

NBER WORKING PAPER SERIES

FRictional Wage Dispersion in Search Models:
A Quantitative Assessment

Andreas Hornstein
Per Krusell
Giovanni L. Violante

Working Paper 13674
<http://www.nber.org/papers/w13674>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2007

We are grateful for comments from Gadi Barlevy, Bjoern Bruegemann, Zvi Eckstein, Chris Flinn, Rasmus Lentz, Dale Mortensen, Giuseppe Moscarini, Fabien Postel-Vinay, Richard Rogerson, Robert Shimer, Philippe Weil, and from many seminar participants. Greg Kaplan and Matthias Kredler provided excellent research assistance. The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Andreas Hornstein, Per Krusell, and Giovanni L. Violante. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Frictional Wage Dispersion in Search Models: A Quantitative Assessment
Andreas Hornstein, Per Krusell, and Giovanni L. Violante
NBER Working Paper No. 13674
December 2007
JEL No. E24,J24,J31,J6,J64

ABSTRACT

Standard search and matching models of equilibrium unemployment, once properly calibrated, can generate only a small amount of frictional wage dispersion, i.e., wage differentials among ex-ante similar workers induced purely by search frictions. We derive this result for a specific measure of wage dispersion -- the ratio between the average wage and the lowest (reservation) wage paid. We show that in a large class of search and matching models this statistic (the "mean-min ratio") can be obtained in closed form as a function of observable variables (i.e., the interest rate, the value of leisure, and statistics of labor market turnover). Various independent data sources suggest that actual residual wage dispersion (i.e., inequality among observationally similar workers) exceeds the model's prediction by a factor of 20. We discuss three extensions of the model (risk aversion, volatile wages during employment, and on-the-job search) and find that, in their simplest versions, they can improve its performance, but only modestly. We conclude that either frictions account for a tiny fraction of residual wage dispersion, or the standard model needs to be augmented to confront the data. In particular, the last generation of models with on-the-job search appears promising.

Andreas Hornstein
Research Department
Federal Reserve Bank of Richmond
POB 27622
Richmond, VA 23261-7622
Andreas.Hornstein@rich.frb.org

Giovanni L. Violante
Department of Economics
New York University
19 W. 4th Street
New York, NY 10012-1119
and NBER
glv2@nyu.edu

Per Krusell
Department of Economics
Princeton University
111 Fisher Hall
Princeton, NJ 08544
and NBER
pkrusell@princeton.edu

1 Introduction

The economic success of individuals is largely determined by their labor market experience. For centuries, economists have been interested in studying the determinants of earnings dispersion among workers. The standard theories of wage differentials in competitive environments are three. Human capital theory suggests that a set of individual characteristics (e.g., individual ability, education, labor market experience, and job tenure) are related to wages because they correlate with productive skills, either innate or cumulated in schools or on the job. The theory of compensating differentials posits that wage dispersion arises because wages compensate for non-pecuniary characteristics of jobs and occupations such as fringe benefits, amenities, location, and risk. Models of discrimination assume that certain demographic groups are discriminated against by employers and, as such, they earn less for similar skill levels.

Mincerian wage regressions based on cross-sectional individual data proxy all these factors through a large range of observable variables, but typically they can explain at most 1/3 of the total wage variation. A vast amount of residual wage variation is left unexplained. In practice, measurement error is large, and the available covariates capture only imperfectly what the theory suggests as determinants of wage differentials. However, even if we could perfectly measure what these competitive theories require, we should not expect to account for all observed wage dispersion.

Theories of frictional labor markets building on the seminal work of McCall (1970), Mortensen (1970), Lucas and Prescott (1974), Burdett (1978), Pissarides (1985), Mortensen and Pissarides (1994), and Burdett and Mortensen (1998) predict that wages can diverge among ex-ante similar workers looking for jobs in the same labor market (e.g., the market for janitors in Philadelphia) because of informational frictions and luck in the search and matching process. We call this type of wage inequality inherently associated to frictions *frictional wage dispersion*.¹

The canonical search model (risk neutrality, no on-the-job search) provides a natural starting point for thinking about frictional wage dispersion. We begin by asking how much frictional wage dispersion this model can generate. We arrive at a surprising answer. This

¹Mortensen (2005), which reviews the theoretical and empirical investigations on the subject, calls it *pure* wage dispersion.

answer is based on a particular measure of frictional wage dispersion: the ratio between the average wage and the lowest (reservation) wage paid in the economy, which we label the “mean-min (Mm) ratio”. The surprise is that optimal search implies a relation between the Mm ratio and a small set of statistics, none of which includes the wage-offer distribution the searching worker is drawing from. Moreover, since there are natural ways of calibrating these statistics—the job finding and separation rates, the interest rate, and the “value of non-market time”—we obtain an equation that predicts a value for our measure of wage dispersion, without having to specify the wage-offer distribution. Thus, with data on frictional wage dispersion that would allow us to measure the Mm ratio, this equation offers a simple and powerful test of the basic theory.

It also turns out that our key equation will hold in all of the three standard versions of the search model (the sequential search model, the island model, and the random matching model). This reflects the important insight that the test we propose here does not rely on a view on how wages are determined in equilibrium, because it only concerns the cornerstone of search theories, namely rational worker search. Indeed, our findings are relevant no matter where the wage offer distribution comes from, be it wage determination mechanisms based on wage posting, bargaining, or efficiency wages.

Our second main finding is quantitative and quite surprising as well: a calibration of the standard search model predicts $Mm = 1.036$. I.e., the model only generates a 3.6% differential between the average wage and the lowest wage paid. The reason is that in the search model workers remain unemployed if the option value of search is high. The latter, in turn, is determined by the dispersion of wage opportunities. The data on unemployment duration show that workers do not wait for very long. Thus, unless they have implausibly high discounting (exceeding 25% per month), or experience an implausibly strong aversion to leisure (exceeding, in dollar terms, five times the average earnings), the observed search behavior of workers can only rationalize a very small dispersion in the wage distribution.

The natural follow-up question is: how large is frictional wage dispersion in actual labor markets? Ideally, one would like to access individual wage observations for ex-ante similar workers searching in the same labor market. These requirements pose several challenges that we can only partially address, given data availability. We exploit three alternative data sources: the 5% Integrated Public Use Microdata Series (IPUMS) sample

of the 1990 U.S. Census, the November 2000 Occupational Employment Survey (OES), the 1967-1996 waves of the Panel Study of Income Dynamics (PSID). Overall, from the empirical work we gauge that the observed Mm ratios are at least twenty times larger than what the model predicts.

How do we reconcile the large discrepancy between the plausibly calibrated search model and the data? On the one hand, one can accept the canonical search model's prediction that frictional wage dispersion is indeed small. One would then attribute most of the residual wage dispersion as measured in the data to measurement error or unobservable, time-varying, heterogeneity that cannot be fully controlled for, given data constraints. This essentially eliminates frictional wage dispersion as an important source of wage inequality, and it implies that a model entirely without frictions in the labor markets suffices to explain almost all of the wage data. On the other hand, one can accept that frictional dispersion as measured is truly large. In this case the canonical search model fails to explain it, and more elaborate search models need to be explored in order to account for our recorded Mm ratios. Thus, in the second part of the paper, we study a number of more elaborate search models through the lens of our instrument of analysis, the Mm ratio.

The first extension introduces risk aversion: risk-averse workers particularly dislike the low-income state (unemployment) and set a low reservation wage, which allows the model to generate a larger Mm ratio. We find that even when we exclude agents from using any form of insurance, a very high risk-aversion coefficient (around eight) is needed to generate an Mm ratio that is comparable to what we see in the data. Moreover, we know that the availability of self-insurance (say, with precautionary saving) would help consumers a lot, thus requiring much higher risk aversion still.

The second extension allows for stochastic wage fluctuations during employment, with endogenous separations. If wages vary over the employment spell, then the wage offer drawn during unemployment is not very informative about the value of a job, so a large dispersion of wages could coexist with a small dispersion of job values which is what drives unemployment duration. This modification is quantitatively not successful since it generates high Mm ratios only if wages are virtually i.i.d. during the employment relationship. Empirically, however, wages are best described by a random walk.

The third set of extensions allows for on-the-job search. The ability to search on the job for new employment opportunities makes unemployed workers less demanding, which reduces their reservation wage and allows the model to generate a higher Mm ratio. This latter modification, in its simplest form, i.e., the basic job-ladder model of Burdett (1978), still falls quite short of explaining the data.² To generate dispersion, this model needs a high arrival rate of offers on the job. However, the high arrival rate implies separations at a frequency that is almost twice that observed. More sophisticated environments with on-the-job search and endogenous search effort (Christensen et al., 2005), ability for the employer to make counteroffers (Postel-Vinay and Robin, 2002), or to offer wage-tenure contracts (Stevens, 1999; Burdett and Coles, 2000) are the ones that show most promise.

In conclusion, we find that for plausible parameterizations a remarkably large class of search models has trouble generating the observed amount of residual wage dispersion, and we explain why. Our analysis is helpful for understanding why the existing empirical structural search literature (see Eckstein and Van den Berg, 2005, for a recent survey) systematically finds very low (negative and large) estimates of the value of non-market time, extremely high estimates of the interest rate, substantial estimates of unobserved worker heterogeneity, or very large estimates of measurement error. Interestingly, the relevance of the value of non-market time also points to tight link between the present analysis and the recently discussed difficulty of replicating the time series of unemployment and vacancies using search theory (see, e.g., Hall, 2005; Shimer, 2005b; Hagedorn and Manovskii, 2006): parameter values that help resolve the former problem make it harder to resolve the latter, and vice versa.

The rest of the paper is organized as follows. Section 2 derives the common expression for the Mm ratio in three canonical search models and quantifies their implications. Section 3 contains the empirical analysis. Section 4 makes a number of first attempts at rescuing the canonical model. Sections 5, 6, and 7 then outline the three significant extensions of the model mentioned above and evaluate them quantitatively. Section 8 discusses the empirical search literature from our perspective. Section 9 concludes the paper. Finally, several of the theoretical propositions in the present paper are proved in

²Since none of our derivations depend on the shape of the wage offer distribution, our results hold in the wage-posting equilibrium versions of the job-ladder model (e.g., Burdett and Mortensen, 1998) as well.

a separate Technical Appendix: Hornstein, Krusell, and Violante (2007).

2 Frictional wage dispersion in canonical models of equilibrium unemployment

The three canonical models of frictional labor markets are the sequential search model developed by McCall (1970) and Mortensen (1970), the island model of Lucas and Prescott (1974), and the random matching model proposed by Pissarides (1985).

In what follows, we show that all three models lead to the same analytical expression for a particular measure of frictional wage dispersion: the mean-min ratio, i.e., the ratio between the average wage and the lowest wage paid in the labor market to an employed worker. Then, we explore the quantitative implications of this class of models for this particular statistic of frictional wage dispersion.

2.1 The basic search model

We begin with the basic sequential search model formulated in continuous time. Consider an economy populated by ex-ante equal, risk-neutral, infinitely lived individuals who discount the future at rate r . Unemployed agents receive job offers at the instantaneous rate λ_u . Conditionally on receiving an offer, the wage is drawn from a well-behaved distribution function $F(w)$ with upper support w^{\max} . Draws are i.i.d. over time and across agents. If a job offer w is accepted, the worker is paid a wage w until the job is exogenously destroyed. Separations occur at rate σ . While unemployed, the worker receives a utility flow b which includes unemployment benefits and a value of leisure and home production, net of search costs. Thus, we have the Bellman equations

$$rW(w) = w - \sigma [W(w) - U] \tag{1}$$

$$rU = b + \lambda_u \int_{w^*}^{w^{\max}} [W(w) - U] dF(w), \tag{2}$$

where $rW(w)$ is the flow (per period) value of employment at wage w , and rU is the flow value of unemployment. In writing the latter, we have used the fact that the optimal search behavior of the worker is a reservation-wage strategy: the unemployed worker accepts all wage offers w above $w^* = rU$, at a capital gain $W(w) - U$. Solving equation (1) for $W(w)$ and substituting in (2) yields the reservation-wage equation

$$w^* = b + \frac{\lambda_u}{r + \sigma} \int_{w^*}^{w^{\max}} [w - w^*] dF(w).$$

Without loss of generality, let $b = \rho \bar{w}$, where $\bar{w} = E[w|w \geq w^*]$. Then,

$$\begin{aligned} w^* &= \rho \bar{w} + \frac{\lambda_u [1 - F(w^*)]}{r + \sigma} \int_{w^*}^{w^{\max}} [w - w^*] \frac{dF(w)}{1 - F(w^*)} \\ &= \rho \bar{w} + \frac{\lambda_u^*}{r + \sigma} [\bar{w} - w^*], \end{aligned} \quad (3)$$

where $\lambda_u^* \equiv \lambda_u [1 - F(w^*)]$ is the job-finding rate. Equation (3) relates the lowest wage paid (the reservation wage) to the average wage paid in the economy through a small set of model parameters.

If we now define the mean-min wage ratio as $Mm \equiv \bar{w}/w^*$ and rearrange terms in (3), we arrive at

$$Mm = \frac{\frac{\lambda_u^*}{r + \sigma} + 1}{\frac{\lambda_u^*}{r + \sigma} + \rho}. \quad (4)$$

The mean-min ratio Mm is our new measure of frictional wage dispersion, i.e., wage differentials entirely determined by luck in the random meeting process. This measure has one important property: it does not depend *directly* on the shape of the wage distribution F . Put differently, the theory allows predictions on Mm without requiring any information on F . The reason is that all that is relevant to know about F , i.e., its probability mass below w^* , is already contained in the job finding rate λ_u^* , which we can measure directly through labor market flows from unemployment to employment and treat as a parameter.

