

Search and Non-Wage Job Characteristics

Paul Sullivan Ted To

Bureau of Labor Statistics

September 2011

Abstract

This paper quantifies the importance of non-wage job characteristics to workers by estimating a structural on-the-job search model. The model generalizes the standard search framework by allowing workers to search for jobs based on both wages and job-specific non-wage utility flows. Within the structure of the search model, data on accepted wages and wage changes at job transitions identify the importance of non-wage utility through revealed preference. The parameters of the model are estimated by simulated minimum distance using the 1997 cohort of the National Longitudinal Survey of Youth (NLSY97). The estimates reveal that utility from non-wage job characteristics plays an important role in determining job mobility, the value of jobs to workers, and the gains from job search. More specifically, non-wage utility accounts for approximately one-third of the total gains from job mobility. These large non-pecuniary gains from search are missed by search models which assume that the wage captures the entire value of a job to a worker.

Keywords: job search, non-wage job characteristics, wage growth, revealed preference, compensating differentials

JEL codes: D1, D9, J4, J6

1 Introduction

Non-wage job characteristics are important determinants of job mobility and choice. Important non-wage job characteristics include employer provided health insurance (Gruber and Madrian, 2004), employer provided retirement benefits, flexible hours (Altonji and Paxson, 1992), paid vacation, occupational choice (Goddeeris, 1988), risk of injury or death (Thaler and Rosen, 1975), commuting time (White, 1988), on-site amenities, or a whole host of other, possibly intangible or heterogeneously valued,¹ job characteristics. Despite their importance, there is relatively little research that estimates search models with non-wage job characteristics and studies their effect on job choice and mobility decisions. The bulk of the the empirical search literature assumes that the wage captures the entire value of a job and the literature that does account for non-wage job characteristics typically focuses on a single job characteristic. For example, Blau (1991), Bloemen (2008), Flabbi and Moro (2010) and Gørgens (2002) estimate models with hours or hours flexibility, Dey and Flinn (2005, 2008) estimate models with health insurance provision, and Sullivan (2010) estimates a model with occupational choice. Instead of focusing on a single observable job characteristic, we estimate a structural search model that allows workers to derive utility from their aggregate valuation of all the non-wage characteristics of a particular job.

The goals of this paper are to estimate the total value that workers place on the non-wage attributes of their jobs and to quantify the importance of non-wage factors in determining individual labor market dynamics. To accomplish this goal, we estimate a search model which augments the standard income maximizing on-the-job search framework (Burdett, 1978) by including utility from non-wage job characteristics. In the model, employed and unemployed workers search across jobs that offer different wages and levels of non-wage utility. When a worker and firm meet, the worker receives a wage offer and also observes a match-specific non-wage utility flow which represents the net value that this particular worker places on all the non-wage job characteristics present at the job. Search frictions are present because both

¹See Bhaskar and To (1999) and Bhaskar et al. (2002).

job offers and layoffs occur randomly, and because both wages and non-wage match values are modeled as random draws from a distribution that is known to the worker. Following a large fraction of the empirical search literature, we adopt a stationary, partial equilibrium framework.² As in the canonical on-the-job search model, wage growth occurs as workers climb a job ladder by moving to higher wage jobs. A novel feature of the model is that it also allows workers to benefit from moving to jobs that offer higher non-wage utility. Depending on the importance of the non-wage side of the model, basing conclusions about the value of job mobility solely on wages could give a misleading view of the gains to job search and mobility. Estimating the structural model is a direct way of quantifying the importance of the wage and non-wage channels in determining the total gains to mobility over the career.

The structural parameters are estimated by simulated minimum distance (SMD) using the 1997 cohort of the National Longitudinal Survey of Youth (NLSY97). The estimates reveal that workers place a substantial value on non-wage job characteristics, and also show that non-wage utility flows vary widely across different worker-firm matches. More specifically, workers who are searching for a job face nearly as much dispersion in job-specific non-wage utility flows as in wage offers. Simulations performed using the estimated model reveal that increases in the utility derived from non-wage job characteristics account for nearly one-third of the total gains from job mobility. This result indicates that standard models of on-the-job search – which are based solely on wages – are missing a key determinant of the value of jobs, the causes of worker mobility, and the gains from job search.

Our use of the non-wage match value as an aggregate measure of the non-wage value of a job is primarily motivated by the goal of estimating the total non-wage value of jobs to workers. In addition, four observations about the information available in standard sources of labor market data on the employer provided benefits, tangible job characteristics, and intangible job characteristics that differentiate jobs are relevant. First, important employer provided benefits such as health insurance and retirement plans are imperfectly measured.³

²See, for example, Flinn (2002), Jolivet et al. (2006), Bloemen (2008) and Dey and Flinn (2008).

³For example, in the 1979 and 1997 cohorts of the NLSY, information is available about whether or not employers offer benefits such as health insurance, but there is no information about take up of benefits,

Second, information about many tangible job characteristics, such as risk of injury or commuting time, is frequently unavailable. Third, measures of intangible job characteristics such as a worker's evaluation of his supervisor, which may be significant determinants of the value of a job to a worker, are typically completely absent. Fourth, and perhaps most importantly, it is likely that workers have heterogeneous preferences over the employer provided benefits and tangible and intangible job characteristics that differentiate jobs. With these facts in mind, rather than attempting to estimate the value of specific job characteristics, we estimate the net value of all non-wage job characteristics to a worker using the non-wage match value.

This paper contributes to a growing literature that demonstrates the importance of accounting for imperfect information, search frictions, and dynamics when estimating the value of non-wage job characteristics. Hwang et al. (1992), Dey and Flinn (2005, 2008), and Gronberg and Reed (1994) all discuss the problems caused by using a static framework to analyze non-wage job characteristics in a dynamic labor market. Most recently, Bonhomme and Jolivet (2009) estimate the value of a number of observed job characteristics using a search model. We take a different approach by estimating the total non-wage value of jobs using the non-wage match value, rather than attempting to identify the value of specific characteristics. Becker (2010) develops a model that focuses on incorporating non-wage utility into the equilibrium wage bargaining framework of Postel-Vinay and Robin (2002) and applying the model to unemployment insurance.

Non-wage utility flows are, of course, not observed by the econometrician, so identification is an important concern. The on-the-job search model provides a natural framework for using data on wages, job acceptance decisions, and employment durations to infer the value that workers place on non-wage job characteristics. Broadly speaking, the intuition behind the identification of the model is that since a standard income maximizing search model is nested within the utility maximizing search model, the importance of non-wage job characteristics is identified by the extent to which an income maximizing model fails to explain the moments

dollar amount of the employer and employee contributions, or plan quality.

used in estimation. More specifically, observed patterns of job mobility and wage changes at transitions between jobs are particularly informative about the importance of non-wage job characteristics. To give a concrete example, a key moment matched during estimation is the proportion of direct job-to-job transitions where workers choose to accept a decrease in wages.⁴ Wage declines at job transitions occur frequently: in the NLSY97 data, reported wages decline for more than one-third of direct transitions between jobs. Taking the structure of the model as given, this type of transition indicates through revealed preference that a worker is willing to accept lower wages in exchange for higher non-wage utility at a specific job.

During estimation, we are careful to account for the two alternative explanations for observed wage decreases at direct transitions between jobs that have dominated the empirical search literature up to this point. Ignoring either of these possible explanations during estimation would lead to an upward bias in the estimated importance of non-wage utility. The first explanation is that if a job ends exogenously, a worker might choose to move directly to a lower paying job to avoid unemployment, if this option is available.⁵ The second explanation is that measurement error in wages might cause some transitions between jobs that are actually accompanied by wage increases to be erroneously shown as wage decreases in the data.⁶ Our model allows for both of these explanations, and also adds a third possible explanation: a worker could choose to move from a high wage job to a lower wage job that offers a higher level of non-wage utility. We incorporate the first explanation into the model by allowing existing jobs to end involuntarily (from the perspective of the worker) in the same time period that a job offer is received from a new employer. The probability that this event occurs is identified using NLSY97 data that identifies direct job-to-job transitions that begin with involuntary job endings. Existing research has not used this type of data to

⁴Throughout the paper, direct transitions refer to transitions between jobs that occur without an intervening spell of unemployment.

⁵For example, Jolivet et al. (2006) assume that all direct job-to-job transitions accompanied by wage decreases are the result of simultaneous job endings and mobility to new jobs.

⁶Flinn (2002) adopts this approach, and Wolpin (1992) allows for both measurement error in wages and simultaneous exogenous job endings and outside job offers.

identify involuntary direct transitions between employers.

We account for measurement error in wages by estimating a parametric model of measurement error jointly along with the other parameters of the model.⁷ Although at first glance it might appear that measurement error in wages and match-specific non-wage utility are observationally equivalent, Section 4.4 of this paper demonstrates that they actually have very different implications for the simulated moments used to estimate the model. More specifically, although measurement error and non-wage utility can both account for observed wage decreases at job transitions, neither feature on its own is capable of simultaneously explaining the extent of variation in wages, amount of wage growth over the career, and frequency of wage declines at direct job transitions in the NLSY97.

