

**DEPARTMENT OF ECONOMICS
YALE UNIVERSITY**

P.O. Box 208268
New Haven, CT 06520-8268

<http://www.econ.yale.edu/>



Economics Department Working Paper No. 43
Cowles Foundation Discussion Paper No. 1645

**Understanding Sectoral Labor Market Dynamics:
An Equilibrium Analysis of the Oil
and Gas Field Services Industry**

Patrick Kline

March 2008

Understanding Sectoral Labor Market Dynamics: An Equilibrium Analysis of the Oil and Gas Field Services Industry*

Patrick Kline
Yale University
patrick.kline@yale.edu

March 2nd, 2008

Abstract

This paper examines the response of employment and wages in the US oil and gas field services industry to changes in the price of crude petroleum using a time series of quarterly data spanning the period 1972-2002. I find that labor quickly reallocates across sectors in response to price shocks but that substantial wage premia are necessary to induce such reallocation. The timing of these premia is at odds with the predictions of standard models—wage premia emerge quite slowly, peaking only as labor adjustment ends and then slowly dissipating. After considering alternative explanations, I argue that a dynamic market clearing model with sluggish movements in industry wide labor demand is capable of rationalizing these findings. I proceed to structurally estimate the parameters of the model by minimum distance and find that simulated impulse responses match key features of the estimated dynamics. I also provide auxiliary evidence corroborating the implied dynamics of some important unobserved variables. I conclude with a discussion of the strengths and weaknesses of the model and implications for future research.

*I am deeply indebted to Chris House for sharing with me the art of formulating and solving dynamic models. I would also like to thank Joe Altonji, John Bound, Charlie Brown, Matias Busso, Kerwin Charles, Michael Elsby, Yuriy Gorodnichenko, Ben Keys, Justin McCrary, Matthew Shapiro, Gary Solon, and Robert Willis for useful comments and corrections. This paper benefitted greatly from the comments of seminar participants at Harvard, MIT, Yale, Chicago, Princeton, Stanford, Boston University, Brown, Wharton, Northwestern, Berkeley, San Diego, and Duke. All errors are my own.

1 Introduction

Perhaps the central tenet underlying the theory of competitive markets is the notion that a freely traded commodity will command a single price. In the study of labor markets, this assumption, the so-called “law of one price,” plays an important role in simplifying theories of wage determination, allowing labor economists to abstract from the complex sectoral composition of the economy. The law of one price is typically invoked under the assumption that, in the long run, any sectoral disturbances leading to wage differentials will eventually be arbitrated away by reallocation across sectors. In textbook models this argument is represented by the assumption that the supply of labor to a sector is, in the long run, infinitely elastic, which implies that any sector-specific disturbances to supply and demand will yield wage effects that are ultimately transitory.

A litany of policy implications follow from the standard model. For example, the incidence of sector-specific commodity taxes or tariffs should not, in the long run, fall upon labor in that sector. Similarly, changes in the relative productivity of a sector should not permanently affect a sector’s relative wages. In general, labor markets are thought sufficiently flexible that, given any disturbance, the welfare of workers in a sector will only be affected for as long as it takes labor to flow across sectors and restore equilibrium. It is surprising then that little attention has been devoted to determining how long this adjustment process takes, since, as Keynes famously quipped, “in the long run we are all dead.”¹

The importance of understanding the dynamics governing labor market adjustment extends far beyond welfare economics to the core of many fields of research in positive economics. Much has been made of the supposition that shifts in the sectoral composition of demand are capable of lowering aggregate output via costly reallocation of capital and labor.² In a series of influential papers, Hamilton (1983, 1988, 2003) has argued that major oil shocks may have caused recessions through such a mechanism. Most of the empirical work in this literature has focused on directly assessing the relationship between oil shocks and output, rather than the links mediating this hypothesized relationship.³ The little existing work that does examine the labor market response to oil shocks either ignores wages (Davis and Haltiwanger, 2001) or relies upon relatively short panels incapable of identifying detailed dynamic responses to shocks (Keane and Prasad, 1995).

¹There is however an extensive literature examining the adjustment of local labor markets to shocks (Bartik, 1991; Blanchard and Katz, 1992; Bound and Holzer, 2000; Topel, 1986).

²See, for example, Lilien (1983), Abraham and Katz (1986), Davis (1987), Brainard and Cutler (1993), and Ramey and Shapiro (1997).

³Bresnahan and Ramey’s (1993) study of the U.S. automobile industry is a notable exception.

This paper examines the dynamic response of wages and employment in the U.S. Oil and Gas Field Services (OGFS) industry to changes in the price of crude petroleum using quarterly data from 1972 to 2002.⁴ The oil industry provides an important case study for a number of reasons. First, given the debate over the potential macroeconomic effects of oil, it is of interest to examine the allocative effects of oil price changes on the industry to which it is most directly tied. Second, OGFS is a non-unionized high-turnover industry, requiring little formal training for production workers. In this sense, it approximates a neoclassical spot market for labor. To the extent that important adjustment rigidities are found in this market, they are likely to represent a lower bound on the sort of rigidities found in more specialized labor markets with important training requirements and durable employment relationships. Finally, the immense changes in the price of crude petroleum over the time period in question provide ample exogenous variation in labor demand with which to examine the performance of standard models of adjustment. The fact that oil prices are well measured, volatile, and difficult to forecast makes them ideal for investigating labor market dynamics since they provide the rare opportunity to trace how well-defined demand shocks propagate throughout a labor market at high frequencies. If dynamic market clearing models are to have any empirical content, they must be capable of explaining the basic stylized facts uncovered in this analysis.

Using a simple econometric specification, I find that labor quickly reallocates across sectors in response to price shocks, but that substantial wage premia are necessary to induce such reallocation. Surprisingly, these wage premia emerge quite slowly, peaking only as labor adjustment ends and then slowly dissipating. This profile of wage effects is inconsistent with traditional market clearing models which predict that wages should jump on news of a price change only to be dissipated away by flows of workers into the sector. Indeed, traditional models view wage premia as signals of market scarcity which serve to direct workers towards one sector or another. But in this market, it appears that wage premia actually lag labor flows, seemingly calling into question the allocative role of market wages.

After considering and discarding stories involving contracting and composition bias, I argue that a dynamic market clearing model with sluggish movements in industry-wide labor demand is, in fact, capable of rationalizing the joint response of industry employment and wages to oil price shocks. The key insight is that forward looking workers will use information over and above current wages to make sectoral choice decisions. In such an environment increases in the current price of oil will signal future increases in the demand

⁴Perhaps the closest analogue to this work is by Carrington (1996), who examines the response of the Alaskan economy to the construction of the Trans-Alaska Pipeline by a single firm.

for oil workers which may lead workers to flow into the sector in anticipation of future wage premia thereby depressing current wages.⁵

To assess the quantitative plausibility of this story, I structurally estimate the parameters of the model by minimum distance and find that reasonable parameter estimates yield simulated impulse responses that recreate key features of the estimated dynamics. I also show that the model yields accurate predictions about the evolution of conventional measures of labor market tightness in response to demand shocks.

The contribution of the paper is threefold. First, I provide the most credible evidence to date on the dynamic effects of well-measured demand shocks on the equilibrium behavior of a relatively homogeneous and competitive sectoral labor market of the sort described in undergraduate textbooks. The analysis reveals that even in a very flexible labor market the long run is quite long (approximately seven years) suggesting that sectoral shocks may have protracted effects on the welfare and decisionmaking of workers and firms. I also document that wage premia substantially lag employment adjustment at quarterly frequencies, a fact one would have difficulty uncovering or believing without observable exogenous demand shifters. Second, I build a formal dynamic general equilibrium model illustrating that market clearing behavior is qualitatively consistent with the sort of dynamics uncovered in the empirical work once one allows for forward looking behavior and adjustment rigidities in labor demand. Finally, by structurally estimating the model, I show that the parameters necessary to quantitatively match the dynamics are quite reasonable. I also show that these parameter estimates yield accurate predictions about other moments not used to fit the model.

The next section provides an overview of the Oil and Gas Field Services Industry. Section 3 describes the data used in the analysis, while Section 4 describes the reduced form empirical results. Section 5 considers various explanations of the estimated dynamics. Section 6 lays out a dynamic market clearing model of sectoral reallocation. Section 7 describes structural estimation of the model and discusses the parameter estimates and simulated impulse responses. Section 8 concludes with a discussion of the generalizability of the results and implications for future work.

⁵This effect is similar to that discussed by Topel (1986) in his model of regional reallocation, though Topel did not have dynamics on the labor demand side.

2 A Brief Overview of the U.S. Oil and Gas Field Services Industry

The Oil and Gas Field Services industry (SIC 138) performs drilling, exploration, and maintenance services on a contract basis for large oil companies. Over the period 1968-2002, oil and gas field services employed, on average, 65% of the production workers in the larger oil and gas extraction industry (SIC 13).⁶ The main distinction in the industry is between exploration and extraction. Using increasingly sophisticated methods, small crews of specialized workers search for geological formations likely to contain oil or gas. Upon discovery of oil, an oil company will install a steel structure known as a derrick to support the drilling equipment and dig for oil. If the site is off shore, the company will install a floating rig to support the drilling operation.

The bulk of OGFS employment is in extraction. According to the 1992 Economic Census drilling and maintenance activities account for approximately 90% of total production worker employment in the industry. Tasks undertaken by maintenance crews (the largest group) include excavating slush pits and cellars, building foundations at well locations, surveying wells, running, cutting, and pulling casings tubes, cementing wells, shooting wells, perforating well casings, acidizing and chemically treating wells, and cleaning out, bailing, and swabbing wells. These tasks involve some skill but are primarily manual in nature. Little formal education is required and most training occurs on the job.

The industry employs a variety of occupations in many different work environments. While roustabouts and construction workers perform physical tasks in rugged outdoor environments, there are a number of executives and clerical workers whose work is performed indoors. Geologists, petroleum engineers, and managers frequently split their time between the office and the field.

According to the 1999 Occupational Employment Statistics published by the Bureau of Labor Statistics, “Construction and Extraction Occupations” constitute the largest occupational group representing 41% of total OGFS employment. Among these workers the most common occupation is the roustabout—a handyman who repairs equipment and performs generalized physical tasks.⁷ Petroleum engineers, while common in the broader oil extraction industry, constitute less than 1% of employment in OGFS. Engineering occupations in

⁶Oil and gas field services encompasses the same tasks (e.g. drilling, exploration, and well maintenance) as the general oil and gas extraction industry. The distinction is that the work in SIC 138 is done on a contract basis usually for large oil companies.

⁷Roustabouts represent 10.4% of total employment in OGFS and a quarter of employment among the Construction and Extraction Occupations which are likely to constitute the bulk of production workers.

general are uncommon, making up less than 5% of total employment. Finally, “Office and Administrative Support Occupations” make up around 10% of total employment.

Employment in the U.S. oil industry is concentrated in a few states. In decreasing order of importance they are: Texas, Louisiana, Oklahoma, and California. Over the 1968-2002 period approximately 40% of industry employment was in Texas. The industry is not unionized. Only 4% of workers in the broader oil extraction industry report being union members in the March Current Population Survey. The OGFS industry consists primarily of small and medium size firms. Most maintenance workers are employed in firms with less than 50 employees and most workers in the drilling subindustry work in firms with less than 250 employees.

2.1 The Market for Oil and Gas

The U.S. oil and gas industries are regulated by state agencies, the most important of which is the Texas Railroad Commission (TRC). Since 1919, the TRC has had the authority to set allowable oil and gas production levels and to grant drilling rights. The mission of TRC’s Oil and Gas Division is to “prevent waste of the state’s natural resources” but its practical role has been to stabilize the price of oil by adjusting supply in response to projected changes in demand. Prior to 1972, allowable production levels were set such that the U.S. oil industry operated substantially below capacity, ensuring a high but stable price. By April 1972, demand outstripped supply and the industry began operating at full capacity, effectively ending the rationing of oil. Also by this time net imports of oil and gas had risen to 27.6% of total production signifying an important dependence on foreign oil supplies.

