Deutsche Bundesbank Discussion Paper No XX/2018

Monetary Policy and the Added Worker Effect*

John Leahy
University of Michigan & NBER

Tereza Ranošová
Deutsche Bundesbank

Abstract

We exploit cross-sectional variation in the response of US states to a monetary policy shock to study how the impact of monetary policy varies with the share of married women who work. We find that the economy's response is more muted the lower the share of married women employed before the shock. We argue that a plausible explanation is a shielded demand response by households, insured by the "added worker effect". When women are only weakly attached to the labor market, they can flexibly enter and exit to supplement household income in times of need, providing a powerful insurance mechanism against aggregate shocks. We provide three additional pieces of evidence. First, monetary policy shocks have a stronger effect in states where married women are more firmly attached to the labor market (making fewer transitions in and out). Second, following an increase in the federal funds rate, married women themselves are comparatively more likely to be employed (and to enter employment) in states where the share of married women working pre-shock is low. Third, in contrast to employment, wages of married women fall more in states where married women have worked less, consistent with a differential labor supply response to the shock.

Keywords: Added-worker effect, intrahousehold insurance, monetary policy

JEL classification: J21, J11, E24, E52

*Contact address: Deutsche Bundesbank, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main, Germany. E-Mail: jvleahy@umich.edu, tereza.ranosova@bundesbank.de. We would like to thank Sara Casella, Mel Stephens, Aditi Thapar, Michele Tertilt and Ana Lamo for helpful comments. We are grateful to the National Science Foundation for their support under grant SES-1919362. The views expressed in this paper are those of the author(s) and do not necessarily coincide with the views of the Deutsche Bundesbank or the Eurosystem.

Introduction

Macroeconomists typically assume that the economy is comprised of a representative agent or a collection of individual agents. Actual economies are comprised of families, and families can adjust to their economic environment in ways that individuals cannot. They can pool income and insure consumption. They can conserve resources by sharing ownership of durables such as homes and automobiles. They can adjust their work effort along both the extensive and intensive margin. They can more easily coordinate work effort across market and home production within a period and over time. The question arises whether these margins of adjustment matter for the behavior of aggregate quantities of interest to macroeconomists.

In this paper we argue that families matter. In particular, we show that the impact of monetary policy depends on the fraction of married women who work. The greater is this fraction, the greater is the impact of monetary policy on aggregate employment. A priori, it is not obvious how the labor force participation of married women should affect monetary policy. On the one hand, women tend to work in industries that are less cyclically sensitive, so more married women working might indicate an economy that is less sensitive to shocks. On the other hand, married women may be less firmly attached to the labor market, in which case greater numbers might exacerbate the impact of shocks to the economy. We argue that neither of these intuitions explain our result. We show that the result survives when we control for industrial structure, and we show that when the fraction of married women choosing to work is high, the married women are also more firmly attached to the labor market.

Instead, our proposed explanation relies on the "added worker effect" (Lundberg (1985), Stephens (2002), Parker and Skoufias (2004), Juhn and Potter (2007)) in which non-working women provide insurance to the family by entering employment when their husbands become unemployed. In this way the family uses female labor supply as a form of insurance against shocks to household income. Our story is that when the labor force participation of married women is high, labor force attachment of married women is also high and their employment decisions are similar to those of men. When the labor force participation of married women is low, however, their labor force attachment is also low. They are closer to the margin of working or not. In this case, their decisions of whether to work respond to the needs of the family, and they enter employment when their husband's employment prospects worsen. In this way, the added worker effect mitigates the impact of monetary policy shocks. Monetary policy shocks therefore have greater impact on the economy when the employment of married women is high and the added worker effect is weak.

We begin by presenting aggregate evidence on the relationship between the employment of married women and the impact of monetary policy. This evidence is interesting in its own right regardless of the underlying economic mechanism. We consider the impact of U.S. monetary shocks across U.S. states. To identify monetary policy shocks, we follow recent practice and consider changes in the federal funds rate over a short window around policy announcements.[1] We ask whether these shocks have greater impact on aggregate employment and earnings in a state when the share of married women that are employed

[1]Prominent examples of this approach include Gurkaynak, Sack, and Swanson (2005), Gertler and Karadi (2015) and Nakamura and Steinsson (2018).

in that state is large. Our preferred measure of the share of married women employed is a relative measure. It is the share of married women that are employed relative to the overall share of the state's population that is employed. The advantage of this relative measure is that it abstracts from state characteristics that impact the labor force participation decisions of all demographic groups. Our results, however, survive when we use an absolute measure.

As stated above, we find that the impact of these shocks on employment and earnings is greater, the greater is the (relative) fraction of married women that were employed before the shock. Interestingly, when we look at individual subgroups (married women, married men, unmarried women and unmarried men) we find that the employment of all groups responds more to monetary policy when married women work more. This may reflect general equilibrium spillovers: when the fraction of married women in employment is large, the economy is more responsive to monetary policy shocks and this response spills over to labor demand for all groups in the economy.

Given all of this co-movement, one might ask, "Why married women?" Why not women in general, or men? While the relative employment share of married women is highly correlated with the relative employment share of women, it turns out not to be highly correlated with overall labor force participation. This is one of the advantages of using the relative measure. And while it is true that monetary policy shocks also have greater impact in states in which women work, when we consider the impact of both married women and women together on monetary policy, the employment share of married women remains significant and impactful, whereas that of women in general does not. This is the reason we focus on married women.