The model's mean-min ratio can thus be written as a function of a four-parameter vector, $(r, \sigma, \rho, \lambda_u^*)$, which we can try to measure independently. Thus, looking at this relation, if we measure the discount rate r to be high (high impatience), for given estimates of σ , ρ , and λ_u^* , an increased Mm must follow. Similarly, a higher measure of the separation rate σ increases Mm (because it reduces job durations and thus decreases the value of waiting for a better job opportunity). A lower estimate of the value of non-market time ρ would also increase Mm (agents are then induced to accept worse matches). Finally, a lower measure of the contact rate λ_u^* pushes Mm up, too (because it makes the option value of search less attractive).

2.2 The search island model

We outline a simple version of the island model as in Rogerson, Shimer, and Wright (2006). Consider an economy with a continuum of islands. Each island is indexed by its productivity level p , distributed as $F(p)$. On each island there is a large number of firms operating a linear technology in labor $y = pn$, where n is the number of workers employed. In every period, there is a perfectly competitive spot market for labor on every island. An employed worker is subject to exogenous separations at rate σ . Upon separation, she enters the unemployment pool. Unemployed workers search for employment and at rate λ_u they run into an island drawn randomly from $F(p)$.

It is immediate that one can obtain exactly the same set of equations (1)-(2) for the worker, while for the firm in each island, we can write its flow value of producing as

$$rJ(p) = pn - wn.$$

Competition among firms drives profits to zero, and thus in equilibrium $w = p$. At this point the mapping between the island model and the search model is complete.³ The search island model yields the same expression for Mm as in (4).

2.3 The basic random matching model

There are three key differences between the search setup described above and the matching model (e.g., Pissarides, 1985). First, there is free entry of vacant firms (or jobs). Second, the flow of contacts m between vacant jobs and unemployed workers is governed by an aggregate matching technology $m(u, v)$. Let the workers' contact rate be $\lambda_u = m/u$ and the firm's contact rate be $\lambda_f = m/v$. Third, workers and firms are ex-ante equal, but upon meeting they jointly draw a value p , distributed according to $F(p)$ with upper support p^{\max} , which determines flow output of their potential match. Once p is realized, they bargain over the match surplus in a Nash fashion and determine the wage $w(p)$. Let β be the bargaining power of the worker; then the Nash rule for the wage establishes that

$$w(p) = \beta p + (1 - \beta) rU, \tag{5}$$

³One can also allow firms to operate a constant returns to scale technology in capital and labor, i.e., $y = pk^\alpha n^{1-\alpha}$. If capital is perfectly mobile across islands at the exogenous interest rate r , then firms' optimal choice of capital allows us to rewrite the production technology in a linear fashion, and the equivalence across the two models goes through.

where rU is the flow value of unemployment.⁴ This equation uses the free-entry condition of firms that drives the value of a vacant job to zero.

From the worker's point of view, it is easy to see that equations (1)-(2) hold with a slight modification:

$$\begin{aligned} rW(p) &= w(p) - \sigma [W(p) - U] \\ rU &= b + \lambda_u \int_{p^*}^{p^{\max}} [W(p) - U] dF(p), \end{aligned}$$

i.e., the value of employment is expressed in terms of the value p of the match drawn; similarly, the optimal search strategy is expressed in terms of a reservation productivity p^* . Rearranging these two expressions, we arrive at an equation for the reservation productivity

$$p^* = b + \frac{\lambda_u \beta}{r + \sigma} \int_{p^*}^{p^{\max}} [p - p^*] dF(p), \quad (6)$$

where we have used the fact that $p^* = rU$. Substituting (5) into (6), we obtain

$$\begin{aligned} w^* &= b + \frac{\lambda_u [1 - F(p^*)]}{r + \sigma} \int_{p^*}^{p^{\max}} [w(p) - w^*] \frac{dF(p)}{1 - F(p^*)} \\ &= b + \frac{\lambda_u^*}{r + \sigma} [\bar{w} - w^*]. \end{aligned}$$

Using the definition $b = \rho \bar{w}$ in the last equation, we again obtain the formula in (4) for the mean-min ratio. Note that nothing in this derivation depends on the shape of the matching function.

One can view the island model or the matching model as providing alternative explanations of the wage offer distribution. The analyses of these models then simply underscore the fact that our key finding—the relation between the Mm ratio and a small set of parameters of the model—merely is an implication of optimal search behavior. In particular, this behavior is independent of the particular equilibrium model used to generate the wage offer distribution F . For example, suppose alternatively that efficiency-wage theory lies behind F : identical workers may be offered different wages because different employers have different assessments of what wage is most appropriate in their firm. From the worker's perspective, then, the end result is still a wage distribution F from which they must sample. Thus, the equation based on the Mm ratio appears to be of remarkably general use in order to understand frictional wage dispersion.

⁴See, for example, Pissarides (2000), Section 1.4, for a step-by-step derivation of this wage equation.

2.4 Quantitative implications for the mean-min ratio

How much frictional wage dispersion can these models generate when plausibly calibrated? We set the period to one month.⁵ An interest rate of 5% per year implies $r = 0.0041$. Shimer (2005a) reports, for the period 1967-2004, an average monthly separation rate σ (EU flow) of 2% and a monthly job finding probability (UE flow) of 39%. These two numbers imply a mean unemployment duration of 2.56 months, and an average unemployment rate of 4.88%.

The OECD (2004) reports that the net replacement rate of a single unemployed worker in the U.S. in 2002 was 56%. The fraction of labor force eligible to collect unemployment insurance is close to 90% (Blank and Card, 1991) which implies a mean replacement rate of roughly 50%. Of course, unemployment benefits are only one component of b . Others are the value of leisure, the value of home production (both positive), and the search costs (negative). Shimer (2005b), weighting all these factors, sets ρ to 41%. As discussed by Hagedorn and Manovskii (2006), this is likely to be a lower bound. For example, taxes increase the value of ρ since leisure and home production activities are not taxed. Given that higher values for b will strengthen our argument, we proceed conservatively and set $\rho = 0.4$. In Section 4 we perform a sensitivity analysis.

This choice of parameters implies $Mm = 1.036$: the model can only generate a 3.6% differential between the average wage and the lowest wage paid in the labor market. This number appears very small. What explains the inability of the search/matching model to generate quantitatively significant frictional wage dispersion? We answer in two ways. First, just mechanically, note that

$$Mm = \frac{\frac{\lambda_u^*}{r+\sigma} + 1}{\frac{\lambda_u^*}{r+\sigma} + \rho} = \frac{\frac{0.39}{0.0241} + 1}{\frac{0.39}{0.0241} + 0.4} = \frac{16.18 + 1}{16.18 + 0.4} = 1.036. \quad (7)$$

This “unpleasant search arithmetic” illustrates that what accounts quantitatively for our finding is that the job finding rate λ_u^* is an order of magnitude larger than the separation rate σ ; hence, the term $\lambda_u^*/(\sigma + r)$ dominates both 1 and ρ , the other two terms of the expression in (7).

⁵The Mm ratio has the desirable property of being invariant to the length of the time interval. A change in the length of the period affects the numerator and denominator of the ratio $\lambda_u^*/(r + \sigma)$ proportionately, leaving the ratio unchanged. The parameter ρ is unaffected by the period length.

The second, more intuitive, interpretation of our finding is that in the search model, workers remain unemployed if the option value of search is high. The latter, in turn, is determined by the dispersion of wage opportunities. The short unemployment durations, as in the U.S. data, thus *reveal* that agents do not find it worthwhile to wait because frictional wage inequality is tiny. The message of search theory is that “good things come to those who wait,” so if the wait is short, it must be that good things are not likely to happen.

We now turn to the obvious next question: how large is frictional wage dispersion in actual labor markets?

3 An attempt to measure frictional wage dispersion

The aim of our analysis is to quantify the empirical counterpart of the model’s mean-min ratio Mm . Ideally, one would like to access individual wage observations for ex-ante similar workers searching in the same labor market. This requirement poses three major challenges that we address by exploiting three alternative data sources: the 5% IPUMS sample of the 1990 U.S. Census, the November 2000 survey of the OES, and the 1967-1996 waves of the PSID.

The first challenge is to define a “labor market”. The most natural boundaries across labor markets are, arguably, geographical, sectoral, and occupational, and possibly combinations of all these dimensions. The PSID sample is too small for a construction of detailed labor markets. OES and Census data, however, allow us to look at the wage distribution in thousands of separate labor markets in the U.S. economy.

Second, differences in annual earnings may reflect differences in hours worked. To avoid this problem, one should focus on hourly wages. However, it is well known that measurement error in hours worked plagues household surveys, and large measurement error will generate an upward bias in estimates of wage dispersion. The OES is an establishment survey where measurement error should be negligible. PSID and Census information on hourly wages, though, may suffer from measurement error bias. In particular, estimates of the reservation wage through the lowest wage observation are especially vulnerable to reporting or imputing errors. For this reason, we estimate the reservation wage from the 1st, 5th, and 10th percentile of the wage distribution. These percentiles

are less volatile, though upward biased, estimators of the reservation wage in the empirical wage distribution. We denote the corresponding mean-min ratios by Mp_1 , Mp_5 , and Mp_{10} , respectively.

Third, we would like to eliminate all wage variation due to ex-ante differences across individuals in observable characteristics (e.g., experience, education, race, gender, etc.), as well as all wage variation due to heterogeneity in unobservables (e.g., innate ability, value of leisure, etc.). The OES does not provide any demographic information on workers. In the Census data, we can control for a wide set of observable characteristics. In the PSID, we can exploit the panel dimension to also purge fixed individual heterogeneity from the wage data.⁶

Overall, none of the three data sets is ideal, but each of them is informative in its own way. At the end of this section, we also quickly review the evidence based on linked employer-employee data.

3.1 Census Data

Our first data source is the 5% Integrated Public Use Microdata Series (IPUMS) sample of the 1990 United States Census. In Appendix A, we outline our sample selection. We run an OLS regression on the log of individual hourly wage to control for the variation in wages due to observable characteristics which are rewarded in the labor market. Among the covariates, we include gender dummies, 3 race dummies, 5 education dummies, and a cubic in potential experience. We weight each observation by its Census sample weight. The regression explains 31% of the total variation in log hourly wages.

Next, we group the (exponent of the) regression residuals by labor markets. Our first definition of labor market is the individual occupation. The Census allows us to distinguish between 487 distinct occupations (variable OCC). Our second definition is the combination of occupation and place of work (for the main job). Our indicator for place of work is constructed as follows. Whenever possible, we use the 329 metropolitan areas (PWMETRO). For rural areas we use a variable (PWPUMA) which defines a geographical area by following the boundaries of groups of counties, or census-defined “places” which

⁶Note also that, say, the Mp_5 estimator of the Mm ratio is robust to the existence of a small group of workers in the population with particularly low labor productivities or values of leisure that would artificially drive up many other measures of wage inequality; see Section 4.5.

contain up to 200,000 residents. Overall, we end up with 799 different geographical areas. For each cell identified as a labor market, we calculate the mean-min ratio and we report the median mean-min ratio across U.S. labor markets.

In the top panel of Figure 1, we show one example of the wage distribution for Janitors and Cleaners, excluding Maids and House Cleaners (code 453), in the Philadelphia metropolitan area (code 616). These are the wage residuals of a regression that controls for the demographics listed above, restricted to those working full time (35-45 hours per week) and full year (48-52 weeks per year) in order to reduce the impact of measurement error. Overall, we have 572 observations. As reported in the figure, the ratio between the mean and the first percentile is 2.24. In the bottom panel of Figure 1 we extend the analysis to the entire U.S. economy. We find that there are local labor markets displaying more and markets displaying less residual dispersion than in Philadelphia, but the bottom line seems to be that, even within a very unskilled occupation such as janitors, and even after selecting the sample to minimize the role of measurement error, wage differentials remain large: the median Mp_1 ratio for janitors across the U.S. is 2.20.

Table 1 reports our results in a number of formats. We use various estimators of the mean-min ratio: Mm , Mp_1 , Mp_5 , and Mp_{10} . We condition on cells with at least 50 or 200 observations. Larger cells usually display higher Mm ratios both because of higher unobserved heterogeneity that we did not capture in the first-stage regression, and because they permit a less biased estimate of the low percentiles of the wage distribution.

To get at the measurement error issue, we condition our analysis on full-time, full-year workers who report weekly hours between 35 and 45 and annual weeks worked between 48 and 52. Going from row (2) to row (3) in Table 1, wage dispersion falls with respect to the full sample, but it remains very high.⁷

To eliminate the importance of individual-specific differences in cumulated skills not perfectly correlated with experience (accounted for in the first-stage regression), we condition on workers with less than 10 years of experience, Table 1 row (4), and on a set

⁷The coefficient of variation falls by 16%. For comparison, Bound and Krueger (Table 6, 1991) compare matched Current Population Survey data to administrative Social Security payroll tax records and find that the measurement error explains between 7% and 19% of the total standard deviation of log earnings. Recall that the standard deviation of the logs has the same scale as the coefficient of variation.

Table 1: Dispersion measures for hourly wage from the 1990 Census
(Median across labor markets)

Labor mkt definition	Min. obs. per cell	Number of labor mkts	Ratio of mean wage to				CV
			min.	1st pct.	5th pct.	10th pct.	
(1) Occupation		487	4.54	2.83	2.13	1.83	0.47
(2) Occ./Geog. Area	(N \geq 50)	13,246	2.94	2.66	2.04	1.76	0.41
	(N \geq 200)	2,321	3.85	2.88	2.13	1.82	0.44
(3) Occ./Geog. Area Full time/Full year	(N \geq 50)	7,195	2.74	2.49	1.92	1.66	0.35
	(N \geq 200)	1,117	3.58	2.68	1.98	1.71	0.37
(4) Occ./Geog. Area Experience \leq 10	(N \geq 50)	2,810	2.64	2.46	1.92	1.68	0.40
	(N \geq 200)	406	3.33	2.57	1.97	1.73	0.44
(5) Occ./Geog. Area Unskilled Occ.	(N \geq 50)	1,152	2.51	2.37	1.98	1.77	0.45
	(N \geq 200)	191	2.95	2.57	2.08	1.83	0.49
(6) Occ./Geog. Area Within cell regression	(N \geq 50)	13,246	2.61	2.39	1.88	1.64	0.39
	(N \geq 200)	2,321	3.33	2.66	2.02	1.75	0.42

of very low-skilled occupations, Table 1 row (5), where occupation and firm-specific skills are arguably not very important.⁸ Once again, going from row (2) to either row (4) or row (5) barely affects the findings, that is, residual wage dispersion remains very high.⁹ Moreover, in models of on-the-job search where employers make counteroffers in case

⁸This list of occupations includes, inter alia, launderers and ironers; crossing guards; waiters and waitresses; food counter, fountain and related occupations; janitors and cleaners; elevator operators; pest control occupations; and baggage porters and bellhops.