On October 17, 1973 the OPEC embargo was announced. With domestic suppliers already operating at peak capacity prices rose dramatically. From this point on, international fluctuations in supply and demand became the primary drivers of the price of oil and gas.⁸ While there is debate about whether the proximate causes of oil price shocks have been geopolitical events or shifts in global demand,⁹ it seems clear that oil prices are not being driven by idiosyncratic labor supply shocks to the U.S. oil industry. Accordingly, variation in the price of crude oil provides an ideal opportunity to examine the response of a well defined industrial labor market to exogenous changes in output price and consequently labor demand.

⁸See Hamilton (1985) for an excellent overview of the historical determinants of oil prices.

⁹See Barsky and Kilian (2002, 2004) and Hamilton (2001).

3 Data

I measure employment and wages in the oil extraction industry using the Current Employment Survey (CES), which is a monthly survey of establishments conducted by the Bureau of Labor Statistics (BLS). The CES reports information on the total hours worked, average weekly earnings, and average weekly hours of production workers as reported by employers. The earnings concept includes overtime pay and bonuses while the hours variable includes overtime and sick days. Average hourly wages are measured as average weekly earnings divided by average weekly hours.¹⁰ The CES employment data are benchmarked annually to the ES-202 series, which contains information on earnings and employment for all establishments covered by unemployment insurance laws. The primary advantages of the CES data over the ES-202 data are that they contain information on hours worked and are publicly available over a longer period of time.¹¹

Because oil shocks influence inflation and other macroeconomic variables, it is useful to focus on wages relative to the outside world rather than nominal wages. For this reason I use the nonmetallic mining industry (SIC 14) as a control group in order to filter out macroeconomic disturbances. Workers in nonmetallic mining perform tasks similar to oil workers, have similar skills, and work in roughly comparable physical environments.¹² Relative wages are measured as the difference in log average hourly wages between the OGFS and nonmetallic mining industries. Prior to the large disruptions in the price of oil beginning in 1972, wages in OGFS and nonmetallic mining are nearly identical. Since wages reflect the quality of the workers employed in an industry, this finding reinforces the notion that workers in the two industries are comparable. There are no large shifts in employment in nonmetallic mining during the sample period, and, because it employs similar workers, the nonmetallic mining industry is likely to share much of the secular variation in labor supply conditions experienced by the oil industry.

Monthly data on oil prices are from the Producer Price Index series for Crude Petroleum.¹³ This series corresponds closely to annual data from the Department of Energy on the domestic first purchase price of crude oil. I deflate the price data by the CPI-U since

¹⁰Note that this implicitly weights average wages by hours worked.

¹¹The ES-202 series is publicly available back to 1975.

¹²Workers in this industry extract sand, stone, granite and other minerals from quarries. The tasks performed by production workers are remarkably similar to those in OGFS as they involve drilling, transporting, and processing raw materials. Unlike coal and metal mining, SIC 14 employment is not geographically concentrated as most states have deposits of stone, clay, and sand. Employment is highest in California, Texas, and Georgia. Like in the OGFS industry most workers are employed in medium and small firms.

¹³BLS Commodity Series WPU0561

firms should be interested in maximizing real profits. Because oil prices are not good indicators of domestic demand during the period of regulation by the TRC, I start my analysis with data from the first quarter of 1972. The analysis is conducted at quarterly frequencies. I use the middle month of each quarter in constructing each series.

4 Empirical Dynamics

Figure 1 shows log employment and relative wages in the oil industry versus log oil prices. The oil price series seems well approximated by a pure random walk. Dickey Fuller GLS tests (Elliott et al., 1996) cannot reject the null that the series contains a unit root against the alternative that it is mean stationary at the 10% level.¹⁴ Moreover, it is not possible to reject the null that price changes are mean zero white noise.¹⁵ This implies that, to first order, oil prices, at least at the quarterly frequencies examined, are a martingale, exhibiting no forecastable short or medium run dynamics. Employment and wages also appear to contain a unit root component, but they all have important short run dynamics as well.¹⁶ This differential property of the time series of oil prices is notable, for any theory purporting to explain the dynamics of wages and employment as a function of oil price shocks must generate short run dynamics on its own.

Clearly, the three series in Figure 1 track each other very closely.¹⁷ Note that employment seems responsive both to price increases and decreases, suggesting that the oil industry is able to end employment relationships fairly quickly.¹⁸ The relative wage series is roughly centered around zero, as we would expect given a good measure of the outside wage, but contains a noticeable downward trend which I correct for in the regressions to come. Like the employment series, relative wages are strongly correlated with oil prices. Restricting attention to the massive price buildup from 1972-1981 we see that employment increased by approximately 370%, while relative wages (after detrending) increased by approximately 18%. Thus, if we were to interpret this behavior as representing shifts along a stable supply curve, we would get a back-of-the-envelope elasticity of about 20.

¹⁴Unit root tests are conducted using the lag length selection procedure of Ng and Perron (2001).

¹⁵A portmanteau Q test using 40 lags has a p-value of .6. The lowest p-value across all lags between 1 and 40 is .094. The mean log price change over the sample period is on the order of 10^{-7} .

¹⁶DFGLS tests cannot reject the presence of a unit root against the alternative of trend stationarity at the 10% level for either variable, but portmanteau tests easily reject that the changes in the series are white noise at the 1% level.

¹⁷The correlation between the employment and oil price variables is .78, while the correlation between relative wages and oil prices is .6 after correcting for a linear trend.

¹⁸It could also be that many firms go out of business during this period. I lack the data necessary to distinguish between the two hypotheses.

Though the figures reveal an obvious long run relationship between log employment, relative wages, and oil prices, it is of interest to investigate the dynamics of the relationship between these variables more carefully. I consider simple distributed lag specifications of the form:¹⁹

$$y_t = \mu + \sum_{k=0}^{24} \beta_k p_{t-k} + \delta t + \theta_{qt} + \varepsilon_t \quad (1)$$

where y_t denotes the outcome of interest (log employment, relative wages, or log hours), the p_{t-k} are lags of log oil prices, t is a linear time trend, the θ_{qt} are quarter effects, and ε_t is a serially correlated error term. Assuming that oil prices are exogenous, the β_k coefficients give the effect of a 1 unit change in log oil prices k periods in the past. The specification assumes that the adjustment process concludes after 6 years.²⁰

In light of the aforementioned persistence of oil prices, it is more informative to estimate the dynamic response of the y_t 's to a permanent change in prices since that is the sort of shock the oil industry seems to be faced with in practice. Denote the partial sum of the distributed lag coefficients by the symbol $\pi_k = \sum_{j=0}^k \beta_j$. The π_k 's give the effect of a permanent 1 unit change in log oil prices after k periods. We can reparameterize the above equation to estimate the π_k 's directly:

$$y_t = \mu + \sum_{k=0}^{23} \pi_k \Delta p_{t-k} + \pi_{24} p_{t-24} + \delta t + \theta_{qt} + \varepsilon_t \quad (2)$$

Since oil prices are non-stationary and the error term is serially correlated, I estimate equation (2) in first differences and use Newey-West standard errors for inference.

Figure 2 plots the estimated π_k 's for employment, relative wages, and hours along with 95% confidence intervals. The estimated instantaneous effect of a permanent 10% increase in oil prices is a 1.5% increase in employment. This instantaneous effect is followed by approximately four more quarters of employment increases after which time hiring slows down and employment levels out at a new equilibrium approximately 7% higher than the old steady state. Unlike employment, wages do not respond instantaneously to oil price changes. In fact, the point estimate for the instantaneous effect of oil prices on wages is negative. Wages grow slowly over the next six quarters, after which they plateau at a peak of approximately 1% above steady state. They remain at this level for approximately two

¹⁹Appropriately parameterized vector autoregressive and autoregressive distributed lag specifications yield virtually identical results.

²⁰The results are robust to the inclusion of additional terms.

years and then slowly begin to fall back to parity with the outside world. Hours per worker jump immediately by approximately 1% in response to price shocks and then slowly decline.

The estimates in Figure 2 are somewhat heavily parameterized. Figure 3 constrains the π_k 's to lie on a 5th order polynomial. This reduces the standard errors somewhat and eases visual interpretation of the results. The same pattern emerges. Wages move slowly, rising only as hiring slows down and dissipating only some time after the industry labor market has stopped growing. Hours jump immediately and remain high until employment adjustment is complete. These regularities constitute the set of stylized facts that the next section seeks to explain.

5 Discussion

The dynamics estimated thus far are puzzling for conventional models of wage determination. Given that the marginal revenue product of workers moves instantaneously in response to price changes, one would expect wages to also move instantaneously when labor is supplied inelastically to the sector. These wage premia should dissipate slowly as workers arbitrage wage differentials across sectors until the marginal revenue product of labor is equalized across industries. Once employment adjustment is complete, wages should have returned to steady state. The behavior is all the more puzzling given that hours respond to shocks immediately and wage data should reflect overtime payments.

Several alternative rationalizations of the facts seem possible. An obvious one is that contracts prevent firms from moving wages rapidly. However there is good reason to believe that contracting is not the culprit in this industry. As mentioned previously, the industry is not unionized. As Table 1 shows, monthly separation and accession rates in the oil and gas field services industry are both on the order of 10% in the years for which data is available. Note that this rate is substantially higher than in the crude petroleum, natural gas, and natural gas liquids industry which was also witnessing a dramatic expansion during this time period. This is attributable to the different skill and occupation composition of the two industries. As their title suggests, roustabouts, the backbone of the oil and gas field services industry, do not enter their profession in order to hold down stable jobs. With turnover rates of this magnitude, formal contracting is likely to be costly and inefficient. And though workers and firms would like to share risk, there is little chance for implicit contracts to emerge in an environment where the employment relationship is so likely to be short.

One might still suspect that firms are unable to adjust wages in the short run for other reasons, such as administrative costs. However, inspection of the raw time series of relative

wages in Figure 1 indicates that wages often do adjust very quickly. For example, starting in the first quarter of 1983, after employment had reached its peak, wages in the oil industry began to plummet precipitously. Likewise wages were able to spike immediately after the onset of the first OPEC crisis in the last quarter of 1973. Thus, wage rigidity of the conventional sort does not seem to provide a satisfying explanation of the patterns in the data.

Furthermore, data from the Quarterly Census of Employment and Wages indicate that the number of OGFS establishments increased by 250% between 1975 and 1980 and then fell by a quarter over the next 5 years. If a substantial part of the changes in industry employment involve the entry or exit of firms (or even establishments), it seems unlikely that contracts or administrative costs would be capable of preventing wages from adjusting.

Another explanation might involve composition bias. If lower quality workers are hired in times of high demand this might depress the observed wage even though the real wage has increased. There are two reasons to suspect this is not what is going on. First, this would require very large short run hiring elasticities. Say for example that new hires are only half as productive as experienced workers and that this is reflected in their wages.²¹ Then we can write mean observed wages as:

$$w = [s + 2(1 - s)] w_0$$

where w is the observed average wage, s is the share of workers that are new and w_0 is the wage of an inexperienced worker. Logarithmically differentiating this equation with respect to oil prices yields:

$$\tilde{w} = \tilde{w}_0 - \frac{s}{2 - s} \tilde{s}$$

where variables with a tilde above them are elasticities. Thus observed wage elasticities equal real wage elasticities minus a component due to increases in the fraction of inexperienced workers. The magnitude of this latter component is increasing in the fraction of workers who are inexperienced. To fix things, say, in keeping with the data in Table 1, that $s = .25$ and $\tilde{w}_0 = .5$. Then in order for wages not to move \tilde{s} would need to equal 3.5 – i.e. a 1% increase in oil prices would need to result in a 3.5% increase in the employment share of inexperienced workers.

