

Has Monetary Policy Accelerated Job Polarization?

Neil White*

July 19, 2022

For the most recent version of this paper, visit
<https://sites.google.com/view/neilwhite/research>

Abstract

Since the late 1960s, the share of U.S. employment in occupations involving primarily routine tasks has declined by about 30 percent, a trend that has largely affected workers with a high-school degree but no college. This paper argues that contractionary monetary policy has accelerated this change. In part by disproportionately affecting industries with high shares of routine occupations, contractionary monetary policy shocks lead to large and very 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.

Keywords: Job polarization; monetary policy; asymmetries; occupational employment; industry employment.

*Amherst College, Department of Economics. 100 Boltwood Ave., Converse Hall, Room 303, Amherst, MA 01002 U.S.A. Email: nwhite@amherst.edu. I would like to thank Olivier Coibion, Saroj Bhattarai, Andrew Glover, Nicolas Petrosky-Nadeau, Hassan Afrouzi, Garrett Hagemann, Cooper Howes, and Julian Ludwig for helpful comments, as well as seminar participants at the Federal Reserve Board of Governors and UT Austin.

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 1](#). Autor et al. (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 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](#), 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 December 2019, that number had dropped more than 20 percentage points to less than 42 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 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), I explore three 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

¹This phenomenon is described in detail in Acemoglu (1999), Autor, Levy and Murnane (2003), Jaimovich and Siu (2014), Foote and Ryan (2015), and others.

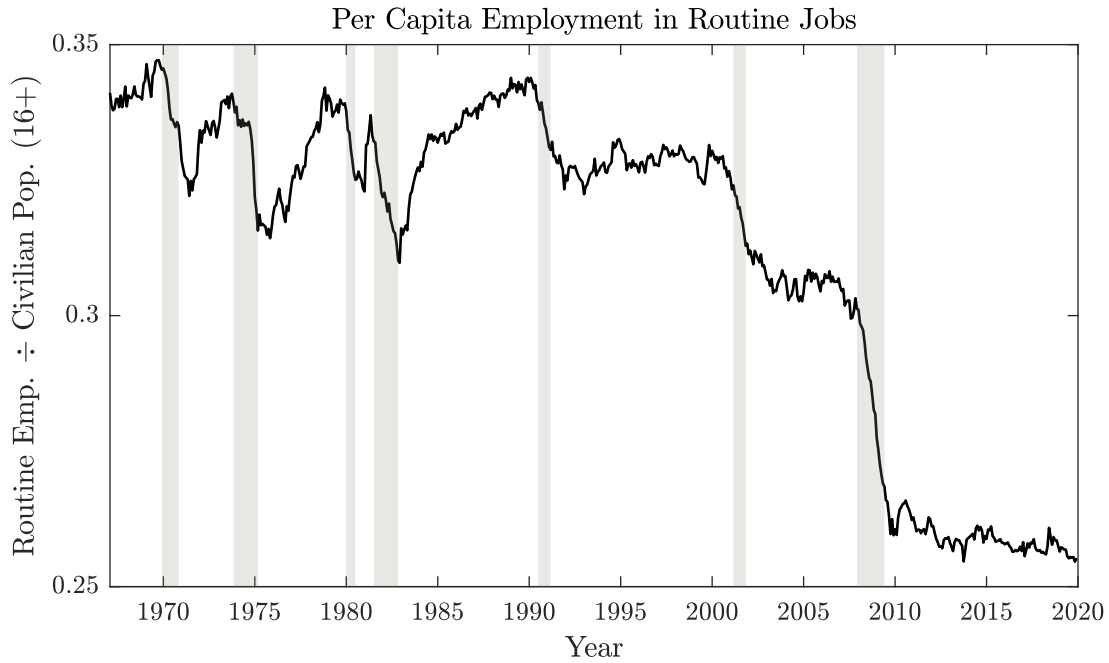


Figure 1: Per capita employment in routine occupations, 1967-2019. Shaded dates are NBER recessions.




Figure 2: Share of total employment in routine occupations, 1967-2019. Shaded dates are NBER recessions.

from labor, an important mechanism that the job polarization literature has highlighted as an explanation for the trends evident in Figures 1 and 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 outside 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 Monetary 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.

Finally, I present some evidence of the drivers of these differing effects by industry. Contractionary shocks lead to large and persistent appreciation of the dollar and declines in exports volume (especially of durable goods), as well as increases in mortgage rates and declines in housing starts and residential investment. The former is related to the role of trade emphasized in the job polarization literature (cite papers), and the latter points to the important role of the housing market in the transmission of monetary policy to the real economy (cite papers).

Nevertheless, this industry effect does not explain the asymmetry *per se*; monetary policy has asymmetric effects on employment in virtually all industries. The source of this asymmetry is likely the persistent endogenous effect of monetary policy.

²It 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 role in job polarization.

³In 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.

In Section 2, I review the recent literature on job polarization and discuss Autor et al.’s (2003) occupation classification system used throughout the paper. In Section 3, I describe the data on employment and monetary policy I use in subsequent sections. In Section 4, I discuss the estimation of linear and asymmetric impulse responses for occupational employment data and present results. In Section 5, I discuss the historical contribution of monetary policy shocks to job polarization. In Section 6, I discuss possible reasons for the results from the previous section. Section 7 concludes.

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 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 et al. (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 et al. (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—especially 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 non-routine 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 Goos and Manning’s (2007) introduction of the term “job polarization.” The trend of job polarization has since been documented in

⁴See Autor and Katz (1999) for an overview.

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 cyclical nature 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 cyclical nature 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. Tenreyro and 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. Angrist, Jordà and Kuersteiner (2016) and Barnichon and 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 here⁷; in the next section, however, I describe in detail the two empirical approaches I take.

⁵Autor (2015) provides an excellent overview of this research.

⁶See for example Auerbach and Gorodnichenko (2012a) and (2012b), Owyang, Ramey and Zubairy (2013), and Ramey and Zubairy (2017).

⁷See 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.

3 Data

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 et al.’s (2003) task-based framework.

Data on investment prices are from 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 Gordon (1990), following Krusell, Ohanian, Ríos-Rull and Violante (2000).

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 Coibion, Gorodnichenko, Kueng and 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:

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

where m indexes FOMC meeting dates, $f f b_m$ 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, $\widetilde{\pi}_{mi}$ denotes forecasts of inflation, \widetilde{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 Gertler and 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 A.1.

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 partic-

⁸This 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.

⁹The 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 Gilchrist and Zakrajsek’s (2012) measure of the excess bond premium.

ularly 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.

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 $h = 1, \dots, H$ I estimate

$$y_{t+h} - y_t = \beta'_h \mathbf{x}_t + \underbrace{\gamma_h \epsilon_t}_{\text{Linear}} + u_{t,h}. \quad (2)$$

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 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 3 displays the IRFs from (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.

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

¹¹To account for serial correlation, the standard errors are adjusted as in Newey and 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)).

¹²Section A.2 discusses alternative lag structures.