The parameter estimates reveal that the variation in non-wage utility flows across worker-firm matches is nearly as large as the variation in the wage offer distribution. This implies that there are substantial gains to workers from job search based on non-wage factors. Although the parameter estimates provide direct evidence on the importance of the non-wage side of the model, perhaps a more informative way of examining the implications of the utility maximizing search model is to study simulated data generated by the estimated model. In these data, wages explain only 68 percent of the total variation in the one-period utility flows that workers receive from employment, which leaves match specific non-wage utility to account for the remainder. On average, measuring the value of a job using only the wage substantially understates the true value of a job to a worker. More specifically, in 88 percent of all jobs in the simulated data, workers place a positive net value on the non-wage characteristics of their job. The fact that workers receive negative non-wage utility flows in only 12 percent of accepted job offers shows that although utility maximizing workers in the model are perfectly willing to accept higher wages in exchange for undesirable job characteristics, the search model reveals a strong tendency for workers to sort into jobs with non-wage job characteristics that they are willing to pay for. The positive average non-wage value of jobs

⁷We find that measurement error accounts for 12.7 percent of the variation in wage offers, which is reassuringly within the range reported by Bound et al. (2001).

is generated by two features of the search environment. First, the reservation utility strategy followed by unemployed agents implies that the accepted job offers observed in the simulated data are truncated from below. Second, on-the-job search implies that workers climb both wage and non-wage utility ladders as they move between employers.

The search model with non-wage job characteristics has important implications for the study of compensating differentials, a topic of considerable interest to economists. Existing papers such as Hwang et al. (1992) and Gronberg and Reed (1994) make the point that in general, estimates of compensating differentials will be biased unless search frictions are taken into account. Our primary contribution to this line of research is to use the estimated structural search model to obtain a direct estimate of the magnitude of the bias caused by estimating compensating differentials using a static framework. Standard hedonic regression approaches to valuing non-wage job characteristics implicitly assume that workers are free to select an optimal job from a perfectly known labor market hedonic wage curve. In contrast, the simulated data from our model contain a sample of wages and non-wage utility received by workers who must search for jobs in a dynamic labor market. When we estimate a standard hedonic regression using these data, the estimated marginal willingness to pay for non-wage job characteristics is biased downward by approximately 50 percent from the true value used to generate the data.

This application of the model offers an explanation for the fact that empirical support for the theory compensating differentials is relatively weak, despite a vast literature on estimating these differentials. The intuition behind the downward bias in estimated compensating differentials is that in a search model, the only information provided by accepted pairs of wages and non-wage utility is that they exceed a reservation utility threshold. In this setting, they do not directly reveal the marginal willingness to pay for non-wage job characteristics as they would in a static, frictionless, perfect information world where workers maximize utility subject to a given labor market hedonic wage locus.

In the following section, we develop a partial equilibrium model of on-the-job search with non-wage utility. In Section 3, we discuss the data set used to estimate our model and in

Section 4 we discuss our econometric methodology and some important identification issues. Section 5 presents our parameter estimates and discusses the effect of non-wage utility on labor market outcomes. Section 6 concludes.

2 The Search Model with Utility from Non-Wage Job Characteristics

This section presents the search model used to estimate the importance of non-wage utility. The model is set in discrete time. Agents maximize the discounted sum of expected utility over an infinite time horizon in a stationary environment. In each time period, individuals occupy one of two states: employment or unemployment.⁸ Agents randomly receive job offers while unemployed and employed, and the employed face a constant risk of exogenous job loss.⁹ When a job ends exogenously, there is no chance of recall. For ease of exposition, this section describes the decision problem facing a single agent. However, we allow for person-specific unobserved heterogeneity when estimating the model (Section 4).

2.1 Preferences and Job Offers

The utility received by an employed agent is determined by the log-wage, w , and the match-specific non-wage utility flow, ξ . The one-period utility from employment is

$$U(w, \xi) = w + \xi, \tag{1}$$

where both w and ξ are specific to a particular match between a worker and employer, and are constant for the duration of the match. A job offer consists of a random draw of (w, ξ) from the distribution $F(w, \xi)$, which is a primitive of the model.

The structure of the search and matching process in the model labor market is as fol-

⁸Following the majority of the search literature, the model does not distinguish between unemployment and non-participation in the labor market.

⁹The terms exogenous job endings and layoff are used interchangeably in the remainder of the paper.

lows. When a worker and firm randomly meet, the worker receives a job offer. At the same time, the worker observes the complete bundle of non-wage job characteristics present at the firm. These characteristics include employer provided benefits (health insurance), tangible job characteristics (risk of injury, commuting time), and intangible job characteristics (friendliness of co-workers). Based on his preferences, which may be heterogeneous across agents, the worker determines the net value of the non-wage job characteristics present at this firm (ξ). The worker then decides whether or not to accept the job offer. Once a job offer is accepted, the wage and non-wage component of the offer remain constant for the duration of the employment spell.¹⁰ Since our primary goal is to estimate the total importance of non-wage job characteristics to workers, which is captured by ξ , we do not attempt to determine how much of the variation in non-wage utility is due to firm level variation in non-wage job characteristics versus preference heterogeneity.¹¹

2.2 Unemployed Search

Unemployed agents search for jobs, which arrive randomly with probability λ_u . Since w and ξ are additively separable in the utility function, it is convenient to define the agent's decision problem in terms of total utility, $w + \xi$, where $U(w, \xi) \equiv U$ and U is distributed as $H(U)$. The discounted expected value of lifetime utility for an unemployed agent is

$$V^u = b + \delta[\lambda_u E \max\{V^u, V^e(U')\} + (1 - \lambda_u)V^u], \quad (2)$$

where b is the one-period utility flow from unemployment, which reflects the value of unemployment benefits and leisure, and δ is the discount factor. The term $V^e(U')$ represents the expected discounted value of lifetime utility for an agent employed in a job with utility level

¹⁰Section 3 explains that the assumption of constant wages within jobs is consistent with the NLSY97 data used to estimate the model.

¹¹We leave decomposing the sources of variation in match-specific non-wage utility as an interesting, although difficult extension for future research. Empirical work along these lines would require detailed data on the complete set of non-wage job characteristics valued by workers along with sufficiently high mobility rates between jobs with different characteristics to identify preference heterogeneity. In addition, data on the firm side of the market, ideally matched worker-firm data, would be useful to control for unobserved, firm-specific variation in working conditions and job amenities.

U' .

The optimal search strategy for an unemployed agent is a reservation utility strategy, which is analogous to the reservation wage strategy found in income maximizing search models. The rule is to accept any job offer which offers a one-period utility flow greater than the reservation level, U^* , and reject all other offers. Appendix A presents the formal derivation of U^* .

2.3 On-the-job Search

In each time period, with probability λ_e an employed agent receives a job offer from an outside firm. The worker may accept the job offer, or reject it and continue working for his current employer. Job matches end with exogenous probability λ_l . When a job ends for this reason, the worker is forced to become unemployed. With probability λ_{le} , a worker's current job exogenously ends and he receives a job offer from a new employer in the same time period. When this happens, the worker can accept the new offer, or become unemployed. Finally, with probability $(1 - \lambda_e - \lambda_l - \lambda_{le})$ the job does not end exogenously and no new offers are received, so the worker is forced to remain in his current job.

The discounted expected value of lifetime utility for a worker who is currently employed in a job with utility level U is

$$\begin{aligned} V^e(U) = & U + \delta[\lambda_e E \max\{V^e(U), V^e(U')\} + \lambda_l V^u \\ & + \lambda_{le} E \max\{V^u, V^e(U')\} + (1 - \lambda_e - \lambda_l - \lambda_{le})V^e(U)]. \end{aligned} \quad (3)$$

The first bracketed term in equation (3), $\lambda_e E \max\{V^e(U), V^e(U')\}$, represents the expected value of the best option available in the next time period for an employed individual who receives a job offer from a new employer.¹² The second bracketed term, $\lambda_l V^u$, corresponds to

¹²The value function reflects the fact that in this model it is never optimal for a worker to quit a job and enter unemployment.

the case where a job exogenously ends and the worker is forced to enter unemployment. The third bracketed term, $\lambda_{le} \max\{V^u, V^e(U')\}$, represents the case where the worker is laid-off but also receives a job offer from a new employer. The final bracketed term represents the case where the worker is neither laid-off nor receives an outside job offer.

In this stationary search environment, optimal decisions for employed agents are based on comparisons of one-period utility flows. When an employed agent receives an offer from an outside firm but does not experience an exogenous job ending, a simple reservation utility strategy is optimal. Since $V^e(U)$ is increasing in U , the rule is to accept the offer if it offers greater utility than the current job ($U' > U$), and reject the offer otherwise ($U' \leq U$). If a worker's job exogenously ends and he receives a new job offer at the same time, which occurs with probability λ_{le} , the situation is identical to the one faced by an unemployed agent who receives a new job offer. As a result, he will choose to accept or reject the offer based on the unemployed reservation utility level U^* .

In the remainder of the paper, we will refer to direct job-to-job transitions that occur as the result of a simultaneous layoff and job offer as “involuntary” transitions between employers. This terminology reflects the fact that although a direct job-to-job transition occurs, the worker's previous job ended involuntarily (exogenously). For agents in the model, voluntary and involuntary transitions are fundamentally different types of job mobility. When a voluntary job-to-job transition occurs, utility increases ($U' > U$). In contrast, when an involuntary transition occurs, the new job offer is preferable to unemployment ($U' > U^*$), but it may offer lower utility than the previous job which exogenously ended ($U' < U$).

