NBER WORKING PAPER SERIES

IMPLICATIONS OF SEARCH FRICTIONS: MATCHING AGGREGATE AND ESTABLISHMENT-LEVEL OBSERVATIONS

Russell Cooper
John Haltiwanger
Jonathan L. Willis

Working Paper 13115
http://www.nber.org/papers/w13115

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

We are grateful to an anonymous referee, Robert King, Shigeru Fujita, Murat Tasci and Eran Yashiv for helpful comments. Comments from Stephen Nickell, Julian Messina and participants at the 2006 Swiss National Bank/JME Conference and seminar participants at Ohio State University, Washington University of St. Louis and the Federal Reserve Banks of Chicago, Dallas, Kansas City and New York for comments and suggestions. The authors thank the NSF for financial support. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City, the NBER or the Federal Reserve System.

Implications of Search Frictions: Matching Aggregate and Establishment-level Observations
Russell Cooper, John Haltiwanger, and Jonathan L. Willis
NBER Working Paper No. 13115
May 2007
JEL No. E24,J6

ABSTRACT

This paper studies hours, employment, vacancies and unemployment at micro and macro levels. It is built around a set of facts concerning the variability of unemployment and vacancies in the aggregate and, at the establishment level, the distribution of net employment growth and the comovement of hours and employment growth. A search model with frictions in hiring and firing is used as a framework to understand these observations. Notable features of this search model include non-convex costs of posting vacancies, establishment level profitability shocks and a contracting framework that determines the response of hours and wages to shocks. The search friction creates an endogenous, cyclical adjustment cost. We specify and estimate the parameters of the search model using simulated method of moments to match establishment-level and aggregate observations. The estimated search model is able to capture both the aggregate and establishment-level facts.

Russell Cooper
Department of Economics
BRB, 2.102A
University of Texas
Austin, TX 78712
and NBER
cooper@eco.utexas.edu

John Haltiwanger
Department of Economics
University of Maryland
College Park, MD 20742
and NBER
haltiwan@econ.umd.edu

Jonathan L. Willis
Federal Reserve Bank of Kansas City
925 Grand Blvd.
Kansas City, MO 64198
jonathan.willis@kc.frb.org

Implications of Search Frictions: Matching Aggregate and Establishment-level Observations *

Russell Cooper†, John Haltiwanger‡, and Jonathan L. Willis§

May 6, 2007

Abstract

This paper studies hours, employment, vacancies and unemployment at micro and macro levels. It is built around a set of facts concerning the variability of unemployment and vacancies in the aggregate and, at the establishment level, the distribution of net employment growth and the comovement of hours and employment growth. A search model with frictions in hiring and firing is used as a framework to understand these observations. Notable features of this search model include non-convex costs of posting vacancies, establishment level profitability shocks and a contracting framework that determines the response of hours and wages to shocks. The search friction creates an endogenous, cyclical adjustment cost. We specify and estimate the parameters of the search model using simulated method of moments to match establishment-level and aggregate observations. The estimated search model is able to capture both the aggregate and establishment-level facts.

1 Motivation

This paper estimates a search model of the labor flows exploiting both microeconomic and macroeconomic data. Aggregate search models tend to focus on aggregate data, largely ignoring microeconomic evidence. For these models, the search and matching process is the key labor market friction. In contrast, many studies of labor adjustment costs emphasize microeconomic observations, but the aggregate implications of these models remain unclear. For these studies, the friction is modeled through a rich specification of costs

*We are grateful to an anonymous referee, Robert King, Shigeru Fujita, Murat Tasci and Eran Yashiv for helpful comments. Comments from Stephen Nickell, Julian Messina and participants at the 2006 Swiss National Bank/JME Conference and seminar participants at Ohio State University, Washington University of St. Louis and the Federal Reserve Banks of Chicago, Dallas, Kansas City and New York for comments and suggestions. The authors thank the NSF for financial support. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System.

†Department of Economics, University of Texas at Austin and the NBER.

‡Department of Economics, University of Maryland and the NBER.

§Research Department, Federal Reserve Bank of Kansas City.

of adjusting employment. This paper attempts to break this dichotomy by estimating the parameters of a search model based upon observations at the establishment-level as well as aggregate data.

A leading study of the aggregate implications of a search model is Shimer (2005). He argues that the standard search model, based upon Mortensen and Pissarides (1994), fails to match certain key features of aggregate data on worker flows.[1] In particular, Shimer reports that the model lacks a mechanism to magnify shocks. The standard deviation of average labor productivity is about equal to the standard deviations of unemployment and vacancies. But in the data, the standard deviations of both vacancies and unemployment are about 10 times the standard deviation of average labor productivity.[2] As discussed by Shimer and others, the productivity impulses are apparently dampened by movements in the real wage, which in turn reduce the incentive for creating more vacancies when productivity increases.

There are equally important, though less frequently cited, facts coming from observations at the establishment level. At the microeconomic level, employment adjustment is sporadic with periods of inactivity in employment adjustment followed by relatively large adjustments in the number of workers.[3] Aggregate search models are routinely evaluated relative to macroeconomic flows without reference to these observations of vacancy, employment and hours variations at the establishment level. This is unfortunate, both because the aggregate models miss important microeconomic facts, and because capturing these features at the micro level may enhance the fit at the aggregate level.

One goal of this paper is to propose and estimate a model of labor adjustment at the microeconomic level that is consistent with observations at both the aggregate and the establishment levels. To match establishment-level observations, our model extends the search model in a couple of important directions.

First, it includes a theory of a producer with multiple jobs. Second, we allow for fixed costs of posting vacancies at the establishment-level. Third, we estimate a process for profitability shocks at the establishment-level to match observations. These ingredients of the model allow us to match observed inactivity in employment flows at the establishment-level.

For the aggregate variables, the greater variability of the establishment-level shocks can increase the variability of labor flows. Accordingly, we include both the standard deviation of unemployment and vacancies as moments to match. Further, in light of the challenge to Shimer's conclusions in Hagedorn and Manovskii (2006), we also include cyclical movements in average aggregate wages in our set of moments.

Another goal of this paper is to bridge the gap between the literatures on search models and labor adjustment costs.[4] One hypothesis is that the literature estimating labor adjustment costs is uncovering the

[1]In an earlier paper, Cole and Rogerson (1999) looked at the implications of the Mortensen and Pissarides matching model relative to facts about aggregate job creation and job destruction. They report some success matching aggregate job flows.

[2]Not surprisingly, this has sparked a considerable discussion concerning these results *per se* as well as alternatives to the standard model meant to better confront these facts. See Yashiv (2006) for a survey of this ongoing research.

[3]See the discussion and references in Cooper, Haltiwanger, and Willis (2004) on job flows. Recent evidence on worker flows draws upon Davis, Faberman, and Haltiwanger (2006b).

[4]See, for example, the discussion in Nickell (1986) and Hamermesh and Pfann (1996) of the sources of these frictions in labor adjustment and empirical specifications to capture them.

implications of search frictions. While the search friction can be viewed as just another form of adjustment costs at the establishment level, it is different from other frictions because it is endogenous. That is, the cost of filling vacancies will generally depend on the tightness of labor markets, which is an equilibrium outcome.

The search model we estimate allows for both search costs and firing costs. Both types of adjustment costs are quite successful in matching both establishment level and aggregate movements. In fact, with the moments we have chosen, it is not possible to conclude that the model with hiring costs is superior to one with firing costs.

The estimated model does not suffer from the magnification puzzle highlighted by Shimer (2005). This is partly due to the role of the idiosyncratic shocks at the establishment-level which, given the non-convexities in the model, are not smoothed by aggregation. Further, the model matches moments on procyclical average wages so that aggregate shocks induce producers to create vacancies. Finally, in our model average labor productivity is endogenous, reflecting both aggregate and idiosyncratic shocks along with labor market frictions. So, part of the resolution of the magnification puzzle comes from the model's implications for the behavior of average labor productivity.

2 Facts

We present empirical evidence on key moments from both aggregate and establishment-level data which form the basis of our empirical analysis. First, section 2.1 provides facts on vacancies, unemployment, and productivity drawn heavily from Shimer (2005). The findings from the latter paper have spawned substantial debate in the literature (see, e.g., Hagedorn and Manovskii (2006)) as to whether standard search and matching models can capture the observed volatility and dynamics of vacancy and unemployment. Second, section 2.2 provides facts on hours and employment dynamics at the aggregate and establishment-level. A key theme of our analysis is to exploit the differences in moments at the aggregate and establishment-level. Third, section 2.3 continues with this theme by presenting facts about the establishment-level employment growth rate distribution and the relationship between this distribution and hours and separations at the establishment-level. Fourth, section 2.4 presents empirical evidence on the cyclicality of real wages. The evidence presented in this section is mostly drawn from the recent literature which we supplement from various databases from the U.S. statistical agencies. Together, this evidence provides a rich empirical characterization of U.S. labor markets and are used to estimate the parameters of our structural search model.

2.1 Unemployment, Vacancy and Productivity Dynamics

Table 1 summarizes the main findings on unemployment, vacancies and labor productivity, as in Shimer (2005).[5] The first part of the table reports U.S. data on unemployment u, vacancies v and average labor

[5] This table was produced by Murat Tasci and appears in Tasci (2006).

productivity, denoted p.[6]

There are three features of the data which deserve emphasis:

- the standard deviations of unemployment and vacancies are both about 10 times the standard deviation of average labor productivity,
- the data exhibit the Beveridge curve: the correlation between unemployment and vacancies is strongly negative,
- both unemployment and vacancies are highly serially correlated.

U.S. DATA	*(Quarterly, 1951Q1-2003Q4)*			
	u	v	v/u	p
Std	0.19	0.20	0.38	0.02
Auto	0.94	0.95	0.95	0.89
Cross Correlations				
u		-0.89	-0.97	-0.42
v			0.97	0.37
v/u				0.40
Standard Search Model				
	u	v	v/u	p
Std	0.01	0.02	0.03	0.02
Auto	0.85	0.74	0.81	0.81
Cross Correlations				
u		-0.87	-0.94	-0.94
v			0.99	0.99
v/u				0.99

For the data moments, the level of unemployment, u, is from the CPS, the level of vacancies, v, is from the Conference Board and average labor productivity, p, is real output per person. These variables are seasonally adjusted and are log deviations from an HP trend.

Table 1: Unemployment and Vacancies

The second part of the table reports moments from a "Standard Search Model" and come from Tasci

[6] These observations are quarterly, seasonally adjusted and detrended using a HP filter. Here unemployment is a level not a rate. See the discussion in Shimer (2005) for more data details.

(2006).[7] The structure of the model generating these moments is summarized in section 3.5.[8]

Comparing the two panels of Table 1, it is clear that the standard search model, as parameterized in Shimer (2005), is unable to create the volatility in unemployment and vacancies observed in the data. We refer to this as the amplification puzzle. Interestingly, the model is able to capture the negative correlation between vacancies and unemployment, the Beveridge curve.

In what follows, we use key moments on the relationship between vacancies and unemployment consistent with the findings in Table 1. In our analysis we focus on vacancy dynamics using the Job Openings and Labor Turnover Survey (JOLTS) data as opposed to the vacancy data from the Help Wanted Index used by Shimer (2005). As many have noted (including Shimer (2005)), the Beveridge curve and volatility of vacancies and unemployment using the JOLTS data (and JOLTS sample period) are similar to that as reported in Table 1.[9] We use the JOLTS data on vacancies since the latter data are monthly and as will be seen below we use the monthly establishment-level distribution of employment growth rates from JOLTS in our analysis. Given our focus on adjustment costs and search and matching frictions, treating a period as one month in our model is more justifiable. The key moments we use are the standard deviation of the unemployment rate, the standard deviation of vacancy rate and the correlation of unemployment and vacancies. From the JOLTS data and sample period, these three statistics are 0.086, 0.116 and -0.954, respectively.

