

Wage Dispersion and Search Behavior *

Robert E. Hall

Hoover Institution and Department of Economics

Stanford University

National Bureau of Economic Research

rehall@stanford.edu

Andreas I. Mueller

Columbia Business School

National Bureau of Economic Research

IZA

amueller@columbia.edu

July 18, 2013

Abstract

We use a rich new body of data on the experiences of unemployed job-seekers to determine the sources of wage dispersion and to create a search model consistent with the acceptance decisions the job-seekers made. From the data and the model, we identify the distributions of four key variables: offered wages, offered non-wage job values, the value of the job-seeker's non-work alternative, and the job-seeker's personal productivity. We resolve the tension between the fairly high dispersion of the values job-seekers assign to their job offers—which suggest a high value to sampling from multiple offers—and the fact that the job-seekers often accept the first offer they receive. An influential recent paper by Hornstein, Krusell, and Violante called attention to this tension. Our resolution rests on the job-ladder model, where unemployed job-seekers accept an offer that beats their non-work value, possibly as an interim job, because they continue to seek jobs while working.

*The Hoover Institution supported Hall's research. The paper is also part of the National Bureau of Economic Research's Economic Fluctuations and Growth Program. We are grateful to Steven Davis, Iourii Manovskii, Tamas Papp, Richard Rogerson, and Robert Shimer for valuable comments. An earlier version was titled "Viewing Job-Seekers' Reservation Wages and Acceptance Decisions through the Lens of Search Theory."

Search theory is firmly established as a useful way to think about unemployment and labor mobility. But it is well known that fitting search models to data on wages and labor flows is a challenge. Wages have huge dispersion among workers with similar observed characteristics. Traditional search theory hypothesized that job-seekers would keep considering wage offers until they found one high in the upper tail of the distribution of available opportunities. But the size of the flow of searchers out of unemployment into jobs suggests that searchers are leaving money on the table by taking jobs long before it is likely that they have adequately sampled the upper tail. The addition of on-the-job search to the model relieves some of this tension, because job-seekers departing unemployment may do so by taking an interim job and continue to search for the dream job in the upper tail while employed in the interim job. Even then, models that calibrate the offer distribution to the distribution of wages across workers find that the exit rate from unemployment to jobs makes sense only if workers find unemployment virtually intolerable, else they would be more picky in their acceptance decisions. These points reflect the suspicions that search economists have harbored for some time. They recently came into sharp focus in an influential article, Hornstein, Krusell and Violante (2011) (HKV).

In this paper, we investigate a new data source with the aim of resolving the conflict between the interpersonal dispersion of wages and the unemployment-to-work flow. That source is a novel survey of unemployed workers in 2009 and 2010 that Alan Krueger and Mueller (KM) carried out—see Krueger and Mueller (2011) and Mueller (2013). The KM data permit a more refined measure of dispersion than do the data sources in earlier work—they liberate search theory from inferring the distribution of opportunities from the residuals of wage regressions. The data include prior wages collected from administrative sources and survey responses about reservation wages each week during a spell of unemployment, and the wages of job offers and of newly accepted jobs. Although earlier surveys have collected cross-section data on reservation wages, the KM survey is the first, as far as we know, that collects panel data on reservation wages. It is also the first U.S. source to match survey data and administrative data, we believe.

By design, the KM survey gathered information on the search decisions of only the unemployed. A re-employed worker can continue to search—job-to-job transitions account for about half of all hires in the U.S. economy. A reasonable strategy for the unemployed is to take an interim job and keep on searching for a more permanent job. Hall (1995) describes

a labor market operating in this mode. The job-ladder model, as in Burdett (1978), Burdett and Mortensen (1998), and Hagedorn and Manovskii (2013), formalizes the process. In that model, provided that the offer rate for employed workers does not fall short of the rate for unemployed job-seekers, the reservation wage for a job-seeker is just the flow value of unemployment. The reservation wage for a worker currently receiving a wage of w is w itself—a costless move to any job that beats the current wage is an improvement. If holding a job results in a lower offer rate or if the job-seeker incurs a cost upon changing jobs, the job-seeker will set a reservation wage above the flow value of unemployment to preserve the option value of unemployment. In section 4, we consider the simple job-ladder model with equal offer probabilities for unemployed and employed workers and no mobility cost, so the reservation wage is the flow value of unemployment. We mention the extension to the case of modest real-option value of the current job resulting from lower offer probabilities for employed job-seekers. Our results do not support an option value, because we find that unemployed job-seekers appear to accept jobs on terms comparable to their non-work options, leaving no gap to be filled with the option value. Because there is a good deal of uncertainty about the benchmark for non-work values inferred from preferences, we are cautious in drawing this conclusion.

The KM data, together with a reasonable set of assumptions, permit a solution to a problem that has significantly impeded research on labor search behavior. HKV, Section II, discuss the challenges in detail with many references. The problem is the lack of information on individual wages relative to personal productivity. Conditional on measures of personal characteristics available to the econometrician, wages have huge dispersion. As HKV observe, research that uses econometric residuals to measure wages relative to personal productivity has failed, so far, to find a model that fits all the restrictions that seem reasonable. In particular, the implied flow value of unemployment as a ratio to earnings, z , is far too low in models that use that approach. The dispersion of wages thought to be available to job-seekers is so high that only a spectacular aversion to unemployment can explain the observed rate at which job-seekers take jobs.

Our key assumption is that the wage-related variables measured in the KM survey are proportional to personal productivity. Under the proportionality assumption for the reservation wage, both the offered wage and the wage in the prior job are proportional to personal productivity. The covariance of their logs reveals the dispersion of personal productivity.

Then the difference between the log variances of the offered wage and the prior wage is the log variance of the offers facing a job-seeker with a given level of productivity. Thus the HM survey permits a new attack on the question that HKV pose but do not answer: Is it possible to build an empirical model conforming to the measured amount of dispersion in the offered wage and a reasonable flow value of unemployment?