3 Data

For several reasons, we use the 1997 rather than the venerable 1979 cohort of the NLSY to estimate our model. To begin with, the NLSY97 is more representative of current labor market conditions. Since we are only interested in young, inexperienced workers, the NLSY97 now has a long enough time series for our purposes. In addition, the NLSY97 design team

incorporated lessons from the NLSY79 and has a more consistent methodology (Pergamit et al., 2001).

The NLSY97 is a nationally representative sample of 8,984 individuals who were between the ages of 12 and 16 on December 31, 1996. Interviews have been conducted annually since 1997, and we use data from interviews up to the 2006 interview to estimate the model. One component of our estimation strategy is to select as uniform a population as is possible. Thus we focus on young, unmarried, low-skilled men who are at the beginning of their careers. The information from the annual interviews is used to construct a weekly employment record for each respondent.

Since the maximum age that an individual could reach during the sample period is only 26 years, our results should be viewed as applying to young workers who tend to be quite mobile during this early phase of their career. Whether the results generalize to older workers, or different cohorts of workers, is an open question. The NLSY97 collects extensive information about labor market behavior and educational experiences which provide the information needed to study the transition from schooling to employment, early career mobility between employers, and the associated dynamics of wages. The model is estimated using a sample of men from the NLSY97. Women are excluded for the usual reason of avoiding the difficulties associated with modeling female labor force participation. Men who are ever married during the sample period are also excluded to avoid issues relating to household search. Since we are interested in low-skilled workers, the sample excludes individuals who attend college at any point during the sample period. Men enter the estimation sample when they stop attending high school. As is standard in the empirical search literature, individuals who ever serve in the military or are self employed are excluded from the sample.

The NLSY97 provides a weekly employment record for each respondent which is aggregated into a monthly¹³ labor force history for the purposes of estimation. First, each individual is classified as unemployed or employed full time for each month depending on whether more weeks were spent employed or unemployed during the month.¹⁴ Next, em-

¹³For tie-breaking purposes, we use a 5-week month.

¹⁴Non-participation and unemployment are considered to be the same state for the purposes of aggregating

ployed individuals are assigned a monthly employer based on the employer that the worker spent the most weeks working for during the month. The monthly wage is the one associated with the monthly employer. The monthly employment record contains a complete record of employment durations, direct transitions between employers that occur without an intervening spell of unemployment, transitions into unemployment, and the growth in wages resulting from mobility between employers.

Since the importance of non-wage job characteristics is identified in part by job-to-job transitions, we are careful to differentiate between those that are voluntary and those that are not. To identify involuntary job-to-job transitions we use the stated reason that a worker left their job. We consider “layoffs,” “plant closings,” “end of a temporary or seasonal job,” “discharged or fired” or “program ended” to be involuntary. While these data may be somewhat noisy, we are reassured by the summary statistics which show that direct transitions we classify as strictly involuntary are more likely to result in a wage decline (Table 1). In addition, on average, workers who make involuntary transitions between employers experience nearly a 2 percent decline in wages. In contrast, wages increase on average by 8 percent at all transitions between employers.

The final issue worthy of discussion regarding the data is the treatment of within-job variation in wages. In the NLSY97, when a job persists across survey interviews, which occur approximately one year apart, a new measurement of the wage is taken. If a job does not last across interview years, only the initial measurement of the wage is available. In principle, it would be possible to allow for within-job variation in wages using these data. However, as discussed by Flinn (2002), jobs with observed wage changes are not a random sample from the population, so there are difficult selection issues which must be confronted when estimating an on-the-job wage process using these data. Even more importantly for our purposes, since the NLSY97 is still a relatively short panel, the majority of jobs do not persist across survey years. For these jobs, it is impossible to observe on-the-job wage growth. More specifically, we only observe a single wage for 72 percent of all jobs in our the data. Full time employment is considered to be jobs that involve at least twenty hours of work per week.

Table 1: Descriptive Statistics: NLSY97 Data

	Job Number within Cycle		
	Job 1	Job 2	Job 3
Mean log-wage	1.979	2.038	2.061
Std. dev. of log-wage	0.425	0.458	0.457
Mean employment spell duration	8.939	9.271	9.738
Number of observations	2614	940	382
	Type of Employer Switch		
	All	Involuntary	
Pr(wage decrease) at job-to-job move	0.364	0.460	
Mean Δw at job-to-job switch	0.081	-0.017	
Mean Δw at job-to-job switch $ \Delta w > 0$	0.359	0.322	
Mean Δw at job-to-job switch $ \Delta w < 0$	-0.327	-0.345	
All Jobs			
Mean unemployment spell duration	5.908		
Mean number of cycles per person	2.878		
Std. dev. of number of cycles per person	1.793		
Mean work experience	40.01		
Fraction of job-to-job transitions that are involuntary	0.151		
Number of people	980		
Mean number of months per person	54.153		

data. In addition, we are unable to reject the null hypothesis that wages are constant within job spells in our estimation sample. Given these features of the data, there is little hope of precisely estimating an on-the-job wage growth process. As a result, we restrict wages to be constant within job spells for the purposes of estimation. When multiple wages are reported for a particular job, we use the first reported wage as the wage for the entire job spell. Moreover, for our application, with our focus on young, unskilled workers during a highly mobile stage of their career, constant wages within jobs does not seem unrealistic.

3.1 Descriptive Statistics

This section highlights the key characteristics of the data used to estimate the importance of non-wage job characteristics in determining employment outcomes. It is convenient to describe the labor market histories in the data and the data generated by the search model in terms of employment cycles, as in Wolpin (1992). An employment cycle begins with

unemployment and includes all of the following employment spells that occur without an intervening unemployment spell. When an individual enters unemployment, a new cycle begins. In the remainder of the paper, whenever a job is referred to by number, it represents the position of the job within an employment cycle.

Table 1 shows the means and standard deviations of key variables from the sample of the NLSY97 used in this analysis. There are 980 individuals in the data who remain in the sample for an average of 54.2 months, and these people experience an average of 2.88 employment cycles. The top section of the table shows that as individuals move between employers within an employment cycle, the average wage and employment duration increase.¹⁵ The middle section of the table shows that although mean wages increase as individuals move directly between jobs, conditional on switching employers without an intervening unemployment spell there is a 36 percent chance that an individual reports a lower wage at his new job.¹⁶ For individuals who report that the direct transition between employers was involuntary, the mean wage change is negative and the probability of a wage decrease rises to 46 percent. Measurement error in wages certainly accounts for some fraction of the observed wage decreases at voluntary transitions between employers. However, the prevalence of these wage decreases and the increased probability of observing a wage decline at an involuntary transition both suggest a role for non-wage job characteristics in determining mobility between jobs.

4 Estimation

The parameters of the model are estimated by simulated minimum distance (SMD). This section begins by specifying the distributional assumptions about the job offer distribution, measurement error in wages, and unobserved heterogeneity needed to estimate the model.

¹⁵Statistics are not reported for more than three jobs within a cycle because only a very small number of people have four or more consecutive jobs without entering unemployment.

¹⁶This number is consistent with existing estimates of the fraction of direct employer-to-employer transitions that involve a wage decrease. Bowlus and Neumann (2006) report that 40 percent of direct transitions involve a wage decrease in the NLSY79.

Then it explains how the simulated data is generated, describes the estimation algorithm, and discusses identification.

4.1 Distributional Assumptions

The Wage Offer Distribution

Estimating the model requires specifying the distribution $F(w, \xi)$, which is a primitive of the model.¹⁷ We assume that log-wage offers and match-specific utility flows are independent, and normally distributed,

$$F(w, \xi) \sim \Omega(w)\Psi(\xi) \tag{4}$$

$$\Omega(w) \sim N(\mu_w, \sigma_w) \tag{5}$$

$$\Psi(\xi) \sim N(0, \sigma_\xi). \tag{6}$$

Relaxing the assumption that offers of w and ξ are independent would require empirically determining whether workers tend to receive high non-wage utility, which is unobserved by the econometrician, from jobs that offer high wages. Arguments can be made for either positive (health insurance) or negative (risk of injury or death) correlation between w and ξ . In the absence of compelling information, we assume zero correlation. Section 5.3 demonstrates that our key empirical results are robust to the independence assumption by adopting the alternative assumption that (w, ξ) are distributed bivariate normal, and re-estimating the model for two different values of the correlation parameter. As a practical matter, we do not believe that there are any moments from the NLSY97 which would identify this correlation. It is also worth noting that although offers of (w, ξ) are independent, the predicted values of (w, ξ) that are accepted by workers will still be correlated because of optimal search behavior (Section 5.4).

¹⁷We do not attempt to endogenize the job offer distribution because our primary goal in this paper is to quantify the relative importance of non-wage utility for workers, taking the offer distribution as given. Developing a tractable partial equilibrium model allows us to focus directly on this issue, as in much of the existing literature that uses search models to quantify the monetary gains to search and mobility (Jolivet et al., 2006; Flinn, 2002; Sullivan, 2010).

Measurement Error in Wages

A large literature surveyed by Bound et al. (2001) finds that wages in typical sources of microeconomic data are measured with error. We account for measurement error by assuming that the relationship between the log-wage observed in the data and the true log-wage is $w^o = w + \varepsilon$, where w^o is the observed log-wage, w is the true log-wage, and $\varepsilon \sim N(0, \sigma_\varepsilon)$ represents measurement error in wages that is independent of the true wage.¹⁸ The parameter σ_ε is estimated jointly along with the other parameters in the model. Section 4.4 discusses how the extent of measurement error in wages is separately identified from the importance of non-wage utility. The addition of measurement error in wages to the model does not change the optimization problem faced by agents because optimal decisions are based on true wages, not observed wages. However, measurement error impacts the simulated data used to estimate the model.