2.2 Aggregate and Establishment-Level Hours and Employment Dynamics

Evidence about aggregate and establishment-level hours and employment dynamics is reported in Table 2. The data set underlying the moments reported in Table 2 contains observations of production workers in each quarter and their total hours worked during the quarter from the Longitudinal Research Database (LRD) for the period 1972-80.[10] This information allows us to compute average hours per worker on a quarterly basis. The LRD only has information on hours per worker on a quarterly basis for this sample period.

We use moments based on the growth (log first difference) of hours per worker and employment at the

[7]The results are very similar to those reported in Shimer (2005).

[8]The model follows that presented in Shimer (2005) and Tasci (2006). Cole and Rogerson (1999) contains a very clear summary of model essentials.

[9]The volatility of vacancies and unemployment is smaller from JOLTS. This may reflect the sample period as well as the data used – many have noted the vacancy volatility from the Conference Board HWI series may be implausibly high – see, e.g., Davis, Faberman, and Haltiwanger (2006a).

[10]The data are for manufacturing plants and are quarterly from 1972-80 and thus do not overlap with the establishment-level data used from JOLTS – see Table 3 below. The data set and its connections to data used in other establishment-level studies is discussed in detail in Cooper, Haltiwanger, and Willis (2004). In the quantitative analysis that follows, we are careful to distinguish between empirical moments from quarterly vs. monthly observations and in turn to take in account time aggregation issues. An open question is the representativeness of the moments for the manufacturing sector for the entire economy. We do know that the variance of establishment-level growth rates for the U.S. manufacturing sector is somewhat lower than, but roughly comparable to, that for the whole U.S. economy. As discussed in Cooper, Haltiwanger, and Willis (2004) the moments based upon aggregate data in Table 2 are robust to using BLS quarterly data on U.S. manufacturing hours and employment from 1972-2003. Thus, while the evidence in Table 2 is confined to manufacturing only it appears that key moments are robust to a longer time period than covered by the LRD.

establishment level in Table 2 to focus on the patterns of adjustment of the labor input rather than the size distribution of the establishments in our sample. In computing these moments at the establishment-level, year and seasonal effects have been removed. Thus these moments characterize the aspects of the cross-sectional distribution of employment and hours growth.[11]

There are three important observations about establishment-level adjustment worth emphasizing as observed in the moments reported in the first column of Table 2.

- the standard deviations of hours growth and employment growth are about the **same**,
- hours growth and employment growth are **negatively** correlated
- hours growth in one period is positively correlated with employment growth in the next.

These facts, particularly the negative correlation in employment and hours growth, were important in distinguishing models of labor frictions in Cooper, Haltiwanger, and Willis (2004) and Caballero, Engel, and Haltiwanger (1997) and play a similar role in this study of search frictions.

The second column of Table 2 is based on aggregating the establishment-level data used in the first column and considering the aggregate time-series variation alone. For the aggregate data,

- the correlation of hours and employment growth is **positive**,
- the standard deviation of employment growth **is almost twice** that of hours growth.
- hours growth tends to lead employment growth [12]

The sharp contrast between the establishment-level and aggregate hours and employment dynamics is a key theme in our analysis. One of the challenges for a model of labor adjustment is to explain the different patterns of the cross-sectional and time-series moments in Table 2.

2.3 Establishment-Level Employment Growth, Hires and Separations

A new establishment-level survey that enables measuring vacancies, hires, separations and employment growth is the Job Openings and Labor Turnover Survey (JOLTS), which samples about 16,000 establishments per month. Respondents report hires and separations during the month, employment in the pay period covering the 12th of the month, and job openings at month's end. They also report quits, layoffs and discharges, and other separations (e.g., retirements). Recent analysis of the establishment-level data from JOLTS is reported in Davis, Faberman, and Haltiwanger (2006b) and Davis, Faberman, and Haltiwanger (2006a). Drawing from Figures 6 and 7 of Davis, Faberman, and Haltiwanger (2006b), we provide summary information about the patterns of worker and job flows at the establishment level which is summarized in

[11] If we compute these moments on a quarterly basis for each of the quarters in our sample, the variation over time is relatively small. Thus we interpret these as cross-sectional moments.

[12] As noted by a referee, hours levels above trend at the industry level also tend to lead employment growth.

Moment	Plant	Aggregate
$\frac{\sigma_{\Delta h}}{\sigma_{\Delta e}}$	0.96	0.55
$Corr(\Delta h, \Delta e)$	-0.296	0.545
$Corr(\Delta h_{-1}, \Delta e)$	0.184	0.519

* Seasonal and Aggregate Effects removed from establishment-level moments.

Seasonal Effects removed from aggregate-level moments

Table 2: Hours and Employment Adjustment: Basic Facts from the LRD*

Table 3.[13] Net employment growth at the establishment-level is characterized by five bins as listed in the first column. The second column shows the share of employment growth in each of these five bins. The remaining columns decompose the employment growth into hires and separations. The column labeled "net" is the average employment growth within each of the bins.

Net Emp. Growth	Share of Emp.	Hires	Sep.	net
<-0.10	0.040	0.025	0.291	-0.266
-0.10 to -0.025	0.083	0.023	0.075	-0.052
-.025 to 0.025	0.745	0.015	0.015	0.000
0.025 to 0.10	0.092	0.079	0.028	0.051
>0.10	0.040	0.296	0.041	0.266

Table 3: Monthly Net Employment Growth Rate Distribution

These moments are size-weighted by employment share. So, for example, about 4.0% of workers were employed in an establishment which had net employment adjustment in excess of 10% in an average month.

From the first three columns of the table and additional examination of the data, a couple of facts stand out:

- there is a significant amount of relatively small net employment adjustment: about 74% of the size-weighted observations entail net employment adjustment between -2.5% and 2.5% in an average

[13] We thank these authors for the summary of the data points underlying these figures. Faberman (2005) provides a detailed discussion of the data set. See the detailed discussion in Davis, Faberman, and Haltiwanger (2006b) regarding the measurement methods to insure the timing of net growth and hires and separations line up.

month,

- there are significant bursts of job creation and destruction: many establishments (8%) either contract or expand employment by more than 10% in a month.

For each employment growth bin, Table 3 presents the average rates of hires and separations.[14] As illustrated in the table, establishments expanding employment by more than 10% do so through high hiring rates, though interestingly, separation rates are also higher for this group.

There is also substantial inaction in employment adjustment. About 32% of the size-weighted observations entail zero adjustment of net employment. The inaction rate on an establishment basis rises to 78%.[15] We also note that about 45% of the size-weighted observations entail zero vacancies.[16]

Our empirical analysis attempts to match the moments reported in the first two columns of Table 3. We also discuss the implications of our model for the observed patterns of hires and separations.

A particularly challenging aspect of matching the patterns of net employment growth is the presence of inaction along with a substantial fraction of observations with relatively small adjustment. With zero adjustment in 32% of the size-weighted observations, the remainder of the 74.5% of the observations with employment growth between -2.5% and +2.5% are relatively small adjustments. As we shall see in discussions of our models, the inaction is consistent with a model in which non-convexities are an important element of adjustment costs. But this model has difficulty explaining the small adjustments.

We interpret some of the very small adjustments as a form of noise reflecting factors outside of our model (and outside the scope of most of the the search and matching literature). For example, in JOLTS an establishment is not supposed to report vacancy postings (job openings) that arise from workers returning from a temporary leave of absence or a hire of a worker onto the payroll who was previously a contract or temporary employee. The point is that such changes in employment status of workers would show up in the measured net employment growth distribution in Table 3 but would not show up in vacancies. Our model like much of the search and matching literature focuses on the costs associated with posting vacancies and making hires – if some of the measured employment changes don't involve such costs then they are outside the scope of our model and much of the literature.[17]

In the analysis that follows we attempt to match the fraction of employment at establishments with very small adjustments but do not distinguish between zero adjustment and very small adjustments. This choice of moments reflects our concern with the measurement issues noted above as well as the fact that observations of large adjustments in employment, along with the correlation of hours and employment, are informative moments with respect to non-convexities in the labor adjustment process.

[14] These are all expressed in rates. Separations equals layoffs plus quits plus other separations – the latter are not reported but are small.

[15] This statistic is from Davis, Faberman, and Haltiwanger (2006b).

[16] This statistic is from Davis, Faberman, and Haltiwanger (2006a).

[17] Davis, Faberman, and Haltiwanger (2006a) present evidence that as many as a third of all hires occur without posted vacancies.

2.4 The Cyclicality of Real Wages

It is well known that quantifying empirically the cyclicality of real wages (see, e.g., Abraham and Haltiwanger (1995)) is a challenge. Empirical patterns are sensitive to time period, the alternative real wage series that can be constructed from aggregate and micro level data, and concerns about cyclical composition bias (Bils (1987) and Barsky, Solon, and Parker (1994)). In addition, the measurement-concept gap between the measurement of real wages in the data relative to the relevant shadow real wage is another challenge.

These measurement and conceptual issues have not been resolved but it is apparent from the recent search and matching literature that accounting for the cyclicality of real wages is essential. Much of the debate between Shimer (2005), Hall (2005) and Hagedorn and Manovskii (2006) about the ability of standard search and matching models to be able to account for the volatility and dynamics of unemployment and vacancies centers on what fraction of the driving processes are absorbed by variations in wages or the labor input.

In our setting, these issues are also relevant and in our analysis we follow the recent literature by using moments on the cyclicality of real wages. Since the measurement and conceptual issues regarding real wage cyclicality have not been resolved, we follow the basic approach of Hagedorn and Manovskii (2006) and explore the sensitivity of our results by considering a range of cyclical elasticities of real wages.

Three alternative measures of aggregate real wages produced by the BLS have been used in the literature (see, e.g., Abraham and Haltiwanger (1995) and for more recent evidence see Hagedorn and Manovskii (2006)). The most commonly used is average hourly earnings (AHE) for nonsupervisory workers from the BLS establishment surveys. An alternative measure that accounts for non-wage compensation and holds the industry and occupation composition of the workforce constant over the cycle is the Employment Cost index (ECI). The more recent literature (see, e.g., Hagedorn and Manovskii (2006) and Rudanko (2006)) focuses on a measure of real wages from the BLS productivity statistics program that is constructed in a manner consistent with the ouput per person and hour data produced and released by the BLS. That is, BLS constructs a measure of hourly compensation based upon the quarterly income from the NIPA attributable to labor.[18] We have examined the real wage cyclicality of all three series (using the CPI to convert nominal wages to real) and while we find they are positively correlated (especially the AHE and ECI series), the series with greater cyclical volatility is the BLS real hourly compensation series (RHC).[19]

In the recent literature (i.e., Hagedorn and Manovskii (2006) and Rudanko (2006)), the elasticity of real

[18]BLS constructs a measure of quarterly output per person and hour data using quarterly measures of output from the National Income and Product accounts (NIPA) and its own constructed measures of employed persons and hours. For purposes of being able to construct a measure of unit labor costs, BLS constructs a measure of hourly compensation based upon the quarterly income from the NIPA attributable to labor. In releasing its output per hour and unit labor cost series, BLS also releases the underlying hourly compensation data (in both nominal and real terms). It is this latter series that has been used in the recent literature – see Hagedorn and Manovskii (2006) and Rudanko (2006).