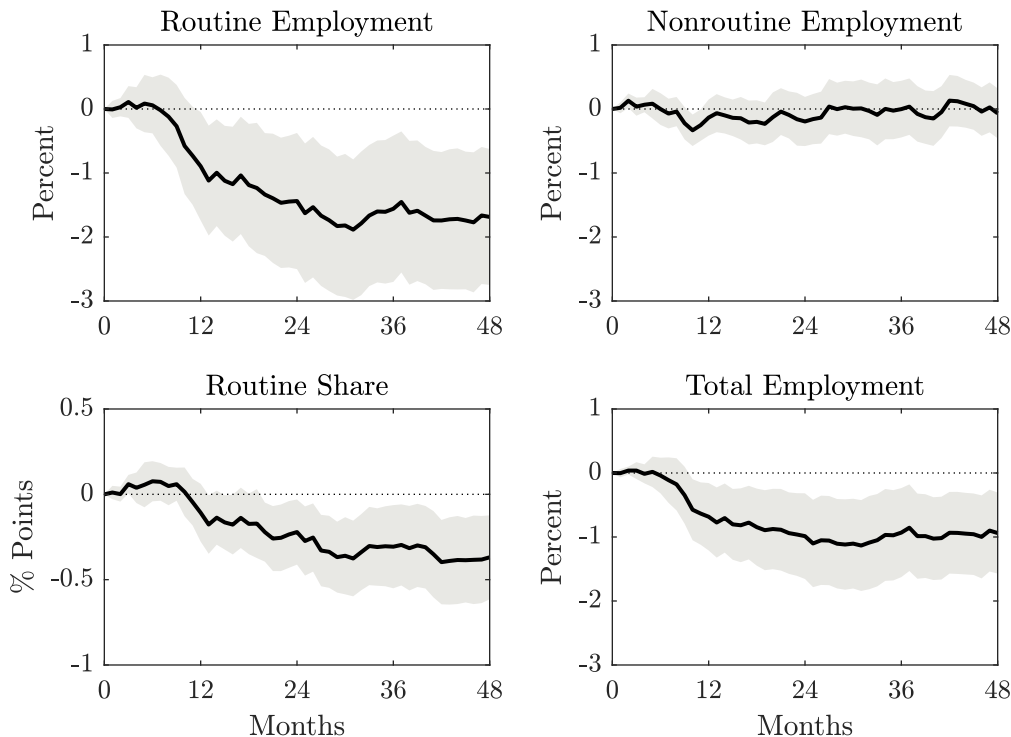


Figure 3: Impulse responses to a 100 b.p. contractionary Romer and Romer (2004) monetary policy shock. Estimated from Equation 2. Shaded areas are 90 percent confidence intervals.

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) 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.

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, 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 Wong (2016). Specifically, I estimate for each horizon $h = 1, \dots, H$,

$$y_{t+h} - y_t = \beta'_h \mathbf{x}_t + \underbrace{\gamma_h^+ \epsilon_t^+}_{\text{Contractionary}} + \underbrace{\gamma_h^- \epsilon_t^-}_{\text{Expansionary}} + u_{t,h}, \quad (3)$$

where $\epsilon_t^+ := \max\{\epsilon_t, 0\}$ and $\epsilon_t^- := \min\{\epsilon_t, 0\}$. To facilitate comparisons with Section 4.1, I include the same controls as in (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 $h = 1, \dots, H$. 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 A.2.

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

$$y_{t+h} - y_t = \beta'_h \mathbf{x}_t + \gamma_{1,h} \epsilon_t + \gamma_{2,h} \epsilon_t^2 + u_{t,h}. \quad (4)$$

This specification dispenses with the assumption in (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 (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 $h = 1, \dots, H$ is somewhat easier to reject than the null hypotheses of symmetry in (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 3 and 4 are displayed in the left and right columns of Figure 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

¹³Indeed, this is reflected in the point-wise asymmetry tests displayed in Figure 5.

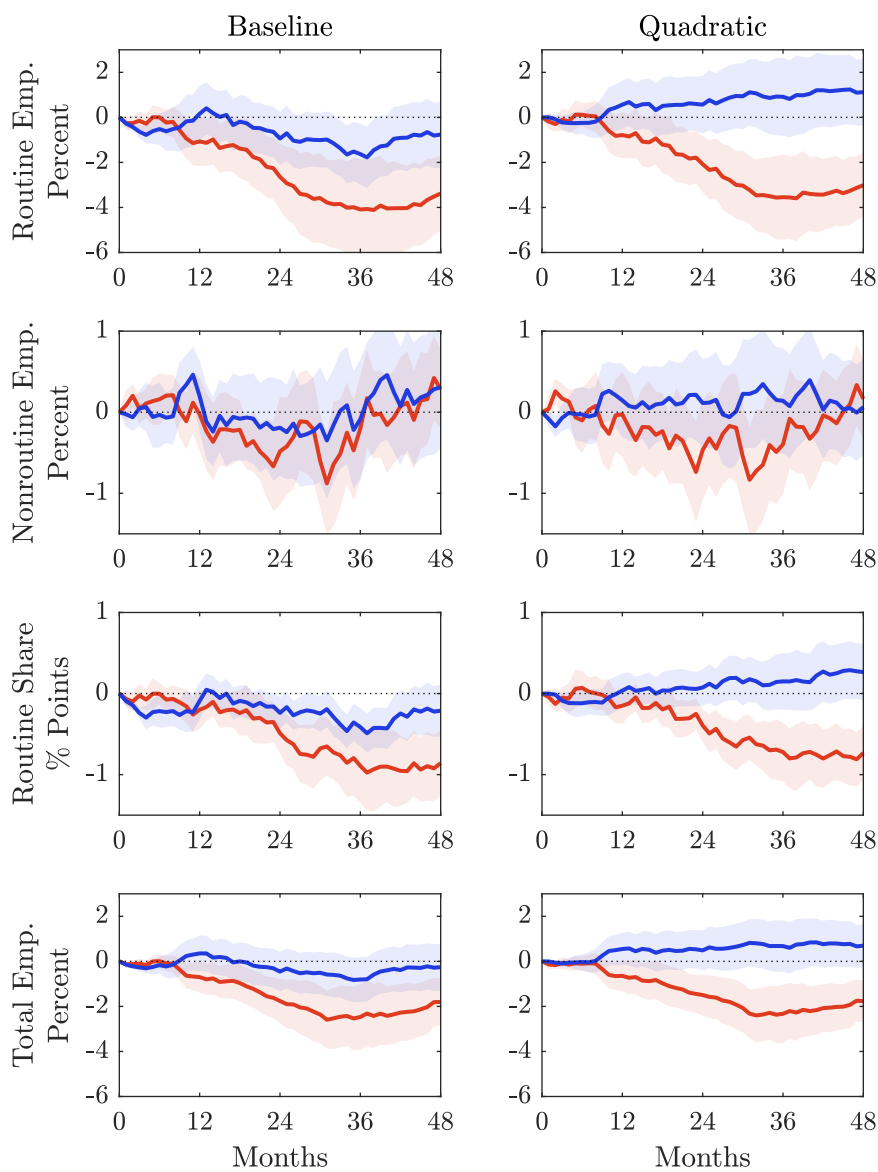


Figure 4: Impulse responses to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock. Shaded areas are 90 percent confidence intervals. *Left column:* baseline asymmetric estimates from Equation 3. *Right column:* alternative quadratic estimates from Equation 4.

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 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 asymmetry.¹⁵ The t-statistics for the quadratic estimates in the right column of [Figure 5](#) are somewhat larger than those 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

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

¹⁴It is worth noting, however, that the results from (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 A.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.

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

¹⁶An 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.