While we have plausibly exogenous monetary shocks, our measures of employment of married women are endogenous. Lacking instruments, we do the next best thing and consider a number of confounding variables. We therefore control for specific state characteristics that evolve over time and have been shown to alter the impact of monetary policy. We consider variables related to the age structure of a state (Leahy and Thapar, 2022) since both marriage and work have a life-cycle component. We consider variables that reflect the industrial structure of a state (Carlino and DeFina, 1998) since working women may work in less cyclically sensitive industries. We also consider state income (Cravino, Lan, and Levchenko, 2020), since the need to insure against shocks to income may be correlated with the level of income. We consider variables that reflect the housing market (Beraja, Fuster, Hurst, and Vavra, 2019) since marriage and home ownership may be correlated. In all cases, the impact of the share of married women working on monetary policy survives. The impact of a monetary policy shock increases with the share of married women working even after controlling for these possible confounding effects. In some sense, this is not surprising since our preferred measure of employment of married women is a relative measure, and it turns out that this relative measure is only weakly correlated with many of these state characteristics.

Given the absence of obvious confounding variables, we explore further the possibility that the added worker effect may explain our results. We do this in two ways. First, if the added worker effect becomes weaker as it becomes more common for married women to work, then we would also expect that employment attachment would become stronger as more married women work. We therefore explore the relationship between the share of married women working and employment attachment. We measure attachment by the

number of transitions between employment, unemployment and out-of-the-labor force, with a larger number of transitions taken to indicate low attachment. We show that this measure of attachment of married women is negatively correlated with the share of married women working. We also show that, consistent with the added worker effect, when this measure of attachment is high, so is the impact of monetary policy on employment and wages.

Second, we look at how monetary policy shocks affect the choices of individual married women and how these choices vary with the relative share of married women working in their state. We find that in states where more married women are working, married women are less likely to be employed, less likely to be in the labor force, and more likely to be unemployed following a monetary policy tightening. These results mirror our aggregate results. We also find that they are less likely to transition into employment and receive higher wages. These results are consistent with the added worker effect being weaker in these states. We also find that the wages of married women fall by less in response to a monetary tightening the greater is the share of married women in the labor force. This is also consistent with the added worker effect. The added worker effect is fundamentally a labor supply response. When the added worker effect is strong, entry by married women in response to a monetary tightening tends to suppress wages.

The next section surveys the relationship between our work and the literature. In Section 2, we discuss our baseline empirical specification, data sources, and identification strategy for monetary policy shocks. Then in Section 3 we present aggregate evidence on the relationship between the share of married women working and the impact of monetary policy. Section 4 focuses on attachment to employment and Section 5 considers the behavior of individuals. Section 6 concludes.

1 Literature

Our paper contributes to several literatures. The first is the literature on the added-worker-effect. Lundberg (1985), Stephens (2002), Parker and Skoufias (2004), Juhn and Potter (2007) study this effect at the household level.[2] They show that when a husband becomes unemployed, his wife is more likely to enter the labor force. In addition, Ellieroth (2019) and Casella (2022) argue that wives respond to increases in unemployment risk not just by working more, but also by quitting less. The literature has also pointed to forces that may increase or reduce the importance of the added worker effect over time. Mankart and Oikonomou (2016) and Mankart, Oikonomou, and Pascucci (2022) argue that idiosyncratic income risk has increased over time and that this increase has made the insurance provided by the added worker effect more valuable. Casella (2022) argues that women have become more firmly attached to the labor market and thus less marginal over time, that this has muted the added worker effect, and thus made the economy more cyclically sensitive. At a broad level, our empirical results illustrate the macroeconomic importance of the added worker effect. More specifically, our results lend support to the mechanism proposed by Casella (2022) by showing that as married women become more

[2]The added worker effect traces back at least to Woytinsky (1942) who argued that labor force participation was boosted by a number of added workers who entered the labor force when their spouses lost their jobs.

attached to the labor market, the importance of this effect diminishes and the economy becomes more sensitive to monetary policy shocks.

Several papers have studied the consequences of the added-worker-effect in general equilibrium models. Mankart and Oikonomou (2017) and Guner, Kulikova, and Valladares-Esteban (2025) show that the added worker effect explains acyclicality of aggregate labor force participation. Blundell, Pistaferri, and Saporta-Eksten (2016) show that insurance within couples is important to smooth consumption. Wang (2019) studies a coordinated joint search model including the added worker effect. Wu and Krueger (2021) show the consequences for optimal tax policies. The two most important papers, from our perspective, are Casella (2022) and Bardoczy (2022). Casella, discussed above, models the interaction between the added worker effect and labor force attachment of married women. Bardoczy (2022) shows that in a heterogeneous household model, the added worker effect only needs to be observed in a small fraction of households for it to have a large impact on the evolution of aggregates. This is consistent with our finding that the greater is the attachment of married women to the labor force (and hence the weaker is the added worker effect), the more volatile is employment for all workers, men and women.

We also contribute to the literature connecting the rapid change in women's labor market attachment to the economy's sensitivity to shocks more broadly. For example, Fukui, Nakamura, and Steinsson (2023), Albanesi (2019) and Olsson (2025) argue that the large trend increase in women's labor force participation over the second half of the 20th century made recoveries from recessions look fast along this transition. In other words, recoveries of men have always been slow, but the aggregate recoveries started to look markedly slower once women already reached a certain level of labor force participation and slowed their convergence.