An important extension of the standard search model adds a non-wage dimension to jobs. In the KM data, a job-seeker frequently accepts a job paying less than the previously stated reservation wage and, less frequently, rejects a job paying more than the reservation wage. We use the observed relation between the acceptance probability and the difference between the offered and reservation wages to infer the distribution of the non-wage value of job offers. This distribution has considerable dispersion. We believe that the principle of compensating differentials implies that offered wages are negatively correlated with non-wage job values, but, as yet, have not succeeded in estimating the magnitude of the correlation.

In our model, wage dispersion arises from five sources:

1. Workers differ in personal productivity
2. Workers with the same personal productivity receive heterogeneous wage offers
3. Workers with the same personal productivity receive job offers with heterogeneous non-wage values; these values are negatively but imperfectly correlated with the corresponding wage offers
4. Workers with the same personal productivity have heterogeneous values in non-market activities, so reservation wages are heterogeneous
5. Workers move up the job ladder according to the random outcomes of on-the-job search

The model posits four underlying distributions of variables that are not measured directly:

1. Personal productivity
2. Wage offers standardized for personal productivity
3. The non-wage value of a job offer standardized for personal productivity
4. The value of non-market activities standardized for personal productivity

The behavioral assumptions in the model are:

1. Unemployed job-seekers accept a job offer if the offered value—wage and non-wage values combined—exceeds the value of alternative non-market activities
2. Employed job-seekers accept a job offer if the offered value exceeds the value of the current job

The variables from the KM survey that we study are:

1. Offered wage
2. Reservation wage reported prior to the receipt of the job offer
3. Whether or not the job-seeker accepted the offer
4. The wage earned before the current spell of unemployment

Our model comes close to matching the empirical distribution of offered wages, reservation wages, and past actual wages, together with the observed frequency of acceptance of offers as a function of the amount by which the offered wage exceeds the reservation wage.

We find robust estimates of dispersion—log standard deviation—for three of the four underlying unobserved variables. These are 0.30 for the offered wage, 0.19 for the non-work value (value of alternative non-market activity), and 0.43 for personal productivity. The bootstrap standard errors of these estimates are all around 0.02. Conditional on an assumed value for the parameter that controls the extent of compensating wage differentials, we estimate the log standard deviation of the component of the non-wage value of job offers that is independent of the offered wage to be 0.56 with a standard error of 0.16. In an alternative specification with no compensating differential, the log standard deviation of the non-wage value is 0.88 with a standard error of 0.26.

HKV note that most empirical search models that appear to rationalize observed unemployment-to-employment flows invoke much too low a flow value of unemployment. The flow value is frequently negative. These models generally infer the value of job search from estimates of the dispersion of wage offers derived from cross-sectional data, where dispersion is high. Sampling from that distribution is highly valuable activity, which implies that people must truly hate unemployment to be in equilibrium while unemployed. In the results in this paper, where job-seekers sample from a distribution of job values with considerable dispersion, the reservation wage for unemployed job-seekers is nonetheless in line with values of the

flow value of unemployment derived from evidence on preferences and on the replacement rate for unemployment benefits.

HKV write, pp. 2894-5 [We have taken the liberty of changing their symbol for the relative value of non-market time, ρ , to z , as generally used in the DMP literature]:

A number of papers in the literature claim that the (on-the-job search) model is successful in simultaneously matching both the wage distribution and labor-market transition data (see, e.g., Christian Bontemps, Robin, and van den Berg 2000; Jolivet, Postel-Vinay, and Robin 2006)...the exercise is incomplete because it neglects the implications of the joint estimates of $F(w)$ and of the transition parameters for the relative value of nonmarket time z . The key additional “test” that we are advocating would thus entail using the estimated $F(w)$ in the reservation-wage equation and, given an estimate of w^* , backing out the implied value for z . In light of our results, we maintain that z would be often negative or close to zero.

This paper carries out the test that HKV recommend. Our specification passes the test nicely.

This version of the paper is preliminary in one important way: we approximate the underlying distributions of the key unobserved variables as log-normal. These are personal productivity, the offered wage given productivity, the independent component of the offered non-wage job value, and the value of non-work given personal productivity. In the concluding section, we describe some directions for additional work.

2 Related Research

See Hornstein et al. (2011) for an extensive discussion and many cites, notably Mortensen (2003), Rogerson, Shimer and Wright (2005), Bontemps, Robin and Berg (2000), Jolivet, Postel-Vinay and Robin (2006), Jolivet (2009), and Postel-Vinay and Robin (2002).

3 The KM Survey

For details on the KM survey, see Krueger and Mueller (2011) and Mueller (2013). The survey enrolled roughly 6,000 job-seekers in New Jersey who were unemployed in September 2009 and collected weekly data from them for several months. The authors also make use of

data from administrative records for the respondents, notably their wages on the jobs they held just prior to becoming unemployed. We follow Krueger and Mueller (2011) and restrict the sample to survey participants of ages 20 to 65.

3.1 Job offers

The KM survey asked respondents each week: “In the last 7 days, did you receive any job offers? If yes, how many?” The respondents in our sample received a total of 2,174 job offers in 37,609 reported weeks of job search. The ratio of the two, 0.058, is a reasonable estimate of the overall weekly rate of receipt of job offers.

For respondents who indicated that they received at least one job offer, the KM survey asked respondents: “What was the wage or salary offered (before deductions)? Is that per year, per month, bi-weekly, weekly or per hour?” In cases where respondents reported that they received more than one offer in a given week, the survey asked the offered wage only for the best wage offer. Figure 1 reports the kernel density of the hourly offered wage for our sample of 1,153 job offers. The sample is restricted to cases where details of the offer (including the wage) and a reservation wage from a previous interview were available. We use the same sample below when we compute the acceptance frequency conditional on the difference between the log of the offered wage and the log of the reservation wage from a previous interview.

To find the underlying dispersion of the offered wage among job-seekers with standardized personal productivity, we remove the dispersion implied by our estimated distribution of personal productivity.