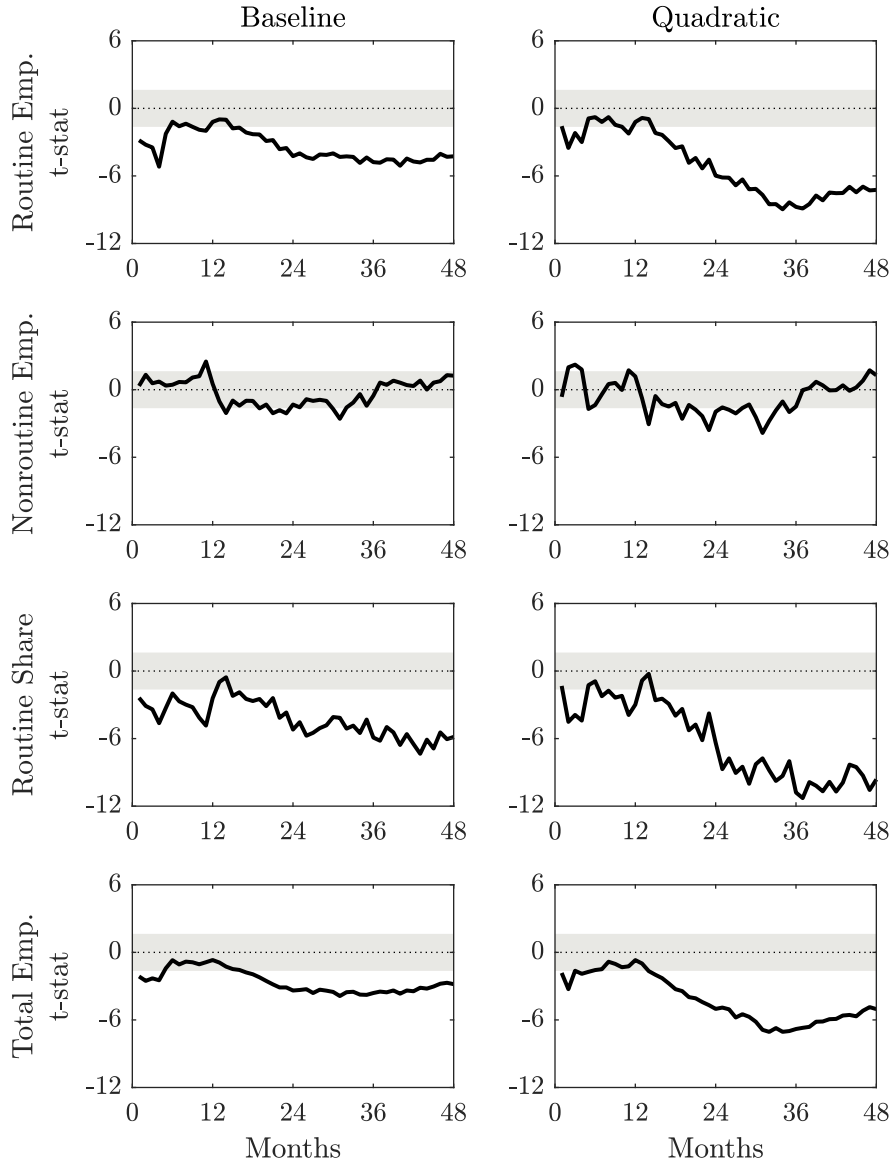


Figure 5: Point-wise t-test against the null of symmetry for sign-dependent IRFs in Figure 4. Shaded areas indicate a 90% confidence region. Lines outside the shaded region indicate the null can be rejected at 10% significance. *Left column*: baseline asymmetric estimates from Equation 3. *Right column*: alternative quadratic estimates from Equation 4.

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.

One approach is to use the one-period-ahead local projection (i.e., (2) or (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), (3), and (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:

$$y_{t+h} - y_t = \tilde{\beta}'_h \tilde{\mathbf{x}}_t + \boldsymbol{\theta}'_h \tilde{\boldsymbol{\epsilon}}_t + u_{t,h}, \quad (5)$$

where $\tilde{\mathbf{x}}_t$ includes only lags of the dependent variable, and $\tilde{\boldsymbol{\epsilon}}_t$ 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 (5) at time t , horizon h is given by

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

The mean squared forecast error (MSFE) at horizon h is then given by $MSFE_h = \mathbb{E}[FE_{t,h}^2] = \mathbb{E}[u_{t,h}^2]$, 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

$$y_{t+h} - y_t = \tilde{\beta}'_h \tilde{x}_t + e_{t,h}. \quad (6)$$

As above, the forecast errors are the estimated regression residuals and the MSFE at horizon h is given by $\mathbb{E}[e_{t,h}^2]$. 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 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 4 are so persistent, this 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.

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

¹⁷I also considered a slightly different comparison that takes into account the existence of monetary policy shocks in estimating the coefficient vector $\tilde{\beta}_h$, 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 A.3. The implied reductions in the FEVs are virtually identical to those from (6).

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 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. I present some evidence that the response of mortgage rates might be decreasing construction of new houses and the exchange rates appreciation could lead to declining demand for durable goods exports.

6.1 Capital-labor substitution and the relative price of investment

As discussed at length in Autor et al. (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 through 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.

¹⁸This 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.

Fisher (2006), for example, finds the relative price of investment goods, particularly investment equipment, to be strongly negatively correlated with output; he argues that shocks to investment-specific technology are a large driver of the business cycle. On the other hand, Beaudry, Moura and 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 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 6 displays impulse responses to monetary policy shocks of these variables estimated from quarterly versions of Equations 2, 3, and 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, the job polarization literature highlights the important role of falling investment prices in driving the long-run trends in declining routine employment.

¹⁹Although 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.

²⁰As 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.

²¹Note that the anomalous responses to expansionary shocks discussed in Section 4.2 are present here as well. As before, these patterns do not appear in the quadratic estimates.