More broadly, our paper relates to the literature studying the role of labor market flows (and specifically of labor supply) in responses to monetary policy. Graves, Huckfeldt, and Swanson (2023) estimate the response of labor market flows to monetary policy shocks and find evidence that, in contrast to the consensus view, a contractionary monetary policy shock leads to a significant increase in labor supply. White, Bhattarai, Glover, Petrosky-Nadeau, Afrouzi, Hagemann, Ryngaert, and Koustas (2018) show that a full stock-flow accounting reveals that job loss is the largest driver of monetary policy's effects on the labor market. Broer, Kramer, and Mitman (2025) show that in Germany, job and income losses rise more following a monetary contraction for lower-earnings households. Faia, Shabalina, and Wiczer (2024) show that within the subpopulation of bellow-median income, monetary policy tightening raises separations more among older, female and non-white individuals. We show that cross-state variation in how much married women tend to change labor market status predicts a differential response of state-level employment to monetary shocks. Moreover, we show that cross-state variation in the share of married women working pre-shock predicts cross-state variation in the entry probability of married women following a monetary shock, and of exit probability of the rest.

Finally, we contribute to the literature that considers features of economies that alter the impact of monetary shocks. Carlino and DeFina (1998) document heterogeneity in the impact of monetary policy shocks across US states, and attribute these differences to industrial mix. Imam (2015) and Leahy and Thapar (2022) argue that the impact of monetary policy shocks depends on the age structure of the economy. Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru, and Yao (2017) and Beraja et al. (2019) argue the

structure of mortgages and recent house price movements may affect the transmission of monetary policy. Cravino et al. (2020) have argued that income may matter as richer individuals consume a greater share of sticky-price goods. We add the employment share of married women to this list.

Our results may seem counter-intuitive at first. A broad literature documents that employment of women is more stable and less cyclical, especially due to the differences in the sectoral composition (see Doepke and Tertilt (2016) for an overview). It might thus seem obvious that economies where women form a larger share of the workforce should be less volatile. In fact, Mennuni (2019) shows that across U.S. states and the time period 1982-2013 higher female labor force participation is associated with lower volatility of output. In contrast, we show that economies where women are already well attached to the labor market, and thus form a higher share of the overall workforce on average, are more sensitive to monetary policy shocks. There is no inherent contradiction between these findings. On the one hand, female-dominated sectors may be less likely to be hit by sector-specific shocks. On the other hand, the responses to a common monetary shock may hit harder states with a larger fraction of married women employed.

2 Aggregate evidence across US states

We begin by considering the impact of monetary policy shocks across US states, and how the impact of these shocks varies with the employment decisions of married women. We show that monetary policy shocks have a greater impact on employment and payroll in states in which a higher share of married women worked pre-shock. This is consistent with the added worker effect if a high share of employment of married women is indicative of high attachment to employment. We investigate the connection between employment attachment and share of married women working in the next section.

We consider the U.S. states to be a panel of economies. This has several advantages. First, since the U.S. is a monetary union all states are subject to the same monetary policy. We do not have to control for differences in monetary policy rules across states, and we do not have to worry that monetary policy might respond to state characteristics such as the share of married women in the workforce. Second, we can use a single common strategy to identify monetary policy shocks across all states. Third, U.S. states have relatively similar legal and institutional structures, removing one potential source of heterogeneity. Problems remain. Immigration and trade connect states. States differ in their industrial structure. These problems, however, are also present among nation states, and we will try to control for them in our robustness checks.

In the remainder of this section we describe our empirical model, discuss the identification of monetary policy shocks, and provide a brief summary of the data that we use to estimate the model. We present our results in the next section.

2.1 Specification

Our baseline specification applies the Jorda projection (Jorda, 2005) in a panel-data context:

$$\Delta \ln y_{s,t,k} = \beta^k \cdot \epsilon_t \cdot S^{MW}_{s,t-1} + \gamma^k \cdot S^{MW}_{s,t-1} + \alpha_t + \alpha_s + u_{s,t+k}. \tag{1}$$

Here s indexes the U.S. state, t indexes time, and k indexes the horizon. $\Delta y_{s,t,k}$ is the percent change in the dependent variable in state s between date $t-1$ and date $t+k$. This dependent variable is typically some measure of employment or total payroll. ϵ_t is a monetary policy shock which we discuss in detail in the next section. $S^{MW}_{s,t-1}$ is a measure of how much married women work. Our main measure is the ratio of the employment rate of married women relative to the employment rate in the economy as a whole:

$$S^{MW}_{s,t} = \frac{E^{MW}_{s,t}}{P^{MW}_{s,t}} \Bigg/ \frac{E_{s,t}}{P_{s,t}}$$

where $P_{s,t}$ is the population of state s at date t, $E_{s,t}$ is the number of employed persons in state s at date t, $P^{MW}_{s,t}$ is the total number of married women in state s at date t, and $E^{MW}_{s,t}$ is the number of married women who are employed in state s at date t. We limit these counts to individuals age 16 and older. To remove seasonality, $S^{MW}_{s,t}$ is a four quarter moving average.[3]

The state fixed effects, α_s, control for persistent differences across states in characteristics such as industrial structure or labor force participation, whereas the time fixed effects, α_t, control for shocks that hit all states at a given point in time such as the state of the aggregate business cycle, the level of interest rates or the exchange rate. The time effects also soak up the first order impact of the monetary policy shock ϵ_t. Finally, $u_{s,t+k}$ is the error term.

We estimate this panel regression for each horizon k. The coefficients of interest are the β^k which capture how the impact of a monetary policy shock in period t on the dependent variable k periods in the future varies with employment decisions of married women. Our monetary policy shocks are associated with increases in interest rates. Positive shocks are therefore contractionary. A positive β^k means that an increase in the share of married women working reduces the impact that a shock to monetary policy has on the growth of the dependent variable. Negative values of β^k, on the other hand, imply that an increase in the share of married women working strengthens the contractionary effect of a monetary policy shock. We can trace out the effect that an increase in the proportion or married women working has on the response of y to a monetary policy shock by estimating 1 for each $k = 0, 1, 2, \ldots, K$. Coefficients are multiplied by 100 and thus the magnitudes represent percentage point changes in the growth of y. S^{MW} is multiplied by 100, thus the magnitudes represent an effect of increasing the relative share of married women working by one percentage point.