[19]In terms of coverage and definition of compensation, the ECI and RHC are more similar to each other than the AHE. However, we find a closer correspondence between the ECI and the AHE. See `http://www.bls.gov/lpc/faqs.htm#Q13` for discussion of the possible sources of differences between the ECI and RHC. We also note that for both the AHE and RHC, real wage cyclicality is greater if the 1960s and 1970s are included in the sample.

wages (measured as the HP-detrended real hourly compensation series) with respect to labor productivity (also measured as the HP-detrended series) is about 0.45. Moreover, Hagedorn and Manovskii (2006) reports that in using PSID data and controlling for composition bias that they obtain an elasticity of real wages at the micro level that is only slightly higher and about 0.47.[20] To explore sensitivity of their results to this elasticity, Hagedorn and Manovskii (2006) use the 95 percent confidence interval of their real wage elasticities. We have obtained the same data and replicate this elasticity of 0.45 with a 95 percent confidence interval of 0.33 to 0.57.[21]

In our analysis, we use the point estimate of 0.45 along with the confidence interval range to explore sensitivity. In considering this range, it is useful to note that this elasticity is on the high end of the patterns observed in the data. We find that the elasticity of the AHE real wage series with respect to productivity growth is 0.26 with a 95 percent confidence interval of 0.21 to 0.32.

We also find that the RHC real wage series yields an elasticity with respect to employment growth of 0.16 with a 95 percent confidence interval of 0.05 to 0.27 while the AHE real wage series yields an equivalent elasticity of 0.20 with a 95 percent confidence interval of 0.16 to 0.25. These patterns show that the RHC real wage series tracks labor productivity more closely than AHE but both track cyclical employment about the same.[22]

3 Model

This section specifies the model. We first provide an overview and then discuss the components in detail.

3.1 Model Overview

There are two types of agents in the model: producers and workers. Producers operate production sites which use labor as an input.[23] The labor input is total hours and thus combines the number of employees and hours worked per employee. There are both aggregate and producer-specific shocks which create revenue from the labor input.

[20] This relatively modest impact of composition bias is roughly consistent with the literature. Barsky, Solon, and Parker (1994) find the most pronounced composition bias and their elasticity of real wages with respect to unemployment changes from 0.006 to 0.116 due to composition bias.

[21] We use a shorter sample period, 1964-2006 than Hagedorn and Manovskii (2006) (who uses the 1951-06 period) but obtain the same elasticity. We use the shorter time period since this permits us to explore sensitivity with other real wage and cyclical indicators (e.g., AHE).

[22] There is an open question whether the correlation between the RHC real wage series and average labor productivity is spuriously high since measurement error in quarterly GDP will yield a positive correlation between RHC and average labor productivity.

[23] In this discussion, producers operate a production site and not a firm. This is consistent with our establishment level observations and assumes that firms with multiple establishments operate them independently, at least with respect to employment decisions.

Workers and producers are brought together through a search process. A worker who is matched with a producer has hours and compensation specified through a state-contingent contract. The worker may lose this job in a subsequent period, thus returning to a state of unemployment. Reflecting the search friction, workers without a job are assumed to find a new job with some probability each period. This probability is exogenous to the worker but is determined in equilibrium.

Producers have, at a point in time, a set of workers with whom they have a contract. In the short-run, the producer responds to variations in a profitability shock, reflecting both productivity and demand, through changes in hours worked per employee. The contract determines the response of hours and compensation to the shock.

Producers also can create vacancies and hence change the number of employees. The process of creating and filling vacancies entails adjustment costs. We allow for both fixed and variable costs of posting vacancies. The presence of these fixed costs distinguishes our model from the existing search literature and defines the boundaries of a producer. Empirically, these fixed costs are important for matching observed inaction in the adjustment of the number of workers.

With this structure in mind, we reconsider the motivation for this exercise. The assumption that producers have a fixed cost of posting vacancies generates some inaction in vacancies and employment adjustment. This is consistent with establishment-level evidence. Further, the observed negative correlation between hours worked and the number of workers could reflect the inaction in posting vacancies. A producer not hiring in the current period responds to higher demand for its product by increasing the hours of its workers. But once the producer decides to post vacancies and hires more workers, average hours worked drops.

From the perspective of matching moments on worker, job and unemployment flows as well as vacancies, there are a couple of points to raise. First, the driving process for the model is establishment-specific profitability. Various studies find that productivity at the establishment level is considerably more variable than in the aggregate. Thus, perhaps one resolution of the magnification problem highlighted in Shimer (2005) is through the presence of volatile establishment-level shocks. The key is that, perhaps, these shocks will create job and worker flows without increasing the measured variability of average labor productivity.[24]

As in Yashiv (2000), the model we consider has costs of adjusting the number of vacancies.[25] As well as creating inaction in vacancy creation and hiring, the non-convexity in adjustment may also increase the volatility of job and worker flows.

3.2 Workers

In general, workers are in one of two states, employed or unemployed. If unemployed, the workers enjoy leisure time and/or the fruits of home production, $U(b(a))$, which is allowed to depend on the level of aggregate

[24] The mapping from a distribution of establishment-specific profitability shocks to an aggregate measure of average labor productivity is likely to entail some smoothing through aggregation and by worker flows across producers.

[25] Yashiv (2000) does not allow for non-convexity in the costs of vacancy creation and thus can estimate parameters from Euler equations. His specification does include differentiable non-linearities in the adjustment costs.

productivity, a. With a positive probability, unemployment workers become employed in the subsequent period. Formally, the value of unemployment for a worker is given by the following:

$$V^u(\bar{u}, \bar{v}, a) = U(b(a)) + \beta E[f(\bar{u}, \bar{v})V^e + (1 - f(\bar{u}, \bar{v}))V^u] \tag{1}$$

where $f(\bar{u}, \bar{v})$ is the job finding rate which depends on the unemployment rate and the aggregate vacancy rate, $\bar{u}$ and $\bar{v}$ respectively. Here there is an expectations operator associated with the future value of employment, V^e and V^u. The future value of employment is random since jobs with different producers may lead to different levels of compensation and hours. The randomness of the future value of unemployment reflects uncertainty in the future value of leisure and in the job finding rate.

Employed workers have a contract for the current period which governs their state-contingent compensation and hours worked. Workers do not save in equilibrium so that compensation and consumption are identical. We specify utility of consumption (c) and hours worked as $U(\omega - g(h))$.[26] Here $U(\cdot)$ is strictly increasing and strictly concave and $g(\cdot)$ is strictly increasing and strictly convex.[27]

The value of being employed is given by

$$V^e = E\{U(\omega - g(h)) + \beta[(1 - S)V^e + SV^u]\} \tag{2}$$

where S is the separation (quits plus fires) rate.[28] The first term is the expected utility given a contract, described below, with a producer. In the following period there is a probability, given by S, of losing their jobs and becoming unemployed. Workers discount at rate β. As noted earlier, both V^e and V^u are state contingent.

The expectation in (2) is taken with respect to idiosyncratic shocks at the producer with whom the worker is employed as well as future state variables. As discussed below, we give all of the bargaining power to producers. Thus the exact details of what information workers have is, in equilibrium, immaterial. The assumption of giving producers all of the bargaining power immensely simplifies the analysis since workers do not care which producer they work for. Otherwise, the different producers would offer different terms and have different retention probabilities, and workers would have to keep track of the entire distribution of producers.

3.3 Producers

Producers have access to a technology which creates output from labor input. The revenue function is given by

[26] The model can accommodate more general utility functions but this specification is particularly tractable.

[27] While worker's do face uncertainty in (net) consumption, $\omega - g(h)$, workers all have the same level of net utility *ex post*. Thus workers have no incentive to trade state contingent consumption. Further, there is no store of value in the model and thus compensation and consumption are the same.

[28] This value may be producer specific but is ignored in the notation.

$$a\varepsilon(eh)^{\alpha} \tag{3}$$

where a is the aggregate (profitability) shock, ε is the producer-specific shock and total labor input is the product of the number of workers, e, and hours per worker, h. We allow for curvature in the revenue function, parameterized by α, which may capture diminishing returns to scale due to fixed factors of production excluded from (3).

We assume two stages in the producer's problem. First, given the aggregate state, the producer contracts with its workers. Second, *ex post*, the producer-specific shock is realized, and state-contingent hours are determined given the contract.

3.3.1 Setting a Contract

A contract is $\delta = (\omega(s), h(s))$ for all s, where $s = (a, \varepsilon, e, \theta)$ is the establishment's state. This state vector includes $\theta \equiv \frac{\bar{v}}{\bar{u}}$ which measures the tightness of labor markets. This aggregate variable, as explained in more detail below, summarizes the state of the labor market and is needed for the producer to predict the ease of hiring workers.

In the state contingent contract, $\omega(s)$ is compensation and $h(s)$ is hours worked. The contract allows compensation and hours to be fully state contingent. In terms of timing, the contract is determined given (a, e) but prior to the determination of ε. This timing is consistent with the literature on risk sharing through labor contracts and is important for matching moments on relative employment and hours variability. All workers with a given producer get the same contract since they are identical and have the same outside option of unemployment.

For the contracting process, assume producers make a take-it-or- leave-it offer to workers. This implies that employed workers get no surplus: $V^e = V^u$. Therefore the value of employment in equilibrium is independent of the producer with whom the worker has a job.

Thus the producer selects the contract to solve

$$\pi(a, e) = max_{\delta} E_{\varepsilon}[a\varepsilon(eh(s))^{\alpha} - e\omega(s)] \tag{4}$$

where the expectation is over the idiosyncratic component of profitability. The constraint is that the expected utility from the contract not be less than the outside option of unemployment, V^u:

$$V^e(a) = E_{\varepsilon}U(\omega(s) - g(h(s))) + \beta E_{a'|a}[(1 - S)V^e(a') + SV^u(a')] \geq V^u(a) \tag{5}$$

where V^e is the value of being employed next period and V^u is the value of being unemployed next period.[29]

[29] These values depend on aggregate productivity through $b(a)$ but do not depend on labor market tightness since $V^e(a) = V^u(a)$ for all a.

Assuming the worker's participation constraint binds, $V^e = V^u$ for all a, $E_\varepsilon U(\omega(s) - g(h(s))) = U(b(a))$ for all s. Given the risk aversion of the workers, optimal risk sharing implies that marginal utility is independent of the realized value of ε. Thus compensation and hours will satisfy the condition

$$U(\omega(s) - g(h(s))) = U(b(a)) \tag{6}$$

for all s.

Interestingly, producer heterogeneity is present in compensation levels, and hours reflect producer-specific state variables and shocks. So there will be a non-degenerate cross-sectional distribution of (ω, h) but a degenerate cross-sectional distribution of utility levels given the aggregate state.

Variations in the aggregate state, acting through $b(a)$, will influence the terms of the contract. In this way, the model produces movements in the aggregate wage without the complexity of allowing the workers to have bargaining power.[30] This added feature is important for matching observations on aggregate wage movements and for tempering the response of vacancies and unemployment to aggregate shocks.

3.3.2 Determining Hours

Once ε is realized, hours are determined by the contract. With (6) holding for all s, workers are fully compensated for hours variations. So the optimal hours choice is easy to characterize. Given (a, e, ε), the producer chooses a level of hours subject to (6), i.e., $\omega = g(h) + b(a)$. The producer's optimization problem for hours is

$$max_h a\varepsilon(eh)^\alpha - eg(h) - eb(a) \tag{7}$$

The hours choice satisfies

$$\alpha a\varepsilon(eh)^{\alpha-1} = g'(h). \tag{8}$$

This first order condition generates a state-dependent policy function, $h(a, e, \varepsilon)$, where hours depends on (a, e, ε). Holding (a, ε) fixed, as e increases, it is clear that h falls. This is relevant for matching a negative correlation in observed hours and employment growth at the establishment level.

3.3.3 Determining the level of Employment

The level of employment is determined by the vacancy-posting decision of the producer.[31] The recruiting decision is made knowing $\tilde{s} \equiv (a, e_{-1}, \varepsilon_{-1}, \theta)$ where e_{-1} is the inherited stock of workers and ε_{-1} is the

[30] We are grateful to the referee and Borghan Narajabad for suggesting this addition to the model. Acemoglu and Hawkins (2006) analyze a model where firms have multiple workers with bargaining. As they note, this approach has many challenges and their model does not have the rich features of hours per worker and adjustment costs that are the focus of our paper.