⁹This conclusion is consistent with the most up-to-date estimates of the returns to firm tenure. Recently, Altonji and Williams (2004) have reassessed the existing evidence (e.g., Topel, 1991; Altonji and Shakotko, 1987), concluding that returns to tenure over 10 years are around 11%, most of which occur in the first 5 to 7 years of the employment spell. If we assume that average tenure is roughly 3.5 years (see Section 7), this factor would account for less than 10% of the wage difference between the average worker (with average tenure) and the lowest paid worker (with zero tenure) in a given occupation/geographical area.

Hagedorn and Manovskii (2005) argue that the bulk of returns to specific human capital is occupation-specific. They find that once occupation is taken into account, returns to human capital specific to industry and employer become virtually zero. Hence, their findings would change the source of wage differentials due to unobserved heterogeneity in specific skills, but not their magnitude.

an employee is contacted by an outside firm, wage growth on the job is observationally equivalent to the accumulation of job-specific skills.¹⁰ Thus at least part of the returns to tenure may be attributable to “luck” and should not be removed from our empirical measure of frictional inequality.

Finally, we also run the first-stage regressions within each occupation/area cell to account for the fact that the role of demographic characteristics in wage determination may be different across occupations. Going from row (2) to row (6), the estimates of dispersion fall by less than 10%.

We conclude that, except for estimates of Mm based on the lowest observed wage that are quite volatile, all the other statistics in Table 1 remain very robust to all these controls and strongly support the view that residual wage dispersion is large. In the top panel of Figure 2, we report the empirical distribution of the Mp_5 across U.S. labor markets corresponding to the sample selection with $N \geq 50$ in row (2) of Table 1.

3.2 OES Data

The second data source we use is the November 2000 survey from the Occupational Employment Statistic (OES). Appendix B contains a description of the survey and of the sample selection criteria we adopt. OES data are collected at the level of the establishment. Each establishment reports the average hourly wage paid within each occupation. To the extent that there are within-establishment differences in wages due to luck or frictions among similar workers, these data underestimate frictional wage dispersion.

We use three different levels of aggregation: nation-wide data by occupation (3-digit), occupation \times metropolitan area data, and occupation (3-digit) \times industry (2-digit) data. The publicly available survey reports the average, the 10th, 25th, 50th, 75th, and 90th percentiles of the hourly wage distribution.

The best possible estimate (which clearly is downward-biased) of the mean-min ratio is the ratio between the average wage and the 10th percentile (Mp_{10}). We exclude all those cells where hourly wage data are not available, and those where the 90th percentile is top coded (at \$70 per hour)—a sign that wages are heavily censored in that cell.¹¹

¹⁰See Section 7.3 for details.

¹¹The first restriction mainly excludes workers in the education sector (25-000), while the second mainly excludes healthcare practitioners (29-000).

Table 2: Dispersion measures from the 2000 OES
(Median across labor markets)

Labor mkt definition	Number of labor mkts	Ratio of mean wage to 10th percentile
Occupation	637	1.68
Occ./Industry	6,293	1.60
Occ./Geog Area	106,278	1.48

In each cell (i.e., a labor market) we compute the Mp_{10} , and then we calculate the median Mp_{10} ratio across labor markets. The median is preferable to the mean because we consistently found that the empirical distributions of mean-min ratios are very skewed to the right, and we are interested in the mean-min ratio of the wage distribution in a typical labor market.

The median Mp_{10} ratio across occupations in the U.S. economy is 1.68. For the classification of labor markets based on 2-digit industry (58 industries) and occupation, the median Mp_{10} ratio is 1.60. Finally, when we define labor markets by metropolitan area (337 areas) and occupation, the median Mp_{10} ratio is estimated to be 1.48. Table 2 summarizes the results.

As the definition of labor market becomes more refined, wage dispersion falls for two reasons. First, there is less worker heterogeneity within a specific occupation in a given industry, or in a given geographical area than at the country level. Second, as we keep disaggregating, the number of establishments sampled within each cell falls. For example, for cells defined by occupation and metropolitan area, we have on average only 11 establishments per cell. With such a low number of observations, the estimate of the 10th percentile could be severely upward biased, and in turn the mean-min ratio underestimated.

3.3 PSID Data

Our third data source is the Panel Study of Income Dynamics (PSID). In Appendix C, we describe our sample selection in detail. With our final sample in hand, for every year in the period 1967-1996, we run an OLS regression on the cross-section of individual log hourly

wages where we control for gender, 3 race dummies (white, non-white, and Hispanic), 5 education dummies (high-school dropout, high-school degree, some college, college degree, and post-graduate degree), a cubic in potential experience (age minus years of education minus five), a dummy for marital status, 6 regional dummies, 25 two-digit occupation dummies, and an interaction between occupation and experience to capture occupation-specific tenure profiles. We face a trade-off in the choice of covariates between (1) the appropriate filtering out of the variation in hourly wages due to observable individual characteristics which are rewarded in the labor market, and (2) the risk of overfitting the data. On average, these year-by-year regressions yield an R^2 between 0.42 and 0.45.¹²

Next, we use the panel dimension of the data and identify individual-specific effects in wages. Let ε_{it} be the residual of the first stage for individual $i = 1, \dots, I$ in year $t = 1, \dots, T$. We limit the sample to those whose number of wage observations in the panel (N_i) is at least ten and estimate $\bar{\varepsilon}_i = \sum_{t=1}^{N_i} \varepsilon_{it}/N_i$ for every individual. The vector $\{\bar{\varepsilon}_i\}_{i=1}^I$ captures the variation in fixed unobserved individual factors (e.g., innate ability or preference for leisure) which affect wages. Let $\tilde{w}_{it} = \exp(\varepsilon_{it} - \bar{\varepsilon}_i)$. For each year t , we then calculate our indexes of residual inequality across workers on \tilde{w}_{it} .¹³

We report our results for the PSID in Figure 3. For comparison with the other data sources, we comment on the values for the last part of the sample (the 1990-1996). The ratio between the mean wage and the lowest wage residual from the basic Mincer regressions is $Mm = 4.47$, but the estimate is clearly very noisy. When the reservation wage is estimated from the 1st and the 5th percentile of the wage distribution, the noise is much reduced and we obtain, respectively, $Mp_1 = 2.73$ and $Mp_5 = 2.08$. The coefficient of variation of the regression residuals is 0.50.

Controlling for individual effects drastically cuts the estimate of the mean-min ratios by more than half. For the period 1990-1996, we estimate $Mm = 3.11$, $Mp_1 = 1.90$, and $Mp_5 = 1.46$. The coefficient of variation of the residuals net of individual effects falls

¹²This R^2 is sizeably higher than the one obtained for the Census regression. The reason is that in the PSID regression we include occupational dummies and occupation-specific experience profiles, which have strong explanatory power. Moreover, the PSID sample is smaller by a factor of 1,500 compared to the Census sample. See Appendices A and C.

¹³In an unreported set of regressions on PSID data, we also allowed fixed differences in (linear and quadratic) time trends across workers—possibly capturing heterogeneity in learning ability—and this did not significantly change our findings either.

to 0.25.¹⁴ One should be cautious in interpreting these results, however. The estimated fixed effects confound worker-specific characteristics with match- (or firm-) specific effects. This is especially true for long-lived matches. Removing the estimated fixed effects may therefore eliminate some of the variation in the data that we want to explain.

To facilitate the comparison with the OES data, we also report ratios between the mean and the 10th percentile. The Mp_{10} for the residuals of the first-stage regression equals 1.77, and the Mp_{10} for the residuals net of individual-specific means is 1.32. The corresponding statistics for the OES data all lie somewhere in between (see Table 2).

An alternative way to control for fixed individual heterogeneity is to compute mean-min ratios within each individual wage history, i.e., on the I samples $\{\exp(\varepsilon_{it})\}_{t=1}^{T_i}$, where T_i denotes the maximum number of wage observations available for individual i in PSID. Next one can plot the distribution of individual mean-min ratios across the population—we do it in the bottom panel of Figure 2—and calculate the median of this distribution. We obtain $Mm = 1.57$. The short length of the individual samples, which are limited by the time coverage of PSID (30 years) implies that the estimator of the individual reservation wage is greatly upward biased. Nevertheless, we still estimate a large amount of residual wage dispersion.¹⁵

3.4 Evidence from employer-employee matched data

Data sets linking firm-level records and individual worker demographics can help effectively control for worker heterogeneity and isolate the variation in the quality of jobs in a given labor market. Abowd and Kramarz (2000) survey the empirical work based on employer-employee linked databases and conclude that, for the U.S., firm effects represent around one third of the total wage variation and, perhaps surprisingly, the estimated correlation between firm and worker effects is close to zero. This evidence is reassuring for two reasons. First, in our PSID exercise, individual effects capture around three quarters of the total variance. Second, if indeed firm and worker effects are orthogonal, a consistent estimate of the firm effect can be obtained by averaging all wages (say, for a certain

¹⁴In passing, we note that Figure 3 is consistent with the views that residual wage dispersion has risen significantly over the period, and that the rise in prices for unobserved innate characteristics is a key component of this phenomenon.

¹⁵We report Mm because, given the short length of individual samples, p_5 is identical to the minimum in almost every individual history, and thus $Mp_5 = Mm$.

occupation) within a firm.

In this context our analysis of wage dispersion in the OES survey data (Table 2) provides estimates of the typical (median) variation in firm-level effects within a labor market defined by occupation and geographical area. Mortensen (2007) performs a similar analysis for a Danish set of matched employer-employee data (IDA). He reports Mp_5 ratios of firm-level effects that vary between 1.32 and 1.65 depending on the occupation. These estimates are slightly below our estimates for the U.S., which is not surprising given that the regulations in Danish labor markets likely imply more wage compression.

3.5 Summary and interpretation

Our three independent data sources offer a fairly consistent view of the size of residual wage dispersion within narrowly defined labor markets. If we focus our attention on the Mp_5 estimate of the mean-min ratio, a review of our findings yields $Mp_5 = 1.46$ from PSID. The PSID estimate could be upward-biased because of measurement error, but at the same time the individual wage demeaning could filter out too much variation, including variation due to “persistent luck” components that should be included in measures of frictional dispersion. From the Census sample restricted to full-time, full-year workers (where measurement error in hours should be negligible) we have estimated $Mp_5 = 1.98$. Given the OES estimate of the Mp_{10} , and the fact that the other two data sets suggest that Mp_5 are roughly 10%-15% larger, we conjecture that the Mp_5 in the OES data may be around 1.67. An average across the three data sets yields 1.70, which we use as a target in the rest of our analysis.

This appraisal of residual wage dispersion—based not on the minimum wage observed, but on the 5th percentile, and hence quite conservative—is about *20 times larger* than what implied by the textbook models of Section 2.

4 Several attempts to rescue the baseline model

4.1 Unemployment vs. wage dispersion

In defense of the baseline search model, one might argue that it is designed to explain unemployment, not wage dispersion. This argument is flawed: in the environments of

Section 2, there is a tight link between the existence of unemployment and the existence of wage dispersion. Unemployment is in large part due to the option value of searching for better wage opportunities. Let us reverse our logic and suppose that, given the amount of frictional wage dispersion observed empirically, we want to predict unemployment duration, i.e., use equation (4) and the empirical value of $Mm = 1.7$ to compute the implied value for λ_u^* . We would obtain $\lambda_u^* = 0.011$. In other words, a search model consistent with the amount of wage dispersion in the data predicts an expected unemployment duration of 91 months, 35 times the average duration in U.S. data.¹⁶

4.2 Targeting European data

In Section 2.4 we have indicated that the short duration of unemployment in the U.S. reveals that frictional wage dispersion must be small. It is well known that in Europe unemployment spells last much longer, on average. For example, Machin and Manning (1999, Table 1) document that in 1995 the proportion of workers unemployed for more than 12 months was less than 10% in the U.S., but over 40% in France, Germany, Greece, Italy, Portugal, Spain, and the United Kingdom. Does this observation mean that the model could be successful in explaining European wage dispersion? To answer this question, recall that in a stationary equilibrium, unemployment is $u = \sigma / (\sigma + \lambda_u^*)$. Using this formula in expression (4) allows one to rewrite the Mm ratio as

$$Mm = \frac{\frac{\sigma}{r+\sigma} \frac{1-u}{u} + 1}{\frac{\sigma}{r+\sigma} \frac{1-u}{u} + \rho} \approx \frac{\frac{1-u}{u} + 1}{\frac{1-u}{u} + \rho},$$

where the “approximately equal” sign is obtained by setting $r = 0$, a step justified by the fact that r is of second order compared to the other parameters in that expression. Setting the unemployment rate to 10%, with $\rho = 0.4$, one obtains $Mm = 1.076$. The reason for the small improvement is that in European data both unemployment duration and employment spells are much longer than in the U.S. labor market. While the first fact is consistent with larger equilibrium wage dispersion, the second implies that unemployed

¹⁶In one of the most commonly used search setups, Mortensen and Pissarides (1994), unemployment duration is not connected to wage dispersion, because in that model, unemployed workers always receive the maximum wage upon employment: they never consider turning down a job offer. There, however, since there are wage shocks on the job, the separation rate is determined by a reservation-wage strategy. Thus, that model instead links wage dispersion to the observed separation rate. We explore this link in Section 6 below.

workers are more selective and set their reservation wage high, which reduces frictional wage dispersion.