2.2 Monetary policy shocks

We use the same monetary policy shock series as Ottonello and Winberry (2020) and Leahy and Thapar (2022). For completeness, we include a discussion of their construction here.

In our baseline specification, we adopt the high-frequency approach to identifying monetary policy shocks used by Gurkaynak et al. (2005), Gertler and Karadi (2015)

[3]Results are qualitatively similar if we instead use $E^{MW}_{s,t}/P^{MW}_{s,t}$ as our measure of labor force participation. $S^{MW}_{s,t}$ has the advantage of controlling for factors that affect employment of both men and women married and unmarried, such as the age structure of the population.

and Nakamura and Steinsson (2018).[4] As noted by Nakamura and Steinsson (2018), the interpretation of these high-frequency shocks raises certain concerns. To address these concerns, we examine the robustness of our results to alternative shock measures in Section A.3 of the Appendix. Ottonello and Winberry (2020) provide a detailed discussion of the construction of these shocks; here, we offer a brief overview.

The construction of the shocks proceeds in two steps. First, we extract raw shocks from movements in the federal funds rate implied by the current-month federal funds futures contract within a narrow window around FOMC announcements, specifically, from 15 minutes before to 45 minutes after the announcement. This short window ensures that the observed changes are unlikely to reflect non-policy news.

Second, following Ottonello and Winberry (2020), we aggregate the high-frequency shocks to a quarterly frequency using a weighted moving average of shocks from the current and previous quarters. The weights account for the fact that employment and income evolve throughout the quarter, while monetary policy affects only decisions made after the shock. They are chosen such that an announcement at the end of quarter $t-1$ is equivalent to one at the beginning of quarter t.[5]

The identified shocks are small but meaningful: over the 1990-2008 sample, the average high-frequency shock is -5 basis points with a standard deviation of 11 basis points. About 40% of the shocks are contractionary and 60% expansionary, with roughly 70% more shocks than actual changes in the federal funds rate, indicating that even unchanged policy rates can represent monetary surprises.

Leahy and Thapar (2022) show that these shocks have the anticipated effect on aggregate activity. Positive shocks lead to a temporary increase in interest rates and a reduction in GDP, income and employment (see Figure 1 of that paper). They therefore look like monetary policy shocks.

2.3 Data

Details on the data are provided in the Data Appendix. We briefly summarize them here.

The U.S. Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) program is our main source for data on employment and payroll. The LEHD provides the number of private-sector jobs and total payroll by state, industry, and demographic group. The dataset forms an unbalanced panel, as states entered the program at different times.[6]

At times we will wish to construct measures of population, employment and number of transitions into and out of employment by demographic groups. For this we use the Current Population Survey (CPS) IPUMS data (Flood, King, Rodgers, Ruggles, Warren, Backman, Chen, Cooper, Richards, Schouweiler, and Westberry (2024)). We also utilize the CPS IPUMS sample to study responses across states on the individual level. For the analysis of employment we use the whole sample in each month. For the analysis of

[4] Rudebusch (1998), Kuttner (2001), Söderström (2001), and Thapar (2008) use interest rate futures to identify monetary policy shocks.

[5] This aggregation introduces some serial correlation in the shocks, which we control for in the estimation.

[6] Massachusetts is excluded from our estimation because it joined the LEHD panel in 2010, after the end of our sample period.

transitions we utilize the panel dimension within CPS.[7] For the analysis of wages, we take advantage of the 'earner study' subsample.[8] The earner study includes the respondent's hourly wage (for those who are paid hourly), usual weekly earnings on the main job (for those not paid hourly) and usual hours worked a week at the main job. We use earnings and hours to fill in an hourly wage measure for the whole sample.

In some of the robustness exercises, we use state-level descriptors from the replication datasets of Leahy and Thapar (2022) (industry shares and personal income per capita) and Boone, Dube, Goodman, and Kaplan (2021) (UI generosity). Finally, we identify high-frequency monetary policy shocks by leveraging information in futures markets, following Ottonello and Winberry (2020). In Appendix A.3 we consider the alternative shock series based on Romer and Romer (2004) narrative evidence, updated through 2007 by Wieland and Yang (2017).

Before presenting the results for estimating equation (1), we briefly discuss some of the properties of S^{MW}. Table 3 in the Appendix presents correlations of the share of married women working with other measures of labor force participation. The first column shows the correlation of these variables with S^{MW}. The second column shows the correlation of these variables with the absolute share of married women working $S^{MW}_{abs} = \frac{E^{MW}_{s,t}}{P^{MW}_{s,t}}$. The two measures are highly correlated. Importantly, controlling for the overall share of the population that is working, greatly reduces the correlation with other measures of labor force participation. For example the correlation between S^{MW} and overall labor force participation is only 0.1.

3 Baseline Results and Robustness

Our baseline sample is a quarterly panel of US states from 1990 Q2 to 2008 Q4. We begin in 1990 Q2 as that is the start of the LEHD. We end our sample period in 2008 Q4 once the federal funds rate hit the zero lower bound. As a robustness exercise, in Appendix A.3 we consider altering the start and end dates.

Figure 1 presents the results for estimating equation (1) taking as the dependent variable state employment or state payroll. Each figure plots the estimated coefficient β^k against the horizon k.