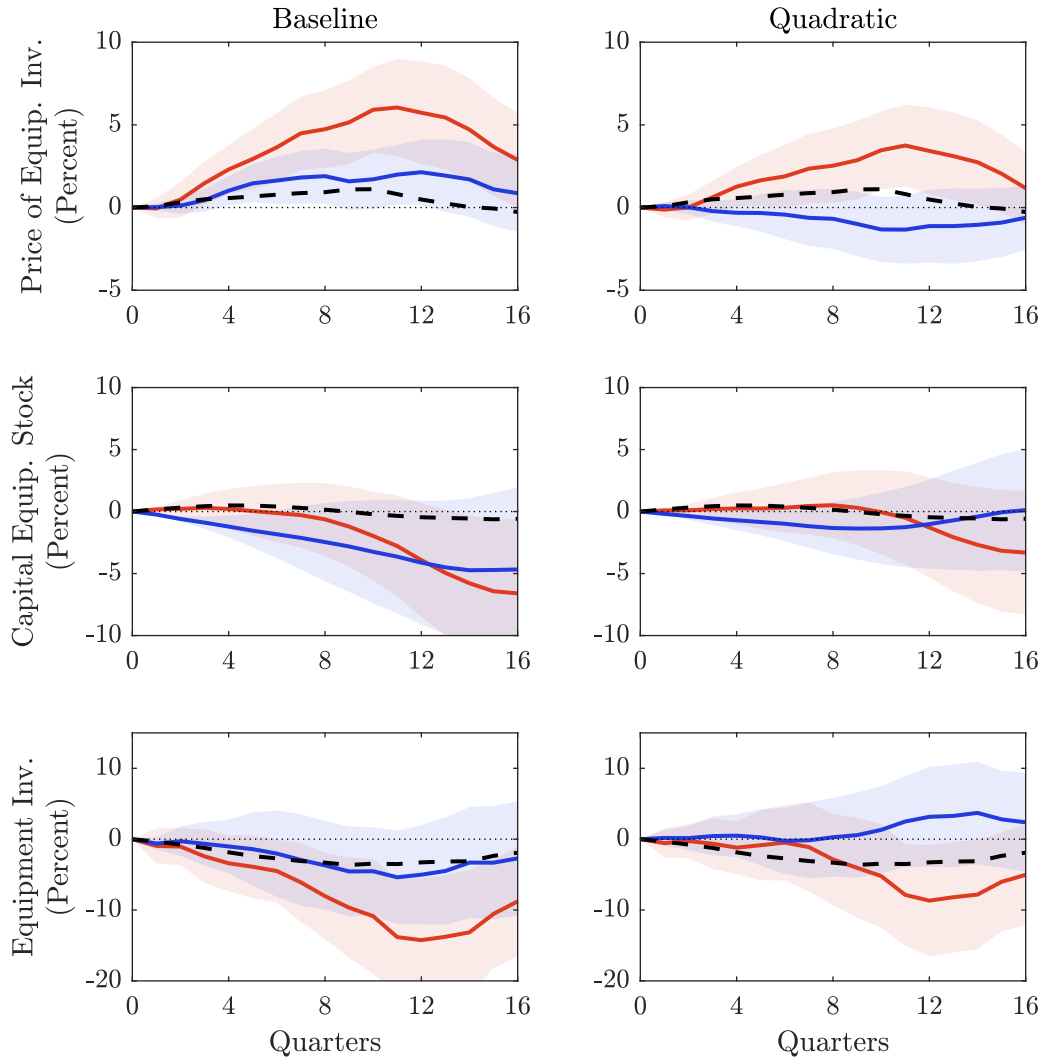


Figure 6: Impulse responses of equipment capital, investment, and prices to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock. Dashed lines are the IRFs of a contractionary shock estimated linearly. Shaded areas are 90 percent confidence intervals. *Left column*: baseline asymmetric estimates from Equation 3. *Right column*: alternative quadratic estimates from Equation 4.

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

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 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 7 displays the IRFs from (2) to a 100 b.p. contractionary Romer and Romer (2004) shock of employment in a variety of industries. While employment in some in-

²²Because 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.

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

²⁴The 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).

²⁵Ideally, 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*.

dustries 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 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 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 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 4](#). It can explain the *large* and *persistent* effects of monetary policy shocks on routine employment; it cannot by itself explain the *asymmetric* effects.

[Figure 10](#) presents some evidence of the channels through which this industry effect might occur. There are large and asymmetric effects of monetary policy shocks on exchange rates, exports, mortgage rates, and exports. The large dollar appreciation that occurs after contractionary monetary policy shocks could be leading to lower foreign demand for U.S.-produced durable goods; indeed exports decline substantially, although the standard errors are large. Increases in mortgage rates could decrease demand for housing, and indeed residential fixed investment declines almost twice as much after contractionary shocks as it rises after expansionary ones.

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 and

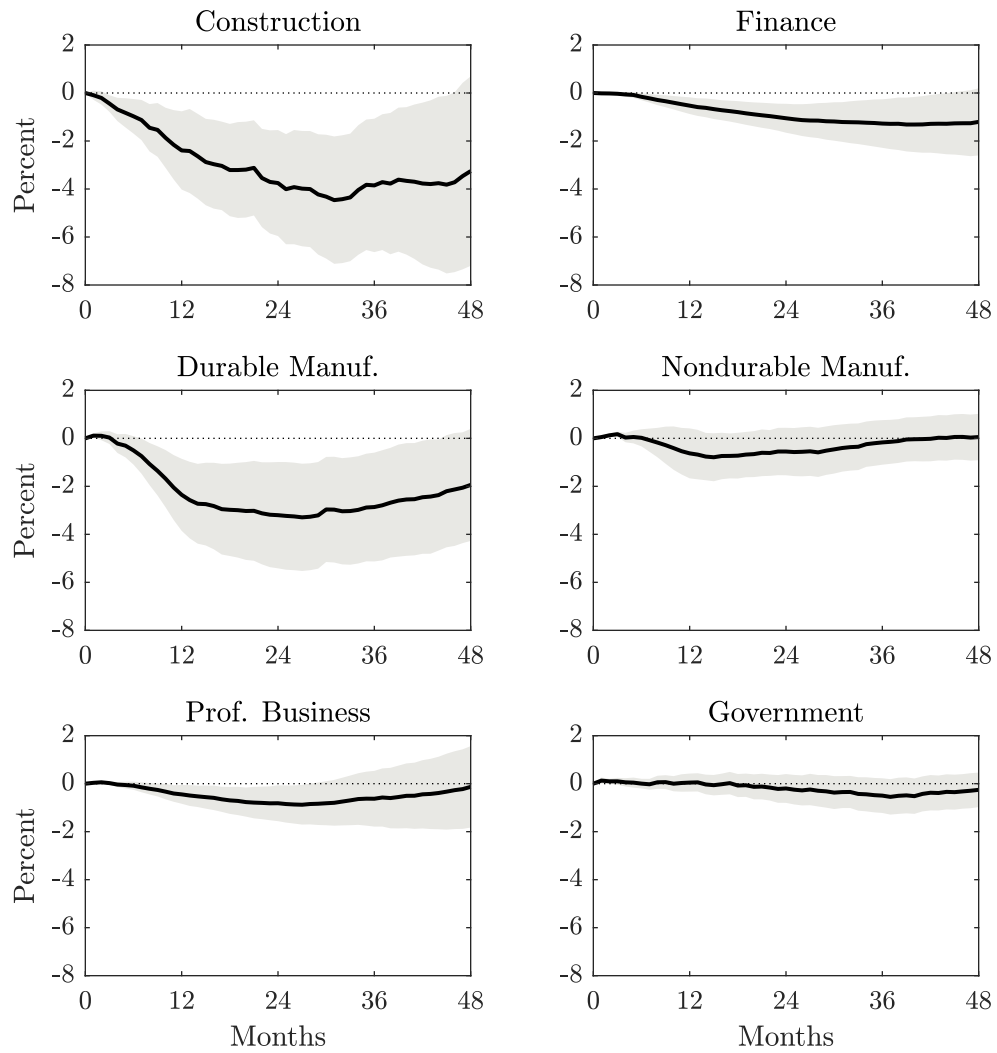


Figure 7: Impulse responses to a 100 b.p. contractionary Romer and Romer (2004) monetary policy shock. Estimated from Equation 2. Shaded areas are 90 percent confidence intervals.