To get a sense of the magnitude of this number, write $s = \frac{I}{T}$ where I is the number of inexperienced workers and T is the total number of workers. If we assume that no experienced

²¹This could be expected to occur if new workers require a period of on-the-job training.

workers leave in times of hiring and that no inexperienced workers become experienced,²² we get the following relationship:

$$\tilde{s} = \left(\frac{1-s}{s} \right) \tilde{T}$$

The regression estimates indicated that a 10% increase in price yields an instantaneous 2% increase in employment ($\tilde{T} = .2$). Hence with $s = .25$ we would only expect $\tilde{s} = .6$, far below the level necessary to prevent wages from moving in this example.

Second, even if one thought that composition biases were large enough to prevent wages from moving instantaneously, it would still be difficult to rationalize the rest of the dynamics found in the previous section. If hiring slows down rapidly and workers only require one period of training then wages could rise slowly in subsequent periods as the fraction of new workers fall. But why should wages remain high for several quarters after hiring has slowed down and industry employment has reached a new steady state? By this time adjustment should be complete and w_0 should have returned to steady state, implying that, if anything, we should expect composition biases to yield a w slightly below steady state after a period of expansion.

5.1 A Forward Looking Alternative

Suppose that potential oil workers are aware of the statistical relationship between oil prices and wages. In such a case workers may be willing to switch into the industry when oil prices increase based upon expectations of future wage increases even if current wages do not move. This shift in the sectoral labor supply curve could in turn put sufficient downward pressure on wages to prevent them from rising in the immediate wake of an oil price increase.

Would such behavior constitute an equilibrium? In the absence of demand side frictions, it would not, for in such a case wages must rise on impact if they are to rise at all. However, if the adjustment of labor demand is sluggish, large preemptive shifts in labor supply may temporarily outweigh the contemporaneous shift in labor demand keeping wages low despite rapid rates of hiring. Such a case is shown in Figure 4. As adjustment continues, however, the number of workers available to work in the industry is drawn down and demand begins to outstrip supply over the medium run, leading wages to eventually rise. But such premia cannot persist indefinitely. Because the outside world is large relative to the oil industry,

²²Relaxing these assumptions will only reinforce the conclusion that composition bias is incapable of explaining the results.

long run labor supply to the sector is highly elastic. Thus, in the long run, any wage premia will eventually be arbitrated away.

Two features of this story are worth pointing out here. First, the usual dichotomy between supply and demand shifters has broken down. Innovations to oil prices shift the contemporaneous labor supply curve *because* they slowly shift the labor demand curve. Workers need not be aware of the manner in which demand moves, only the resulting reduced form statistical relationship between wages and prices.

Second, the belief by agents that oil price innovations will result in future wage premia is part of why the premium is delayed. Were agents totally myopic or ignorant of the relationship between oil prices and wages, the supply curve would not shift out on impact, and wages would inevitably rise. Thus, the beliefs about delayed compensation are self-confirming.

One naturally wonders whether such a story could be quantitatively plausible. How large of a future premium would agents need to expect in order for wages not to move on impact? How predictable would demand need to be? In the next section I lay out a dynamic structural model of sectoral reallocation that can be used to help answer these questions.

6 An Equilibrium Model of Sectoral Reallocation

The previous section argued that the slow response of wages and quick response of employment in the oil industry to oil price shocks may be the result of rational forward looking behavior on the part of workers and adjustment rigidities on the part of firms. This section formalizes a dynamic market clearing model of sectoral choice in the spirit of Lucas and Prescott (1974) capable of recreating dynamics of the sort discussed previously.²³

Workers can be in one of two sectors: the oil industry (O) or a nearby sector (N). Let the symbols L_t^o and L_t^n represent the number of workers in a given period employed in the oil industry and the “nearby” sector respectively. We can think of the nearby sector as a reduced form for search behavior. It is the number of workers considering entering the oil industry in the next period. The size of this group will adjust based upon how attractive sector O is at any given time relative to the rest of the economy.

²³Other multi-sector labor market models in this tradition include Rogerson (1987), Chan (1996), and Phelan and Trejos (2000).

6.1 Labor Supply

Workers have utility that is linear in wages and a random taste shock ε_t^s which varies across sectors and over time. Migrating between sectors is costly with cost equal to the "distance" (d) between sectors. Each period workers in the two sectors observe the price of oil (P_t), their draw of the taste shocks ($\varepsilon_{it}^o, \varepsilon_{it}^n$), the number of workers in each sector at the beginning of the period (L_{t-1}^o, L_{t-1}^n), and the current period's wage in the oil industry (w_t^o). I refer to this composite information set as $\Omega_{it} = \{L_{t-1}^o, L_{t-1}^n, w_t^o, P_t, \varepsilon_{it}^o, \varepsilon_{it}^n\}$. With this information workers make migration decisions and work for the remainder of the period. The Bellman equations for workers in sectors O and N are:

$$\begin{aligned} V_t^o(\Omega_t) &= w_t^o + \beta E_t \max \{V^o(\Omega_{t+1}) + \varepsilon_{it+1}^o, V^n(\Omega_{t+1}) - d + \varepsilon_{it+1}^n\} \\ V_t^n(\Omega_t) &= w_t^n + \beta E_t \max \{V^o(\Omega_{t+1}) - d + \varepsilon_{it+1}^o, V^n(\Omega_{t+1}) + \varepsilon_{it+1}^n\} \end{aligned}$$

where (w_t^o, w_t^n) are the wages paid to workers in sectors O and N respectively and d is the distance between sectors. For simplicity, I assume that w_t^n is a constant w^n representing the flow payoff to search.²⁴

The alternative specific taste shocks ($\varepsilon_{it}^o, \varepsilon_{it}^n$) are meant to represent random fluctuations in the utility of employment in the two sectors.²⁵ Examples include random beginnings or ends of romantic relationships, shifts in tastes, the expiration of a lease or contract, or the death or relocation of friends and relatives. Define $\varepsilon_{it}^s = \sigma v_{it}^s$ where σ is a scale parameter reflecting the variance of the underlying idiosyncratic shocks. I assume that the v_{it}^j are independently and identically distributed according to a standard Type I Extreme Value distribution.²⁶

Define $\tilde{V}_t^{s,S}$ as the value to a worker starting in sector s of choosing to move to sector $S \in \{s, s'\}$ in period t minus the sectoral taste shock ε_{it+1}^s so that a generic worker's Bellman

²⁴Endogenizing w_t^o so that it varies over time does not qualitatively change the results.

²⁵Permanent differences in mobility costs could easily be incorporated into this framework by introducing heterogeneity across workers in the distances between sectors. Such additions are not necessary for my purposes.

²⁶A more realistic model would allow the taste shocks to be serially correlated or even to have different intercepts across workers. Both of these extensions would make the short run labor supply response of workers to expected wage gaps dependent upon additional state variables capturing either the distribution of initial conditions of workers or the distribution of permanent tastes across sectors. While interesting, these extensions are difficult to implement for they require keeping track of additional states. Moreover, they are likely to be difficult to distinguish from more basic labor supply parameters such as σ which characterize the general sensitivity of migration decisions to expected wage differentials.

equation may be written as:

$$V_t^s = w_t^s + \beta E_t \max \left\{ \tilde{V}_{t+1}^{s,s} + \varepsilon_{it+1}^s, \tilde{V}_{t+1}^{s,s'} + \varepsilon_{it+1}^{s'} \right\}$$

Note that the $\tilde{V}_t^{s,S}$ do not contain an i subscript because the expected continuation values are common across workers. Using the above notation we can find the probability in any period that a worker switches sectors. A worker in, for example, sector N will choose to move to sector O if and only if

$$\tilde{V}_t^{n,o} + \varepsilon_{it}^o > \tilde{V}_t^{n,n} + \varepsilon_{it}^n$$

which is equivalent to the following condition:

$$v_{it}^o - v_{it}^n > -(V^o(\Omega_t) - V^n(\Omega_t) - d) / \sigma$$

Define a selection variable $D_t^{s,S} = 1$ if a worker moves from sector s to sector S in period t . Then given our distributional assumptions on v it follows from standard results (e.g. McFadden, 1978) that the probability of switching sectors can be written:

$$p_t^{s,S} \equiv P(D_t^{s,S} = 1) = \frac{e^{\tilde{V}_t^{s,S}/\sigma}}{e^{\tilde{V}_t^{s,s}/\sigma} + e^{\tilde{V}_t^{s,s'}/\sigma}} \quad (3)$$

Note that the current migration probabilities implicitly depend upon the equilibrium distribution of continuation values $V^o(\Omega_{t+1})$ and $V^n(\Omega_{t+1})$. Making use of the properties of Extreme Value distributions documented in McFadden (2003) we can simplify the expression for these values by means of iterated expectations, integrating out the component due to the taste shocks. Note that

$$\begin{aligned} & E_t \max \left\{ \tilde{V}_{t+1}^{s,s} + \varepsilon_{it+1}^s, \tilde{V}_{t+1}^{s,s'} + \varepsilon_{it+1}^{s'} \right\} \\ &= \sigma E_t \max \left\{ \tilde{V}_{t+1}^{s,s}/\sigma + v_{it+1}^s, \tilde{V}_{t+1}^{s,s'}/\sigma + v_{it+1}^{s'} \right\} \\ &= \sigma E_t \left[E_{v_{t+1}} \max \left\{ \tilde{V}_{t+1}^{s,s}/\sigma + v_{it+1}^s, \tilde{V}_{t+1}^{s,s'}/\sigma + v_{it+1}^{s'} \right\} \right] \\ &= \sigma E_t \left[\gamma + \ln \left(\exp \left(\tilde{V}_{t+1}^{s,s}/\sigma \right) + \exp \left(\tilde{V}_{t+1}^{s,s'}/\sigma \right) \right) \right] \end{aligned}$$

where $\gamma \approx .5772$ is Euler's constant and $E_{v_{t+1}}$ denotes the expectation next period with respect to the fundamental taste shocks (v_{it+1}^o, v_{it+1}^n). Substituting the above expression into the original Bellman equations yields value functions that are recursive and differentiable in

their future values:

$$V_t^o(\Omega_t) = w_t^o + \beta\sigma E_t [\gamma + \ln (\exp (V_{t+1}^o/\sigma) + \exp ((V_{t+1}^n - d) / \sigma))] \quad (4)$$

$$V_t^n(\Omega_t) = w_t^n + \beta\sigma E_t [\gamma + \ln (\exp (V_{t+1}^n/\sigma) + \exp ((V_{t+1}^o - d) / \sigma))] \quad (5)$$

Together, (3), (4), and (5) characterize the evolution of the migration probabilities $p^{s,S}$ in response to beliefs about the future path of wages. These in turn map into aggregate migration flows by means of the following identity:

$$m_t^{s,S} = p_t^{s,S} L_{t-1}^s \quad (6)$$

With expressions for the flow of workers between sectors O and N we turn now to the laws of motion characterizing the evolution of sector sizes:

$$L_t^o = L_{t-1}^o + m_{n,o} - m_{o,n} \quad (7)$$

$$L_t^n = L_{t-1}^n + m_{o,n} - m_{n,o} + F_t$$

The term F_t represents net flows to sector N from the outside world. Were this economy closed, so that $F_t = 0$, permanent wage premia would be necessary to sustain increases in sector size. Indeed, it is easy to show that in a closed economy the larger sector must pay a higher wage since in any period only half of the workers in the economy will prefer to be employed in one sector or another based solely upon their taste shocks.