The solid lines in the figure illustrate the effect that a one percent increase in S^{MW} has on the response of the outcome variable to an increase in the federal funds rate. This response is over and above the common effect that monetary policy has on all states. This common component has been removed from the regression by time fixed effects. Leahy and Thapar (2022) show that this common effect has the standard properties of a monetary policy shock: employment and payroll both decline temporarily in response to an increase in interest rates. Negative values in the figure therefore indicate that a higher employment rate for married women causes a variable to fall further than average,

[7] Each respondent is interviewed for four consecutive months, then leaves the sample for an eight month break and is then added again for four months.

[8] The 'earner study' subsample consists of individuals who are about to rotate out of the panel, i.e. those who have been in the sample for 4 or 8 months respectively, who are currently employed and over 15 years old (self-employed individuals are excluded). No earnings or income information is available in the core survey. Extensive income information is included in the March CPS ASEC sample, however this would only allow analysis on an annual level.

Figure 1: Effect of an Increase in the Share of Married Women Working (relative to the average) on the Economy's Response to a Monetary Policy Shock

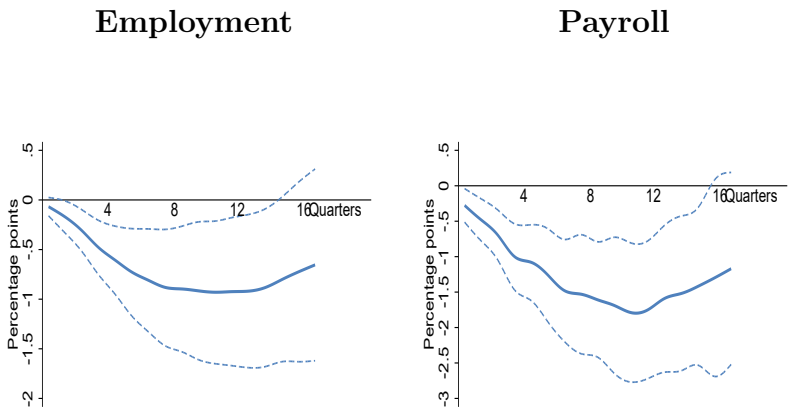

Each figure depicts the effects over time of a one percentage point increase in S^{MW} on the responsiveness of the relevant variable (state personal employment or payroll) to a monetary policy shock. Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

whereas positive values typically indicate that the variable falls by less. The dashed lines in the figure represent the 95 percent confidence intervals based on Driscoll-Kraay (1998) standard errors. These are meant to be robust to general forms of spatial and temporal correlation in the error terms.

The figure tells us that when the employment rate of married women is higher in a state, monetary policy shocks have a greater effect on labor market outcomes in that state. The impact on employment and payroll are qualitatively the same. The impact on payroll is somewhat larger, reflecting the fact that wages fall along with employment. The shape of the response in Figure 1 is very similar to the aggregate response in Leahy and Thapar indicating that a higher share of married women working amplifies the typical impact of monetary policy.[9]

Table 2 in Appendix A.1 presents coefficient estimates and standard errors. The main takeaway is that the effect on payroll is statistically significant at the one percent level between one and 11 quarters after the policy shock.

What explains the greater sensitivity to monetary policy in the face of strong female employment? In the next subsection, we dismiss several uninteresting explanations such as compositional effects and confounding variables. We also investigate the robustness of our results. Then in the next section we present evidence consistent with the added worker effect.

[9]The impulse response in Leahy and Thapar peaks at 10 quarters and dies out after 16 quarters.

3.1 Compositional effects, confounding variables, and robustness

To better understand our result, we consider a number of extensions. Some of these experiments consider alternative dependent variables. Others add potential confounding variables interacted with the shocks.

Women or Married Women

We have shown that the impact of monetary policy is larger in states in which the share of married women who work is large relative to the overall share of the population that is employed. A natural question is whether this result stems from behavior of married women or whether there is something about the labor market attachment of women in general that makes monetary policy more potent for aggregate employment. To answer this question we define S^W to be the (relative) share of women employed in the state,

$$S^W_{s,t} = \frac{E^W_{s,t}}{P^W_{s,t}} \Big/ \frac{E_{s,t}}{P_{s,t}}$$

and we add S^W to our baseline regression:

$$\begin{aligned} y_{s,t,k} = & \quad \beta^k \cdot \epsilon_t \cdot S^{MW}_{s,t-1} + \gamma^k \cdot S^{MW}_{s,t-1} + \\ + & \quad \theta^k \cdot \epsilon_t \cdot S^W_{s,t-1} + \delta^k \cdot S^W_{s,t-1} + \alpha_t + \alpha_s + u_{s,t,k} \end{aligned} \tag{2}$$

The coefficients of interest are β^k and θ^k which capture how S^{MW} and S^W effect the impact of monetary policy shocks. One can think of this as a horse race between the two variables.

Figure 2 presents the results of this experiment. The left-hand panel shows the impact of S^{MW}. It looks very similar to our baseline results. Adding S^W does little to alter the impact of S^{MW} on monetary policy. The right-hand panel shows the impact of S^W. S^W has little impact on monetary policy once we control for S^{MW}.[10] We interpret these results as justification for our focus on married women.

Compositional effects

In this subsection we consider subsets of the population. We take $\Delta \ln y$ in our baseline specification to be the change in employment of married women or the change in employment of all other adults in the state and ask if employment responds similarly for both groups.

Figure 3 presents the results. The responses of employment for the two subgroups are similar in magnitude and we cannot reject the hypothesis that responses are equal. Moreover, the responses of employment are similar to the aggregate response in Figure 1. Note that the data in this figure is aggregated from the IPUMS data, which is a bit more noisy than the LEHD data used in Figure 1. This may contribute to the wider error bands.