For the argument to be fully convincing, one would need to document the extent of residual wage dispersion, as well as the magnitude of the value of non market time, in European countries. A systematic investigation goes beyond the scope of this paper, but conventional wisdom would suggest that while inequality is lower in Europe, social benefits for the unemployed are much more generous.¹⁷

4.3 Alternative parameterizations

To calibrate the pair (λ_u^*, σ) , we used the UE and EU flow data. One could argue that we should also incorporate flows in and out of the labor force. Taking this into account, Shimer (2005a) reports the monthly separation rate to be 3.5% and the monthly job finding rate to be 61%. For the same values of r and ρ used in the baseline calibration, we obtain $Mm = 1.038$.¹⁸

With respect to the interest rate r , we have used a standard value, but it is possible that unemployed workers, especially the long-term unemployed, face a higher effective interest rate if they wanted to borrow. Much less is known about ρ . To assess the robustness of our conclusions to the choice of values for these two parameters, in Figure 4, we plot the pairs (r, ρ) which are consistent with an Mm ratio of 1.7, together with the region of “plausible pairs” based on our prior. This region covers the area where $\rho \in (0, 1)$ and r is at most 27% per year.

The results are striking and suggest the baseline model cannot be rescued. Positive net values of non-market time are consistent with the observed wage dispersion only for interest rates beyond 1,350% per year.¹⁹ Even for annual interest rates around 40% per year, one would need agents to value one month of time away from the market the

¹⁷For example, Hansen (1998) calculates benefits replacement ratios with respect to the average wage up to 75% in some European countries.

¹⁸Since, under the new parameterization, both the job finding rate and the separation rate increase, the effect on the Mm ratio is negligible.

¹⁹A number of authors in the health and social behavioral sciences have argued that unemployment can lead to stress-related illnesses due, e.g., to financial insecurity or to a loss of self-esteem. This psychological cost would imply an additional negative component in b . Economists have argued that this empirical literature has not convincingly solved the serious endogeneity problem underlying the relationship between employment and health status and even have, at times, reached the opposite conclusion, i.e., that there is a positive association between time spent in non-market activity and health status (e.g., Ruhm 2003).

equivalent of *minus* three times the average monthly wage. To be concrete, consider that, for a worker earning \$500 per week, this means the following: in order to avoid unemployment, she would be willing to *work for free* for a week, *pay* \$1,500, and at the end of the week draw a job offer from the *same* distribution she would face if she had remained unemployed. This appears economically implausible.

4.4 Endogenous search effort

The typical calibration of the value of non-market time as a fraction of the mean wage, $\rho = 0.4$ (e.g., Shimer, 2005b), is mainly based on average UI replacement rates. Therefore, it does not take into account search costs directly. Here we explore the possibility that large search costs improve the model's performance.

Consider an extension of the baseline model where search effort is endogenous. Let $c_u(\lambda_u)$, with $c'_u > 0$ and $c''_u > 0$, be the effort cost as a function of the offer arrival rate λ_u , the endogenous variable chosen by the unemployed worker. The flow value of employment remains as in equation (1), while the flow value of unemployment in equation (2) now contains the extra term $-c_u(\lambda_u)$ on the right-hand side. The same derivations as in Section 2 lead to the reservation-wage equation

$$w^* = b - c_u(\lambda_u^o) + \frac{\lambda_u^*}{r + \sigma} (\bar{w} - w^*), \quad (8)$$

where λ_u^o denotes the optimal individual choice, and $\lambda_u^* \equiv \lambda_u^o [1 - F(w^*)]$. The FOC for optimal search effort is

$$c'_u(\lambda_u^o) = \frac{1}{r + \delta} \int_{w^*}^{w^{\max}} (w - w^*) dF(z). \quad (9)$$

We follow Christensen et al. (2005) and choose the isoelastic functional form $c_u(\lambda_u) = \kappa_u \lambda_u^{1+1/\gamma}$ for the effort cost, with $\gamma > 0$ denoting the elasticity of the optimal search effort with respect to the expected return from search. Using the relationship between marginal and average search cost which follows from this specification, we arrive at the net return from search relative to the average wage

$$\frac{b - c_u(\lambda_u^o)}{\bar{w}} = \rho - \frac{\lambda_u^*}{r + \sigma} \frac{\gamma}{1 + \gamma} \left(1 - \frac{1}{Mm} \right). \quad (10)$$

Combining (8) and (10), and rearranging, we obtain

$$Mm = \frac{\frac{\lambda_u^*}{r+\sigma} \frac{1}{1+\gamma} + 1}{\frac{\lambda_u^*}{r+\sigma} \frac{1}{1+\gamma} + \rho},$$

which highlights the role of search costs. If $\gamma = 0$, then optimal effort (and the associated search cost) is zero and the model collapses to the baseline. The larger is γ , the less sensitive are search costs to the chosen offer arrival rate, and optimal search effort λ_u^o rises. As a result, the search cost increases, lowering the net-of-search-cost value of non-market time $b - c_u(\lambda_u^o)$ and raising the Mm ratio.

The new parameter needed to quantitatively evaluate this extension is the elasticity γ . For a quadratic search cost, $\gamma = 1$, the wage dispersion, $Mm = 1.071$, remains small relative to its data counterpart.²⁰ To match the observed wage dispersion, $Mm = 1.7$, one would need $\gamma = 34.5$. From equation (10), it then follows that the search cost of spending one *month* unemployed is almost equal to average *annual* earnings. Thus the net-of-search-cost value of non-market time relative to average wages, which is -6.1 , does not change relative to the model without endogenous search; see Figure 4.

4.5 Ability differences

Wage inequality can, of course, naturally arise from ability differences. A very simple illustration, extending the above search setting, goes as follows: there are two worker types, and type 1 is more productive than type 2 by μ percent in the following sense: $F_2(w) \equiv F(w)$ and $F_1(w) = F(w/(1 + \mu))$ for all w . Suppose also that the workers have the same values for ρ , r , σ , and λ_u . Then

$$w_i^* = \rho \bar{w}_i + \frac{\lambda_u [1 - F_i(w_i^*)]}{r + \sigma} (\bar{w}_i - w_i^*) \quad (11)$$

for each type i . It is easy to show that if w_2^* and $\bar{w}_2 = \int_{w_2^*} w F(dw) / [1 - F(w_2^*)]$ solve (11) for $i = 2$, then, using the assumed symmetry, $w_1^* = (1 + \mu)w_2^*$ implies $F_1(w_1^*) = F_2(w_2^*)$ and $\bar{w}_1 = \int_{w_1^*} w F_1(dw) / [1 - F(w_1^*)] = (1 + \mu)\bar{w}_2$, and therefore solves (11) for $i = 1$. Thus, in this model the observed wage distribution for the type-1 worker is a μ -percent

²⁰Christensen et al. (2005) estimate $\gamma = 1.19$ for Danish data. We are not aware of estimates for γ from U.S. data.

scaling up of that of type-2 workers, and $\bar{w}_1/w_1^* = \bar{w}_2/w_2^*$, which will be a small number given the above analysis.

However, for the population, $\bar{w} = \alpha\bar{w}_1 + (1 - \alpha)\bar{w}_2$, where α is the share in population of type 1. So the population-wide mean-min ratio, which the econometrician observes, will equal

$$Mm = \frac{\alpha\bar{w}_1 + (1 - \alpha)\bar{w}_2}{w_2^*} = \left(\alpha \frac{\bar{w}_1}{\bar{w}_2} + 1 - \alpha \right) \frac{\bar{w}_2}{w_2^*} = (1 + \alpha\mu) \frac{\bar{w}_2}{w_2^*}.$$

Thus, if μ is large and α is not too small, we can obtain large population mean-min values, even though mean-min values within groups are small. In particular, large enough ability difference will generate any desired mean-min ratio for the overall population. The model just described, however, is *not* one of frictional wage inequality, but rather one of ability-driven wage differences.

A related model could also be constructed assuming that the types are random; perhaps they follow a Markov chain, depicting the evolution of human capital on (and also off) the job. Similar settings have been used by Ljungqvist and Sargent (1998) and Kamboorov and Manovskii (2004). One can use arguments along the lines of those above to demonstrate that such settings also allow large overall wage dispersion, but only because of the skill differences between types being large: within a given type, frictional wage inequality—as we view it and define it—is still tiny.

4.6 Relation between Mm and other dispersion measures

Although the mean-min ratio is not a commonly used index of dispersion, it has desirable features. The axiomatic approach to “ideal” inequality indexes is discussed in depth by Cowell (2000), who lists five standard axioms: anonymity, the population principle, scale invariance, the principle of transfers, and decomposability.²¹ The Mm ratio always satisfies the first three axioms and weak versions of the last two. In addition, it should

²¹Anonymity requires that the inequality measure be independent of any characteristic of individuals other than income. According to the population principle, measured inequality should be invariant to replications of the population. Scale invariance requires the inequality measure to be constant when each individual’s income changes by the same proportion. According to the principle of transfers, an income transfer from a poorer person to a richer person should imply a rise (in its strict version) or at least no change (in its weak version) in inequality. Decomposability requires that if the same distribution, say F , is mixed with G and G' , then ordering of the resulting mixture distribution is determined solely by the ordering of G and G' . See Cowell (2000) for details.

be pointed out that our statistical index of dispersion has the same properties as quantile ratios (e.g., the 90th – 10th percentile ratio), a class of indexes that is commonly used in the empirical inequality literature.

Could it be that even though the model fares poorly in terms of our Mm ratio statistic, its implications for more common measures of dispersion, such as the coefficient of variation (cv), is satisfactory? To answer this question, we need to make further assumptions about the equilibrium wage distribution. Given a parametric specification for this distribution, we can map predicted mean-min ratios into cv 's, i.e., we can determine the value of the cv corresponding to a certain value for the mean-min ratio.

The Gamma distribution (see Mood et al., 1974, for a standard reference) is a convenient choice because it is a flexible parametric family and has certain properties that are useful in our application. Let wages w be distributed according to the density

$$g(w; w^*, \alpha, \gamma) = \frac{\left(\frac{w-w^*}{\alpha}\right)^{\gamma-1} \exp\left(-\frac{w-w^*}{\alpha}\right)}{\alpha\Gamma(\gamma)}, \quad (12)$$

with $\alpha, \gamma > 0$, and with $\Gamma(\gamma)$ denoting the Gamma function. The value for w^* is the location parameter and determines the lowest wage observation, α is the scale parameter determining how spread out the density is on its domain, whereas γ is the parameter that determines the shape of the function (e.g., exponentially declining, bell-shaped, etc.).²²

The mean and standard deviation of a random variable distributed with $g(w; w^*, \alpha, \gamma)$ are given by, respectively, $\bar{m} = w^* + \alpha\gamma$ and $sd = \alpha\sqrt{\gamma}$. Recalling that $Mm = \bar{m}/w^*$, it is easy to obtain a relation between the coefficient of variation cv and the mean-min ratio, and this relation only depends on the shape parameter γ :

$$cv = \frac{1}{\sqrt{\gamma}} \left[\frac{Mm - 1}{Mm} \right]. \quad (13)$$

The empirical analysis of section (3) suggest that $cv = 0.30$ and $Mm = 1.7$ are reasonably conservative estimates for the coefficient of variation and the mean-min ratio of the wage distribution within labor markets.²³ From equation (13), this implies $\gamma = 1.88$.

A search model generating a mean-min ratio of 1.036, under the Gamma wage distribution assumption, would generate a coefficient of variation for hourly wages of 0.025,

²²The Gamma family is very flexible: it includes the Weibull (and, hence, the exponential) distribution for $\gamma = 1$ and, as $\gamma \rightarrow \infty$, the lognormal distribution.

²³This value for cv is an average between the PSID estimate (0.25) and the estimate on Census sample of full-time, full-year workers (0.35).

i.e., 1/12th of the coefficient of variation in the wage data. We conclude that the failure of the model generalizes to more common measures of dispersion.

4.7 Non-pecuniary job attributes

In many jobs, wages are only one component of total compensation. In a search model where a job offer is a bundle of a monetary component and a non-pecuniary component, short unemployment duration can coexist with large wage dispersion, as long as non-pecuniary job attributes are *negatively correlated* with wages so that the dispersion of total job values is indeed small.

This hypothesis, which combines the theory of compensating differentials with search theory, does not show too much promise. First, it is well known that certain key non-monetary benefits such as health insurance tend to be positively correlated with the wage, e.g., through firm size.²⁴ Second, illness or injury risks are very occupation-specific and our measures of frictional wage dispersion apply within-occupation indexes. Third, differences in work shifts and part-time penalties are quantitatively small. Kostiuk (1990) shows that genuine compensating differentials between day and night shifts can explain at most 9% of wage gaps. Manning and Petrongolo (2005) calculate that part time penalties for observationally similar workers are around 3%.

4.8 Taking stock

How can we resolve this striking discrepancy between the size of measured residual wage dispersion and the model-implied frictional wage dispersion? We discuss three possible reactions to our finding.

One reaction is that the actual wage data hide large differentials due to unobservable skills that we cannot fully control for, given data constraints. This is possible, though one has to bear in mind that despite controlling for unobserved heterogeneity that is fixed over time, measured residual inequality remains large. Thus, for worker heterogeneity to explain large Mm ratios, one would need to invoke time-varying, unobserved skills, or preferences, which influence remuneration in the labor market. Such heterogeneity

²⁴For example, the mean wage in jobs offering health insurance coverage is 15%-20% higher than in those not offering it; see Dey and Flinn (2006).

cannot be ruled out a priori; when better data sources become available, they may shed more light on its quantitative importance. If it is indeed found that our measurement of residual wage dispersion is accurate, it would conversely also imply that the bulk of the observed differences in individual wage outcomes could be accounted for by a model with time-varying human capital priced within a frictionless labor market.²⁵

A different reaction is to argue that there is nothing wrong with the baseline model but that our estimates of ρ , perhaps due to psychological and other hard-to-measure costs of unemployment, are not accurate: frictional wage dispersion is large, and with a large enough disutility of being unemployed, the model can account for it. Such a conclusion, of course, would also have far-fetching implications for the rest of macroeconomics. First, we would have to “add unemployment as an argument of utility functions”, thus radically altering available analyses of aggregate labor supply. Second, a large value of ρ would damage the matching model as a model of unemployment fluctuations. In particular, Hall (2005), Shimer (2005b), and Hagedorn and Manovskii (2006) point to the inability of the matching model to generate enough time-series fluctuations in aggregate unemployment and vacancies. They show that, even with the incorporation of real-wage rigidity, the model requires a high ρ in order to produce realistic movements in vacancy and unemployment rates; without real-wage rigidity, the value needs to be very close to one. Thus, in addition to the arguments we put forth above, available analyses of aggregate data also suggest that large negative values of ρ are implausible.

