The Dissertation Committee for Neil Ware White IV
certifies that this is the approved version of the following dissertation:

Essays on Labor Markets,
Monetary Policy, and Uncertainty

Committee:

Olivier Coibion, Supervisor

Saroj Bhattarai

Andrew Glover

Nicolas Petrosky-Nadeau

Essays on Labor Markets, Monetary Policy, and Uncertainty

by

Neil Ware White IV

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

May 2018

Dedicated to everyone who has ever taught me anything.

Acknowledgments

I wish to thank Olivier Coibion, Saroj Bhattarai, Andy Glover, and Nicolas Petrosky-Nadeau for their guidance, advice, and support. I also benefited greatly from discussions with Hassan Afrouzi, Chris Boehm, Nitya Pandalai-Nayar, Zheng Liu, Mary Daly, and innumerable seminar participants at the University of Texas at Austin, the Federal Reserve Bank of San Francisco, the Federal Reserve Board of Governors, Louisiana State University, Texas A&M University, the Federal Reserve Bank of Cleveland, the University of Oklahoma, Amherst College, Williams College, the University of South Carolina, and Purdue University.

**Essays on Labor Markets,
Monetary Policy, and Uncertainty**

by

Neil Ware White IV, Ph.D.
The University of Texas at Austin, 2018

Supervisor: Olivier Coibion

This dissertation examines the impacts on the labor market of monetary policy and macroeconomic uncertainty.

The first chapter examines how monetary policy shocks in the U.S. affect the flows of workers among three labor market categories—employment, unemployment, and nonparticipation—and assesses each flow's relative importance to changes in labor market "stock" variables like the unemployment rate. I find that job loss accounts for the largest portion of monetary policy's effect on labor markets. I develop a New Keynesian model that incorporates these channels and show how a central bank can achieve welfare gains from targeting job loss, rather than output, in an otherwise standard Taylor rule.

The second chapter examines the role of monetary policy in "job polarization." I argue that contractionary monetary policy has accelerated the decline of employment in routine occupations, which largely affected workers with a high-school degree but no college. In part by disproportionately affecting industries with high shares of routine occupations,

contractionary monetary policy shocks lead to large and persistent shifts away from routine employment. Expansionary shocks, on the other hand, have little effect on these industries. Indeed, monetary policy's effect on overall employment is concentrated in routine jobs. These results highlight monetary policy's role in generating fluctuations not only in the level of employment, but also the composition of employment across occupations and industries.

The third chapter introduces new direct measures of uncertainty derived from the Michigan Survey of Consumers. The series underlying these new measures are more strongly correlated with economic activity than many other series that are the basis for uncertainty proxies. The survey also facilitates comparison with response dispersion or disagreement, a commonly used proxy for uncertainty in the literature. Dispersion measures have low or negative correlation with direct measures of uncertainty and often have causal effects of opposite sign, suggesting that they are poor proxies for uncertainty. For the measures based on series most closely correlated with economic activity, positive uncertainty shocks are mildly expansionary. This result is robust across identification and estimation strategies and is consistent with "growth options" theories of the effects of uncertainty.

Table of Contents

List of Tables

List of Figures

Chapter 1

Gross Worker Flows, Job Loss, and Monetary Policy

1.1 Introduction

How does monetary policy affect the labor market? A common answer is that by directly changing interest rates, monetary policy changes the desired levels of consumption and investment and, consequently, affects output. Because short-run changes in production come primarily from changes in employment, monetary policy is ultimately able to influence the labor market. This is how the Federal Reserve describes monetary policy's ability to influence the real economy,[1] and much of the recent literature has focused on how monetary policy affects output and financial variables, rather than employment.[2]

In this paper, I take a different approach to examining how monetary policy influences the labor market. Rather than focusing on output measures and arguing that employment tracks output, I study the effects of monetary policy on the full set of gross *flows* of workers into and out of employment, unemployment, and nonparticipation. This approach allows me to address questions such as: Does monetary policy affect the labor market by influencing job finding rates? Does it affect job loss probabilities? Do participation decisions—workers'

[1]See `https://www.federalreserve.gov/aboutthefed/files/pf_complete.pdf`, pp.27–31.

[2]The literature quantifying the effects of monetary policy shocks on output, prices, and financial variables is too long to review in detail. Lawrence J. Christiano, Martin Eichenbaum and Charles Evans (1999) provide an overview of the early literature. Valerie A. Ramey (2016) provides a review of more recent developments.

choices to enter or leave the labor force altogether—drive monetary policy's impact on labor markets? As I demonstrate in this paper, the answers to these questions matter for understanding both the limits to and efficacy of monetary policy's ability to offset macroeconomic shocks.

I find that job loss—that is, workers moving from employment to unemployment—is the most important driver of the responses of employment, unemployment, and labor force participation after monetary policy shocks; job finding plays a secondary role. Moreover, ignoring the participation margin—as is common in the literature—hides quantitatively important labor force composition effects driven by job loss. To demonstrate this, I first give a detailed characterization of the effects of monetary policy shocks on the movement of workers among employment, unemployment, and nonparticipation. I then assess the relative importance of these flows in the transmission of monetary policy shocks to the following "stock" variables: the employment-to-population ratio (EP), unemployment rate (UR), and labor force participation rate (LFP). To do so, I examine the impulse responses of worker flows and exploit the fact that they sum up to the responses of the stock variables to construct decompositions of the stock responses into the underlying flow responses. I show that the flow of workers from employment to unemployment is the most important flow driving the responses of all three stock variables. Alternative decompositions (for example, into groups of flows) point toward the same result: job loss drives the labor market response to monetary policy shocks.

Throughout the paper, I focus on three mutually exclusive labor market states as defined by the Current Population Survey (CPS)—employment (E), unemployment (U), and nonparticipation (N)—and the flows of workers across these three states. To be specific, a worker who is employed one month and unemployed the next would constitute an *EU* flow, while a worker who is out of the labor force this month but unemployed next month would be an *NU* flow; other flows are defined in the same way. As is evident in Figure 1.1, substantial "churning" underlies the more familiar labor market stock measures. Although their dynamics are substantially more complicated, the flows *uniquely* pin down the stock variables and not *vice versa*; it is this appealing characteristic of the flows that I exploit in my analysis. The first contribution of this paper, therefore, is to quantify the effects of monetary policy shocks on the movement of workers between employment, unemployment, and nonparticipation.

I find that the pattern of the responses of gross worker flows to contractionary monetary policy shocks is distinctive, and qualitatively similar to the pattern observed during most recessions (see Figure 1.1). In particular, the *EU* flow rises rapidly but the response is relative short-lived; the UE and UN flows decline much more slowly and stay lower for four years; and EN and NE decline modestly while *NU* rises. As discussed above, because fluctuations in the gross flows uniquely pin down changes in the stock measures, it is relatively straightforward to assess flows' relative contributions to stock variable fluctuations. I find that the decline in EP and LFP and the increase in UR after a contractionary shock are primarily accounted for by the *EU* flow and to a lesser extent by the UE and UN flows. Half of the magnitude of the peak responses of the stock variables is accounted for by the *EU* flow alone.

Motivated by these empirical results, I develop a model that embeds a Diamond-Mortensen-Pissarides (DMP) labor market in a New Keynesian sticky-price framework in which worker flows are well defined. The model includes endogenous job separations and

Figure 1.1: Labor Market Stocks and Flows, 1967–2017

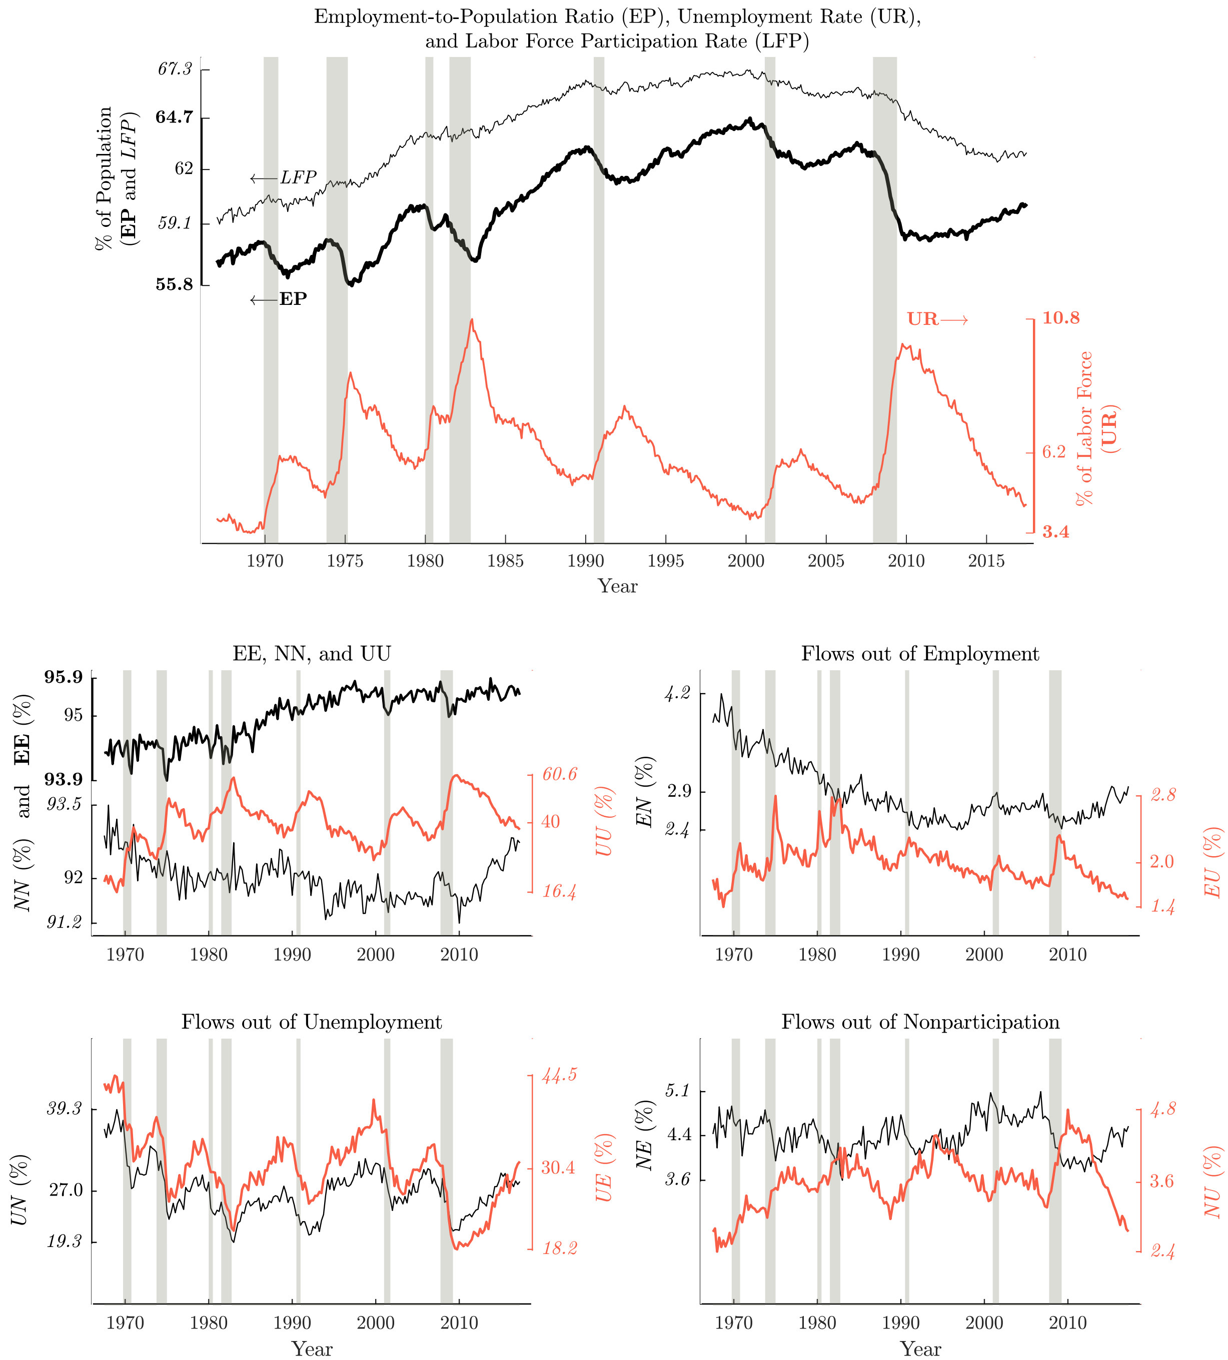

Stock (upper panel, monthly) and flow (lower panel, quarterly averages) measures, June 1967 to June 2017. Vertical axes indicate minimum, mean, and maximum values of these variables over the sample period. *Red* lines are measured on the right axis, while *black* and **bold black** lines are measured on the left axis. Shaded dates are NBER recessions.

labor force participation decisions. Workers differ in their degree of labor force attachment—that is, their propensity to exit the labor force. I show that this specific type of heterogeneity is necessary for the model to match the empirical flow responses. I verify that this type of heterogeneity is indeed present in the data.[3] Job loss due to monetary policy shocks produces changes in the *composition* of workers who are unemployed—in particular, the composition shifts towards workers with higher attachment—and this composition effect is an important driver of not only the flow variables, but also the stock variables.[4] The model is able to replicate qualitatively the responses of the labor market stocks and flows to monetary policy shocks that I estimate in the first part of the paper.

I then use the model to illustrate how a central bank that targets job loss in a simple Taylor-type rule is able to improve on welfare outcomes relative to either strict inflation targeting or a standard Taylor rule that targets both inflation and output or unemployment gaps. I find that the optimal simple Taylor-type rule targets inflation, the output gap, and the EU gap—that is the gap between EU flows and their efficient level. Moreover, the losses relative to this optimum from a simple rule that targets *only* the EU gap and inflation are negligible, while those from a rule that targets only the output gap and inflation, a commonly used Taylor rule in the literature, are significant.[5] I compare this optimal simple rule to other proposed rules in similar models that do not include some of the key features necessary to

[3] Specifically, after a contractionary monetary policy shock, a larger share of the stock of unemployed is made up by workers with low dropout propensities conditional on observable characteristics such as age, gender, marital status, and reason for unemployment.

[4] Although worker heterogeneity has been cited elsewhere as a source of persistence and propagation of shocks in the DMP model (e.g., Federico Ravenna and Carl E. Walsh (2012)), this particular source of heterogeneity, described in detail in Section 1.4, is a novel one.

[5] David Berger, Ian Dew-Becker, Konstantin Milbrandt, Lawrence D.W. Schmidt and Yuta Takahashi (2016), using different data sources and a different model, argue that the Fed should target layoffs; my paper also shows that it indeed *can* affect layoffs.

match the conditional moments I identify the in the empirical section of the paper. These rules, despite being optimal in similar models, deliver significantly worse outcomes than the optimal policy I derive.

More broadly, this paper contributes to at least three strands of the macroeconomic literature. The first is the literature on labor force flows. While previous papers, such as Robert Shimer (2012) and Michael W. L. Elsby, Bart Hobijn and Ayşegül Şahin (2015) have studied their unconditional moments, and Regis Barnichon and Christopher J. Nekara (2012) have explored their usefulness for forecasting labor market variables, the flows' conditional moments have yet to be explored. This paper is the first to examine the conditional moments of the full set of gross worker flows.[6] The methodology I introduce can easily be extended to study the effects on the labor market of other macroeconomic shocks; it can also provide a set of conditional moments that can be used to discipline a model in the spirit of Robert E. Lucas (1980) and Emi Nakamura and Jón Steinsson (2017).[7]

The second contribution is to the literature on the effects of monetary policy. To the extent that the empirical monetary policy literature has considered labor market effects directly, it has done so mostly by examining the responses of the stock measures only (EP, UR, and LFP). One notable exception is Helge Braun, Reinout De Bock and Riccardo DiCecio (2009), who examine the response to various shocks (including a monetary demand shock) of job separation and job finding rates; however, the exclusion of the participation

[6] Claudio Michelacci and David Lopez-Salido (2007) and Fabio Canova, David Lopez-Salido and Claudio Michelacci (2007) consider the effect of technology shocks on job creation, destruction, finding, and separation, but because they do not estimate the effects on *all* the flows, they are not able to construct decompositions to assess the relative importance of each flow in the responses of other labor market variables.

[7] Indeed, I use the empirical results on the effects of monetary policy shocks on worker flows for exactly this purpose.

margin from their analysis leads to qualitatively different results from those of this paper.

Finally, this paper contributes to the literature on optimal macroeconomic stabilization policy, specifically optimal monetary policy. The study of optimal policy in simple New Keynesian models like those described in Michael Woodford (2003) and Jordi Galí (2008) has evolved into finding optimal policy in settings with steady-state distortions along with various real and nominal frictions, as in Stephanie Schmitt-Grohé and Martin Uribe (2006) and 2007, or with a more realistic treatment of the labor market, such as in Ester Faia (2008), Jordi Galí (2011), and Federico Ravenna and Carl E. Walsh (2011). I find that simple rules that target the flow of workers from employment to unemployment improve welfare significantly relative to standard Taylor rules targeting output or unemployment gaps or strict inflation targeting.

The remainder of the paper is organized as follows. Section 1.2 describes the labor market flow data, from Shimer (2012) and Elsby, Hobijn and Şahin (2015), and monetary policy shock series from Christina D. Romer and David H. Romer (2004) I use throughout the paper. In this section I also discuss previous research that uses these detailed flow data series. Section 1.3 uses single-equation regressions and Romer and Romer's (2004) shock series to estimate the effects of monetary policy shocks on worker transition probabilities and conducts decompositions to quantify the relative importance of each flow. Section 1.4 describes the model, and Section 1.5 discusses its implications for optimal policy. Section 1.6 concludes.

1.2 Data

This section describes the construction of gross worker flow series from the CPS, as well as the estimation of Romer and Romer's (2004) monetary policy shock series. Alternative VAR-based shock identification strategies and the data used in these estimates are discussed in the appendix. The results are broadly consistent with the baseline estimates, which follow Romer and Romer (2004).

1.2.1 Measures of worker flows

The data on worker flows I use below are measures of transition probabilities based on monthly "gross flow" data from the CPS.[8] Approximately three-quarters of households interviewed as part of the CPS in a given month are re-interviewed the next month, which allows individuals' labor force states to be tracked, which in turn allows for the calculation of the total number of transitions among the three labor force states between months.

Examining these month-to-month counts directly, however, can potentially be misleading. If for example, someone is out of the labor force when interviewed in June, begins looking for work after being interviewed, and by the July interview has found a job, she would be counted as not in the labor force in June and employed in July, even though there was a period between the interview dates during which she would have been considered unemployed, as she did not have a job and was actively seeking one. Because the survey

[8] The quarterly flow data I use in the empirical work below are freely available on Robert Shimer's website, https://sites.google.com/site/robertshimer/research/flows. Data from January 1976 forward were constructed by Robert Shimer. For additional details, see Shimer (2012). Data from June 1967 to December 1975 were tabulated by Joe Ritter and made available by Hoyt Bleakley. Monthly flow data back to February 1990 are available from the Bureau of Labor Statistics (BLS) website, and back to June 1967 on Bart Hobijn's website, http://www.barthobijn.net.

interviews occur at discrete dates, the unadjusted gross flow data in this case "miss" a transition, counting the flow pattern N to U to E as simply N to E.

The data I use have been corrected for this time-aggregation problem following Shimer (2012), and therefore represent the probability of a transition from one state to another in a given month.[9] In what follows, I use the quarterly average of the monthly transition flow probabilities from 1967:Q2 to 2007:Q3. I use these years primarily because Romer and Romer's (2004) extended monetary shock series is available from 1969 to 2008, but also because the flow data for this period have been previously tabulated in a consistent manner by Shimer (2012). I use the quarterly averages because the flow data are noisy on a month-to-month basis.[10] Although not used in its entirety in the baseline estimates, I have constructed an extended series of the flow transition probabilities from 1967:Q2 to 2017:Q2, which is displayed in the bottom panel of Figure 1.1. The details of the construction of the data are discussed in Appendix A.1.

The obvious cyclical patterns in the flow data have been discussed at length elsewhere;[11] nevertheless a few such patterns are worth emphasizing here. During recessions, there is typically a short-lived spike in EU as workers are laid off, accompanied by relatively smaller declines in EE and EN as workers delay retirement or find periods of recession comparatively less attractive for non-market activities. Recessions also see a slow, hump-shaped decline in UE and a symmetric increase in UU, as unemployed workers are less likely to

[9]This adjustment from the monthly data as published by the BLS is discussed in detail in Appendix A.1.

[10]In the appendix, I use a high-frequency identification strategy along with the monthly flow data to relax both these restrictions. The results are essentially unchanged from the baseline.

[11]See, for example, Olivier J. Blanchard and Peter Diamond (1990), Shimer (2012), Elsby, Hobijn and Şahin (2015), and Per J. Krusell, Toshihiko Mukoyama, Richard Rogerson and Ayşegül Şahin (2016).

transition to employment; UN actually *declines* during a recession, despite the conventional wisdom that widespread discouragement of job seekers leads more of the unemployed to drop out of the labor force. As Elsby, Hobijn and Şahin (2015) point out, this is likely due to a compositional effect: as a recession progresses, a larger share of the pool of unemployed is made up by workers with greater labor-force attachment than in normal times, implying a lower overall UN transition probability.[12] In Section 1.3, I show this same composition effect occurs after a monetary policy shock. Along the participation margin, NE declines and NU increases during recessions, reflecting a relative increase in the probability of entering the labor force as unemployed conditional on a transition into the labor force.

1.2.2 Previous research on worker flows

Early research on worker flow data from the CPS focused on the technical problems involved in actually calculating the gross flows between labor force states. John M. Abowd and Arnold Zellner (1985) noticed substantial misreporting of labor force statuses in the CPS, leading to measured transitions that were not, in fact, occurring. They and James M. Poterba and Lawrence H. Summers (1986) proposed different correction methods for this problem; however, as noted in Shigeru Fujita and Gary Ramey (2009) and Elsby, Hobijn and Şahin (2015), although correcting for potential misclassification can affect the levels of the flows, it does not alter their fluctuations or relative contributions to stock variables.

Olivier J. Blanchard and Peter Diamond (1989) and 1990 examine the trends and cyclicality of the gross flows among employment, unemployment, and nonparticipation. In

[12]Note, however, that, despite the lower UN transition probability, because the actual number of unemployed workers is larger during a recession, this pattern is still consistent with the declines in labor force participation that occur during recessions.

their later paper, combining their analysis with other data on manufacturing employment, they characterize much of the cylical patterns discussed above. They also find that lower employment during recessions is due more to high rates of job destruction than low rates of job creation, while "booms" are the result more of low rates of job destruction than high rates of job creation. The question of the relative importance of worker flows in the cylicality of labor market variables was later debated with Robert E. Hall (2005a) and 2005b and Robert Shimer (2005b) and 2012 attributing nearly all of the rise of unemployment during downturns to declines in job finding, and Shigeru Fujita and Gary Ramey (2006), 2007, and 2009 attributing a much larger share of the rise in unemployment to the job separation margin. This strand of the literature has tended, however, to focus on two labor market states: employment and unemployment (or non-employment), abstracting from the participation margin. Elsby, Hobijn and Şahin (2015) decompose historical fluctuations in the unemployment rate into the component flows, and find that slightly more of the long-run variance in the unemployment rate is attributable to UE transitions than EU transitions; they also find that UN transitions are just as important as EU transitions. In contrast to these studies, I focus on *conditional* moments of worker flows.

Only recently have three-state models been developed explicitly to match the flows observed in the data. Per J. Krusell, Toshihiko Mukoyama, Richard Rogerson and Ayşegül Şahin (2011) are able to match the average long-run values of the flows in a steady-state equilibrium with persistent idiosyncratic shocks meant to represent events such as wage shocks or health shocks that generate long-time separation from or attachment to the labor force. In a later paper, Krusell et al. (2016) build a similar model that—via idiosyncratic productivity shocks, random job matchings and separations, indivisible labor, and incomplete markets—reproduces the cyclical patterns in the data noted above. The driver of the cycle in this model is a pattern of correlated aggregate shocks in partial equilibrium. In neither of these models, however, does monetary policy play any role.

Perhaps closest to this paper is Braun, De Bock and DiCecio (2009) who use a structural VAR to assess the impact on job finding and job separation rates of supply and demand shocks (including non-monetary and monetary demand shocks) identified via sign restrictions. Importantly, they do not consider the effects on labor force participation. Although measuring the complete response of worker flows to a monetary shock is not the aim of their paper, the absence of the labor force participation channel in their approach leads to conclusions on the relative importance of job finding and job separations that differ from my results accounting for labor force participation. Specifically, they find that nearly all of the increase in unemployment following a contractionary monetary policy shock is due to the decline in the job finding rate,[13] while I demonstrate that when one considers the full set of worker flows—including flows into and out of the labor force—jobs loss is a larger driver of employment and unemployment.

1.2.3 Monetary policy shocks

In Section 1.3, I use a measure of monetary policy shocks developed by Romer and Romer (2004) and extended through 2007 to estimate the response of worker flows. Romer and Romer (2004) identify monetary policy shocks as changes to the Federal Funds target rate

[13]Although they find that job separation rates contribute almost one-half of the impact effect of the shock on unemployment, the impact effect on unemployment is small and the relative contribution of job separations quickly dies out. I find that the increase in job separations alone, once one accounts for the participation margin, also results in a persistent, hump-shaped response of unemployment.

that are not predictable by the economic information in the Federal Reserve's "Greenbook" forecasts.[14]

The focus on monetary policy shocks is motivated by three observations. First, as shown in Olivier Coibion (2012), monetary policy shocks can account for a fairly large share of the historical fluctuations in the unemployment rate. Second, the long and well-established literature on monetary policy shocks encompasses a variety of methods for identification and estimation, which makes it a natural candidate for providing a set of "identified moments" for distinguishing between different models as suggested by Nakamura and Steinsson (2017). Finally, and related to the previous point, although empirical methods can be used to estimate the effects of *exogenous* changes in monetary policy, they cannot be used to directly estimate the responses of macroeconomic variables to *systematic* monetary policy changes; for that, a model is required. The implications of a model, such as optimal policy rules, might reasonably be thought to be more valid if that model is able to replicate the moments of interest that can be identified empirically. As Lucas (1980) put it, "The more dimensions on which the model mimics the answers actual economies give to simple questions, the more we trust its answers to harder questions."

1.3 Single-equation regressions

This section uses a flexible single-equation specification, making use of the shock series constructed by Romer and Romer (2004) extended through 2007, discussed in Section 1.2.3, to quantify the effects of monetary policy shocks on worker flows. The specification

[14] The exact identification strategy is discussed in the appendix. I also discuss alternative methods of shock identification.

below regresses the period-t value of the variable of interest on its own lags as well as lags of the monetary policy shock. It is identical to that used in Romer and Romer's (2004) estimation of the effect of a monetary policy shock on industrial production. If y_t denotes the dependent variable in time t, the equation to be estimated is

$$\Delta y_t = c + \sum_{j=1}^{J} \beta_j \Delta y_{t-j} + \sum_{i=1}^{I} \gamma_i \hat{s}_{t-i} + \epsilon_t, \tag{1.1}$$

where $\hat{s}_t$ is the value of Romer and Romer's (2004) monetary policy shock series in time t. The number of lags of the dependent variable and the shock in the estimation below are, respectively, $J = 8$ and $I = 12$. The sample period is 1969Q1 through 2007Q3.[15] As I show in the appendix, the results from estimating the impulse responses using Òscar Jordà's (2005) method of local projections are essentially identical.

Equation 1.1 is estimated independently for each variable of interest by ordinary least squares (OLS). The objects of interest are again impulse response functions. The cumulative response of the dependent variable one month after the shock is γ_1; two months after the shock it is $\gamma_1 + \gamma_2 + \beta_1\gamma_1$, and so on. Standard errors for the impulse responses can be constructed by drawing repeatedly from the asymptotic distribution of the coefficients estimated by OLS, computing the impulse responses for each draw, and taking the standard deviation of these simulated responses at each horizon.[16]

[15] These lag lengths are the quarterly equivalents of those in Romer and Romer's (2004) original estimation. The choice of sample period is dictated in part by the feasibility of the Romer and Romer (2004) shock estimation, which is not amenable to the ZLB period. In the appendix, I use high-frequency shock identification methods to include portions of the ZLB period. The results are essentially unchanged.

[16] Because the error bands of impulse responses implicitly test whether the response is different from zero, as discussed in Adrian Pagan (1984), the standard errors remain valid despite the presence of a generated regressor.

I estimate two sets of impulse responses: one for the stock variables and one for the flow transition probabilities. The stock variables I use are the employment-to-population ratio (EP), the unemployment rate (UR), and the labor force participation rate (LFP); the flow variables are the six off-diagonal transition probabilities.[17]

1.3.1 Results

The impulse response functions of the stock variables are displayed in the top row of Figure 1.2. EP and LFP fall, while UR increases, all with hump-shaped responses that peak about two years after the shock; EP and UR recover fully five years after the shock. The shock has a peak effect on EP and UR of more than six-tenths of a percentage point; LFP falls by only slightly more than 0.15 percentage point, but remains at that level throughout the five-year IRF horizon. Using their average values over the time period considered, this corresponds to a decline in EP from 60.4 to 58.8 percent, an increase in UR from 6.0 to 6.6 percent, and a decline in LFP from 65.0 to 64.9 percent. The responses of EP and UR are fairly large, while that of LFP is little more than rounding error.[18] As shown in Section 1.3.2, however, the almost negligible response of labor force participation does *not* imply that the participation channel can be innocuously ignored in the understanding of monetary policy's effects on the labor market.

The responses of the transition probabilities from the flow regressions are shown in the bottom two rows of Figure 1.2. Most noteworthy and statistically significant are the

[17] The diagonal transitions (i.e., EE, UU, and NN) are constructed as residuals from the other flows, so estimating them in addition is redundant.

[18] The use of Romer and Romer's (2004) shock series produces larger effects than those estimated using standard VAR approaches. See Coibion (2012).

Figure 1.2: Impulse Responses of Stocks and Flows to a Monetary Policy Shock

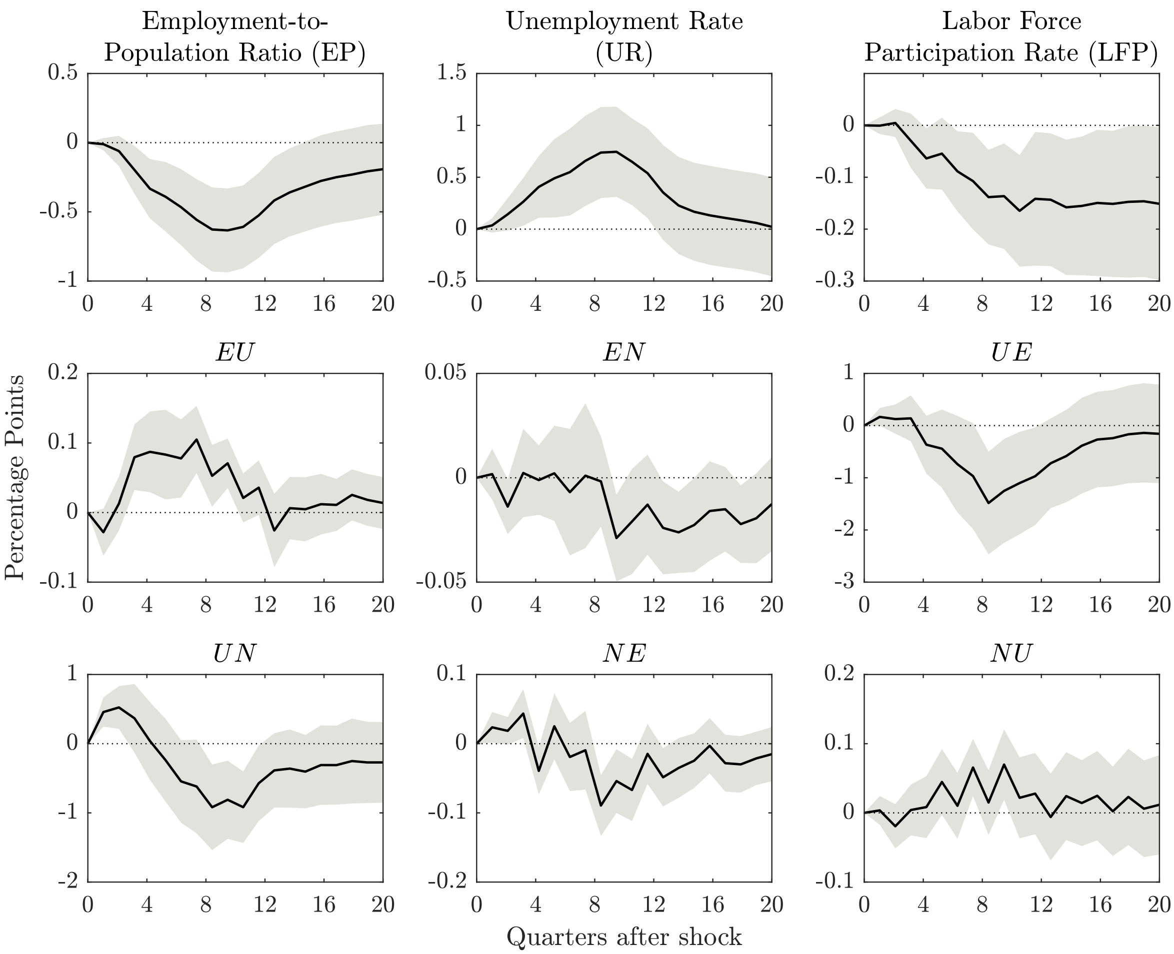

Impulse responses of stock (top row) and flow (bottom two rows) variables to a 100 b.p. contractionary Romer and Romer (2004) monetary policy shock, 1969Q1–2007Q3. Autoregressive Distributed Lag (ADL) model (Equation 1.1). Shaded areas are one standard deviation intervals from a bootstrap.

rapid increase in EU and subsequent steady decline, and the slow, hump-shaped declines of UE and UN. Qualitatively, these responses are similar to the typical cyclical pattern noted in Section 1.2.1 and displayed in Figure 1.1. Also similar to the typical cyclical pattern of the flow rates are the responses of EN and NE, both of which decline slightly, while NU increases; the responses of these three flows are at most only modestly significant.[19]

The response of the UN flow is worth discussing more, as it motivates portions of the structure of the model presented in Section 1.4. As mentioned above, Elsby, Hobijn and Şahin (2015) argue that the *unconditional* cyclical pattern in UN is caused by larger numbers of workers with high labor force attachment being driven into unemployment during recessions. Indeed, a similar phenomenon occurs *conditional* on a monetary policy shock. To assess this composition effect, I estimate impulse responses of the shares of unemployment made up by different groups of workers—specifically, for groups that have high labor force participation rates and low average UN transition rates over the entire sample period. The groups I consider are prime-age workers (ages 25-54), married workers, those seeking full-time employment, and job losers, i.e., those who report being unemployed because they lost an existing job.