This is a rather undesirable property of a model of the market for oil workers for two reasons. First, the empirical evidence illustrated in Figures 2 and 3 suggests that wages do eventually asymptote back towards parity with the outside world after expansions. Second, it is hard to believe *a priori* that large wage premia could persist indefinitely in a market consisting of jobs with very low requirements for entry. Although tastes for employment in a sector vary across individuals, the number of people capable of eventually working in this industry is extremely large. As working in the oil industry becomes more lucrative, the profile of the sector will rise and the number of workers considering entering the sector will grow. This will put downward pressure on the compensating differential required for any given level of employment.

To model the feedback from wages to the visibility of a sector, I assume that net flows from the outside world to sector N are proportional to the size of sector N and the proportional

deviation of the oil wage w_t^o from its steady state value \bar{w} .²⁷ I also allow for some sluggishness in the response to represent lags in the diffusion of knowledge about market conditions. This is formalized by the following autoregressive specification of flows:

$$F_t = \rho F_{t-1} + k L_{t-1}^n (\ln(w_t^o) - \ln(\bar{w})) \quad (8)$$

The parameter k controls the rate at which additional searchers enter sector N in response to increases in the oil wage, while ρ controls how quickly flows into N ratchet up or down if the wage premium remains fixed. As long as k is positive, the steady state value of w_t^o will have to equal \bar{w} , for otherwise sector N would grow or shrink without bound.

To simplify matters, I choose $\bar{w} = w^n$ so that wages are equal across sectors in steady state.²⁸ This choice makes the (nonstochastic) steady state particularly simple. Equality of wages implies equality of sector values since V^o and V^n are merely symmetric nonlinear functions of future wages in the two sectors. It is easy to see from (3) that $V^o = V^n$ implies that the migration probabilities $p^{o,n}$ and $p^{n,o}$ must be equal, which in turn implies that the sector sizes L^o and L^n must also be equal for otherwise the gross flows between sectors would not be perfectly offsetting. Thus, the steady state of the labor market is perfectly symmetric in all of the state variables.

Moreover, several variables exhibit steady states independent of the long run value of P . Increases in P will yield larger sector sizes. But as long as steady state labor demand is downward sloping, oil wages, and V^o and V^n along with them, will be independent of steady state oil prices. Thus, in keeping with the static intuition of elastic long run labor supply, (8) imposes a long run restriction that the impact of disturbances on oil wages will ultimately be transitory.

6.2 Labor Demand

We turn now to specifying the demand side of the model. Firms are price takers on the input market. Because in such an environment wages are determined only by industry-wide demand, I will not attempt to model the microeconomic details of oil production nor the attendant heterogeneity across firms in productivity, resources, or stocks of labor and

²⁷One could also allow the feedback to be proportional to V_t^o or V_t^n . These setups all yield essentially similar results. Unsurprisingly, when F is made proportional to either continuation value higher values of ρ are necessary to explain the behavior found in the distributed lags, for these forward looking values jump on news of a price shock.

²⁸Put differently, I choose $w^o = w^c$ and solve for V^o and V^n using equations (4) and (5). It follows directly from the symmetry of these equations that for such a choice $V^o = V^n = \bar{V}$.

capital. The key idea for the current discussion is that sectoral labor demand should respond sluggishly to shocks. This could be accomplished by means of capital adjustment costs, employment adjustment costs, gross hiring costs, risk aversion, learning, or any other number of familiar stories. In this paper I focus on employment adjustment costs because they are simple to model and likely to be important in this industry.²⁹

I use a standard representative firm framework to capture the behavior of industry-wide movements in the demand for oil production workers.³⁰ The firm produces output using a production technology with quadratic labor adjustment costs.³¹ The profit function is given by

$$\Pi(\Omega_t) = P_t A F(L_t) - w_t L_t - \frac{\chi}{2} (L_t^o - L_{t-1}^o)^2 + \beta E_t \Pi(\Omega_{t+1})$$

where χ is a parameter governing the cost of adjusting the size of the firm's workforce and A is a scale parameter.

The first order condition for employment is

$$w_t = P_t A F'(L_t) - \chi (L_t^o - L_{t-1}^o) + \beta \chi E_t [L_{t+1}^o - L_t^o]$$

A useful alternative representation of this labor demand curve is the following

$$L_t^o - L_{t-1}^o = \frac{1}{\chi} \lambda_t \tag{9}$$

$$\lambda_t = P_t A F'(L_t) - w_t + \beta E_t [\lambda_{t+1}]$$

In words, the desired change in employment is proportional to the discounted stream of gross marginal profits (λ_t) expected to be earned by permanently increasing the size of the firm's workforce.³² Without adjustment costs employment would be set so that λ_t always equals zero. With adjustment costs λ_t only equals zero in steady state.

A parametric form for the production function remains to be chosen. The distributed

²⁹In a previous version of this paper, I included capital in the production function and found qualitatively similar results. Without time series data on capital adjustment, such additions add little to the empirical work.

³⁰While it is by now well recognized that the microeconomic details of the adjustment costs faced by firms can influence the aggregate dynamics of factor demand (Caballero et al. 1993, 1997), the gains from modeling such processes are likely to be small in this situation since the phenomenon in question is quite generally applicable to sluggish demand shifts.

³¹Classic examples of the use of quadratic labor adjustment costs under rational expectations include Sargent (1978) and Shapiro (1986).

³²This representation of dynamic labor demand is similar to the q-theory representation of investment as expounded in, for example, Hayashi (1982).

lags in Figure 3 indicate a long run employment price elasticity of around .75. A suitable production function would be capable of accommodating this behavior. Steady state wages are:

$$w = PAF'(L)$$

Totally differentiating the above and imposing long run wage equalization yields a long run employment price elasticity of

$$\epsilon_{l,p} = -\frac{F'(L)}{F''(L)L}$$

which for the case of a Cobb-Douglas production function of the form $F(L) = L^a$ can be shown to equal $\frac{1}{1-a}$. Because this elasticity is bounded below by one for $a > 0$, it will not do for this analysis. Instead, I use the following "isoelastic" generalization of a single input Cobb Douglas function capable of exhibiting sufficient concavity to yield long run elasticities below one:

$$F(L) = C - \frac{L^{-\alpha}}{\alpha} \tag{10}$$

where C is a positive constant and α is allowed to vary over the entire real line.³³ It is straightforward to show that this function exhibits a long run elasticity of $\frac{1}{1+\alpha}$ which will lie below unity for $\alpha > 0$ and exceed it for $\alpha < 0$. The parameter C , which is necessary only to ensure that output is positive at all employment levels, falls out of the first order conditions for employment since the marginal product of labor is simply $F'(L) = L^{-(1+\alpha)}$.

6.3 Equilibrium

Having laid out the equations governing labor supply and demand we now attempt to characterize the resulting equilibrium. The migration probabilities expressed in (3) make clear that the flows between sectors are a function of both current and future wage premia. The upper panel of Figure 5 illustrates gross flows into and out of the oil industry as a function of the current oil wage conditional on beliefs about the future path of wages. Each flow curve has a logistic shape reflecting the functional form of the choice probabilities. The particular shape and position of each curve depends upon the number of workers in the originating sector, the distance between sectors, and the scale σ of the taste shocks. Flow curves from large sectors will have shapes that appear to be stretched vertically, since small changes in probability will yield large changes in flows. Similar effects will ensue if σ or d is small. In the steady state, the two curves will be symmetric (because both sectors are the same size)

³³By L'Hopital's rule, as α approaches zero, $\frac{L^{-\alpha}}{\alpha}$ approaches $-\ln(L)$. Values of α above zero are more concave than a simple logarithm, while values below zero are less concave.

and will cross at wage w^n , at which point net flows will be zero.

In the wake of an oil price increase V_t^o will rise relative to V_t^c on expectations of future changes in the oil wage. This will lead the inflows curve to shift to the right and the outflows curve to shift to the left, thereby motivating large net flows into the oil sector equal to the horizontal distance between the two curves at the going wage. This increase in the size of sector O has important feedback effects on the system.

First, net in-migration will put downward pressure on wages as the marginal product of labor is gradually reduced. Second, as sector O expands, the base population at risk of emigrating from sector O increases, shifting the outmigration curve to the right. Finally, the realized wage increases eventually cause sector N to grow thereby offsetting the effects on the immigration curve of the decreases in wages. The new steady state equilibrium illustrated in the bottom panel of Figure 5 has larger gross flows in both directions, larger sectoral workforces, and wages equal to their original steady state level.

It is convenient to illustrate the dynamics of the equilibrium in terms of the behavior of net flows ΔL_t^o to the oil industry since we may also graphically represent demand in such a space by means of equation (9). A key feature of this model is that the gross migration curves and consequently L_t^o depend upon expectations of future changes in demand. If labor demand were expected to increase in the future but for some reason had not yet shifted, we would actually expect to see wages *decrease* in response to a price shock.³⁴ Even if demand did shift contemporaneously, if the future changes in demand were expected to result in substantial wage premia, the supply curve might shift enough for wages to fall on impact.

Figure 6 illustrates such a case graphically. Here we graph the supply and demand for net migration to the oil industry in wage quantity space. We start at the steady state where wages are such that $\Delta L_t^o = 0$ meaning that gross outmigration equals gross immigration. Oil prices increase raising the expected continuation value of being in sector O and causing both the supply and demand curves to shift out to S' and D'. This leads wages to fall very slightly but results in large flows into the sector. As the sector grows, the demand for additional hires falls and the demand curve shifts to the left. However, the supply curve of net migrants also shifts to the left. This happens for two reasons. One is that the number of workers in nearby sector N is drawn down causing the immigration curve $m_{n,o}$ to shift leftward. Second, as sector O grows, outmigration becomes more common since more workers are at risk of emigrating. This serves to diminish net flows into the sector and consequently for demand

³⁴Topel (1986) finds an analogous result in his analysis of the migratory response to predictable changes in local labor market conditions.

to outstrip supply and for wages to rise. Labor demand continues to ratchet to the left as λ_t is driven down by increases in L_o and the demand curve approaches its steady state. Labor supply also continues to shift to the left as the temporary wage changes are realized leading the expected continuation value of residence in sector O to fall. These shifts lead wages to settle down to an equilibrium near their old level.

7 Structural Estimation

The key element of the model rationalizing the slow movement of wages is that workers interpret oil price shocks as signals of future changes in the wage. Because of this oil prices can, in the short run, motivate very large flows into the sector at a constant wage. Competitive firms delay paying such wages because adjustment costs lead aggregate labor demand to shift out slowly. Once the wage premia emerge, they are driven back towards parity with the outside world by increases in sector size. This section asks what sort of parameter values are necessary to rationalize this behavior.

After specifying a stochastic process for oil prices, the equations in (3), (4), (5), (6), (7), (9), and (10) collectively characterize the dynamic stochastic process generating the labor market variables. To solve the system I use Dynare++ 1.3.5, which is a C++ routine for numerically simulating Dynamic Stochastic General Equilibrium (DSGE) models via perturbation methods.³⁵ Policy functions are obtained by calculating Taylor series approximations to the decision rules implied by a series of dynamic stochastic first order conditions. Because these approximations are made around a deterministic steady state I specify log oil prices to be a near unit root so that a proper steady state can be said to exist.³⁶ The specification used is

$$\ln(P_t) = .001\mu + .999 \ln(P_{t-1}) + \xi_t \quad (11)$$

where $\mu = 3.91$ is the log deflated oil price in the first quarter of 1972. The simulations assume that ξ is a normally distributed i.i.d. shock with variance equal to 0.02, the empirical variance of log oil price changes.