[31] Here vacancies must be reposted each period. See Fujita and Ramey (2005) for a model, where vacancies are a state variable.

shock last period used to predict the current one. The state vector $\tilde{s}$ is similar to s, except for the timing of decisions on e and the realization of the idiosyncratic profitability shock.

$Q(\tilde{s})$ is the value of the establishment in state $\tilde{s}$ and is given by

$$Q(\tilde{s}) = max\{Q^h(\tilde{s}), Q^n(\tilde{s}), Q^f(\tilde{s})\}. \tag{9}$$

In this optimization problem, $Q^h(\tilde{s})$, $Q^n(\tilde{s})$ and $Q^f(\tilde{s})$ relate to the hiring, no-adjustment and firing options.

The value of hiring workers is given by

$$Q^h(\tilde{s}) = max_v E_{e,\varepsilon}\pi(a, \varepsilon, e) - F^+ - C^+(v) + \beta EQ(\tilde{s}') \tag{10}$$

When the producer posts v vacancies, the evolution of employment is

$$e = e_{-1}(1 - \bar{q}) + H(\bar{u}, \bar{v})v \tag{11}$$

where $\bar{q}$ is the quit rate and $H(\cdot)$ is the rate at which a vacancy is filled. In this formulation, the vacancy filling rate depends on $\bar{u}$, the unemployment rate, and $\bar{v}$, the aggregate vacancy rate. This is where the aggregate state of the economy influences the magnitude of the adjustment cost for employment. For some specifications, as the labor market tightens ($\bar{u}$ is low relative to $\bar{v}$), producers find that filling vacancies is expensive so that more of the labor input variation arises through hours rather than workers.

For a given aggregate state of the economy, $H(\cdot)$ and $\bar{q}$ are deterministic from the producer's viewpoint. That is, stochastic aspects of the matching process and quits are not part of the uncertainty facing a producer.

There are two types of costs of posting vacancies in the model: a fixed cost component, F^+, and a variable cost component, $C^+(v)$. A familiar interpretation of this type of specification is based upon recruiting in Economics. The fixed cost appears in the form of reading numerous files, flying a committee to interview and so forth. The variable cost is related to the number of interviews and fly-outs. In terms of matching the moments, these two costs are relevant for capturing inaction, through F^+, and partial adjustment, through $C^+(v)$.

The value of firing workers is given by

$$Q^f(\tilde{s}) = max_f E_\varepsilon \pi(a, \varepsilon, e_{-1}(1 - \bar{q} - f)) - F^- - C^-(f) + \beta EQ(\tilde{s}'). \tag{12}$$

Here the level of employment reflects quits and fires. There are fixed, F^-, and variable costs, $C^-(f)$, of firing workers.

The value of inaction is given by

$$Q^n(\tilde{s}) = E_\varepsilon \pi(a, \varepsilon, e_{-1}(1 - \bar{q})) + \beta EQ(\tilde{s}'). \tag{13}$$

Here inaction means no hiring and no firing so that employment at the establishment level falls due to quits.

We assume that any profits realized by producers are consumed by entrepreneurs who own the production process. These agents are risk neutral and thus are the natural suppliers of insurance to workers. Producers discount at the same rate, β, as do workers. There is no free entry in the model.

The solution of this optimization problem generates two decision rules. First, on the extensive margin, there is the issue of adjustment of the stock of workers. Second, on the intensive margin, if the producer decides to hire (fire) workers, there is the choice of the magnitude of the adjustment.

3.4 Equilibrium

An equilibrium for this economy requires optimization by producers and workers and consistency conditions. For the optimization problems, the components of an equilibrium are:

- an optimal labor contract which solves (4) subject to the participation constraint of the workers,
- a state-contingent hours schedule which solves (8),
- a decision rule for employment adjustment which solves (9),
- a decision rule for workers entailing acceptance or rejection of the contract, as in (5).

With regards to the consistency conditions, the vacancy-filling rate appearing in the employment transition constraint in the optimization problem of the producers, (11), depends on θ, an aggregate variable determined in equilibrium. This function, which is taken as given in the optimization problem of the producer, must be consistent with the relationship generated by the model and the data. As described below, this consistency is enforced in our estimation.

Finally, the unemployment rate follows $\bar{u}' = (1-\bar{u})S(\bar{u},\bar{v}) + (1-f(\bar{u},\bar{v}))u$ where $S(\bar{u},\bar{v})$ is the separation rate (quits plus layoffs) and $f(\bar{u},\bar{v})$ is the job finding rate. In equilibrium, $0 \leq \bar{u} \leq 1$. Because all workers are either employed or unemployed in our model, we use the transition equation for total employment to generate an unemployment series to evaluate the moments reported in Table 1.

3.5 Comparison to Standard Search Model

Section 2.1 discusses key moments of worker flows and, following Shimer (2005), points to key differences between the data and the standard search model. Here we briefly outline the standard search model relative to the model we consider. This discussion draws upon Shimer (2005) and Tasci (2006).

There are a couple of key differences between the models. The standard search model assumes each producer has at most one worker. Firms without a worker post a vacancy at a (flow) cost and that vacancy is either filled or not.

Further, there is no movement on the intensive margin in the standard model. That is, variations in hours are not studied, and thus there is no state-contingent contract governing the response of hours to shocks. So matching hours variations is not possible.

Shocks in the standard model are common across producers.[32] Thus the models are not equipped to match cross-sectional observations.

In the standard model, wages are determined by Nash bargaining. For the moments in Table 1, the bargaining weight for workers was set at 0.72.[33] The bargaining weight clearly determines the split of the surplus. Further, compensation responds to the aggregate state of the economy through the bargaining process. In our setting, we create cyclicality of real wages through cyclical variations in the opportunity cost of workers' time.

Once workers have a share of the surplus, then the utility flow during a period of unemployment becomes more important to the analysis. Shimer (2005) assumes a value of leisure at 40% of average productivity. Hagedorn and Manovskii (2006) take a more general view of the value of leisure beyond the replacement rate from unemployment insurance and, by matching moments of labor flows and the elasticity of wages with respect to productivity variations, set the value of leisure at 0.955 (as a fraction of average productivity) and the bargaining share of workers at only 0.05. As seen in their Table 4, this parameterization resolves the problem of the relative standard deviation of unemployment and vacancies.

In many respects, the optimal contracting structure, along with the bargaining power held by producers we explore in our model, is much closer to the parameterization of Hagedorn and Manovskii (2006). We have given the producer all of the bargaining power so that the workers have a zero weight. Further, the optimal contract stabilizes worker's *ex post* utility at the state-contingent value of leisure. The property that workers receive insurance over employment status is a common feature of optimal contracting models.[34] In our model, we allow work sharing so that hours vary. Hence, there are no *ex post* employment variations within a period, and thus inclusion of unemployment insurance in the optimal contract does not arise.

Also, the fact that the utility flow to workers is determined by an exogenous value of leisure can generate a form of "wage stickiness" as in Hall (2005). The extent of this stickiness depends on the relationship between the value of leisure and the aggregate shock, $b(a)$ and not the effect of the aggregate shock on productivity. So here, the degree of "wage stickiness" is endogenous rather than imposed. This property of our model reflects the optimal insurance features of the labor contract rather than an assumption of inflexible real wages. Further, *ex post* wages *per se* are not allocative in our framework. Instead, hours respond to the state and satisfy (8), and compensation adjusts so that utility is equal to $U(\omega - g(h)) = U(b(a))$ *ex post*. Thus the model has implications for the cross-sectional distributions of hours and compensation which we do not exploit.

Finally, the standard model has a free entry condition, which pins down the value of a vacancy at zero. Instead, we have producers optimally choosing the number of vacancies to post. Our focus is more on the intensive margin of adjustment in the number of vacancies per producer rather than the number of producers.

[32] Mortensen and Pissarides (1994) have job-specific shocks, which they use to generate job destruction. Cole and Rogerson (1999) also allow for idiosyncratic shocks, but find that these shocks do not have aggregate effects.

[33] See Shimer (2005) and Hagedorn and Manovskii (2006) for a discussion of the parameterization of the bargaining weight.

[34] See Azariadis (1975) for details as well as a model in which there are unemployment risks due to the absence of severance pay and work sharing.

4 Numerical Analysis

The optimization problem for an individual producer is solved through value function iteration of (9) and the functions used to define this value. There is no household optimization to consider: as long as $V^e = V^u$, workers are indifferent between accepting a job or not.

To solve the producer's optimization problem, we need a number of parameterized functions. We discuss the functions here and then summarize them, along with parameterizations in Table 4. We specify and parameterize the model at a monthly frequency. Thus the moments associated with worker flows are not time aggregated.

4.1 Functional Forms

As argued above, the wage, $\omega(s)$, will satisfy $U(\omega(s) - g(h(s))) = U(b(a))$ so that $\omega(s) = b(a) + g(h(s))$ for all s. We parameterize the disutility of work, $g(h)$, so that

$$\omega(s) = b(a) + \omega_1 h(s)^{\zeta} \tag{14}$$

is the compensation required to guarantee the utility level $U(b(a))$ in all states, s. Here ζ is important for determining the utility cost of variations in hours. Generally, if ζ is low, then variations in hours are inexpensive, so that much of the adjustment of the labor input will be on the intensive margin and not through variations in the number of workers.

We have allowed the workers value of leisure to depend on the aggregate shock, a. Thus we assume $b(a) = b_0 a^{b^1}$. The parameter b^1 thus governs the sensitivity of the worker's outside option to aggregate productivity.

For the hiring and firing of workers, we assume that the variable cost of vacancies is given by $C(v) = c_0^+ v^{c_1^+}$. The variable cost of firing workers is similarly parameterized, $C(f) = c_0^- f^{c_1^-}$.

Following the literature, the matching function is specified with constant returns to scale and is given by

$$m = \mu \bar{u}^{\gamma} \bar{v}^{1-\gamma} = \mu \bar{v} \theta^{-\gamma} \tag{15}$$

where $\theta \equiv \frac{\bar{v}}{\bar{u}}$ measures the tightness of labor markets. From this relationship, we obtain two additional functions: the vacancy filling rate for producers and the job finding rate for workers. Let $H = \frac{m}{\bar{v}}$ be the vacancy filling rate. Using the specification of the match rate, (15),

$$H = \mu \theta^{-\gamma}. \tag{16}$$

Let $f = \frac{m}{\bar{u}}$ be the job finding rate for workers. Using (15) again,

$$f = \mu \theta^{1-\gamma}. \tag{17}$$

Of course, these three functions are related by common parameters.

The solution of the producers optimization problem requires the vacancy filling rate in the transition equation for employment, (11). Using (16), the vacancy filling rate depend on the current state of labor market tightness, θ. Thus the solution of the producer's optimization problem requires it to know the evolution of θ in the equilibrium of the model economy. Given the heterogeneity of producers, forecasting labor market tightness requires knowledge of the cross sectional distribution of workers and producer specific profitability as well as the aggregate variables. This is a complex forecasting problem computationally.

To simplify the analysis, we assume a form of bounded rationality or limited information by producers: they forecast θ using an AR(1) process. For this reason, we have included θ in the state vector of the producer.

We estimate an AR(1) representation of θ from our data. From that regression, the AR(1) coefficient on θ is 0.93. The R^2 from the regression is 0.96. The simple autoregressive structure does an outstanding job of capturing the dynamics of labor market tightness in our sample.

For our simulations, we impose beliefs on the producer. In the estimation of the model, we add moments to guarantee that the process for θ is consistent with beliefs of producers.

4.2 Choosing Parameters

We estimate the parameters for the vacancy filling rate directly from data. We also calibrate some of the other parameters. This is for the purpose of the simulations of the model. We return to this discussion in the context of estimating the model.

4.2.1 Estimated Functions

This section reports our estimates of the parameters for the vacancy filling rate function. This estimation is done directly from the data and is thus outside of the solution of the producer's dynamic optimization problem.