Figure 1.3 displays the impulse responses of the share of unemployment made up by these groups, estimated from (1.1). After a contractionary shock, these workers make up a larger share of the unemployed, with most groups' response peaking at 1 percentage point. The share of unemployed made up by job losers increases by almost 2.5 percentage points. The average UN rate of job losers over this period is 13.3 percent, compared to

[19] Responses of similar shape and magnitude are obtained using Romer and Romer's (2004) shock series and Jordà's (2005) local projection method of estimating impulse responses, described in the appendix.

27.0 percent for all unemployed workers. A back-of-the-envelope calculation suggests that the composition effect of the 2.4 percentage-point increase in job losers alone can account for two-thirds of the decline in UN.[20] After a contractionary shock, the pool of unemployed shifts in composition towards workers with a *lower* propensity to exit the labor force, thereby driving the UN flow rate down.

Because of the difficulty in interpreting the importance of the impulse responses of the flow variables, in the next section I construct decompositions of the stock impulse responses into the underlying flows.

1.3.2 Decompositions

The quantitative importance of the effects of a monetary shock on flow variables is unclear from the impulse responses of the transition probabilities alone. The peak response of EU is an increase of 0.1 percentage point, which is approximately a 4% increase relative to its average value over the sample period. Similarly, the peak declines of UE and UN are also around 3-4% of their average values. To assess the impact of a monetary shock on these flows in terms of more familiar variables, one can exploit the fact that the flows sum to the stocks to calculate the *stock* impulse responses that would occur if only specific flows responded to the shock.

I do this in a manner similar to Shimer's (2012) assessment of the historical contribu-

[20] This calculation considers the 2.4 percentage-point increase in job losers as a share of the unemployed and applies the average UN transitions over the full time sample. Ideally, one would estimate the transition rates for each group directly; unfortunately, small sample sizes prevents doing this at monthly or quarterly frequencies (the full set of flows for all workers, even before examining different groups, shrinks the sample sizes to one-twelfth of the full CPS sample).

Figure 1.3: Impulse Responses of Unemployment Shares to Monetary Policy Shocks

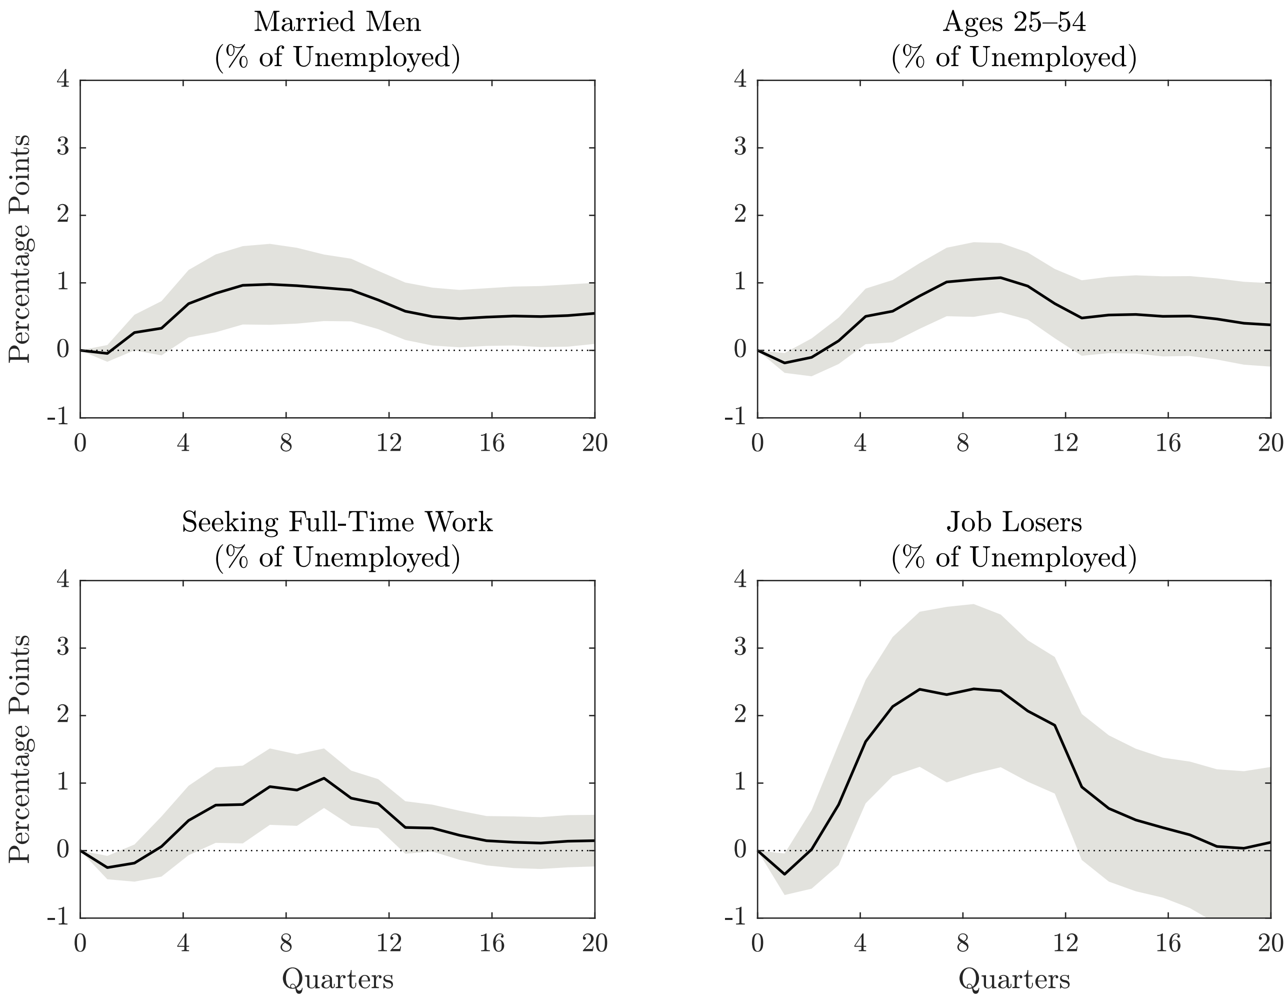

Impulse responses of the *share* of unemployed workers made up by various groups to a 100 b.p. contractionary Romer and Romer (2004) monetary policy shock, 1969Q1–2007Q3. Autoregressive Distributed Lag (ADL) model (Equation 1.1). Shaded areas are one standard deviation intervals from a bootstrap.

tion of each flow probability to the unemployment rate. Specifically, I shut down the effect of the shock on particular flows by fixing them at their average values throughout the horizon of the impulse response. Other flows move according to the responses estimated from (1.1). I then re-calculate the implied stock responses from this new pattern of flow responses with one or more flows fixed at their average values. These counterfactual responses answer the question, "What would the stock impulse response look like if *only* certain flows responded to the shock?"[21]

The first set of decompositions I consider looks at the contribution of each flow individually to the responses of the three stock variables. The results are presented in Table 1.1. Each row within a panel of the table displays, at yearly horizons, the value of the impulse response of the stock variable if *only* the indicated flow responded to the shock. The final column visually depicts the same information across all horizons. Within each panel, each column sums to the total response of the stock variable at that horizon. The importance of the EU flow is immediately evident. It contributes the most to the decline in EP for nearly the entire horizon; the most to the increase in UR for the first two years, after which it contributes almost equally to UE; and the most to the decline in LFP two to four years out. The UE flow contributes to the responses of these variables for only the later periods of the IRFs. The decomposition of stock variables into the individual flows demonstrates that

[21] These exercises are similar to those conducted by Christopher A. Sims and Tao Zha (2006) and Ben S Bernanke, Mark Gertler and Mark Watson (1997) to assess the importance of monetary policy's endogenous responses to macroeconomic shocks, and more recently by Rüdiger Bachmann and Eric R. Sims (2012) to understand the role of confidence in the transmission of government spending shocks to output; however, because a given time path of the flows "adds up" to that of the corresponding stocks, the single-equation approach here is both simpler and "model-free." Constructing counterfactual IRFs in a structural VAR setting requires the researcher to take a stronger stand on the relevant interactions with other variables in the economy.

the EU flow is the most important driver of labor market variables' responses to monetary policy.

It should also be noted that the offsetting forces of the many of the flows are evident in the decomposition of EP and, especially, LFP. As discussed in Section 1.3, the small response of LFP to a monetary policy shock masks the importance of participation; the flat response of the stock is the outcome of large, but offsetting flow responses.

Other decompositions can also be considered. The top and bottom panels of Figure 1.4 display the decomposition of stock variables into, respectively, inflows and outflows. Similar to Table 1.1, the red counterfactual IRFs in each column sum to the total response (in black). As above, these counterfactuals also depict the response of stock variables if only these flows responded to the shock. Inflows into unemployment ($EU + NU$) account for about half the decline in EP (more at early horizons, less at later) and nearly all of the increase in UR, while inflows to employment contribute roughly equally to EP and nothing to UR. Outflows from employment ($EU + EN$) account for most of the decline in EP and more than half of the increase in UR; outflows from unemployment produce very little change in *any* stock variable. In these alternative decompositions, it is inflows to unemployment and outflows from employment that drive most of the response of labor market variables to monetary policy shocks. These "grouped" decompositions are both consistent with job loss being the primary driver of these responses.[22]

Figure 1.5 shows the contribution of flows into and out of the labor force. While the majority of the response of EP is accounted for by EU and UE flows, participation decisions

[22] In the last columns, the offsetting contributions of flows to the response of LFP is again evident.

Table 1.1: Contribution of Individual Flows to Stock Impulse Responses

Stock	Flow	Contribution of flows to stock IRFs at each horizon (Years) 1	2	3	4	5	IRF
EP ratio	*EU*	**-0.13**	**-0.30**	**-0.29**	**-0.17**	-0.12	
	EN	-0.03	-0.07	-0.02	0.03	0.04	
	UE	-0.03	-0.15	-0.23	-0.17	-0.11	
	UN	-0.01	0.01	0.07	0.09	0.08	
	NE	-0.01	-0.08	-0.14	-0.12	**-0.12**	
+	*NU*	0.02	0.07	0.13	0.12	0.08	
=	*Total*	*-0.20*	*-0.53*	*-0.48*	*-0.23*	*-0.16*	
Unemp. Rate	*EU*	**0.18**	**0.28**	0.16	0.04	0.02	
	EN	0.00	0.01	0.00	0.00	-0.01	
	UE	0.04	0.16	**0.17**	**0.06**	**0.02**	
	UN	-0.02	0.07	0.09	0.02	0.00	
	NE	0.00	0.01	0.02	0.02	0.02	
+	*NU*	0.06	0.12	0.09	0.00	-0.03	
=	*Total*	*0.27*	*0.68*	*0.56*	*0.14*	*0.02*	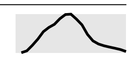
LFP	*EU*	-0.01	**-0.13**	**-0.20**	**-0.16**	-0.11	
	EN	**-0.03**	-0.07	-0.02	0.03	0.04	
	UE	-0.01	-0.05	-0.13	-0.14	-0.11	
	UN	-0.02	0.06	0.14	0.11	0.08	
	NE	-0.01	-0.08	-0.13	-0.12	**-0.12**	
+	*NU*	0.06	0.16	0.20	0.13	0.07	
=	*Total*	*-0.03*	*-0.11*	*-0.14*	*-0.15*	*-0.15*	

Notes: **Bold** numbers indicate largest contributions, red numbers indicate gross contribution in the opposite direction of stock IRF. Rows in each columns sum to totals. Rightmost column depicts the full counterfactual impulse response. See Section 1.3.2 for details.

Figure 1.4: Contribution of Inflows and Outflows to Stock Impulse Responses

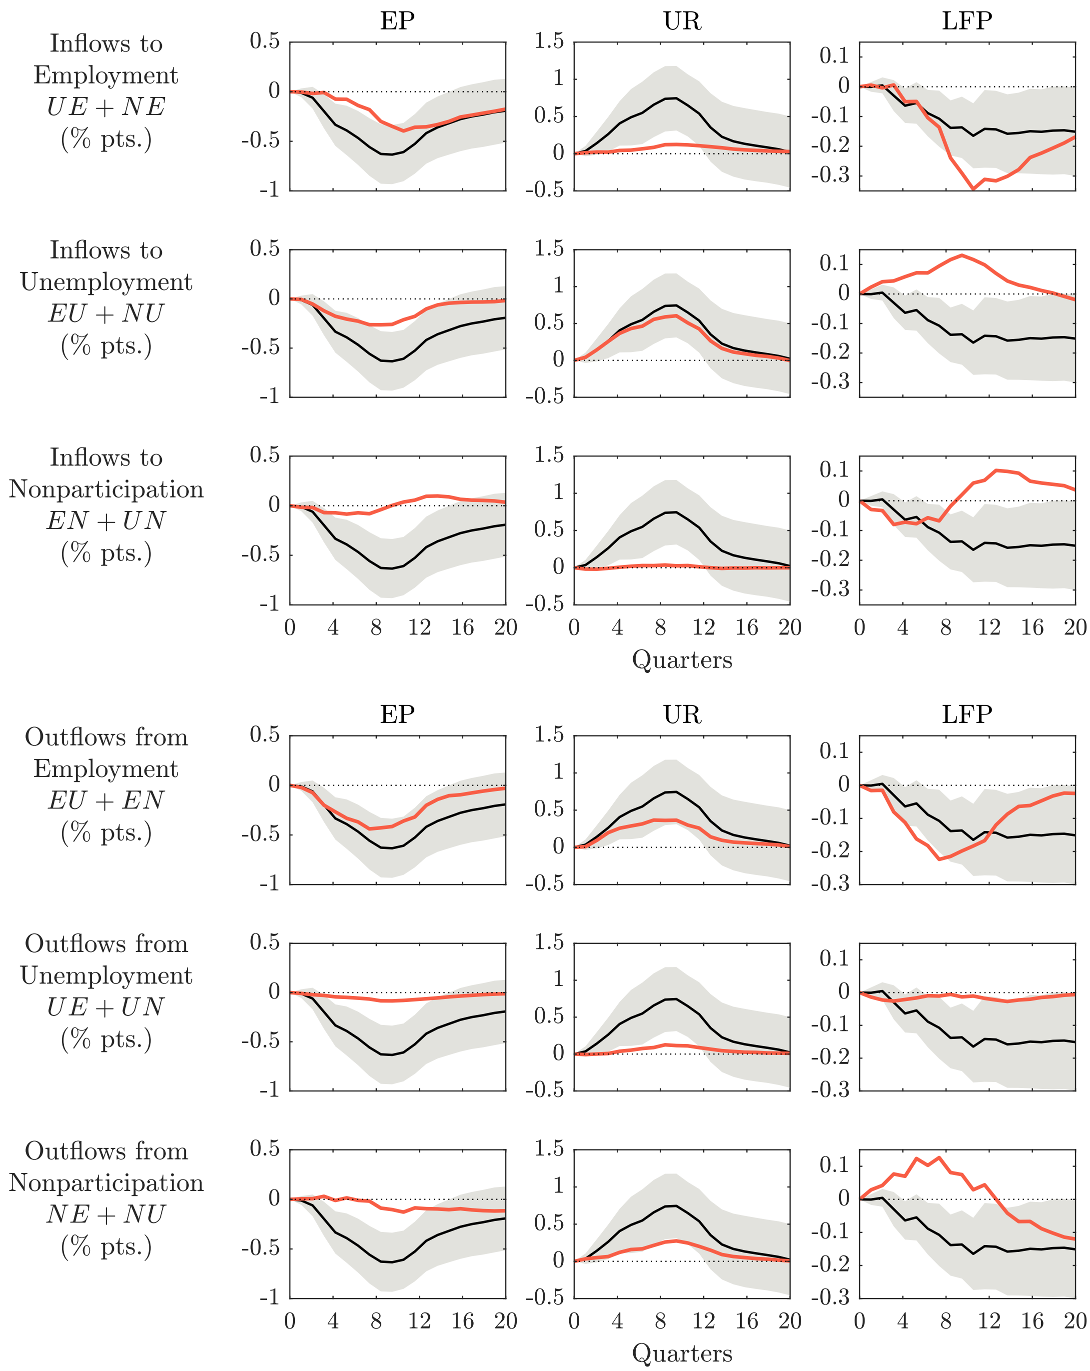

Contribution of inflows (upper panel) and outflows (lower panel) to stock variable responses to a 100 b.p. Romer and Romer (2004) monetary policy shock. In each row, red lines are counterfactual responses in which only those flows are responding. Within each pannel, the red lines in each column sum to the **black** lines (baseline stock IRFs). See Section 1.3.2 for details.

play a larger role in the response of UR, accounting for about one third of its increase in response to contractionary monetary policy shocks. The two groups of flows have partially offsetting effects on LFP.

The last decomposition I consider divides the flows into two groups: the first group, $EU+EN+UN$, are labeled "separation"; the second, $UE+NE+NU$, are labeled "finding." This division is motivated by the results from Section 1.3 that the response of UN is largely driven by a compositional change in the pool of unemployed workers driven by job loss. I therefore group it with separations from employment. I group NE and NU with UE because their responses measure the rate of transitioning to employment conditional on entering the labor force. The results are displayed in Figure 1.6. The "separations" group accounts for most of the decline in EP and half the increase in UR. "Finding" only contributes to the EP decline at the end of the response. As with the other decompositions considered above, these two groups have partially offsetting effects on LFP.

The decompositions considered in this section point to three conclusions. First, job loss—regardless of whether one considers the EU flow alone, inflows into unemployment, outflows from employment, or EU, EN, and UN all together—is the most important driver of the labor market's response to monetary policy. Second, flows into and out of the labor force account for roughly one third of the response of UR to monetary policy shocks. And, finally, flows tend to have partially offsetting effects on LFP; a focus on the stock variable alone masks large responses in the underlying flows, indicating that monetary policy shocks can have large effects on participation *decisions* despite its modest impact on overall LFP.

The first result—that job loss drives the labor market's response to monetary policy shocks—is particularly striking considering the importance of the job finding rate in

Figure 1.5: Contributions of Participation Flows to Stock Impulse Responses

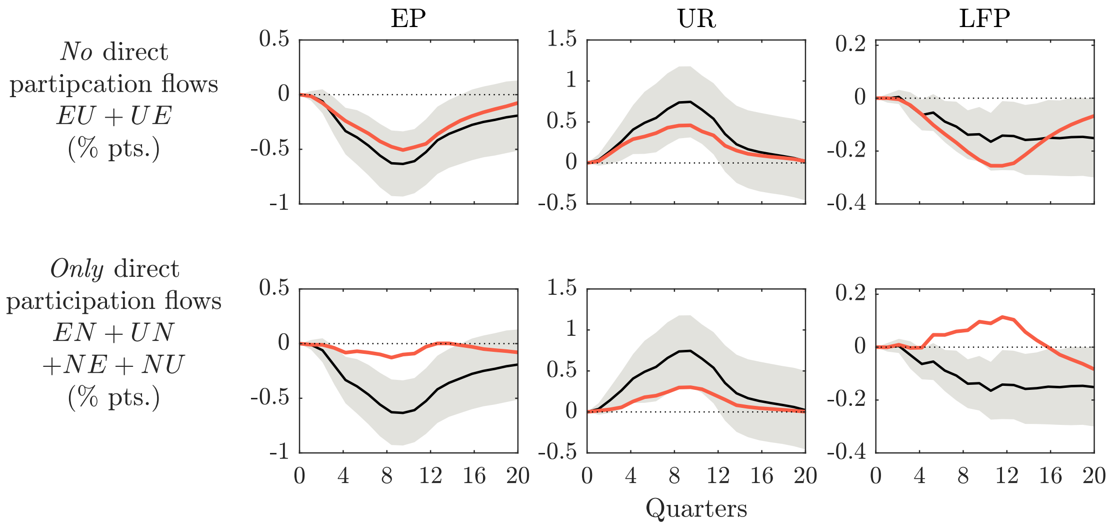

Actual and counterfactual IRFs. *Upper panel*: Only flows between E and U ($EU + UE$) respond. *Lower panel*: Only direct participations flows ($EN + UN + NE + NU$) respond. Shaded areas are one standard deviation bootstrap intervals. (Equation 1.1).

accounting for *unconditional* labor market moments.[23] My results are not necessarily in conflict with that literature, however, as I focus on *conditional* moments. The results that flows into and out of the labor force contribute to about one third of the increase in UR after a contractionary monetary policy shock and that flows have offsetting effects on LFP are, however, in line with the unconditional decompositions described in Elsby, Hobijn and Şahin (2015).

The results from this section motivate the construction of a model that can accurately replicate these conditional moments in order to conduct policy experiments. The structure of the model I describe below is informed by the results on the importance of job loss,

[23]For discussions on the importance of the job finding channel, see, for example, Hall (2005*a*), Hall (2005*b*), Shimer (2005*b*), and Shimer (2012).

Figure 1.6: Contribution of Job Separation and Job Finding to Stock Impulse Responses

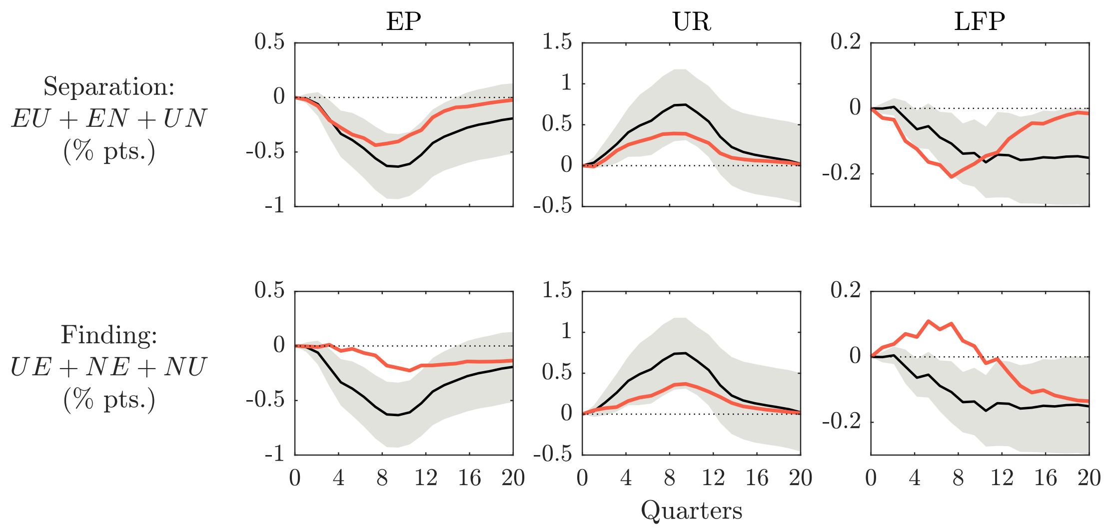

Actual and counterfactual IRFs. *Upper panel*: Only job separation ($EU + EN + UN$) responds. *Lower panel*: Only job finding ($UE + NE + NU$) responds. Shaded areas are one standard deviation bootstrap intervals. (Equation 1.1).

participation decisions, and composition effects in accounting for both the stock and flow responses to monetary policy shocks. The model I build therefore includes a search model of the labor market with endogenous separations, labor supply decisions that include a nontrivial role for nonparticipation, heterogeneity in labor force attachment, and a role for monetary policy to affect the real economy. The next section describes this model in detail and discusses its policy implications.

1.4 Model

The model consists of a representative household made up of two types of workers indexed by $i \in \{h, \ell\}$ and representative firms. *Ex ante*, the two types of workers differ only according to their participation in the labor market. Workers of each type are subject to search frictions and idiosyncratic productivity shocks. "Wholesale" firms produce using labor of each type and sell their output in competitive product markets to monopolistically competitive "retail" firms that are subject to price stickiness.

1.4.1 Households

The setup of the household sector follows Monika Merz's (1995) "large family" construct, in which workers are able to perfectly insure their consumption against idiosyncratic shocks. The household receives utility from an aggregate consumption good C_t and leisure from nonparticipants. Each worker is either employed (E), unemployed and searching for work (U) or not in the labor force (N). Each period, each non-employed worker of type i draws a nonparticipation utility x_t from a distribution with c.d.f. $G^i(x)$. The household receives wages from employed workers, which depends on workers' match-specific productivity drawn from $F^i(a)$ when employed, and an unemployment benefit from unemployed workers, and trades in risk-free nominal bonds. Both productivity and nonparticipation utility shocks are assumed to be serially uncorrelated.

The household solves the following problem:

$$\max_{C_t, B_t, x_t^{i*}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ u(C_t) - \sum_i \gamma^i \left[\chi E_t^i - (1 - E_t^i) \int_{x_t^{i*}}^{\infty} v(\chi)\, dG^i(\chi) \right] \right\}$$

$$\text{s.t. } P_t C_t + Q_t B_t = D_t + B_{t-1} + P_t T_t,$$

and the laws of motion of the labor market, described below, where γ^i is the measure of type i workers, B_t is the risk-free nominal bond with price Q_t equal to the inverse of the gross nominal interest rate, T_t are transfers, including lump-sum taxes and firm profits, and P_t is the aggregate price level.[24] Consumption C_t is a Dixit-Stiglitz aggregate across retail goods. Importantly, x_t^{i*} is the reservation utility for workers of type i, the utility above which they do not participate in the labor market. D_t is total labor market income given by

$$D_t = P_t \sum_i \gamma^i \left[bU_t^i + \frac{E_t^i}{1 - F^i(a_t^{i*})} \int_{a_t^{i*}}^{\infty} w_t^i(\alpha)\, dF^i(\alpha) \right],$$

where b is the (real) unemployment benefit, which is assumed to be independent of worker type, $w_t^i(a)$ is the wage function of type-i workers with match-specific productivity a, and a_t^{i*} is the threshold productivity level, below which employed workers separate endogenously from their matches.[25]

1.4.2 Labor market

The labor market is characterized by DMP-style search frictions. At the start of any period workers are subject to idiosyncratic separation and utility shocks; at the end of any period, a worker is in one of three states: employment, unemployment, and nonparticipation.

At the start of period t, unemployed workers are matched with firms according to an aggregate matching function $m(U_{t-1}^i, V_{t-1}^i)$, which matches unemployed workers to vacant firms.[26] After these matches occur, all matched workers of type i—including both these newly

[24]Note that utility to the household from nonparticipation is simply a transformed equation for the expected value of the utility draw conditional on its being above the threshold x_t^{i*} times the number of nonparticipants.

[25]As will be seen below, because of Nash-bargained wages, this threshold will be the value such that the total match surplus is zero.

[26]The timing assumption implies that searching workers and firms are only matched and able to produce in the period *after* they begin their search. The timing convention follows Michael U. Krause and Thomas A. Lubik (2010), who study parameter regions for determinacy in this class of models.

matched workers and workers employed from the previous period—separate with exogenous probability δ^i. Those who survive this separation draw a new idiosyncratic productivity from $F^i(a)$ and they separate *endogenously* if this draw is sufficiently low. The law of motion for employment of type-i workers can, therefore, be written as

$$E_t^i = (1-\delta^i)(1-F^i(a_t^{i*}))\left[E_{t-1}^i + f_{t-1}^i U_{t-1}^i\right],$$

where $f_t^i \equiv f(\theta_t^i)$ is the probability a type-i unemployed worker is matched with a vacant firm, which is a function of $\theta_t^i \equiv \frac{V_t^i}{U_t^i}$.

After exogenous and endogenous separations occur, all unmatched workers draw nonparticipation utilities from c.d.f. $G^i(x)$. Workers who get a sufficiently high draw exit (or remain out of) the labor force, and the household gets their utility realization but forgoes the opportunity of forming a match next period. Workers who draw a low nonparticipation utility become (or remain) unemployed, and the household forgoes their utility draw, obtains the unemployment benefit b, and has the opportunity to be matched next period. Unemployment and nonparticipation in period t are, therefore, given by

$$U_t^i = G^i(x_t^{i*})(1-E_t^i),$$

$$N_t^i = (1-G^i(x_t^{i*}))(1-E_t^i).$$

The model uniquely pins down all the flows described in Section 1.2 with the exception of NE.[27] All flows vary endogenously via the matching function and endogenous thresholds a_t^{i*} and x_t^{i*}.

[27] As discussed above, although in reality a very large number of workers enter employment *directly* from nonparticipation, this flow is of minor importance in the transmission of monetary policy shocks. Therefore, the model presented here abstracts from this mechanism.

1.4.3 Firms

Firms are owned by the household. The wholesale firm produces using only labor, which is assumed to be perfectly substitutable across types.[28] It hires labor by posting vacancies for each type of labor (V_t^i) in separate matching markets. Wholesale output is linear in labor, and the firm sells its output to retail firms at relative price $\frac{P_t^w}{P_t}$. Let $\mu_t \equiv \frac{P_t}{P_t^w}$ denote the retail markup over wholesale prices. Firms maximize profits discounted by the household's stochastic discount factor.

The wholesale firm's time-t problem can be written as

$$\max_{V_t^i} \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{\lambda_{t+j}}{\lambda_t} \sum_i \left[E_{t+j}^i \left(\frac{\widetilde{A}_{t+j}^i Z_{t+j}}{\mu_{t+j}} - \widetilde{w}_{t+j}^i \right) - \kappa^i V_{t+j}^i \right],$$

$$\text{s.t.} \quad E_{t+j}^i = (1-\delta^i)(1-F^i(a_t^{i*})) \left[E_{t+j-1}^i + q_{t+j-1}^i V_{t+j-1}^i \right],$$

where $q_t^i \equiv q(\theta_t^i)$ denotes the probability a type-i vacancy is matched with a searching worker, λ_t is the household's marginal utility of consumption, and $\widetilde{w}_t^i \equiv \mathbb{E}\left[w_t^i(a) | a \geq a_t^{i*}\right] = (1 - F^i(a_t^{i*}))^{-1} \int_{a_t^{i*}}^{\infty} w_t^i(\alpha)\, dF^i(\alpha)$, and $\widetilde{A}_t^i \equiv \mathbb{E}\left[a | a \geq a_t^{i*}\right]$ is evaluated similarly. Z_t is an aggregate labor productivity shock and κ^i is the flow vacancy cost. The assumption of perfect substitutability of the wholesale good produced by different types of labor ensures that μ_t does not depend on i.

[28] Perfect substitutability and linear production assures that the "one firm, many workers" assumption made here is identical to a "one firm, one worker" assumption.

1.4.4 Wage setting

Wages are determined by period-by-period generalized Nash bargaining, which gives rise to the surplus sharing rule

$$\eta_t \mathcal{S}_t^{iF} = (1-\eta_t)\mathcal{S}_t^{iH},$$

where $\eta_t \in (0,1)$ is the time-varying bargaining power of workers, and $\mathcal{S}_t^{iF}$ and $\mathcal{S}_t^{iH}$ denote the match surplus to the firm and household, respectively.[29] Shocks to the worker's bargaining weight enter as "cost-push" shocks to the New Keynesian Phillip's Curve and are the source of inefficient fluctuations in the optimal policy experiment in Section 1.5.2.

1.4.5 Retail firms and monetary policy

There is a continuum of monopolistically competitive retail firms, indexed by j, that transform the wholesale good into the retail good according to the production function

$$Y_t^r(j) = \sum_i Y_t^i(j).$$

Retail firms are subject to staggered price setting as in Guillermo A. Calvo (1983); in each period retail firms can reset their prices with probability $1-\phi$. These firms maximize expected discounted profits subject to price stickiness and a demand function arising from the Dixit-Stiglitz consumption aggregator. Thus, the retail firm's problem is identical to that of the standard New Keynesian model[30] with linear production and real marginal cost equal to $1/\mu_t$; therefore, a detailed derivation is omitted.

[29] The surplus in the Nash bargain is given by the marginal value of employed workers in the household's and firm's problems.

[30] Described, for example, in Galí (2008).

Finally, in the baseline model, the central bank is assumed to follow a Taylor rule of the following form:[31]

$$\beta(1+i_t) = (\beta(1+i_{t-1}))^{\rho_i} \left[(1+\pi_t)^{\phi_\pi} \left(\frac{Y_t}{Y^{ss}} \right)^{\phi_Y} \right]^{1-\rho_i} \epsilon_t^{mp},$$

where i_t is the nominal interest rate, π_t is inflation, Y^{ss} is steady-state output of the final good, $\phi_Y > 0$ and $\phi_\pi > 1$, $\rho_i \in [0, 1)$ is an interest-rate smoothing parameter, and ϵ_t^{mp} is a monetary policy shock.

1.4.6 Parameterization

The model as described above is fairly general and flexible, accommodating potentially many types of workers differing in their average productivities and drawing idiosyncratic shocks from several combinations of distributions. While most parameters can be calibrated to standard values from the literature, there is, however, a degree of freedom in choosing the distributions of idiosyncratic shocks, $G^i(x)$ and $F^i(a)$, and the degree to which types differ from one another. In order to the keep the model as simple as possible, I make stark assumptions on these differences.

I focus exclusively on the case with two types of workers, denoted "high" and "low." The conceptual mapping between the model and the data is that high-type workers correspond to prime-age workers, those between 25 and 54, while low-type workers correspond to those aged 16-24 or over 55. In the U.S., almost exactly half of the civilian noninstitutional population over 16 is between the ages of 25 and 54, so I assume equal measure of each type of worker in the household.[32]

[31] Section 1.5.2 considers alternative simple Taylor rules

[32] The mapping of worker types to age groups follows the approach taken in Ravenna and Walsh (2012).

Furthermore, because the labor force participation rate of workers between 25 and 54 is very high in the U.S. (above 80 percent), while that of the other group is roughly half that (44.6 percent in April, 2017), I assume that *only* low-type workers are able to receive utility from nonparticipation. This difference in labor force attachment (low-type workers choose whether or not to participate in the labor market, while high types alway participate) is the only difference between the two types of workers.

Separating the participation decision across worker types in this way will ensure that the model produces cyclical variation in the composition of the unemployed, which was shown to respond to monetary policy shocks above in Section 1.3, and which Elsby, Hobijn and Şahin (2015) argue is an important driver of the cyclicality of the flow of workers from unemployment to nonparticipation. In particular, after a contractionary monetary policy shock, the value of employment falls, which in turn may decrease the threshold value of nonparticipation x_t^*, causing low type workers to exit the labor force. Thus, after a contractionary shock, a larger share of the pool of unemployed is made up by high-type workers, just as in the data.[33]

The distributions $F^{\ell}(a) = F^h(a)$ are assumed to be Type II Pareto distributions.[34] This distribution has two attractive features. First, the distribution of income is often estimated to have a Pareto tail. Second, as illustrated in Figure 1.7, the elasticity of separations with respect to the (endogenous) idiosyncratic productivity threshold a_t^{i*} is *higher* for lower values of the threshold. This gives the model a reasonable chance to match Mueller's (2017)

[33] Andreas I. Mueller (2017) argues that the pool of unemployed workers shifts towards high-productivity, high-wage workers during recessions, while I focus on compositional changes with respect to labor force attachment.