There are nine structural parameters in the system: $\beta, w_n, A, \chi, \sigma, d, k, \rho, \alpha$. I calibrate $\beta = .95$ and impose that steady state wages in the oil industry w_n equal their 1972 value of

³⁵For details see Juillard (1996) and Collard and Juillard (2001a, 2001b).

³⁶In fact, it is hard to believe that oil prices, even in logarithms, follow a pure random walk. It is well acknowledged that the best forecast for oil prices over the very long run is somewhere near the historical mean of approximately 20 dollars a barrel.

\$8.86/hr. I also calibrate the production scale parameter A to guarantee that steady state employment in the oil industry equals 99.6 thousand, its value in 1972, which is roughly the modal size of the oil workforce experienced over the sample period.

This leaves six parameters $(\chi, \sigma, d, k, \rho, \alpha)$ free to estimate. Taking a minimum distance approach, I treat the distributed lag coefficients from the previous regressions of relative wages and employment on oil prices as moments to be matched to theoretical impulse responses (IRFs) to oil shocks generated by Dynare. Thus the relevant moment conditions are of the form

$$E[g(\gamma)] = E\left[\widehat{\beta}_{ols} - \beta(\gamma)\right] = 0 \quad (12)$$

where $\beta_{ols} = \begin{bmatrix} \widehat{\beta}_{ols}^L \\ \widehat{\beta}_{ols}^w \\ \widehat{\beta}_{ols} \end{bmatrix}$ is the stacked 50x1 vector of distributed lag coefficients estimated from OLS equations of the form given in (2) subject to the additional constraint that the effect of oil prices on wages dies out after six years.³⁷ The vector $\beta(\gamma)$ contains the impulse responses generated by the model with structural parameter vector γ . I choose parameter estimates $\widehat{\gamma}$ to minimize the quadratic form

$$Q(\gamma) = g(\gamma)' \widehat{V}^{-1} g(\gamma)$$

where \widehat{V} is a 50x50 heteroscedasticity and autocorrelation consistent (HAC) estimate of the covariance matrix of the reduced form distributed lag coefficients.³⁸

The simulated IRFs were generated by a first order approximation method. Higher order methods yield nearly identical results. The parameters were transformed before conducting the minimization process in order to make the minimization unconstrained. Four parameters: $\chi, \sigma, d,$ and k were transformed into logarithms. A logit transform was used to map the values of ρ into the real line. The variance-covariance matrix of the transformed parameters was computed via the formula $\tilde{V} = \left(G' \widehat{V}^{-1} G\right)^{-1}$, where G is the 50x6 numerical gradient of the moment conditions with respect to the structural parameters. Standard errors for the untransformed parameters were recovered via the Delta method.

7.1 Parameter Estimates

The estimated parameter values and standard errors for the model are shown below:

³⁷This is implemented by differencing the final lag of p_t in the wage equation.

³⁸The HAC covariance estimates are computed using a Quadratic Spectral kernel on residuals pre-whitened by a VAR in conjunction with the automated bandwidth selection procedure detailed in Andrews (1991).

Minimum Distance Parameter Estimates and Standard Errors

Parameter	χ	σ	k	d	α	ρ	χ^2
Interpretation	Adjustment Cost	Variance of Taste Shocks	Short Run Sensitivity of Outside World	Switching Cost	Concavity of Production	Sluggishness of Outside Response	Over-Identification Statistic
Estimated Value	2.51 (0.68)	0.55 (0.22)	0.16 (0.20)	0.98 (0.27)	0.26 (0.28)	0.93 (0.22)	186.46

Standard errors in parentheses

All of the parameters save for k and α are estimated quite precisely. The chi-squared goodness-of-fit test which has only 44 degrees of freedom, easily rejects that the differences between the simulated IRFs and the reduced form moments are due to chance. The primary difficulties come from matching the near linear estimated ramp-up of employment and wages in the immediate aftermath of an oil price shock. This behavior presents a challenge for models of quadratic adjustment costs which naturally prefer geometric to linear employment adjustment.

Turning now to the parameters, the fundamental metric of costs in this setup is quarterly dollars per hour. For example, moving from sector N to O yields an estimated mobility cost d equal to approximately \$0.98/hr. in wages paid over the duration of a quarter or, roughly, a week and a half's worth of steady state earnings.³⁹ The standard deviation of the taste shocks can be shown to be $\frac{\pi}{\sqrt{6}}\sigma$ and is also measured in quarterly dollars per hour. The estimate of $\sigma = .55$ implies that the standard deviation of the transitory taste shocks is equivalent to \$0.71/hr. in quarterly wages or about a week's worth of steady state earnings.

The adjustment cost parameter χ indicates the marginal cost to the firm in quarterly dollars per hour of expanding the workforce by 1,000 laborers. The estimated value of $\chi = 2.51$ implies that the 1,000th worker hired costs the firm \$2.51/hr., equal to around a fourth of that worker's quarterly earnings in steady state. If the average worker works 45 hours a week and thirteen weeks a quarter this means the total dollar value of the adjustment cost is $45 \times 13 \times 2.51 = \1468 . The average per capita cost of hiring 1,000 workers is half this amount at \$653.

In keeping with equation (9), we can also think of χ as a mapping between discounted gross marginal profits on labor and desired rates of increase in the oil labor force. A value of $\chi = 1$ would imply that a dollar per hour gross profit on the marginal worker (with total capitalized dollar value of $45 \times 13 = \$585$) motivates the representative firm to expand the workforce by 1,000 workers or approximately 1% near steady state over the next quarter.

³⁹Steady state wages are \$8.86. Assume 13 weeks in a quarter. $0.98/8.86 = 0.111 \approx 1.5 \times \frac{1}{13}$.

By contrast, the estimated value of χ indicates that a dollar per hour gross profit would lead the sector to demand only $\frac{1}{2.51} \times 1000 \approx 400$ additional workers over the quarter.

The parameter α is a measure of the concavity of the representative firm's production function. Recall from earlier that the long run employment price elasticity of this production technology is $\epsilon_{l,p} = \frac{1}{1+\alpha}$. Because the empirical value of $\epsilon_{l,p} \approx .75$ one would expect that any attempt to match the long run behavior of the distributed lags would require $\alpha \approx 1/3$. The estimated value of α is slightly less than this fraction at .26 which implies a long run elasticity $\epsilon_{l,p}$ of approximately .8.

The parameters k and ρ govern the rate at which workers flow into sector N in response to expected wage premia in sector O. A value of $k = .16$ indicates that a one percent premium in the oil wage motivates a contemporaneous .16% increase in the number of workers searching for work in the oil industry. The estimates of ρ indicate how sluggish the migratory responses are with higher values indicating slower responses. The long run response to a permanent one percent premium in the wage is $\frac{\kappa}{1-\rho} L^n$ so (conditional on κ) higher values of ρ also indicate larger long run responses. The value of ρ is estimated to be quite high indicating some fairly substantial delay in the migratory response to oil shocks. We will see in impulse responses that the delayed response of F_t to price shocks implied by these estimates allow for the number of workers nearby to be driven down to the point where wages eventually rise due to increases in market tightness (i.e. due to a scarcity of workers in sector N).

7.2 Impulse Responses

At least as interesting as the estimated parameters are the impulse responses they imply. Figure 7 shows the simulated response of oil sector wages and employment to a permanent price shock at the estimated parameter values along with two alternative paths implied by negative perturbations of d and χ of two standard errors. The price innovation is one log point in magnitude and the impulse responses are in logarithmic deviations from steady state so they may be read as elasticities.

The minimum distance estimates match the behavior of the distributed lag coefficients illustrated in Figure 2 relatively well. Wages fall slightly on impact and then begin to rise, peaking roughly two and a half years after the shock and then slowly declining. Employment jumps on impact and proceeds to ramp up rapidly towards its new steady state level. Although wages peak somewhat later and at lower values than in the estimated distributed lags, the largest discrepancy between the simulations and Figure 2 comes from the failure of employment to actually converge to its new steady state over the simulation horizon. This

occurs not because adjustment costs are prohibitively high, but rather because the eventual decreases in the wage lead to small increases in the quantity of labor demanded by firms.

To illustrate how the impulse response curves depend on the model parameters two alternative IRFs are displayed: one with lower firm adjustment costs (χ) and the other with lower migration costs (d) on the part of workers. Lowering χ has relatively straightforward effects, labor demand shifts out more swiftly leading oil employment to jump more on impact and ramp up faster while wages are also driven up faster and to greater peaks. Lowering worker migration costs (d) has more subtle effects. The employment response to the price shock is essentially unchanged by decreases in migration costs, but the wage response differs. Wages no longer fall on impact and now ratchet up to a lower peak that occurs earlier than before. To understand this note that in the absence of migration costs workers would never be willing to switch into the oil sector when there is a drop in current wages even if future wages were expected to rise because they can always wait until wages actually do rise to switch. But with migration costs, a worker with a high enough draw of the taste shocks ($v_{it}^o - v_{it}^n$) in a given period will be willing to switch into the oil industry and take a wage cut in anticipation of future wage gains because the effective switching cost $d - \sigma(v_{it}^o - v_{it}^n)$ is *temporarily* low. Thus the migration costs d interact with the variance of the taste shocks σ to determine the degree to which workers are willing to substitute intertemporally.

As already discussed, net changes in sectoral employment are the sum of the gross migration flows between sectors. Figure 8 illustrates the response of gross flows between sectors O and N to an oil shock. In all cases, flows into the sector jump on impact then ratchet down as the number of workers in sector N is depleted. Eventually however, as flows from the outside world to sector N ramp up, flows to sector O ramp up along with them and reach a new larger steady state. The exact timing of this inflection depends primarily upon the parameters k and ρ which govern how quickly sector N is drawn down and refills. Gross flows out of sector O, which can also be thought of as industry-wide separations, fall on impact but then ratchet up as the sector grows.

Figure 9 shows the simulated behavior of some other key variables in the model. The value of being in sector O jumps on news of the price shock and slowly intensifies as the largest wage premia draw closer. As wages begin to dissipate so do the value functions of which they are a forward moving average. The value of search V_t^n also jumps on impact by virtue of the option of switching to sector O and then decreases towards its previous steady state. Net flows from the outside world fall very slightly on impact with the wage, but then slowly ramp up in response to emergence of substantial wage premia. With more forward

looking agents in the outside world, flows would actually ramp up in anticipation of those premia.

7.3 Auxiliary Evidence

We have seen that a model of sectoral migrations with slow moving demand can recreate the qualitative features of wages and employment we initially set out to explain. We ask now whether any auxiliary evidence can be brought to bear on the mechanisms generating the employment and wage dynamics in the model. Fluctuations in the wage are ultimately driven by the dynamic scarcity of labor. Wages rise because at some point insufficiently many cheap workers are available. The structural model allows us to simulate the dynamics of labor scarcity by examining the response of the size of sector N to shocks. When L_t^n is below steady state, the labor market will be tight and it will be hard to attract large numbers of workers without wage premia.

A key question then is whether L_t^n has any empirical analogue. The traditional measure of market tightness is the unemployment rate. Since in this model sector N is meant to represent some notion of the number of workers engaged in (directed) search, unemployment may not be a bad proxy for L_t^n . I construct for use as a proxy the log difference⁴⁰ in unemployment rates between Texas and the entire U.S. from two sources: the Employment and Training Administration's insured unemployment rate, which is derived from state Unemployment Insurance (UI) data available beginning in 1987, and the BLS Local Area Unemployment Statistics (LAUS) series, which is based upon CPS data available starting in 1976. Both series are seasonally adjusted. Since the UI data are not based upon a sample, and are not smoothed over using a time series model, they are more likely to be able to accurately capture some of the high frequency movements in state level unemployment that might be disregarded as noise in the LAUS estimates.⁴¹ As with relative wages, log deviating the unemployment rate in Texas from the national rate is meant to reduce the influence of any macroeconomic disturbances on the analysis.