Accounting for Unobserved Heterogeneity

The search model presented in Section 2 assumes that all individuals are *ex ante* identical at the start of their careers, which implies that all differences in wages and employment outcomes are driven by randomness in the labor market. Although the sample of workers from the NLSY97 used in estimation consists of a fairly homogeneous group in terms of observable characteristics, it is possible that there are permanent differences between workers that are unobserved to the econometrician. In general, ignoring unobserved heterogeneity during estimation will lead to biased parameter estimates if unobserved heterogeneity is actually present.

In this application, the specific concern is that ignoring unobserved differences between workers could lead to an overstatement of the importance of non-wage utility. For example, suppose that a worker remains in a job with a wage in the bottom 5 percent of the wage distribution over the entire sample period. If workers are assumed to be homogeneous, then

¹⁸Accounting for measurement error in this way is standard in the search literature. See, for example, Stern (1989), Wolpin (1992), and Eckstein et al. (2009).

the model will tend to explain the long duration of this low wage job as a situation where the worker has a large draw of ξ , so he is willing to remain in the low wage job because it provides a high level of utility. However, if there is heterogeneity across workers in ability, low ability workers could choose to remain in jobs that offer low wages relative to the overall wage distribution because these jobs are actually high paying relative to their personal (low ability) wage distribution.

We account for person-specific unobserved heterogeneity in ability by allowing the mean of the wage offer distribution (μ_w) to vary across workers. Heterogeneity in preferences for leisure is captured by allowing the reservation utility level (U^*) to vary across workers. In addition, we allow the job offer arrival rates while unemployed (λ_u) and employed (λ_e), the layoff probability (λ_l), and the simultaneous layoff-offer probability (λ_{le}) to vary across workers to allow for the possibility that workers face different amounts of randomness in job offer arrivals and exogenous job endings.¹⁹ Following Keane and Wolpin (1997), and a large subsequent literature, we assume that the joint distribution of unobserved heterogeneity is a mixture of discrete types. Assume that there are J types of people in the economy, and let π_j represent the proportion of type j in the population. The parameters of the distribution of unobserved heterogeneity, $\{\mu_w(j), U^*(j), \lambda_e(j), \lambda_l(j), \lambda_{le}(j), \lambda_u(j), \pi_j\}_{j=1}^J$, are estimated jointly along with the other parameters of the model.

4.2 Data Simulation

As discussed in Section 2, the optimal decision rules for the dynamic optimization problem can be described using simple static comparisons of one-period utility flows. It is straightforward to simulate data from the model using these optimal decision rules without numerically solving for the value functions that characterize the optimization problem.

The first step when simulating the model is to randomly assign each individual in the data to one of the J discrete types that make up the population distribution of unobserved

¹⁹Eckstein et al. (2009) also take the approach of allowing for heterogeneity in offer arrival and layoff probabilities.

heterogeneity. Next, a simulated career is formed for each individual in the NLSY97 estimation sample by randomly generating job offers and exogenous job endings, and then assigning simulated choices for each time period based on the reservation value decision rules. The number of time periods that each simulated person appears in the simulated data is censored to match the corresponding person in the NLSY97 data. Measurement error is added to the simulated accepted wage data based on the assumed measurement error process.

4.3 Simulated Minimum Distance Estimation

Simulated minimum distance estimation finds the vector of structural parameters that minimizes the weighted difference between vectors of statistics estimated using two different data sets: the NLSY97 data, and simulated data from the model. We use the terminology simulated minimum distance to make it clear that during estimation we match moments from the data (as in the simulated method of moments) and the parameters of an auxiliary model (as in indirect inference).²⁰ In this application, the auxiliary parameters are the parameters of a reduced form wage regression. In the remainder of the paper, for brevity of notation we refer to all of the statistics from the data that are matched during estimation as moments.

Let $\theta = \{\sigma_w, \sigma_\xi, \sigma_\varepsilon, \mu_w(j), U^*(j), \lambda_e(j), \lambda_l(j), \lambda_{le}(j), \lambda_u(j), \pi_j\}_{j=1}^J$ represent the parameter vector that must be estimated.²¹ The search model is used to simulate S artificial datasets, where each simulated dataset contains a randomly generated employment history for each individual in the sample. The simulated and actual data are each summarized by K moments. The SMD estimate of the structural parameters minimizes the difference between the simulated and sample moments. Let m_k represent the k th moment in the data, and let $m_k^S(\theta)$ represent the k th simulated moment, where the superscript S denotes averaging across the S artificial datasets. The vector of differences between the simulated and actual

²⁰See Stern (1997) for a survey of simulation based estimation, and Smith (1993) for the development of indirect inference. Recent examples of papers that use this approach to estimating search models include Eckstein et al. (2009) and Yamaguchi (2010).

²¹The parameter vector that is estimated does not include the discount factor (δ), or the one period utility flow from unemployment (b). However, by choosing a value for the discount factor, and fixing the other parameter values at the SMD estimates, we can use the structure of the model to compute b using equation (10).

moments is $g(\theta)' = [m_1 - m_1^S(\theta), \dots, m_K - m_K^S(\theta)]$, and the simulated minimum distance estimate of θ minimizes the following objective function,

$$\Phi(\theta) = g(\theta)'Wg(\theta), \tag{7}$$

where W is a weighting matrix. We use a diagonal weighting matrix during estimation, where each diagonal element is the inverse of the variance of the corresponding moment. We estimate W using a nonparametric bootstrap with 300,000 replications. Bootstrapping the matrix W is convenient because it is not necessary to update the weighting matrix during estimation. Parameter estimates are obtained by minimizing the objective function shown in equation (7) using simulated annealing.²² Simulated moments are averaged over $S = 25$ simulated datasets. The standard errors are computed using a nonparametric bootstrap using 900 draws from the NLSY97 data.

4.4 Choice of Moments and Identification

This section discusses the moments targeted during estimation and provides an informal discussion of how they identify the parameters of the structural model. Table 2 lists the 44 moments from the NLSY97 that are used to estimate the model. Most of the parameters, with the exception of those governing the importance of non-wage utility, are identified as they would be in a standard on-the-job search model by moments that summarize wages, employment and unemployment durations, and transition rates. For example, the moments shown in Panel 1 of Table 2 describe the first three employment spells within employment cycles using the mean and standard deviation of accepted wages, and the mean employment spell duration. The wage moments provide information about the parameters of the wage offer distribution. Recall that employers within a cycle represent a sequence of direct transitions between employers that occur without an intervening spell of unemployment, so mean

²²See Goffe et al. (1994) for a discussion of the simulated annealing algorithm and FORTRAN source code to implement the algorithm. The primary advantage of this algorithm is that it is a global search algorithm that can escape local optima.

wages conditional on employer number also provide information about wage growth from job search. Also, the increase in the average employment duration as workers move between jobs reflects improvements in job match values from on-the-job search. The moments shown in Panel 2 of Table 2 provide additional information about unemployment spell durations, job durations, and transition probabilities which help to identify the job offer and layoff rates.

Table 2: Moments of the NLSY97 Data and Simulated Data

Moment	NLSY97 Data	Simulated
Cycle Moments (Panel 1)		
Mean log-wage (employer 1)	1.9791	1.9871
Std. dev. of log-wage (employer 1)	0.4249	0.3968
Mean job duration (employer 1)	8.9392	9.1354
Mean log-wage (employer 2)	2.0377	2.0537
Std. dev. of log-wage (employer 2)	0.4582	0.4151
Mean job duration (employer 2)	9.2713	10.0262
Mean log-wage (employer 3)	2.0608	1.9988
Std. dev. of log-wage (employer 3)	0.4572	0.4686
Mean job duration (employer 3)	9.7382	9.2036
Transition and Duration Moments (Panel 2)		
Mean unemployment spell duration	5.9087	5.7809
Std. dev(unemp. duration)	7.7319	7.5213
Pr(transition from employment to unemployment)	0.0364	0.0383
Pr(job-to-job transition)	0.0364	0.0383
Pr(unemp. duration = 1)	0.2375	0.3196
Pr(unemp. duration = 2)	0.1697	0.1150
Pr(unemp. duration = 3)	0.1092	0.0928
Pr(emp. duration = 1)	0.1423	0.1488
Pr(emp. duration = 2)	0.1412	0.1296
Pr(emp. duration = 3)	0.1209	0.1037
Mean number of firms per cycle	1.6983	1.7777
Std. dev. of number of firms per cycle	1.1513	1.0586
Mean number of firms over career	4.3755	4.2010

Continued on next page

Table 2 – continued from previous page

Moment	NLSY97 Data	Simulated
Job-to-job Transition Moments (Panel 3)		
Mean Δw at job-to-job switch	0.0812	0.1097
Mean Δw at job-to-job switch $ \Delta w > 0$	0.3592	0.3841
Mean Δw at job-to-job switch $ \Delta w < 0$	-0.3273	-0.3099
Pr(wage decrease at job-to-job transition)	0.3640	0.3871
Pr(wage decrease at involuntary job-to-job transition)	0.4601	0.4979
Mean Δw at involuntary job-to-job switch	-0.0168	-0.0289
Mean Δw at involuntary job-to-job switch $ \Delta w > 0$	0.3224	0.3341
Mean Δw at involuntary job-to-job switch $ \Delta w < 0$	-0.3454	-0.3649
Fraction of job-to-job transitions that are involuntary	0.1505	0.1592
Wage Regression (Panel 4)		
Constant	1.9311	1.9268
Experience	0.0058	0.0047
Experience ² /100	-0.0021	-0.0032
Covariances (Panel 5)		
Within-person cov(wages)	0.0448	0.0510
Cov(1st wage, 1st unemp. duration)	-0.1439	-0.1582
Cov(1st unemp. duration, 1st emp. duration)	-1.4050	-1.6091
Within-person cov(mean(w),mean(% months unemp.))	-0.0332	-0.0346
Cov(wage,transition into unemp.)	-0.0560	-0.0589
Other Moments (Panel 6)		
Across-person std. dev(w)	0.3131	0.3191
Across-person std. dev(unemp. duration)	5.9004	4.2334
Across-person std. dev. of number of firms	2.9437	3.0766
Across-person mean % months unemp.	0.2745	0.2953
Across-person std. dev(% months unemp.)	0.2587	0.2446

Notes: Simulated data is generated using the estimated model.