[34] That is, Pareto distributions shifted to have support on $[0, \infty)$.

observation that the separation rate of prime-age workers' is lower but responds *more* to the business cycle than those of other age groups. I choose $F^i(\cdot)$ to have an unconditional mean of 1 and variance of 3.[35]

Given this functional form for the idiosyncratic productivity shocks, other parameters are chosen to match U.S. labor market data. Utilities from nonparticipation for low types are drawn from $G^\ell(x)$, which is assumed to be a uniform distribution with support on $[\underline{x}, \bar{x}]$. The lower bound is assumed to be zero, and the upper bound is chosen to match the U.S. labor force participation rate (see below). Because the idiosyncratic shocks are serially uncorrelated, it is necessary to have low values in the support of these distributions to induce endogenous flows into and out of the labor force and endogenous separations.[36] Because high-type workers do not receive nonparticipation utility draws, the exact functional form of $G^\ell(\cdot)$ matters little; the first and second moments are what matter for participation decisions in the model.

Finally, I let $u(C_t) = \log(C_t)$ and $v(x) = \psi \log(x)$, and the matching function is assumed to be the standard, Cobb-Douglas form:

$$m(U_t, V_t) = mU_t^\alpha V_t^{1-\alpha},$$

so that the vacancy matching rate is given by

$$q(\theta_t) = m\theta_t^{-\alpha}$$

[35]This implies a scale parameter of 2 and a curvature parameter of 3, so the CDF is given by $F^i(x) = 1 - \left(1 + \frac{x}{2}\right)^{-3}$.

[36]If, for example, the lower bound of the support of idiosyncratic productivity shocks was sufficiently high, workers would never separate endogenously, implying a constant EU flow.

Table 1.2: The Calibration of Parameters in the Baseline Model.

Variable	Description	High Value	Low Value
β	Discount factor	0.99	0.99
δ^i	Separation rate	0.02	0.02
α	Matching function elasticity	0.5	0.5
χ	Employment disutility	0.266	0.266
ψ	Nonparticipation utility coeff.	N/A	2
b	Unemployment benefit	0	0
κ	Vacancy cost	0.2	0.2
η	Worker Share in Nash bargain (S.S.)	0.5	0.5
$\bar{x}$	Maximum utility from nonparticipation	N/A	2
ϕ	Calvo parameter	0.67	0.67
ϵ	Elasticity of substitution	6	6
ϕ_π	Taylor rule coeff. on inflation	1.5	1.5
ϕ_Y	Taylor rule coeff. on output	0.125	0.125
ρ_i	Interest rate smoothing	0	0

Notes: The column labeled "High Value" displays the parameter values assigned to high-attachment workers, while the column labled "Low Value" displays the values assigned to low-attachment workers.

and $f(\theta_t) = \theta_t q(\theta_t) = m\theta_t^{1-\alpha}$.

Many parameters are calibrated to values standard in the literature; this facilitates straightforward welfare comparisons across comparable models discussed in Section 1.5.2. The full parameterization is presented in Table 1.2.

The parameters of the distributions of the idiosyncratic shocks are chosen so that the steady state stock variables match their counterparts in the data. However, because high-attachment workers in the model always participate, I target the *relative* participation rate of prime-age workers to non-prime-age workers. Between 1990–2017, the labor force participation rate of prime-age workers averaged 83 percent while that of non-prime-age workers averaged 45 percent. To match the relative participation rates, I therefore target a steady-state participation rate for low types of 54 percent. The matching efficiency parameter m is chosen to match a target unemployment rate. Following Antonella Trigari (2009), because labor force participation is higher in the model than in the data, I target a broader measure of unemployment than the official unemployment rate. Specifically, I target the U-6 unemployment rate, which includes not only the unemployed, but also marginally attached workers and those working part time for economic reasons. This measure has an average value of 10.6 percent since 1994, the first year in which it was published.

1.4.7 Monetary policy shocks

The baseline model is solved by a second-order approximation around a zero-inflation, non-stochastic steady state.[37] The experiment I first consider is a monetary policy shock that raises the (annualized) nominal interest rate by 100 basis points on impact.

After a contractionary shock, labor market tightness θ_t^i falls in both markets. This reduces the job-finding rates f_t^i but increases the vacancy-matching rates q_t^i since there are more unemployed workers. The match surplus is lower for a firm (because value of being vacant is higher due to the higher q_t^i). Wages fall and, because wholesale prices are flexible, so does relative price of the labor produced intermediate good. This is the marginal cost of the retails firm, so the markup of prices over marginal cost μ_t rises. Prices must fall to return to the desired level of markups, producing deflation. Lower wages and job-finding rates reduce the value of employment—and, consequently, unemployment—causing low-type workers to exit the labor force. Lower match surplus increases the threshold for endogenous

[37] A second-order approximation is required for the welfare calculations in Section 1.5.2.

Figure 1.7: The Mechanism for Endogenous Separations Following a Monetary Policy Shock

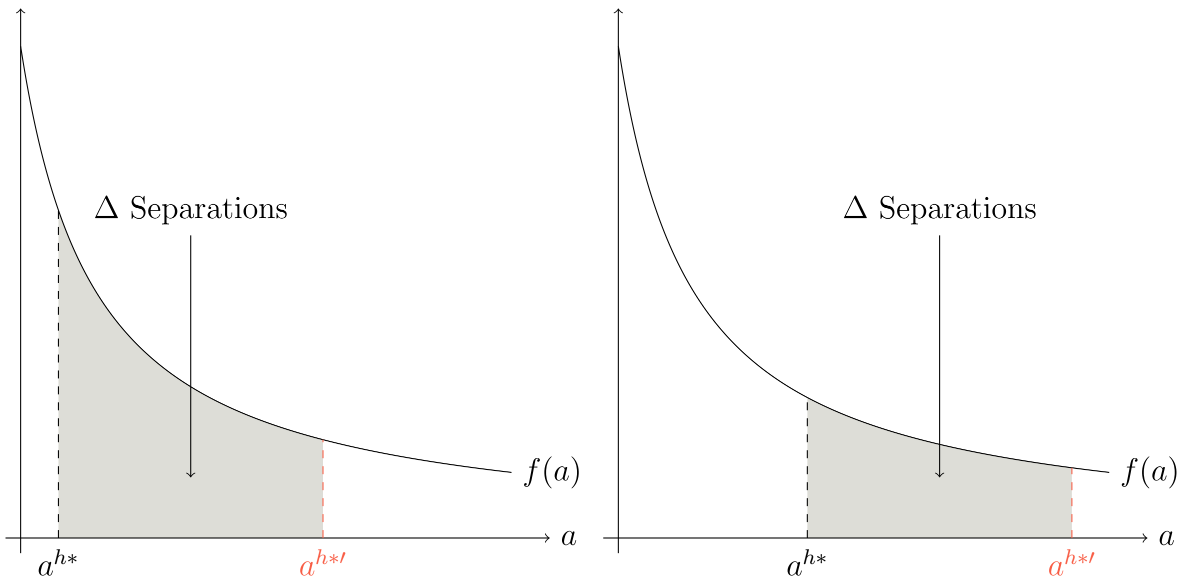

The separations resulting from a given increase in the idiosyncratic productivity thresholds a^* for a lower initial value (left) and a higher initial value (right), under the assumption that the productivities are drawn from a Type II Pareto distribution with density $f(a)$. The separations resulting from the same increase in the threshold are higher for a lower initial value of a^*.

separations for both types. Because these thresholds differ for each type of worker, responses of these thresholds—even of the same magnitude—imply changes in separations of differing magnitudes across types. Figure 1.7 illustrates this mechanism. These separations combined with low types exiting the labor force cause a change in the composition of the unemployed.

The model impulse responses for stock variables—EP, UR, and LFP—and the empirical estimates from Section 1.3 are presented in Figure 1.8. EP and LFP fall while UR rises, matching, in a qualitative sense, their empirical counterparts. The magnitude of the responses are also roughly equivalent, although somewhat smaller in the model than in the

Figure 1.8: Impulse Responses of Stock Variables—Data vs. Model

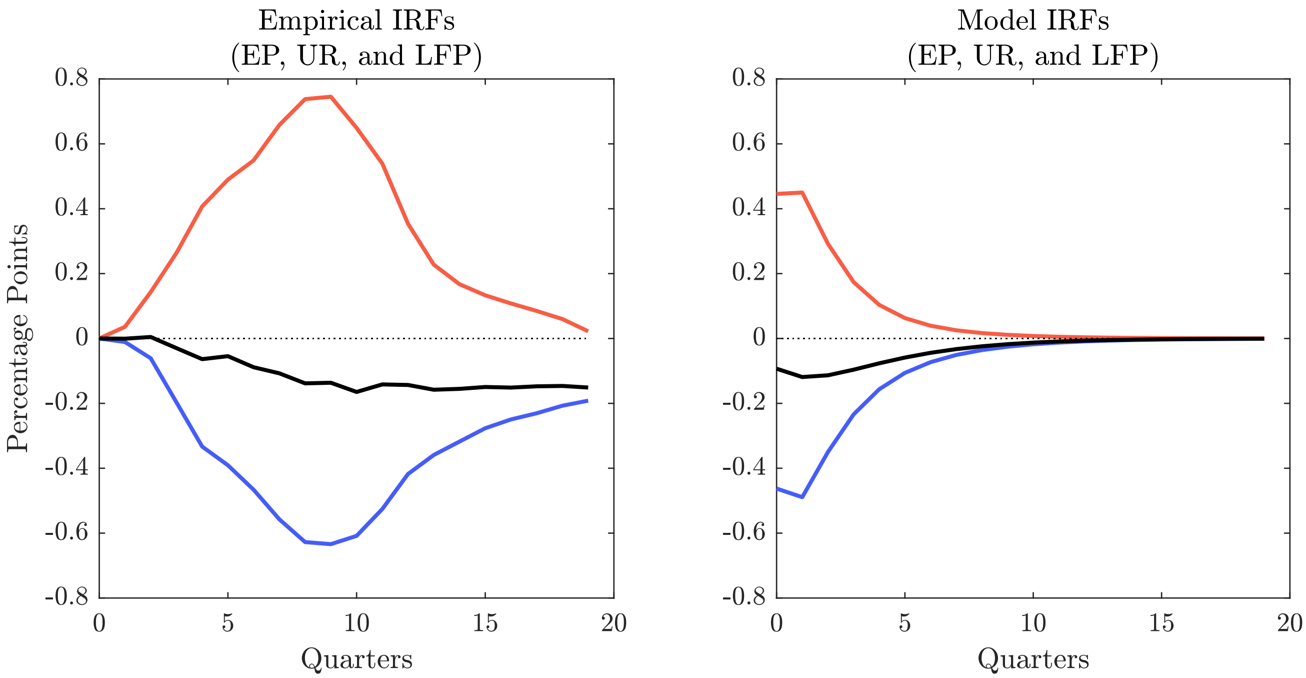

Model impulse responses of the EP, UR and **LFP** to a contractionary monetary policy shock that raises the nominal interest rate by 100 basis points. Percentage-point deviation from steady state.

data, at roughly two-thirds the peak magnitudes.[38]

The model impulse responses for the flow variables EU, EN UE, and UN are displayed in Figure 1.8. All match, qualitatively, their empirical counterparts. EU increases and returns quickly to steady state, while UE and UN decline and are more persistent and hump-shaped, an unsurprising result since these two measures are driven by compositional changes. The EN response is modest. The initial increase in UN (also present in the data)

[38] As is well known, the very persistent and hump-shaped empirical responses are difficult to match without a variety of additional frictions, as discussed in, for example, Lawrence J. Christiano, Martin Eichenbaum and Charles Evans (2005).

is the initial exit from the labor force of unmatched low-type workers. As seen in Figure 1.10, the share of the unemployed made up by high types increases in response to the shock, as does the share of high-type vacancies among all vacancies. This shift in the makeup of vacancies increases the volatility of labor market tightness in the low-attachment sector, which in turn impacts the cyclicality of low-attachment workers' participation in the labor market.

This highlights a new role for heterogeneity in contributing to the persistence of monetary policy shocks. Although Ravenna and Walsh (2012) emphasize the role of heterogeneity in their model's persistent response to shocks, the mechanism is slightly different. In their model, the unemployment pool shifts systematically toward low-productivity workers in responses to a contractionary shock, reducing firms' incentives to post vacancies. In the model presented here, on the other hand, compositional shifts in the pool of unemployed induce firms to *shift* their vacancies toward different types of workers. After a contractionary shock, high-type workers make up a larger share of all unemployed workers, increasing the probability of a vacancy being matched with a worker in the high-type market relative to the low-type market. Persistence in this model comes from the resulting long-lasting employment matches of firms and high-type workers, rather than reduced incentives to post vacancies as in Ravenna and Walsh (2012).

1.5 Welfare

The baseline model is able to replicate some of the key moments identified in the empirical section of the paper. It produces reasonably accurate responses to monetary policy shocks of both the stock and flow variables in the labor market. In this section I use the

Figure 1.9: Impulse Responses of Flow Variables—Data vs. Model

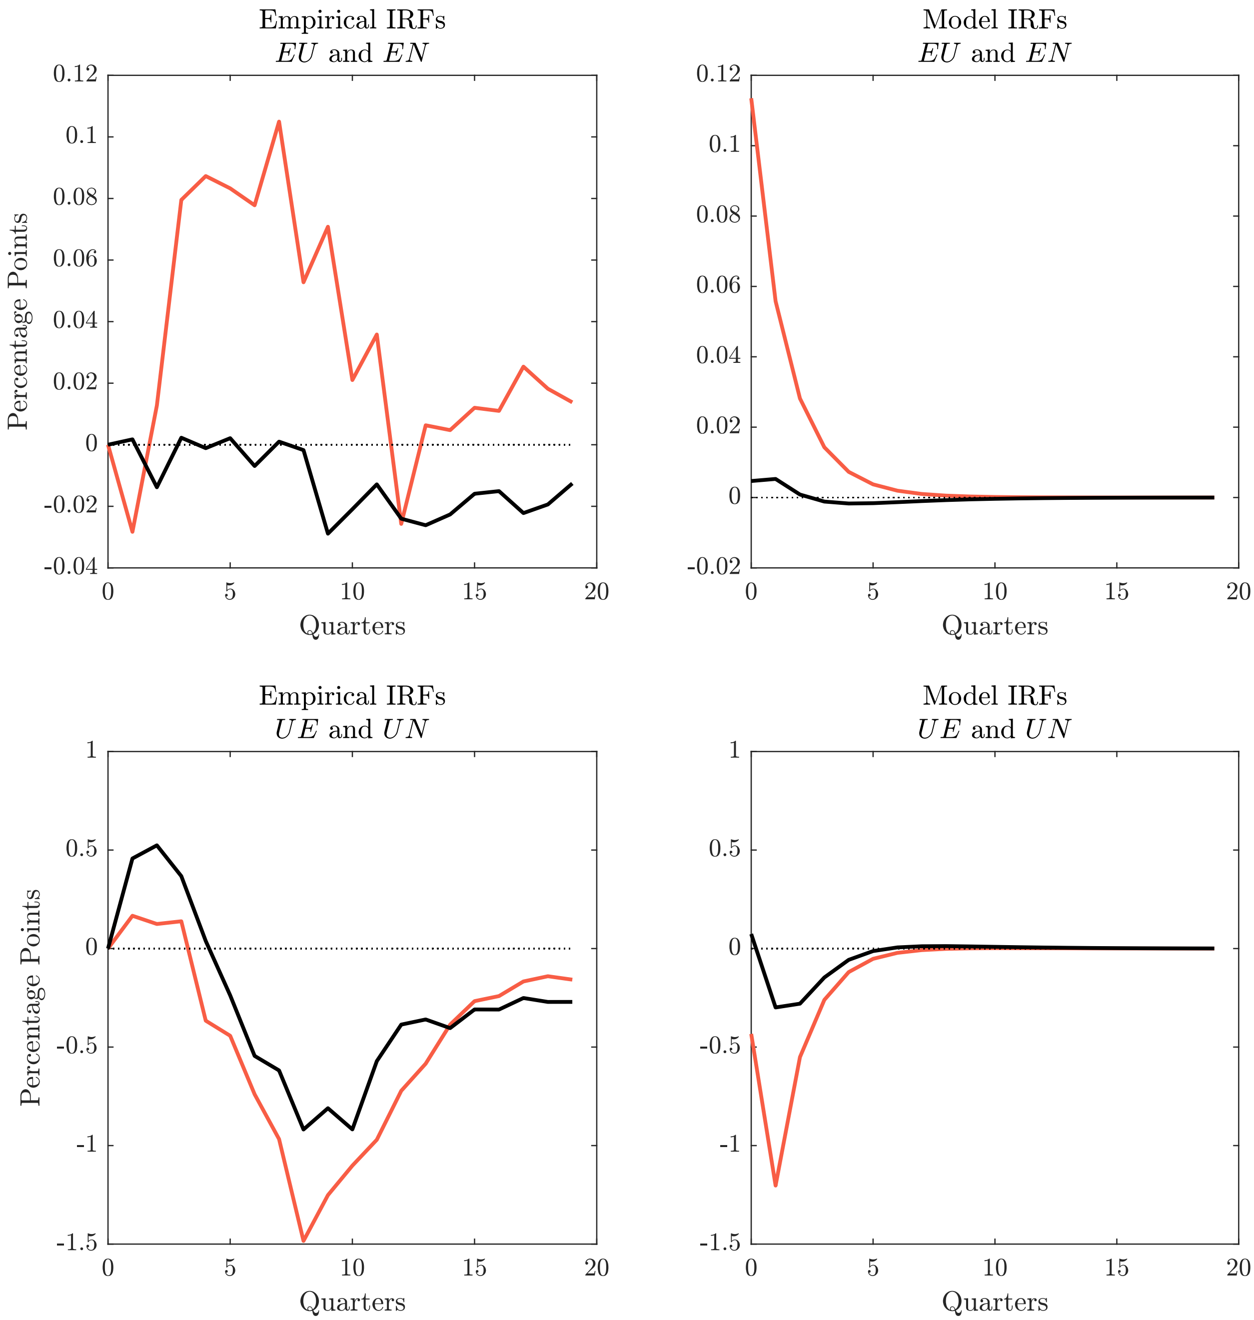

Top row: Empirical and model impulse responses of EU and EN flows to a contractionary monetary policy shock that raises the nominal interest rate by 100 basis points. Percentage-point changes. *Bottom row*: Empirical and model impulse responses of UE and UN.

Figure 1.10: Model Impulse Responses of Unemployment Shares

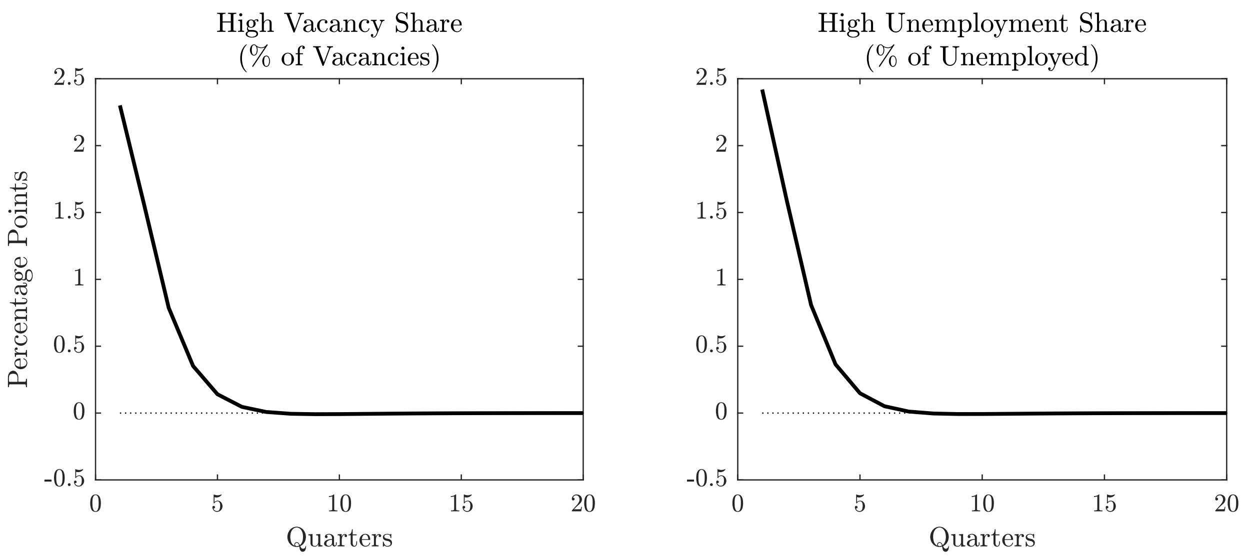

Model impulse responses of the share of the unemployed made up by high-type workers and the share of vacancies made up by high-type vacancies to a contractionary monetary policy shock that raises the nominal interest rate by 100 basis points. Percentage-point deviation from steady state.

model as a normative tool to solve for optimal monetary policy within a class of simple Taylor-type rules of the form described in Section 1.4.

1.5.1 Sources of inefficiency

There are two sources of inefficiencies in the steady state of the model. The first, a familiar source in the labor search literature, occurs if the worker's share in the Nash bargain differs from the elasticity of the matching function with respect to vacancies (i.e., $\eta \neq \alpha$). This is the well-known Arthur J. Hosios (1990) condition. The second is inefficiency low output do to monopolistic competition, which occurs in standard New Keynesian models in the absence of a production subsidy.

If these two conditions are met in the steady state, and if the only shocks are to aggregate productivity (Z_t), it is straightforward to show that a strict inflation targeting rule is able to replicate the social planner's solution and achieve first best. Following Ravenna and Walsh (2011), I introduce exogenous shocks to η, the worker's share in the Nash bargain. These shocks are meant to capture some of the inefficiencies introduced from deviations from the Hosios condition without taking a stance on the particular type of wage rigidity or bargaining differences present. Because the Hosios condition holds in the steady state, however, these shocks are more similar to the wage rigidities in Hall (2005*a*) or Olivier J. Blanchard and Jordi Galí (2010) than to Marcus Hagedorn and Iourii Manovskii's (2008) alternative calibration of Robert Shimer's (2005*a*) model.

1.5.2 Wage bargaining shocks and welfare

This section considers the economy described above in an efficient steady state that is hit with shocks to productivity Z_t and η_t the worker's bargaining weight in wage determination and compares welfare across simple Taylor-type rules. I calibrate the variance and persistence of these shocks to values suggested in Ravenna and Walsh (2011), where the standard deviations of innovations are 0.32 percent for productivity, and 3.87 percent for the bargaining weight, in terms of deviations from steady state.

I compare the baseline Taylor rule with a number of alternatives. To do so, I approximate the model to second order around the non-stochastic steady state and simulate the model for 10,000 periods. I compare alternative policies according to the share of average consumption[39] the household would be willing to forgo (or would need to receive) in order to be indifferent between the policies. Specifically, letting variables with tildes denote the values in the solution to a social planner's problem and those without denote the competitive equilibrium values, I solve for the value of Λ that satisfies

$$\begin{aligned}\widetilde{\mathcal{V}}_0 \equiv \sum_{t=0}^{\infty} \beta^t \left\{ u(\widetilde{C}_t) - \sum_i \gamma^i \left[\chi \widetilde{E}_t^i - (1-\widetilde{E}_t^i) \int_{\tilde{x}_t^{i*}}^{\infty} v(\chi)\, dG^i(\chi) \right] \right\} \\ = \sum_{t=0}^{\infty} \beta^t \left\{ u((1+\Lambda)C_t) - \sum_i \gamma^i \left[\chi E_t^i - (1-E_t^i) \int_{x_t^{i*}}^{\infty} v(\chi)\, dG^i(\chi) \right] \right\}.\end{aligned}$$

Under the assumption of log utility from consumption, it is straightforward to show that

$$\Lambda = \exp\left\{ \left(\widetilde{\mathcal{V}}_0 - \mathcal{V}_0 \right) (1-\beta) \right\} - 1,$$

[39] Because the model is solved by a second-order approximation, certainty equivalence fails, and average consumption does not equal consumption in the deterministic steady state.

where $\mathcal{V}_0$ denotes the time-zero expected present discounted value of utility under the competitive equilibrium allocation. The value of Λ is the share of average equilibrium consumption the household would have to be given to be indifferent between the decentralized equilibrium (under a given policy rule) and the social planner's constrained efficient allocation.

1.5.3 Optimal simple rules

I first find the optimal simple Taylor-type rules. This exercise is equivalent to computing a Ramsey-optimal policy in which the policymaking instrument is restricted to be a class of simple Taylor-type reaction functions to variables within the model. I consider four classes of Taylor rules. Each allows for interest-rate smoothing and include inflation targets,[40] in addition to the following variables:

1. Output or unemployment gaps.[41]
2. The employment-to-unemployment flow (EU) gap.
3. Output or unemployment gaps *and* the EU gap.
4. *Changes* in output or unemployment.

The first of these classes of rules is the standard Taylor-type rule. The second and third are motivated by the empirical findings from Section 1.3 that monetary policy affects the labor

[40]Reactions to inflation are necessary to ensure that there exists a unique equilibrium.
[41]That is, the difference between their equilibrium and efficient levels.

market primarily through the *EU* flow. The last class of rules are motivated by Schmitt-Grohé and Uribe's (2006) observation that a desirable aspect of policy is that it be based on *observables*, as opposed to gaps of variables relative to some unknown steady-state or efficient value.

The results are displayed in the upper panel of Table 1.3. The column titled "Absolute loss" displays the loss relative to the efficient allocation expressed as a percent of average consumption ($100 \times \Lambda$, in the notation of the previous section). The "Relative loss" column displays the percentage loss relative to the optimal simple rule. The optimal rule responds modestly to inflation, and reacts to both the output gap and the *EU* gap. The best rule that targets *only* the *EU* gap features losses of about two-thirds of one percent relative to the optimal rule, while the best rule that targets *only* the output gap has losses of more than 20 percent relative to the optimal simple rule.[42]

I also compare the welfare losses across other Taylor-type rules commonly used in the literature. These include Taylor rules that respond to inflation and output gaps (relative to both steady state and the efficient level of output), a regime of strict inflation targeting, and optimal simple rules derived in the models of Faia (2008) and Galí (2011), both of which feature labor market frictions. These results are displayed in Table 1.4. A few patterns emerge from Tables 1.3 and 1.4. First, interest-rate smoothing can be welfare improving for sub-optimal rules, but the optimal rules of various forms feature no interest-rate smoothing. Second, the optimal rules derived in models only slightly different from the one described

[42] As is typically the case, the losses due to deviations from the efficient allocation are relatively small; the optimal rules of each form all involve losses of less than two-tenths of one percent of average per-period consumption. There are, however, substantial differences across rules relative to the optimum.

Table 1.3: Optimal Simple Rules and Associated Welfare Losses

Rule	ϕ_π	ϕ_Y	ϕ_U	ϕ_{EU}	ρ_i	Abs. Loss (%)	Rel. Loss (%)
Y-gap	1.1	2.18	–	–	0	0.0182	20.52
U-gap	1.09	–	-3.3	–	0	0.0175	15.89
ΔY	1.06	3.3	–	–	0	0.019	25.83
ΔU	1.01	–	-0.81	–	0.2	0.0189	25.17
Y-gap, EU-gap	1.1	0.56	–	-0.33	0	0.0151	0
U-gap, EU-gap	1.01	–	0	-0.17	0	0.0152	0.66
EU-gap	1.01	–	–	-0.17	0	0.0152	0.66

Notes: Optimal simple rules and associated welfare losses compared to the efficient allocation ("Abs. Loss") and relative to the optimal simple rule ("Rel. Loss"), expressed as percentages of average consumption (see Section 1.5.2 for details).

above deliver outcomes worse than even simple, commonly used Taylor rules. Finally, aside from ensuring determinacy, rules that feature a strong inflation response deliver inferior welfare outcomes. Indeed, a regime of strict price stability performs the *worst* of all the alternative rules considered. The optimal rules all feature inflation responses on the low end of the region of determinacy.

This last result is particular striking, since a strong inflation response is a common feature of *optimal* policy in other models. In the optimal simple rule in Schmitt-Grohé and Uribe (2006), the coefficient on inflation is 3, while that on output is 0.01. The optimal policies derived in Faia (2008) and Galí (2011) both involve strong inflation responses, and Ravenna and Walsh (2011) find that strict price stability is very close to the fully optimal policy with commitment. Ravenna and Walsh's (2011) model is almost identical the one presented above, but it does not feature endogenous separations, labor force participation, or heterogeneity. The relative simplicity of their model facilitates an algebraic derivation

Table 1.4: Alternative Simple Rules and Associated Welfare Losses

Rule	ϕ_π	$\phi_{Y^{ss}}$	ϕ_Y	ϕ_U	ρ_i	Absolute Loss (%)	Relative Loss (%)
Taylor (SS)	1.5	0.125	–	–	0	0.0258	70.86
Taylor (SS)	1.5	0.125	–	–	0.8	0.0220	45.70
Taylor (SS)	1.5	0.5	–	–	0	0.0238	57.62
Taylor (SS)	1.5	0.5	–	–	0.8	0.0216	43.05
Taylor (gap)	1.5	–	0.125	–	0	0.0256	69.54
Taylor (gap)	1.5	–	0.125	–	0.8	0.0218	44.37
Taylor (gap)	1.5	–	0.5	–	0	0.0232	53.64
Taylor (gap)	1.5	–	0.5	–	0.8	0.0209	38.41
Strict inflation	∞	–	–	–	0	0.0268	77.48
Faia (2008)	3	–	0	-0.15	0	0.0265	75.50
Galí (2011)	1.51	–	-0.1	-0.025	0	0.0264	74.83

Notes: Alternative simple rules and associated welfare losses compared to the efficient allocation ("Abs. Loss") and relative to the optimal simple rule ("Rel. Loss"), expressed as percentages of average consumption (see Section 1.5.2 for details).

of the optimal policy using a linear-quadratic approach, but the absence of the key model features described above leads to a very different policy implication. Indeed, a policy that is nearly optimal in their model ranks among the worst of the policies I consider.

While these results highlight the importance of model specification in deriving optimal policies, another way to approach the question is to ask how these rules perform when the model outcomes are ranked according to a *different* welfare criterion. Table 1.5 displays the losses from two alternative loss functions: the quadratic loss functions from the textbook New Keynesian model and Ravenna and Walsh's (2011) simple search model. The loss function from these models all involve parameters and endogenous variables that have counterparts the baseline model. Specifically, I simulate the model using a given rule and compute the welfare losses implied by the loss functions from these other models. Strict inflation targeting delivers the best outcome for *both* of these alternative loss functions, while the actual optimal policy ranks very near the bottom for each alternative measure. This highlights the importance of comparing policies using model-consistent welfare measures.

1.6 Conclusion

The experience of the U.S. labor market during the Great Recession has highlighted the importance of a full characterization of labor market dynamics in understanding the business cycle. In this paper, I examine the effects of monetary policy shocks on three labor market states—employment, unemployment, and nonparticipation—and the flows of workers among them. A close examination of these labor market dynamics reveals the importance of job loss in understanding the effects of monetary policy on the labor market. Decompositions of labor market variables reveal that the flow of workers from employment to unemployment (EU) is the largest contributor to the increase in the unemployment rate and declines in the employment-to-population ratio and labor force participation rate after a contractionary monetary policy shock. Other decompositions lead to the same conclusion: job loss drives the labor market's response to monetary policy.

These decompositions also demonstrate the important role played by participation decisions. Although labor force participation responds relatively little to monetary policy, this is the outcome of large but offsetting responses of worker flows into and out of the labor force. Of particular interest is the response of flows from unemployment to nonparticipation, which are driven by a composition effect. After a contractionary monetary policy shock, job loss drives workers with high labor force attachment (low dropout propensities) into the pool of unemployed, lowering the overall U-to-N transition rate. In addition, transitions into and

Table 1.5: Losses from Alternative Welfare Measures

Rule	ϕ_π	$\phi_{Y^{ss}}$	ϕ_Y	ϕ_U	ϕ_{EU}	ρ_i	Loss NK	Loss R&W
Strict inflation	∞	–	–	–	–	0	0.014	0.5053
Opt. Y-gap	1.1	–	2.18	–	–	0	0.6142	0.5112
Opt. U-gap	1.09	–	–	-3.3	–	0	1.3226	0.515
Opt. ΔY	1.06	–	3.3	–	–	0	0.5026	0.51
Opt. ΔU	1.01	–	–	-0.81	–	0.2	0.4624	0.5093
Opt. Y-gap, EU-gap	1.1	–	0.56	–	-0.33	0	1.7645	0.5221
Opt. U-gap, EU-gap	1.01	–	–	0	-0.17	0	1.7741	0.5219
Opt. EU-gap	1.01	–	–	–	-0.17	0	1.7741	0.5219
Galí (2011)	1.51	–	-0.1	-0.025	–	0	0.1683	0.5076
Faia (2008)	3	–	0	-0.15	–	0	0.0484	0.5061
Taylor (SS)	1.5	0.125	–	–	–	0	0.1848	0.5076
Taylor (SS)	1.5	0.125	–	–	–	0.8	0.1472	0.5070
Taylor (SS)	1.5	0.5	–	–	–	0	0.3251	0.5086
Taylor (SS)	1.5	0.5	–	–	–	0.8	0.2512	0.5076
Taylor (gap)	1.5	–	0.125	–	–	0	0.159	0.5075
Taylor (gap)	1.5	–	0.125	–	–	0.8	0.1223	0.5068
Taylor (gap)	1.5	–	0.5	–	–	0	0.2123	0.5077
Taylor (gap)	1.5	–	0.5	–	–	0.8	0.1344	0.5067

Notes: Welfare losses from alternative models. "Loss NK" is the implied loss from the simple New Keynesian model described in Galí (2008), while "Loss R&W" is the loss from Ravenna and Walsh's (2011) model, which includes labor market frictions, but no heterogeneity, endogenous separations, or participation. Both "Loss NK" and "Loss R&W" are in units of utility. Strict inflation targeting is the best simple-rule policy by both of these criteria, while it ranks *worst* among simple rules by the model-consistent welfare criterion.

out of the labor force account for one third of the increase in the unemployment rate after a contractionary monetary policy shock. These results show that participation decisions also play an important role in the labor market's response to monetary policy.