A third reaction is to conclude that the data reveal a sizeable amount of frictional wage dispersion, and that the baseline model, plausibly calibrated, cannot account for it. Then, alternative specifications of search theory would need to be explored. We do so here by looking at the quantitative implications of three elements that are missing in the baseline search model above: risk aversion, stochastic wages during the employment relationship, and on-the-job search. In particular, the analysis of on-the-job-search allows us to inspect some of the most recent, and promising, contributions to the literature.

²⁵Incidentally, a large class of quantitative macroeconomic models of the Bewley-Huggett-Aiyagari style implicitly takes this view, presuming idiosyncratic risks which are modeled as a stochastic process for efficiency units of labor, priced in a frictionless labor market. An example of this approach in the search literature is Ljungqvist and Sargent (1998).

5 Risk aversion

Risk-averse workers particularly dislike states with low consumption, like unemployment. Compared to risk-neutral workers, *ceteris paribus*, they lower their reservation wage in order to exit unemployment rapidly, thus allowing Mm to increase.

Let $u(c)$ be the utility of consumption, with $u' > 0$, and $u'' < 0$. To make progress analytically, we assume that workers have no access to storage, i.e., $c = w$ when employed, and $c = b$ when unemployed. It is clear that this model will give an *upper bound* for the role of risk aversion: with any access to storage, self-insurance or borrowing, agents can better smooth consumption, thus becoming effectively less risk-averse.

To obtain the reservation-wage equation with risk aversion, observe that in the Bellman equations for the value of employment and unemployment, the monetary flow values of work and leisure are simply replaced by their corresponding utility values. The reservation-wage equation (3) then becomes

$$u(w^*) = u(\rho\bar{w}) + \frac{\lambda_u^*}{r + \sigma} [E_{w^*} [u(w)] - u(w^*)], \quad (14)$$

with $E_{w^*} [u(w)] = E[u(w) | w \geq w^*]$. A second-order Taylor expansion of $u(w)$ around the conditional mean \bar{w} yields

$$u(w) \simeq u(\bar{w}) + u'(\bar{w})(w - \bar{w}) + \frac{1}{2}u''(\bar{w})(w - \bar{w})^2.$$

Take the conditional expectation of both sides of the above equation and arrive at

$$E_{w^*} [u(w)] \simeq u(\bar{w}) + \frac{1}{2}u''(\bar{w}) \text{var}(w). \quad (15)$$

Let $u(w)$ belong to the CRRA family, with θ representing the coefficient of relative risk aversion. Then, using (15) in (14), and rearranging, we obtain

$$Mm \equiv \frac{\bar{w}}{w^*} \simeq \left[\frac{\frac{\lambda_u^*}{r+\sigma} \left(1 + \frac{1}{2}(\theta - 1)\theta cv^2(w)\right) + \rho^{1-\theta}}{\frac{\lambda_u^*}{r+\sigma} + 1} \right]^{\frac{1}{\theta-1}}. \quad (16)$$

It is immediate to see that, for $\theta = 0$, the risk-neutrality case, the expression above equals that in equation (4).

To assess the quantitative role of risk aversion, we start with the parameterization of the risk-neutral case, and based on the evidence provided in section 3, we set $cv(w) = 0.30$.

Figure 5 plots the pairs of (ρ, θ) consistent with $Mm = 1.70$. For $\theta = 8.3$, the model can match the data.²⁶

Recall, however, the upper-bound nature of our experiment: in fact, plausibly calibrated models of risk-averse individuals who have access to a risk-free bond for saving and borrowing are much closer to full insurance than to autarky (see, e.g., Aiyagari, 1994). For example, it is well known that as $r \rightarrow 0$, the bond economy converges to complete markets (Levine and Zame, 2002). In the Technical Appendix we study a search environment where preferences display constant absolute risk aversion and workers can borrow and save through a riskless asset. We show that relative risk aversion around 100 is needed in order to match the observed mean-min ratio with positive replacement rates.

6 Wage shocks during employment

We now extend the basic search model by allowing wages to fluctuate stochastically along the employment spell.²⁷ Unemployed workers draw wage offers from the distribution $F(w)$ at rate λ_u , but now these wage offers are not permanent. At rate δ , the wage changes, and the worker draws again from $F(w)$. Draws are i.i.d. over time. Separations are now endogenous and will occur at rate $\sigma^* \equiv \delta F(w^*)$, where w^* is the reservation wage.

The reason why this generalization can potentially generate a larger Mm ratio is that the particular value drawn from $F(w)$ by an unemployed worker is not a good predictor of the continuation value of employment, if the wage is very volatile. Unemployed workers will therefore be more willing to accept initially low wage offers, which reduces w^* and increases dispersion.

The Bellman equations for employment and unemployment are, respectively,

$$\begin{aligned} rW(w) &= w + \delta \int_{w^*}^{w^{\max}} [W(z) - W(w)] dF(z) - \delta F(w^*) [W(w) - U] \\ rU &= b + \lambda_u \int_{w^*}^{w^{\max}} [W(z) - U] dF(z). \end{aligned}$$

²⁶It is easy to derive third- and fourth-order approximations of the reservation-wage equation involving the coefficients of skewness and kurtosis. At the values for these coefficients estimated in the data, our conclusions are extremely robust. For example, with the third- and fourth-order approximation one needs, respectively, $\theta = 8.4$ and $\theta = 7.8$ to replicate the observed Mm ratio.

²⁷The Technical Appendix contains the details of all the derivations in this and the next sections.

In comparison with equation (1), the value of employment is modified in two ways. First, the endogenous separation rate is now $\delta F(w^*)$. Second, there is a surplus value from accepting a job at wage w which is given by the second term on the right-hand side. Exploiting the definition of reservation wage $W(w^*) = U$, integrating by parts, and using $W'(w) = 1/(r + \delta)$, we arrive at

$$w^* = b + \frac{\lambda_u - \delta}{r + \delta} \int_{w^*}^{w^{\max}} [1 - F(z)] dz.$$

Now, using the definition of the conditional mean wage \bar{w} , we obtain

$$w^* = b + \frac{(\lambda_u - \delta) [1 - F(w^*)]}{r + \delta} (\bar{w} - w^*).$$

Therefore, imposing $b = \rho \bar{w}$, and rearranging, we can write the Mm ratio in this model as

$$Mm = \frac{\bar{w}}{w^*} = \frac{\frac{\lambda_u^* - \delta + \sigma^*}{r + \delta} + 1}{\frac{\lambda_u^* - \delta + \sigma^*}{r + \delta} + \rho}. \quad (17)$$

As $\delta \rightarrow 0$, the Mm ratio converges to equation (4) with $\sigma^* = 0$, since without any shock during employment every job lasts forever. As $\delta \rightarrow \lambda_u$, unemployed workers accept every offer above b since being on the job has an option value equal to being unemployed.

The parameter δ maps into the degree of persistence of the wage process during employment. In particular, in a discrete time model where $\delta \in (0, 1)$ the autocorrelation coefficient of the wage process is $1 - \delta$.²⁸ Individual panel data suggest that residual wages are very persistent, indeed near a random walk, so plausible values of δ are close to zero.

We repeat the exercise in Section 4 on the (r, ρ) pair. Given values for $(\lambda_u^*, \sigma^*, r)$ one can search for the values of (δ, ρ) that generate Mm ratios of the observed magnitude. Figure 6 reports the results. Once again, the model is very far from the data for plausible values of ρ and of the degree of wage persistence of wage shocks. For example, for $\delta = 0.1$ (corresponding to an annual autocorrelation coefficient of 0.9) the model requires $\rho = -13$. Only for virtually i.i.d. wage shocks ($\delta \approx 1$) would the model succeed.

The setup of Mortensen and Pissarides (1994) is similar to that described here, with one difference: upon employment, all workers start with the highest wage, w^{\max} , and

²⁸It is easily seen that a discrete time version of this model leads exactly to equation (17).

thus they only sample from $F(w)$ while employed. This model, however, does not offer higher mean-min ratios: in the Technical Appendix we show that the resulting Mm ratio is strictly bounded above by that in equation (17).

7 On-the-job search

Allowing on-the-job search goes, qualitatively, in the right direction for reasons similar to those applying in the model with stochastic wages. If the arrival rate of offers on the job is high, workers are willing to leave unemployment quickly since they do not entirely forego the option value of search. This property breaks the link between duration of unemployment and wage dispersion that dooms the baseline model. However, for a proper test, we now need to explore the implied labor-market flows, which unlike before include employment-to-employment transitions.

We generalize the model of Section 2 and turn it into the canonical job ladder model outlined by Burdett (1978). A worker employed with wage \hat{w} encounters new job opportunities w at rate λ_e . These opportunities are drawn from the wage offer distribution $F(w)$ and they are accepted if $w > \hat{w}$.

A large class of equilibrium wage posting models, starting from the seminal work by Burdett and Mortensen (1998), derives the optimal wage policy of the firms and the implied equilibrium wage offer distribution as a function of structural parameters. It is not necessary, at any point in our derivations, to specify what $F(w)$ looks like. Our expression for Mm will hold in any equilibrium wage posting model that satisfies the following two assumptions. First, employed workers accept any wage offer above their current wage; second, every worker (employed or unemployed) faces the same wage offer distribution. Moreover, without loss of generality, to simplify the algebra, we posit that no firm would offer a wage below the reservation wage w^* ; thus, $F(w^*) = 0$.

The flow values of employment and unemployment are:

$$\begin{aligned} rW(w) &= w + \lambda_e \int_w^{w^{\max}} [W(z) - W(w)] dF(z) - \sigma [W(w) - U] \\ rU &= b + \lambda_u \int_{w^*}^{w^{\max}} [W(z) - U] dF(z), \end{aligned}$$

and the reservation-wage equation becomes

$$w^* = b + (\lambda_u - \lambda_e) \int_{w^*}^{w^{\max}} \frac{1 - F(z)}{r + \sigma + \lambda_e [1 - F(z)]} dz. \quad (18)$$

It is easy to see that, in steady state, the cross-sectional wage distribution among employed workers is

$$G(w) = \frac{\sigma F(w)}{\sigma + \lambda_e [1 - F(w)]}. \quad (19)$$

Using this relation between $G(w)$ and $F(w)$ in the reservation-wage equation (18), and exploiting the fact that the average wage is

$$\bar{w} = w^* + \int_{w^*}^{w^{\max}} [1 - G(z)] dz, \quad (20)$$

we arrive at the new expression for the Mm ratio,

$$Mm \approx \frac{\frac{\lambda_u - \lambda_e}{r + \sigma + \lambda_e} + 1}{\frac{\lambda_u - \lambda_e}{r + \sigma + \lambda_e} + \rho}, \quad (21)$$

in the model with on-the-job search.²⁹

Note that, if the search technology is the same in both employment states and $\lambda_e = \lambda_u$, the reservation wage will be equal to b , since searching when unemployed gives no advantage in terms of arrival rate of new job offers. Indeed, for $\lambda_e > \lambda_u$, unemployed workers optimally accept jobs below the flow value of non-market time b in order to access the better search technology available during employment.

The crucial new parameter of this model is the arrival rate of offers on the job λ_e . To pin down λ_e , note that average job tenure in the model is given by

$$\tau = \int_{w^*}^{w^{\max}} \frac{dG(w)}{\sigma + \lambda_e [1 - F(w)]} = \frac{\sigma + \lambda_e/2}{\sigma(\sigma + \lambda_e)} \in \left(\frac{1}{2\sigma}, \frac{1}{\sigma} \right). \quad (22)$$

Since we set the monthly separation rate σ to 0.02, the model can only generate average tenures between 25 and 50 months. Based on the CPS Tenure Supplement, the BLS

²⁹The ‘‘approximately equal’’ sign originates from one step of the derivation where we have set

$$\frac{r + \sigma}{r + \sigma + \lambda_w [1 - F(z)]} \approx \frac{\sigma}{\sigma + \lambda_w [1 - F(z)]},$$

a valid approximation since, for plausible calibrations, r is negligible compared to σ .

(2006, Table 1) reports that *median* job tenure (with current employer) for workers 16 years old and over, from 1983-2004, was 3.64 years, or 43.7 months.³⁰

An alternative way to restrict the choice of λ_e is to compute the average separation rate χ implied by the model, which is given by

$$\chi = \sigma + \lambda_e \int_{w^*} [1 - F(w)] dG(w) = \frac{\sigma(\lambda_e + \sigma) \log\left(\frac{\sigma + \lambda_e}{\sigma}\right)}{\lambda_e}, \quad (23)$$

where the first equality states that the separation rate equals the EU flow rate (σ) plus the EE flow rate (the integral). Recent evidence from CPS data post 1994 (Fallick and Fleischman, 2004; Nagypal, 2005) puts the data counterpart of χ to be just below 4% per month.³¹ It should be noted here that Job Openings and Labor Turnover Survey (JOLTS) data for the period 2000-2007 indicate that the total separation rate, which includes workers who exit the labor force, is less than 4%. Ongoing research is addressing the sources of downward bias in the JOLTS estimate.

In what follows, we ask whether the on-the-job search model can generate the observed Mm ratio while, at the same time, being consistent with tenure lengths higher than 43.7 months or, alternatively, a monthly average separation rate of 4%. Since the Mm ratio is increasing in λ_e , it should be clear that the model will have a good chance to generate a large Mm ratio for low job tenures/high separation rates (which imply high values of λ_e).