An important distinction between this paper and the existing literature is that we allow for three possible explanations for observed wage declines at direct job-to-job transitions. The possible explanations are measurement error in wages, involuntary job endings that occur at the same time as outside job offers, and non-wage utility. To the best of our knowledge, this is the first paper to build a model that incorporates all of these explanations, and estimates the model to quantify the importance of each.²³ The most straightforward of these three possible

²³Flinn (2002) allows for measurement error in wages, Wolpin (1992) allows for both measurement error in wages and simultaneous exogenous job endings and outside job offers, and Bonhomme and Jolivet (2009) allow for simultaneous job endings and job offers in a search model where specific non-wage job characteristics enter the utility function.

explanations from the perspective of identification is involuntary job endings that occur at the same time as outside job offers. The probability that this event occurs is represented in the model by the parameter λ_{le} . As discussed in Section 3, we use data from the NLSY97 on the reason that jobs end to distinguish between voluntary and involuntary direct transitions between employers. If an individual reports that a job ends involuntarily, and he moves to a new job without experiencing an intervening spell of unemployment, then a simultaneous exogenous job ending and accepted outside offer has occurred. It follows that the fraction of direct job-to-job transitions in the data that are involuntary (Panel 3 of Table 2) identifies λ_{le} .

To see the importance of accounting for involuntary transitions between employers during estimation, note that within the structure of the model, a voluntary transition to a lower wage job can only be explained by measurement error or non-wage utility. In contrast, if a job exogenously ends and a new offer is received, a worker could move to a job that offers lower utility than his previous job because it is preferable to unemployment. More concretely, suppose that a job paying wage w exogenously ends, and the worker simultaneously receives a new outside job offer (w', ξ') , where $w' < w$. The worker will accept a wage decrease equal to $(w' - w)$ instead of becoming unemployed if $U(w', \xi') > U^*$. If the presence of involuntary transitions in the data was ignored during estimation, it would force the model to account for all negative wage changes at job transitions in the data with either measurement error in wages or non-wage utility.

Taking the identification of λ_{le} as given, it remains to discuss identification of the importance of measurement error in wages, true variation in wage offers, and non-wage utility. At first glance, it might appear difficult to distinguish the effects of non-wage utility from measurement error without relying on validation data to identify misreported wages. However, the parameters that determine measurement error in wages (σ_ϵ), true variation in wage offers (σ_w), and variation in non-wage utility (σ_ξ) actually have very different implications for the moments used during estimation.

To understand how σ_ε and σ_ξ are separately identified it is useful to begin with the standard income maximizing on-the-job search model which assumes that $\sigma_\xi = 0$. In such a model, job acceptance is governed by a reservation wage w^* , as opposed to the reservation utility level (U^*) in our utility maximizing search model. In the standard model, unemployment durations and observed accepted wages identify $\mu_w, \sigma_w, \lambda_u$, and w^* . Given μ_w and σ_w , job-to-job transitions identify λ_e , and transitions from employment to unemployment identify λ_l . Since job mobility choices are based only on wages, voluntary moves to lower wage jobs must be attributed to measurement error in wages, so this feature of the data identifies σ_ε . In order to match the frequency of wage declines at job transitions shown in Panel 3 of Table 2, σ_ε must be relatively large. However, as σ_ε increases, σ_w must decrease, or else the model will generate too much variation in observed wages relative to the data (Panel 1 of Table 2). In other words, when the amount of measurement error in the model is high, the amount of true variation in wage offers must be low in order to match the observed variation in wages in the NLSY97. This property of the model is demonstrated in columns 1 and 2 of Table 3, which show the parameter estimates for a restricted version of the model which assumes that $\sigma_\xi = 0$ along with the estimates for the unrestricted model. Finally, it is important to note that as σ_w decreases, the model generates lower wage gains from job search. This happens because as σ_w decreases, there is a lower chance that an employed worker will receive a higher outside wage offer.

Next, consider estimating a model with non-wage utility ($\sigma_\xi > 0$). Using only the job transition moments shown in Panel 3 of Table 2, it is impossible to separately identify σ_ε and σ_ξ , because a voluntary job-to-job move to a lower observed wage could be due to either mis-measurement of the wage or unobserved non-wage utility. The key to identifying the full model with non-wage utility is that during estimation, we simultaneously match moments that capture three features of the data. First, the amount of variation in wages (Panel 1 of Table 2). Second, the frequency and magnitude of wage declines at job transitions (Panel 3 of Table 2). Third, the amount of wage growth as reflected in the mean wages from the first, second and third jobs within job cycles (Panel 1 of Table 2), and the coefficients

from a reduced form regression of wages on work experience (Panel 4 of Table 2). For the reasons discussed in the previous paragraph, the model without non-wage utility is unable to simultaneously match these three features of the data. However, the full model is able to match these moments. This is the case because in the model with non-wage utility, as the parameter σ_ξ increases from zero, the amount of measurement error needed to explain the frequency of wage declines at voluntary job transitions falls. As a result, σ_w can be relatively large without causing the model to over predict the amount of wage dispersion in the data. As σ_w increases, the model is able to match the amount of wage growth, the extent of variation in wages, and the frequency of wage declines at job transitions found in the NLSY97 data.

The final two groups of moments shown in Panels 5 and 6 of Table 2 identify the importance of person-specific heterogeneity in the model, which is captured by the parameters $\{\mu_w(j), U^*(j), \lambda_e(j), \lambda_l(j), \lambda_{te}(j), \lambda_u(j), \pi_j\}_{j=1}^J$. For example, the within-person covariance in wages identifies the person-specific component of wages, just as it would in a simpler panel data model of wages. When there is no heterogeneity in μ_w , the model generates a within-person covariance of zero between wages on employers that are separated by unemployment spells. Similarly, the covariances between wages and unemployment durations, and unemployment durations and employment durations provide information about the covariances between the various person-specific heterogeneity terms. Finally, the across-person standard deviations of wages and unemployment durations (Panel 6) provide additional information about the importance of person-specific unobserved heterogeneity relative to the importance of true randomness in determining wages and unemployment durations.

5 Empirical Results

This section discusses the estimated structural model. It begins with a discussion of the estimated parameters and their impact on labor market outcomes. Next, we discuss the ability of the estimated model to fit the data. Finally, we check the robustness of our

estimates by re-estimating the model using alternative assumptions regarding the correlation between log-wage offers and non-wage utility.

5.1 Parameter Estimates

This section presents the estimation results based on the theoretical model discussed in Section 2 and the econometric methodology described in Section 4. The estimated parameters for two specifications of the model are given in Table 3. In both specifications, we allow for unobserved worker heterogeneity. The first and preferred specification (1) allows for utility from non-wage job characteristics. In the second specification (2), workers are wage maximizers so that non-wage job characteristics are unimportant to workers' job choice and mobility decisions. There are two key results that are worth highlighting.

Table 3: Parameter Estimates

Parameter	Notation	Specifications	
		(1)	(2)
Stand. dev. of wage offer	σ_w	0.3176 (0.0039)	0.2671 (0.0026)
Stand. dev. of non-wage match	σ_ξ	0.2657 (0.0063)	0.0000 (—)
Stand. dev. of measurement error	σ_ε	0.1777 (0.0243)	0.2439 (0.0167)
<u>Type 1</u>			
Mean wage offer	$\mu(1)$	1.0619 (0.0521)	1.3169 (0.0476)
Reservation utility	$U^*(1)$	1.0892 (0.0820)	0.3893 (1.5078)
Pr(offer while unemployed)	$\lambda_u(1)$	0.0753 (0.0484)	0.0441 (0.0063)
Pr(layoff)	$\lambda_l(1)$	0.0159 (0.0118)	0.1240 (0.0191)
Pr(offer while employed)	$\lambda_e(1)$	0.9060 (0.0354)	0.0016 (0.0070)
Pr(offer and layoff)	$\lambda_{le}(1)$	0.0780 (0.0321)	0.2615 (0.0464)
<u>Type 2</u>			
Mean wage offer	$\mu(2)$	1.7105 (0.0069)	1.7015 (0.0041)