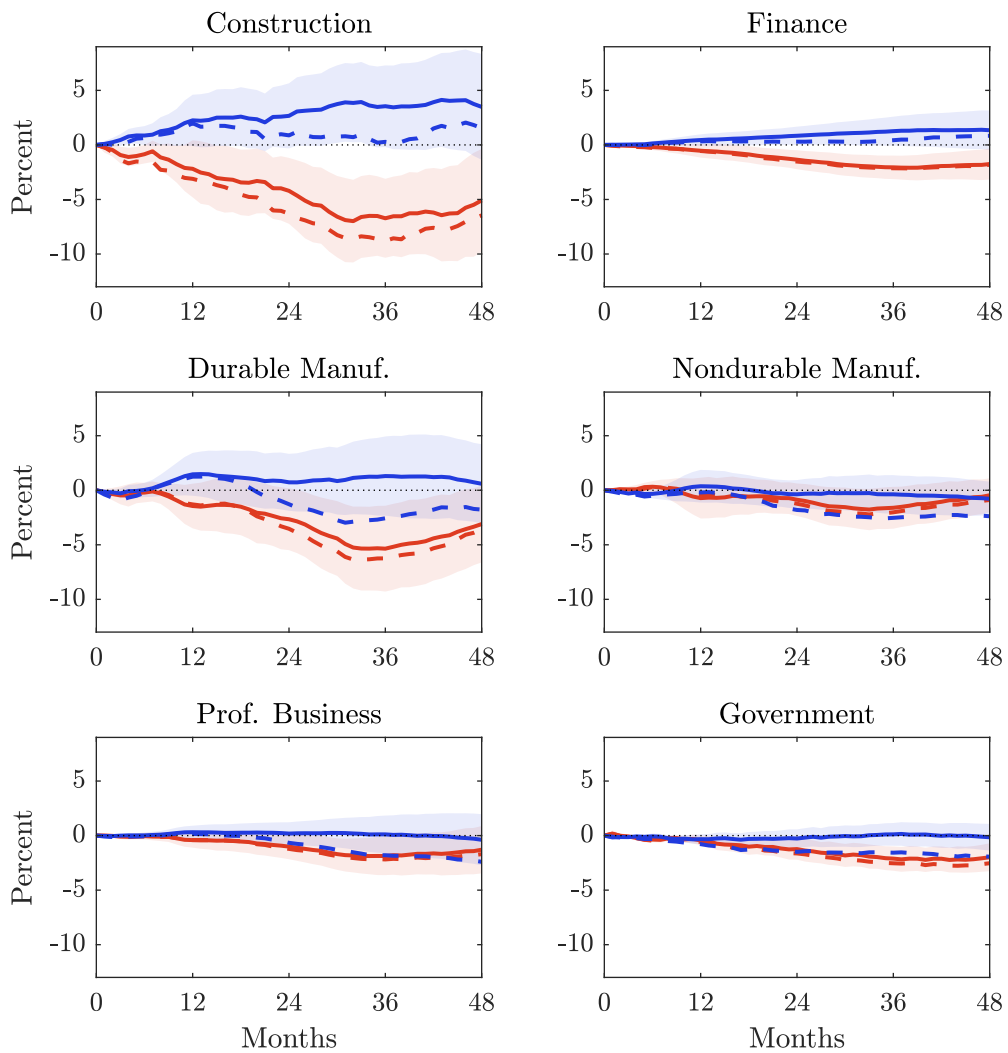


Figure 8: Impulse responses to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock. Estimated from Equation 4. Dashed lines are point estimates from 3. Shaded areas are 90 percent confidence intervals.

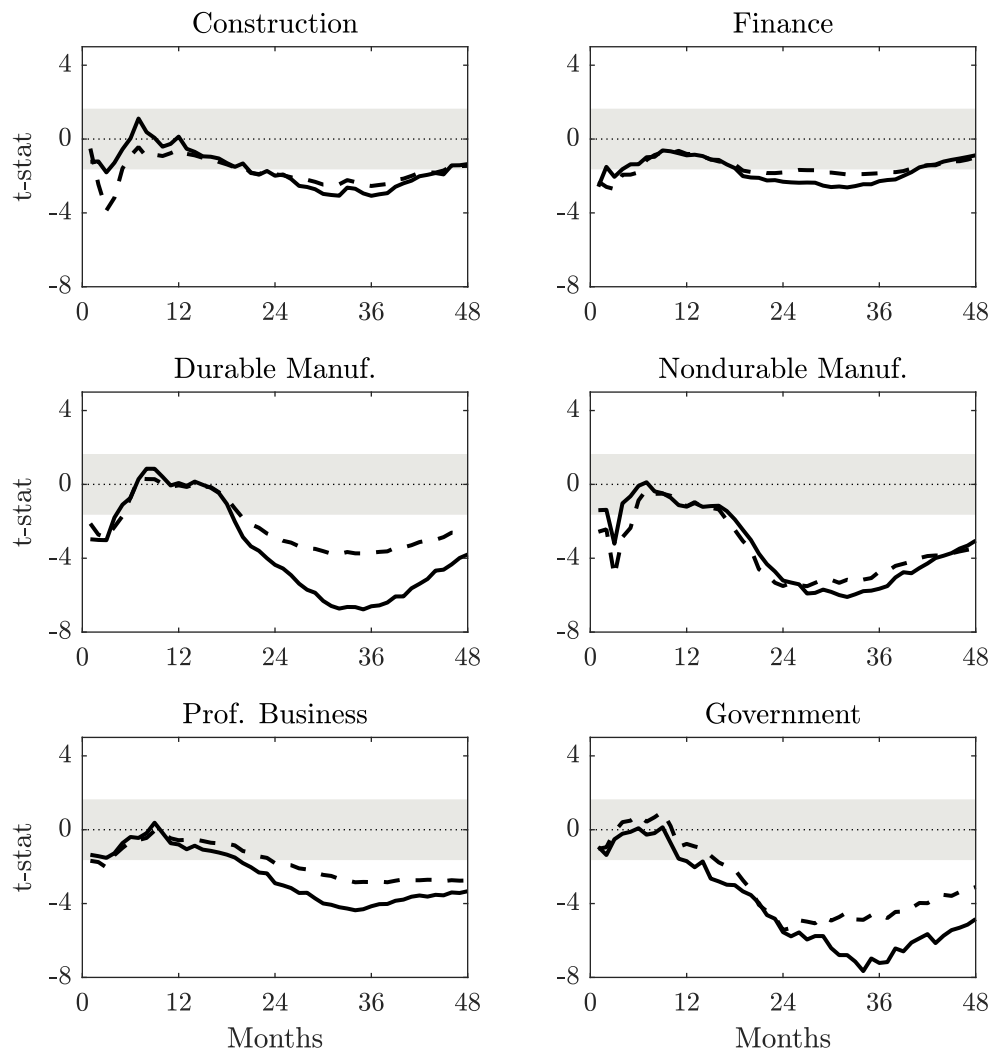


Figure 9: Point-wise t-tests against the null of symmetry for sign-dependent IRFs in Figure 8. Shaded areas indicate 90% confidence region. Solids lines from are from (4), dashed lines from (3). Lines outside the shaded region indicate the null can be rejected at 10% significance.