A model designed to match these empirical conditional moments features a search-and-matching labor market with endogenous separations and nontrivial participation decisions, heterogeneity in labor force attachment, and sticky prices to allow monetary policy to affect the real economy. Policy experiments in this model suggest that a central bank can target job loss (or EU transitions) in a simple Taylor-type rule to achieve better welfare outcomes in response to macroeconomic shocks than under either strict inflation targeting or a standard Taylor rule that targets both inflation and output. This policy rule is attractive not only for its theoretical simplicity, but also in a practical sense since the data on job loss are available from multiple sources and at high frequencies. Layoffs, therefore, are a measure of the labor market that the Federal Reserve can target to achieve better outcomes, while staying within its mandate to promote full employment.

Obviously, the questions this paper addresses are not motivated solely by theoretical curiosity; they have arisen precisely because these issues have become increasingly important in real-world monetary policy decisions. Federal Reserve Chairs Janet Yellen and Ben Bernanke and other Federal Open Market Committee members have repeatedly made reference to worker flows—in particular, the relative contributions of flows into and out of the labor force to changes in the unemployment rate—when discussing the efficacy of monetary policy in internal deliberations, policy speeches, press conferences, and congressional testimony. The results of this paper demonstrate that a solid foundation for understanding the effects of monetary policy on labor market flows is not merely an academic pursuit, but rather a real-world necessity.

Chapter 2

Has Monetary Policy Accelerated Job Polarization?

2.1 Introduction

Job polarization refers to the "hollowing out" of middle-skill, routine occupations that has occurred in the United States since 1980.[1] Over the last 40 years, jobs that involve routine tasks (e.g., assembly line workers, data entry technicians, office administrators) have been replaced with those involving nonroutine tasks (e.g., management and personal services), as illustrated in Figure 2.1. Autor, Levy and Murnane (2003) and others have argued that this trend is due to routine-biased technological change (RBTC). Routine tasks are amenable to automation, and, in recent decades, employers have substituted away from labor and toward capital to accomplish these tasks.

As Jaimovich and Siu (2014) show, and as is evident from Figure 2.1, the majority of job losses in these occupations since the mid-1980s occurred during, just before, or after the last three NBER recessions. They link this phenomenon to the so-called "jobless recoveries" that accompanied these recessions. They also highlight that the decline was *not* driven primarily by changes in industry composition.

The long-term trend decline in routine employment is even more evident in Figure 2.2,

[1] This phenomenon is described in detail in Daron Acemoglu (1999), David H. Autor, Frank Levy and Richard J. Murnane (2003), Nir Jaimovich and Henry E. Siu (2014), Christopher L. Foote and Richard W. Ryan (2015), and others.

Figure 2.1: Per Capita Employment in Routine Jobs

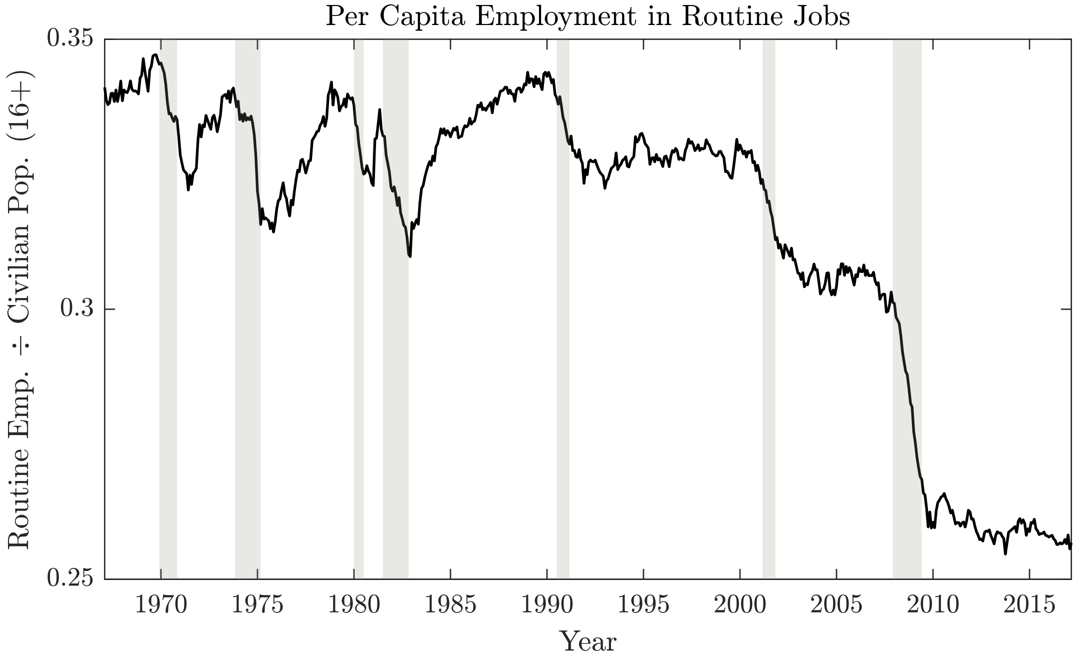

Per capita employment in routine occupations, 1967-2017. Occupational employment data from the Current Population Survey. Shaded dates indicate NBER recessions.

which depicts how the relative decline in routine-task employment has translated into the *share* of total employment in routine jobs. In February 1967, routine-task occupations made up 62 percent of employment. By February 2017, that number had dropped almost 20 percentage points to 43 percent.

In this paper, I argue that monetary policy—specifically, *contractionary* monetary policy—accelerates this process. In particular, a contractionary monetary policy shock that increases the Fed Funds Rate by 100 basis points (b.p.) produces a persistent decline in the share of employment in routine occupations that peaks at 1 percentage point. An expansionary policy shock of the same size, on the other hand, does not produce a statistically

Figure 2.2: Routine Jobs as a Share of Total Employment

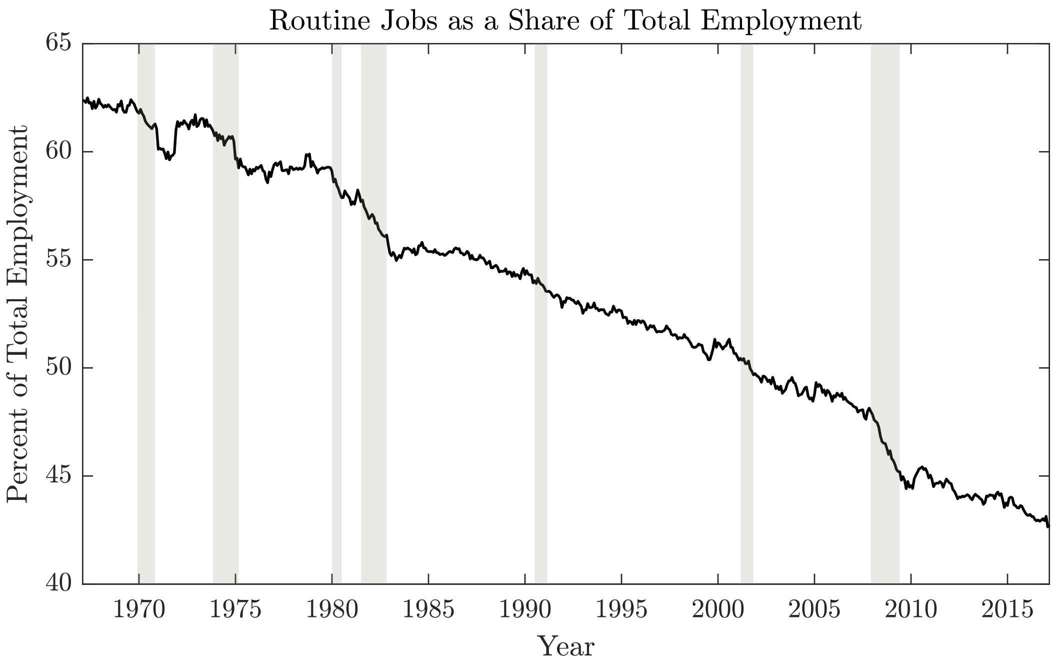

Share of total civilian employment in routine occupations, 1967-2017. Occupational employment data from the Current Population Survey. Shaded dates indicate NBER recessions.

significant change. Moreover, the decline in employment in routine jobs drives almost all of the decline in total employment after a contractionary monetary policy shock. I find that monetary policy shocks account for up to 40 percent of the changes in the share of employment in routine occupations over a two- to three-year horizon.

After establishing that monetary policy shocks have large and asymmetric effects on employment in routine jobs (and essentially no effect on nonroutine employment), this paper explores two possible channels for this pattern. One possibility is that monetary policy shocks affect the relative price of investment in a way that induces substitution toward capital and away from labor, the very mechanism that the job polarization literature has highlighted as an explanation for the trends evident in Figures 2.1 and 2.2. I find little evidence that monetary policy works through this channel. In fact the contractionary monetary policy shocks that drive the declines in routine employment *increase* the price of investment goods. Although there is some evidence of substitution toward *existing* capital, it is qualitatively different from the type of substitution toward new technologies emphasized in the job polarization literature.

Another possible way for monetary policy to have an outsize role on routine-task employment is that its effects differ by industry, and industries differ in the share of employment made up by routine-task jobs. In contrast to Jaimovich and Siu's (2014) analysis of routine job loss during recessions, I find this industry composition mechanism to be an important driver of the effects of monetary policy shocks on routine employment.[2] Mone-

[2]It is worth emphasizing that these results are not contradictory. I find only that industry composition is important in accounting for the effects of monetary policy, while they show that long-run changes in the industry makeup of the U.S. economy cannot explain the majority of job polarization. Foote and Ryan (2015), on the other hand, do argue that the decline of U.S. manufacturing employment played an important

tary policy shocks have larger effects on total employment in construction and durable goods manufacturing, and employment in each of these industries is concentrated in routine-task occupations.[3] By having stronger effects overall on those industries that utilize a greater share of routine-task employment—and even stronger effects when policy is contractionary—monetary policy has persistent effects on the mix of occupations in economy and potentially disproportionate welfare consequences for the "middle-skill" workers affected. Nevertheless, this industry effect does not explain the asymmetry *per se*; monetary policy has asymmetric effects on employment in virtually all industries.

In Section 2.2, I review the recent literature on job polarization and discuss Autor, Levy and Murnane's (2003) occupation classification system used throughout the paper. In Section 2.3, I describe the data on employment and monetary policy I use in subsequent sections. In Section 2.4, I discuss the estimation of linear and asymmetric impulse responses for occupational employment data and present results. In Section 2.5, I discuss the historical contribution of monetary policy shocks to job polarization. In Section 2.6, I discuss possible reasons for the results from the previous section. Section 2.7 concludes.

2.2 Literature review and occupation classifications

The literature on job polarization grew out of an earlier literature on earnings inequality, skill-biased technical change (SBTC), and international trade.[4] Research on these issues

role in job polarization.

[3]In 1972, the earliest year for which annual occupation-by-industry data are available from the CPS, more than 80 percent of employment in these industries was in routine employment, and these industries made up 17.2 percent of all nonroutine employment.

[4]See David H. Autor and Lawrence F. Katz (1999) for an overview.

typically focused on two groups—low- and high-skill workers—and studied the differential effects of SBTC, off-shoring, and immigration on these types of workers.

Autor, Levy and Murnane (2003) introduced a more nuanced categorization of jobs with the "task-based" framework. Their system highlights the tasks involved in performing a job, rather than the characteristics of the person holding that job. They delineate occupations along two dimensions: manual vs. cognitive, and routine vs. nonroutine. The first category describes whether the job's tasks are primarily physical or mental; the second, whether the job consists of "carrying out a limited and well-defined set of ... activities, those that can be accomplished by following explicit rules" (Autor, Levy and Murnane (2003)). They document a compositional shift away from routine occupations beginning in the 1970s, coincidental with the start of rapid computerization. Indeed, the change occurred most rapidly in those industries that more quickly adopted computing technology. Computers substituted for labor in performing routine tasks, while they complemented labor in nonroutine—escpecially nonroutine cognitive—occupations.

Subsequent research noted that routine occupations, both cognitive and manual, were held by workers in the middle of the wage and educational distribution. Cognitive nonroutine jobs (such as journalists, professors, or CEOs) were typically held by high-wage, high-education workers, while low-wage, low-education workers were in manual nonroutine jobs (such as janitors, barbers, groundskeepers). The growth of low-skill and high-skill jobs and simultaneous decline of middle-skill jobs, led to Maarten Goos and Alan Manning's (2007) introduction of the term "job polarization." The trend of job polarization has since been documented in many European countries and has been confirmed in the U.S. across numerous datasets and methodologies.[5]

More recently, researchers have begun to examine the cyclicality of job polarization. Jaimovich and Siu (2014) document that the decline in middle-skill jobs in the U.S. occurred almost entirely during recessions and never recovered during subsequent expansions. They show that this phenomenon can partially explain the so-called "jobless recoveries" after the three most recent recessions. Foote and Ryan (2015) also examine middle-skill employment over the business cycle and argue that its cyclicality is driven by middle-skill jobs' concentration in highly cyclical industries like construction and manufacturing. They also argue job polarization can partially explain recent declines in labor force participation, especially for men.

This paper links the *asymmetry* documented by Jaimovich and Siu (2014)—that routine employment falls during economic downturns but does not recover—with the *composition channel* that Foote and Ryan (2015) study—that routine jobs are concentrated in cyclical industries—and examines them in the context of monetary policy. The approach taken here is closely related to the literature examining asymmetric and state-dependent effects of shocks. While much emphasis has been given to the state-dependence of fiscal and tax shocks,[6] a more recent literature has considered the state- and sign-dependence of monetary policy shocks. Silvana Tenreyro and Gregory Thwaites (2016) find that monetary policy shocks have smaller effects on the real economy during recessions than in booms, and that contractionary shocks have larger effects than expansionary ones. Joshua D. An-

[5] David H. Autor (2015) provides an excellent overview of this research.

[6] See for example Alan Auerbach and Yuriy Gorodnichenko (2012*a*) and 2012*b*, Michael T. Owyang, Valerie A. Ramey and Sarah Zubairy (2013), and Valerie A. Ramey and Sarah Zubairy (2017).

grist, Òscar Jordà and Guido M. Kuersteiner (2016) and Regis A. Barnichon and Christian Matthes (2017), using very different methods, also find that contractionary monetary policy shocks have larger effects. This paper confirms these findings with respect to employment and highlights that the asymmetry exists only in employment in routine occupations.

The literature on the real effects of monetary policy is too long to review in detail here[7]; in the next section, however, I describe in detail the two empirical approaches I take.

2.3 Data

2.3.1 Occupational, industry, and investment data

Aggregated monthly employment by industry is available directly from the U.S. Bureau of Labor Statistics (BLS). Historically comparable occupational employment data from the Current Population Survey (CPS) are available from the BLS only back to 1983. For the years 1967–1982, I compiled monthly occupational employment data from tables in the BLS's monthly *Employment and Earnings* publication. Although the detailed occupational categories differ across the time period, the division of occupations into routine and non-routine occupations can be constructed in a consistent way, another appealing feature of Autor, Levy and Murnane's (2003) task-based framework.

Data on investment prices are from Riccardo DiCecio (2009) and are updated on a quarterly basis available at Federal Reserve Bank of St. Louis's FRED database. Data on the equipment capital stock are constructed from Bureau of Economic Analysis (BEA) data, adjusted as suggested by Robert J. Gordon (1990), following Per Krusell, Lee E. Ohanian,

[7]See Christiano, Eichenbaum and Evans (1999) for an overview of the early literature on monetary policy shocks, and Ramey (2016) for an overview of more recent developments.

José-Víctor Ríos-Rull and Giovanni L. Violante (2000).

2.3.2 Measures of monetary policy

To address the question of how monetary policy affects employment by occupation and industry, as a baseline I use monetary policy shocks identified by Romer and Romer (2004) and extended to 2008 by Olivier Coibion, Yuriy Gorodnichenko, Lorenz Kueng and John Silvia (2012). Romer and Romer (2004) identify monetary policy shocks as changes to the Federal Funds target rate not predictable by the economic information in the Federal Reserve's "Greenbook" forecasts. Specifically, their monetary policy shock series is given by the residuals of the following regression:

$$\begin{aligned}\Delta ff_m = \alpha + \beta ffb_m \quad &+ \quad \sum_{i=-1}^{2} \gamma_i \widetilde{\Delta y}_{mi} + \sum_{i=-1}^{2} \lambda_i \left(\widetilde{\Delta y}_{mi} - \widetilde{\Delta y}_{m-1,i} \right) \\ &+ \quad \sum_{i=-1}^{2} \varphi_i \tilde{\pi}_{mi} + \sum_{i=-1}^{2} \theta_i \left(\tilde{\pi}_{mi} - \tilde{\pi}_{m-1,i} \right) + \rho \tilde{u}_{m0} + \varepsilon_m, \end{aligned} \tag{2.1}$$

where m indexes FOMC meeting dates, ffb_m denotes the level of the Federal Funds target rate at the time of meeting m, $\widetilde{\Delta y}_{mi}$ denotes forecasts of real output growth, $\tilde{\pi}_{mi}$ denotes forecasts of inflation, $\tilde{u}_{m0}$ denotes forecasts of current unemployment, and ε_m, the residual, is the monetary policy shock. The index i is the horizon of the forecast, and horizon $i = -1$ may be a true forecast or a realized value of the variable, depending on when the actual data were available. The shock series is from 1969–2008, which determines the sample period in all the estimates that follow.

An independently identified shock series—as opposed to one identified in a structural vector autoregression (SVAR)—facilitates estimation of nonlinear impulse responses (in this case, asymmetric responses to contractionary versus expansionary shocks). The use of this series is not without drawbacks, however. For example, the series ends in 2008, but Romer and Romer's (2004) methodology is not amenable to extension to the zero lower-bound period since it estimates the shock based on the change in the Fed Funds Rate, which exhibits essentially zero variation for almost a decade after 2008. Moreover, ignoring the post-2008 period necessitates dropping the largest decline in routine-task employment in the data.

Therefore, as an alternative, I use a monetary policy shock identified using a hybrid of high-frequency and SVAR methods, as in Mark Gertler and Peter Karadi (2015). Their proxy SVAR method uses Fed Funds Futures as external instruments to identify a monetary policy shock. While pure high-frequency identification strategies are limited by data availability, the hybrid method of Gertler and Karadi (2015) facilitates the identification of shocks over a longer period. Specifically, they use high-frequency data to estimate the linear relationship between reduced-form and structural shocks in the VAR system. They then estimate the SVAR over a longer period under the assumption that the estimated relationship between reduced-form and structural shocks is valid for the entire period.[8] I use their method to estimate a monetary policy shock over a long horizon beginning in 1973.[9] Because the results from this exercise are broadly similar with the baseline linear results, I present the results from the hybrid identification approach in Appendix B.1.

[8] This may not be the case, for example, in the pre- and post- Paul Volker chairmanship periods. The results presented here, however, are robust to the exclusion of pre-Volcker years.

[9] The details of the construction of this shock measure are discussed in Gertler and Karadi (2015), so I omit a detailed description here. The start date of 1973 is dictated by the availability of Simon Gilchrist and Egon Zakrajsek's (2012) measure of the excess bond premium.

2.4 Estimation methodology

In the baseline estimates, I use Jordà's (2005) method of local projections to estimate impulse response functions (IRFs) of occupational employment to a Romer and Romer (2004) monetary policy shock. The local projection method is extremely flexible and is particularly amenable to externally-identified shocks. Not only does it allow for simple linear estimates,[10] but Jordà's (2005) method also allows for a variety of nonlinearities, including sign-dependence.

In what follows, I first estimate linear responses both as a baseline for comparison against the asymmetric IRFs estimated in what follows and to facilitate comparison across shock identification techniques. I then exploit the flexibility of local projections to compute *sign-dependent* impulse responses for two different specifications.

2.4.1 Linear estimates

The baseline linear estimates are obtained as follows. For a shock ϵ_t and controls $\mathbf{x}_t$, and for horizons $h = 1, \ldots, H$ I estimate

$$y_{t+h} - y_t = \boldsymbol{\beta}_h' \mathbf{x}_t + \underbrace{\gamma_h \epsilon_t}_{\text{Linear}} + u_{t,h}. \tag{2.2}$$

The impulse response at horizon h to a one-unit shock, which corresponds to a 100 b.p. contractionary Romer and Romer (2004) shock, is simply given by the estimated coefficient $\hat{\gamma}_h$, and the width of the error bands is determined from the standard errors of these coefficient

[10]That is, simple relative to estimates from a linear VAR, the IRFs for which are highly nonlinear functions of estimated parameters.

estimates.[11] As a baseline, the vector $\mathbf{x}_t$ includes one year of lags each of the dependent variable and shock, as well as a constant and linear time trend.[12]

Figure 2.3 displays the IRFs from (2.2) of total employment as well as employment in routine- and nonroutine-task jobs, both in absolute terms and as a share of total non-farm employment. A contractionary shock reduces employment in routine occupations by almost 2 percent, and is statistically significant for more than four years after the shock. The IRF of nonroutine employment is insignificantly different from zero at nearly all horizons. This is reflected in the response of employment shares and total employment: a contractionary shock produces a decline in the share of employment in routine occupations of about 0.4 percentage points and a decline in total employment of about 1 percent, both of which are significant for more than four years.

The results from this section highlight that nearly *all* the decline in employment after a contractionary monetary policy shock comes from the response of employment in routine jobs.

The next section examines how these same variables respond when the linearity assumptions of (2.2) are relaxed. This is motivated by the findings of Jaimovich and Siu (2014), as well as the literature that finds significantly asymmetric effects of monetary policy discussed in Section 2.2.

[11] To account for serial correlation, the standard errors are adjusted as in Whitney K. Newey and Kenneth D. West (1987). In addition, since the error bands for IRFs implicitly test the null hypothesis of zero effect, inference based on the standard errors of these coefficients is valid despite the presence of a generated regressor (see Pagan (1984)).

[12] Section B.2 discusses alternative lag structures.

Figure 2.3: Impulse Responses to a Monetary Policy Shock – Linear Estimates

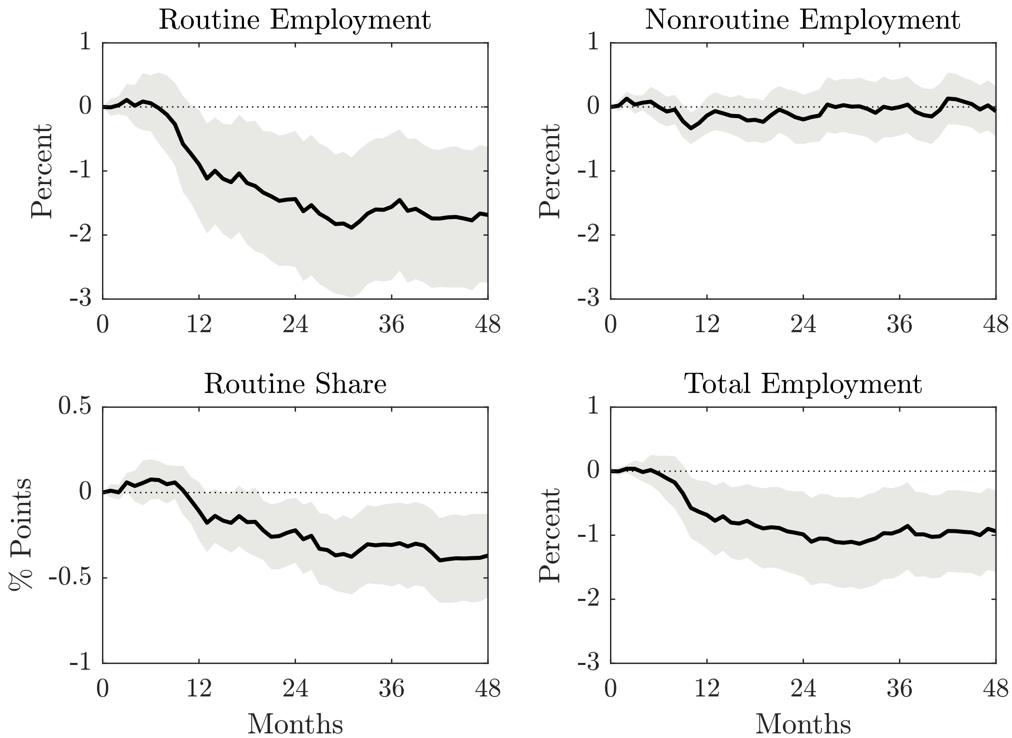

Impulse responses of occupational employment to a 100 b.p. contractionary Romer and Romer (2004) monetary policy shock. Estimated from Equation 2.2. Shaded areas are 90 percent confidence intervals.

2.4.2 Asymmetric estimates

As discussed above, Romer and Romer's (2004) shock series combined with Jordà's (2005) flexible method of local projections allows for convenient estimation of asymmetric impulse responses. Equation 2.2, in keeping with the traditional VAR literature, imposes the assumption that impulse responses are symmetric; that is the responses to positive and negative shocks are identical except for their sign. Local projections facilitate a number of ways of dispensing with this assumption. First, I estimate a specification that makes the simple assumption that positive and negative shocks are qualitatively different and allows for responses to differ depending on the sign of the shock. This simple specification is identical to that used in Tenreyro and Thwaites (2016) and Arlene Wong (2016). Specifically, I estimate for each horizon $h = 1, \ldots, H$,

$$y_{t+h} - y_t = \boldsymbol{\beta}_h' \mathbf{x}_t + \underbrace{\gamma_h^+ \epsilon_t^+}_{\text{Contractionary}} + \underbrace{\gamma_h^- \epsilon_t^-}_{\text{Expansionary}} + u_{t,h}, \tag{2.3}$$

where $\epsilon_t^+ := \max\{\epsilon_t, 0\}$ and $\epsilon_t^- := \min\{\epsilon_t, 0\}$. To facilitate comparisons with Section 2.4.1, I include the same controls as in (2.2), and estimate the IRFs to a four-year horizon. This specification allows for straightforward tests of the null hypothesis of symmetry, i.e., H_0 : $\gamma_h^+ = \gamma_h^-$ for $h = 1, \ldots, H$. One potential drawback of this specification, however, is that it imposes linearity on the impulse response conditional on the sign of the shock. This linearity assumption produces some problematic results discussed briefly below and in more detail in Appendix B.2.

Therefore, as an alternative, I also estimate asymmetric IRFs using a specification that is quadratic in the shock. Specifically, for $h = 1, \ldots, H$, I estimate

$$y_{t+h} - y_t = \boldsymbol{\beta}_h' \mathbf{x}_t + \gamma_{1,h} \epsilon_t + \gamma_{2,h} \epsilon_t^2 + u_{t,h}. \tag{2.4}$$

This specification dispenses with the assumption in (2.3) that shocks have linear effects conditional on their sign, instead allowing for more flexible asymmetries. One potential drawback of dropping this assumption involves the test for asymmetry. While the baseline specification in (2.3) restricts shocks to differ *only* with respect to their sign, the quadratic specification allows for other nonlinearities as well. For example, including the square of the shock allows for the impulse response to differ nontrivially (i.e., not simply a proportional scale change) with the size of the shock. Therefore, the test of the null hypothesis H_0 : $\gamma_{2,h} = 0$ for $h = 1, \ldots, H$ is somewhat easier to reject than the null hypotheses of symmetry in (2.3); that is, the null hypothesis is not that there are no asymmetries, but rather that there are no *nonlinearities* (up to second order), a null that is easier to reject.[13]

The impulse responses estimated from Equations 2.3 and 2.4 are displayed in the left and right columns of Figure 2.4, respectively. In response to a contractionary shock of 100 b.p., routine employment falls by 4 percent, while the point estimate of the response to an expansionary shock is at most 1 percent and is statistically insignificant. In contrast, the linear estimates implied a peak decrease (increase) of about 2 percent to a contractionary (expansionary) shock.

The response of employment shares is even starker. A contractionary shock produces a full percentage-point decline in the share of employment made up by routine occupations—a response that lasts for the full estimated horizon—while an expansionary shock produces a peak effect of at most less than 0.3 percentage points and is insignificant. By comparison the peak linear response was less than 0.4 percent points.[14]

The point-wise tests for asymmetries in the impulse responses are displayed in Figure 2.5. Under both specifications, the null hypothesis of symmetry can be rejected at almost all horizons for all variables except for nonroutine employment, which displays relatively little

[13]Indeed, this is reflected in the point-wise asymmetry tests displayed in Figure 2.5.

[14]It is worth noting, however, that the results from (2.3) imply that nominally "expansionary" shocks—i.e., shocks that produce declines in the Fed Funds Rate—are actually *contractionary*, at least for some periods for the variables included here. As discussed in Appendix B.2, this pattern is mitigated with the inclusion of more lags of the shock and exacerbated with the inclusion of fewer lags. As is evident from these specifications, as well as the quadratic specification, this anomaly is *not* a robust feature of the data.

Figure 2.4: Impulse Responses to a Monetary Policy Shock – Asymmetric Estimates

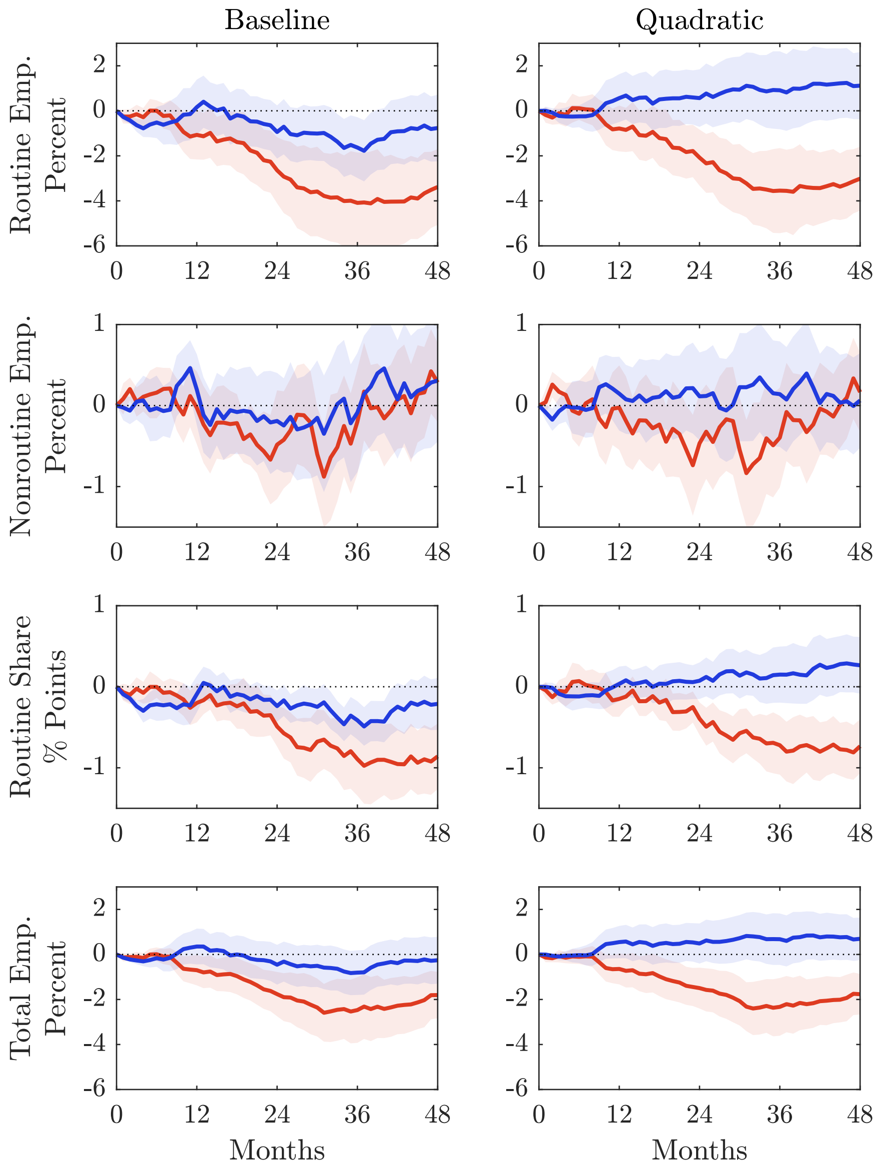

Impulse responses of occupational employment to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock. Shaded areas are 90 percent confidence intervals. *Left column*: baseline asymmetric estimates from Equation 2.3. *Right column*: alternative quadratic estimates from Equation 2.4.

Figure 2.5: Point-wise Tests for Asymmetry

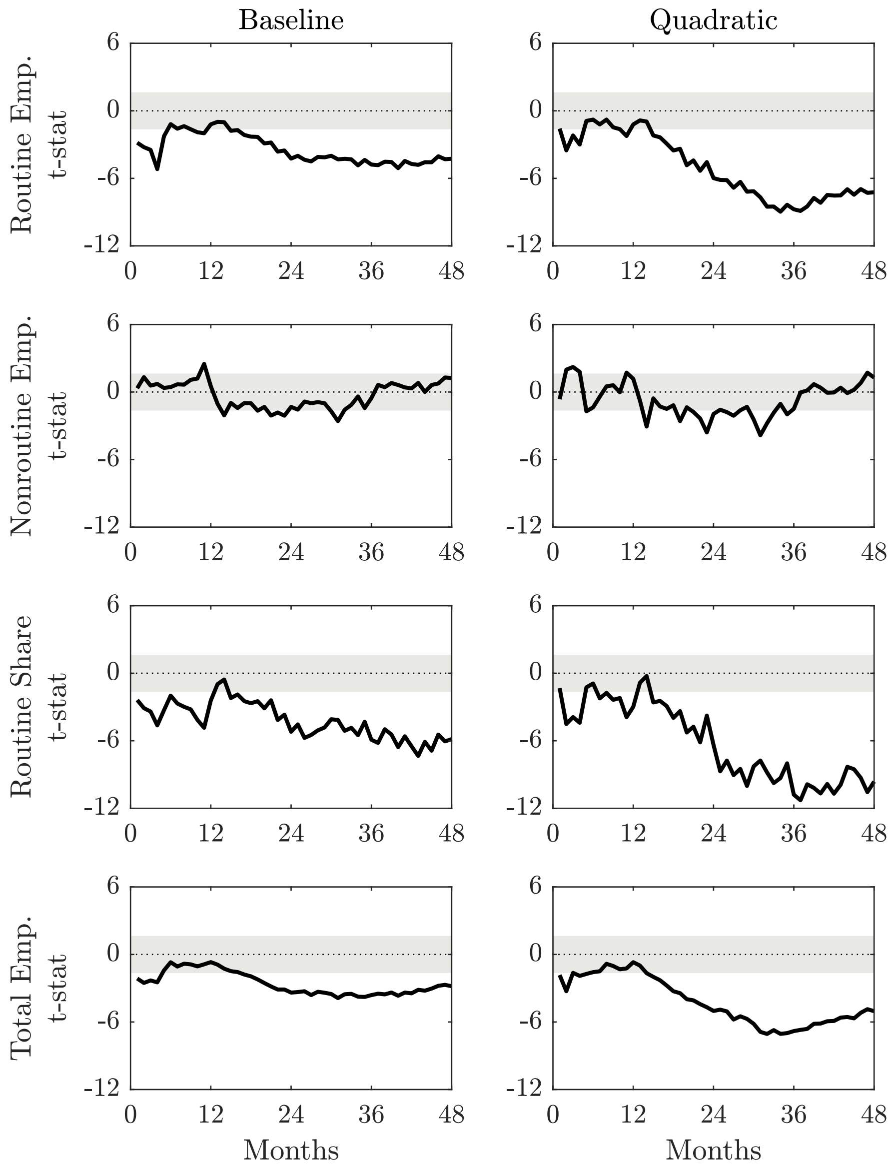

Point-wise t-test against the null of symmetry for sign-dependent IRFs in Figure 2.4. Shaded areas indicate a 90% confidence region. Lines outside the shaded region indicate the null can be rejected at 10% significance. *Left column*: baseline asymmetric estimates from Equation 2.3. *Right column*: alternative quadratic estimates from Equation 2.4.

asymmetry.[15] The t-statistics for the quadratic estimates in the right column of Figure 2.5 are somewhat larger than those of of the baseline asymmetric specification; this is consistent with the previous observation that the quadratic specification allows for other nonlinearities than just asymmetry. The overall pattern of asymmetry is nevertheless quite similar across specifications, with larger degrees of asymmetry at longer horizons.[16]

2.5 Contribution of monetary policy to job polarization

Thus far, the analysis has focused on the conditional responses of occupational employment to monetary policy shocks. Impulse responses, while they reveal the causal effects of monetary policy, do not by themselves, however, paint a full picture of the economic significance of shocks. Of equal interest is how much monetary policy shocks have contributed historically to variations in occupational employment.