We draw two conclusions from this exercise. First, our baseline results do not appear to be simply due to the composition of the labor force. It is conceivable that the reason monetary policy has a greater impact in states with more married women is that monetary

[10] If we simply replace S^{MW} with S^W in our baseline specification, then the impact of S^W on monetary policy is qualitatively similar to the impact of S^{MW}.

Figure 2: Horse-race between share of married women working and share of all women working

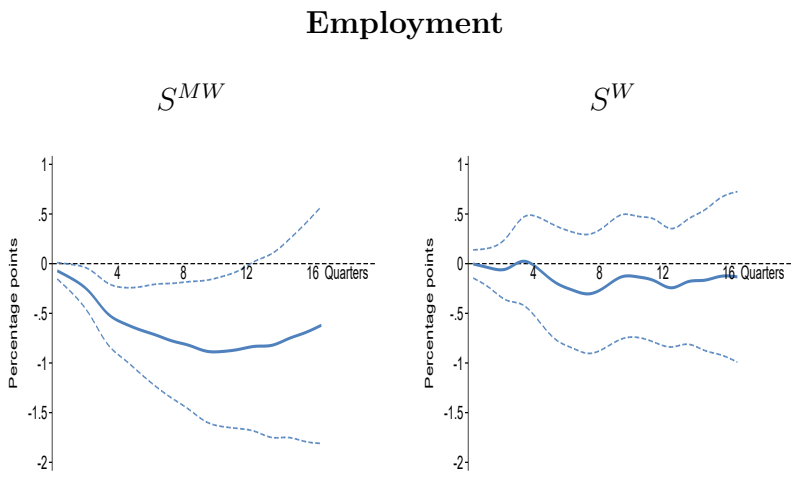

Each figure depicts the effects over time of a one percentage point increase in S^X on the responsiveness of employment to a monetary policy shock, where both S^{MW} and S^W (both relative measures) are included in the same regression (2). Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

Figure 3: Effect of an Increase in the Share of Married Women Working on Labor Market outcomes of Married Women and All Others

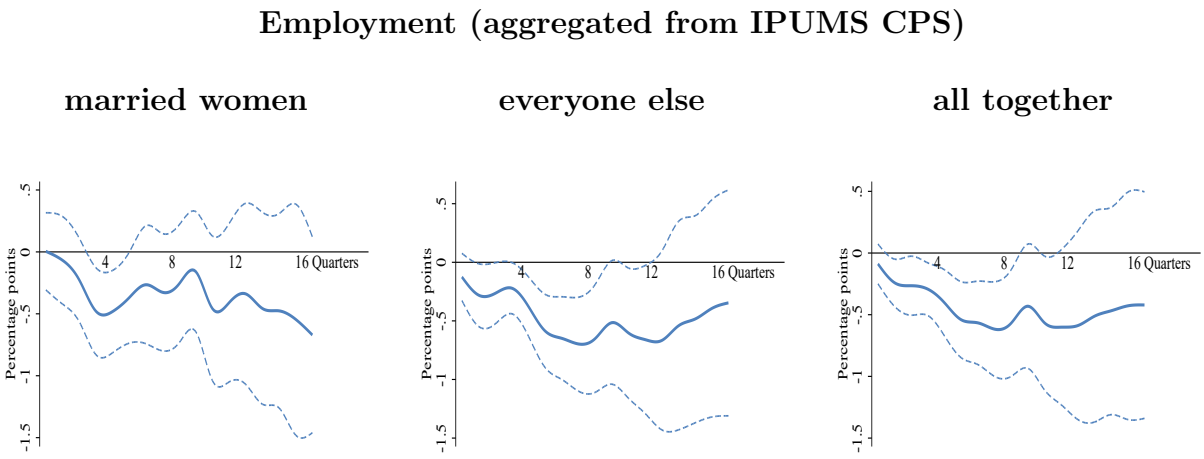

Using the relative measure of S^{MW}.

policy has a stronger impact on the employment of married women and there are more married women in these states. If this were the sole reason behind the results in Figure 1, however, we would not see any impact in figure 3.[11][12]

Second, whatever is the reason that the impact of monetary policy depends on the employment share of married women, the response of all others indicates that changes in this share influence how monetary policy impacts all groups in the economy. Changes in the employment share of married women appear to change the way in which the economy as a whole responds to monetary policy shocks.

Other Robustness Checks

It is possible that the effect of the share of married women on monetary policy is driven by some third variable correlated with the share. Definitively demonstrating a causal link between female employment and monetary policy requires an instrument. Given that we lack such instruments, our next set of experiments attempts to control for possible confounders.

These regressions take the form:

$$\Delta \ln y_{s,t,k} = \beta^k \cdot \epsilon_t \cdot S^{MW}_{s,t-1} + \gamma^k \cdot S^{MW}_{s,t-1} + \theta^k \cdot \epsilon_t \cdot X_{s,t-1} + \delta^k \cdot X_{s,t-1} + \alpha_t + \alpha_s + u_{s,t+k}. \quad (3)$$

where X is the confounding variable. Figures associated with these experiments are contained in appendix A.2. Here we summarize the results. In all cases, the impact of the employment share of married women on monetary policy remains unchanged.

- Our focus on married women might confound the impact of marriage with the impact of gender or labor force participation. Above we controlled for the share of women in the workforce. We also control for the share of men in the workforce and labor force participation. The results are shown in Figure (9). The impact of S^{MW} survives.