For this analysis as well as the estimation, we construct a monthly vacancy series using JOLTS. The monthly unemployment data is from the CPS. Relative to Shimer (2005), our data are higher frequency and we use the JOLTS vacancy series. For monthly labor productivity, we construct a series using the Industrial Production Index from the FRB and total hours using employment and hours data from the BLS. Labor market tightness, θ, is the log of the vacancy-unemployment ratio. All series are converted to logs and then HP filtered.

The primary relationship estimated directly from these data is the matching function, where the logarithm of the match rate is regressed on the logarithm of labor market tightness. In keeping with a large part of the literature we impose constant returns to scale. With this restriction we estimate $\gamma = 0.36$, which is considerably lower than the estimate of 0.72 reported in Shimer (2005) and closer to the estimate of 0.235 reported in Hall (2005). The difference in estimates may reflect the use of the JOLTS data rather than the Conference Board vacancies numbers and the use of different sample periods. The logarithm of the constant

Relationship	Parameter
discount rate (β)	0.9966
curvature of profit function (α)	0.65
elasticity of disutility of hours (ζ)	2.90
serial correlation of aggregate shocks (ρ_a)	0.95
standard deviation of innovation to aggregate shocks (σ_a)	0.0016
curvature of matching function (γ)	0.36

Table 4: Calibrated Parameters

is estimated at 0.0072. Hence μ from (15) is estimated at 1.0072.[35]

To check on these estimates, we measure the monthly job finding rate of workers at 0.606. Using (17), the estimate of $\gamma = 0.358$ and the mean value of $\theta = 0.46$ imply $\mu = 1.0009$. These are very close to the point estimate from the regression.

4.2.2 Calibrated Parameters

We calibrate a subset of our parameters, summarized in Table 4. We use this calibration for the simulation results in Section 5 and in the estimation reported in Section 6.

As the model is monthly, we set $\beta = 0.9966$. There are some parameters, b_0 and ω_1 which are set to match average establishment size and average hours.[36] The value of α is in the range for this parameter as estimated in Cooper, Haltiwanger, and Willis (2004). The calibrated value of ζ follows Caballero and Engel (1993).

A couple of these parameters are particularly important for the analysis. The wage elasticity, ζ, governs the cost of adjusting hours and is relevant for matching the relative standard deviation of hours growth to employment growth.

The common component of the profitability shock, a, is modeled as a log normal AR(1) process. We set the serial correlation, ρ_a, and standard deviation of the innovation to the process, σ_a, to bring the models predictions close to the serial correlation and variance of aggregate employment.

5 Simulations

To get a sense of how the model performs, we conduct a preliminary analysis by simulating the model for some two specifications of adjustment costs. In particular we look at the following special cases:

[35]The standard error on the estimate of γ and μ equals 0.02.

[36]For our model, the value of b_0 influences the size of producers, not the difference in utility between employed and unemployed agents. We set $b_0 = 1.12$ and $\omega_1 = 0.000013$ so that average establishment size is 60 and hours equal 40 on average. The average size of 60 is consistent with the data set, described in Cooper, Haltiwanger, and Willis (2004), used to estimate α.

- **NC**: No Costs of Adjusting the labor force,
- **H-Fix**: Fixed Costs of Posting Vacancies,

The exact specification of the vacancy posting costs are provided below. The other parameters are set as described in Table 4.

For each specification of adjustment costs, we do two experiments with the sensitivity of the value of leisure with respect to the aggregate shock. In one case, we set $b^1 = 0$ so that the value of leisure is constant. In the other, we set $b^1 = 0.5$ so that the value of leisure and thus aggregate wages are procyclical.

Further, for these simulations we set producers beliefs directly by assuming a functional relationship between the vacancy filling rate and a. That is, we impose $H = \nu_0 a^{\nu_1}$ to approximate $H(\theta)$ and the dependence of θ on a.[37]

For these simulations, the number of establishments was set at 8000 and the simulation length was 360 months. There were 21 points in the space of idiosyncratic shocks. The simulation results did not change much when these settings were expanded to more time periods, more establishments and a finer grid.

5.1 Key Moments

The simulation exercise is structured around three set of moments from the data. These moments are of interest partly because they are associated with empirical puzzles. Here we discuss these key moments and then study how our model addresses them.

The first set of key moments concerns aggregate variables. We focus on the variability of unemployment and vacancies as well as their comovement. These moments are summarized in Table 5 and differ from those in Table 1 as we are using monthly JOLTS data to compute these statistics. In addition, we look at the comovement of aggregate wages and productivity.

The second key moment is the observed negative correlation between employment and hours growth at the establishment level, as indicated in Table 2. The third set of moments is the cross-sectional distribution of net employment growth, summarized in Table 3.

The challenge is to find a model capable of matching these disparate observations from micro and aggregate data. As illustrated in the simulations below, our model's inclusion of non-convex costs of vacancy creation along with substantial dispersion in establishment-specific shocks provides a basis for matching these moments.

5.2 Simulation Results

Table 5 presents the main moments of interest for the cases listed above. Almost all of these moments are calculated on a monthly basis, the same frequency as most of the data. The exception are the two quarterly

[37] For the simulations we set $\nu_0 = 1.3, \nu_1 = -30.00$. This produces a serial correlation of θ around 0.85 across the treatments and thus is a good approximation to the estimated process of θ for these simulations.

moments from the LRD, $\sigma_{\tilde{h}}/\sigma_{\tilde{e}}$ and $corr(\tilde{h}, \tilde{e})$, which are calculated for each quarter by sampling from the simulated monthly data.

The first specification, NC, has no costs of hiring and firing. The consequence of this is substantial volatility in unemployment, matching and job finding, relative to the other cases. In contrast, there is relatively little variability in vacancies. This difference in variability of unemployment relative to vacancies reflects the fact that unemployment is a state variable. The model does generate a Beveridge curve.

Due to the assumption that hours but not employment respond immediately to the idiosyncratic shock, adjustment of both hours and employees occurs. But without adjustment costs, much of the variation is in the form of employment growth, so $\sigma_{\tilde{h}}/\sigma_{\tilde{e}}$ is much higher in the data than in the model. Interestingly, the model without adjustment costs reproduces the negative $corr(\tilde{h}, \tilde{e})$ in the data.

Looking at employment adjustment, it is clear that in the absence of adjustment costs, the volatility of employment growth at the establishment level is excessive relative to observation. In the simulation it is not uncommon to see employment growth, in absolute value, in excess of 10%. This finding should not be surprising given the volatility of the idiosyncratic shocks. Clearly a key issue in the estimation is the identification of the costs of creating vacancies and the other adjustment costs from the variance of the idiosyncratic shocks.

The second row of the table retains the assumption of no adjustment costs but assumes that the value of home production is procyclical: $b^1 = 0.5$. This creates positive correlation between the aggregate component of profitability and the average wage rate.

Relative to these results, the introduction of costs of posting vacancies and firing are relevant for reducing the variability of job and worker flows. To illustrate, the H-fix specification introduces a fixed cost, $F^+ = 1$, into the model. At this value of F^+, the average adjustment cost incurred, given in the "AC" column, is about 4.7% of monthly gross profits (revenues less compensation to workers). With this relatively high adjustment cost, there is relatively more inaction, i.e. more employment growth in the $(-2.5, 2.5)$ bin of the employment distribution and less frequent bursts of job creation and destruction. Further, the distribution of employment growth is skewed: there are few small positive employment growth rates but more small negative employment growth rates. In this sense, the adjustment cost alters the distribution of employment variation.

Since the quit rate is assumed positive (at 1.7%), there are no observations of zero net employment growth. There is substantial inaction on vacancies: 93% of the observations entail zero vacancies compared to 44.6% in the NC specification.

The Beveridge curve is still present in the simulated data. The fixed cost increases the variability of vacancies, relative to the no-adjustment cost case. There is a negative correlation between hours and employment growth at the establishment level and employment growth remains more variable than hours growth.

Model	Wages and ALP	Unemployment and Vacancies			LRD: $\Delta e, \Delta h$		JOLTS: Δe					AC
	$\omega \mid p$	$\sigma(u)$	$\sigma(v)$	$corr(u,v)$	$\sigma_{\tilde{h}}/\sigma_{\tilde{e}}$	$corr(\tilde{h},\tilde{e})$	< -0.10	$-0.10 < -0.025$	$-0.025 < 0.025$	$0.025 < 0.10$	> 0.10	
Data	0.45	0.09	0.12	-0.95	0.96	-0.30	0.04	0.08	0.75	0.09	0.04	NaN
NC	0.00	0.21	0.09	-0.69	0.30	-0.65	0.31	0.01	0.13	0.01	0.53	0.00
NC, $b^1 = 0.5$	0.28	0.18	0.09	-0.53	0.30	-0.65	0.31	0.01	0.13	0.02	0.53	0.00
H-Fix	0.00	0.19	0.10	-0.71	0.38	-0.53	0.22	0.03	0.57	0.00	0.17	4.72
H-Fix,$b^1 = 0.5$	0.21	0.15	0.10	-0.56	0.38	-0.53	0.22	0.03	0.57	0.00	0.17	4.72

Table 5: Simulation Results

Note: $\sigma(x)$ is the standard deviation of x. Unemployment comes from the CPS and vacancies are directly calculated from JOLTS at a monthly frequency. $\omega|p$ is the coefficient from regressing log of hourly compensation on the log of average labor productivity. The variables are in logs and are HP filtered. The employment adjustment moments are size weighted using employment.

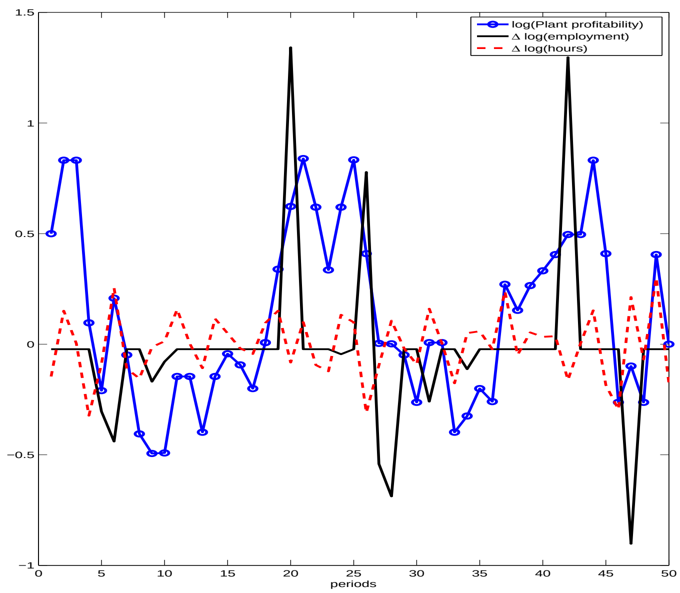

..

Figure 1: Employment and Hours: Establishment Level Growth Rates

There is a second specification with fixed costs of posting vacancies in which we allow the flow value of unemployment to depend on a: $b(a) = b_0 a^{b^1}$ with $b^1 = 0.5$. In this model, labeled **H-Fix,** $b^1 = 0.5$, variations in a are positively correlated with the value of leisure so that in high productivity states, wages rise as well. In theory, this ought to reduce the value of creating vacancies in high productivity states.

We find that, relative to the H-fix case, the variability of vacancies falls by around 5% while the variability of unemployment falls by nearly 20%. Further, the correlation between unemployment and vacancies is clearly reduced. This parameterization of the model does a better job of matching the correlation between wages and average labor productivity.

Interestingly, none of the plant specific moments are influenced by the state dependent outside option of the worker. This reinforces the view that the plant-level moments are reflecting cross sectional rather than time series variation.

5.3 Inspecting the Mechanism

This is a rather rich model, and the mapping from parameters to moments is not immediately clear. To help build further intuition about the models mechanics, we explore in more detail the specification with a fixed costs of posting vacancies, the H-fix in Table 5. We do so here by presenting two figures related to key moments.