In a structural VAR, it is fairly straightforward to calculate the historical contribution of a shock to the time series of a variable of interest because the VAR is a structural econometric model. These very structural assumptions, however, are what lead to the inflexibility of VAR-based impulse responses relative to those estimated using local projections. On the other hand, because IRFs estimated by local projections impose few structural assumptions, the estimates provide little information on how a shock propagates through the economy.

[15]Perhaps an unsurprising fact, given that neither impulse response is significantly different from zero at virtually any horizon.

[16]An alternative approach is to estimate all horizons of the impulse response jointly, allowing for correlation of the errors across horizons, and to perform a joint significance test of the relevant coefficients. In such a test, the null hypothesis of symmetry is very easy to reject. The point-wise tests displayed here are biased against finding asymmetries and are consistent with the point-wise confidence intervals displayed in the IRFs throughout the paper.

One approach is to use the one-period-ahead local projection (i.e., (2.2) or (2.3) estimated at $h = 1$)—which is equivalent to an autoregressive distributed lag (ADL) model—as the structural model; however, this effectively ignores the information contained in the other $H-1$ projections. The approach I take here, based on forecast error variances (FEVs), exploits the full set of estimates from (2.2), (2.3), and (2.4), at the cost of not being able to discuss particular historical episodes.

For notational simplicity, consider a more general series of local projections that nests the linear and both asymmetric specifications considered above:

$$y_{t+h} - y_t = \widetilde{\boldsymbol{\beta}}_h' \widetilde{\mathbf{x}}_t + \boldsymbol{\theta}_h' \widetilde{\boldsymbol{\epsilon}}_t + u_{t,h}, \tag{2.5}$$

where $\widetilde{\mathbf{x}}_t$ includes only lags of the dependent variable, and $\widetilde{\boldsymbol{\epsilon}}_t$ includes current and lagged values of the monetary policy shock, potentially distinguishing between ϵ_t^+ and ϵ_t^-, or ϵ_t and ϵ_t^2, depending on the specification. The forecast error from (2.5) at time t, horizon h is given by

$$\begin{aligned} FE_{t,h} &\equiv y_{t+h} - y_t - \mathbb{E}\left[y_{t+h} - y_t \,|\, \widetilde{\mathbf{x}}_t, \widetilde{\boldsymbol{\epsilon}}_t\right]. \\ &= u_{t,h}. \end{aligned}$$

The mean squared forecast error (MSFE) at horizon h is then given by $MSFE_h = \mathbb{E}\left[FE_{t,h}^2\right] = \mathbb{E}\left[u_{t,h}^2\right]$, which can be estimated using regression residuals.

To assess the importance of monetary policy shocks, I compare this MSFE to that of a specification in which monetary policy shocks are not included in the regression at all. The degree to which the specifications that do include non-zero monetary policy shocks improve on the forecast errors of those that do not include them can be interpreted as the contribution of monetary policy shocks to changes in the dependent variable at a given horizon.

The most straightforward comparison is to consider the local projections that do not include monetary policy shocks at all.[17] That is

$$y_{t+h} - y_t = \widetilde{\boldsymbol{\beta}}_h' \widetilde{\mathbf{x}}_t + e_{t,h}. \tag{2.6}$$

As above, the forecast errors are the estimated regression residuals and the MSFE at horizon h is given by $\mathbb{E}\left[e_{t,h}^2\right]$. Since the forecast errors are mean zero by construction, a direct comparison of the MSFEs from the two specifications at a given horizon gives the share of the forecast error variance explained by monetary policy shocks.

Table 2.1 displays the share of the FEV of occupational employment variables at a given horizon that is due to the monetary policy shock, for both the linear and asymmetric case and for the alternative forecast discussed above. Even in the linear case, monetary policy shocks account for relatively large shares of the FEV for total employment and routine employment; they account for a negligible portion of the FEV for nonroutine employment, however. In the asymmetric cases, the role of monetary policy shocks is even larger, accounting for 40 percent of the variance in the share of employment made up by routine jobs over a two- to three year horizon.

That monetary policy shocks can explain such a large share of the short-run fluctuations in routine employment is surprising. Together with the fact that the effects on routine employment of contractionary shocks displayed in Figure 2.4 are so persistent, this

[17] I also considered a slightly different comparison that takes into account the existence of monetary policy shocks in estimating the coefficient vector $\widehat{\boldsymbol{\beta}_h}$, but assumes that monetary policy shocks are always at their mean value. This comparison is straightforward in the linear case, since the shock series is mean zero, but the asymmetric estimates introduce some small complications because there are terms with non-zero conditional means that must be accounted for. The details and results are described in Appendix B.3. The implied reductions in the FEVs are virtually identical to those from (2.6).

Table 2.1: Share of Forecast Error Variance Due to Monetary Policy Shocks

	Horizon (months)	Linear Eqn. 2.2	Baseline Eqn. 2.3	Quadratic Eqn. 2.4
Routine Employment	12	0.18	0.23	0.23
	24	0.18	0.32	0.33
	36	0.11	0.23	0.24
	48	0.09	0.16	0.17
	Max.	*0.19*	*0.33*	*0.33*
Nonroutine Employment	12	0.03	0.07	0.07
	24	0.01	0.04	0.04
	36	0.01	0.05	0.03
	48	0.01	0.06	0.04
	Max.	*0.05*	*0.10*	*0.11*
Routine Share	12	0.16	0.22	0.20
	24	0.20	0.38	0.37
	36	0.13	0.32	0.32
	48	0.10	0.23	0.22
	Max.	*0.22*	*0.40*	*0.37*
Total Employment	12	0.18	0.24	0.24
	24	0.15	0.27	0.29
	36	0.08	0.18	0.19
	48	0.07	0.11	0.12
	Max.	*0.20*	*0.28*	*0.29*

Notes: The table presents the share of the forecast error variance of occupational employment variables due to monetary policy shocks, for both the linear and asymmetric baseline estimates.

observation raises the question of what causes monetary policy shocks to have such large and long-lasting effects on routine employment. Two different mechanisms are discussed in the next section.

2.6 Two possible mechanisms

This section discusses two possible channels through which monetary policy may have a disproportionate effect on routine employment. This discussion is by no means meant to be exhaustive, but it does highlight that the driving force of short- to medium-run changes in routine employment is very different from the long-run technological forces emphasized in the job polarization literature.

The first possibility I consider is that monetary policy shocks might lead indirectly to changes in routine employment by affecting the price of capital goods, which are a substitute in production for routine employment. If, for example, a contractionary monetary policy shock leads to a decline in the price of capital goods, substitution of new capital for routine labor could be driving the results in the previous sections. As discussed in Section 2.1, the trend decline in the price of equipment capital, reflecting technological progress, is widely thought to be an important driver of the long-run decline in routine employment. This "price-of-investment" mechanism is therefore a natural candidate to examine in the context of monetary policy.

Another possibility is a simple industry composition effect. Monetary policy might disproportionately affect routine employment because it has stronger effects on those industries in which routine employment is concentrated.

I find no evidence that the "price-of-investment" mechanism is driving the results from the previous section. A contractionary shock leads to an *increase* in the cost of capital. Thus, substitution between capital and labor cannot by itself explain the strong response of routine employment to monetary policy shocks.[18] I find modest evidence in support of the "industry composition" channel, however. Monetary policy shocks disproportionately affect employment in construction and durable goods manufacturing, industries in which employment is concentrated in routine occupations. The asymmetry of monetary policy shocks (as opposed to the *magnitude*) is fairly consistent across all industries, however.

2.6.1 Capital-labor substitution and the relative price of investment

As discussed at length in Autor, Levy and Murnane (2003) and Autor (2015), routine-task occupations are particularly susceptible to automation. Although the process of automation is typically thought of as a long-run trend phenomenon, it can potentially impact the cyclical dynamics of routine employment as well, at least to the extent that capital and labor are substitutable in the short run. One particular channel though which short-run substitution of capital for labor might occur is via the relative price of investment. If a contractionary monetary policy shock lowers the price of capital for which routine labor is a substitute in production, then one would expect to see very persistent declines in routine employment as firms substitute toward capital.

Jonas D. M. Fisher (2006), for example, finds the relative price of investment goods, particularly investment equipment, to be strongly negatively correlated with output; he

[18] This does not necessarily mean that substitutability between capital and routine-task labor plays *no* role; there might still be substitution toward capital relative to labor with both inputs declining.

argues that shocks to investment-specific technology are a large driver of the business cycle. On the other hand, Paul Beaudry, Alban Moura and Franck Portier (2015) find that the relative price of investment to be acyclical or even procyclical during certain time periods. Thus, it is not obvious *a priori* how monetary policy shocks will affect investment prices.

To test this notion, I consider the effect of a Romer and Romer (2004) monetary policy shock on measures of capital, investment, and the relative price of investment goods. Following Krusell et al. (2000), I focus on equipment investment and capital.[19] The focus on capital equipment—as opposed to structures—is meant to get closer to the exact types of capital for which routine labor is a substitute; for example, an assembly line worker might be substituted for by a robotic system, but a new factory is not a (direct) substitute for that worker. This focus on equipment is motivated by the effects of automation on routine employment from the job polarization literature, but the results below hold if one considers both equipment and structures.

I construct series on the capital stock, investment flow, and relative price of investment for equipment goods on a quarterly basis to cover the same time period as the data in Section 2.3, 1969-2012. The source data are from the National Income and Product Accounts (NIPA) Table 5.3.5 and DiCecio's (2009) quality-adjusted price measures based on Krusell et al. (2000). Following Krusell et al. (2000), I construct a series of the equipment capital stock using base-year estimates from Gordon (1990).

Figure 2.6 displays impulse responses to monetary policy shocks of these variables esti-

[19]Although Krusell et al. (2000) focus on the differences between skilled and unskilled labor, their basic framework extends to the routine vs. nonroutine distinction, as discussed in Section 2.2.

Figure 2.6: IRFs to a Monetary Policy Shock — Capital, Investment, and Prices.

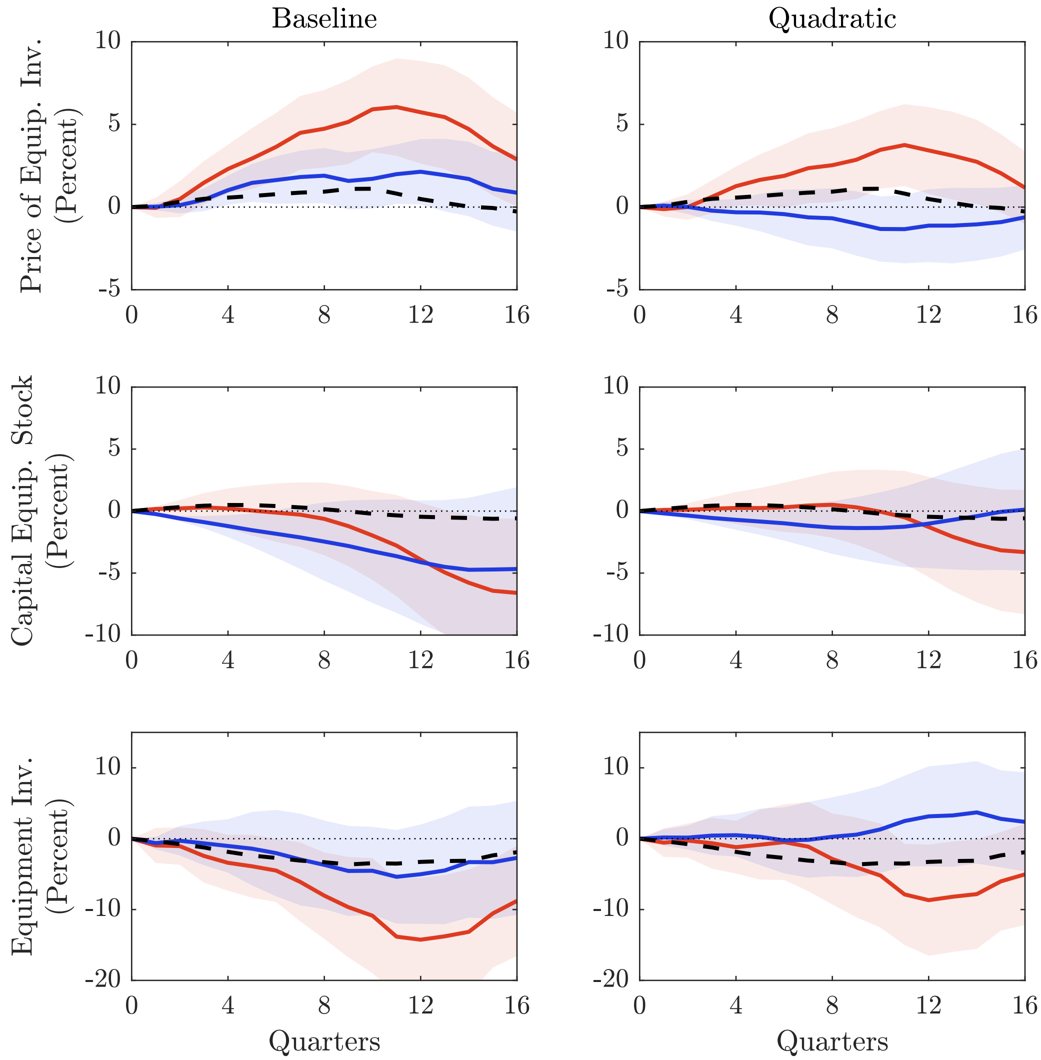

Impulse responses of equipment capital, investment, and prices to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock. Dashed lines are the IRFs of a contractionary shock estimated linearly. Shaded areas are 90 percent confidence intervals. *Left column*: baseline asymmetric estimates from Equation 2.3. *Right column*: alternative quadratic estimates from Equation 2.4.

mated from quarterly versions of Equations 2.2, 2.3, and 2.4.[20] A contractionary shock leads to large and significant increase of about 4 to 5 percent in the price of investment equipment. It is unsurprising, given this result, that investment in equipment declines dramatically, between 9 and 15 percent, depending on the specification, and the stock of equipment capital declines gradually. Linear estimates also predict this pattern in response to a contractionary shock. As with the previous estimates, expansionary shocks have little effect.[21]

As discussed in Section 2.2, the job polarization literature highlights the important role of falling investment prices in driving the long-run trends in declining routine employment. The results in this section demonstrate that this argument does *not* extend to the short-run fluctuations in routine employment driven by monetary policy shocks. The large declines in routine employment after contractionary monetary policy shocks are not driven by lower prices of investment and substitution toward new technologies. In fact, the price of investment *increases* in response to a contractionary monetary policy shock, and investment in equipment declines. Firms are not substituting away from routine labor and toward capital; indeed, both routine labor and the stock of capital equipment decline by around 4 percent at a three- to four-year horizon.[22]

[20]As in the baseline monthly calculations, the vector of control variables $\mathbf{x}_t$ includes a one-year lag of both the dependent variable and the shock.

[21]Note that the anomalous responses to expansionary shocks discussed in Section 2.4.2 are present here as well. As before, these patterns do not appear in the quadratic estimates.

[22]Because routine employment declines more rapidly than the capital stock, firms may effectively be substituting away from labor and toward *existing* capital in the very short run. This is a qualitatively different mechanism, however, than substitution toward *new* capital that is an important driver of the long-run trends.

2.6.2 Industry effects

Another way that monetary policy might have strong effects on routine employment is via industry effects. In particular, monetary policy shocks might have larger effects on some industries than others; if those strongly impacted industries employ a large number of routine-task workers, the responses of routine employment in Section 2.4.2 may be driven by these industry differences. This section presents some evidence in support of this mechanism.

I find that total employment in construction and durable goods manufacturing respond strongly to contractionary shocks and only modestly to expansionary ones. In 1972,[23] routine jobs made up 86 and 82 percent of total employment in these industries, and these two industries accounted for 17 percent of total routine employment.[24] Other industries' employment responses are also asymmetric, but no other industries' responses are as large in magnitude as construction and durable goods manufacturing.[25]

Figure 2.7 displays the IRFs from (2.2) to a 100 b.p. contractionary Romer and Romer (2004) shock of employment in a variety of industries. While employment in some industries is hardly affected, other industries—in particular, construction and durable goods manufacturing—decline by large amounts, about 3 to 4 percent. As a comparison, the

[23]The first year for which annual employment by occupation *and* industry are available.

[24]The same figures in 2017 are 79 percent, 66 percent, and 8 percent, indicating that these industries played an important role in the overall trend decline in routine employment, as argued in Foote and Ryan (2015).

[25]Ideally, this effect would be tested directly by examining data by occupation *and* industry. Unfortunately, because of several industry and occupation reclassifications in the Current Population Survey (CPS), it is difficult to construct consistent time series of employment by industry and occupation; moreover, the monthly data are available from the CPS for a shorter time period and smaller sample sizes within each industry-occupation cell introduces additional uncertainty. Therefore, I focus on comparing impulse responses in construction and durable goods manufacturing with other industries that differ in the share of their employment made up by routine jobs. This exercise is supportive of the notion that the direction of causation is from industry characteristics to occupational composition and not *vice versa*.

response of employment in finance, another industry that might reasonably be thought to be interest-rate sensitive, is statistically significant, but its peak effect is about four times smaller than that of construction. The linear estimates by themselves, given the concentration of routine employment in these two industries, suggest that monetary policy's differential impact across industries might drive its strong effects on routine employment.

Asymmetric industry-level employment estimates are displayed in Figure 2.8. Similar to the linear case, construction and durable goods manufacturing have larger responses than other industries, but here it is evident that this result is only true for contractionary shocks. In response to an expansionary shock, the increase in construction and durable goods manufacturing employment is modest and statistically insignificant at most horizons; contractionary shocks, on the other hand, lead to peak declines of about 6 to 7 percent, twice as large as the corresponding linear estimates.

Tests of asymmetry by industry are displayed in Figure 2.9. Here it is evident that the asymmetric effects of monetary policy shocks on employment are broad-based across industries. Contractionary shocks have significantly larger effects than expansionary shocks at two- to four-year horizons for all the industries displayed. The industry mechanism described in this section can, therefore, only explain a portion of the results on occupational employment in Section 2.4. It can explain the *large* and *persistent* effects on monetary policy shocks on routine employment; it cannot by itself explain the *asymmetric* effects.

2.7 Conclusion

This paper connects two phenomena previously not thought to be related: job polarization and the asymmetry of monetary policy shocks. Monetary policy shocks have large

Figure 2.7: IRFs to a Monetary Policy Shock — Industries, Linear

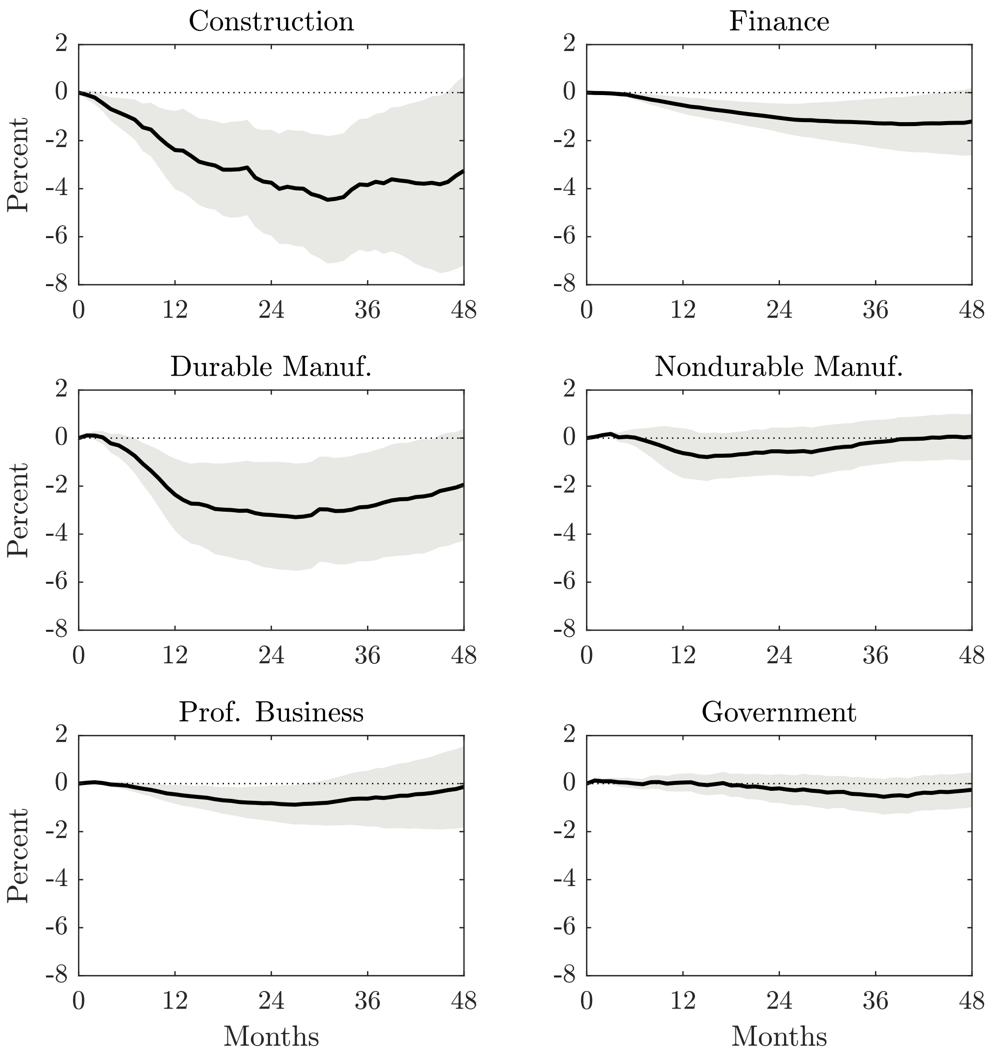

Impulse responses to a 100 b.p. contractionary Romer and Romer (2004) monetary policy shock. Estimated from Equation 2.2. Shaded areas are 90 percent confidence intervals.

Figure 2.8: IRFs to a Monetary Policy Shock — Industries, Asymmetric

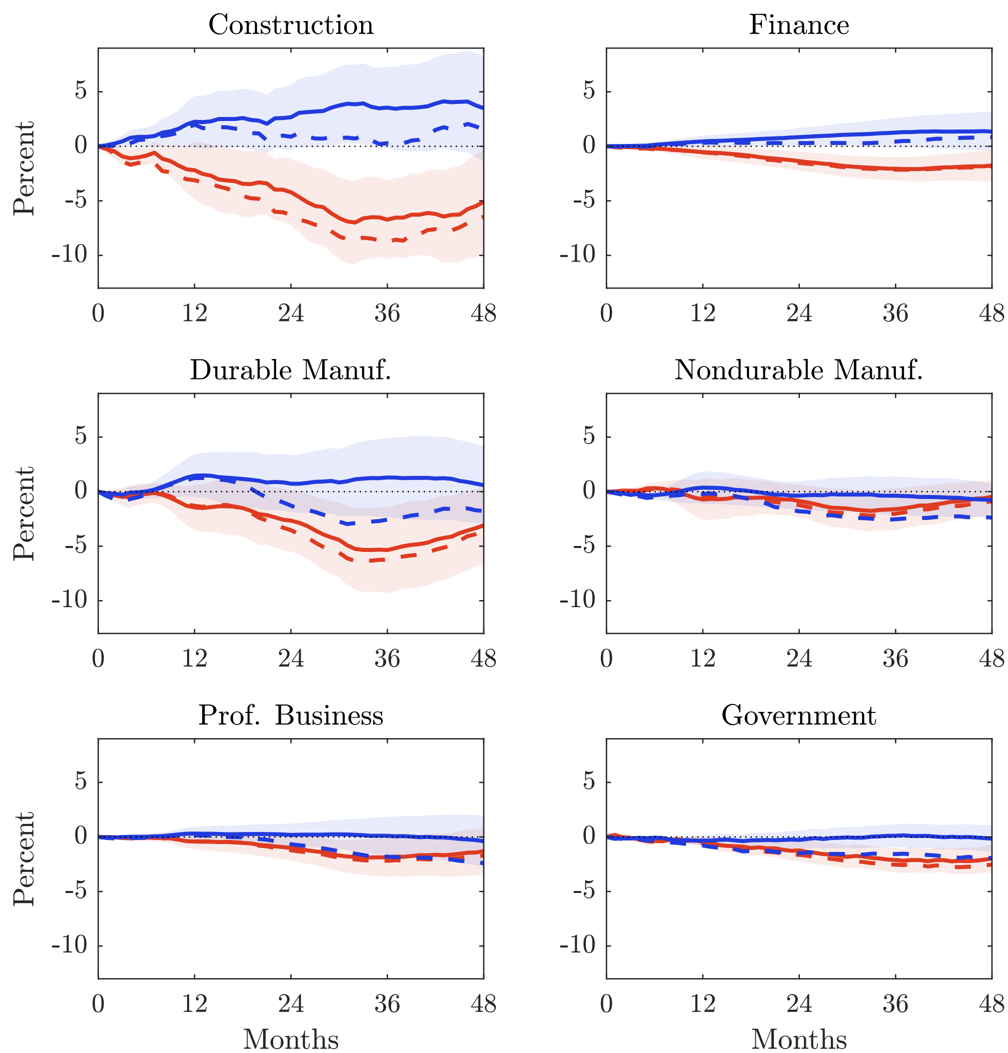

Impulse responses to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock. Estimated from Equation 2.4. Dashed lines are point estimates from 2.3. Shaded areas are 90 percent confidence intervals.

Figure 2.9: Point-wise Tests for Asymmetry — Industries

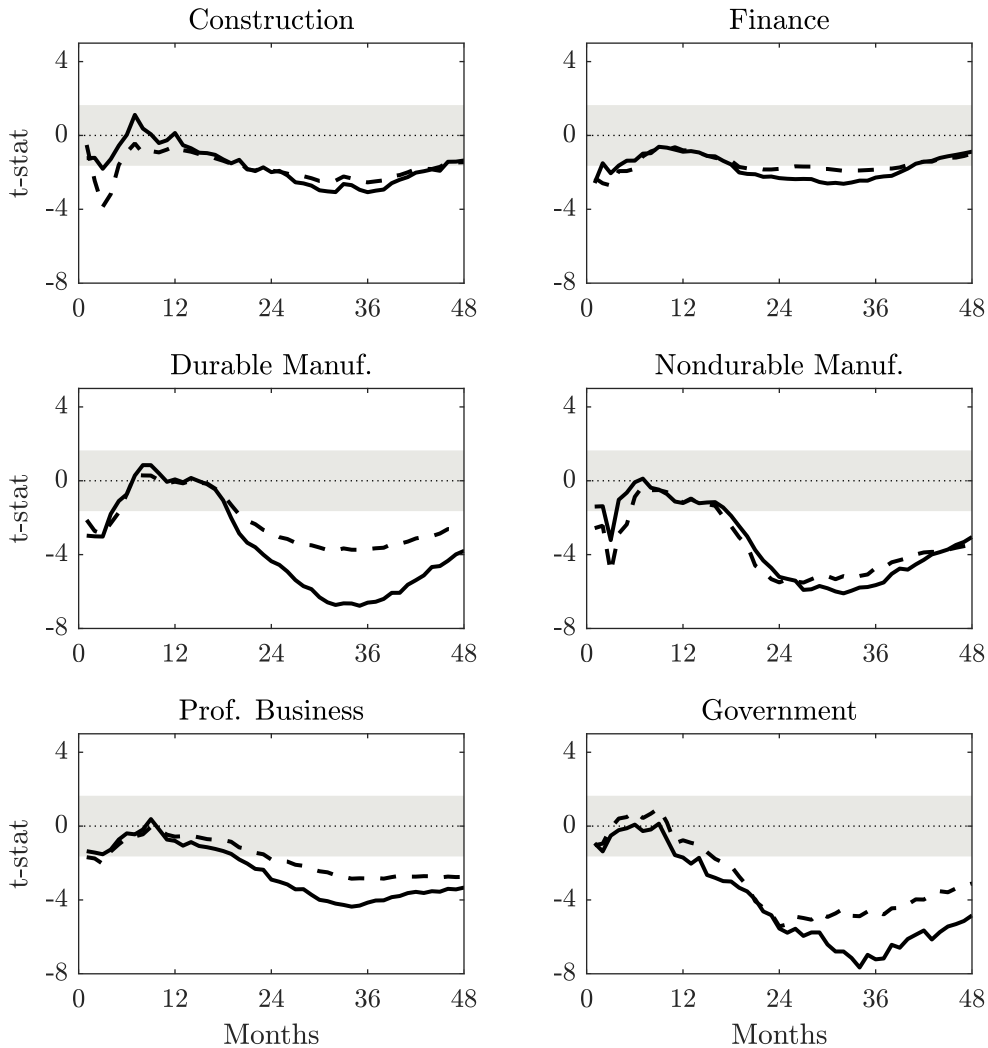

Point-wise t-tests against the null of symmetry for sign-dependent IRFs in Figure 2.8. Shaded areas indicate 90% confidence region. Solids lines from are from (2.4), dashed lines from (2.3). Lines outside the shaded region indicate the null can by rejected at 10% significance.

and persistent effects on employment in routine-task occupations—jobs that are typically held by workers with a high school degree and no college—and this effect is highly asymmetric: contractionary shocks lead to large declines in employment while expansionary shocks produce only modest increases. Moreover, the pattern of employment in routine jobs explains essentially all of the effects of monetary policy on total employment; the effects on nonroutine employment are small and insignificant. Monetary policy shocks can explain up to 40 percent of the forecast error variance of the share of employment in routine occupations over a two-year horizon. Because the effects of contractionary monetary policy shocks on routine employment are *large* and *persistent*, while those of expansionary shocks are essentially insignificant, monetary policy shocks have on average accelerated job polarization: contractionary shocks lead to almost permanent declines in routine employment, while expansionary shocks have almost no impact.

Although much of the trend decline in routine employment can be explained by increases in technology and the consequent declines in investment prices driving substitution toward capital and away from routine labor, the large share of short-run fluctuations in routine employment due to monetary policy shocks cannot be explained by this mechanism. In fact, contractionary monetary policy shocks lead to *increases* in the price of investment goods. I have presented some evidence, however, that monetary policy shocks drive short run fluctuations in routine employment because the industries on which monetary policy has large effects primarily employ workers in routine-task occupations. Employment in construction and durable goods manufacturing, in particular, respond very strongly and asymmetrically to monetary policy shocks, and between 65 and 85 percent of total employment in those industries is made up by routine jobs. Although the large and persistent effects on employment are unique to these industries, asymmetric responses occur in most industries. A fuller understanding of the sources of the asymmetric effects of monetary policy on routine employment is an important goal for future research.

The results presented in this paper may help explain the particularly deep decline in routine employment during the Great Recession. In particular, to the extent that monetary policy was constrained by the zero lower bound on interest rates, then the steep decline in routine employment is consistent with the results presented above. These results might also be used to help understand monetary policy's effects on inequality as described in Coibion et al. (2012). An understanding of monetary policy's disparate impacts on different groups—in the labor market and in capital and financial markets—is an important area of research. This paper is a first step toward a fuller understanding of monetary policy's contributions toward job polarization and other longer-term trends in the labor market.

Chapter 3

Are Uncertainty Shocks Expansionary? Evidence from the Michigan Survey of Consumers

3.1 Introduction

In recent years, much attention has been given to the role of uncertainty in macroeconomic fluctuations. Since macroeconomic uncertainty is not directly observed, a large portion of the literature has been focused on finding proxies for uncertainty. These have included measures based on stock market volatility, disagreement among survey respondents, and forecast errors from large macroeconomic and financial datasets.[1] In this paper, I construct new measures of uncertainty from the Michigan Survey of Consumers based on the share of respondents who say they are uncertain about various aspects of the economy. The responses to these survey questions are strongly correlated with future economic activity, leading strong credence to the use of the uncertainty measures derived from them.

The Michigan Survey is a monthly survey of approximately 500 individuals. It asks questions about respondents' views of business conditions over the next year(s) and the favorability of conditions for buying new houses or cars or large household durable items.[2]

[1]Nicholas Bloom's (2009) seminal contribution, and many subsequent papers, use stock market volatility as a measure of uncertainty. Rüdiger Bachmann, Steffen Elstner and Eric R. Sims (2013) construct their measure based on dispersion of responses about future economic activity from manufacturing businesses. Kyle Jurado, Sydney C. Ludvigson and Serena Ng (2015) and Sydney C. Ludvigson, Sai Ma and Serena Ng (2018) construct macroeconomic and financial uncertainty measures using large datasets.

[2]Sylvain Leduc and Zheng Liu (2016) also use a measure of uncertainty based on the Michigan Survey.

The exact responses coded in the survey differ depending on the question (I describe these in detail below), but in all cases respondents can answer that they are uncertain, they see both pros and cons to the situation, or they simply do not know. It is these responses, as well as the disagreement among respondents, that I use to construct measures of uncertainty. It is a particular strength of the Michigan Survey that both direct and dispersion measures of uncertainty can be constructed, allowing for an informal test of the external validity of Bachmann, Elstner and Sims's (2013) approach of using survey dispersion or disagreement as proxies for uncertainty.