- Carlino and DeFina (1998) argue differences in the impact of monetary policy across states are due in part to differences in the mix of industries. The response of industries to monetary policy may be correlated with the gender composition of their workforce. We therefore consider including the share of manufacturing, construction, services, and also financial services as confounders. Again the impact of S^{MW} survives.

- Leahy and Thapar (2022) argue that the age structure of the state is correlated with the impact of monetary policy. Age may also be correlated with marriage and the employment decisions of married women. We consider as confounders the share of the population 25-40, the share of the population 40-65, and the share of the population 65 and up Figure (10). Again the impact of S^{MW} survives.

[11] It should be noted that such a composition effect seems unlikely on the face of it. A broad literature documents that employment of women is more stable and less cyclical (see Doepke and Tertilt (2016) for an overview). While cyclical sensitivity is not the same as responsiveness to monetary policy shocks, it seems unlikely that monetary policy has a stronger impact on the employment of married women.

[12] It should also be noted that Figure 3 does not test whether monetary policy has a greater effect on female or male labor market outcomes in the time series. It shows only how the effect of monetary policy changes with changes in the employment of married women. The average response that these changes operate off of is soaked up by the time fixed effects.

- Cravino et al. (2020) argue that higher income individuals consume goods with stickier prices and that the consumption of these goods is more sensitive to monetary policy shocks. States with more higher income might therefore respond more to monetary policy. As marriage may also be correlated with income, we consider state personal income per capita as a confounding variable in Figure (9). Again the impact of S^{MW} survives.

- Beraja et al. (2019) argue that there is a relationship between mortgage refinance and the impact of monetary policy. As marriage and ownership are correlated, we consider the home-ownership rate as a confounder. Again the impact of S^{MW} survives.

As an additional experiment, we included region fixed effects interacted with the monetary policy shock. Including these fixed effects should control for regional differences in the responsiveness to monetary policy shocks that might be correlated with cross-state differences in labor-market attachment of married women. The results appear in Figure 8 in appendix A.2. The impulse responses are similar to our baseline although the measured effects are slightly smaller and a bit more noisy.[13]

In some sense, it is not surprising that our results survive including these confounding variables. Table 3 in the appendix lists some of these variables and their correlation with S^{MW}. Many are only weakly correlated with the employment share of married women. The variable most highly correlated with S^{MW} is the share of the population between 40 and 65 years of age. This correlation is 0.41.

In addition to investigating the impact of confounders, we performed a number of additional robustness exercises. We show that the baseline estimates are robust to shifts in the sample and to dropping small states. We also show that they are robust to alternative methods of identifying monetary shocks. In particular, we show that the estimating (1) with the Romer and Romer (2004) shocks yields very similar results. Lastly, we show in Figure 11 that the results are robust to using the absolute measure of S^{MW}. These robustness checks can be found in Appendix A.3.

While we cannot establish causality definitively, the evidence leaves room for a causal role for the employment share of married women. We now present evidence that the added worker effect might be the mechanism by which the share of married women impacts monetary policy.

4 Marginality

Our proposed explanation for why states where married women work less on average are more insulated from monetary shocks is that when married women work less they are only marginally attached to the labor market so that they can time their participation as a response to the household's needs of income smoothing. If this hypothesis is correct, then it should be the case that married women enter and exit the employment more frequently in states in which fewer married women work. In this section we take a closer look at the

[13] If we consider state fixed effects interacted with the monetary policy shocks, the impact of the share of working women largely vanishes. Our power appears to come more from the cross section than the time series.

relationship between employment of married women and their attachment to the labor market.

To measure labor market attachment, we use the panel dimension of the CPS to measure how likely a married women is to change labor market status during a quarter. Let t^m index time as measured in months, while t remains a quarterly time index. Let $I_{i,t^m} = 1$ if a married woman i is employed in month t^m; $I_{i,t^m} = 2$ if she is unemployed; and $I_{i,t^m} = 3$ if she is out of the labor force. Let

$$T^{MW}_{s,t} = \sum_{t^m \in t, i \in s} \sum_{j=1,2,3} (I_{i,t^m} = j)(I_{i,t^m-1} \neq j)$$

denote the number of transitions from one employment status to another during quarter t. $T^{MW}_{s,t}$ is then the count of all monthly transitions between any pair of employment, unemployment and out-of-labor force observed for married women in state s and quarter t. Our first measure of mobility is the average number of such transitions per married woman in a quarter: $M^{MW}_{s,t} = \frac{T^{MW}}{P^{MW}}_{s,t}$[14]. Similarly to S^{MW} we also construct a relative measure of marginality for comparison $M^{MW}_{rel} = \frac{M^{MW}_{s,t}}{M^{all}_{s,t}}$ where $M^{all}_{s,t}$ is constructed in the same manner as $M^{MW}_{s,t}$ but using the universe of adults in the sample in period t rather than focusing on married women.

It turns out most transitions in the CPS microdata are incredibly short lived. We therefore exploit the fact that individuals in the CPS are interviewed for four months in a row, and construct a measure of 'stable' transitions in which a person is observed for two months in one employment status and then for two months in another[15] We denote this stable transition measure $M^{MW}_{rel,sticky}$

Table 1 compares these measures of marginality to our measure of the employment of married women, S^{MW}. Consistent with our hypothesis, these measures of marginality correlate negatively with the overall employment share of married women in a state. In other words, states where married women are less likely to work are often states where married women are more likely to make transitions in their labor market status.

With these measures we run cross-state panel regressions analogous to our baseline results:

$$\Delta \ln y_{s,t,k} = \beta^k \cdot \epsilon_t \cdot MX^{MW}_{s,t-1} + \gamma^k \cdot MX^{MW}_{s,t-1} + \alpha_t + \alpha_s + u_{s,t+k} \tag{4}$$

[14]As with $S^{MW}_{s,t}$, we use four quarter moving averages of T and P to minimize seasonality and small sample noise.