Figure 7 illustrates that, once λ_e is chosen so that the model can generate transitions as in the U.S. data, the model will produce high Mm ratios only for negative values of ρ . Although the model is still far from a success, it is clear that the introduction of on-the-job search represents a significant quantitative improvement. Compare Figure 4 with Figure 7. In the former graph, one needs $\rho = -5$ to reproduce $Mm = 1.70$, whereas in the latter setting $\rho = -0.5$ allows the model to roughly match both the separation and the wage dispersion facts.³² However, the independent evidence on average job tenure seems to put more strain on the model: wage dispersion and average tenure are jointly replicated only for values of ρ around -4 .

³⁰For two reasons, this number is a lower bound as the empirical counterpart of τ . First, the implied average job tenure would be longer, since tenure distributions are notoriously skewed to the right. Second, the BLS data refer to the truncated tenure distribution. A higher estimate for average tenure strengthens our argument.

³¹There is no inconsistency between average tenure of 43.7 months and an average separation rate of 4%, since the employment hazard is not constant.

³²For $\rho = 0.4$, the model calibrated to match the aggregate separation rate gives rise to $Mm = 1.10$.

This result contains an important lesson. One may be tempted to think that, given the relationship between $F(w)$ and $G(w)$ in equation (19), it suffices to assume enough dispersion in firms' productivities, and hence in the wage offer distribution $F(w)$, to generate large wage differentials among employed workers. Equation (18), instead, tells us that a high variation in $F(w)$ is consistent with a low reservation wage, and thus a high Mm ratio, only for negative values of b , or high values of r .

Arguably, models with on-the-job search today represent the most vibrant research area within search theory. Next, we analyze three of the most recent theoretical developments of this class of models and evaluate their potential to generate frictional wage dispersion.

7.1 Reallocation shocks

Micro data indicate that a non-negligible fraction of job-to-job movers receive a wage cut. In line with this observation, some recent contributions (e.g., Jolivet et al., 2006) advocate that on-the-job search models, to be successful, must introduce “reallocation shocks”, i.e., a situation where employed workers receive a job offer with an associated wage drawn from $F(w)$ that cannot be rejected. In other words, under the scenario of a reallocation shock the outside option of a worker is not keeping the current job (and the current wage), but becoming unemployed which is always dominated by accepting any new job offered. This category of employment to employment (EE) transitions may include, for example, search activity during a notice period, or a geographical move for non-pecuniary motives.

If we let ϕ be the arrival rate of such an event, it is easy to see that the mean-min ratio in the model becomes

$$Mm \approx \frac{\frac{\lambda_u - \lambda_e - \phi}{r + \sigma + \lambda_e + \phi} + 1}{\frac{\lambda_u - \lambda_e - \phi}{r + \sigma + \lambda_e + \phi} + \rho}, \quad (24)$$

a magnitude that is increasing in ϕ .³³ Ceteris paribus, the reallocation shock shortens the life of a job and makes unemployed workers less picky, and thus they reduce their reservation wage, which helps raising frictional wage dispersion. Moreover, it is straightforward to derive that the average separation rate in this model is exactly as in (23) with σ replaced by $\sigma + \phi$.

³³The approximately equal sign holds here for the same reason as in equation (21).

Even though this extension goes in the right direction, quantitatively it falls short of succeeding because, as in the standard on-the-job search model, the model with forced job-to-job mobility cannot reconcile the observed wage dispersion with labor turnover data. To replicate $Mm = 1.70$ with positive values of ρ , the separation rate would have to be implausibly high.

A variant of the reallocation shock story which shows more promise is suggested by Nagypal (2006). In her model, employed workers move within the wage distribution in two ways: through job-to-job movements following on-the-job search with contact rate λ_e , and through shocks occurring at rate ϕ that change the wage during an employment relationship without inducing a separation: at rate ϕ , workers make another wage draw from $F(w)$ that is uncorrelated to their current wage and that has to be accepted, or else separation takes place. Effectively, this model is a combination of the canonical on-the-job search model and the model with stochastic wages we presented in Section 6.

The mean-min ratio in this model has exactly the same form as in (24), but now raising ϕ does translate into a wage change without increasing the separation rate. To gauge the quantitative implications of this version of the on-the-job search model, we perform the following exercise: keeping the values for (σ, ρ) unchanged relative to our previous analysis, we set the pair (λ_e, ϕ) to match the aggregate separation rate (4%) and $Mm = 1.70$. Then, we calculate the average arrival rate ζ of a wage cut among employed workers in the model. It is given by the formula

$$\zeta = \phi \int F(w) dG(w) = \phi \left[1 + \frac{\sigma + \phi}{\lambda_e} - \frac{(\sigma + \phi)(\lambda_e + \sigma + \phi) \log\left(\frac{\sigma + \phi + \lambda_e}{\sigma + \phi}\right)}{\lambda_e^2} \right].$$

The model implies that 9.5% of the workforce is subject to a wage cut every month, i.e., 70% of the stock of employees experience at least one salary cut from month to month, in any given year.³⁴ This extreme implication is similar to our finding of Section 6, where we concluded that the model with wage shocks (but without on-the-job search) generates individual wage profiles that are too volatile.

³⁴The calibration yields a value for the Poisson arrival rate ζ of 10% which implies a monthly probability of receiving a wage cut of $1 - \exp(-\zeta) = 9.5\%$.

7.2 On-the-job search with endogenous search effort

Recently, Christensen et al. (2005, CLMNW hereafter) have extended Burdett's on-the-job search model allowing for the optimal choice of the wage offer arrival rates, both off and on the job.³⁵

Let $c_i(\lambda_i) = \kappa_i \lambda_i^{1+1/\gamma}$ be the search effort cost, as a function of the chosen arrival rate, in employment state $i \in \{u, e\}$. During unemployment, the optimal effort choice is a scalar λ_u^o . During employment, the search cost is independent of current earnings w , whereas the return to search is declining in these earnings. Thus, the optimal effort choice, $\lambda_e^o(w)$, is decreasing in w . Therefore the contact rate at the reservation wage is higher than the average contact rate in the economy, and the latter, in turn, is higher than $\lambda_e^o(w^{\max})$ where w^{\max} is the upper bound of the earnings distribution. Indeed, optimality requires $\lambda_e^o(w^{\max}) = 0$.

In the Technical Appendix we use these inequalities to derive bounds on the Mm ratio implied by the model.³⁶ In particular, if we let $\lambda_u \equiv \lambda_u^o$ and $\lambda_{w^*} \equiv \lambda_e^o(w^*)$ for notational simplicity, we have

$$\frac{\frac{\lambda_u - \lambda_{w^*}}{(1+\gamma)(r+\sigma)} + 1}{\frac{\lambda_u - \lambda_{w^*}}{(1+\gamma)(r+\sigma)} + \rho} \leq Mm \leq \frac{\frac{\lambda_u - \lambda_{w^*}}{(1+\gamma)(r+\sigma + \lambda_{w^*})} + 1}{\frac{\lambda_u - \lambda_{w^*}}{(1+\gamma)(r+\sigma + \lambda_{w^*})} + \rho}. \quad (25)$$

CLMNW estimate this model for Danish data under the assumption that the search cost functions are the same in the two employment states. In this case $\lambda_u = \lambda_{w^*}$, and $Mm = 1/\rho$. For $\rho = 0.4$ this yields substantial wage dispersion: $Mm = 2.5$. The assumption of identical search cost functions, however, appears to be inconsistent with other aspects of the Danish labor market. In particular, CLMNW obtain a tight estimate for the job finding rate of workers with the reservation wage, $\lambda_{w^*} = 0.07$, that is substantially lower than the job finding rate for the unemployed, $\lambda_u = 0.11$, reported by Rosholm and Svarer (1999).

Suppose that in the U.S. labor market λ_{w^*} is also roughly 64% of λ_u . Then, given $\lambda_u = 0.39$, a plausible estimate for the U.S. would set λ_{w^*} at 0.25. Using this value for λ_{w^*} in (25), together with the values for the other parameters already discussed in the

³⁵In Section 4.4 we studied the version of this model without on-the-job search.

³⁶We thank Dale Mortensen for providing us with these derivations.

paper, we obtain a lower bound of 1.20 and an upper bound of 1.95 for Mm . Therefore, it appears that the model can match the data.

Given that CLMNW estimate their search cost function, i.e., select it in order to match the wage data, one might want to independently assess the implied magnitude of the costs actually incurred by searching workers; recall from Section 4.4 that search costs, which make unemployment unattractive relative to employment, can be thought of as a low ρ . The bounds on the Mm ratio, together with the first-order condition for optimal search effort during unemployment, allow us to construct bounds for the search cost as a fraction of the average wage. In the Technical Appendix, we show that

$$\frac{\lambda_u}{r + \sigma + \lambda_{w^*}} \frac{\gamma}{1 + \gamma} \left(1 - \frac{1}{Mm}\right) \leq \frac{c_u(\lambda_u)}{\bar{w}} \leq \frac{\lambda_u}{r + \sigma} \frac{\gamma}{1 + \gamma} \left(1 - \frac{1}{Mm}\right). \quad (26)$$

Assuming that the model does match an Mm ratio of 1.7, the inequalities in (26) yield a lower bound of 0.32 and an upper bound of 3.62 for the normalized search cost when unemployed. With $\rho = 0.4$, the net-of-search-cost value of non-market time necessary to produce the observed frictional dispersion is then close to zero or, most likely, negative. This calculation echoes one of the central observations of our paper: to generate sizeable wage dispersion the search model needs a very large disutility of non-market time during unemployment.

7.3 Counteroffers and wage-tenure contracts

We end this section by analyzing two very recent developments in the search literature. By focusing on the Mm ratio, we show that both models hold much promise for the task of generating sizeable frictional wage dispersion.

In Postel-Vinay and Robin (2002), firms are allowed to make counteroffers when one of their employees is contacted by an outside firm. The two employers then engage in a Bertrand bidding game which may result either in a separation or in a salary increase on the current job, depending on the productivity gap between the two competing firms. It is immediate that, through this mechanism, the model has a much weaker link between wage dispersion and job-to-job flows, which is the central problem of the standard on-the-job search environment.

Consider a simple version of the counteroffer model where all firms have equal productivity p . The wage determination mechanism implies that firms extract all the surplus

from the unemployed worker, but an employed worker who is contacted by an outside firm extracts all the surplus from the current employer and receives wage $w = p$. It is easy to see that the reservation-wage equation (derived in the Technical Appendix) is

$$w^* = b - \frac{\lambda_e}{r + \sigma} (p - b), \quad (27)$$

which implies that $w^* < b$. The worker expects to earn a higher wage after she has been hired, and thus she is willing to accept a low entry-wage. As a result, $Mm > 1/\rho$. Note that two features are key for generating wage dispersion here. First, the bargaining position of the worker is extremely volatile: when she faces one firm, her surplus is zero, and when she faces two firms she captures all of the surplus. Second, the firm is able to commit to pay $w = p$ forever, following the first contact, even after the outside firm is gone.³⁷

Although the counteroffer model is undoubtedly a good representation for certain high-skill occupations (e.g., academic jobs), it does not appear to be a widespread mechanism for wage setting in the labor market at large.³⁸ An additional limit of this model is that, for high values of λ_e , the entry wage w^* may be negative. Entry fees into jobs are not commonly observed.

Stevens (1999) generalizes the wage-posting model of Burdett and Mortensen (1998) in order to allow employers to offer long-term wage-tenure contracts on a take-it-or-leave-it basis.³⁹ Consider the same simple economy with homogenous firms, and restrict attention to a family of contracts (w^*, w_1, λ) such that the entrant worker is paid w^* (optimally determined by the firm in equilibrium) and then, at rate λ , i.e., after an expected tenure length of $1/\lambda$, her wage jumps up to w_1 . Clearly, an optimal contract that minimizes costly turnover for the firm sets $w_1 = p$.

The long-term wage tenure contract is observationally equivalent to the counteroffer model if $\lambda = \lambda_e$. This equivalence means that in search environments with wage-tenure

³⁷Kiyotaki and Lagos (2006) develop an equilibrium model of two-sided, on-the-job search with no commitment (i.e., with continuous renegotiation) where the worker, even when facing only one employer, always has the chance of making a take-it-or-leave-it offer with a fixed probability.

³⁸One reason is that in many labor markets asymmetric information may prevent the firm from being able to verify the outside offer. See Mortensen (2005) for a discussion of why counteroffers are uncommon in actual labor markets.

³⁹Burdett and Coles (2001) extend it further to allow for risk-averse workers who are unable to save and borrow.

contracts, frictional wage dispersion can be large, even allowing $Mm > 1/\rho$. An advantage of this class of models is that it avoids the problem of a negative entry wage. Contracts can be restricted so that $w^* = \omega \in (0, b)$ and equation (27) determines the expected tenure $1/\lambda$ at which the worker sees his salary increasing to p . A possible drawback of this class of models is the rather heavy reliance on the assumption that firms can commit to long-term contracts; in practice, it appears that renegotiation is frequent.

8 Structural estimation of search models

Since the pioneering effort of Flinn and Heckman (1983), a rather vast literature on structural estimation of search models has developed (see Eckstein and van den Berg, 2005, for a recent survey). We view these contributions as having generated many valuable insights. From our perspective here, it is important to comment on how the literature has dealt with the baseline model's apparent inability to generate frictional wage dispersion. We summarize our reading of the literature as follows: it has either (1) simply "accepted" implausible parameter estimates for the value of non-market time and the interest rate, or (2) introduced unobserved skill heterogeneity and measurement error that soak up the large wage residuals in the data. We now proceed to discussing a number of selected papers in more detail.