I use these new measures to estimate the effects of uncertainty shocks using a number of structural vector autoregression (SVAR) models. Shocks to direct uncertainty measures and shocks to disagreement or response dispersion give estimates of the *opposite sign*. In particular, I find that increases in direct measures of uncertainty are expansionary, leading to modest declines in unemployment, while increases in disagreement are contractionary. The latter finding is consistent with results based on survey dispersion or disagreement, but the former suggests that these survey disagreement-based measures may not be good proxies for uncertainty.

The notion that uncertainty is contractionary is motivated primarily by the observation that most proxies for uncertainty tend to rise during recessions. The economic theory of uncertainty, however, is ambiguous regarding the sign of its effects.[3] As discussed in

As I describe below, the particular measure they choose appears to be incorrectly identified as "uncertainty" as a result of a classification error on the Michigan Survey's website.

[3]Contractionary effects of uncertainty include "real options" or "wait-and-see" effects, as discussed in Ben S. Bernanke (1983) and Bloom (2009). Expansionary effects include "growth options" in which downside effects of investment risk are limited, and exist in neoclassical growth models (see Simon Gilchrist and John C. Williams (2005) and Susanto Basu and Brent Bundick (2017)). Nicholas Bloom (2014) provides a broad

Ludvigson, Ma and Ng (2018), who also find shocks to macroeconomic uncertainty to be expansionary, estimating the *causal* effects of uncertainty depends crucially on both measurement and identification; that is, existing proxies are imperfect, and common identifying assumptions for uncertainty shocks (typically based on Cholesky decompositions in VAR models) are problematic. This paper attempts to makes progress on both fronts, introducing new measures of uncertainty and estimating causal effects using external instruments. In line with their results, I find that shocks that increase these new measures of uncertainty are expansionary. Results from recursively identified VARs and models using external instruments for identification are broadly consistent, at least for a subset of the new measures I construct.

In Section 2, I describe the data from the Michigan Survey and the new measures of uncertainty derived from it. In this section I also compare these new measures with existing uncertainty proxies. In Section 3, I estimate the effects of uncertainty shocks using SVARs. Section 4 concludes.

3.2 Survey-based measures of uncertainty

This section describes the Michigan Survey of Consumers and various measures of uncertainty derived from it and other surveys. In the first subsection below, I describe the Michigan Survey and the direct measures of uncertainty that I construct from it. I then discuss existing measures of uncertainty from survey data: Leduc and Liu's (2016) measure from the Michigan Survey and Bachmann, Elstner and Sims's (2013) method of

overview of the theory.

measuring survey response dispersion or disagreement. The former measure is a particularly problematic measure of uncertainty, as I discuss below. The latter method can be applied in a straightforward way to construct dispersion (or disagreement) measures from the Michigan Survey.

3.2.1 Direct measures from the Michigan Survey of Consumers

The Michigan Survey of Consumers is a monthly survey of individuals in the United States[4] conducted by the University of Michigan. Each month about 500 individuals are interviewed about their own individual financial situation and their views of the broader economy.[5] Data are available on a monthly basis since 1978 and quarterly since 1960.

Among the questions asked of survey respondents are questions about general business conditions and whether now is a good time to purchase a house, a car, or large durable goods. In this paper, I will focus on the following questions:

> (BUS): *"Now turning to business conditions in the country as a whole—do you think that during the next 12 months we'll have good times financially, or bad times, or what?"*

> (VEHIC): *"Speaking now of the automobile market—do you think the next 12 months or so will be a good time or a bad time to buy a vehicle, such as a car, pickup, van, or sport utility vehicle?"*

---

[4]Excluding Alaska and Hawaii.

[5]In the early years of the survey, often as many as 1400 individuals were interviewed. Since the late 1980s, however, the sample size has typically been between 500 and 600.

(DUR): *"Generally speaking, do you think now is a good or a bad time for people to buy major household items?"*

(HOUSE): *"Generally speaking, do you think now is a good time or a bad time to buy a house?"*

Responses are coded into six categories for the BUS question: Good times; good with qualifications; pros and cons or uncertain; bad with qualifications; bad times; or don't know. The other three questions do not have "with qualifications" responses coded separately, but are otherwise the same.

What I use to construct measures of uncertainty is the number of respondents who are unable to decide one way or another on the questions above. A period in which many of those being surveyed respond that they are uncertain or that there are pros and cons is considered to be a period of high uncertainty. To assess the reasonableness of these measures, as a first pass, Figure 3.1 displays the share of positive responses—that is that times are good, or that buying conditions are good—to these questions with real GDP growth. All series are procyclical, declining in recessions and increasing during expansions. Table 1 displays the correlations with GDP and industrial production growth and the unemployment rate at various leads and lags. It also includes the same correlations for other measures from which uncertainty proxies are constructed. All four series are correlated with future growth and lower unemployment, especially the response about general economic activity (BUS). This measure is more strongly correlated with output growth than the stock market or the survey responses used in Bachmann, Elstner and Sims (2013). Because they *predict* economic

Figure 3.1: Output Growth and Expectations from the Michigan Survey, 1961–2017

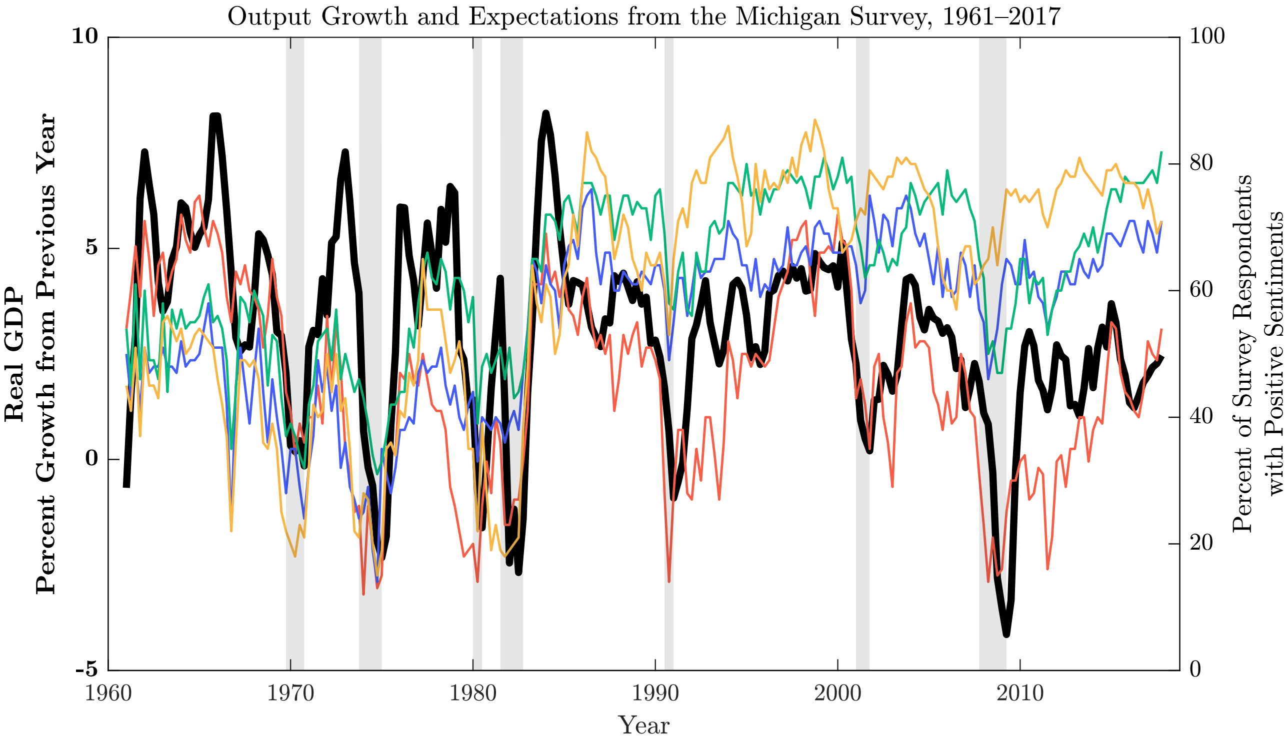

Left axis: **real GDP growth**. Right axis: positive answers to BUSS, VEHIC, DUR, and HOUSE. Shaded areas indicate NBER recessions, quarterly, 1961–2017.

activity so well, the measures proposed above are well-suited as proxies for macroeconomic uncertainty.

3.2.2 Alternative uncertainty measures

After the questions VEHIC, DUR, and HOUSE, respondents are asked why they gave the answer they did. It is on this follow-up question to VEHIC that Leduc and Liu (2016) base their measure of uncertainty. Table 38 of the Michigan Survey's time series

Table 3.1: Correlations of Expectations with Output Growth

	Leads & lags real GDP growth (quarters)								
Measure	-4	-3	-2	-1	0	+1	+2	+3	+4
BUS_GOOD	0.28	0.37	0.46	0.57	0.65	0.69	0.65	0.57	0.46
VEHIC_GOOD	-0.10	-0.06	-0.01	0.04	0.09	0.13	0.14	0.12	0.08
DUR_GOOD	0.10	0.14	0.19	0.24	0.25	0.25	0.19	0.11	0.01
HOUSE_GOOD	-0.10	-0.07	-0.03	0.01	0.06	0.12	0.14	0.12	0.07
BOS_INCR (1968–2017)	-0.53	-0.51	-0.43	-0.28	-0.07	0.14	0.31	0.42	0.44
Consumer Sentiment	0.31	0.39	0.48	0.57	0.63	0.65	0.60	0.51	0.39
S&P 500	-0.16	-0.15	-0.06	0.12	0.31	0.45	0.49	0.42	0.28
	Leads & lags of industrial production growth (quarters)								
	-4	-3	-2	-1	0	+1	+2	+3	+4
BUS_GOOD	0.20	0.26	0.33	0.42	0.52	0.57	0.56	0.50	0.39
VEHIC_GOOD	-0.17	-0.13	-0.09	-0.04	0.02	0.09	0.12	0.13	0.10
DUR_GOOD	0.00	0.05	0.10	0.16	0.21	0.23	0.19	0.13	0.01
HOUSE_GOOD	-0.12	-0.09	-0.06	-0.01	0.04	0.11	0.15	0.15	0.11
BOS_INCR (1968–2017)	-0.49	-0.47	-0.42	-0.33	-0.14	0.06	0.24	0.38	0.43
Consumer Sentiment	0.20	0.27	0.35	0.43	0.52	0.55	0.52	0.45	0.32
S&P 500	-0.16	-0.19	-0.12	0.03	0.24	0.45	0.54	0.51	0.38
	Leads & lags of the unemployment rate (quarters)								
	-4	-3	-2	-1	0	+1	+2	+3	+4
BUS_GOOD	-0.15	-0.21	-0.27	-0.35	-0.44	-0.53	-0.60	-0.65	-0.67
VEHIC_GOOD	0.17	0.14	0.10	0.05	-0.01	-0.07	-0.13	-0.18	-0.22
DUR_GOOD	0.11	0.04	-0.04	-0.12	-0.21	-0.28	-0.33	-0.36	-0.37
HOUSE_GOOD	0.20	0.17	0.13	0.08	0.02	-0.05	-0.11	-0.16	-0.21
BOS_INCR (1968–2017)	0.27	0.36	0.44	0.49	0.49	0.46	0.40	0.33	0.26
Consumer Sentiment	-0.21	-0.28	-0.35	-0.44	-0.55	-0.63	-0.70	-0.74	-0.76
S&P 500	0.24	0.26	0.25	0.19	0.09	-0.02	-0.12	-0.19	-0.22

Notes: Correlations of the share of positive responses to Michigan Survey and Business Outlook Survey (BOS) questions, the Index of Consumer Sentiment, and the S&P 500 with leads and lags of output measures: year-over-year growth in real GDP and industrial production or the unemployment rate (in levels). Quarterly, 1961–2017, except where otherwise indicated.

data[6] lists twelve categories of responses, one of which is labeled as "Bad Time—Uncertain Future." Leduc and Liu (2016) use this as their measure of uncertainty, the share of survey participants responding that now is a bad time to buy a car because the future is uncertain. The response categories in this table, however, are not the same as the coded responses that survey conductors record. There are, in fact, 77 different coded responses to this follow-up question.[7] Many of the twelve categories in the table are "bins" of these underlying responses grouped together; others are simply labeled differently. The underlying response to which "Bad Time—Uncertain Future" corresponds is actually coded as "People should save money, bad times ahead." While this response could indeed be capturing some measure of uncertainty, it is at least equally plausible that it is instead measuring mostly bad news about the future, rather than uncertainty *per se*. It is not perhaps not surprising, then, that shocks to this measure of "uncertainty" are strongly contractionary. The measures I construct are arguably more reasonable measures of uncertainty, and I compare them with Leduc and Liu's (2016) measure below.

The qualitative nature of the responses in the Michigan Survey of Consumers is similar to the Federal Reserve Bank of Philadelphia's Business Outlook Survey (BOS), which Bachmann, Elstner and Sims (2013) use to construct an uncertainty proxy based on disagreement. Businesses in that survey are asked whether they think general business activity will increase, decrease, or stay the same (importantly, the survey does *not* record mixed responses or any direct measure of uncertainty). Bachmann, Elstner and Sims (2013) use

[6] Available online at https://data.sca.isr.umich.edu/data-archive/mine.php.

[7] More detailed data from the survey beyond the headline time series can be accessed at https://data.sca.isr.umich.edu/sda-public/.

the following measure of disagreement as a proxy for uncertainty:

$$Uncertainty_t = \sqrt{Incr_t + Decr_t - (Incr_t - Decr_t)^2} \tag{3.1}$$

where $Incr_t$ ($Decr_t$) is the share of businesses in month t responding that general business activity will increase (decrease). It is straightforward to construct analogous measures of disagreement from the questions asked in the Michigan Survey on which my direct measures of uncertainty are based. This allows for a convenient comparison of direct uncertainty measures with disagreement-based proxies.

The dispersion measures constructed from the Michigan Survey have low (or negative) correlations with the direct measures. For BUS, VEHIC, DUR, and HOUSE the correlations are, respectively, -0.09, -0.24, -0.06, and 0.01. Although this is already suggestive that dispersion measures something different from uncertainty, I provide additional results in the next section on the estimated effects of shocks to these measures.

3.3 Estimating the effects of uncertainty shocks

In this section, I describe estimates of the effects of shocks to uncertainty, using these new measures, on the macroeconomy. I first describe identification based on a recursive ordering in a VAR, starting from Leduc and Liu's (2016) specification and assessing a number of extensions.[8] I then estimate the effects of uncertainty shocks using a "proxy SVAR," which uses external instruments, as in James H. Stock and Mark W. Watson (2012), Karel Mertens and Morten O. Ravn (2013), and Gertler and Karadi (2015). Identification in this setup is

[8]Most of the literature studying the effects of uncertainty shocks makes use of recursive ordering schemes. See, for example, Bloom (2009), Bachmann, Elstner and Sims (2013), Jurado, Ludvigson and Ng (2015), and Scott R. Baker, Nicholas Bloom and Steven J. Davis (2016).

achieved by isolating the variation in survey-based estimates of uncertainty that is due to fluctuations in broader measures of uncertainty taken from the literature.

3.3.1 Recursive VARs

As mentioned above, the most common approach to identifying uncertainty shocks has been by imposing short-run timing restrictions in VAR systems using the Cholesky decomposition. Bloom (2009) achieved identification in a medium-size VAR by arguing that uncertainty—as measured by stock market volatility—responds to stock-market fluctuations within a period (a month), but uncertainty shocks do not affect the stock market within a period, while all macroeconomic variables are allowed to respond contemporaneously to the uncertainty shock. Bachmann, Elstner and Sims (2013) estimate the effects of uncertainty shocks using both a series bivariate VARs with uncertainty measures order first and Bloom's (2009) larger VAR. Leduc and Liu (2016), in a smaller VAR, order their Michigan Survey-based uncertainty measure first.

As a first step and a point of comparison, I re-estimate Leduc and Liu's (2016) VAR, with the time period extended through 2017.[9] The system includes their measure of uncertainty, ordered first, followed by the unemployment rate, CPI inflation, and the three-month treasury rate. The data are monthly from 1978–2017. Impulse responses to a one standard deviation uncertainty shock are displayed in Figure 3.2. The shock leads to an increase in unemployment of about 0.15 percentage points and declines in inflation and short-term interest rates of similar magnitudes. The results are essentially identical to Leduc

[9] The sample period in their original VAR was 1978 through October 2013. The results are essentially unchanged by the inclusion of these additional years in the sample.

Figure 3.2: Impulse Responses to an Uncertainty Shock, Leduc and Liu (2016)

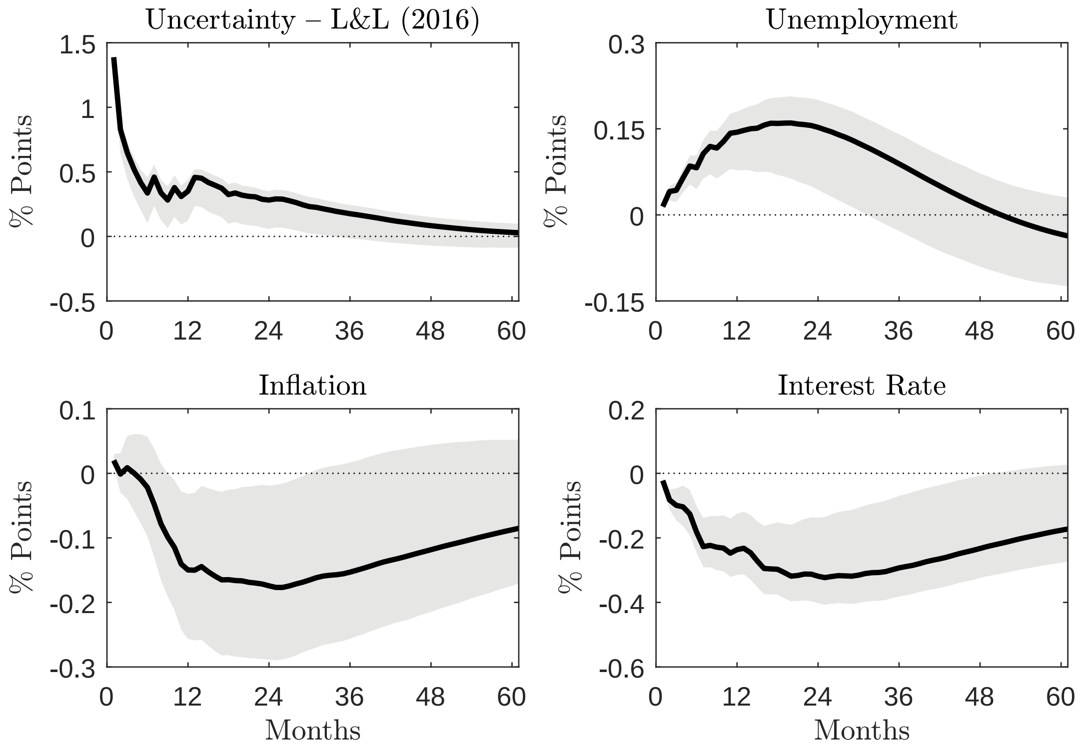

Impulse responses to an one-standard deviation shock to uncertainty, using Leduc and Liu's (2016) measure—the share of survey respondents answering that now is bad time to buy a car because "people should save more" or there are "bad times ahead." Error bands are 95% confidence intervals from a bootstrap.

and Liu's (2016) baseline estimates.

I next estimate a similar system using the measures of uncertainty from the Michigan Survey that were described above. Figure 3.3 displays impulse responses to the same system, using instead the share of "uncertain" responses to the questions described above, BUS, VEHIC, DUR, and HOUSE. An uncertainty shock using these measures is *expansionary* leading to a decline in unemployment and an increase in inflation and interest rates of similar magnitude to the *contractionary* effects estimated using Leduc and Liu's (2016) uncertainty measure.

Bachmann, Elstner and Sims (2013) provide an alternative survey-based measure of uncertainty. Using the Federal Reserve Bank of Philadelphia's Business Outlook Survey, they construct an uncertainty measure based on response dispersion. Their index of dispersion for a particular question is given by Equation 3.1. The Business Outlook Survey does not have a direct measure of uncertainty, however. The questions on the Michigan Survey allow for the construction of analogous dispersion indices, and comparison with the direct measures described above to assess the reliability of dispersion-based uncertainty proxies. This is of interest since dispersion of survey responses is not necessarily indicative of uncertainty. Survey response dispersion could also indicate differing—but precise—forecasts. Comparison with direct measures of uncertainty gives some indication of how good a proxy dispersion measures are for uncertainty.

Figure 3.4 displays impulse responses using Bachmann, Elstner and Sims's (2013) measure of dispersion from the Philadelphia Fed's Business Outlook Survey in the same VAR specification and time period as the estimates above.[10] A shock to uncertainty using their survey dispersion measure produces a small increase in the unemployment rate and small and insignificant changes in inflation and interest rates.

Figure 3.5 displays the impulse responses to analogous dispersion measures for the questions asked in the Michigan Survey. In three of the four cases, shocks to dispersion lead to increases in unemployment, as in Bachmann, Elstner and Sims (2013), while shocks to the direct measures of uncertainty are uniformly expansionary. This discrepancy suggests

[10]Note that the sample period differs from Bachmann, Elstner and Sims (2013), whose sample period runs from 1968 to 2011.

Figure 3.3: Impulse Responses to Uncertainty Shocks, New Direct Measures

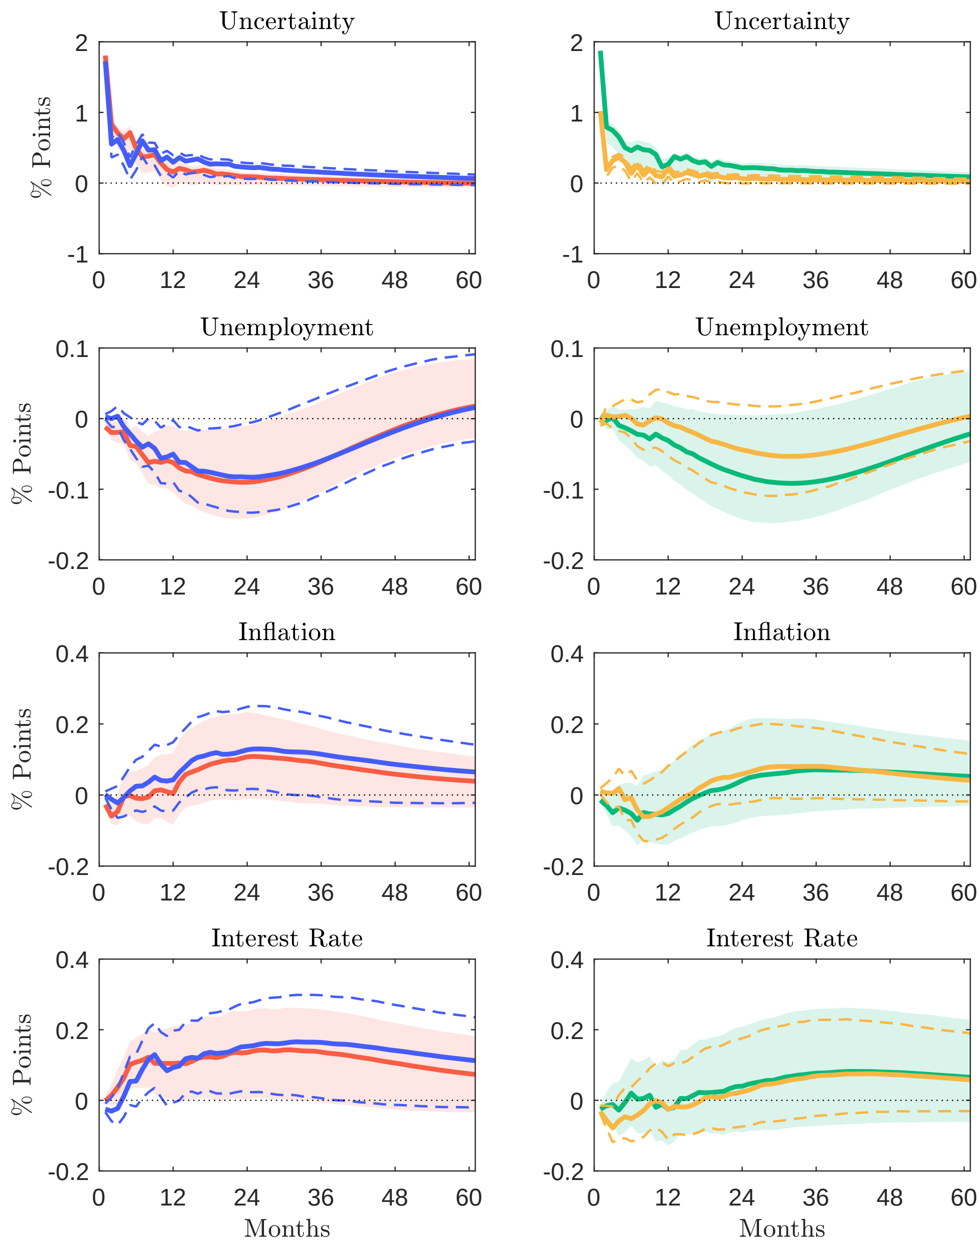

Impulse responses to a one-standard deviation shock to the share of BUSS and VEHIC (*left column*) and DUR and HOUSE (*right column*) respondents answering "uncertain," "don't know," or "pros and cons." Error bands are 95% confidence intervals from a bootstrap.

Figure 3.4: Impulse Responses to an Uncertainty Shock, Bachmann, *et al.* (2013)

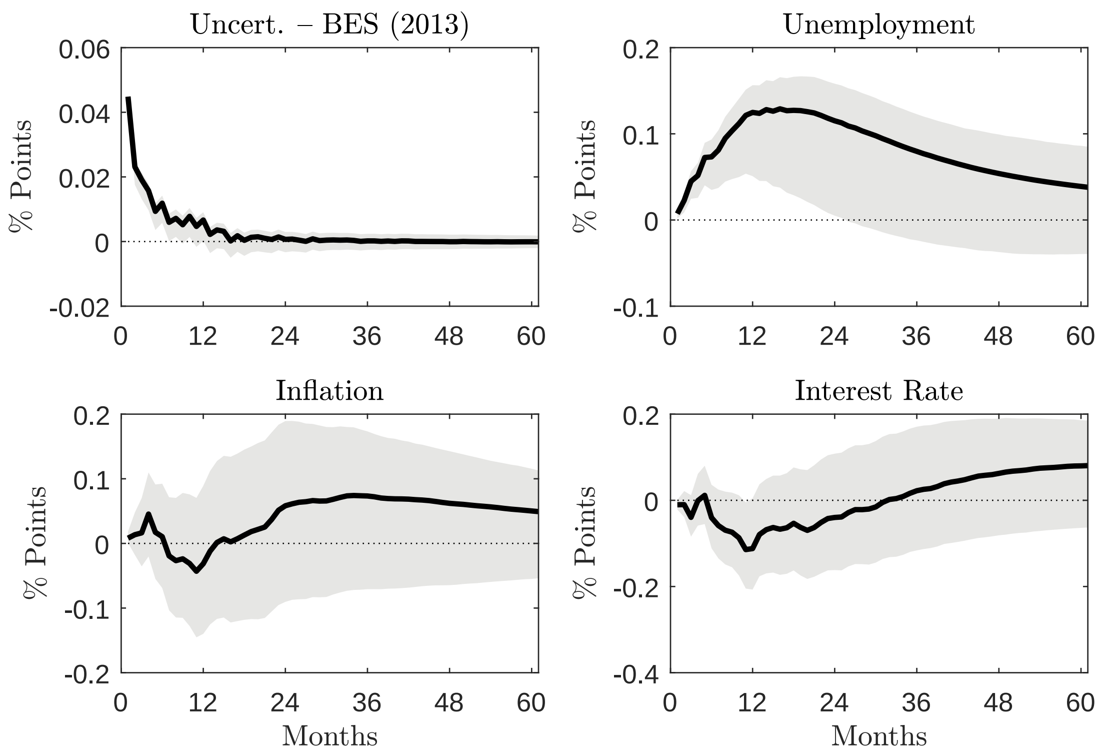

Impulse responses to a one-standard deviation shock to uncertainty, using Bachmann, Elstner and Sims's (2013) measure of survey dispersion from the Business Outlook Survey, 1978–2017. Error bands are 95% confidence intervals from a bootstrap.

caution using dispersion measures as proxies for uncertainty.[11]

One potential issue with these small-scale VARs is that they do not account fully for "first-moment" information contained in uncertainty measures. To account for this, Bloom (2009) includes the S&P 500 index in his VAR, while Leduc and Liu (2016) and Baker, Bloom and Davis (2016) include the Michigan Consumer Sentiment Index, constructed from the Michigan Survey, in their VARs. Estimates of the real effects of uncertainty shocks decline when these first-moment variables are included, indicating that uncertainty measures contain some first-moment information.

This same phenomenon occurs using the new measures of uncertainty constructed above. Indeed the effect of uncertainty shocks on unemployment are more modest and largely insignificant (although point estimates are still of the same sign). Figure 3.6 displays the impulse responses to the unemployment rate when including either the S&P 500 or Consumer Sentiment Index in the baseline VAR above. Bloom (2009) orders the S&P 500 before his volatility-based measure of uncertainty, while Leduc and Liu (2016) order sentiment after their uncertainty measure. Following Baker, Bloom and Davis (2016), I display the results for either ordering.[12] In nearly all cases, for either ordering or first-moment measure, the estimated effects from above are diminished.

These results suggest the importance of including first-moment measures in estimating and identifying uncertainty shocks. The restrictions implicit in the recursive identification scheme involving these variables, however, are strong and arguably unreasonable. When

[11] These results are in line with the finding of Robert Rich and Joseph Tracy (2010), who document differences between uncertainty and disagreement among professional forecasters.

[12] Neither condition—that the stock market or sentiment can affect uncertainty within a period but not *vice versa*, or the reverse restriction—is completely satisfactory. See Section 3.3.2, below.

Figure 3.5: Impulse Responses to Uncertainty Shocks, New Dispersions Measures

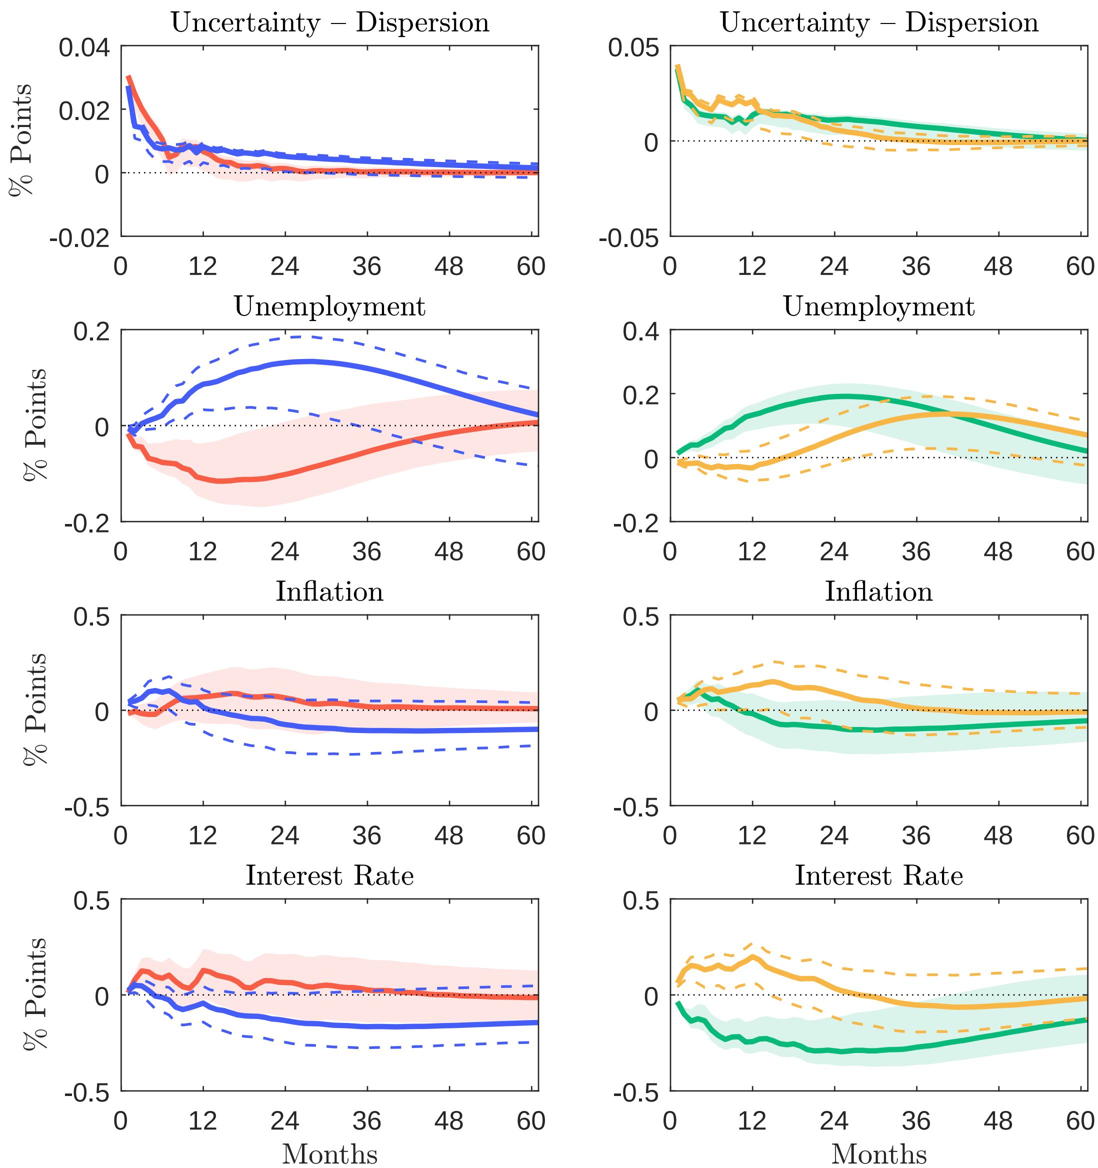

Impulse responses to a one-standard deviation shock to response dispersion of BUSS and VEHIC (*left column*) and DUR and HOUSE (*right column*), as computed from Equation 3.1. Error bands are 95% confidence intervals from a bootstrap.