3.2 Reservation wage

Each week, the respondents in the KM survey answered a question about their reservation wages: “Suppose someone offered you a job today. What is the lowest wage or salary you would accept (before deductions) for the type of work you are looking for?” We only use the first reservation wage observation available for each person in the survey so that the sample is representative of the cross-section of unemployed workers. We apply the same sample restrictions as Krueger and Mueller (2011) who exclude survey participants who reported working in the last seven days or already accepted a job offer at the time of the




Figure 1: Kernel Density of the Log Hourly Offered Wage, y

interview. Figure 2 shows the kernel density of hourly reservation wage for our sample of 4,138 unemployed workers.

The model interprets the reservation wage of an unemployed job-seekers as the value of the non-work option. The survey reveals the dispersion of the non-work value across all respondents. Much of that dispersion arises from the dispersion in personal productivity, which we take to influence the non-work value in the same proportion as for work. To find the underlying dispersion of the non-work value among people with standardized personal productivity, we remove the dispersion implied by our estimated distribution of personal productivity.

3.3 Acceptance

Notice that many job-seekers accept job offers that pay less than the job-seeker’s previously reported reservation wage and that some do the reverse, rejecting an offer that pays more than the reservation wage. Our model accounts as a general matter for the fact that the offered wage does not control the acceptance decision by invoking a non-wage value—job-seekers accept jobs paying less than the reservation wage because these jobs are desirable in




Figure 2: Kernel Density of the Log Hourly Reservation Wage, r

other respects. The model accounts for the bias toward acceptance by treating the reported reservation wage as referring to a job with below-normal non-wage value.

We study the acceptance probability as a function of the difference between the log of the offered wage and the log of the reservation wage. We use the reservation wage reported in a previous interview to exclude the possibility that survey participants changed their reservation wage based on the job offer. Mueller (2013) gives a detailed analysis of the the acceptance frequency in the survey. The job acceptance frequency rises with $d = y - r$. The average frequency of job acceptance in our sample is 71.9 percent. In 20.9 percent of the cases, respondents indicated that they had not yet decided whether to accept the job offer or not.

To deal with the problem of missing data for acceptance of some job offers, we make use of administrative data on exit from unemployment insurance. UI exit is a potentially useful but imperfect indicator of acceptance, for four reasons: (1) A delay occurs between job acceptance and UI exit. (2) An exit from the UI system may relate to a different offer from the one reported in the survey. (3) UI exit data are censored at the point of UI exhaustion, as the data do not track recipients after they exhaust benefits. (4) An unemployed worker may perform limited part-time work while receiving benefits and thus acceptances of such offers




Figure 3: Smoothed Acceptance Frequency, A , as a Function of the Difference between the Log Hourly Offered Wage and the Log Hourly Reservation Wage, $d = y - r$

will not be reflected in an exit from the UI system. Mueller shows that the rate of UI exit for those were undecided was almost exactly half way between the rate of UI exit for those who accepted the offer and the rate of UI exit those who rejected the offer. Notwithstanding the imperfect relation between exits and acceptances of offers, we believe that this estimate is the best available, so we create an indicator variable A that takes on the value zero for a rejected offer, 0.5 for an offer for which the respondent was undecided, and 1 for an accepted offer.

Figure 3 shows the acceptance frequency smoothed in two ways: (1) as the fitted values from a regression of A on a 6th-order polynomial in $y - r$ and (2) as the fitted values from a locally weighted regression (LOWESS) with bandwidth 0.3. The figure runs from first-percentile value of d to the 99th percentile value. Values outside that range are inherently unreliable for any smoothing method.

In our approach to estimation, the acceptance function identifies the dispersion of the non-wage value. The fact that many jobs are accepted that pay well below the reported reservation shows that fairly large positive non-wage values are common. We characterize




Figure 4: Kernel Density of the Log Hourly Wage on the prior Job, w

the function by the acceptance rate at two values of d . These are sufficient to identify the bias and standard deviation of the non-wage job value.

3.4 Prior wage

Our model views the prior wage as the result of search during an earlier spell of unemployment, the acceptance of the first job offered that exceeded the reservation job value (combining wage and non-wage components), and the acceptance of later offers, each of which exceeded the job value of the prior job.

Figure 4 shows the kernel density of the hourly wage on the prior job for our sample of 4,138 unemployed workers. The wage is computed from administrative data on weekly earnings during the base year, which typically consist of the first four of the five quarters before the date of the UI claim, and survey data on weekly hours for the previous employment. Hours on the previous job might not perfectly overlap with the period of the base year. Moreover, roughly 15 percent of the respondents answered that hours varied on their previous jobs and we imputed their hours based on demographic characteristics as in Krueger and Mueller (2011). For these reasons, the *hourly* previous wage includes some measurement error despite the fact that weekly earnings are taken from administrative data.

<i>Moment</i>	<i>Symbol</i>	<i>Value</i>
Mean previous wage	$m_{\hat{w}}$	2.87
Mean offered wage	$m_{\hat{y}}$	2.75
Mean reservation wage	$m_{\hat{r}}$	2.82
Standard deviation of previous wage	$s_{\hat{w}}$	0.583
Standard deviation of offered wage	$s_{\hat{y}}$	0.525
Standard deviation of reservation wage	$s_{\hat{r}}$	0.474
Covariance of offered wage and reservation wage	$c_{\hat{y},\hat{r}}$	0.183
Covariance of offered wage and previous wage	$c_{\hat{y},\hat{w}}$	0.183
Covariance of reservation wage and previous wage	$c_{\hat{r},\hat{w}}$	0.199
Acceptance frequency at $d_1 = -1$	\hat{A}_1	0.262
Acceptance frequency at $d_2 = 0.5$	\hat{A}_2	0.856

Table 1: Moments to Match

In the model, the distribution of the prior wage depends on all four unobserved distributions. We carry out a rather complicated calculation of the distribution and match it to the observed one. We update the wage by 3.1 percent to adjust for the time elapsed between the measurement of the respondents' earnings in March 2008 to the median survey month, November 2009, based on the Bureau of Labor Statistics Employment Cost Index, which is adjusted for changes in the composition of employment.