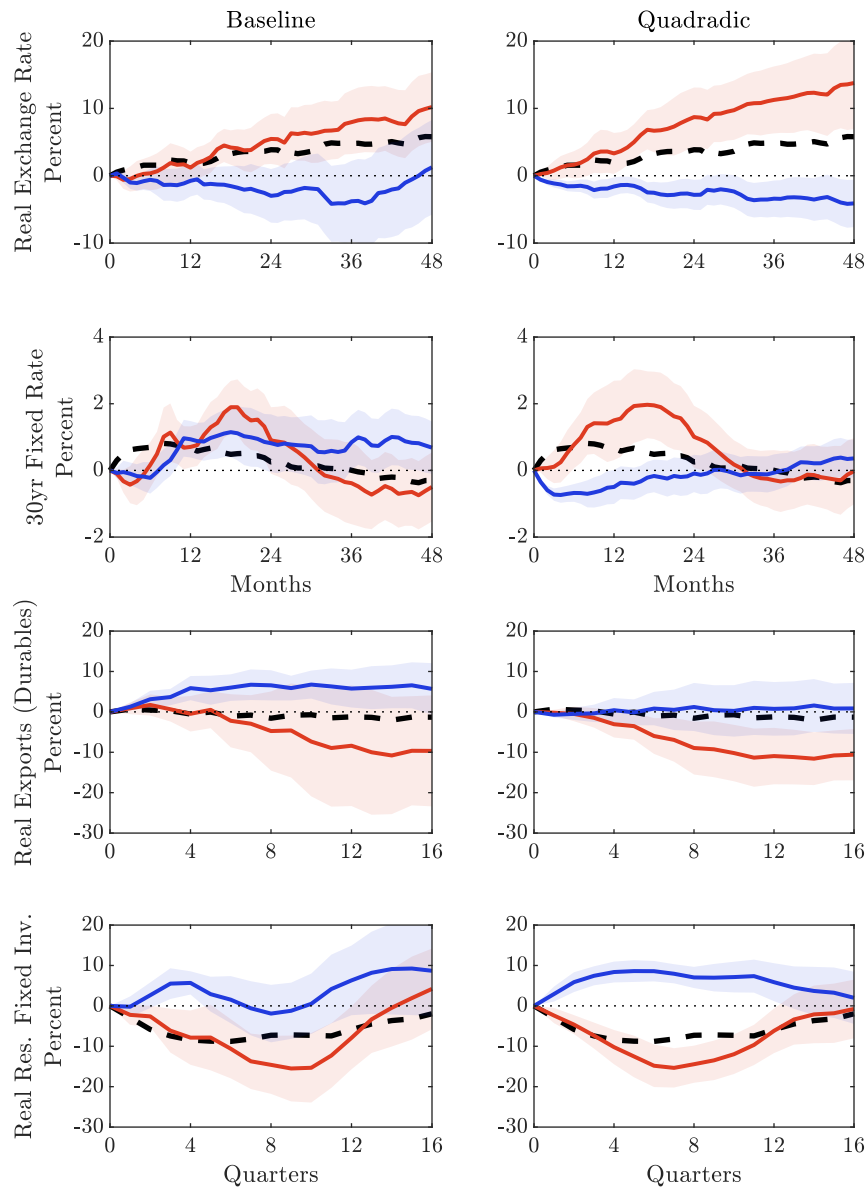


Figure 10: Impulse responses to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock. Estimated from Equation 4. Dashed lines are point estimates from 3. Shaded areas are 90 percent confidence intervals.

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.

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

	Horizon (months)	Linear Eqn. 2	Baseline Eqn. 3	Quadratic Eqn. 4
Routine	12	0.18	0.23	0.23
Employment	24	0.18	0.32	0.33
	36	0.11	0.23	0.24
	48	0.09	0.16	0.17
	<i>Max.</i>	<i>0.19</i>	<i>0.33</i>	<i>0.33</i>
	Nonroutine	12	0.03	0.07
Employment	24	0.01	0.04	0.04
	36	0.01	0.05	0.03
	48	0.01	0.06	0.04
	<i>Max.</i>	<i>0.05</i>	<i>0.10</i>	<i>0.11</i>
	Routine	12	0.16	0.22
Share	24	0.20	0.38	0.37
	36	0.13	0.32	0.32
	48	0.10	0.23	0.22
	<i>Max.</i>	<i>0.22</i>	<i>0.40</i>	<i>0.37</i>
	Total	12	0.18	0.24
Employment	24	0.15	0.27	0.29
	36	0.08	0.18	0.19
	48	0.07	0.11	0.12
	<i>Max.</i>	<i>0.20</i>	<i>0.28</i>	<i>0.29</i>

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.

References

- Acemoglu, D. (1999), ‘Changes in unemployment and wage inequality: An alternative theory and some evidence’, *American Economic Review* **89**(5), 1259–1278.
- Angrist, J. D., Jordà, O. and Kuersteiner, G. M. (2016), ‘Semiparametric estimates of monetary policy effects: String theory revisited’, *Journal of Business and Economic Statistics* pp. 1–17.
- Auerbach, A. and Gorodnichenko, Y. (2012a), Fiscal multipliers in recession and expansion, in A. Alesina and F. Giavazzi, eds, ‘Fiscal Policy after the Financial Crisis’, University of Chicago Press.
- Auerbach, A. and Gorodnichenko, Y. (2012b), ‘Measuring the output responses to fiscal policy’, *American Economic Journal: Economic Policy* **4**(2), 1–27.
- Autor, D. 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, D. H. and Katz, L. F. (1999), Changes in the wage structure and earnings inequality, in O. Ashenfelter and D. Card, eds, ‘Handbook of Labor Economics, Vol. 3’, Elsevier Science B.V., pp. 1463–1555.
- Autor, D. H., Levy, F. and Murnane, R. J. (2003), ‘The skill content of recent technological change: an empirical exploration’, *Quarterly Journal of Economics* **118**(4), 1279–1333.
- Barnichon, R. A. and Matthes, C. (2017), Functional approximations of impulse responses: New insights into the asymmetric effects of monetary policy. Working paper.
- Beaudry, P., Moura, A. and Portier, F. (2015), ‘Reexamining the cyclical behavior of the relative price of investment’, *Economic Letters* **135**, 108–111.
- Christiano, L. J., Eichenbaum, M. and Evans, C. (1996), ‘The effects of monetary policy shocks: Evidence from the flow of funds’, *Review of Economics and Statistics* **78**(1), 16–34.
- Christiano, L. J., Eichenbaum, M. and Evans, C. (1999), ‘Monetary policy shocks: What have we learned and to what end?’, *Handbook of Macroeconomics* **1**, 65–148.
- Coibion, O. (2012), ‘Are the effects of monetary policy shocks big or small?’, *American Economic Journal: Macroeconomics* **4**(2), 1–32.
- Coibion, O., Gorodnichenko, Y., Kueng, L. and Silvia, J. (2012), Innocent bystanders? monetary policy and inequality in the u.s. NBER Working Paper No. 18170.
- DiCecio, R. (2009), ‘Stick wages and sectoral labor comovement’, *Journal of Economic Dynamics & Control* **33**, 538–553.

- Fisher, J. D. M. (2006), ‘The dynamic effects of neutral and investment specific technology shocks’, *Journal of Political Economy* **114**(3), 413–451.
- Foote, C. L. and Ryan, R. W. (2015), Labor-market polarization over the business cycle, *in* ‘Macroeconomics Annual’, Vol. 29(1).
- Gertler, M. and Karadi, P. (2015), ‘Monetary policy surprises, credit costs, and economic activity’, *American Economic Journal: Macroeconomics* **7**(1), 44–76.
- Gilchrist, S. and Zakrajsek, E. (2012), ‘Credit spreads and business cycle fluctuations’, *American Economic Review* **102**(4), 1692–1720.
- Goos, M. and Manning, A. (2007), ‘Lousy and lovely jobs: the rising polarization of work in Britain’, *The Review of Economics and Statistics* **89**(1), 118–133.
- Gordon, R. J. (1990), *The Measurement of Durable Goods Prices*, University of Chicago Press.
- Jaimovich, N. and Siu, H. E. (2014), The trend is the cycle: Job polarization and jobless recoveries. Working paper.
- Jordà, O. (2005), ‘Estimation and inference of impulse responses by local projection’, *American Economic Review* **95**(1), 161–182.
- Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V. and Violante, G. L. (2000), ‘Capital-skill complementarity and inequality: A macroeconomic analysis’, *Econometrica* **68**(5), 1029–1053.
- Newey, W. K. and West, K. D. (1987), ‘A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix’, *Econometrica* **55**(3), 703–708.
- Owyang, M. T., Ramey, V. A. and Zubairy, S. (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, A. (1984), ‘Econometric issues in the analysis of regressions with generated regressors’, *International Economic Review* **25**(1), 221–247.
- Ramey, V. A. (2016), Macroeconomic shocks and their propagation, *in* J. B. Taylor and H. Uhlig, eds, ‘Handbook of Macroeconomics, Vol. 2’, Elsevier, pp. 71–162.
- Ramey, V. A. and Zubairy, S. (2017), Government spending multipliers in good times and in bad: Evidence from U.S. historical data. *Journal of Political Economy*, forthcoming.
- Romer, C. D. and Romer, D. H. (2004), ‘A new measure of monetary shocks: Derivation and implications’, *American Economic Review* **94**(4), 1055–1084.
- Tenreyro, S. and Thwaites, G. (2016), ‘Pushing on a string: US monetary policy is less powerful in recessions’, *American Economic Journal: Macroeconomics* **8**(4), 43–74.