Figure 3.6: Impulse Responses of Unemployment to Uncertainty Shocks with First-Moment Controls

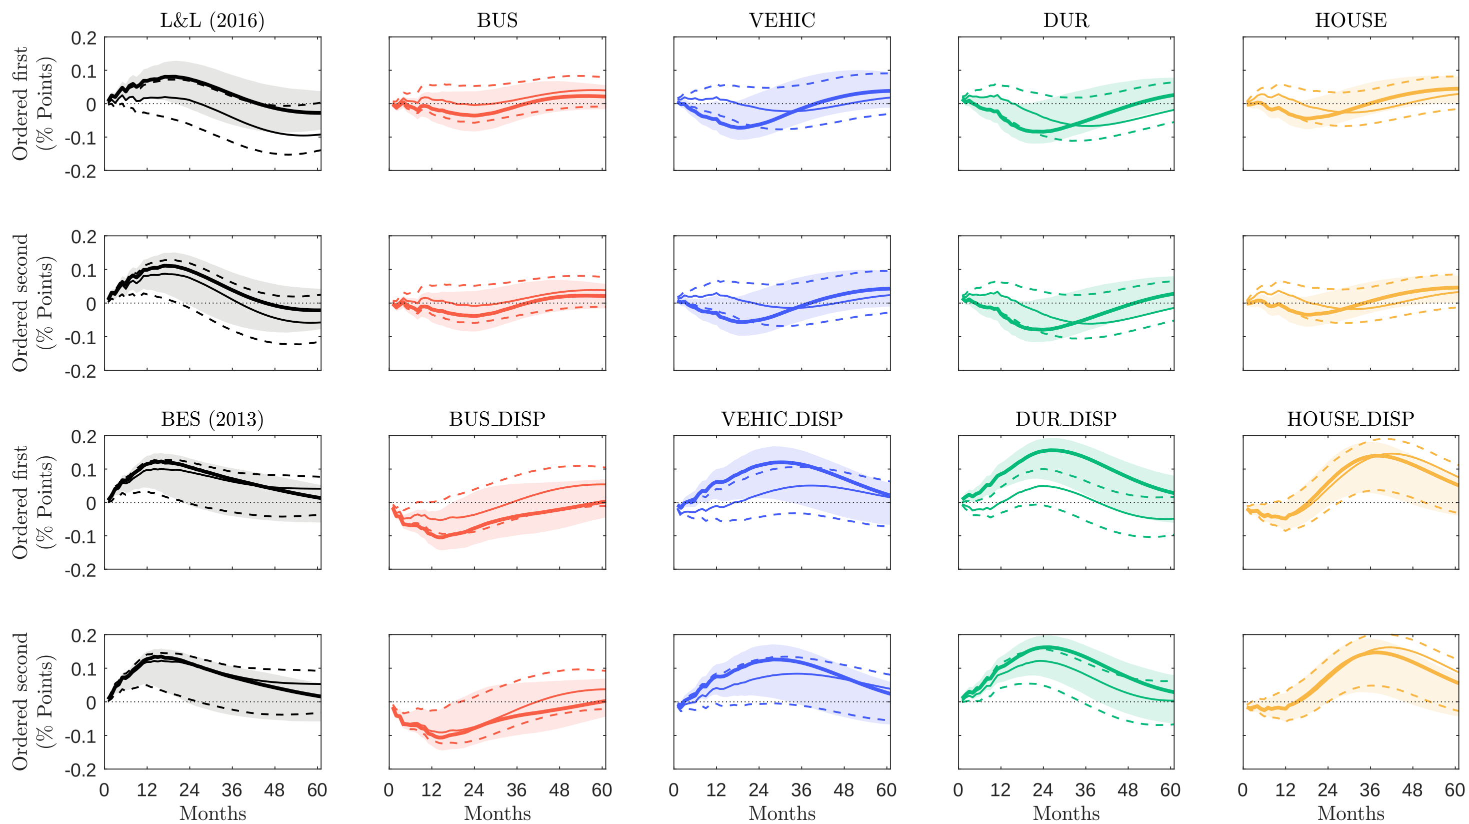

Impulse responses of the unemployment rate to uncertainty shocks with "first-moment" variables included. First and third rows have first-moment variables ordered first (before uncertainty). Second and fourth rows have first-moment variables ordered second (after uncertainty). Thick lines and shaded error bands: S&P 500. Thick lines and dashed error bands: Index of Consumer Sentiment. *Upper panel:* Leduc and Liu (2016) and direct measures introduced above. *Lower panel:* Bachmann, Elstner and Sims (2013) and dispersion measures described above. Error bands are 95% confidence intervals from a bootstrap.

uncertainty is ordered after the first-moment variable, uncertainty shocks are assumed not to affect that variable within the month, but the first-moment variable's level is controlled for in estimating the shock; when uncertainty is ordered before the first-moment variables, that variable's level is *not* accounted for in estimating the shock, but uncertainty shocks can affect it within the period. Neither of these conditions is entirely realistic. Financial variables such as the stock market or expectations-based variables like consumer sentiments are likely to respond quickly to any shock. Therefore, in the next section I make use of external instruments to estimate proxy SVARs in which these unsatisfactory restrictions can be relaxed.

3.3.2 Proxy VARs

The restrictions underlying recursive identification schemes in the previous section can be dispensed with if one uses the "external instruments" approach pioneered by Stock and Watson (2012) and Mertens and Ravn (2013).[13] This approach requires the use of an instrument for the structural shock obtained from outside the VAR system. External instruments obviate the need for arguably unrealistic timing restrictions in recursive identification schemes; in particular, in the setting described above, they allow for the incorporation of multiple fast-moving "first-moment" variables such as consumer sentiment, stock market levels, or measures of credit conditions.

The appealing features of this approach do not come without costs, however; external

[13]Many authors have used external instruments in other contexts, including Gertler and Karadi (2015), who identify monetary policy shocks using a hybrid method of high-frequency identification and VAR estimation, and James H. Stock and Mark W. Watson (2018) who discuss the use of external instruments for identification in macroeconomics more generally.

instruments must satisfy conditions analogous to those encountered in the usual setting of estimating causal effects by instrumental variables. In particular, they must satisfy instrument relevance and exclusion restriction conditions. Letting $\mathbf{Z}_t$ denote a vector of instrumental variables and $\boldsymbol{\varepsilon}_t = [\varepsilon_t^u \; \boldsymbol{\varepsilon}_t^o]'$ denote a partitioned vector of the structural shocks affecting the system, where ε_t^u is the structural shock of interest, these conditions can be written as

$$\mathbb{E}\left[\mathbf{Z}_t \varepsilon_t^{u\prime}\right] = \boldsymbol{\Sigma}$$

$$\mathbb{E}\left[\mathbf{Z}_t \boldsymbol{\varepsilon}_t^{o\prime}\right] = \mathbf{0}.$$

The first condition says that the instrument is correlated with the shock of interest, while the second says that it is orthogonal to all other structural shocks.

With a set of valid instruments in hand, estimation is straightforward. First, estimate the VAR system to obtain reduced-form residuals $\mathbf{u}_t$, then regress the residuals associated with the shock measures (i.e., uncertainty), u_t^u on the instrument set to obtain $\widehat{u}_t^u$. Intuitively, this isolates the portion of the residual that is due to the (unobserved) structural shock ε_t^u via the instruments $\mathbf{Z}_t$. Combining estimated coefficients from regressing $\mathbf{u}_t^o$ on $\widehat{u}_t^u$ with a variance normalization identical to that in recursively identified systems gives an estimated linear relationship between the reduced-form residuals and the structural shock: $\varepsilon_t^u = \widehat{\boldsymbol{\gamma}}'\mathbf{u}_t$. Construction of impulse response functions and other objects of interest follows the same process as in other SVARs.[14]

The difficulty, however, is find a valid set of instruments for the shock of interest. This is particularly true in the case of uncertainty, most measures of which are expected to move

[14] The method is described in detail by Mertens and Ravn (2013) and Gertler and Karadi (2015). Inference involves incorporating the first stage into a wild bootstrap.

endogenously in response to other "first-moment" shocks, evidence of which was presented in Section 3.3.1. Gertler and Karadi (2015) are able to identify monetary policy shocks using price changes of Fed Funds futures contracts around a 30-minute window of Federal Reserve policy announcements, while Stock and Watson (2012) use a variety of externally-identified shock series, including uncertainty. Since there is no obvious high-frequency series to instrument for uncertainty, I follow the latter approach but do so making use of a new dataset that was unavailable to Stock and Watson (2012).

As instruments for uncertainty, they use innovations to stock market volatility and Baker, Bloom and Davis's (2016) index of policy uncertainty. In contrast, I use the macroeconomic uncertainty indices of Jurado, Ludvigson and Ng (2015) (JLN) as instruments for uncertainty shocks.[15] The JLN series is constructed from a large dataset of hundreds of macroeconomic times series. They motivate the construction of their series by the observation that economic decisions are not made based on volatility or dispersion, but rather whether the economy "has become more or less *predictable*; that is, less or more uncertain."(Jurado, Ludvigson and Ng, 2015). Their measures of uncertainty are weighted aggregates of the expected squared forecast errors of many macroeconomic times series at various horizons. They argue that this measure isolates the purely unforecastable component of the macroeconomy.

To be sure, these are not perfect instruments. To the extent that uncertainty responds endogenously to other events in the economy, it is difficult to argue that they are perfectly orthogonal to other macroeconomic shocks. Therefore, I follow Stock and Watson (2012)

[15] I have also used the VIX and the Economic Policy Uncertainty index as instruments, but they are very weak and give impulse responses that are insignificantly different from zero at all horizons.

and use the residuals from estimated AR(2) processes on the one-, three-, and twelve-month horizon JLN series as instruments for a shock to uncertainty. I estimate both the baseline four-variable VAR from Leduc and Liu (2016) and a larger VAR system that includes the S&P 500 or the Michigan Index of Consumer Sentiment and Gilchrist and Zakrajsek's (2012) excess bond premium—a measure of credit conditions—in addition to the variables in the baseline VARs estimated above.

Figure 3.7 displays the impulse responses from the baseline four-variable VARs with shocks identified using JLN's macroeconomic uncertainty indices as external instruments. For the two series most correlated with economic activity—DUR and BUS—unemployment falls and interest rates rise or remain unchanged. For BUS, inflation increases slightly, while it falls for DUR. For both VEHIC and HOUSE, unemployment increases, while inflation rises and interest rates remain largely unchanged. The expansionary effects of uncertainty shocks measured by BUS or DUR are consistent with Ludvigson, Ma and Ng (2018), who argue that recursively identified systems are invalid and who also find that shocks to macroeconomic uncertainty are mildly expansionary.[16] The impulse responses of unemployment in the larger proxy VAR are displayed in Figure 3.8, and are essentially the same as in the baseline proxy VAR. In all cases, first-stage F-statistics are indicative of weak instruments; this is especially true for the specification using HOUSE as the uncertainty measure, which has a first-stage F-statistic less than one.[17]

[16] They argue that the increases in macroeconomic uncertainty during recessions are mostly an endogenous response to other shocks, but they do find a large role for *financial* uncertainty shocks.

[17] In all cases, first-stage F-statistics are below the typical threshold of ten. However, as Gertler and Karadi (2015) observe, impulse responses are little changed when using similar, but weaker, instruments with F-statistics around 3 or 4.

Figure 3.7: Impulse Responses to Uncertainty Shocks, Small Proxy VAR

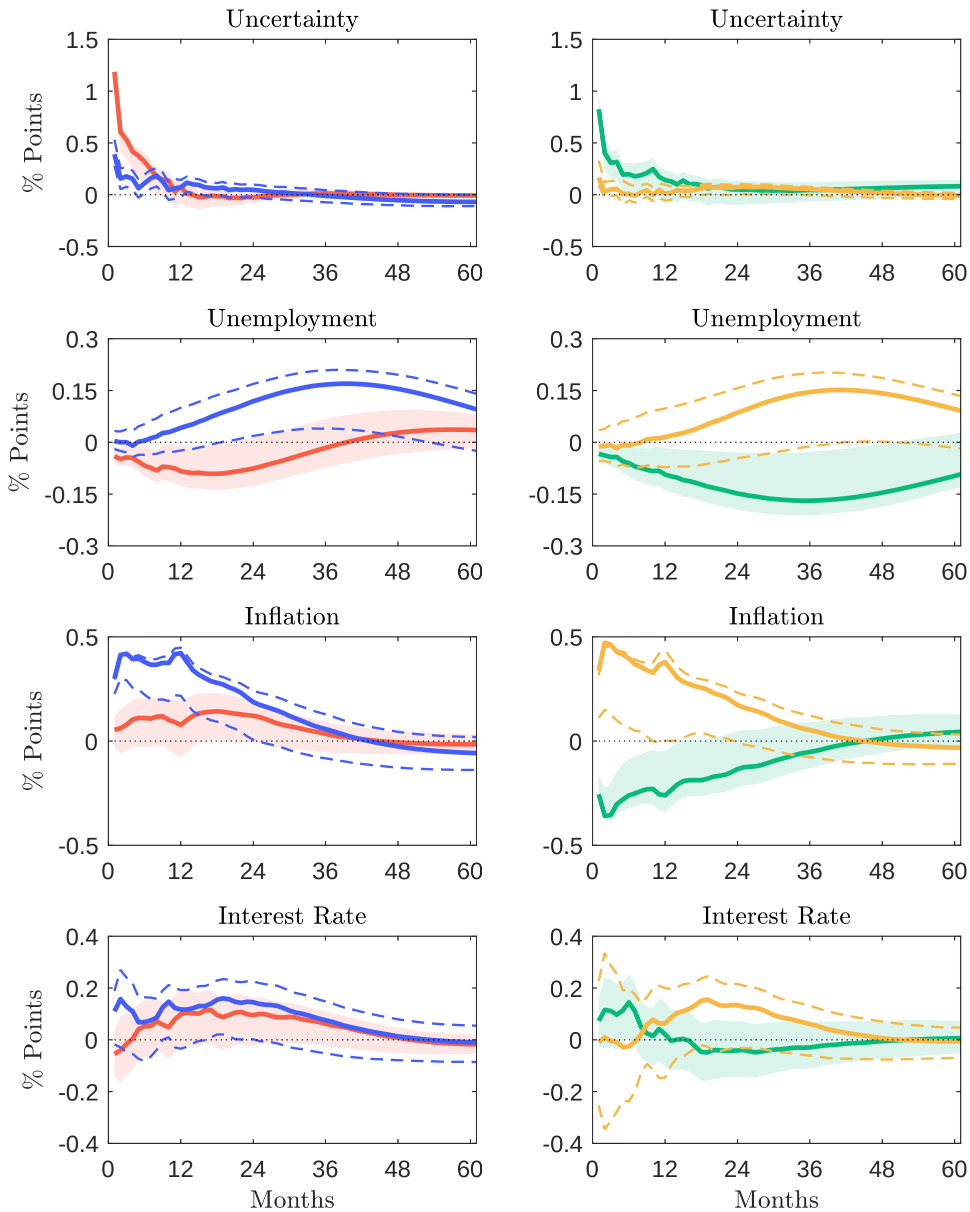

Responses to uncertainty shocks measured by BUS, VEHIC, DUR, HOUSE in a small proxy VAR. External instruments are JLN's macroeconomic uncertainty index at one-, three-, and twelve-month horizons, residuals from an estimated AR(2) process. First-stage F-statistics: 4.37, 1.12, 2.31, 0.23. Error bands are 95% confidence intervals from a bootstrap.

Figure 3.8: Impulse Responses of Unemployment to Uncertainty Shocks, Large Proxy VAR

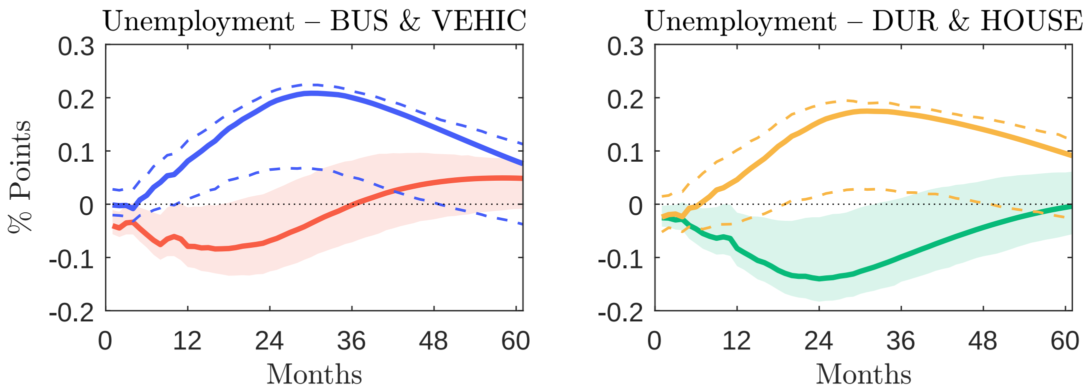

Responses of the unemployment rate to uncertainty shocks measured by BUS, VEHIC, DUR, HOUSE in a larger proxy VAR, including the S&P 500 and Gilchrist and Zakrajsek's (2012) excess bond premium. External instruments are JLN's macroeconomic uncertainty index at one-, three-, and twelve-month horizons, residuals from an estimated AR(2) process. First-stage F-statistics: 3.64, 1.43, 2.89, 0.68. Error bands are 95% confidence intervals from a bootstrap.

Overall, the results from VARs are mixed. In recursively-identified systems, shocks to any of the new measures of uncertainty results in lower unemployment, although the magnitude of the decline is somewhat sensitive, and the identifying assumptions are dubious. In systems estimated by external instruments, shocks to BUS and DUR result in lower unemployment, but the reverse is true for VEHIC and HOUSE. However, in addition to the difficulty of finding valid instruments to begin with, these estimates suffer from weak instruments.

Taken altogether, the results from this section are suggestive of mildly expansionary effects of uncertainty shocks, at least with respect to the series most closely correlated with economic activity. These results are consistent with "growth options" theories of uncertainty. In such settings, if downside risk is bounded while upside risk is potentially unbounded, increases in uncertainty can be expansionary—intuitively, a widening distribution of outcomes leads to a heavier right tail, while the lower bound limits the widening of the left tail.

3.4 Conclusion

In this paper, I introduced new measures of uncertainty based on the Michigan Survey of Consumers. They are direct measures of uncertainty, in the sense that they represent the share of respondents who said they were uncertain about macroeconomic conditions. Responses to questions about business conditions or durable goods are strongly correlated with future economic activity; positive views of business conditions are more strongly correlated with current and future economic activity than the S&P 500 or positive responses to the Business Outlook Survey, two series from which other uncertainty measures have been constructed.

I also construct measures of survey response dispersion from the Michigan Survey. Comparison with the direct measures suggests that dispersion is a poor proxy for uncertainty. Dispersion measures have low (or negative) correlation with uncertainty when both are constructed from the same series. In addition, the effects of uncertainty shocks measured by dispersion or the direct measures are of opposite sign.

Evidence from SVARs points to uncertainty shocks having mildly expansionary effects. In recursively identified systems shocks that increase uncertainty lead consistently to small declines in unemployment and increases in inflation and interest rates. Proxy VAR systems using external instruments yield mixed results. I find expansionary effects of uncertainty shocks when using the measures most correlated with economic activity, but contractionary effects for other measures. These mixed results and their sensitivity to different identifying assumptions should motivate future research to focus on cleaner identification of uncertainty shocks.

Appendices

Appendix A

Appendix to Chapter 1

A.1 Data Adjustments for Time Aggregation

Previous work has attempted to account for the time-aggregation problem in the following way. If one views the transition rates (that is, the number of transitions that occur to each state expressed as a percentage of the pool of workers in the beginning state) as discrete-time Markov transition probabilities, while the underlying true flow pattern is governed by a continuous-time Markov chain, as shown by Shimer (2012), generically, there is a one-to-one mapping between the implied transition matrices. That is, the instantaneous flow hazard rates can be backed out from the directly observed gross flows.

Following Shimer (2012), assume the *observed* month-to-month transition rates define a 3×3 Markov matrix, so that the evolution of the observed numbers of workers employed, unemployed, and not in the labor force evolve as a discrete-time Markov chain:

$$s_{t+1} := \begin{bmatrix} E_{t+1} \\ U_{t+1} \\ N_{t+1} \end{bmatrix} = \begin{bmatrix} \pi_t^{EE} & \pi_t^{UE} & \pi_t^{NE} \\ \pi_t^{EU} & \pi_t^{UU} & \pi_t^{NU} \\ \pi_t^{EN} & \pi_t^{UN} & \pi_t^{NN} \end{bmatrix} \begin{bmatrix} E_t \\ U_t \\ N_t \end{bmatrix} =: \Pi_t s_t, \tag{A.1}$$

where π_t^{XY} denotes the observed transition rate in month t from state X to state Y. However, these are just the discrete observations. Assume that the true (but unobserved) labor market evolves according to a continuous-time Markov chain

$$\dot{s}_{t+\tau} := \frac{d}{d\tau} \begin{bmatrix} E_{t+\tau} \\ U_{t+\tau} \\ N_{t+\tau} \end{bmatrix} = \lambda_t s_{t+\tau},\ \forall \tau \in [0, 1), \tag{A.2}$$

where λ_t is such that it produces the observed discrete-time process given by Equation A.1. The mapping between the matrix of observed month-to-month transition rates (Π_t) and the flow transition hazards that are the off-diagonal terms of λ_t is derived below; it is based on a simple decomposition of Π_t into a product of matrices of its own eigenvalues and eigenvectors. Given the matrices λ_t, the implied month-to-month transition probabilities after correcting for this time-aggregation issue are therefore given by $1 - \exp(-\lambda_t^{ij})$ for $i \neq j$.

Derivation of λ from Π

The derivation of λ_t, the matrix of instantaneous transition hazards, from Π_t, the matrix of (discretely) observed month-to-month transition rates, follows Shimer (2012), with the notation slightly adjusted to that I use in this paper. Beginning with the discrete-time process in Equation A.1, divide each time period into Δ intervals where the relationship between states $\frac{1}{\Delta}$ time apart for some $\tau \in [0, 1)$ is given by

$$s_{t+\tau+\frac{1}{\Delta}} = \Pi_{t,\Delta} s_{t+\tau}, \tag{A.3}$$

for some matrix $\Pi_{t,\Delta}$. Starting at $\tau = 0$ and iterating this expression forward Δ times gives

$$s_{t+1} = \Pi_{t,\Delta}^{\Delta} s_t, \tag{A.4}$$

establishing a relationship between Π_t and $\Pi_{t,\Delta}$, i.e., $\Pi_t = \Pi_{t,\Delta}^{\Delta}$. Subtracting s_t from both sides of Equation A.3 at $\tau = 0$ and dividing both sides by $\frac{1}{\Delta}$ gives

$$\frac{s_{t+\frac{1}{\Delta}} - s_t}{\frac{1}{\Delta}} = \frac{\left(\Pi_{t,\Delta} - I_{[3\times 3]}\right) s_t}{\frac{1}{\Delta}}, \tag{A.5}$$

where $I_{[n\times n]}$ denotes the $n \times n$ identity matrix. The limit as $\frac{1}{\Delta} \to 0$ of the left-hand side is simply $\dot{s}_t$. Therefore, the limit of the term multiplying s_t on the RHS is λ_t. That is, re-writing in terms of the *observed* Π_t,

$$\lambda_t = \lim_{\frac{1}{\Delta}\to 0} \frac{\Pi_t^{\frac{1}{\Delta}} - I_{[3\times 3]}}{\frac{1}{\Delta}}. \tag{A.6}$$

The matrix Π_t can be written as $\Pi_t = V_t \Lambda_t V_t^{-1}$ where Λ_t is a diagonal matrix of the eigenvalues of Π_t and V_t is the matrix whose columns are the corresponding eigenvectors. Plugging in this for Π_t above and pre- and post-multiplying each side by V_t^{-1} and V_t, respectively, gives

$$V_t^{-1}\lambda_t V_t = \lim_{\frac{1}{\Delta}\to 0} \frac{\Lambda_t^{\frac{1}{\Delta}} - I_{[3\times 3]}}{\frac{1}{\Delta}}. \tag{A.7}$$

Since Λ_t is diagonal, all off-diagonal terms on the right-hand side will be zero in the limit as well, and each diagonal term is given by $\lim_{\frac{1}{\Delta}\to 0} \frac{(\Lambda_t^{ii})^{\frac{1}{\Delta}} - 1}{\frac{1}{\Delta}} = \ln\left(\Lambda_t^{ii}\right),\ i = 1, 2, 3$. Therefore, we have

$$\lambda_t = V_t \Lambda_t^{\log} V_t^{-1}, \tag{A.8}$$

where $\Lambda_t^{\log}$ is the matrix with diagonal terms given by $\ln\left(\Lambda_t^{ii}\right)$ and zeros everywhere else. This establishes the mapping between the observed transition rates and the unobserved transition probabilities that are used in the regressions in this paper.

A.2 Derivation of Stock Responses from Flows

This section is reproduced from and closely follows Elsby, Hobijn and Şahin (2015), adjusted to the notation I use in throughout this paper.

The first step is to normalize $E_t + U_t + N_t \equiv 1$ for all t, so each stock variable is a share of the total population, which is normalized to 1. Next, rewrite Equations A.1 and

A.2 in terms of the new state vector containing only E and U. These equations become

$$\tilde{s}_{t+1} := \begin{bmatrix} E_{t+1} \\ U_{t+1} \end{bmatrix} = \begin{bmatrix} \pi_t^{EE} - \pi_t^{NE} & \pi_t^{UE} - \pi_t^{NE} \\ \pi_t^{EU} - \pi_t^{NU} & \pi_t^{UU} - \pi_t^{NU} \end{bmatrix} \tilde{s}_t + \begin{bmatrix} \pi_t^{NE} \\ \pi_t^{NU} \end{bmatrix} =: \tilde{\Pi}_t \tilde{s}_t + \xi_t, \tag{A.9}$$

and

$$\dot{\tilde{s}}_t := \begin{bmatrix} \dot{E}_t \\ \dot{U}_t \end{bmatrix} = \begin{bmatrix} \lambda_t^{EE} - \lambda_t^{NE} & \lambda_t^{UE} - \lambda_t^{NE} \\ \lambda_t^{EU} - \lambda_t^{NU} & \lambda_t^{UU} - \lambda_t^{NU} \end{bmatrix} \tilde{s}_t + \begin{bmatrix} \lambda_t^{NE} \\ \lambda_t^{NU} \end{bmatrix} =: \tilde{\lambda}_t \tilde{s}_t + \epsilon_t, \tag{A.10}$$

where λ^{EE} and λ^{UU} are such that the columns of λ (the matrix of transition hazards in the 3-state, un-normalized setting) sum to zero. Equation A.9 gives the evolution of the observed unemployment and employment levels. The impulse responses, however, give the response of $1 - \exp(-\lambda^{ij})$, the probability a transition from I to J occurs during a month, conditional on having begun the month in state I. These obviously define unique responses of each λ^{ij}, so all that is needed is to translate these responses to the π^{ij} terms. This is straightforward using the mapping from Π to λ above.

Specifically, since the impulse response functions are for the variables $x^{ij} = 1 - \exp(-\lambda^{ij})$, we can write $\lambda^{ij} = -\ln(1 - x^{ij})$. Then the response of λ^{ij} can be found by noting that up to first order

$$\Delta\lambda^{ij} = \frac{1}{1 - x^{ij}} \Delta x^{ij} = \frac{\Delta x^{ij}}{\exp(-\lambda^{ij})}. \tag{A.11}$$

Using $\lambda_t^{ij} = \lambda_{t-1}^{ij} + \frac{\Delta x_t^{ij}}{\exp\left(-\lambda_{t-1}^{ij}\right)}$ as the entries for λ_t and using the mapping from the previous section gives the sequence of π_t^{ij} terms from which $\tilde{\Pi}_t$ and ξ_t can be calculated directly, giving the implied paths for $\tilde{s}_t$ and $\tilde{s}_t^*$ (since $\tilde{s}_t^* = \left(I_{[2\times 2]} - \tilde{\Pi}_t\right)^{-1} \xi_t = -\tilde{\lambda}_t^{-1} \epsilon_t$).

A.3 Romer and Romer's (2004) Monetary Policy Shock Series

In Section 1.3, I use a measure of monetary policy shocks developed by Romer and Romer (2004) and extended through 2007 to estimate the response of worker flows. Romer and Romer (2004) identify monetary policy shocks as changes to the Federal Funds target rate not predictable by the economic information in the Federal Reserve's "Greenbook" forecasts. Specifically, their monetary policy shock series is given by the residuals of the following regression:

$$\begin{aligned}\Delta ff_m = \alpha + \beta ffb_m \quad &+ \quad \sum_{i=-1}^{2} \gamma_i \widetilde{\Delta y}_{mi} + \sum_{i=-1}^{2} \lambda_i \left(\widetilde{\Delta y}_{mi} - \widetilde{\Delta y}_{m-1,i} \right) \\ &+ \quad \sum_{i=-1}^{2} \varphi_i \tilde{\pi}_{mi} + \sum_{i=-1}^{2} \theta_i \left(\tilde{\pi}_{mi} - \tilde{\pi}_{m-1,i} \right) + \rho \tilde{u}_{m0} + \varepsilon_m, \end{aligned} \tag{A.12}$$

where m indexes FOMC meeting dates, ffb_m denotes the level of the Federal Funds target rate at the time of meeting m, $\widetilde{\Delta y}_{mi}$ denotes forecasts of real output growth, $\tilde{\pi}_{mi}$ denotes forecasts of inflation, $\tilde{u}_{m0}$ denotes forecasts of current unemployment, and ε_m, the residual, is the monetary policy shock. The index i is the horizon of the forecast, and horizon $i = -1$ may be a true forecast or a realized value of the variable, depending on when the actual data were available.

A.4 Alternative Estimation and Identification

A.4.1 Estimation by Jordà's (2005) Method of Local Projections

As a robustness check, I have estimated the responses to monetary policy shocks using Jordà's (2005) method of local projections. The results are displayed below and are broadly similar to those estimated in the main text.

Figure A.1: Impulse Responses Estimated by Local Projections

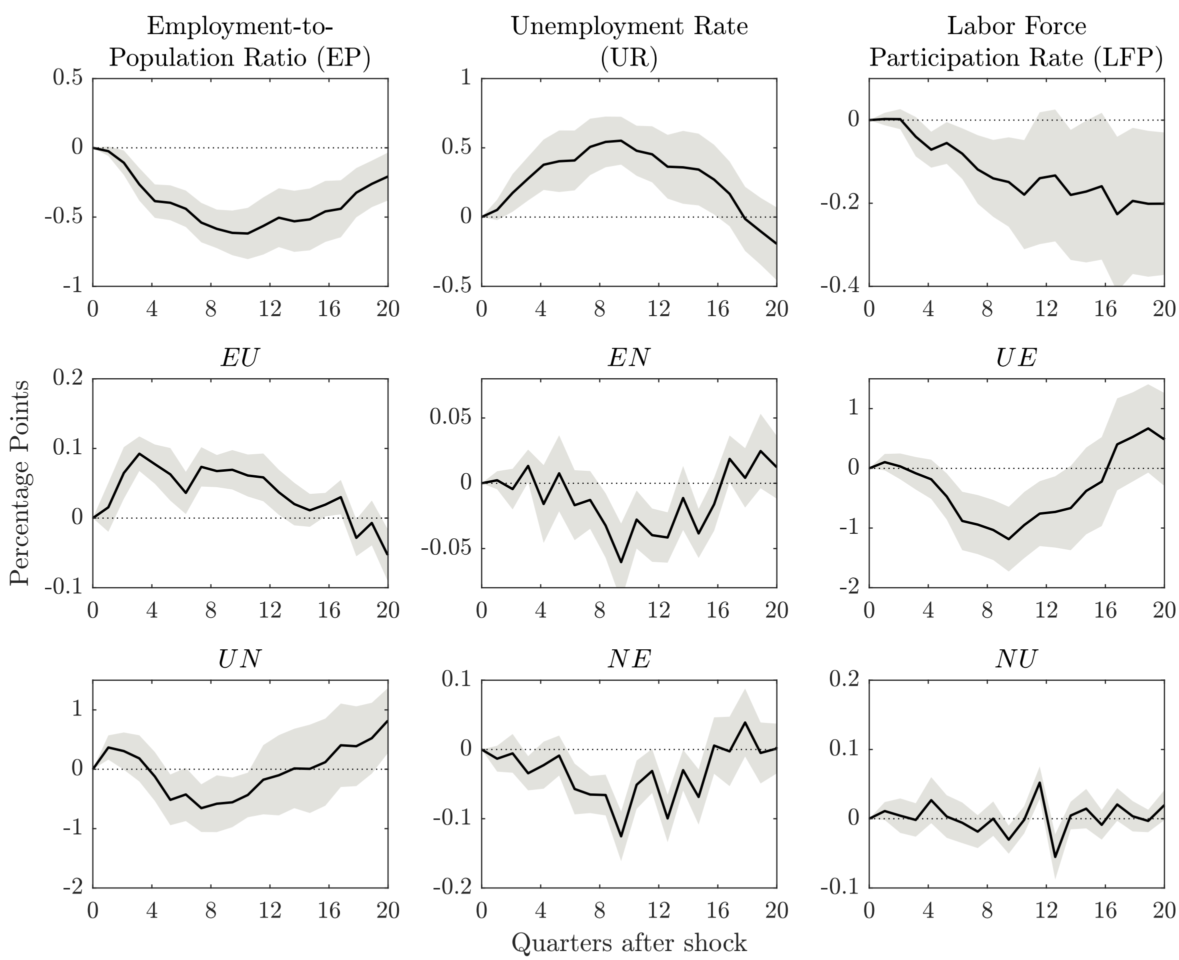

Impulse responses of stock (top row) and flow (bottom two rows) variables to a 100 b.p. contractionary Romer and Romer (2004) monetary policy shock, 1969Q1–2007Q3. Local projection estimates. Shaded areas are one standard deviation intervals.

A.4.2 High-Frequency Identification and Estimation by Proxy VAR

As another robustness check, I use Gertler and Karadi's (2015) hybrid-VAR method that utilizes high-frequency identification strategies combined with standard VAR techniques. I estimate the responses to a shock that has the same effect on impact to the one-year Treasury rate as the Romer and Romer (2004) shock. I include each measure in an independently estimated VAR in addition to the variables in Gertler and Karadi's (2015) VAR. The results are displayed below and are broadly consistent with the results in the main text.