3.5 Moments

Table 1 shows the moments of the data that we try to match with the model. The last two moments are taken from the smoothed acceptance frequency, and more precisely, the predicted values of the polynomial of degree 6 evaluated at two values of d .

4 Model

We formulate the model in terms of the logs of the variables. We let x be the log of the personal productivity of a worker and assume it is known to employers and to the worker. When we refer to standardization for productivity, we mean the population with $x = 0$. Much of our discussion involves standardized variables. Where necessary to avoid confusion, we

use a hat ($\hat{\cdot}$) when we are referring to variables or their distributions in the entire population, not standardized for productivity. We denote the log of the reported reservation wage as r , the log of the offered wage as y , the log of the non-wage value as n , and the log of the non-work value as h . We proceed under the *proportionality-to-productivity hypothesis*:

The distributions of $\hat{r} - x$, $\hat{y} - x$, $\hat{n} - x$, and $\hat{h} - x$ in the population with personal productivity x are the same as the distributions of r , y , n , and h in the population with $x = 0$.

The most controversial aspect of this hypothesis is that non-market productivity is higher by the entire amount of market productivity in the population with higher values of x . Low- x populations are not systematically more choosy about taking jobs than are high- x populations. While this assumption obviously fails if applied across the entire population, it appears reasonable in a sample of workers eligible for unemployment compensation.

We use the term “offer” to describe a job-seeker’s encounter with a definite opportunity to take a job. Nothing in this paper requires that employers make firm job offers and that job-seekers then make up-or-down decisions. The job-seeker’s decision problem, upon finding a job opportunity with a joint surplus for the job-seeker and employer, is the same whether the employer is making a single firm offer, or the parties make a Nash bargain, or they engage in alternating-offer bargaining. In the bargaining cases, the job-seeker knows in advance what the bargain will be, so the job-seeker decides whether to engage in bargaining just as he would if he faced a firm offer. That said, the survey included a question about the nature of the job offer and in the majority of cases, the employer did make a firm offer.

4.1 Acceptance

We treat the reported reservation wage as the lowest wage a job-seeker will accept for a job with a reference level of its non-wage value of \bar{n} . The lowest job value an unemployed job-seeker will accept is h , so we have

$$r + \bar{n} = h. \tag{1}$$

The principle of compensating wage differentials suggests that the correlation between wage offers y and non-wage values n should be negative—employers offer lower wages for jobs with favorable non-wage values. The correlation is not perfect, however, because there is a personal dimension to the non-wage value that the firm may ignore, under a posted-wage

policy, or respond to only partially, in a bargained-wage policy. For example, commuting cost varies across individual workers. For this reason, we assume that the non-wage value n comprises (1) a component η that is uncorrelated with the other fundamentals and (2) a component that is the negative of a fraction κ of the offered wage minus its mean:

$$n = \eta - \kappa(y - \mu_y). \quad (2)$$

Recall that $d = y - r$ is the difference between the offered wage and the reservation wage. The probability of acceptance of a job offer is

$$\begin{aligned} \tilde{A}(d, y) &= \text{Prob}[y + n \geq h] = \text{Prob}[\eta \geq \bar{n} - [d - \kappa(y - \mu_y)]] \\ &= 1 - F_\eta(\bar{n} - d + \kappa(y - \mu_y)). \end{aligned} \quad (3)$$

Integrating over y and h such that $h = y - d + \bar{n}$ gives

$$\begin{aligned} A(d) &= \int \tilde{A}(d, y) dF_y(y | y - h = d - \bar{n}) \\ &= \int [1 - F_\eta(\bar{n} - d + \kappa(y - \mu_y))] \frac{dF_h(y - d + \bar{n}) dF_y(y)}{\int dF_h(y - d + \bar{n}) dF_y(y)}. \end{aligned} \quad (4)$$

Note that in the case where $\kappa = 0$, the probability of acceptance of a job offer is

$$A(d) = \text{Prob}[y + n \geq h] = \text{Prob}[y - r \geq \bar{n} - n] = \text{Prob}[n \geq \bar{n} - d] \quad (5)$$

Then we can write

$$A(d) = 1 - F_n(\bar{n} - d) \quad (6)$$

Thus, letting $n = \bar{n} - d$, we can calculate F_n directly from the acceptance function:

$$F_n(n) = 1 - A(\bar{n} - n). \quad (7)$$

In this case, F_n , a theoretical distribution referring to individuals with productivity $x = 0$, turns out to be equal to the observed function $A(y - r)$ relating the acceptance probability to the gap between the offered wage y and the reservation wage r , both observed. This identification rests on the proportionality-to-productivity hypothesis and our assumption about what respondents mean by their reservation wages. Notice that $A(y - r)$ is the same for all values of x , because subtracting x from both y and r leaves the difference unchanged.

5 Results

We present our results in two steps. The first does not make use of data on the prior wage and, correspondingly, does not use the job-ladder model to calculate the implied moments of the distribution of the prior wage. The second extends the analysis to include the job-ladder model to match the moments of the prior wage.

We take the distributions of the four fundamental variables, y , η , h , and x to be log-normal. Two of them, η and x , are normalized to have means of zero. The other two means, μ_y and μ_h , are parameters to estimate. The standard deviations, σ_y , σ_η , σ_h , and σ_x , are also parameters. The reference level of the non-wage job value, \bar{n} and the relation of the non-wage value n to the offered wage y , κ , are the final two parameters, for a total of 8.

We use the following 7 data moments: the means $m_{\hat{y}}$ and $m_{\hat{r}}$, standard deviations $s_{\hat{y}}$ and $s_{\hat{r}}$ of the two directly observed variables, the covariance $c_{\hat{y},\hat{r}}$, and the two values \hat{A}_1 and \hat{A}_2 of the acceptance frequency. We infer the moments of η by picking two values d_1 and d_2 and solving equation (4) for \bar{n} and σ_η .