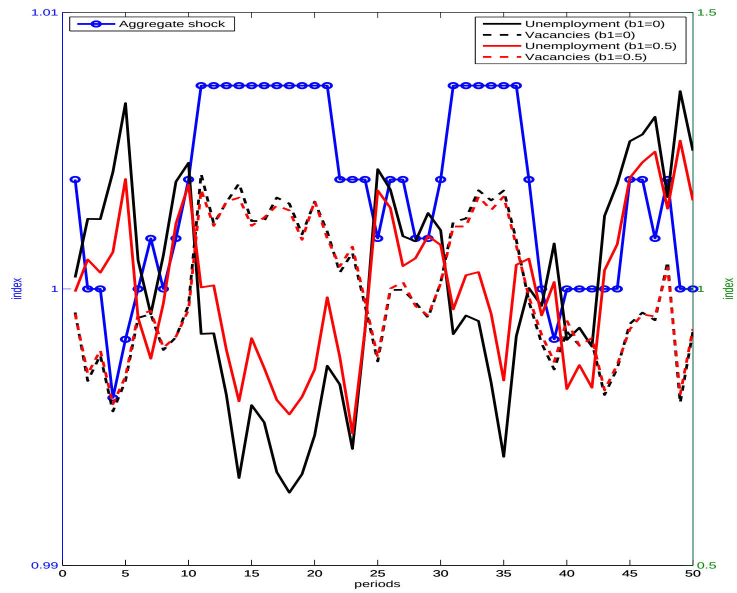

..

Figure 2: Unemployment and Vacancies: Aggregate

Figure 1 shows simulation results at the establishment level. The point here is to understand how an establishment, in the presence of fixed costs of posting vacancies, responds to variations in profitability. These simulations assume $b^1 = 0$.[38]

Two points are illustrated in Figure 1. First, hours and employment are negatively correlated. When the producer is subject to an increase in profitability, starting around period 17 in Figure 1, hours respond immediately. In period 19, the adjustment cost is paid and employment adjusts to a higher level. Hours are reduced as employment expands, and this produces the negative correlation between hours and employment growth. When profitability falls around period 25, first hours and then employment is reduced.

Second, due to the adjustment costs, employment does not always respond to variations in profitability. This is evident from period 34-40 in Figure 1. Though profitability has risen in this period, there is no employment response. Instead, fluctuations in profitability are met by variations in hours. Thus variations in adjustment are sporadic.

Figure 2 shows simulation results for unemployment, vacancies and the aggregate component of the profitability shock. These aggregate variables are obtained by the aggregation of the establishment level results for the same simulation shown in Figure 1. There are two cases shown here: $b^1 = 0$ and $b^1 = 0.50$.

The Beveridge curve is apparent in these simulations. When there is a positive aggregate shock, such as around period 10, there is an immediate response in the creation of vacancies. Unemployment falls as

[38] As noted earlier, behavior at the producer level is essentially independent of b^1.

vacancies are filled. The strength of this response depends partly on the cost of creating vacancies and on the rate in which vacancies are filled.

The main difference effect of having the value of leisure (home production) sensitive to the aggregate state, i.e. $b^1 = 0.50$, is that unemployment falls less rapidly in response to a positive aggregate shock. This is consistent with the moments presented in Table 5: it is the standard deviation of unemployment which responds to the increase in b^1.

There are inherent differences in the dynamics of the response of unemployment and vacancies to the shock. In most of these search and matching models, unemployment is a state variable but vacancies are not.[39] As reported in Table 1, the serial correlations of unemployment and vacancies are about the same. As is evident from Figure 2, the serial correlation of vacancies created by the model is substantially less than the serial correlation of unemployment.

6 Estimation

The key parameters in our study are those determining the costs of hiring and firing as well as the driving process for the shocks at the establishment level. These parameters are estimated through a simulated method of moments procedure. Other parameters are calibrated at the values in Table 4.

6.1 Methodology

The estimation entails finding the vector of structural parameters, Λ, to minimize the (weighted) distance between moments from the data, Γ^d, and moments produced from a simulation of the model given a vector of parameters, $\Gamma^s(\Lambda)$. Thus our estimate of Λ minimizes $\pounds(\Lambda)$ where

$$\pounds(\Lambda) \equiv (\Gamma^d - \Gamma^s(\Lambda))W(\Gamma^d - \Gamma^s(\Lambda))' \tag{18}$$

and W is a weighting matrix.[40]

This minimization problem is solved by simulation to create a mapping from Λ to the moments. The methodology is as follows. Given vector Λ, we solve the producer's dynamic optimization problem using value function iteration. From this and the solution to (8), we generate policy functions at the producer level for employment, vacancies and hours. The model is solved at a monthly frequency. Using these policy functions, we create a simulated data set at the producer level. Given this, we compute the microeconomic moments directly from the data and, by aggregation, compute aggregate flows of vacancies and unemployment as well. In this manner, we obtain $\Gamma^s(\Lambda)$.

The simulated data set consists of 8000 establishments simulated over 360 months. The results are robust to increasing the number of establishments and time periods. The number of points in the grid for the idiosyncratic shock was 21.

[39] As noted earlier, one exception is Fujita and Ramey (2005).

[40] In the discussion which follows, W is an identity matrix that produces consistent estimates of Λ.

For our analysis, the parameter vector we estimate includes: $\Lambda = (F^{+}, c_0^{+}, F^{-}, c_0^{-}, b^1, \rho_\varepsilon, \sigma_\varepsilon)$. The first three parameters represent the cost of posting vacancies and the second three are firing costs. The parameters $(\rho_\varepsilon, \sigma_\varepsilon)$ characterize the log normal AR(1) process for the establishment-specific profitability shocks, where σ_ε is the standard deviation of the innovation to the process. The final parameter, b^1, captures the sensitivity of the worker's value of leisure to variations in aggregate profitability. Other parameters are set at the values used in the simulation exercises, summarized in Table 4.

We separate the moments to match into four categories:

- **Equilibrium**: The estimated model must mimic the regression results for AR(1) representation of labor market tightness, denoted as ρ_θ.
- **Unemployment and Vacancies**: The key moments are $\sigma(\bar{u})$, $\sigma(\bar{v})$ and $corr(\bar{u}, \bar{v})$. The moments are reported in Table 5.
- **Hours and Employment**: The key moments are $corr(\Delta e, \Delta h)$ and $\sigma_{\tilde{h}}/\sigma_{\tilde{e}}$. These are in terms of growth rates and are measured in the data and simulation quarterly, as reported in Table 2 for the establishment level.
- **Worker and Job Flows**: The key moments are from the distribution of Δe reported in Table 3.
- **Aggregate Wages and Productivity**: The key moment is the coefficient from a regression of (log) average hourly wages on (log) average labor productivity. This moment, which has no structural interpretation, is denoted $\omega|p$.

These moments are chosen largely because they characterize basic aspects of worker and job flows at both the microeconomic and aggregate levels. This is in accord with the point of our analysis: to investigate a search model capable of jointly explaining both microeconomic and macroeconomic facts.

The inclusion of the AR(1) coefficient for labor market tightness implies that the beliefs of the producers in the model, which conform with the empirically observed serial correlation, match the data. In this way, we also ensure that we have an equilibrium: the beliefs of the producers are mutually consistent.

6.2 Results

The estimation is undertaken for four cases: two with hiring costs and two with firing costs. While the model focuses on the significance of search and thus hiring costs, we also study the implications of firing costs. This provides a perspective on the role of search costs and also on the more general topic of the source of labor adjustment costs in general.

Each of the two hiring cost specifications includes a fixed cost of vacancies in conjunction with either a quadratic or linear cost of vacancies. Each of the two firing cost cases includes a fixed cost of firing in conjunction with either a quadratic or linear firing cost.

Specification	F^+	c_0^+	F^-	c_0^-	b^1	ρ_ε	σ_ε	AC	$\pounds(\Lambda)$
Hiring Costs									
Fixed, Quadratic ($c_1^+ = 2$)	0.056	0.001	0	0	0.369	0.338	0.186	0.550	0.059
Fixed, Linear ($c_1^+ = 1$)	0.135	0.007	0	0	0.501	0.331	0.227	0.775	0.024
Firing Costs									
Fixed, Quadratic ($c_1^- = 2$)	0	0	0.220	0.858	0.575	0.894	0.115	0	0.048
Fixed, Linear ($c_1^- = 1$)	0	0	0.112	0.024	0.487	0.459	0.196	0.235	0.024

Table 6: Estimation Results: Parameters

The results are reported in the following two tables. The parameter estimates are reported in Table 6. The moments of the models relative to data are summarized in Table 7 and discussed in section 7.1.1.

Table 6 shows the parameter estimates for the four specifications, adjustment costs as a percentage of gross profits (AC) and the fit of the model from (18). For both hiring and firing costs, the fixed and quadratic specifications do not do as well as the fixed and linear adjustment cost cases. The specification with linear and fixed costs of hiring workers fits the moments essentially as well as the specification with linear and fixed hiring costs. We have not been able to improve the fit using a specification with both hiring and firing costs.

Note that this does not mean the model with hiring costs is observationally equivalent to the model with firing costs. Looking at the distribution of net employment growth, the specifications clearly have very different implications for these moments.

For the specification with fixed and linear hiring costs, the estimated costs of adjustment (paid) are around 0.77% of gross profits. The costs are much lower with firing costs since workers quit at an exogenous rate and thus producers can avoid this cost. The idiosyncratic profitability shocks are serially correlated and much more variable than aggregate shocks. The value of b^1 in both specifications is around 0.5 indicating a strong response of the value of leisure to aggregate profitability.

7 Evaluation of Results

This section summarizes our findings. We first discuss how well the estimated model matches key moments. We then discuss other implications of the estimated model.

7.1 Explaining the Moments and Additional Facts

From Table 6, the specification with fixed and linear hiring costs and with fixed and linear firing costs seem to fit the moments equally well. Given the focus of this paper on hiring costs, we call this the "best fit" specification in the discussion which follows though we often refer to the firing cost results as well. Here we discuss how that model matches the moments in more detail and also touch on other moments.

7.1.1 Matching the Moments

Table 7 summarizes the moment implications for the cases including the best fit. Both of the specifications with linear adjustment costs do well matching moments in many dimensions. Both models, largely through b^1 are able to match the regression coefficient of wages on average labor productivity, $\omega|p$. With regard to aggregate facts, the models matches the variability of unemployment and vacancies and produce a Beveridge curve. At the establishment level, the estimated models match the relative volatility of hours and employment growth as well as the negative correlation between hours and workers.

Further, both models reproduce the serial correlation in θ from the data. Thus in both cases, the beliefs of producers about the evolution of labor market tightness and thus the vacancy filling rate are consistent with the data and the outcome of the model economy.

Differences between the specifications emerge in the distribution of employment growth at the producer level. The model with hiring costs predicts too little hiring and an excessive level of net employment growth smaller than 10% as there are no firing costs. Conversely, the model with firing costs produces fewer burst of firing but some intermediate rates of positive net employment growth. Neither model alone is fully capable of capturing the entire employment growth distribution.

