed Q4	-0.0200	0.0026	-7.6000	<.0001
		Sum of		
Source	DF	Squares	Mean Square	F Value
Model	19	310.60719	16.34775	179.97
Error	143,465	13,031.52	0.09083	
Corrected Total	143,484	13,342.13		
	R-Square	Coeff Var	Root MSE	LHS Variable Mean
	0.02328	399.5266	0.301387	0.075436

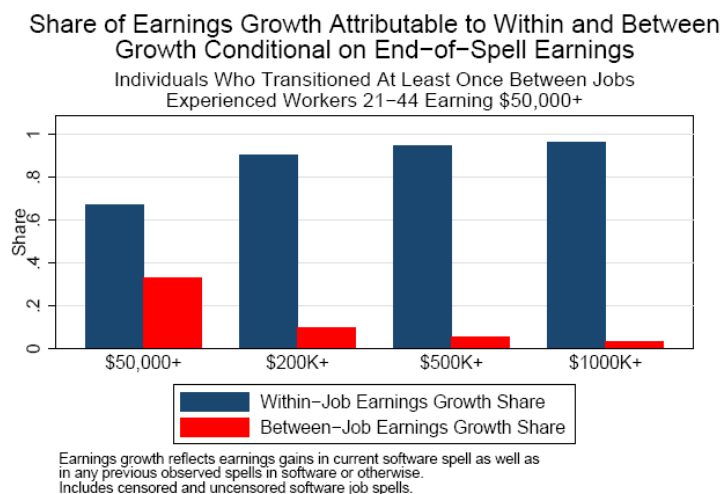
Table A5: Annualized Between-Job Earnings Growth

Dependent Variable: Annualized Between-Job Earnings Growth				
Observations	59,538			
Parameter	Estimate	Standard Error	t Value	Pr> t
Intercept	1.6536	0.0382	43.3100	<.0001
Age	-0.0588	0.0019	-30.4600	<.0001
Age Squared	0.0006	0.0000	23.6500	<.0001
Tenure in Previous Job	-0.0491	0.0045	-10.8900	<.0001
Tenure in Previous Job Squared	0.0016	0.0002	7.2700	<.0001
Left Censored Tenure in Previous Job	-0.1683	0.0130	-12.9400	<.0001
Tenure in Previous Job x Left Censored Previous Job Spell	0.0482	0.0054	8.9400	<.0001
Tenure in Previous Job Squared x Left Censored Previous Job Spell	-0.0012	0.0003	-4.2500	<.0001
Age x Tenure in Previous Job	0.0005	0.0001	4.8000	<.0001
Age x Tenure in Previous Job x Left Censored Previous Job Spell	-0.0007	0.0001	-7.1100	<.0001
Switch Within Industries	-0.1002	0.0068	-14.8200	<.0001
Switch Between Industries	(omitted)	.	.	.
Separated Q1	-0.0132	0.0073	-1.8000	0.0716
Separated Q3	-0.0038	0.0063	-0.6100	0.5427
Separated Q4	-0.0253	0.0078	-3.2400	0.0012
Accessed Q1	0.0051	0.0066	0.7700	0.4411
Accessed Q3	-0.0049	0.0064	-0.7600	0.4463
Accessed Q4	0.0712	0.0071	10.1000	<.0001
		Sum of		
Source	DF	Squares	Mean Square	F Value
Model	16	1,581.45	98.84088	293.17
Error	59,521	20,066.95	0.33714	
Corrected Total	59,537	21,648.40		
	R-Square	Coeff Var	Root MSE	LHS Variable Mean
	0.073052	294.0657	0.580638	0.197452

Appendix 3: Supplementary Analysis of Wage Growth

As Figure 5 in Section 5 suggests, individuals with relatively high earnings levels who were in the midst of software spells in 1997 experienced much of their earnings growth through within-job earnings growth over the course of their entire employment histories including the software spell and prior spells. However, one might be suspicious of these results given that some spells may be censored and that those who are not observed to transition between jobs necessarily have shares attributable to within-job earnings growth equal to one. Yet, as Figure A1 reveals, even for those who are observed to transition at least once, the share of growth attributable to within-job earnings growth is substantially higher among the highest earners.

Figure A1



The following two figures reinforce this message. They depict the share of earnings growth attributable to within-job earnings growth and between-job earnings growth by number of spells. While the share of earnings growth attributable to within-job earnings growth falls with the number of job transitions observed, those with the highest end-of-spell earnings tend to be individuals observed in relatively few spells.

Figure A2




Figure A3