In one of the first attempts at a full structural estimation, Eckstein and Wolpin (1990) estimate the Albrecht and Axell (1984) search model with worker heterogeneity in the value of non-market productivity and conclude that their model cannot generate any significant wage dispersion, and that almost all of the observed wage dispersion is explained through measurement error. Eckstein and Wolpin (1995) reach a far better match between model and data, by introducing a five-point distribution of unobserved worker heterogeneity within each race/education group (8 groups in total). In spite of such heterogeneity, however, for many of the groups the estimates of b remain extremely small or negative (see their Table 7, page 284). In this work, thus, wage dispersion is for the most part accounted for by heterogeneity in observable and unobservable characteristics. In our view, this procedure, which is quite frequent in this literature, can perhaps be categorized more as part of the human-capital theory of wages: it delivers wage inequality, but this

inequality is not frictional in nature.⁴⁰

Negative estimates of the net value of non market time are quite common. The survey paper by Bunzel et al. (2001) estimates several models with on-the-job search on Danish data. When firms are assumed to be homogeneous, the point estimate for ρ is -2 . With heterogeneity in firms' productivity it increases to just about 0. Only the model with measurement error produces a large and positive estimate of ρ .⁴¹ Flinn (2006) estimates a Pissarides-style matching model of the labor market, without on-the-job search, to evaluate the impact of the minimum wage on employment and welfare. In his model, as is typical in estimation exercises, the pair (ρ, r) is not separately identified. Setting r to 5% annually in his model implies roughly $\rho = -4$.⁴²

Another example of extreme parameter estimates that can be found in Postel-Vinay and Robin (2002). Under risk neutrality, their estimates of the discount rate r always exceed 30% per year in every occupational group, reaching 55% for unskilled workers, where they find no role for unobserved heterogeneity. Risk aversion (in the form of log utility, without storage) does not help: the estimates of r decline by just 3 percentage points on average. Recall, from Figure 4, that a negative value for ρ and a high value for r are two sides of the same coin.

Whenever authors have restricted (r, ρ) to plausible values ex-ante, not surprisingly given the results in the present paper, they end up finding that frictions play a minor role. For example, Van den Berg and Ridder (1998) estimate the Burdett-Mortensen model on Dutch data allowing for measurement error and observed worker heterogeneity (58 groups defined by education, age and occupation). They set r to zero and b to equal the average unemployment benefit for each group, i.e., roughly 60% of the average wage. They conclude that observed heterogeneity and measurement error account for over 80% of the empirical wage variation. Moscarini (2003) develops an equilibrium search model where workers learn about their match values, based on Jovanovic (1979). When the

⁴⁰A theoretical argument has also been raised against this kind of model of frictional wage dispersion. Gaumont et al. (2006) demonstrate that wage dispersion in an Albrecht and Axell (1984) model with worker heterogeneity in the value of leisure is fragile. As soon as an arbitrarily small search cost is introduced, the equilibrium unravels and we are back to the "Diamond paradox", i.e., to an equilibrium with a unique wage.

⁴¹These values for ρ are obtained from Bunzel et al. (2001) by dividing the estimates of b for the entire sample, in Tables II and V, by the average wage from Table I.

⁴²Calculations are available upon request.

model is calibrated, r is set to 5% annually and ρ to 0.6. His model generates a Mm ratio of just 1.16 (Moscarini, 2005, Table 2).

One may note that a number of papers in the literature do claim that the (on-the-job search) model is successful in simultaneously matching the wage distribution and labor market transition data (see, e.g., Bontempo et al., 2000, and Jolivet et al., 2005). These claims of success, clearly, need to be properly reinterpreted in light of our findings. The typical strategy in these papers is, first, to estimate the wage distribution $G(w)$ non-parametrically without using the search model. Next, the model is used to predict the wage-offer distribution $F(w)$ through a steady-state relationship like (19), where the structural parameters of the relationship ($\sigma, \lambda_u, \lambda_e$) are estimated by matching transition data. Success is then expressed as a good fit (in some specific metric). However, the exercise cannot be considered entirely a success because it neglects the implications of the joint estimates of $F(w)$ and of the transition parameters for the relative value of leisure ρ (or, similarly, for the interest rate r). The key additional “test” that we are advocating would thus entail using the estimated $F(w)$ in the reservation-wage equation (18) and, given an estimate of w^* (for example, the bottom-percentile wage observed), backing out the implied value for ρ . In light of our results, we maintain that ρ would be negative or close to zero.

In conclusion, while we definitely recognize substantial progress in this empirical literature, the success is really only partial: the literature has not yet fully managed to match the data with plausible parameter values. In short, important parameters such as b and r (the value of leisure and the discount rate, respectively), are considered free parameters, and values that are far from what we view as plausible are routinely “accepted” in the estimation. Alternatively, unobserved heterogeneity or measurement error must be introduced, also with amounts that are free parameters, in order to match the data. Our contribution in this context is to provide a new tool of analysis, the mean-min ratio, that makes quantitative analysis more transparent: one can easily calibrate the model and thereby shed light on and explain a range of seemingly puzzling findings in the literature on structural estimation of search models.

9 Conclusions

Search theory maintains that similar workers looking for jobs in the same labor market may end up earning different wages according to their luck in the matching process. This paper has proposed a simple, but widely applicable, strategy for evaluating the quantitative ability of search models to generate frictional wage dispersion. The strategy is based on a particular measure of wage differentials, the mean-min ratio, that arises very naturally, in closed form, from the reservation-wage equation, the cornerstone of a vast class of search models.

We have demonstrated that, when plausibly calibrated, the textbook search and matching model (risk neutrality, no on the job search) implies that frictions play virtually no role in the labor market: the mean-min ratio is less than 4%.

The data we have analyzed tell a striking story: residual wage dispersion among observationally similar workers in narrowly defined labor markets (by occupation and geographical area) is twenty times larger than what predicted by the model. This leaves two possible, but radically different, conclusions on the table.

First, residual wage inequality in the data is attributable to unobserved (and time-varying) skills that are remunerated in a near frictionless labor market; put differently, the role of frictions is actually small in the data as well. This conclusion is viewed as plausible by many, but it does call for more careful testing. In particular, one must look for more individual-specific information (like significant events altering the value of leisure, e.g., need for child or health care in individuals' lives or the value of skills, e.g., occupational human capital becoming obsolete) in order to reduce the current residual dispersion. One may also be able to look more deeply at implications for on-the-job wage variability, which should be quite large if indeed individual-specific and time-varying skill heterogeneity accounts for the wage dispersion. Alternatively, the second possible conclusion which puts more faith in the measures of frictional wage dispersion reported here, is that our basic search theory needs further development. Distinguishing between these two conclusions should be a central task in macroeconomics and labor economics: the relative roles of skills and luck in the labor market lie at the heart of policy design.

The mean-min ratio also proved very useful in assessing a set of extensions to the basic search setting. Risk aversion can be successful only if one believes that self insurance is

unimportant, but a decade of quantitative investigations of Bewley-type models speaks against this possibility. Volatility in wages during employment can be successful only if one believes that wages are as volatile during an employment spell as in the cross-section. The simplest version of the on-the-job-search model also fails to produce large wage dispersion while, at the same time, being consistent with the observed labor market transition data. There is much current work on developing theories of unemployment and wage setting, however. Among them, the very recent theory of “rest unemployment” in Alvarez and Shimer (2007) also seems to generate too little wage dispersion, consistently with our findings for the benchmark theory of worker search. However, the most recent developments of on-the-job-search models, including those with endogenous search effort, matching of outside offers, or firm posting of wage-tenure contracts, seem to have more potential in terms of generating wage dispersion.⁴³

An implicit but interesting implication of the present work is closely connected to the recent debate, fueled by Hall (2005) and Shimer (2005b), on the inability of the matching model to generate enough time-series fluctuations in aggregate unemployment and vacancies. There, it is pointed out that the matching model, at least if one is to avoid the incorporation of significant real-wage rigidity, requires a very high benefit of non-market time (denoted ρ above, expressing this benefit as a *fraction* of the average wage) in order to produce sharp movements in vacancy and unemployment rates. But, as we showed here, the higher is the value of ρ , the more difficult it is to explain wages cross-sectionally. More specifically, the time-series facts necessitate a value of ρ close to one to explain the data, and our cross-sectional facts demand a value significantly below zero. We believe that it is important in future work to keep both “puzzles” in mind while developing and using search-matching theories of the labor market.

⁴³Using Austrian data on unemployment durations and worker asset holdings, Card, Chetty, and Weber (2007) recently argue that models with endogenous search intensity improve significantly on pure reservation-wage models of unemployment. Interestingly, they also find the mean-min wage ratio to be a key statistic in their analysis.

10 Appendix

A: THE 5% IPUMS SAMPLE OF THE 1990 CENSUS

The 1990 Census of the U.S. population uses a single long-form questionnaire for sample questions completed by one half of persons in locations with a population under 2,500, one sixth of persons in other tracts and block numbering areas with fewer than 2,000 housing units, and one eighth of persons in all other areas. Overall, about one sixth of all housing units complete a long form. Within each state, the Bureau divides the sample questionnaires into an appropriate number of 1-percent samples. For example, if 20 percent of the population of a state completed long forms, the sample questionnaires for that state are divided into twenty subsamples of equal size. The 5-percent files are then selected at random from the 1-percent subsamples for each state. Weights are attached to each case representing the number of individuals in the general population represented by any particular case in the sample.

The original data set contains over 12,500,000 person-level observations. To create our sample, first we exclude every person below 20 and above 60 years old, as well as every individual currently in school, self-employed, or disabled, which leaves 4,636,759 individual records in the sample. Next, we exclude all individuals who report zero wage income or zero weeks worked over the year, and individuals whose annual earnings are top-coded, i.e. higher than \$140,000 (19,890 cases). We also remove altogether occupational groups where more than 3% of all workers are top coded. Eleven occupations are excluded based on this criterion, e.g., Airplane pilots, Athletes, Dentists, Physicians, Podiatrists, and Judges. Finally, we eliminate individuals for whom estimated hourly wages are below the federal minimum wage (\$3.35 in 1989), i.e., roughly 400,000 observations.⁴⁴ We are then left with 3,923,744 individual records in our final sample.

We construct hourly wage as annual wage and salary income (variable INCWAGE) divided by the product between the number of weeks worked last year (WKSWRK1) and usual weekly hours worked (UHRSWORK).⁴⁵

B: THE OCCUPATIONAL EMPLOYMENT SURVEY

The Occupational Employment Statistics (OES) program collects data on employees in approximately 200,000 non-farm establishments to produce employment and wage estimates for 821 occupations classified based on the Standard Occupational Classification (SOC).

Since November 2003, the program samples 200,000 establishments semi-annually. Before then, it sampled 400,000 once a year. The OES survey is designed to produce occupational wage and employment estimates using six panels (3 years) of data. The BLS Employment Cost Index is used to adjust survey data from prior panels before combining them with the current panel's data. The full six-panel sample of 1.2 million establishments allows the construction of occupational estimates at detailed levels of geography and industry. Estimates based on geographic areas are available at the National, State, and Metropolitan Area levels. Industry

⁴⁴Through validation with CPS data, Baum-Snow and Neal (2006) find that a significant fraction of workers report usually working 8 hours per week on the census' long form when they actually usually worked 40 hours per week; thus, they respond as if the question meant to report "usual hours per day". However, this type of measurement error which plagues the 1980 Census is much less frequent in the 1990 census. Moreover, these respondents are excluded from our selection criteria on annual hours.

⁴⁵See Ruggles et al. (2000) for a detailed explanation of the IPUMS Census data.

classifications correspond to 3, 4, and 5-digit North American Industry Classification System (NAICS) industry groups.

The OES survey form sent to establishments defines wages as straight-time, gross pay, exclusive of premium pay. Base rate, cost-of-living allowances, guaranteed pay, hazardous-duty pay, incentive pay including commissions and production bonuses, tips, and on-call pay are included. Excluded are back pay, jury duty pay, overtime pay, severance pay, shift differentials, non-production bonuses, employer cost for supplementary benefits, and tuition reimbursements.

The OES survey groups wages in 12 discrete intervals. In November 2004, the lowest interval was “Under \$6.75” and the highest was “\$70 and over”. Mean hourly wage rate for an occupation equals total wages that all workers in the occupation earn in an hour divided by the total employment of the occupation. The same concept applies to more disaggregated levels such as occupations within metropolitan areas or industries. The mean and percentiles for an occupation are calculated by uniformly distributing the workers inside each wage interval, ranking the workers from lowest paid to highest paid.

C: THE PANEL STUDY OF INCOME DYNAMICS

The Panel Study of Income Dynamics (PSID) began collecting information on a sample of approximately 5,000 households in 1968. Of these, about 3,000 are representative of the U.S. population as a whole (core sample), while the rest are low-income families (SEO sample). Since then, both the original households and their split-offs (members of the original household forming a family of their own) have been followed over time. Questions on labor income and hours are retrospective: information collected in the year t wave refers to the calendar year $t - 1$.

Our initial sample comprises every head and spouse between 20-60 years old in the 1968-1997 waves of the PSID core sample. We then exclude individuals currently in school, self-employed, or disabled, and those with annual hours below 520 and above 5096 to reduce the role for measurement error in hours, which leaves 80,979 individual/year records in the sample.⁴⁶ Next, we exclude all individuals whose earnings are top-coded or whose hourly wage (computed as annual earnings divided by annual hours) is below the federal minimum wage, which eliminates around 4,141 individual/years observations. At the end of this selection, we are left with 76,848 individual/year observations in our final sample, i.e., roughly 2,500 per year.

The estimation of the individual fixed effect described in the main text is based on workers with at least ten reported wage observations which reduces the sample size to 49,010 individual/year observations, i.e., 1,633 per year on average.

⁴⁶French (2005) uses the PSID Validation Study to assess the size of measurement error in hourly wages. He concludes that it accounts for 24% of the standard deviation of log wages. By trimming the hours distribution below 520 and above 5096, we eliminate many outliers that are due to reporting errors.