The model has 8 parameters to estimate: μ_y , μ_h , σ_y , σ_h , σ_x , σ_n , \bar{n} , and κ . We consider two values of κ and estimate the remaining 7 parameters by solving the exactly identified system. The values are $\kappa = 0$, for no compensating variations—no tendency for higher wage offers y to come with jobs with lower non-wage values n —and $\kappa = 0.5$, a moderate amount of compensating variation in wages to offset non-wage values.

The observed moments and their counterparts in the model are:

$$m_{\hat{y}} = \mu_y \tag{8}$$

$$m_{\hat{r}} = \mu_h - \bar{n} \tag{9}$$

$$s_{\hat{y}} = \sqrt{\sigma_y^2 + \sigma_x^2} \tag{10}$$

$$s_{\hat{r}} = \sqrt{\sigma_h^2 + \sigma_x^2} \tag{11}$$

$$c_{\hat{y},\hat{r}} = \sigma_x^2 \tag{12}$$

$$A(d_i) = \int (1 - \Phi(\bar{n} - d_i + \kappa(y - \mu_y), 0, \sigma_\eta)) \frac{\phi(y - d_i + \bar{n}, \mu_h, \sigma_h) \phi(y, \mu_y, \sigma_y) dy}{\int \phi(y - d_i + \bar{n}, \mu_h, \sigma_h) \phi(y, \mu_y, \sigma_y) dy}, i = 1, 2 \quad (13)$$

We solve for μ_y , σ_y , σ_h , and σ_x in the obvious ways. Then we substitute $\mu_h = m_{\hat{r}} + \bar{n}$ into both instances of equation (35) and solve them jointly for σ_η and \bar{n} . Table 2 shows the results for the two candidate values of κ .

To measure sampling variation, we use a bootstrap strategy. In our actual estimation procedure, we compute our moments from two different samples: we take the moments $m_{\hat{r}}$ and $s_{\hat{r}}$ from the first interview for all unemployed workers in the survey who were not working or had not yet accepted a job offer, whereas we take $m_{\hat{y}}$, $s_{\hat{y}}$, $c_{\hat{y}, \hat{r}}$, \hat{A}_1 and \hat{A}_2 from the sample of 1,153 job offers with information on the offered wage and on the lagged reservation wage. The standard bootstrap strategy applies to single samples. Accordingly, we use only the smaller sample; the resulting bootstrap distribution provides an upper bound on the dispersion of our actual sampling distribution. This smaller sample appears not to be biased, as $m_{\hat{r}} = 2.83$ and $s_{\hat{r}} = 0.47$, which are almost identical to the estimates in the bigger sample. In the smaller sample, $m_{\hat{w}} = 2.86$, which is also very close to the estimate in the bigger sample, and $s_{\hat{w}} = 0.61$, which is a little higher than in the bigger sample. For the bootstrap, we thus sample with replacement from the 1,153 job offers, and compute the moments in the data and in the model for 20 draws. The resulting sampling distribution has higher sampling dispersion than our actual results, which draw in part from additional data not simulated in the bootstrap.

Table 2 shows the results for the two designated values of κ . The bootstrap standard errors are relatively large for μ_h , σ_η and \bar{n} because the estimation of these moments relies on the acceptance function. The standard errors of the other moments are relatively small.

Our findings for the case of no compensating wage differentials ($\kappa = 0$) are:

1. The mean non-work value is 37 log points below the mean of the offered wage—as we discuss later, this value is in the range calculated in Hall and Milgrom (2008) and implies a reasonable high value of the ratio of the non-work value to the marginal product of labor. Unlike earlier work, we find no contradiction between the dispersion of offered job values and the frequency of acceptance.
2. The dispersion in the offered wage among people with the same personal productivity is moderate: $\sigma_y = 0.304$.

<i>Parameter</i>	<i>Explanation</i>	<i>Value of compensating differential coefficient</i>			
		$\kappa = 0$		$\kappa = 0.5$	
		Estimate	(s.e.)	Estimate	(s.e.)
μ_y	Mean of wage offers	2.75	(0.02)	2.75	(0.02)
μ_h	Mean of non-work values	2.38	(0.21)	2.51	(0.13)
σ_y	Standard deviation of the offered wage	0.304	(0.019)	0.304	(0.019)
σ_h	Standard deviation of the non-work value	0.190	(0.028)	0.190	(0.028)
σ_η	Standard deviation of the independent component of non-wage value of job offer	0.882	(0.264)	0.559	(0.161)
σ_x	Standard deviation of personal productivity	0.428	(0.018)	0.428	(0.018)
\bar{n}	Reference level of the non-wage value when forming the reservation wage	-0.44	(0.21)	-0.31	(0.13)

Table 2: Parameter Estimates without Matching Moments of Prior Wage

3. The dispersion in the value of non-work among people with the same personal productivity is low: $\sigma_h = 0.190$.
4. The dispersion of the non-wage job value is substantial: $\sigma_\eta = 0.882$.
5. The dispersion of personal productivity is moderate: $\sigma_x = 0.428$.
6. The reference value of the non-wage value of a job when reporting the reservation wage is 44 percent below the median: $\bar{n} = -0.44$.

Our findings for the case of positive compensating wage differentials ($\kappa = 0.5$) are:

1. The mean non-work value is 24 log points below the mean of the offered wage—this also appears to be in the range calculated in Hall and Milgrom (2008).
2. The dispersion in the offered wage among people with the same personal productivity is the same as in the previous case: $\sigma_y = 0.304$.
3. The dispersion in the value of non-work among people with the same personal productivity is the same as in the previous case: $\sigma_h = 0.190$.
4. The dispersion of the non-wage job value is high but quite a bit lower than in the previous case: $\sigma_\eta = 0.559$.
5. The dispersion of personal productivity is the same as in the previous case: $\sigma_x = 0.428$.
6. The reference value of the non-wage value of a job when reporting the reservation wage is 31 percent below the median, probably a more reasonable value than in the previous case: $\bar{n} = -0.31$.