Figure 10 plots the simulated impulse responses of L_t^n under the various parameter schemes against distributed lags of both measures of unemployment on oil prices. Given that none of the parameters used in the simulations were estimated taking these distributed

⁴⁰Logarithms are used to place the empirical distributed lags in the same units as the simulations which show the behavior of $\ln(L_t^n)$.

⁴¹The LAUS series uses a state-space smoother to distinguish signal from noise in the monthly state level data. This smoothing may result in some of the high frequency variation in the signal being lost. See <http://www.bls.gov/lau/laumthd.htm>.

lags into account, the similarity between these plots is remarkable. All show a U shaped impulse response with minimums between one and two years after the shock. While it is not surprising that unemployment falls in Texas relative to the U.S. average when the price of oil increases, there is no reason to suppose that the timing and magnitude of the response would look like this. The resemblance between the UI coefficients and the simulated IRFs is especially striking. Not only is the turning point of the response similar, but both the simulated and estimated IRF's exhibit overshooting in the sense that the pool of unemployed actually increases permanently after the shock.⁴²

The general resemblance of relative unemployment in Texas to the simulated path of the model's unobservables neither confirms the assumptions underlying the model nor suggests that unemployment is the primary element of market tightness. Many of the workers considering entering the oil industry are probably employed in other sectors. Moreover, oil shocks likely affect Texas unemployment through mechanisms over and above changes in the demand for oil workers. What this exercise tells us however is that the dynamic relationship between Texas labor market conditions and oil prices follows just the pattern necessary to rationalize the empirical response of oil industry wages and employment to price shocks through the model. This is a provocative finding and one that should stimulate further thinking on local labor market dynamics.

8 Conclusion

This paper has investigated the dynamics of a single sectoral labor market. The empirical finding that, on average, wages lag employment in response to exogenous shocks to labor demand is at odds with the predictions of conventional market clearing models. I have shown that a forward looking model of sectoral reallocation with standard adjustment rigidities can rationalize this behavior with sensible parameter values.

An obvious question is the extent to which such a model might generalize to other industries or environments. The recipe for equilibrium behavior of the sort under study is clear: an industry should be subject to large and recurrent shocks to product demand, firms in the industry should face substantial adjustment rigidities, and there should exist a pool of low cost workers capable of entering the industry in the short run. Whether large segments of the U.S. economy meet these criteria is an open question. In manufacturing industries

⁴²Similar results are obtained in the empirical distributed lags if one examines unemployment levels instead of rates. This should be relatively unsurprising since the log difference in unemployment rates is the sum of the difference in log unemployment levels and the difference in log labor forces. The latter quantity is very slow moving and exhibits relatively little variation in response to oil shocks.

other factors such as unionization are likely to complicate the wage setting process, while in higher skilled industries such as engineering, lags in the training of workers are likely to add additional dynamics to the labor supply decision.

Caveats aside, there is evidence of similar employment and wage dynamics in a few important settings. The first is Carrington's (1996) study of the building of the Alaskan pipeline, which found a slow increase in the earnings of workers in construction and related industries which eventually reverted to trend.⁴³ Second, Blanchard and Katz (1992) find hump shaped wage responses to seemingly permanent labor demand shocks in panels of U.S. states. In both papers most of the employment adjustment seems to occur before wages peak. Finally, in a closely related paper, Topel (1986) finds evidence of state level wages falling in response to predictable changes in local labor demand. The evidence in these papers raises the possibility that labor market dynamics of the sort modeled here may be found in settings far more general than the oil industry.

A few additional points are worth taking away from this exercise. First, the labor market under study is extremely flexible. Between 1978 and 1982 employment in oil and gas field services doubled while over the next four years employment fell back to its 1978 level. These adjustments highlight the ability of well functioning markets to effectively match workers to jobs. However, a well functioning matching process does not imply that reallocations are costless. The behavior of wages over the course of these dramatic shifts in labor demand suggests that sectoral flows impose substantial costs on both workers and firms. A researcher armed with detailed longitudinal microdata might take seriously the task of estimating the social costs associated with such high frequency intersector reallocations.⁴⁴

Second, despite the flexibility of the oil labor market, permanent demand shocks appear to be associated with wage premia that persist for several years, even when the system is begun in steady state. A series of persistent shocks such as those experienced by the oil industry can keep wages out of steady state for decades at a time. To the extent that these sorts of persistent expansions and contractions are present in other industries, an important component of sectoral choice is likely to involve market conditions. More work is needed linking standard models of sectoral choice with dynamic market models.⁴⁵ Particularly

⁴³See figs. 3,7, 8 and especially fig. 9 in that paper.

⁴⁴Lee and Wolpin (2006) provide a detailed analysis of the social costs of the long run reallocation of labor between the service and manufacturing sectors.

⁴⁵The literature on sectoral and occupational choice is too large to document here. Starting with Roy's original (1951) contribution there have been several notable attempts to estimate models of selection in the labor market. Famous examples include Willis and Rosen (1979) and Heckman and Sedlacek (1985, 1990). One of the great challenges in linking selection models with dynamic market models is disentangling

fertile ground for research may be found in developing countries heavily invested in exporting commodities subject to large persistent price risks. Labor supply decisions in such countries are likely to be fundamentally influenced by expectations and uncertainty regarding the future path of commodity prices. Learning more about the dynamics of these decisions and how they interact with individual heterogeneity may provide important insights for the crafting of effective industrial and labor market policies.

Finally, the model presented here is applicable to labor markets defined in spaces more general than output sectors. A natural parallel is to local and regional labor markets. As in Topel (1986), the analysis presented here has been one of spatial equilibrium. But unlike with Topel's model, the implication has been that the dynamic linkages between markets are governed in part by the "distance" separating them. Thinking carefully about networks of local labor markets that are (perhaps unequally) dynamically interrelated holds the promise of revealing deeper insights into how labor markets adjust to shocks.

heterogeneity and dynamics. Recent advances in statistical modeling and data availability may soon yield great progress in this area.

References

1. Abraham, Katherine, James Spletzer, and Jay Stewart. 1998. "Divergent Trends in Alternative Wage Series." *Labor Statistics Measurement Issues*, eds. Haltiwanger, Manser, and Topel. Chicago: University of Chicago Press.
2. Andrews, Donald. 1991. "Heteroscedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica*, 59(3):307-345.
3. Barsky, Robert and Lutz Kilian. 2002. "Do We Really Know That Oil Caused the Great Stagflation? A Monetary Alternative." *NBER Macroeconomics Annual* 16:137-183.
4. Barsky, Robert and Lutz Kilian. 2004. "Oil and the Macroeconomy Since the 1970s." *Journal of Economic Perspectives* 18(4):115-134.
5. Bartik, Timothy. 1991. *Who Benefits From State and Local Economic Development Policies?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.
6. Blanchard, Olivier and Lawrence Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity* 1992(1):1-75.
7. Bound, John and Harry Holzer. 2000. "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s." *Journal of Labor Economics* 18(1):20-54.
8. Brainard, Lael S. and David M. Cutler. 1993. "Sectoral Shifts and Cyclical Unemployment Reconsidered." *Quarterly Journal of Economics* 108(1):219-243
9. Bresnahan, Timothy and Valerie A. Ramey. 1993. "Segment Shifts and Capacity Utilization in the U.S. Automobile Industry." *American Economic Review Papers & Proceedings* 83(2):213-218.
10. Bureau of Labor Statistics. 2005. "Career Guide to Industries: Oil and Gas Extraction" obtained online at: <http://www.bls.gov/oco/cg/cgs005.htm>.
11. Caballero, Ricardo, Eduardo Engel, and John Haltiwanger. 1997. "Aggregate Employment Dynamics: Building from Microeconomics." *American Economic Review* 87(1):115-137.
12. Caballero, Ricardo and Eduardo Engel. 1993. "Microeconomic Adjustment Hazards and Aggregate Dynamics." *Quarterly Journal of Economics* 108(2):359-383.

-
13. Carrington, William J. 1996. "The Alaskan Labor Market During the Pipeline Era." *Journal of Political Economy* 104(1):186-218.
 14. Chan, William. 1996. "Intersectoral Mobility and Short Run Labor Market Adjustments." *Journal of Labor Economics* 14(3):454-471.
 15. Collard, Fabrice and Michel Juillard. 2001a. "A Higher-Order Taylor Expansion Approach to Simulation of Stochastic Forward-Looking Models with an Application to a Non-Linear Phillips Curve." *Computational Economics* 17(2):125-139.
 16. Collard, Fabrice and Michel Juillard. 2001b. "Accuracy of Stochastic Perturbation Methods: The Case of Asset Pricing Models." *Journal of Economic Dynamics and Control* 25(7):979-999.
 17. Davis, Steven J. 1987. "Fluctuations in the Pace of Labor Reallocation." *Carnegie-Rochester Conference Series on Public Policy* 27:335-402.
 18. Davis, Steven and John Haltiwanger. 2001. "Sectoral Job Creation and Destruction Responses to Oil Price Changes." *Journal of Monetary Economics* 48(3):465-512.
 19. Elliott, Graham, Thomas Rothenberg and James Stock. 1996. "Efficient Tests for an Autoregressive Unit Root." *Econometrica* 64(4):813-836.
 20. Hamilton, James D. 1983. "Oil and the Macroeconomy Since World War II." *Journal of Political Economy* 91(2):228-248.
 21. Hamilton, James D. 1985. "Historical Causes of Postwar Oil Shocks and Recessions." *The Energy Journal* 6(1):97-116.
 22. Hamilton, James D. 1988. "A Neoclassical Model of Unemployment and the Business Cycle." *Journal of Political Economy* 96(3):593-617.
 23. Hamilton, James D. 2003. "What is an Oil Shock?" *Journal of Econometrics* 113:363-398.
 24. Hayashi, Fumio. 1982. "Tobin's Marginal q and Average q: A Neoclassical Interpretation." *Econometrica* 50(1):214-224.
 25. Heckman, James J. and Guilherme L. Sedlacek. 1985. "Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-Selection in the Labor Market." *Journal of Political Economy* 93(6):1077-1125.