Continued on next page

Table 3 – continued from previous page

Parameter	Notation	Specifications	
		(1)	(2)
Reservation utility	$U^*(2)$	2.0995 (0.0067)	1.9385 (0.0034)
Pr(offer while unemployed)	$\lambda_u(2)$	0.9887 (0.0462)	0.9196 (0.0383)
Pr(layoff)	$\lambda_l(2)$	0.0001 (0.0026)	0.0518 (0.0092)
Pr(offer while employed)	$\lambda_e(2)$	0.4972 (0.0224)	0.5157 (0.0124)
Pr(offer and layoff)	$\lambda_{le}(2)$	0.1056 (0.0054)	0.0643 (0.0045)
<u>Type 3</u>			
Mean wage offer	$\mu(3)$	2.1752 (0.0369)	2.1544 (0.0335)
Reservation utility	$U^*(3)$	0.9074 (1.4014)	1.4305 (1.6555)
Pr(offer while unemployed)	$\lambda_u(3)$	0.0712 (0.0896)	0.0767 (0.0425)
Pr(layoff)	$\lambda_l(3)$	0.0011 (0.0012)	0.0038 (0.0018)
Pr(offer while employed)	$\lambda_e(3)$	0.0472 (0.0081)	0.0501 (0.0093)
Pr(offer and layoff)	$\lambda_{le}(3)$	0.0005 (0.0009)	0.0016 (0.0019)
<u>Type Probabilities</u>			
Pr(type 1)	π_1	0.1141 (0.0138)	0.1010 (0.0072)
Pr(type 2)	π_2	0.5377 (0.0252)	0.5488 (0.0127)
Pr(type 3)	π_3	0.3483 (0.0188)	0.3502 (0.0137)

Notes: Bootstrapped standard errors in parenthesis.

First, the estimate of the standard deviation of measurement error in wages for our preferred specification (1) is very plausible. In particular, the estimated standard deviation of measurement error is 0.1777 and for wage offers is 0.3176. Along with the variation in mean wages due to unobserved heterogeneity, these coefficients imply that about 12.7 percent of the variation in observed wages is due to the presence of measurement error. On the other hand,

for specification (2) where we impose the restriction that $\sigma_\xi = 0$, the estimated standard deviations of measurement error and wages are 0.2439 and 0.2671, which implies that 24.2 percent of the variation in observed wages is due to measurement error. More importantly, as alluded to in our discussion of identification (Section 4.4), without modeling utility from non-wage job characteristics, the higher relative importance of measurement error depresses the estimated standard deviation in wage offers from 0.3176 to 0.2671. Taken together, these two specifications of the model clearly demonstrate that empirical search models which ignore non-wage utility will provide an upward biased estimate of the extent of measurement error in wages, and will also provide a downward biased estimate of the amount of variation in the wage offer distribution.

The implications of adding non-wage utility to the standard on-the-job search model are made even more apparent by comparing simulated data from specifications (1) and (2) of the model. Table 4 presents mean wages and utility, and percent changes in wages and utility, by employer number within a job cycle. These simulated data are obtained by simulating careers for one million artificial agents over a 300 month time horizon. A comparison of the mean wage changes between columns (1) and (2) of Table 4 demonstrates that the compressed wage offer distribution in the model that ignores non-wage utility generates lower wage growth from job mobility. Notice that between employers 1 and 2, between employers 2 and 3, and over all voluntary transitions, the mean percentage increases in wages are 11.7, 8.8 and 9.5 for our preferred specification (1). But for specification (2), where the estimated wage offer distribution has been compressed by the assumption that $\sigma_\xi = 0$, the mean percentage increases are only 9.6, 7.5 and 8.1. That is, the compressed wage offer distribution results in lower wage growth due to mobility.

Second, and perhaps more importantly from our perspective, utility from non-wage job characteristics is a very important factor as workers evaluate both unemployed and employed job offers. In particular, Table 3 shows that the variation in the utility from the non-wage match is only slightly less than the variation in the wage offer distribution (0.2657 vs. 0.3176).

Table 4: Steady State Cross Section of Simulated Wages and Utility by Employer

		Specification	
		(1)	(2)
Employer 1	Mean wage	2.09	2.11
	Mean utility	2.36	2.11
Employer 2	Mean wage	2.27	2.29
	Mean % change in wage	11.7%	9.6%
	Mean utility	2.55	2.29
	Mean % change in utility	16.0%	9.6%
	mean($\Delta\xi$)/mean(Δ utility)	0.28	0.0
Employer 3	Mean Wage	2.40	2.40
	Mean % change in wage	8.8%	7.5%
	Mean utility	2.68	2.40
	Mean % change in utility	12.4%	7.5%
	mean($\Delta\xi$)/mean(Δ utility)	0.40	0.0
All Employers	Mean wage	2.25	2.37
	Mean % change in wage	9.5%	8.1%
	Mean utility	2.53	2.37
	Mean % change in utility	13.2%	8.1%
	mean($\Delta\xi$)/mean(Δ utility)	0.34	0.0

Notes: Specifications (1) and (2) refer to the estimates in Table 3.

As a result, the standard deviation of the total utility of a job offer is 0.4141. This is in comparison to the standard deviation of total utility (i.e., the log-wage) of 0.2671 when non-wage utility is omitted from the model. These results imply that the standard on-the-job search model underestimates the gains from mobility in two ways. First, as pointed out in the prior paragraph, in the absence of non-wage utility, estimated measurement error must rise and consequently, the estimated variation in wage offers must fall. As a result, wage growth from job-to-job transitions will also be depressed. Second, ignoring non-wage utility misses an important component of job-to-job utility increases. The first column of Table 4 shows that, on average, the mean percentage increases in total utility between employers 1 and 2, 2 and 3, and over all voluntary transitions are 16.0, 12.4 and 13.2. Quantitatively, mobility to higher non-wage matches is an important component of the total gains to workers from job mobility: increases in ξ account for 28 percent of the total gains in utility ($\text{mean}(\Delta\xi)/\text{mean}(\Delta U)$) at the first transition between employers, 40 percent of the total gains at the second transition, and 34 percent of the total gains across all transitions. In aggregate, the benefits to workers

Table 5: Career and Steady State Outcomes: Baseline vs. Counterfactual ($\sigma_\xi = 0$) Simulations

Statistic	Model Specification			
	Notation	Baseline	Counterfactual	% change
Mean disc. sums over career				
log wages	$\frac{1}{N} \sum_{t=0}^T \delta^t w_t$	148.4	142.4	-4.0%
employment utility	$\frac{1}{N} \sum_{t=0}^T \delta^t (w_t + \xi_t)$	164.6	142.4	-13.5%
lifetime utility	$\frac{1}{N} \sum_{t=0}^T \delta^t U_t$	210.1	201.9	-4.0%
Steady state summary statistics				
mean wage		2.26	2.45	
employment utility		2.53	2.45	
employment rate		0.75	0.67	

Notes: Baseline simulation uses the estimates from specification (1).

Counterfactual simulation uses the estimates from specification (1), but imposes $\sigma_\xi = 0$.

from job mobility are considerably understated by models that ignore non-wage utility.

As discussed in the previous paragraph, differences in the simulated data between specifications (1) and (2) arise from two sources. First, a direct effect from eliminating non-wage utility, and second, an indirect effect from changes in other estimated parameters (such as the wage offer distribution) caused by estimating the model under the restriction that $\sigma_\xi = 0$. To isolate the direct effect of non-wage utility on workers, we simulate a counterfactual dataset which uses the parameter estimates from specification (1), but imposes the condition that workers do not search over non-wage job characteristics by setting $\sigma_\xi = 0$.²⁴ Summary statistics for this counterfactual experiment are shown in Table 5.

The counterfactual experiment demonstrates that even with an appropriate estimate of the wage offer distribution from specification (1), eliminating non-wage utility affects workers' optimal search strategies. In particular, if workers do not search over non-wage job characteristics, they have less flexibility in the choice over job offers and as a result, are optimally more selective (U^* increases). Also, in the baseline model workers are willing to accept low wage job offers that offer high non-wage utility, but in the counterfactual model

²⁴This simulation introduces an additional complication because the estimated reservation utility depends on the distribution of ξ , as shown in equation (10). Thus in order to simulate the model under the assumption that $\sigma_\xi = 0$, we need to first recalculate U^* after imposing $\sigma_\xi = 0$.

they are not able to make this trade off. Both of these effects increase the counterfactual mean steady-state log-wage relative to the baseline specification (2.45 vs 2.26).

Despite the fact that the mean wage rises when there is no variation in non-wage utility, the average discounted sum of log-wages over the career is actually lower (142.4 vs 148.4 for specification (1)) in the counterfactual model. This happens because without the consideration of non-wage job characteristics, workers do not take poorly paid jobs, so the rate of employment actually falls (0.67 vs 0.75 for specification (1)). That is, the presence of non-wage job characteristics provides workers with greater flexibility in the selection of jobs, resulting in less time spent unemployed, greater lifetime earnings, and greater overall utility.

Table 5 also shows the discounted expected value of lifetime utility in the baseline and counterfactual simulated labor markets. The average discounted expected value of lifetime utility is an estimate of the value function, and can be used to quantify the welfare impact of search over non-wage utility. On average, the mean discounted sum of lifetime utility when workers can search over non-wage job characteristics is 210.1 whereas when workers only search over wages, mean lifetime utility is 201.9. This difference of 8.2 “utils” implies that an average worker living in our counterfactual world could be made just as well off as an average worker from specification (1) with an additional weekly payment of \$40.66 (assuming a 40 hour work week). With a mean steady-state wage of about \$9.54 or weekly earnings of about \$381, \$40.66 amounts to nearly 11 percent of a typical worker’s weekly earnings. In other words, the average worker is willing to pay 11 percent of his earnings to obtain the option value of searching over non-wage job characteristics.