Model	wage and ALP	AR(1) of θ	Unemployment and Vacancies			LRD		JOLTS: Δe				
	$\omega \| p$	ρ_θ	$\sigma(u)$	$\sigma(v)$	$corr(u,v)$	$\sigma_{\tilde{h}}/\sigma_{\tilde{e}}$	$corr(\tilde{h},\tilde{e})$	< -0.10	$-0.10 < -0.025$	$-0.025 < 0.025$	$0.025 < 0.10$	> 0.10
Data	0.45	0.93	0.09	0.12	-0.95	0.96	-0.30	0.04	0.08	0.75	0.09	0.04
Hiring Costs												
Fixed, Quadratic	0.41	0.89	0.13	0.08	-0.82	0.92	-0.43	0.06	0.11	0.72	0.00	0.11
Fixed, Linear	0.44	0.90	0.13	0.10	-0.90	0.99	-0.36	0.05	0.07	0.79	0.00	0.08
Firing Costs												
Fixed, Quadratic	0.39	0.94	0.12	0.07	-0.91	0.88	-0.26	0.00	0.00	0.88	0.04	0.07
Fixed, Linear	0.45	0.94	0.12	0.08	-0.93	0.97	-0.39	0.01	0.00	0.78	0.10	0.10

Table 7: Estimation Results: Moments

Note: $\sigma(x)$ is the standard deviation of x. Unemployment comes from the CPS and vacancies are directly calculated from JOLTS at a monthly frequency. The variables are in logs and are HP filtered. The employment adjustment moments are size weighted using employment.

As noted earlier, we set $\omega|p$ at 0.45. If we reduce the sensitivity of aggregate wages to average labor productivity and set $\omega|p = 0.26$, our conclusions do not change very much. The models with linear hiring and firing still fit the moments best. The lower $\omega|p$ leads to a lower estimate of b^1. Consistent with the simulation results, the variability of both unemployment and vacancies is higher. As noted in the discussion of the simulation results, the establishment level moments are not very sensitive to changes in b^1. Hence changing the $\omega|p$ moment does not have a significant effect on our estimates of adjustment frictions.

7.2 TFP vs. Average Labor Productivity: Unemployment and Vacancies

One of the more interesting points of difference between the empirical literature on aggregate search models and the macroeconomic literature based on the stochastic growth model is the measurement of productivity. The search literature largely looks at average labor productivity (ALP), while the macroeconomics literature focuses on total factor productivity (TFP) as an exogenous shock.

In this paper, we are closer to the tradition of the stochastic growth model. We have treated profitability as exogenous where the profitability shock summarizes both technological and demand factors influencing the revenues of a producer. If the model is competitive, then the profitability shock is most naturally interpreted as a variation in technology.[41]

In our model, revenues are specified in (3). For the competitive economy, revenues are the same as output. If labor is freely mobile (no adjustment costs and no rigidities due to timing assumptions) between production sites, then the cross sectional distribution of ALP will be degenerate. Further, fluctuations in the average product of labor will reflect the dependence of worker opportunities on aggregate productivity, b^1 in the model.

Generally though, TFP and ALP are not the same. There are two economic forces that separate these measures of productivity: (i) frictions in the adjustment process and (ii) idiosyncratic shocks.

To see these influences, think about two extreme economies. In one, suppose that labor flows freely across producers, and in the second, suppose there are frictions in labor flows. Suppose the distribution of idiosyncratic shocks is the same in the two economies and is fixed over time.

As aggregate TFP varies, ALP varies in both of these economies. For fixed TFP in the second economy, ALP increases as labor flows from less productive to more productive producers. Thus, variations in ALP will generally reflect both TFP and frictions.

Still, ALP is much easier to measure and thus plays a prominent role in the empirical literature. One of the advantages of our simulation environment is that we can use our model to generate a measure of ALP, for producer i in period t as

$$a_t \varepsilon_{i,t}(e_{i,t}h_{i,t})^{\alpha-1}. \tag{19}$$

[41] Allowing monopolistic competition allows for an interpretation of α as including a markup and the shocks to revenue including a relative demand disturbance. As the model studied here does not include product differentiation, α and the shocks should be more narrowly interpreted.

The value of α used for our analysis is 0.65. In the context of the model, $\alpha < 1$ reflects the presence of fixed factors of production such as managerial ability, structures and predetermined components of the equipment stock.

With these differences between TFP and ALP in mind, we return to a discussion of moments. The puzzle posed by Shimer (2005) concerns the standard deviation of unemployment and vacancies relative to ALP, as shown in Table 1. Our estimation, in contrast, has not focused on this moment *per se* but rather we target the absolute standard deviations on unemployment and vacancies.

Given our estimated model, we simulate average labor productivity and relate it to unemployment and vacancies. Our results are summarized in Table 8 along with relevant moments from the JOLTS data.

Here our measure of average labor productivity is denoted p. For the simulated data, it was computed at the establishment level using (19). The aggregate measure is total output divided by aggregate hours. For the actual data, ALP was computed using industrial production as a measure of output and total hours, eh, for production workers as the measure of labor input. Though the data on vacancies are from JOLTS, rather than the Conference Board, and are at a monthly frequency, the standard deviations of unemployment and vacancies are similar in magnitude to those reported in Shimer (2005).

Moment	JOLTS data	Estimated Model
$\sigma_{\bar{u}}/\sigma_a$	-	28.04
$\sigma_{\bar{v}}/\sigma_a$	-	21.49
$\sigma_{\bar{u}}/\sigma_p$	7.64	43.36
$\sigma_{\bar{v}}/\sigma_p$	10.34	33.21
$\text{corr}(\theta, p)$	-0.695	0.439
$\text{corr}(a, p)$	-	0.598

Table 8: Unemployment, Vacancies and Average Labor Productivity

From this table there is clearly a substantial difference between a and p: the correlation is 0.60. As discussed earlier, this reflects the interaction of frictions in labor flows and the idiosyncratic shocks in the estimated model.

There are two measures of the standard deviations of $\bar{u}$ and $\bar{v}$ relative to productivity. Looking at a, both unemployment and vacancies are substantially more volatile than productivity. This is also true when we use p as a productivity measure. The magnification puzzle isolated by Shimer (2005) is not present in this model.

The volatility of unemployment and vacancies relative to the variability in a stems from a few features of the labor market in our model. First, producers have all of the bargaining power. Second, the outside option of workers depends on a through the parameter b^1. Recall that b^1 is estimated in the model and is useful in matching the relationship between compensation and productivity, the $\omega|p$ moment.

To see the linkages, an increase in a leads to an increase in the profitability of hiring, which creates incentives on both the extensive and the intensive margins: a larger fraction of producers to post vacancies and post more vacancies. But, as $b^1 > 0$, the outside option of the workers is more valuable and thus average compensation rises with a. This effect dampens incentives for vacancy creation.

It is also worth noting that the time-series volatility of p is much less than the volatility of a. This is consistent with the intuition above that, in effect, with the estimated value of b^1 the economy is operating on a relatively flat labor supply curve so that average productivity fluctuates less than the aggregate shock.

Hagedorn and Manovskii (2006) argue that the elasticity of wages to average labor productivity is one for Shimer (2005). To study the importance of this elasticity for our results, we estimated a version of our model with $\omega|p = 1.0$ instead of $\omega|p = 0.0.45$. This change in the moments increased the estimate of b^1 from 0.5 to 0.87. For this case, there is still no amplification puzzle: $\sigma_{\bar{u}}/\sigma_p = 26.80$ and $\sigma_{\bar{v}}/\sigma_p = 24.7$. Apparently, our model can resolve the magnification puzzle with the same response of aggregate wages to productivity as assumed in Shimer (2005).

The behavior of ALP and labor market tightness poses another challenge to the analysis. As shown in Table 8, there is a negative correlation between labor market tightness and average labor productivity in the sample period data. But in the estimated model, as reported in the table, $corr(\theta, p) > 0$. In Shimer (2005), the relationship between productivity and labor market tightness is not stable across sub-samples of the data set. For his overall sample, as in our model, the correlation is positive, but for the sub-samples covered by JOLTS this correlation is also negative.

This difference between model and data may reflect the fact that average labor productivity is endogenous. Thus reverse causality may be an important consideration.

One link between ALP and θ comes from the effects of a, the aggregate exogenous driving force in the model. As a increases, so does p, as indicated by the positive correlation in the simulated data. Further, the correlation between a and θ is about 0.8 in the simulated data. This leads to the positive correlation between θ and p and is captured in the model.

The correlation between ALP and θ may also reflect other influences on ALP. For example, suppose there are variations in θ independent of a. A reduction in θ leads, based upon the regression results in section 4.2.1, to an increase in the vacancy filling rate. In this way, search frictions may be endogenously reduced, leading to a more efficient allocation of labor across producers. This reallocation is reflected in a higher value of p and a negative correlation between θ and p. In fact, in the data p is positively correlated with the vacancy filling rate.

Understanding the instability between ALP and labor market tightness, as well as the factors influencing the interaction, remains an area for further work.

7.3 Is Search the Basis for Models of Adjustment Costs?

There is a vast literature on the implications of labor adjustment costs for employment dynamics.[42] But to some extent, this literature suffers from the problem that the adjustment costs appear as a black box. When pressed, researchers sometimes use a search model as a basis for the costs of creating new jobs.

Given our study of a search model with costs of posting vacancies, it is natural to ask how the results of the model relate to findings on adjustment costs of labor. Superficially, the mapping seems solid. The fixed costs of posting vacancies match with the fixed cost of adjusting labor, as in the structural model of Cooper, Haltiwanger, and Willis (2004). The quadratic adjustment costs, which generate partial adjustment, may simply reflect the matching process in which only a fraction of vacancies are filled each period.[43]

Though it is outside this paper, the mapping between the search model and the labor adjustment cost models can be studied more formally. One possibility, along the lines of indirect inference, is to simulate data for the model estimated here. Then, consider a structural model of labor adjustment costs, along the lines of Cooper, Haltiwanger, and Willis (2004), which incorporates both convex and non-convex costs of adjustment. One could estimate the labor adjustment cost parameters to match relevant moments from the data simulated from the search model.

Clearly the results reported above suggest that the search model is capable of capturing many but not all aspects of the data. We find evidence that adjustment costs associated with the firing of workers seem important as well. Thus search alone is likely not to be the entire story.

7.4 Idiosyncratic Shocks

An important part of the specification of the model is the presence of producer specific profitability shocks. The standard deviation of the innovation to the idiosyncratic component of profitability is estimated at 0.227. Table 9 shows how some of the key moments respond to a reduction in σ_ε from this estimated value to 0.057. A key lesson from these results is that the moments of the model economy do depend on the distribution of establishment shocks: these shocks are not smoothed out by aggregation.

As is evident from the table, the correlation between unemployment and vacancies becomes less negative for lower values of σ_ε. This is indicative of the fact that parameters describing microeconomic objects can impact aggregate variables in this framework. Further, the distribution of employment adjustment is more condensed as σ_ε falls. This makes sense since there are fewer large draws of the idiosyncratic shock and so, given adjustment costs, less variability in employment growth.

Overall, the results of Table 9 help illustrate the importance of simultaneously matching the micro and the macro moments. To match the micro moments, adjustment costs and a high variance of idiosyncratic shocks must both be present. Moreover, without the high variance of idiosyncratic shocks but with the

[42] See Hamermesh and Pfann (1996) and Nickell (1986) for example.

[43] Matching that feature would require the model to have more of a stochastic matching structure. In our model, there is a deterministic relationship between vacancies and employment at the establishment level.

σ_ε	$corr(u,v)$	$-0.10 > \Delta e$	$-0.10 < \Delta e < -0.025$	$-0.025 < \Delta e < 0.025$
0.2266	-0.9022	0.0519	0.0712	0.7938
0.17	-0.8413	0.0182	0.0585	0.8493
0.1133	-0.6833	0.0002	0.0352	0.8918
0.057	-0.5382	0	0.0007	0.9249

Table 9: Effects of Reducing σ_ε

adjustment costs, we cannot match the macro moments such as the Beveridge curve.

7.4.1 Other Implications

There are a couple of additional facts, used to motivate this study, which were not included in the set of moments. Table 3, for example, includes the hires and separation patterns across growth rate bins showing that the margin for adjustment for expanding (contracting) firms is primarily hires (separations). For the estimated model, we find analogous patterns. For example, in the growth rate bin for establishments expanding by more than 10%, our estimated model implies the hiring rate is 21% and the separation rate is 2%. These estimates compare to the JOLTS estimates of a 30% hiring rate and 4% separation rate for the same growth rate bin from Table 3.