6 Estimation Including the Prior Wage and the Job-Ladder Model

Now we extend the model to include the distribution of the prior wage and we add the prior wage to the moments to be matched.

6.1 The distribution of wages in the job-ladder model

We let $F_w(w)$ be the cdf of wages among workers with $x = 0$. An individual draws a non-work value h at the outset. A personal state variable records whether the individual is unemployed or employed. The flow value of the current job, $v = w + n$, is a second personal state variable for the employed. The probability of receiving a job offer in a given period is λ . An offer comprises a wage y and a non-wage value n . Jobs end because of the arrival of a better offer or through exogenous separation and a drop to the bottom of the ladder. The latter occurs with fixed probability s and sends the worker into unemployment at the bottom of the ladder.

Define

$$F_v(v) = \int f_{y,n} \left(\frac{v - \eta - \kappa \bar{y}}{1 - \kappa}, \eta \right) d\eta, \quad (14)$$

the cdf of a job offer with value v . Here $f_{y,n}(y, n)$ is the joint density of y and n . The probability in one period that an unemployed worker with non-work value h will remain unemployed in the next period is

$$T_{uu}(h) = 1 - \lambda(1 - F_v(h)). \quad (15)$$

The probability than an unemployed individual will be at work in the succeeding period with a job value not greater than v' is

$$T_{ue}(v'|h) = \lambda(F_v(v') - F_v(h)). \quad (16)$$

The probability that an employed worker will be unemployed in the next period is

$$T_{eu} = s. \quad (17)$$

The probability than an employed individual will remain employed at the same job value with value v is

$$T_{ee}(v|v) = (1 - s)[1 - \lambda(1 - F_v(v))]. \quad (18)$$

The probability than an employed individual will move to a better job with value $v' > v$ is

$$T_{ee}(v'|v) = (1 - s)\lambda(F_v(v') - F_v(v)). \quad (19)$$

Let q be the compound state variable combining a binary indicator for unemployment/employment and the job value v and let $T(q'|q, h)$ be its transition cdf derived above. The

stationary distribution of q , $F_q(q|h)$ satisfies the invariance condition,

$$F_q(q'|h) = \int T(q'|q, h) dF_q(q|h). \quad (20)$$

Throughout, an integral without limits of integration is over the support of the integrand. The ergodic distribution of the job value for employed workers, $F_v(v|h)$, is the conditional distribution of v for values of q for employed workers.

The cdf of the wage, w , conditional on the job value v , is

$$F_w(w|v) = \frac{\int^w f_{y,\eta}(y, v - y(1 - \kappa) - \kappa\bar{y}) dy}{\int f_{y,\eta}(y, v - y(1 - \kappa) - \kappa\bar{y}) dy}. \quad (21)$$

The implied ergodic distribution for the wage is

$$F_w(w|h) = \int F_w(w|v) dF_v(v|h). \quad (22)$$

Finally, the distribution in the population with $x = 0$ is the mixture,

$$F_w(w) = \int F_w(w|h) dF_h(h) \quad (23)$$

and the distribution in the overall population is the mixture,

$$F_{\hat{w}}(\hat{w}) = \int F_w(\hat{w} - x) dF_x(x). \quad (24)$$

6.2 Calibration

We take the offer arrival rate to be $\lambda = 0.058$ from the survey. We calculate the entry rate to unemployment, s , as

$$s = \frac{u\lambda a}{1 - u} = 0.0041 \text{ per week}, \quad (25)$$

the weekly rate consistent in stationary stochastic equilibrium with an unemployment rate of $u = 0.09$ and the observed job-finding rate.

6.3 Results

We use all of the 11 data moments in Table 1: the means $m_{\hat{w}}$, $m_{\hat{y}}$, and $m_{\hat{r}}$, the standard deviations $s_{\hat{w}}$, $s_{\hat{y}}$, and $s_{\hat{r}}$ of the three directly observed variables, the covariances $c_{\hat{y},\hat{r}}$, $c_{\hat{y},\hat{w}}$, and $c_{\hat{r},\hat{w}}$, and the two values \hat{A}_1 and \hat{A}_2 of the acceptance frequency.

The model has 8 parameters to estimate: μ_y , μ_h , σ_y , σ_h , σ_x , σ_η , \bar{n} , and κ . We let $\mu_w(\theta)$ and $\sigma_w(\theta)$ denote the mean and standard deviation of the wage among the population with

$x = 0$ and $\gamma(\theta)$ denote the covariance of the prior wage and the reservation wage implied by the wage-ladder model with parameter vector θ .

The observed moments and their counterparts in the model are:

$$m_{\hat{y}} = \mu_y \quad (26)$$

$$m_{\hat{r}} = \mu_h - \bar{n} \quad (27)$$

$$m_{\hat{w}} = \mu_w(\theta) \quad (28)$$

$$s_{\hat{y}} = \sqrt{\sigma_y^2 + \sigma_x^2} \quad (29)$$

$$s_{\hat{r}} = \sqrt{\sigma_h^2 + \sigma_x^2} \quad (30)$$

$$s_{\hat{w}} = \sqrt{\sigma_w^2(\theta) + \sigma_x^2} \quad (31)$$

$$c_{\hat{y}, \hat{r}} = \sigma_x^2 \quad (32)$$

$$c_{\hat{y}, \hat{w}} = \sigma_x^2 \quad (33)$$

$$c_{\hat{w}, \hat{r}} = \gamma(\theta) \quad (34)$$

$$A(d_i) = \int (1 - \Phi(\bar{n} - d_i + \kappa(y - \mu_y), 0, \sigma_\eta)) \frac{\phi(y - d_i + \bar{n}, \mu_h, \sigma_h) \phi(y, \mu_y, \sigma_y) dy}{\int \phi(y - d_i + \bar{n}, \mu_h, \sigma_h) \phi(y, \mu_y, \sigma_y) dy}, i = 1, 2 \quad (35)$$

We ask, how well do the previous parameter estimates result in matching the additional moments? Table 3 answers this question. Again we consider two values of the parameter κ .