Unobserved worker heterogeneity is also an important consideration in estimating our model. However, its primary importance for our purposes is as a control so that we obtain unbiased estimates of the parameters of primary interest.²⁵ Nevertheless, it is worth discussing the parameter estimates that capture unobserved worker heterogeneity since the estimates provide evidence of substantial unobserved differences across workers. Since it is

²⁵The minimized value of the SMD objective function is 718.3 for a one type model which assumes that all workers are homogeneous. The minimized objective function falls to 332.9 and 248.9 for models which allow for two and three discrete types.

Table 6: Mean Outcomes by Type in Simulated Data

	Type 1	Type 2	Type 3	All Types
Employment duration	5.38	6.77	42.84	10.51
Unemployment duration	17.53	5.29	7.59	7.12
Log-Wage	1.33	2.07	2.29	2.00
Utility (while employed)	1.54	2.35	2.38	2.25
Disc. EV lifetime utility	153.37	214.15	222.06	210.07
Proportion	0.11	0.54	0.35	1.00

Notes: Durations in months.

sometimes difficult to see how differences in these primitive parameters across types translate into labor market outcomes, Table 6 summarizes outcomes in the simulated data conditional on type. This table also shows the discounted expected value of lifetime utility for each type. Although we do not directly compute the value function during estimation, it is straightforward to approximate the value function using the average discounted lifetime utility realized in a large number of simulated careers generated by the model.²⁶

Unobserved heterogeneity across workers has a large effect on labor market outcomes. Type 1 workers have the shortest average employment duration, the longest average unemployment duration, and by far the lowest average log-wage. Type 3's have by far the longest average employment duration of nearly 43 months, but also experience the second longest mean unemployment duration because they have the lowest unemployed job offer probability. Type 3's on average receive a log-wage that is nearly double that of type 1's, and is about 11 percent larger than that of a Type 2 worker. The differences in parameter values across types translate into moderate differences in the expected discounted value of lifetime utility. The discounted expected value of lifetime utilities for Type 2 and 3 workers are fairly similar and are about 40 percent higher than the discounted expected value of lifetime utility for the frequently unemployed, low wage, Type 1 workers.

²⁶The value function estimates are based on 1,000,000 simulated individuals over 300 months.

5.2 Model Fit

It is also interesting to note that the model does an excellent job fitting the data. Table 2 shows the moments from the NLSY97 data used to estimate the model along with the simulated moments generated by the parameter estimates shown in Table 3 for specification (1) of the model. Overall, the fit of the model to these moments is very good. In particular, the model captures the upward trend in average wages and employment durations as individuals move between employers within employment cycles. The model also generates wage dispersion, conditional on employer number, that closely matches the standard deviation of wages in the NLSY97 data. Panel 4 of Table 2 shows that the model slightly under-predicts the reduced form wage-experience profile, which is perhaps not surprising because the search model does not allow for wage growth due to human capital. However, the difference between predicted and actual wage growth is relatively small: at the mean level of experience in the NLSY97 data, the predicted and actual mean log-wages differ by only 0.019.

The model also does a good job matching the transition and duration moments shown in Panels 2 and 3 of Table 2. In particular, the model closely tracks patterns in mean wage changes in the NLSY97 data. Importantly, given the focus of the paper, the model predicts that 38.7 percent of job-to-job transitions will involve a wage decrease, while 36.4 percent of job-to-job transitions in the NLSY97 data have this feature. Finally, the model is in general successful in fitting the within person covariance moments which are generated by unobserved heterogeneity, and in fitting the across-person moments shown in Panels 5 and 6 of Table 2.

5.3 Sensitivity Analysis: Correlation Between Wage Offers and Non-Wage Utility

This section examines the sensitivity of our estimates to the assumption that log-wage and non-wage utility offers, (w, ξ) , are independent. As discussed in Section 4, arguments can be made in favor of either a positive or negative correlation. For example, amenities such

Table 7: Parameter Estimates Conditional on Fixed $\rho(\xi, w)$

	Notation	Parameter Estimates		
		$\rho(\xi, w) = 0$	$\rho(\xi, w) = -0.20$	$\rho(\xi, w) = 0.20$
Stand. dev. of wage offer	σ_w	0.3176 (0.0039)	0.3318	0.2986
Stand. dev. of non-wage match	σ_ξ	0.2657 (0.0063)	0.3246	0.2716
Stand. dev. of measurement error	σ_ε	0.1777 (0.0243)	0.1197	0.2046

Notes: The estimates shown in the column headed by $\rho(\xi, w) = 0$ correspond to the baseline version of the model (shown for reference purposes). The final two columns show the parameter estimates from versions of the model that assume a non-zero value of $\rho(\xi, w)$. All model parameters are re-estimated in each specification, but only three parameters are shown here in the interest of readability.

as employer provided health insurance and retirement benefits are positively correlated with wages. On the other hand, unpleasant or dangerous jobs (low ξ jobs) are associated with higher wages. Most importantly, the correlation between w and ξ is not identified using available data. As such, we exogenously impose different correlations between w and ξ , re-estimate the model, and examine the sensitivity of the key model parameters.

Table 7 presents an abbreviated table of parameter estimates that includes the standard deviations of the wage, non-wage match, and measurement error. In this exercise, we assume that (w, ξ) are distributed bivariate normal with correlation parameter $\rho(w, \xi)$, and present estimates for correlations of -0.20 and 0.20 . Most importantly for our purposes, we find that the estimate of σ_ξ is quite robust to changes in the correlation parameter of this magnitude. In fact, the results of this experiment suggest that our preferred estimate of $\sigma_\xi = 0.2657$ represents a lower bound on the importance of non-wage utility.

Also notable is that the amount of measurement error falls for a negative correlation and rises for a positive correlation. The intuition is straightforward: when wages and non-wage

utility are negatively correlated, low wage offers tend to be balanced by good non-wage job characteristics, so less measurement error is needed to rationalize voluntary moves to lower wage jobs. The converse is true with a positive correlation. In order to compensate for lower (higher) levels of measurement error when $\rho = -0.20$ ($\rho = 0.20$), the standard deviation of wage offers is slightly higher for a negative (positive) correlation.

5.4 Implications for Compensating Wage Differentials

The empirical literature on compensating wage differentials often yields mixed results. For example, Brown (1980) showed that even controlling for individual characteristics, hedonic estimates are “often wrong signed or insignificant.” In this section, we provide a brief discussion of this topic to illustrate the implications of incorporating non-wage job characteristics into a search model to standard, static hedonic estimates of compensating wage differentials. In Sullivan and To (2011), we show in detail how and why search over wages and non-wage job characteristics can yield biased compensating wage differential estimates.

Broadly speaking, a bias arises because optimal search behavior by workers implies that accepted job offers are truncated from below. Since total utility for employed workers is typically greater than their reservation utility level, observed job choices do not directly reveal the willingness to pay for non-wage job characteristics.²⁷ Given our additively separable utility function, $U = w + \xi$, the known willingness to pay is one-to-one so that in a frictionless labor market, the compensating wage differential should be -1 . Although our simulated dataset reveals a trade-off between wages and non-wage job characteristics, the fact that most jobs offer total utility strictly above the reservation level biases willingness-to-pay estimates towards zero. Indeed, estimating a traditional hedonic regression of simulated log-wages on ξ yields a slope coefficient of only -0.48 . Search frictions result in a severely attenuated compensating wage differential estimate.

In general, when search frictions are an important feature of the labor market, compen-

²⁷Indeed, with a continuous wage offer distribution, it must be the case that acceptable job offers almost always yield utility greater than the reservation level.

sating wage differential estimates will be biased. This is similar to the findings of Jolivet et al. (2006), who estimate a model with several non-wage job characteristics. They find strong preferences for amenities but little evidence of compensating differentials in their simulated data. Our aggregate approach with choice over just two dimensions draws a clear picture of precisely how search frictions interact to bias compensating wage differential estimates. In our model, biased compensating wage differential estimates arise because search frictions imply that acceptable jobs typically provide utility greater than the reservation level. This explanation differs from that of Hwang et al. (1992), who show that compensating wage differential estimates are biased because in equilibrium, total job valuation is correlated with the amenity level.

6 Concluding Remarks

This paper develops and estimates an on-the-job search model which allows workers to search across jobs based on both wages and job-specific non-wage utility flows. Estimating the model provides a direct test of the widespread assumption that workers act as pure income maximizers. We estimate the structural model by simulated minimum distance using the NLSY97. The importance of non-wage utility is revealed through voluntary job-to-job moves, wage changes at transitions, and job durations. Measurement error in wages is separately identified from non-wage utility because incorrectly attributing events not explained by observed wages to measurement error compresses the estimated wage offer distribution and as a result, causes the model to generate too little wage growth relative to the data.

The empirical results show that workers place a substantial value on non-wage job characteristics. When searching for a job, workers face nearly as much dispersion in non-wage utility matches as in wage offers. Furthermore, utility from non-wage job characteristics accounts for approximately one-third of the total gains to workers from job mobility. Standard income maximizing models of on-the-job search, which are frequently used to quantify the gains to mobility, are missing a sizable fraction of the gains from search.