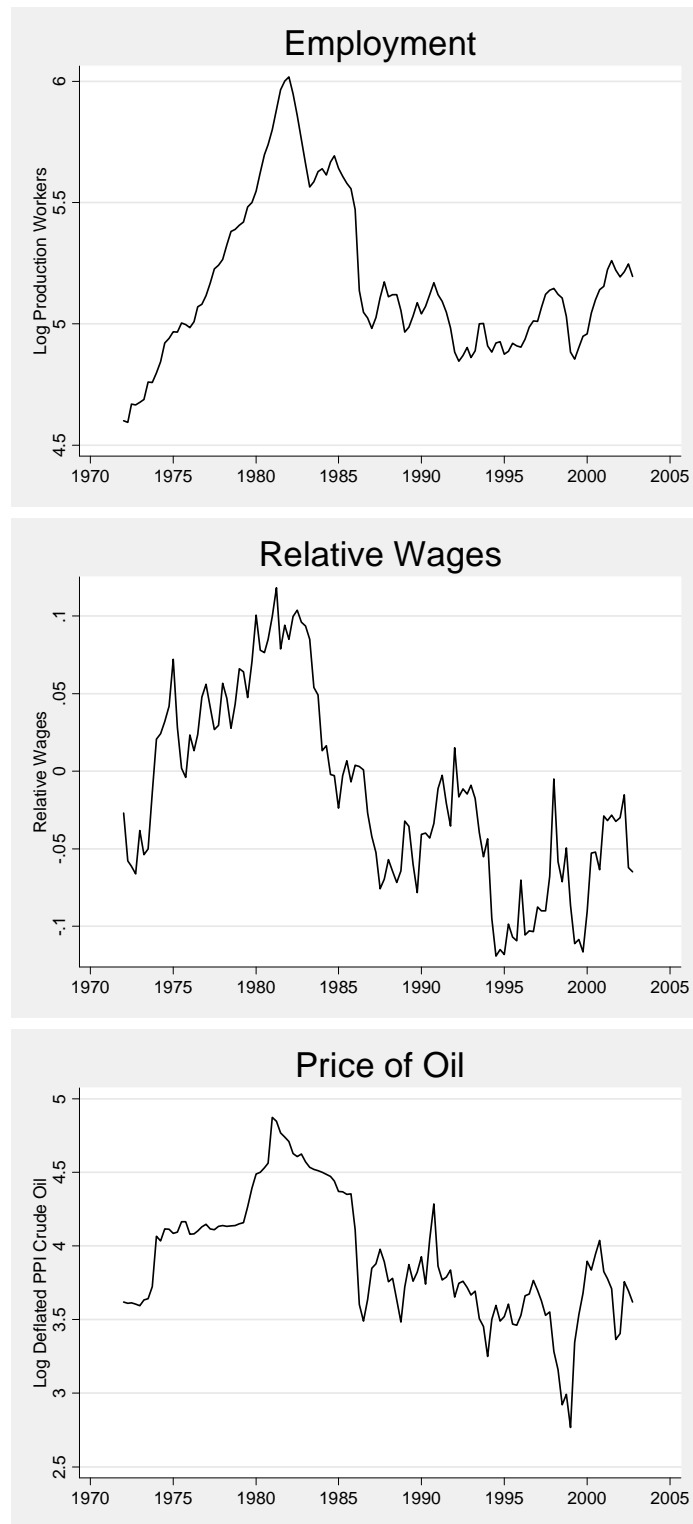
-
26. Heckman, James J. and Guilherme L. Sedlacek. 1990. "Self-Selection and the Distribution of Hourly Wages." *Journal of Labor Economics* 8(1):S329-S363.
 27. Juillard, Michel. 1996. "DYNARE: A Program for the Resolution and Simulation of Dynamic Models with Forward Variables through the Use of a Relaxation Algorithm." CEPREMAP, Couverture Orange No. 9602.
 28. Keane, Michael and Eswar Prasad. 1996. "The Employment and Wage Effects of Oil Price Changes: A Sectoral Analysis." *Review of Economics and Statistics* 78(3):389-400.
 29. Lee, Donghoon and Kenneth Wolpin. 2006. "Intersectoral Labor Mobility and the Growth of the Service Sector." *Econometrica* 47(1):1-46.
 30. Lucas, Robert E. and Edward C. Prescott. 1974. "Equilibrium Search and Unemployment." *Journal of Economic Theory* 7(2):188-209.
 31. McFadden, Daniel. 1978. "Modeling the Choice of Residential Location." in *Spatial Interaction Theory and Planning Models* ed. Karlqvist, Lundqvist, Snickars, and Weibull. Amsterdam: North Holland.
 32. McFadden, Daniel, Dennis Bolduc, and Michel Bierlaire. 2003. "Characteristics of Generalized Extreme Value Distributions." Mimeo.
 33. Newey, Whitney and Kenneth West. 1994. "Automatic Lag Selection in Covariance Matrix Estimation." *Review of Economic Studies* 61(4):631-653.
 34. Ng, Serena and Pierre Perron. 2001. "Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power." *Econometrica* 69(6):1519-1554.
 35. Phelan, Christopher and Alberto Trejos. 2000. "The Aggregate Effects of Sectoral Reallocations." *Journal of Monetary Economics* 45(2):249-268.
 36. Ramey, Valerie and Matthew Shapiro. 1998. "Costly Capital Reallocation and the Effects of Government Spending." *Carnegie-Rochester Conference Series on Public Policy* 48:145-194.
 37. Ramey, Valerie and Matthew Shapiro. 2001. "Displaced Capital: A Study of Aerospace Plant Closings." *Journal of Political Economy* 109(5):958-992.
 38. Rogerson, Richard. 1987. "An Equilibrium Model of Sectoral Reallocation." *Journal of Political Economy* 95(4):824-834.

-
39. Roy, A.D. 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers* 3(2):135-146.
 40. Sargent, Thomas. 1978. "Estimation of Dynamic Labor Demand Schedules Under Rational Expectations." *Journal of Political Economy* 86(6):1009-1044.
 41. Shapiro, Matthew. 1986. "The Dynamic Demand for Capital and Labor." *Quarterly Journal of Economics* 101(3):513-542.
 42. Topel, Robert. 1986. "Local Labor Markets." *Journal of Political Economy* 94(3):S111-S143.
 43. Willis, Robert and Sherwin Rosen. 1979. "Education and Self-Selection." *Journal of Political Economy* 87(5):S7-S36.

Table 1: Monthly Turnover Rates for Selected Mining Industries

Oil and Gas Field Services						
	Accession rates			Separation rates		
Date	Total	Hires	Recalls	Total	Quits	Layoffs
July 1978	9.6	7.7	0.7	8.0	5.9	0.2
September 1979	10.8	8.8	1.5	9.8	7.3	0.3
August 1980	10.8	9.0	1.3	10.7	8.8	0.3
May 1981	11.3	9.6	1.1	9.5	6.8	0.6
September 1981	10.6	9.7	0.7	10.4	7.5	0.3
Mean	10.62	8.96	1.06	9.68	7.26	0.34
Crude Petroleum, Natural Gas, and Natural Gas Liquids						
	Accession rates			Separation rates		
Date	Total	Hires	Recalls	Total	Quits	Layoffs
July 1978	2.4	1.8	0.5	1.7	1.1	0.1
September 1979	2.6	1.9	0.4	2.5	1.3	0.2
August 1980	2.4	1.6	0.3	2.7	1.8	0.1
May 1981	4.0	3.4	0.4	1.5	1.1	0.1
September 1981	2.5	1.9	0.4	2.7	1.7	0.3
Mean	2.78	2.12	0.4	2.22	1.4	0.16
Nonmetallic Mining						
	Accession rates			Separation rates		
Date	Total	Hires	Recalls	Total	Quits	Layoffs
July 1978	4.1	3.6	0.4	3.2	2.2	0.4
September 1979	3.2	2.6	0.3	3.1	2.2	0.3
August 1980	2.6	1.8	0.7	4.2	1.8	1.5
May 1981	4.4	2.2	1.8	2.4	1.1	0.8
September 1981	2.1	1.5	0.5	2.7	1.4	0.7
Mean	3.28	2.34	0.74	3.12	1.74	0.74
Source: BLS Employment and Earnings						

Figure 1: Employment, Relative Wages, and the Price of Oil



Note: Relative wages are the log of the ratio of average production wages in the oil industry to average production wages in "nonmetallic mining. Oil prices are deflated using the CPI-U series.

Figure 2: Estimated Response to a Permanent Unit Log Increase in Oil Prices (Unrestricted Distributed Lag)



Figure 3: Estimated Response to a Permanent Unit Log Increase in Oil Prices (5th Order Polynomial Distributed Lag)

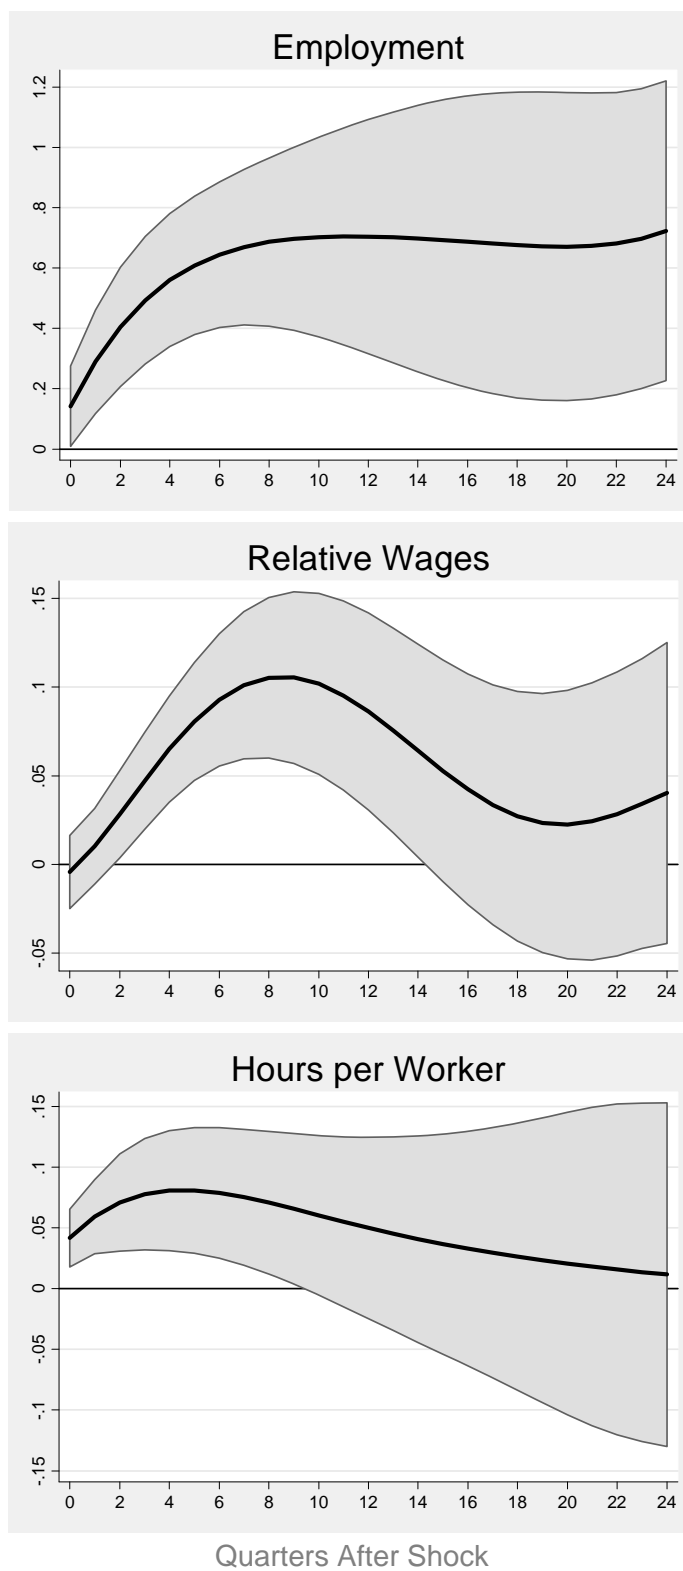


Figure 4: Equilibrium Response to a Sluggish Demand Shock




Figure 5: Gross and Net Migration Curves





Figure 6: Equilibrium Response to an Increase in the Price of Oil




Figure 7: Simulated Response of Employment and Wages to a Permanent Unit Log Increase in Oil Prices