1. We are not far off in matching $m_{\hat{w}}$ with $\kappa = 0.5$.
2. The one free parameter, κ , does not seem capable of matching the standard deviation of the prior wage, $s_{\hat{w}}$. The fitted value is about 0.06 log points below the actual value of the moment for both values of κ .

<i>Moment</i>	<i>Symbol</i>	<i>Actual value</i>	<i>Fitted value</i>			
			$\kappa = 0$		$\kappa = 0.5$	
			Estimate	(s.e.)	Estimate	(s.e.)
Mean previous wage, adjusted for intervening wage growth	$m_{\hat{w}}$	2.90	2.90	(0.06)	2.87	(0.05)
Standard deviation of previous wage	$s_{\hat{w}}$	0.583	0.521	(0.015)	0.522	(0.016)
Covariance of offered wage and previous wage	$c_{\hat{y},\hat{w}}$	0.183	0.183	(0.018)	0.183	(0.018)
Covariance of reservation wage and previous wage	$c_{\hat{r},\hat{w}}$	0.199	0.184	(0.015)	0.184	(0.017)

Table 3: Actual and Fitted Values of the Four Moments not Included in Table 2

3. The fitted value of the covariance of the offered and prior wages, $c_{\hat{y},\hat{w}}$ fits the observed value perfectly.
4. The fitted value of the covariance of the reservation and prior wages, $c_{\hat{r},\hat{w}}$ fits the observed value reasonably closely.
5. The bootstrap dispersion of the fitted values is quite small in all cases.

7 The Flow Values of Non-Work to Work

Hornstein et al. (2011) take earlier authors to task for failing to note that search models imply an extremely low, even negative, value of non-work. The essential point is that the dispersion of offered wages is high enough to justify sampling a large number of offers before picking the best, so that the observed time to acceptance only makes sense if waiting to go to work is painful. They note that the problem remains, though less acute, with on-the-job search.

In the search-and-matching literature whose canon is Mortensen and Pissarides (1994), a variable often called z describes the relation between the flow value of remaining out of the labor market and the flow value of participating in the market. z is often taken as a parameter in these models. It is the ratio of the flow value of non-work to the mean of the marginal product of labor.

		$\kappa=0$	$\kappa=0.5$
1	Value of non-work at median for $x=0$, $\exp(\mu_h)$	10.82	12.28
2	Earnings while employed, median for $x=0$, $\exp(m_{\hat{w}})$	18.14	17.62
3	Ratio of earnings to marginal product	0.985	0.985
4	Implied marginal product	18.42	17.88
5	Ratio of value of non-work to marginal product	0.59	0.69

Table 4: Ratio of the Flow Value of Non-Work to the Marginal Product of Labor

7.1 The implied value of z

Table 4 shows the calculation of z for the two cases in Table 2. Line 1 shows the value of non-work as estimated in that table, expressed in dollars per hour at the median of the distribution of h . Line 2 shows the median wage. Line 3 gives an estimate of the ratio of the wage to the marginal product from a typical calibration of an MP-type model, from Hall and Milgrom (2008). Line 4 gives the implied value of the marginal product of labor. Line 5 reports the resulting value of z , the ratio of the value of non-work to the marginal product. The values are robustly positive. They are close to the value of z found in Hall and Milgrom, which is 0.71. Our view of wage dispersion easily passes the test that Hornstein et al. (2011) advocate, that the overall picture not impute an unrealistic dislike of time off work.

We avoid taking much credit for the fairly close match to the Hall-Milgrom calibration, because that calibration is sensitive to the replacement rate for unemployment insurance. Our sample is drawn from workers who receive benefits, so the replacement rate is likely to be higher than the 25 percent that Hall and Milgrom assume. The corresponding value of z is much higher—about equal to the median wage—with the 50-percent replacement rate we believe is more realistic. We do not believe that z could possibly be that high, because it would imply, in the context of our model, that almost half the employed workers were earning less than their non-work values. Rather, it shows that the calibration does not give reasonable results with a higher replacement rate.

To the extent that our estimate that z in the range of 0.6 to 0.7 coincides or falls short of values derived from preferences, our assumption of the simple job-ladder model, with no option value to remaining unemployed, receives support. An alternative model would have a reservation wage for unemployed searchers higher than the flow value of time spent unemployed, reflecting the option value of job search. In that model, search would be more effective for the unemployed than for the employed.

7.2 The reservation wage over the period of unemployment

Earlier research on reservation wages had found that across job-seekers of varying amounts of elapsed unemployment, those with longer durations had lower ratios of reservation wages to past wages. This finding suggested that individual job-seekers lowered their reservation wages as time passed without taking a new job. The KM survey found that this suggestion does not hold among job-seekers in general. In regressions with the log of the ratio of the reservation wage to the prior wage as the left-hand variable, and with fixed effects for respondents, there is almost no downward trend in the ratio. Mueller (2013) finds that the downward trend is 0.6 percent for each 10 weeks of additional time since job loss, with a standard error of 0.6 percentage points. Among job-seekers aged 51 to 65, the downward trend is 2.6 percent for each 10 weeks, with a standard error of 0.7 percentage points.

The stability of the reservation wage during a spell of unemployment supports the view that job-seekers regard the flow value of unemployment as the reservation wage. The perceived option value of unemployment, if substantially positive, would surely decline with disappointing search results, so job-seekers would lower their reservation wages as unemployment continued.