Wong, A. (2016), Transmission of monetary policy to consumption and population aging.
Working paper.

A Appendix

The appendix includes a presentation of alternative monetary policy shock identification techniques; a discussion of the sensitivity of estimates from Equation 3 to different lag specifications; and a FEV analysis for alternative MSFEs.

A.1 Alternative 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 thirty-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).²⁶

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 1973,²⁷ but I include additional variables one by one, as in Christiano, Eichenbaum and 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 11. The estimated magnitudes are smaller than, but broadly consistent with, the IRFs estimated in Section 4.1.

²⁶Since the VARs estimated here are essentially identical to the baseline in Gertler and Karadi (2015), apart from trivially adding an additional variable to the system, for the details of the estimation, the reader is referred to their paper.

²⁷The results are essentially identical for the Gertler and Karadi time period.

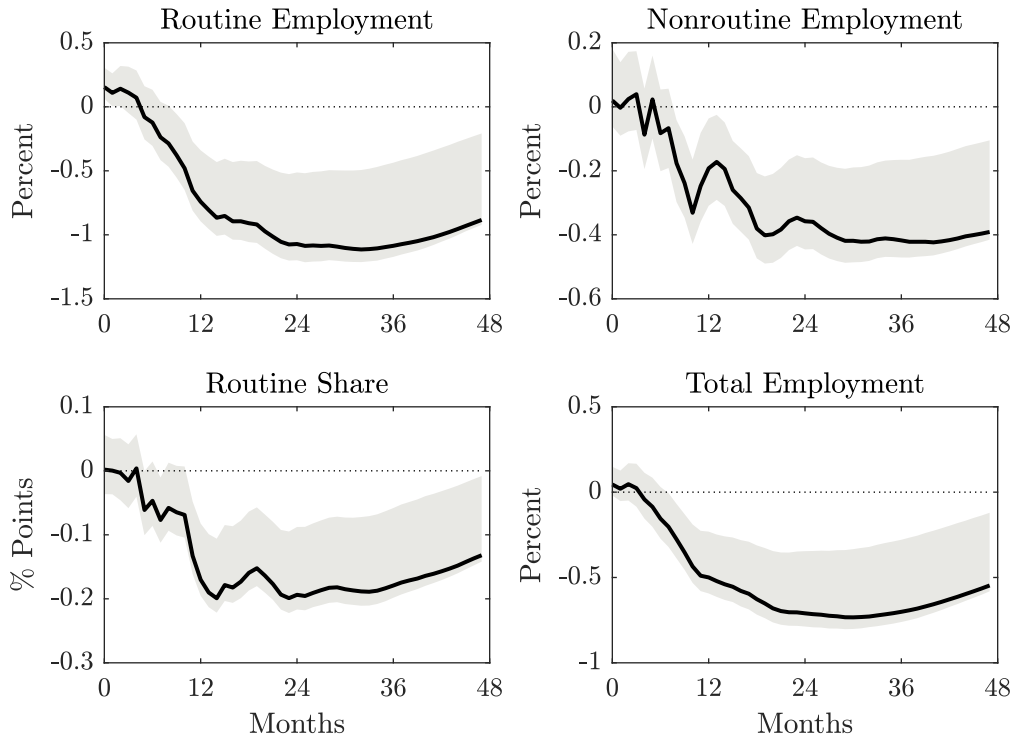


Figure 11: Impulse responses of occupation employment to a contractionary Gertler and Karadi (2015) monetary policy shock that increases the one-year rate 100 b.p. on impact. Estimated from the five-variable VAR discussed in Section A.1. Shaded areas are 90 percent bootstrap confidence intervals.

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).

A.2 Alternative lag structures

This section considers some alternative lag structures for estimates obtained from Equation 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.

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 12, 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 13, 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 (3); it does not appear in any polynomial specification, regardless of the number of lags included.

A.3 Alternative MSFE comparisons

The coefficients estimated from (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 (5) with (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 at their mean values.²⁸ As an alternative, therefore, I consider the forecast error from (5) conditional on

²⁸That is, zero in the linear or polynomial cases.

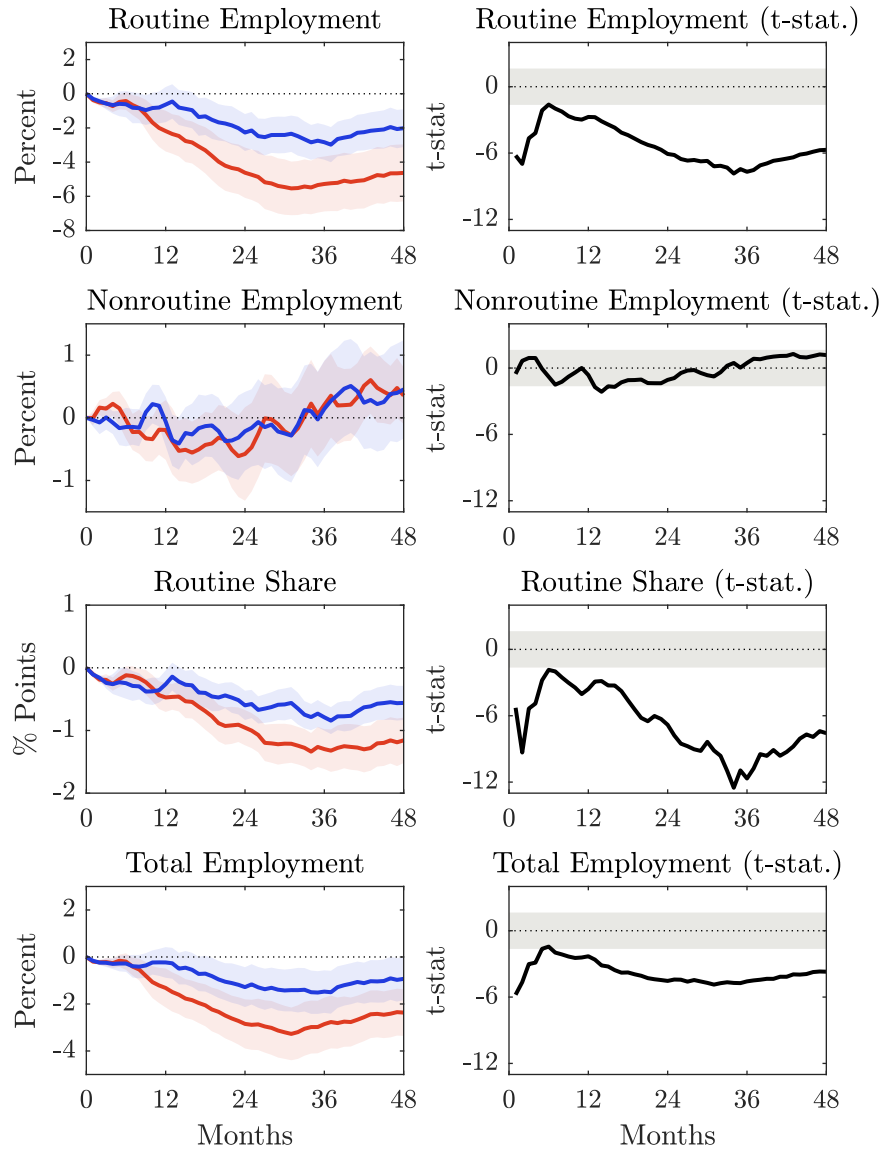


Figure 12: **Fewer lags:** *Left column:* Impulse responses of occupation employment to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock, no lags of the shock. Estimated from Equation 3. *Right column:* The same responses estimated from Equation 4. Shaded areas are 90 percent confidence intervals.