References

- [1] Abowd J., and F. Kramarz (1999). “The Analysis of Labor Markets Using Matched Employer-Employee Data,” Handbook of Labor Economics, O. Ashenfelter and D. Card eds., North-Holland, pages 2629-2710.
- [2] Aiyagari, R. (1994). “Uninsured Idiosyncratic Risk and Aggregate Saving,” Quarterly Journal of Economics, vol. 109(3), pages 659-684.
- [3] Albrecht J., and B. Axell (1984). “An Equilibrium Model of Search Unemployment,” Journal of Political Economy, vol. 92, pages 824-840.
- [4] Altonji, J., and R. Shakotko (1987). “Do Wages Rise With Job Seniority?” Review of Economic Studies, vol. 65(3), pages 437-460.
- [5] Altonji, J., and N. Williams (2004). “Do Wages Rise with Job Seniority? A Reassessment,” forthcoming in the Industrial and Labor Relations Review.
- [6] Baum-Snow, N., and D. Neal (2006). “Data Problems and the Measurement of Racial and Gender Wage Gaps Over Time,” mimeo, University of Chicago.
- [7] Blank, R., and D. Card (1991). “Recent Trends in Insured and Uninsured Unemployment: Is There an Explanation?” Quarterly Journal of Economics, vol. 106(4), pages 1157-1189.
- [8] Bontemps, C., J.-M. Robin, and G. J. Van den Berg (2000). “Equilibrium Search with Continuous Productivity Dispersion Theory and Non Parametric Estimation,” International Economic Review, vol. 41(2), pages 305-358.
- [9] Bound, J., and A. Krueger (1991). “The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?,” Journal of Labor Economics, vol. 9(1), pages 1-24.
- [10] Burdett, K. (1978). “A Theory of Employee Job Search and Quit Rates,” American Economic Review, vol. 68, pages 212-220.
- [11] Burdett, K. and M. Coles (2000). “Equilibrium Wage-Tenure Contracts,” Econometrica, vol. 71, pages 1377-1404.
- [12] Burdett, K., and D. T. Mortensen (1998). “Wage Differentials, Employer Size, and Unemployment,” International Economic Review, vol. 39(2), pages 257-273.
- [13] Bureau of Labor Statistics (2006). “Employee Tenure Summary,” <http://www.bls.gov>.
- [14] Bunzel, H., B. J. Christensen, P. Jensen, N. M. Kiefer, L. Korsholm, L. Muus, G. R. Neumann, and M. Rosholm (2001). “Specification and Estimation of Equilibrium Search Models,” Review of Economic Dynamics, vol. 4, pages 90-126.
- [15] Christensen, B. J., R. Lentz, D. T. Mortensen, G. R. Neumann, and A. Wervatz (2005). “On the Job Search and the Wage Distribution,” Journal of Labor Economics, vol. 23(1), pages 31-58.

- [16] Cowell, F. (2000). "Measurement of Inequality," Handbook of Income Distribution, A. Atkinson and F. Bourgignon eds., North Holland, pages 87-166.
- [17] Dey, M. and C. Flinn (2006). "Household Search and Health Insurance Coverage," mimeo, NYU.
- [18] Eckstein, Z., and G. Van den Berg (2005). "Empirical Labor Search: A Survey," IZA Discussion paper.
- [19] Eckstein, Z., and K. Wolpin (1990). "Estimating a Market Equilibrium Search Model from Panel Data on Individuals," *Econometrica*, vol. 4, pages 783-808.
- [20] Eckstein, Z., and K. Wolpin (1995). "Duration to First Job and the Return to Schooling: Estimates from a Search-Matching Model," *Review of Economic Studies*, vol. 62, pages 263-286.
- [21] Fallick, B. and C. A. Fleischman (2004). "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows" Finance and Economics Discussion Series 2004-34. Washington: Board of Governors of the Federal Reserve System.
- [22] Flinn, C. (2006). "Minimum Wage Effects on Labor Market Outcomes Under Search, Bargaining, and Endogenous Contact Rates," *Econometrica*, vol. 74, pages 1013-1062.
- [23] Flinn, C. and J. Heckman (1983). "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?" *Journal of Labor Economics*, vol. 1(1), pages 28-42.
- [24] French, E. (2005). "The Labor Supply Response to Mismeasured (but Predictable) Wage Changes," forthcoming in the *Review of Economics and Statistics*.
- [25] Gaumont, D., M. Schindler, and R. Wright (2006). "Alternative Theories of Wage Dispersion," *European Economic Review*, vol. 50(4), pages 831-848.
- [26] Hagedorn, M., and I. Manovskii (2006). "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," mimeo, University of Pennsylvania.
- [27] Hall, R. E. (2005). "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, vol. 95(1), pages 50-65.
- [28] Hansen, H. (1998). "Transition from Unemployment Benefits to Social Assistance in Seven European OECD Countries," *Empirical Economics*, vol. 23(1), pages 5-30.
- [29] Hornstein, A., P. Krusell and G.L. Violante (2006). "Technical Appendix for Frictional Wage Dispersion in Search Models: A Quantitative Assessment," available at <http://www.econ.nyu.edu/user/violante/research.htm>
- [30] Kambourov, G. and I. Manovskii (2004). "Occupational Mobility and Wage Inequality," mimeo.
- [31] Jolivet, G., F. Postel-Vinay and J.-M. Robin (2006). "The Empirical Content of the Job Search Model: Labor Mobility and Wage Dispersion in Europe and the U.S.," *European Economic Review*, vol. 50(4), pages 877-907

- [32] Jovanovic, B. (1979). "Job Matching and the Theory of Turnover," *Journal of Political Economy*, vol. 87(5), pages 972-990.
- [33] Kostiuk, P. (1990). "Compensating Differentials for Shift Work," *Journal of Political Economy*, vol. 98, pages 1054-1075.
- [34] Lagos, R. and N. Kiyotaki (2006). "A Model of Job and Worker Flows," mimeo, NYU.
- [35] Levine, D. and W. Zame (2002). "Does Market Incompleteness Matter?" *Econometrica*, vol. 70(5), pages 1805-1839.
- [36] Lise, J. (2006). "On-the-Job Search and Precautionary Savings: Theory and Empirics of Earnings and Wealth Inequality," mimeo, Queens University.
- [37] Ljungqvist, L. and T.J. Sargent (1998). "The European Unemployment Dilemma," *Journal of Political Economy*, vol. 106, pages 514-550.
- [38] Lucas, R. E., and E. Prescott (1974). "Equilibrium Search and Unemployment," *Journal of Economic Theory*, vol. 7(2), pages 188-209.
- [39] Machin, S., and A. Manning (1999). "The Causes and Consequences of Long-Term Unemployment in Europe," in *Handbook of Labor Economics*, (edited by Ashenfelter, O., and D. Card), North-Holland.
- [40] Manning, A., and B. Petrongolo (2005). "The Part-Time Pay Penalty," CEP Discussion Paper No. 679.
- [41] McCall, J. (1970). "Economics of Information and Job Search," *Quarterly Journal of Economics*, vol. 84(1), pages 113-126.
- [42] Mood, A., F. Graybill, and D. Boes (1974). *Introduction to the Theory of Statistics*, McGraw-Hill
- [43] Mortensen, D. T. (1970). "Job Search, the Duration of Unemployment, and the Phillips Curve," *American Economic Review*, vol. 60(5), pages 847-62.
- [44] Mortensen, D. T. (2005). *Wage Dispersion: Why Are Similar Workers Paid Differently?*, MIT Press.
- [45] Mortensen, D. T. (2007). Discussion of "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," by A. Hornstein, P. Krusell, and G.L. Violante. EFG Meetings, San Francisco, February.
- [46] Mortensen, D. T., and C. Pissarides (1994). "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, vol. 61(3), pages 397-415.
- [47] Moscarini, G. (2003). "Skill and Luck in the Theory of Turnover," mimeo, Yale University.
- [48] Nagypal, E. (2005). "Worker Reallocation Over the Business Cycle: The Importance of Job-to-Job Transitions," mimeo, Northwestern University.

- [49] Nagypal, E. (2006). "On the Extent of Job to Job Transitions," mimeo, Northwestern University.
- [50] OECD (2004). *Taxes and Benefits in the United States*, OECD.
- [51] Pissarides, C. (1985). "Short-run Equilibrium Dynamics of Unemployment Vacancies, and Real Wages," *American Economic Review*, vol. 75(4), pages 676-90.
- [52] Pissarides, C. (2000). *Equilibrium Unemployment Theory*, MIT Press.
- [53] Postel-Vinay, F. and J.-M. Robin (2002). "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, vol. 70(6), pages 2295-2350.
- [54] Rogerson, R., R. Shimer, and R. Wright (2006). "Search Theoretic Models of the Labor Market," *Journal of Economic Literature*, vol. 43 (4), pages 959-988.
- [55] Rosholm, M. and M. Svarer (2005). "Endogenous Wage Dispersion in a Search-Matching Model," *Labor Economics*, vol. 11(5), pages 623-645.
- [56] Ruggles, S., M. Sobek, A. Trent, C. A. Fitch, R. Goeken, P. K. Hall, M. King, and C. Ronnander (2004). "Integrated Public Use Microdata Series: Version 3.0," Minneapolis, MN: Minnesota Population Center. <http://www.ipums.org>.
- [57] Ruhm, C. (2003). "Good Times Makes You Sick," *Journal of Health Economics*, vol. 22(4), pages 637-658.
- [58] Shimer, R. (2005a). "Reassessing the Ins and Outs of Unemployment", mimeo, University of Chicago.
- [59] Shimer, R. (2005b). "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, vol. 95(1), pages 25-49.
- [60] Stevens, M. (1999). "Wage-Tenure Contracts in a Frictional Labour Market: Firms' Strategies for Recruitment and Retention," *Review of Economic Studies*, vol. 71, pages 535-551.
- [61] Topel, R. (1991). "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority," *Journal of Political Economy*, vol. 99(1), pages 145-76.
- [62] Van den Berg, G., and G. Ridder (1998). "An Empirical Equilibrium Search Model of the Labor Market," *Econometrica*, vol. 66(5), pages 1183-1221.





Figure 1: Top panel: Residual wage distribution for full-time, full-year janitors and cleaners in the Philadelphia area. Bottom panel: Distribution of mean-min ratios for full-time, full-year janitors and cleaners across U.S. geographical areas. Source: authors' computation on 1990 Census.





Figure 2: Top panel: Distribution of Mp_5 ratios across labor markets in the U.S. economy. Source: authors' computation on 1990 Census. Bottom panel: Distribution of Mm ratios for individual labor market histories. Source: authors' computation on PSID.







Figure 3: Empirical analysis on PSID data. The first-stage residuals refer to the regression on observable covariates. The second-stage residuals are the first-stage residuals demeaned individually.




Figure 4: Pairs of the value of non-market time and the interest rate that can generate $Mm = 1.70$.

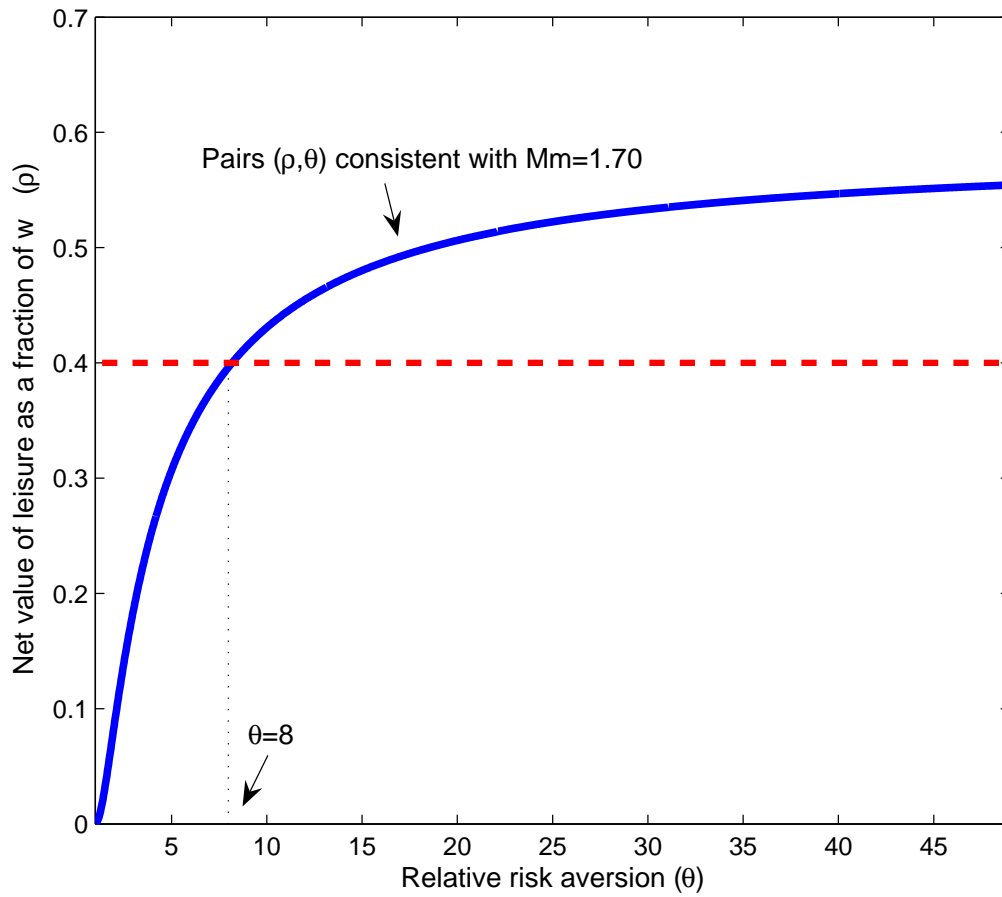


Figure 5: Pairs of the value of non-market time and risk aversion that can generate $Mm = 1.70$.




Figure 6: Pairs of the value of non-market time and the wage autoregression coefficient that can generate $Mm = 1.70$.







Figure 7: Top panels: pairs of the value of non-market time and the job offer rate during employment that can generate $Mm = 1.70$. Bottom panels: mapping between the arrival rate of offers on the job and average job tenure (left panel), and separation rate (right panel).