8 Concluding Remarks and Next Steps

The KM data support a view of the labor market based on the job-ladder model that appears to meet all the criteria that labor economists and macroeconomists have developed in earlier work on job search. Job-seekers adopt reservation values that govern their acceptance decisions. After adjustment for an upward bias, the reservation wages seem to be rational, in that they reflect both the value job-seekers enjoy when not working and the fact that there is a low real option value to search, in the sense that search can continue after starting a job. That job will have an interim character if the value of the job is low, because the

likelihood of a better offer is fairly high. The problem that has arisen in much past work does not arise in this view. Earlier, as HKV stress, there seemed to be a yawning discrepancy between the high dispersion of wages across workers, even after adjustment for observed characteristics, and the high flow of job-seekers into jobs, which suggests that job-seekers see offered wages as close to uniform, so that declining an offer has little chance of resulting in a better subsequent offer.

Obviously, one ingredient in the resolution of the discrepancy is the invocation of the simple job-ladder model, which robs job search of a real-option character and makes it rational to take the first job that comes along that beats the flow value of being unemployed. That point is well known. The demonstration that the simple job-ladder model generates a stationary distribution of earnings with reasonable dispersion is a step forward in establishing the plausibility of the job-ladder model. And the finding that reported reservation wages have an upward bias relative to the actual acceptance decisions of job-seekers is important in making all the pieces fit together.

The simple job-ladder model makes the potentially extreme assumptions that holding a job is not a major impediment to searching for a better one and that job mobility has a low joint cost to worker and employer. The first assumption may be reasonable. Time-use data, including the detailed data in the HM survey, suggest that search rarely consumes much time relative to normal amounts of work, even for jobless searchers. Most search seems to take the form of applying for jobs and waiting to see what happens. Many observers have suggested that applicants already holding jobs receive favorable treatment from potential new employers—unemployment is an adverse signal of an applicant’s quality.

This discussion follows HKV, except that they believe that the amount of implied dispersion in the offer distribution in the job-ladder model is small, whereas we find that it is enough, possibly, to resolve the puzzle that arises in the traditional search model. They write, “Therefore, through the lenses of models with on-the-job search, deviations from the law of one price are more significant, albeit still fairly minor in absolute size” (page 2875). Our review of the KM data finds deviations from the law of one price—that is, dispersion of the purely personal or frictional component of the wage—to have a log-standard deviation of 0.30, a fairly large amount, though much smaller than the cross-sectional standard deviation of 0.58.

The immediate next step in our research is to calculate the best-fitting parameters for the full model. Though this is a straightforward exercise, we do not think it likely to go much beyond our current findings.

We believe it should be possible to match the observed standard deviation of the prior wage in a realistic way by adding measurement errors in the wage estimates.

So far, we have not found a way to identify κ and shed light on the question of the magnitude of compensating wage differentials. We are considering additional moments and alternative specifications with errors of measurement that may help on this point.

Even if we cannot pin down κ , we have found robust estimates of the dispersion of three of the four distributions of the model and have bracketed the dispersion of the fourth, the non-wage value of job offers. And our model answers the challenge in Hornstein et al. (2011) to reconcile the dispersion of offered job values to the acceptance decisions of unemployed searchers.

We have mapped out a strategy for more flexible parametrization of the four underlying distributions, as mixtures of log-normal distributions. We do not expect that the results will be much different from those presented in this version of the paper.

References

- Bontemps, Christian, Jean-Marc Robin, and Gerard J. van den Berg, “Equilibrium Search with Continuous Productivity Dispersion: Theory and Nonparametric Estimation,” *International Economic Review*, 2000, 41 (2), 305–358.
- Burdett, Kenneth, “A Theory of Employee Job Search and Quit Rates,” *American Economic Review*, 1978, 68 (1), 212–220.
- and Dale T. Mortensen, “Wage Differentials, Employer Size, and Unemployment,” *International Economic Review*, May 1998, 39 (2), 257–273.
- Hagedorn, Marcus and Iourii Manovskii, “Job Selection and Wages over the Business Cycle,” *American Economic Review*, 2013. Forthcoming.
- Hall, Robert E., “Lost Jobs,” *Brookings Papers on Economic Activity*, 1995, 1995 (1), 221–273.
- and Paul R. Milgrom, “The Limited Influence of Unemployment on the Wage Bargain,” *American Economic Review*, September 2008, 98 (4), 1653–1674.
- Hornstein, Andreas, Per Krusell, and Giovanni L. Violante, “Frictional Wage Dispersion in Search Models: A Quantitative Assessment,” *American Economic Review*, 2011, 101 (7), 2873–98.
- Jolivet, Grégory, “A Longitudinal Analysis of Search Frictions and Matching in the U.S. labor market,” *Labour Economics*, 2009, 16 (2), 121–134.
- , Fabien Postel-Vinay, and Jean-Marc Robin, “The Empirical Content of the Job Search Model: Labor Mobility and Wage Distributions in Europe and the US,” *European Economic Review*, 2006, 50 (4), 877–907.
- Krueger, Alan B. and Andreas I. Mueller, “Job Search, Emotional Well-Being, and Job Finding in a Period of Mass Unemployment: Evidence from High-Frequency Longitudinal Data,” *Brookings Papers on Economic Activity*, Spring 2011, (1), 1–70.
- Mortensen, Dale T., *Wage Dispersion: Why Are Similar Workers Paid Differently?*, Cambridge, Massachusetts: MIT Press: Zeuthen Lecture Book Series, 2003.

— and Christopher Pissarides, “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 1994, *61*, 397–415.

Mueller, Andreas I., “Reservation Wages over the Spell of Unemployment: Evidence from High-Frequency Longitudinal Data,” February 2013. Columbia University.

Postel-Vinay, Fabien and Jean-Marc Robin, “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 2002, *70* (6), 2295–2350.

Rogerson, Richard, Robert Shimer, and Randall Wright, “Search-Theoretic Models of the Labor Market: A Survey,” *Journal of Economic Literature*, December 2005, *43*, 959–988.

```
Bibtex: @UNPUBLISHED{Hall:Mueller,author = {Robert E.Hall and Andreas Mueller},
title = {Wage Dispersion and Search Behavior},note = {Hoover Institution, Stanford
University},
month = {March},year = {2013}}
```