Figure A.2: Impulse Responses Estimated by Proxy VAR

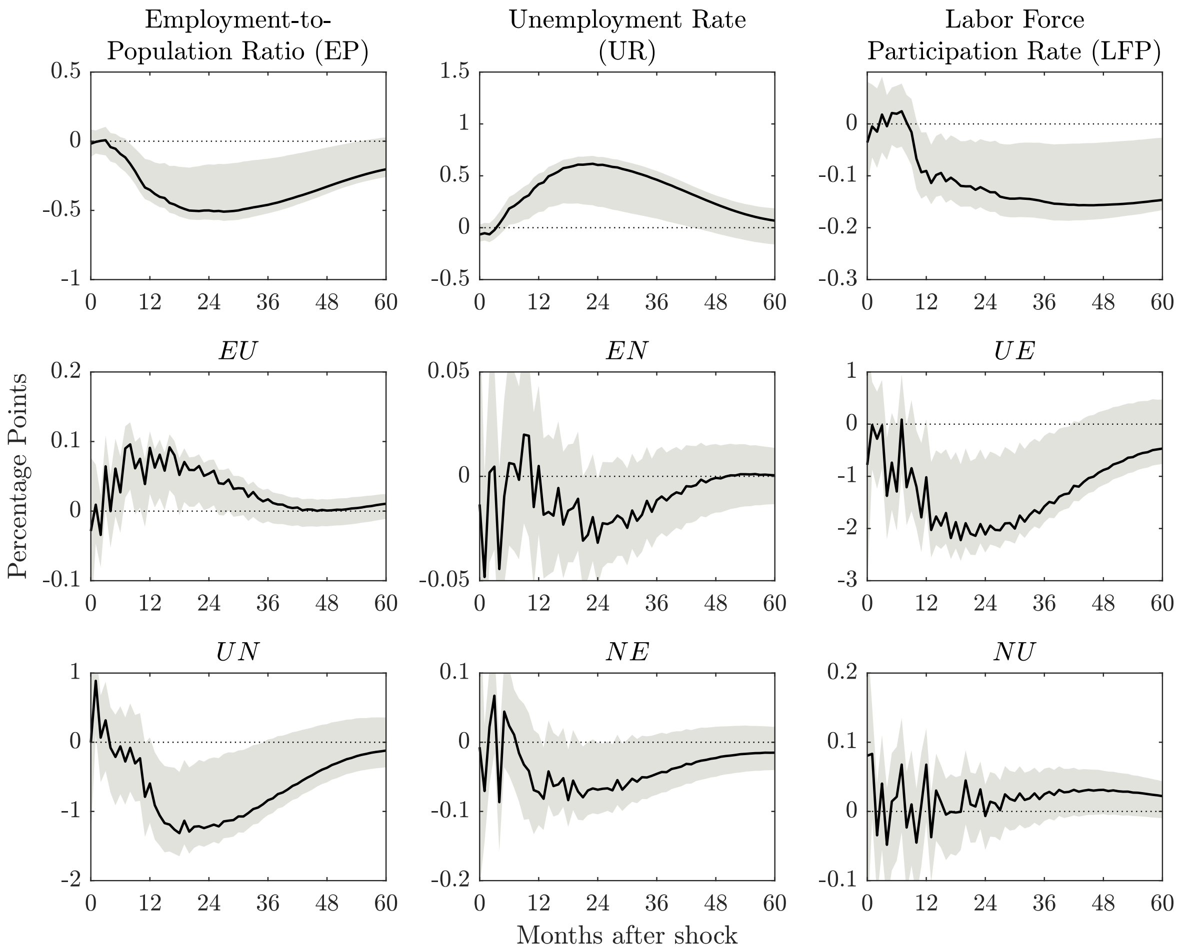

Impulse responses of stock (top row) and flow (bottom two rows) variables to a contractionary monetary policy shock from Gertler and Karadi (2015) that increases the 1-year nominal interest rate by 100 b.p. on impact. July, 1979 through June, 2012, monthly data. Shaded areas are 95 percent confidence intervals from a bootstrap.

Appendix B

Appendix to Chapter 2

B.1 Alternative Monetary Policy Shock Identification

As a robustness check to the linear local projections, as well as to include the post-ZLB period in the estimates, I estimate a VAR, identifying a monetary policy shock using external instruments as in Gertler and Karadi (2015).

The Gertler and Karadi (2015) approach to identifying a monetary policy shock is to use "surprise" changes to Fed Funds futures in a thiry-minute window around Federal Open Market Committee (FOMC) policy announcements as instruments for changes in a short-term interest rate (the Fed Funds Rate, the one-year or two-year treasury); such changes in the short-term rate were unanticipated by markets, and are therefore exogenous and due to policy announcements. The data on the futures contracts used as instruments is only available from the mid-1990s, however. Gertler and Karadi's (2015) innovation over other high-frequency approaches was to apply the high-frequency identification strategy to a longer horizon. They do so by estimating the linear relationship between the residuals from the reduced form VAR estimates from the shorter period (1991–2012), then assume that this relationship between reduced form residuals and the structural monetary policy shock is the same for the extended period (1979–2012).[1]

[1] Since the VARs estimated here are essentially identical to the baseline in Gertler and Karadi (2015),

Their baseline VAR includes the one-year treasury rate, (log) industrial productions, (log) consumer price index, and Gilchrist and Zakrajsek's (2012) excess bond premium, a measure of credit conditions; the time period they consider is July 1979 through June 2012. For the results below, I estimate this same system over a longer period that begins in January 1973,[2] but I include additional variables one by one, as in Lawrence J. Christiano, Martin Eichenbaum and Charles Evans (1996), reestimating the VAR (including shock identification) each time.

The impulse responses to a contractionary Gertler and Karadi (2015) shock that raises the one-year Treasury rate by 100 b.p. impact are displayed in Figure B.1. The estimated magnitudes are smaller than, but broadly consistent with, the IRFs estimated in Section 2.4.1. The response of routine employment is more than twice as large as the response of nonroutine employment. The difference in magnitudes between responses identified in a VAR and those identified using a narrative approach as in Romer and Romer (2004) are discussed in Coibion (2012).

B.2 Alternative Lag Structure

This section considers some alternative lag structures for estimates obtained from Equation 2.3 to understand the somewhat anomalous responses to expansionary shocks. The results are qualitatively similar to the baseline estimates, aside from the strength of this anomaly.

---

apart from trivially adding an additional variable to the system, for the details of the estimation, the reader is referred to their paper.

[2] The results are essentially identical for the Gertler and Karadi time period.

Figure B.1: Impulse Responses Estimated by Proxy VAR—Occupations

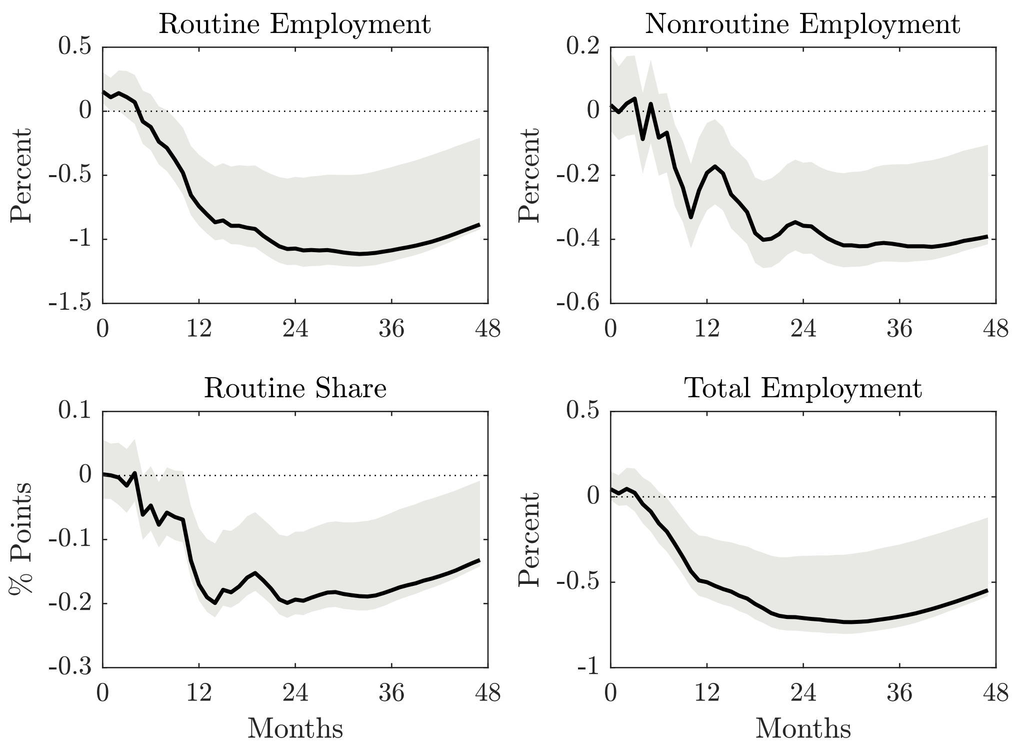

Impulse responses of occupation employment to a contractionary Gertler and Karadi (2015) monetary policy shock that increases the one-year rate 100 b.p. on impact. Estimated from the five-variable VAR discussed in Section B.1. Shaded areas are 90 percent bootstrap confidence intervals.

Recall that the baseline estimates include in the vector of controls $\mathbf{x}_t$ one year of lagged changes in the dependent variable and one quarter of lags of the shocks. Here I consider two alternative sets of controls. Specifically, I consider a vector of controls $\mathbf{x}_t^{12,1}$ which includes a year of lags of the dependent variable and no lags of the shock, and another alternative $\mathbf{x}_t^{24,36}$, which includes two years of lags of the dependent variable and three years of lags of the shock. This lag structure is based on Romer and Romer's (2004) original impulse response estimation for output.

As is evident in Figure B.2, with fewer lags of the shock included, the contractionary effect of nominally "expansionary" shocks is exacerbated. "Expansionary" shocks are even more contractionary. This naturally leads to easier rejection of the null hypothesis of symmetry, as is evident in the right column. In Figure B.3, with more lags, the problem is essentially gone. The effects of contractionary shocks are essentially the same regardless of the lag structure. This result is specific to the linear asymmetric specification in (2.3); it does not appear in any polynomial specification, regardless of the number of lags included.

B.3 Alternative Mean Squared Forecast Error Comparisons

The coefficients estimated from (2.6) suffer from omitted variable bias relative to the estimates that include monetary policy shocks. Although this bias is irrelevant from a pure forecasting perspective, comparing (2.5) with (2.6) will understate the contribution of monetary policy shocks relative to the arguably more relevant counterfactual in which the effects of monetary policy shocks are accounted for in the regression, but are assumed to be

Figure B.2: IRFs and Tests for Asymmetry—Fewer Lags

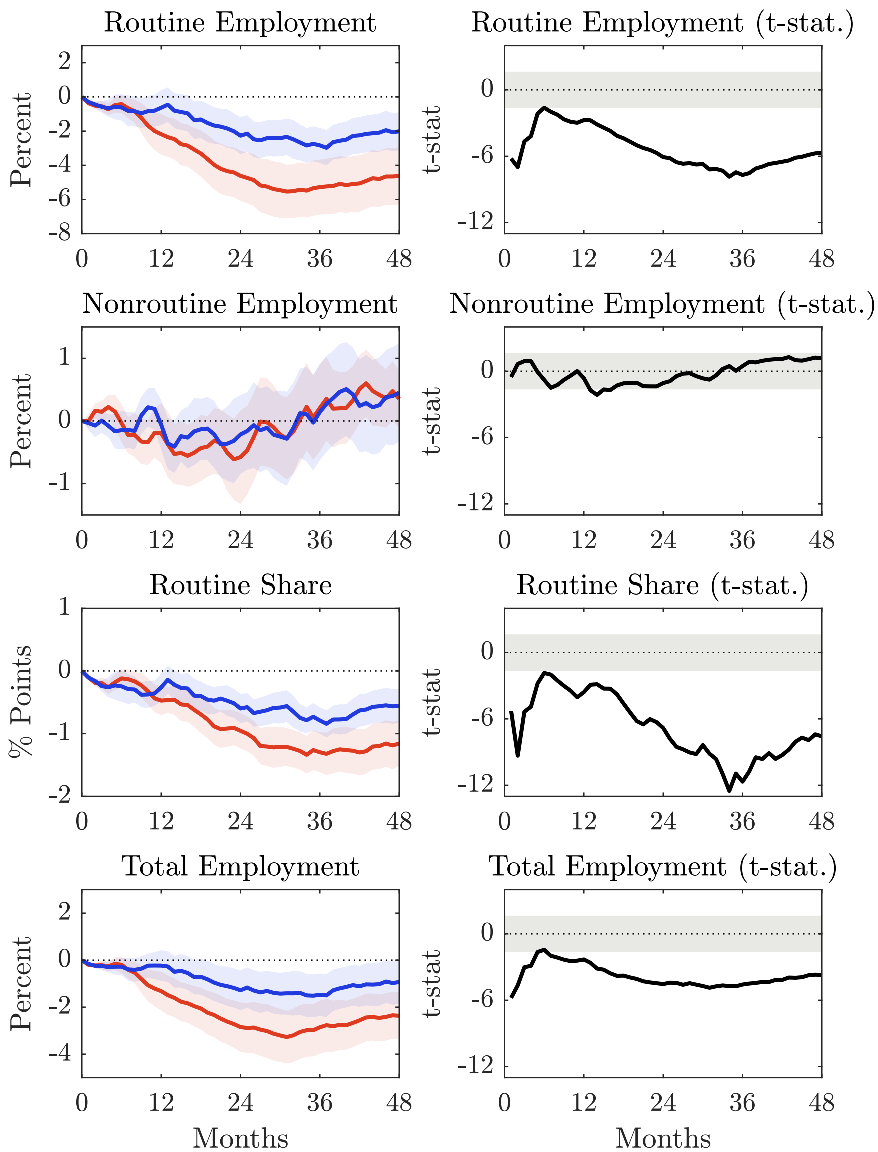

Fewer lags: *Left column*: Impulse responses of occupation employment to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock, no lags of the shock. Estimated from Equation 2.3. *Right column*: Associated t-tests for asymmetry. Shaded areas are 90 percent confidence regions.

Figure B.3: IRFs and Tests for Asymmetry—More Lags

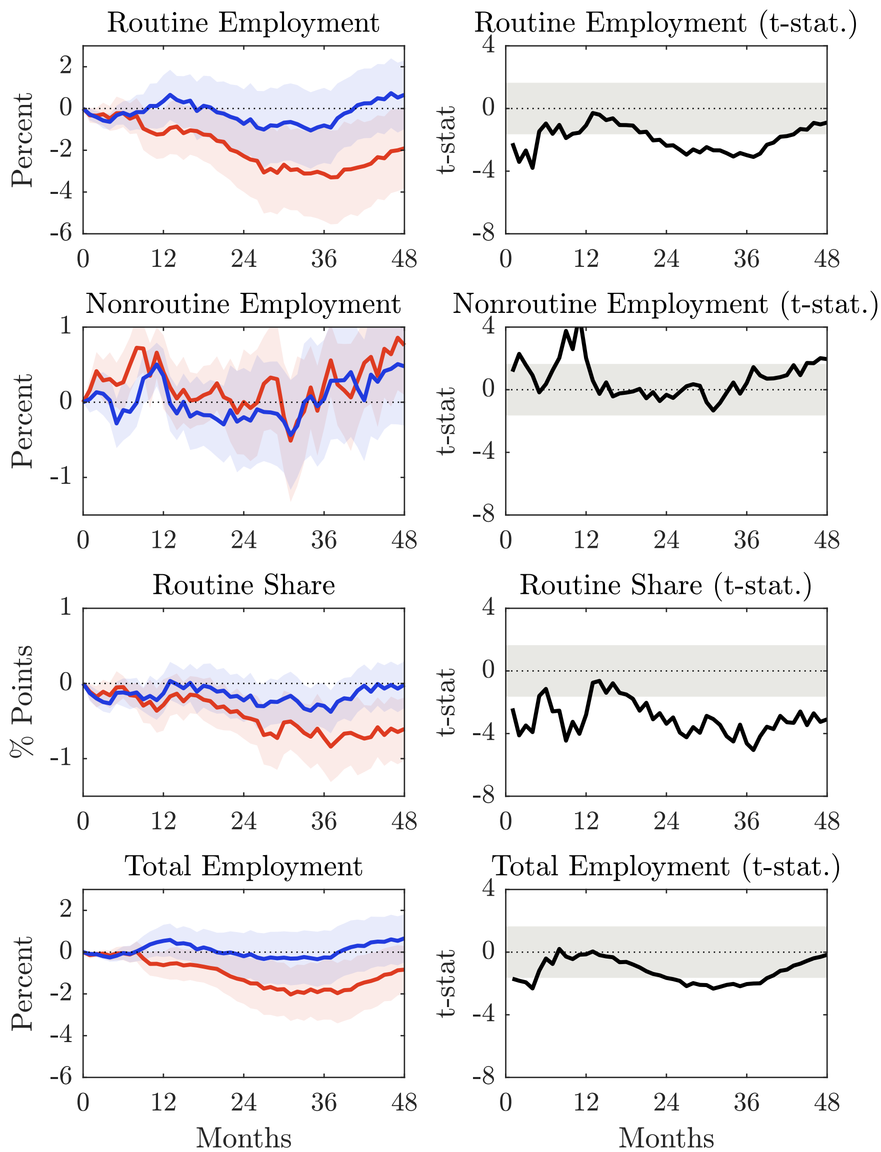

More lags: *Left column*: Impulse responses of occupation employment to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock, with 2 years of lagged dependent variables and 3 years of lagged shocks. Estimated from Equation 2.3. *Right column*: Associated t-tests for asymmetry. Shaded areas are 90 percent confidence regions.

at their mean values.[3] As an alternative, therefore, I consider the forecast error from (2.5) conditional on $\widetilde{\mathbf{x}}_t$ only:

$$\begin{aligned}
\widetilde{FE}_{t,h} &\equiv y_{t+h} - y_t - \mathbb{E}\left[y_{t+h} - y_t \,|\, \widetilde{\mathbf{x}}_t\right] \\
&= y_{t+h} - y_t - \widetilde{\boldsymbol{\beta}}'_h \widetilde{\mathbf{x}}_t - \boldsymbol{\theta}'_h \mathbb{E}\left[\widetilde{\boldsymbol{\epsilon}}_t \,|\, \widetilde{\mathbf{x}}_t\right] \\
&= u_{t,h} + \boldsymbol{\theta}'_h \left[\widetilde{\boldsymbol{\epsilon}}_t - \mathbb{E}\left[\widetilde{\boldsymbol{\epsilon}}_t \,|\, \widetilde{\mathbf{x}}_t\right]\right]. \qquad \text{(B.1)}
\end{aligned}$$

Then $\widetilde{MSFE}_h = \mathbb{E}\left[\widetilde{FE}^2_{t,h}\right]$.

For the linear local projections in (2.2), the term $\mathbb{E}\left[\widetilde{\boldsymbol{\epsilon}}_t \,|\, \widetilde{\mathbf{x}}_t\right]$ theoretically should be zero since the monetary policy shock is unconditionally mean zero and should not be forecastable by macroeconomic variables.[4] For the asymmetric projections in (2.3), however, it will not be zero since ϵ_t^+ and ϵ_t^- have positive and negative means, respectively. In practice, however, the conditional expectation in (B.1) is numerically identical to the unconditional mean—that is, the shocks are not forecastable. I estimate $\mathbb{E}\left[\widetilde{\boldsymbol{\epsilon}}_t \,|\, \widetilde{\mathbf{x}}_t\right]$ from the linear projection

$$\widetilde{\boldsymbol{\epsilon}}_t = \boldsymbol{\Gamma}\widetilde{\mathbf{x}}_t + \boldsymbol{\nu}_t. \qquad \text{(B.2)}$$

Note that $\widetilde{\mathbf{x}}_t$ includes deterministic terms (in the baseline, a constant and linear time trend). The conditional expectation in (B.1) is then just given by $\hat{\boldsymbol{\Gamma}}$, the estimate of the coefficient matrix in (B.2).

It is straightforward to verify that both (2.6) and (B.1) have a mean of zero, so that for each the MSFE is variance of the forecast error. Therefore, for a given horizon, the

[3]That is, zero in the linear or polynomial cases.
[4]I have verified that this is true in practice as well.

relative difference in the FEV between (2.5) and either (2.6) or (B.1) can be interpreted as the share of the FEV explained by the monetary policy shock.

Table B.1 displays the share of the FEV explained by monetary policy shocks for this alternative specification as well as the baseline in the main text. This alternative modestly increases the explanatory power of monetary policy shocks, but the overall pattern across horizons remains unchanged.

Table B.1: Share of FEV due to monetary shocks — Occupations

	Horizon (months)	Linear (2.6)	Linear (B.1)	Baseline (2.6)	Baseline (B.1)	Quadratic (2.6)	Quadratic (B.1)
Routine Employment	12	0.18	0.19	0.23	0.24	0.23	0.23
	24	0.18	0.19	0.32	0.34	0.33	0.33
	36	0.11	0.12	0.23	0.24	0.24	0.24
	48	0.09	0.10	0.16	0.17	0.17	0.17
	Max.	*0.19*	*0.20*	*0.33*	*0.35*	*0.33*	*0.34*
Nonroutine Employment	12	0.03	0.03	0.07	0.07	0.07	0.07
	24	0.01	0.01	0.04	0.04	0.04	0.04
	36	0.01	0.01	0.05	0.07	0.03	0.04
	48	0.01	0.01	0.06	0.08	0.04	0.04
	Max.	*0.05*	*0.05*	*0.10*	*0.10*	*0.11*	*0.11*
Routine Share	12	0.16	0.16	0.22	0.22	0.20	0.20
	24	0.20	0.20	0.38	0.40	0.37	0.37
	36	0.13	0.14	0.32	0.33	0.32	0.32
	48	0.10	0.10	0.23	0.23	0.22	0.22
	Max.	*0.22*	*0.23*	*0.40*	*0.41*	*0.37*	*0.37*
Total Employment	12	0.18	0.19	0.24	0.24	0.24	0.24
	24	0.15	0.16	0.27	0.28	0.29	0.29
	36	0.08	0.08	0.18	0.18	0.19	0.19
	48	0.07	0.07	0.11	0.11	0.12	0.12
	Max.	*0.20*	*0.21*	*0.28*	*0.29*	*0.29*	*0.29*

Note: The table presents the share of the forecast error variance of occupational employment variables due to monetary policy shocks, for both the linear and asymmetric baseline estimates. Within each heading, the number in parentheses indicates the alternative forecasting specification the baseline is compared with, as discussed in Section 2.5.

Bibliography

Abowd, John M., and Arnold Zellner. 1985. "Estimating Gross Labor-Force Flows." *Journal of Business and Economic Statistics*, 3(3): 254–283.

Acemoglu, Daron. 1999. "Changes in unemployment and wage inequality: An alternative theory and some evidence." *American Economic Review*, 89(5): 1259–1278.

Angrist, Joshua D., Òscar Jordà, and Guido M. Kuersteiner. 2016. "Semiparametric Estimates of Monetary Policy Effects: String Theory Revisited." *Journal of Business and Economic Statistics*, 1–17.

Auerbach, Alan, and Yuriy Gorodnichenko. 2012*a*. "Fiscal Multipliers in Recession and Expansion." In *Fiscal Policy after the Financial Crisis.* , ed. Alberto Alesina and Francesco Giavezzi. University of Chicago Press.

Auerbach, Alan, and Yuriy Gorodnichenko. 2012*b*. "Measuring the output responses to fiscal policy." *American Economic Journal: Economic Policy*, 4(2): 1–27.

Autor, David H. 2015. "Why are there still so many jobs? The history and future of workplace automation." *Journal of Economic Perspectives*, 29(3): 3–30.

Autor, David H., and Lawrence F. Katz. 1999. "Changes in the wage structure and earnings inequality." In *Handbook of Labor Economics, Vol. 3.* , ed. Orley Ashenfelter and David Card, 1463–1555. Elsevier Science B.V.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The skill content of recent technological change: an empirical exploration." *Quarterly Journal of Economics*, 118(4): 1279–1333.

Bachmann, Rüdiger, and Eric R. Sims. 2012. "Confidence and the Transmission of Government Spending Shocks." *Journal of Monetary Economics*, 59(3): 235–249.

Bachmann, Rüdiger, Steffen Elstner, and Eric R. Sims. 2013. "Uncertainty and Economic Activity: Evidence from Business Survey Data." *American Economic Journal: Macroeconomics*, 5(2): 217–249.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. "Measuring Economic Policy Uncertainty." *The Quarterly Journal of Economics*, 131(4): 1593–1636.

Barnichon, Regis A., and Christian Matthes. 2017. "Functional Approximations of Impulse Responses: New Insights into the Asymmetric Effects of Monetary Policy." Working paper.

Barnichon, Regis, and Christopher J. Nekara. 2012. "The Ins and Outs of Forecasting Unemployment: Using Labor Force Flows to Forecast the Labor Market." *Brookings Papers on Economic Activity*, 83–131.

Basu, Susanto, and Brent Bundick. 2017. "Uncertainty Shocks in a Model of Effective Demand." *Econometrica*, 85(3): 937–958.

Beaudry, Paul, Alban Moura, and Franck Portier. 2015. "Reexamining the cyclical behavior of the relative price of investment." *Economic Letters*, 135: 108–111.

Berger, David, Ian Dew-Becker, Konstantin Milbrandt, Lawrence D.W. Schmidt, and Yuta Takahashi. 2016. “Layoff risk, the welfare cost of business cycles, and monetary policy.” Working paper.

Bernanke, Ben S. 1983. “Irreversibility, Uncertainty, and Cyclical Investment.” *Quarterly Journal of Economics*, 98(1): 85–106.

Bernanke, Ben S, Mark Gertler, and Mark Watson. 1997. “Systematic Monetary Policy and the Effects of Oil Price Shocks.” *Brookings Papers on Economic Activity*, 1997(1): 91–157.

Blanchard, Olivier J., and Jordi Galí. 2010. “Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment.” *American Economic Journal: Macroeconomics*, 2(2): 1–30.

Blanchard, Olivier J., and Peter Diamond. 1989. “The Beveridge Curve.” *Brookings Papers on Economic Activity*, 1989(1): 1–60.

Blanchard, Olivier J., and Peter Diamond. 1990. “The Cyclical Behavior of the Gross Flows of U.S. Workers.” *Brookings Papers on Economic Activity*, 1990(2): 85–143.

Bloom, Nicholas. 2009. “The Impact of Uncertainty Shocks.” *Econometrica*, 77(3): 623–685.

Bloom, Nicholas. 2014. “Fluctuations in Uncertainty.” *Journal of Economic Perspectives*, 28(2): 153–176.

Braun, Helge, Reinout De Bock, and Riccardo DiCecio. 2009. "Supply Shocks, Demand Shocks, and Labor Market Fluctuations." *Federal Reserve Bank of St. Louis Review*, 91(3): 155–178.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics*, 12(3): 383–398.

Canova, Fabio, David Lopez-Salido, and Claudio Michelacci. 2007. "Labor Market Effects of Technology Shocks." Banco de España Research Paper No. WP-0719.

Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans. 1996. "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds." *Review of Economics and Statistics*, 78(1): 16–34.

Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans. 1999. "Monetary Policy Shocks: What Have We Learned and to What End?" *Handbook of Macroeconomics*, 1: 65–148.

Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.

Coibion, Olivier. 2012. "Are the Effects of Monetary Policy Shocks Big or Small?" *American Economic Journal: Marcoeconomics*, 4(2): 1–32.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2012. "Innocent Bystanders? Monetary Policy and Inequality in the U.S." NBER Working Paper No. 18170.

DiCecio, Riccardo. 2009. "Stick wages and sectoral labor comovement." *Journal of Economic Dynamics & Control*, 33: 538–553.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2015. "On the Importance of the Participation Margin for Labor Market Fluctuations." *Journal of Monetary Economics*, 72: 64–82.

Faia, Ester. 2008. "Optimal Monetary Policy Rules with Labor Market Frictions." *Journal of Economic Dynamics and Control*, 32(5): 1600–1621.

Fisher, Jonas D. M. 2006. "The Dynamic Effects of Neutral and Investment Specific Technology Shocks." *Journal of Political Economy*, 114(3): 413–451.

Foote, Christopher L., and Richard W. Ryan. 2015. "Labor-market polarization over the business cycle." In *Macroeconomics Annual*. Vol. 29(1).

Fujita, Shigeru, and Gary Ramey. 2006. "The Cyclicality of Job Loss and Hiring." Federal Reserve Bank of Philadelphia Working Paper 06-17.

Fujita, Shigeru, and Gary Ramey. 2007. "Reassessing the Shimer Facts." Federal Reserve Bank of Philadelphia Working Paper 07-02.

Fujita, Shigeru, and Gary Ramey. 2009. "The Cyclicality of Separation and Job Finding Rates." *International Economic Review*, 50(2): 415–430.

Galí, Jordi. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework.* Princeton University Press.

Galí, Jordi. 2011. "Monetary Policy and Unemployment." In *Handbook of Monetary Economics, Vol. 3A.* , ed. Benjamin Friedman and Michael Woodford, 487–546. Elsevier B.V.

Gertler, Mark, and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics*, 7(1): 44–76.

Gilchrist, Simon, and Egon Zakrajsek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review*, 102(4): 1692–1720.

Gilchrist, Simon, and John C. Williams. 2005. "Investment, Capacity, and Uncertainty: A Putty-Clay Approach." *Review of Economic Dynamics*, 8(1): 1–27.

Goos, Maarten, and Alan Manning. 2007. "Lousy and lovely jobs: the rising polarization of work in Britain." *The Review of Economics and Statistics*, 89(1): 118–133.

Gordon, Robert J. 1990. *The Measurement of Durable Goods Prices.* University of Chicago Press.

Hagedorn, Marcus, and Iourii Manovskii. 2008. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited." *American Economic Review*, 98(4): 1692–1706.

Hall, Robert E. 2005*a*. "Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market." *Review of Economics and Statistics*, 87(3): 397–407.

Hall, Robert E. 2005*b*. "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years." *NBER Macroeconomics Annual*, 101–137.

Hosios, Arthur J. 1990. "On the Efficiency of Matching and Related Models of Search and Unemployment." *The Review of Economic Studies*, 57(2): 279–298.

Jaimovich, Nir, and Henry E. Siu. 2014. "The trend is the cycle: Job polarization and jobless recoveries." Working paper.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projection." *American Economic Review*, 95(1): 161–182.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng. 2015. "Measuring Uncertainty." *American Economic Review*, 105(5): 1177–1216.

Krause, Michael U., and Thomas A. Lubik. 2010. "Instability and Indeterminacy in a Simple Search and Matching Model." *Federal Reserve Bank of Richmond Economic Quarterly*, 96(3): 259–272.

Krusell, Per J., Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. 2011. "A Three State Model of Worker Flows in General Equilibrium." *Journal of Economic Theory*, 146(3): 812–844.

Krusell, Per J., Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. 2016. "Gross Worker Flows Over the Business Cycle." Working Paper, Sept. 2016.

Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante. 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica*, 68(5): 1029–1053.

Leduc, Sylvain, and Zheng Liu. 2016. "Uncertainty Shocks are Aggregate Demand Shocks." *Journal of Monetary Economics*, 82: 20–35.

Lucas, Robert E. 1980. "Methods and Problems in Business Cycle Theory." *Journal of Money, Credit, and Banking*, 12(4): 696–715.

Ludvigson, Sydney C., Sai Ma, and Serena Ng. 2018. "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" Working Paper.

Mertens, Karel, and Morten O. Ravn. 2013. "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States." *American Economic Review*, 103(4): 1212–1247.

Merz, Monika. 1995. "Search in the labor market and the real business cycle." *Journal of Monetary Economics*, 36(2): 269–300.

Michelacci, Claudio, and David Lopez-Salido. 2007. "Technology Shocks and Job Flows." *The Review of Economic Studies*, 74(4): 1195–1227.

Mueller, Andreas I. 2017. "Separations, Sorting and Cyclical Unemployment." *American Economic Review, forthcoming.*

Nakamura, Emi, and Jón Steinsson. 2017. "Identification in Macroeconomics." *Journal of Economic Perspectives*, forthcoming.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–708.

Owyang, Michael T., Valerie A. Ramey, and Sarah Zubairy. 2013. "Are Government Multipliers Greater during Periods of Slack: Evidence from Twentieth-Century Historical Data." *American Economic Review: Papers & Proceedings*, 103(3): 129–134.

Pagan, Adrian. 1984. "Econometric Issues in the Analysis of Regressions with Generated Regressors." *International Economic Review*, 25(1): 221–247.

Poterba, James M., and Lawrence H. Summers. 1986. "Reporting Errors and Labor Market Dynamics." *Econometrica*, 54(6): 1319–1338.

Ramey, Valerie A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics, Vol. 2.* , ed. John B. Taylor and Harald Uhlig, 71–162. Elsevier.

Ramey, Valerie A., and Sarah Zubairy. 2017. "Government Spending Multipliers in Goods Times and in Bad: Evidence from U.S. Historical Data." *Journal of Political Economy, forthcoming.*

Ravenna, Federico, and Carl E. Walsh. 2011. "Welfare-Based Optimal Monetary Policy with Unemployment and Sticky Prices: A Linear-Quadratic Framework." *American Economic Journal: Macroeconomics*, 3(2): 130–162.

Ravenna, Federico, and Carl E. Walsh. 2012. "Screening and Labor Market Flows in a Model with Heterogeneous Workers." *Journal of Money, Credit, and Banking*, 44(2): 32–71.

Rich, Robert, and Joseph Tracy. 2010. "The Relationships Among Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts." *The Review of Economics and Statistics*, 92(1): 200–207.

Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94(4): 1055–1084.

Schmitt-Grohé, Stephanie, and Martin Uribe. 2006. "Optimal Fiscal and Monetary Policy in a Medium-Scale Macroeconomic Model." In *NBER Macroeconomics Annual 2005, Volume 20.* , ed. Mark Gertler and Kenneth Rogoff, 383–425.

Schmitt-Grohé, Stephanie, and Martin Uribe. 2007. "Optimal simple and implementable monetary and fiscal rules." *Journal of Monetary Economics*, 54(6): 1702–1725.

Shimer, Robert. 2005*a*. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review*, 95(1): 25–49.

Shimer, Robert. 2005*b*. "The Cyclicality of Hires, Separations, and Job-to-Job Transitions." *Federal Reserve Bank of St. Louis Review*, 87(4): 493–507.

Shimer, Robert. 2012. "Reassessing the Ins and Outs of Unemployment." *Review of Economic Dynamics*, 15(2): 127–148.

Sims, Christopher A., and Tao Zha. 2006. "Does Monetary Policy Generate Recessions?" *Macroeconomic Dynamics*, 10(2): 231–272.

Stock, James H., and Mark W. Watson. 2012. "Disentangling the Channels of the 2007–09 Recession." *Brookings Papers on Economic Activity*, 42(1): 81–135.

Stock, James H., and Mark W. Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics using External Instruments." NBER Working paper 24216.

Tenreyro, Silvana, and Gregory Thwaites. 2016. "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions." *American Economic Journal: Macroeconomics*, 8(4): 43–74.

Trigari, Antonella. 2009. "Equilibrium Unemployment, Job Flows, and Inflation Dynamics." *Journal of Money, Credit and Banking*, 41(1): 1–33.

Wong, Arlene. 2016. "Transmission of Monetary Policy to Consumption and Population Aging." Working paper.

Woodford, Michael. 2003. *Interest and Prices.* Princeton University Press.