We estimate several different specifications of the structural model to examine the sensitivity of the results to alternative assumptions about the correlation between wage offers and non-wage match values. Although there are some shifts in the parameter estimates, the key parameters (σ_w , σ_ξ , and σ_ε) are relatively stable, and the changes in these parameters between specifications have intuitive explanations. Most importantly, the sensitivity analysis suggests that the baseline specification of the model provides a lower bound estimate of the importance of non-wage utility.

Our model also provides a framework for understanding the difficulty that economists have had in estimating compensating wage differentials. In a frictionless competitive labor market, equally able workers must receive the same total compensation and the estimated wage differential for a job attribute will equal the workers' willingness-to-pay for that attribute. In contrast, in a labor market with search frictions, total utility will in general exceed a worker's reservation utility and different, equally able workers will receive different compensation packages, biasing estimates of compensating wage differentials.

A Appendix: Derivation of Reservation Utility

The reservation utility level for unemployed agents, U^* , solves $V^e(U) = V^u$. To derive U^* , we must first rearrange (3) and (2) so that common terms can be collected when evaluated at $U = U^*$. Subtracting $\delta V^e(U)$ from both sides of (3):

$$(1 - \delta)V^e(U) = U + \delta[\lambda_e E \max\{0, V^e(U') - V^e(U)\} + \lambda_l(V^u - V^e(U)) \\ + \lambda_{le} E \max\{V^u - V^e(U), V^e(U') - V^e(U)\}.$$

Evaluating this at $U = U^*$:

$$\begin{aligned}
(1 - \delta)V^e(U^*) &= U^* + \delta \left[\lambda_e \int_{U^*}^{\infty} [V^e(U') - V^u] dH(U') + \lambda_{le} \int_{U^*}^{\infty} [V^e(U') - V^u] dH(U') \right] \\
&= U^* + \delta(\lambda_e + \lambda_{le}) \int_{U^*}^{\infty} [V^e(U') - V^u] dH(U')
\end{aligned} \tag{8}$$

Similarly, subtracting δV^u from both sides of (2),

$$\begin{aligned}
(1 - \delta)V^u &= b + \delta\lambda_u E \max\{0, V^e(U') - V^u\} \\
&= b + \delta\lambda_u \int_{U^*}^{\infty} [V^e(U') - V^u] dH(U').
\end{aligned} \tag{9}$$

Evaluating at $U = U^*$, we can equate (8) and (9), integrate by parts and solve to get:

$$\begin{aligned}
U^* &= b + \delta[\lambda_u - (\lambda_e + \lambda_{le})] \int_{U^*}^{\infty} [V^e(U') - V^u] dH(U') \\
&= b + \delta[\lambda_u - (\lambda_e + \lambda_{le})] \int_{U^*}^{\infty} V^{e'}(U') [1 - H(U)] dU' \\
&= b + \delta[\lambda_u - (\lambda_e + \lambda_{le})] \int_{U^*}^{\infty} \frac{1 - H(U')}{(1 - \delta) + \delta\{\lambda_e[1 - H(U')] + \lambda_l + \lambda_{le}\}} dU'.
\end{aligned} \tag{10}$$

When $\lambda_u > \lambda_e + \lambda_{le}$ (the probability of receiving an offer while unemployed is greater than that when employed), an unemployed worker's reservation wage exceeds the one-period utility flow from unemployment.

References

Altonji, J. G. and C. H. Paxson, "Labor Supply, Hours Constraints, and Job Mobility,"

Journal of Human Resources, 27(2):256–278, 1992.

Becker, D., "Non-Wage Job Characteristics and the Case of the Missing Margin," 2010,

working paper.

Bhaskar, V., A. Manning and T. To, “Oligopsony and Monopsonistic Competition in Labor Markets,” *Journal of Economic Perspectives*, 16(2):155–174, 2002.

Bhaskar, V. and T. To, “Minimum Wages for Ronald McDonald Monopsonies: A Theory of Monopsonistic Competition,” *The Economic Journal*, 109:190–203, 1999.

Blau, D. M., “Search for Nonwage Job Characteristics: A Test of the Reservation Wage Hypothesis,” *Journal of Labor Economics*, 9(2):186–204, 1991.

Bloemen, H., “Job Search, Hours Restrictions, and Desired Hours of Work,” *Journal of Labor Economics*, 26(1):137–179, 2008.

Bonhomme, S. and G. Jolivet, “The Pervasive Absence of Compensating Differentials,” *Journal of Applied Econometrics*, 24(5):763–795, 2009.

Bound, J., C. C. Brown and N. Mathiowetz, “Measurement Error in Survey Data,” in E. E. Learner and J. J. Heckman, eds., “Handbook of Econometrics,” pp. 3705–3843, North Holland Publishing, New York, 2001.

Bowlus, A. J. and G. R. Neumann, “The Job Ladder,” in H. Bunzel, B. J. Christensen, G. R. Neumann and J.-M. Robin, eds., “Structural Models of Wage and Employment Dynamics,” Elsevier, 2006.

Brown, C., “Equalizing Difference in the Labor Market,” *Quarterly Journal of Economics*, 94(1):113–134, 1980.

Burdett, K., “A Theory of Employee Job Search and Quit Rates,” *American Economic Review*, 68(1):212–220, 1978.

Dey, M. and C. Flinn, “Household Search and Health Insurance Coverage,” *Journal of Econometrics*, 145(1–2):43–63, 2008.

- Dey, M. S. and C. J. Flinn, "An Equilibrium Model of Health Insurance Provision and Wage Determination," *Econometrica*, 73(2):571–627, 2005.
- Eckstein, Z., S. Ge and B. Petrongolo, "Job and Wage Mobility with Minimum Wages and Imperfect Compliance," *Journal of Applied Econometrics*, forthcoming, 2009.
- Flabbi, L. and A. Moro, "The Effect of Job Flexibility on Female Labor Market Outcomes: Estimates from a Search and Bargaining Model," 2010, working paper.
- Flinn, C., "Labour market structure and inequality: A comparison of Italy and the US," *Review of Economic Studies*, 69(3):611–645, 2002.
- Goddeeris, J. H., "Compensating Differentials and Self-Selection: An Application to Lawyers," *Journal of Political Economy*, 96(2):411–428, 1988.
- Goffe, W. L., G. D. Ferrier and J. Rogers, "Global optimization of statistical functions with simulated annealing," *Journal of Econometrics*, 60(1-2):65 – 99, 1994, ISSN 0304-4076, doi:DOI: 10.1016/0304-4076(94)90038-8.
- Gørgens, T., "Reservation Wages and Working Hours for Recently Unemployed US Women," *Labour Economics*, 9(1):93–123, 2002.
- Gronberg, T. J. and W. R. Reed, "Estimating Workers' Marginal Willingness to Pay for Job Attributes Using Duration Data," *The Journal of Human Resources*, 29(3):911–931, 1994.
- Gruber, J. and B. C. Madrian, "Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature," in C. McLaughlin, ed., "Health Policy and the Uninsured," pp. 97–178, Urban Institute, Washington, DC, 2004.
- Hwang, H.-s., W. R. Reed and C. Hubbard, "Compensating Wage Differentials and Unobserved Productivity," *Journal of Political Economy*, 100(4):835–858, 1992.
- Jolivet, G., F. Postel-Vinay and J.-M. Robin, "The Empirical Content of the Job Search Model: Labor Mobility and Wage Distributions in Europe and the US," *European Economic Review*, 50:877–907, 2006.

- Keane, M. P. and K. I. Wolpin, "The Career Decisions of Young Men," *Journal of Political Economy*, 105(3):473–522, 1997.
- Pergamit, M. R., C. R. Pierret, D. S. Rothstein and J. R. Veum, "The National Longitudinal Surveys," *Journal of Economic Perspectives*, 15(2):239–253, 2001.
- Postel-Vinay, F. and J.-M. Robin, "The Distribution of Earnings in an Equilibrium Search Model with State Dependent Offers and Counter-Offer," *International Economic Review*, 43(4):989–1016, 2002.
- Smith, A., "Estimating Nonlinear Time-Series Models Using Simulated Vector Autoregressions," *Journal of Applied Econometrics*, 8(S):S63–84, 1993.
- Stern, S., "Estimating a Simultaneous Search Model," *Journal of Labor Economics*, 7(3):348–369, 1989.
- , "Simulation-based estimation," *Journal of Economic Literature*, 35(4):2006–2039, 1997.
- Sullivan, P., "A Dynamic Analysis of Educational Attainment, Occupational Choices, and Job Search," *International Economic Review*, 51(1):289–317, 2010.
- Sullivan, P. and T. To, "Search and Compensating Wage Differentials," 2011, working paper.
- Thaler, R. and S. Rosen, "The Value of Saving a Life: Evidence from the Labor Market," in N. E. Terleckyj, ed., "Household Production and Consumption," pp. 265–302, National Bureau of Economic Research, New York, 1975.
- White, M. J., "Location Choice and Commuting Behavior Cities with Decentralized Employment," *Journal of Urban Economics*, 24:129–152, 1988.
- Wolpin, K. I., "The Determinants of Black-White Differences in Early Employment Careers: Search, Layoffs, Quits, and Endogenous Wage Growth," *The Journal of Political Economy*, 100(3):535–560, 1992.

Yamaguchi, S., "Job Search, Bargaining, and Wage Dynamics," *Journal of Labor Economics*, forthcoming, 2010.