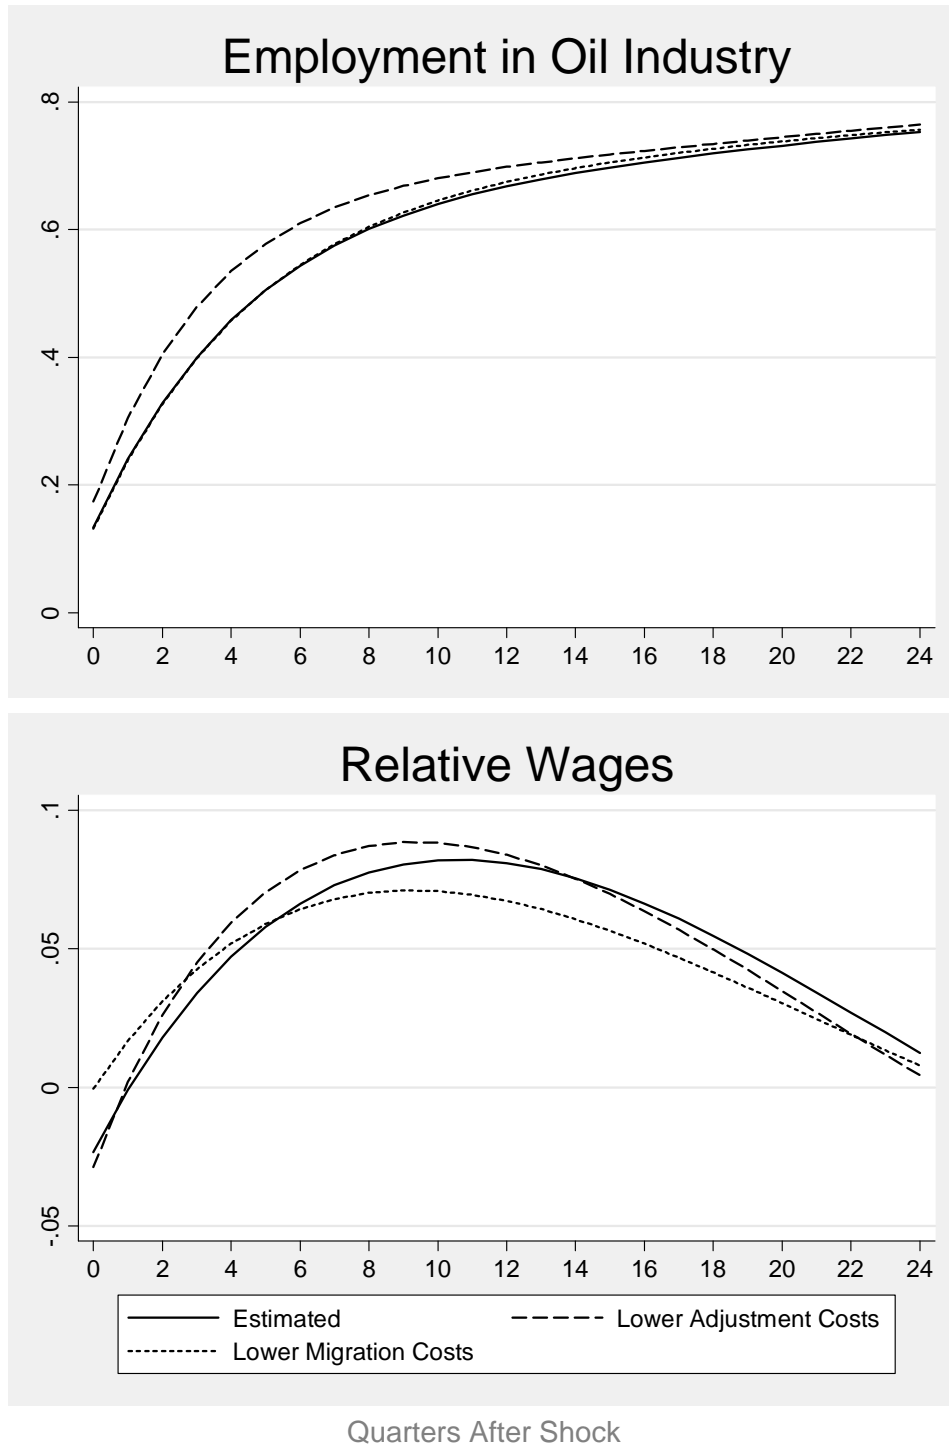


Figure 8: Simulated Response of Gross Migrations to a Permanent Unit Log Increase in Oil Prices

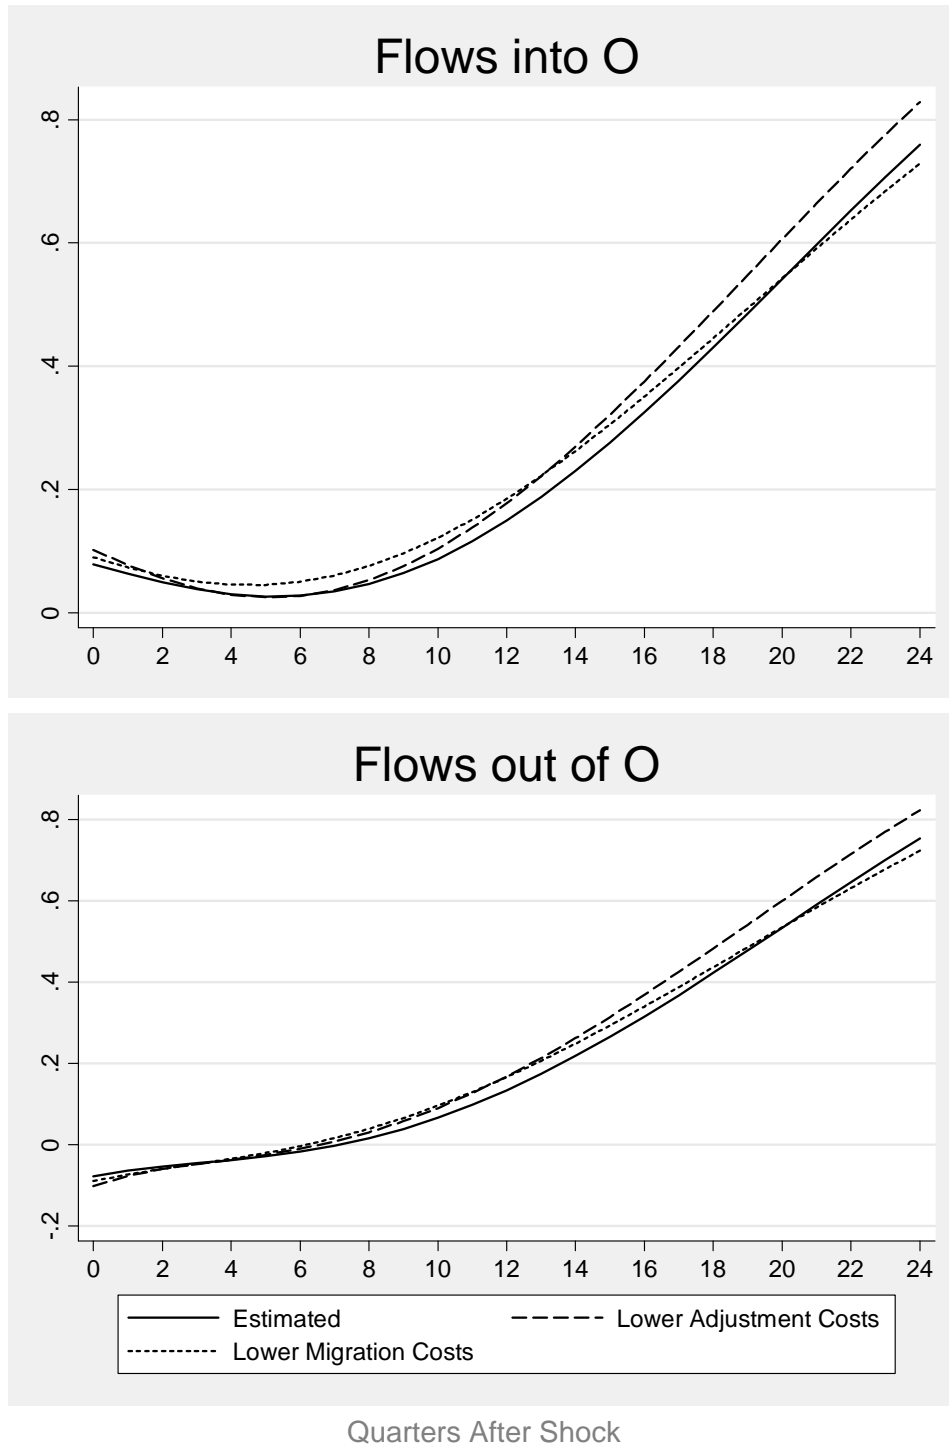


Figure 9: Simulated Response of Other Variables to a Permanent Unit Log Increase in Oil Prices

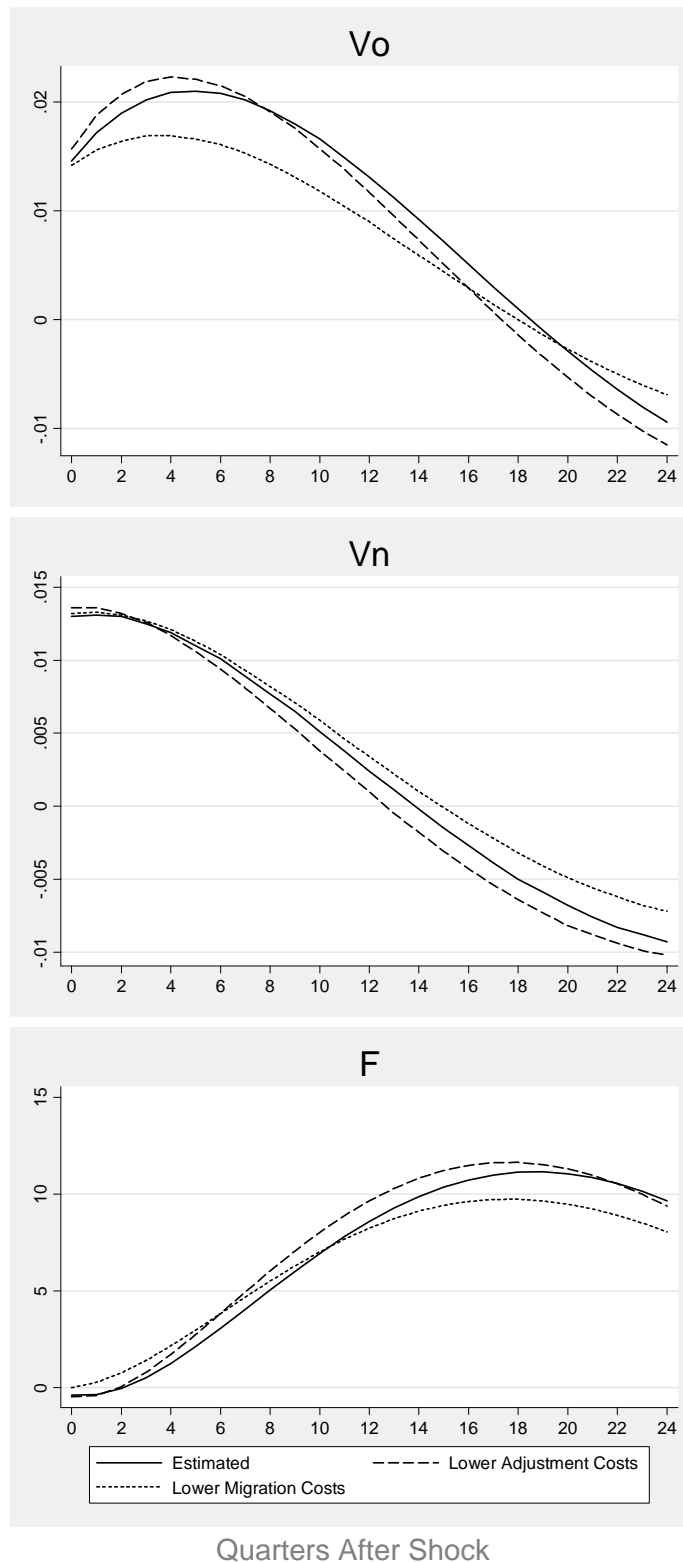


Figure 10: Estimated/Simulated Response of Market Tightness to a Permanent Unit Log Increase in Oil Prices