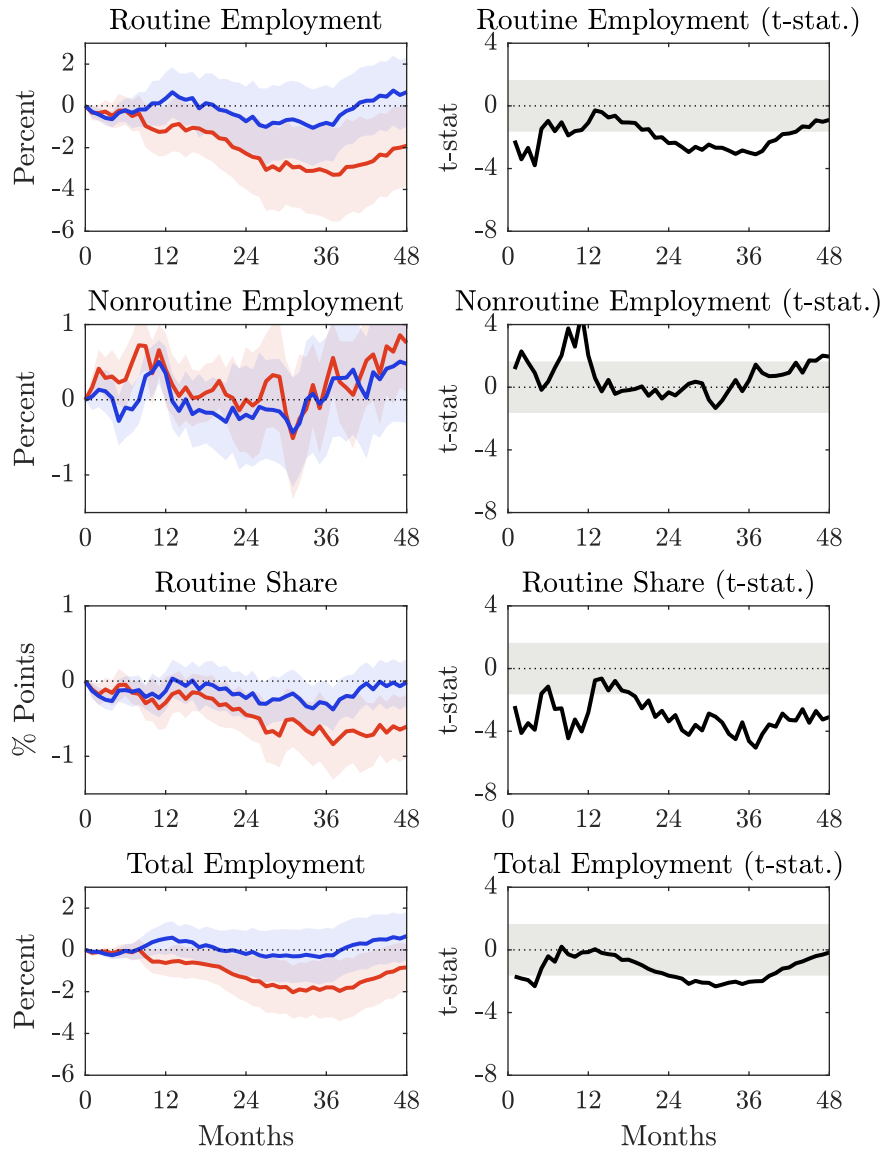


Figure 13: **More lags:** *Left column:* Impulse responses of occupation employment to a 100 b.p. contractionary (red) and expansionary (blue) Romer and Romer (2004) monetary policy shock, with 2 years of lagged dependent variables and 3 years of lagged shocks. Estimated from Equation 3. *Right column:* The same responses estimated from Equation 4. Shaded areas are 90 percent confidence intervals.

$\tilde{\mathbf{x}}_t$ only:

$$\begin{aligned}
\widetilde{FE}_{t,h} &\equiv y_{t+h} - y_t - \mathbb{E}[y_{t+h} - y_t \mid \tilde{\mathbf{x}}_t] \\
&= y_{t+h} - y_t - \widetilde{\boldsymbol{\beta}}_h' \tilde{\mathbf{x}}_t - \boldsymbol{\theta}'_h \mathbb{E}[\tilde{\boldsymbol{\epsilon}}_t \mid \tilde{\mathbf{x}}_t] \\
&= u_{t,h} + \boldsymbol{\theta}'_h [\tilde{\boldsymbol{\epsilon}}_t - \mathbb{E}[\tilde{\boldsymbol{\epsilon}}_t \mid \tilde{\mathbf{x}}_t]].
\end{aligned} \tag{7}$$

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

For the linear local projections in (2), the term $\mathbb{E}[\tilde{\boldsymbol{\epsilon}}_t \mid \tilde{\mathbf{x}}_t]$ theoretically should be zero since the monetary policy shock is unconditionally mean zero and should not be forecastable by macroeconomic variables.²⁹ For the asymmetric projections in (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 (7) is numerically identical to the unconditional mean—that is, the shocks are not forecastable. I estimate $\mathbb{E}[\tilde{\boldsymbol{\epsilon}}_t \mid \tilde{\mathbf{x}}_t]$ from the linear projection

$$\tilde{\boldsymbol{\epsilon}}_t = \boldsymbol{\Gamma} \tilde{\mathbf{x}}_t + \boldsymbol{\nu}_t. \tag{8}$$

Note that $\tilde{\mathbf{x}}_t$ includes deterministic terms (in the baseline, a constant and linear time trend). The conditional expectation in (7) is then just given by $\hat{\boldsymbol{\Gamma}}$, the estimate of the coefficient matrix in (8).

It is straightforward to verify that both (6) and (7) have a mean of zero, so that for each the MSFE is variance of the forecast error. Therefore, for a given horizon, the relative difference in the FEV between (5) and either (6) or (7) can be interpreted as the share of the FEV explained by the monetary policy shock.

Table 2 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.

²⁹I have verified that this is true in practice as well.

Table 2: Share of FEV due to monetary shocks — Occupations

	Horizon (months)	Linear		Baseline		Quadratic	
		(6)	(7)	(6)	(7)	(6)	(7)
Routine	12	0.18	0.19	0.23	0.24	0.23	0.23
Employment	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
	<i>Max.</i>	<i>0.19</i>	<i>0.20</i>	<i>0.33</i>	<i>0.35</i>	<i>0.33</i>	<i>0.34</i>
Nonroutine	12	0.03	0.03	0.07	0.07	0.07	0.07
Employment	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
	<i>Max.</i>	<i>0.05</i>	<i>0.05</i>	<i>0.10</i>	<i>0.10</i>	<i>0.11</i>	<i>0.11</i>
Routine	12	0.16	0.16	0.22	0.22	0.20	0.20
Share	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
	<i>Max.</i>	<i>0.22</i>	<i>0.23</i>	<i>0.40</i>	<i>0.41</i>	<i>0.37</i>	<i>0.37</i>
Total	12	0.18	0.19	0.24	0.24	0.24	0.24
Employment	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
	<i>Max.</i>	<i>0.20</i>	<i>0.21</i>	<i>0.28</i>	<i>0.29</i>	<i>0.29</i>	<i>0.29</i>

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 5.