As noted earlier, there is substantial inaction on vacancies and employment growth at the establishment level. In 32% of the (size weighted) observations, net employment growth is zero. Further, 45% of the observations entail zero vacancies created. The best fit model with hiring costs did not fit these facts. It predicts no observations of zero net employment growth and too much inaction in vacancies: there are no vacancies in about 93% of the simulated observations. The best fit model with firing costs has some inaction in employment growth and vacancy inaction in 73% of the observations.

This inability to match the inaction in net employment reflects the underlying assumption of the search model: adjustment costs are associated with gross rather than net hires and quits are deterministic. In our specification, the non-convex cost of posting vacancies leads to zero gross employment growth in about 78% of our observation. This does not translate into inaction in net employment growth because of the exogenous quits.

Table 2 contained moments on employment and hours beyond the moments emphasized in our estimation. The model was estimated to fit the relative standard deviations of hours to employment growth and the correlation between them at the establishment level on a quarterly basis. The estimated model does well matching these moments.

This table shows some additional moments at both the establishment and aggregate levels. In terms of microeconomic observations, the correlation between lagged hours growth and employment is 0.184 in the plant-level data. Using our best estimates, the simulated data shows a slightly negative correlation. This seems to reflect time aggregation. In keeping with data, we report quarterly moments even though our model is solved on a monthly basis. If instead we look at monthly simulated data, the correlation between lagged hours growth and employment rises to 0.43. This correlation reflects our basic intuition about the dynamic interaction between hours and employment growth: the initial response of a producer to a shock is through hours with the employment adjustment (if it occurs) arising later. The problem is that at the quarterly frequency, this relationship is obscured by the lack of persistence of the idiosyncratic shocks and the multiple decision periods between lagged hours and current employment adjustment.

Regarding aggregate implications, Table 2 highlights differences between comovement and variability of hours and employees at the establishment and aggregate levels. Our model does not match these moments well at the quarterly frequency. But, as noted above for the establishment-level moments, if we look at monthly frequency, the correlation between lagged hours growth and and employment growth is 0.396.

Moment	Plant		Aggregate	
	LRD	Est. Model	LRD	Est. Model
$\frac{\sigma_{\Delta h}}{\sigma_{\Delta e}}$	0.96	0.99	0.55	0.93
$Corr(\Delta h, \Delta e)$	-0.296	-0.363	0.545	0.013
$Corr(\Delta h_{-1}, \Delta e)$	0.184	-0.019	0.519	-0.0983

* Seasonal and Aggregate Effects removed from establishment-level moments. Seasonal Effects removed from aggregate-level moments. All moments are at a quarterly frequency.

Table 10: Hours and Employment Adjustment: Basic Facts from the LRD*

Both unemployment and vacancies are serially correlated. Those moments were not used in the estimation. In the JOLTS data, the serial correlation of $\bar{u}$ is 0.92 and of $\bar{v}$ is 0.91. From simulations of the estimated model with hiring costs, we find that the serial correlation of unemployment is 0.96 and the serial correlation of vacancies is 0.81. Thus the model does not quite generate the required serial correlation of vacancies.

8 Extensions

Here we briefly discuss some aspects of the model and empirical approach which deserve additional consideration. The focus is on variants of the model to bring it closer to observations.

8.1 Size Distribution of Establishments

Our model has a rich cross-sectional distribution of profitability. This translates into cross-sectional distributions of a variety of variables. While we have looked as some of these in detail, such as employment growth and hours growth, these moments have all been size weighted, using employment. These same moments can be calculated as simple, not size-weighted, statistics. A comparison of these two calculations is then informative about the size distribution of establishments. Moreover, the frictions (e.g., adjustment costs) likely vary by the size of the firm, and exploring the variation across the size distribution of establishments could be provide further information about these frictions.

8.2 Capital

As with most of the search models, there is no explicit capital in the model. There are three implications of this structure worth noting.

First, the absence of capital makes interpreting α more difficult. One might assume that capital flows freely across production sites so that underlying the revenue function is a capital decision. Or, one might assume that capital is determined at the time the establishment is created. In that case, α would be interpreted as labor's share. Including capital explicitly in the model would make the interpretation as well as the calibration/estimation of α more transparent.

Second, it is natural to think that there are costs of adjusting capital as well. How capital adjustment costs interact with labor demand and how the costs of labor and capital adjustment interact remains unexplored.

Third, workers in the model consumed their compensation. If there was capital, then workers could save.

8.3 Bargaining Power for Workers

The model assumes that producers make workers a take-it-or-leave-it offer so that, in equilibrium, employed and unemployed workers receive the same payoff, $U(b(a))$. This assumption considerably simplifies the analysis since workers need not consider the distribution of employment opportunities in the economy. To do so would require the state space for workers' decisions to include the cross-sectional distribution of establishment profitability shocks. In addition, solving for an equilibrium would require the determination of the market-clearing equilibrium level of expected utility on a state-by-state basis.

Clearly there are gains to tractability through these assumptions. But, the model does imply that unemployed and employed workers have the same levels of expected utility. If state-contingent severance pay was not feasible, the bargaining power of workers could create a gap in utility between employed and unemployed workers. This is a common property of search models.

As discussed in Hagedorn and Manovskii (2006), both the bargaining weight and the utility of unemployed relative to employed workers impact the cyclical properties of the model. This is partially accommodated in the model by allowing $b(a)$.

Understanding the quantitative implications of relaxing the assumptions made on the bargaining process would be a useful extension of this framework. Among other things, this would imply that the cross-sectional distribution of productivity would be part of the state vector and thus another source for richer dynamics.

8.4 Quits

The facts about inaction in employment adjustment and vacancies are revealing about adjustment costs, quits and the process of hiring. If posting vacancies is the only way to hire workers, as in our model, then inaction on vacancies must imply zero new hires. If this is coupled with zero quits, then zero net employment growth arises. But, if there are quits, then the frequency of zero net employment adjustment will be less than the frequency of zero vacancy posting.

The model we study allows for quits that are deterministic at the establishment level. In fact, we have imposed a constant quit rate at each establishment since the model does not have a rich theory of quits given the result that $V^e = V^u$.

In reality, quits are stochastic and endogenous. For small establishments, the randomness in quits is not nullified by the law of large numbers. Thus, on a monthly basis, the quit rate ought to be modeled as stochastic rather than deterministic.

As discussed in Davis, Faberman, and Haltiwanger (2006a), quits are not independent of the employment growth at the establishment level. When employment growth is negative, the quit rate is apparently higher.

Extending the model to include stochastic and endogenous quits may be important for matching observations on the frequency of zero net employment growth, about 30% of observations, and the frequency of zero vacancies, about 45% of observations. With deterministic quits and non-convex adjustment costs, it is not possible to explain inaction in both vacancies and net employment growth. Still, it is not clear how much impact a richer model of quits would have on the aggregate implications of the model.

9 Conclusions

The goal of this paper has been to study the implications of a search model for observed movements in

- unemployment and vacancies at the aggregate level,
- employment and hours variations at the establishment level,
- the distribution of net employment growth at the establishment level.

While each of these aspects of the data have been explored independently in other studies, it is valuable to look jointly at these observations. The micro evidence guides and disciplines the models built to match aggregate observations, and the models based on the establishment level ought to be challenged to match aggregates.

Our framework is an extension of the standard search model with two features. First, in order to match establishment level observations, we introduce non-convexities into the process of posting vacancies along with convex adjustment costs. Second, we introduce hours variations into the search model through an *ex ante* labor contracting structure. Both of these features require us to create a nontrivial model of the producer. Finally, as in Mortensen and Pissarides (1994), establishment specific shocks play a prominent role in the analysis and in the moments generated by the model.

The returns on these novel modeling features accrue from insights into the costs of vacancies and firing and the ability of the model to resolve some of the puzzling aspects of the data. In that regard, we find the following:

- fixed and linear adjustment costs are necessary to match observations,
- the model with hiring cost and the one with firing costs do equally as well matching key moments,
- the model is able to match observed co-movements in hours and employment growth at the establishment level,
- the model does not suffer from the amplification problem highlighted in Shimer (2005), even with a high elasticity of wages to productivity,
- the model does match some of the inaction and bursts reported for vacancies and employment adjustment at the establishment level.

The paper concludes with a list of extensions to consider. These are intended to further shrink the gap between the model and facts about labor markets.

References

Abraham, K., and J. Haltiwanger (1995): "Real Wages and the Business Cycle," *Journal of Economic Literature*, 33(3), 1215–1264.

Acemoglu, D., and W. Hawkins (2006): "Equilibrium Unemployment in a Generalized Search Model," `http://troi.cc.rochester.edu/~whawkin2/papers/acehaw.pdf`.

Azariadis, C. (1975): "Implicit Contracts and Underemployment Equilibria," *Journal of Political Economy*, 83(6), 1183–1202.

Barsky, R., G. Solon, and J. Parker (1994): "Measuring the Cyclicality of Real Wages: How Important is the Composition Bias?," *Quarterly Journal of Economics*, 109, 1–25.

Bils, M. (1987): "Real Wages over the Business Cycle: Evidence from Panel Data," *Journal of Political Economy*, 85, 666–689.

Caballero, R., and E. Engel (1993): "Microeconomic Adjustment Hazards and Aggregate Dynamics," *Quarterly Journal of Economics*, 108, 313–58.

Caballero, R., E. Engel, and J. Haltiwanger (1997): "Aggregate Employment Dynamics: Building From Microeconomic Evidence," *American Economic Review*, 87, 115–37.

Cole, H., and R. Rogerson (1999): "Can the Mortensen-Pissarides Matching Model Match Business Cycle Facts?," *International Economic Review*, 40, 933–60.

Cooper, R., J. Haltiwanger, and J. Willis (2004): "Dynamics of Labor Demand: Evidence From Plant-Level Observations and Aggregate Implications," NBER Working Paper #10297.

Davis, S., R. Faberman, and J. Haltiwanger (2006a): "The Establishment-Level Behavior of Vacancies and Hiring," mimeo.

——— (2006b): "The Flow Approach To Labor Markets: New Data Sources And Micro-Macro Links," *Journal of Economic Perspectives*, 20(3), 3–26.

Faberman, R. J. (2005): "Analyzing the JOLTS Hires and Separations Data," *Proceedings of the 2005 American Statistical Association.*

Fujita, S., and G. Ramey (2005): "The Dynamic Beveridge Curve," `http://www.phil.frb.org/files/wps/2005/WP05-22.pdf`, Federal Reserve Bank of Philadelphia, Working Paper #05-22.

Hagedorn, M., and I. Manovskii (2006): "The cyclical behavior of equilibrium unem- ployment and vacancies revisited," `http://www.econ.upenn.edu/~manovski/papers/BCUV.pdf`, Mimeo, University of Frankfurt and University of Pennsylvania.

HALL, R. (2005): "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95, 53–69.

HAMERMESH, D., AND G. PFANN (1996): "Adjustment Costs in Factor Demand," *Journal of Economic Literature*, 34, 1264–92.

MORTENSEN, D. T., AND C. A. PISSARIDES (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61, 397–415.

NICKELL, S. (1986): "Dynamic Models of Labor Demand," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and R. Layard. Amsterdam: North Holland.

RUDANKO, L. (2006): "Idiosyncratic and Aggregate Risk in a Frictional Labor Market," mimeo.

SHIMER, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25–49.

TASCI, M. (2006): "On the Job Search and Labor Reallocation," `http://www.eco.utexas.edu/graduate/Tasci/ojsandlmr.pdf`.

YASHIV, E. (2000): "The Determinants of Equilibrium Unemployment," *American Economic Review*, 90, 1297–1322.

——— (2006): "Evaluating the Performance of the Search and Matching Model," *European Economic Review*, 50, 909–36.