[15]For example, if we focus on transitions to employment between the second and third month of participation in the survey (where we have information on 1 month prior and 1 month post the transition), then only 39.8% of these transitions are 'stable'. In all other cases there is either a transition between months one and two or a transition between months three and four. Similarly, only 30.3% of transitions to labor force are 'stable'. Given two months is itself a very short period of time, we find it unlikely that many of the raw transitions can be interpreted as true changes in an optimized labor supply decisions. Rather they probably reflect being employed in a very temporary job or classification errors. This suggest the raw measure of transitions in the CPS is likely mixing up genuine sustainable switches with quick moves in and out that are more difficult to interpret. This is similar to the standard "de-NUN-ification" of typically applied to CPS data, as it is commonly noted that some transitions between unemployment and out-of-labor force are likely errors. We extend this logic to other transitions, especially the ENE and NEN transitions.

Table 1: Correlation between share of married women working and marginality of married women

		S^{MW}	M^{MW}	M^{MW}_{rel}
Share of mw working (rel to pop), 16+	S^{MW}	1.00		
Marginality (all Ts) of married women	M^{MW}	-0.44	1.00	
Marginality (all Ts) of married women relative to population	M^{MW}_{rel}	-0.37	0.96	1.00
Marginality (sticky Ts) of married women relative to population	$M^{MW}_{\text{rel, sticky}}$	-0.26	0.14	0.13

Where $MX^{MW} = M^{MW}, M^{MW}_{rel}, M^{MW}_{rel,sticky}$.

Figure 4 presents the results. All three measures of marginality tell a qualitatively similar story.[16] States where married women are more marginal with respect to the labor market exhibit weaker responses to monetary policy shocks, consistent with our hypothesis that labor force participation decisions of married women in these states may insulate them from these shocks. Reassuringly, these look similar when marginality measures are constructed with a restricted sample of married women who are 25-50 in order to abstract from retirement or first-time entrants into the labor market, as well as with more restricted marginality measures based on only moves to and from employment or to and from labor force. These results are presented in Appendix table 16. We also conducted a horse race between S^{MW} and both M^{MW}_{rel} and $M^{MW}_{rel,sticky}$. The impact of S^{MW} becomes slightly less significant but holds up pretty well. $M^{MW}_{rel,sticky}$ also holds up well, but the effects are smaller. M^{MW}_{rel} becomes insignificant, once we add S^{MW}. This is consistent with all three being imperfect proxies for the same mechanism, and with sticky transitions being more informative than raw transitions. These results are presented in Appendix table 17.

[16]Note that the range of the marginality measures differs from measure to measure so we cannot compare the quantitative impact.

Figure 4: Effect of an Increase in the 'Marginality' of Married Women on the Economy's Response to a Monetary Policy Shock

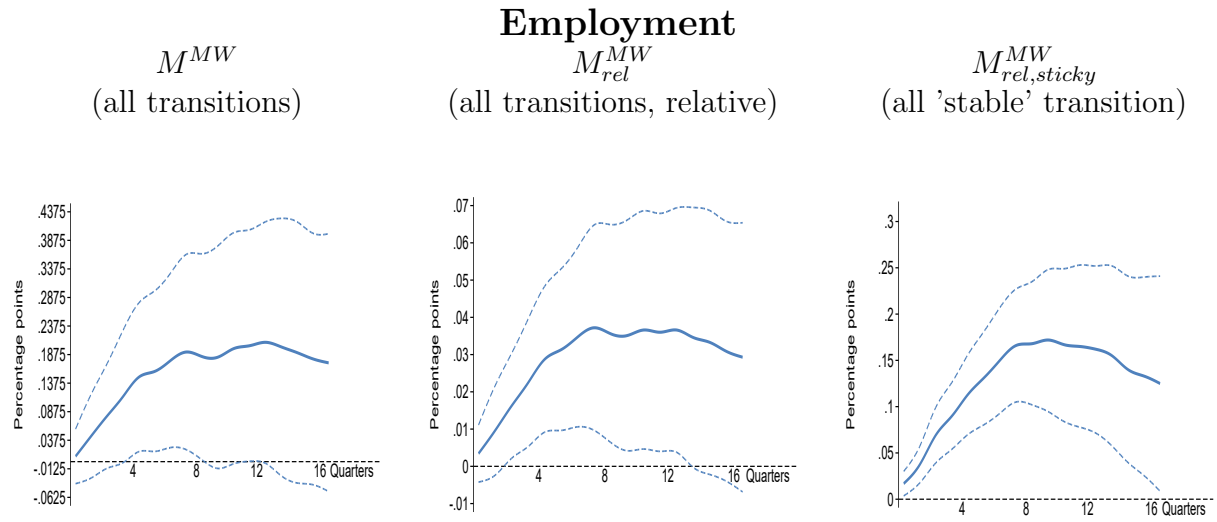

Each figure depicts the effects over time of a one percentage point increase in a marginality measure on the responsiveness of the employment to a monetary policy shock. Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

5 Individual-level responses

In this section we use the CPS IPUMS to look more closely at individual responses to a monetary policy shock and ask whether these responses differ depending on whether the respondent lives in a state that has a higher or lower fraction of married women in the work force. The individual level data allows us to compare the behavior of married women to other members of the workforce. It also allows us to focus on transitions between labor market states, as well as labor market participation.

We consider regressions similar to our baseline regression, but with outcomes at the individual level. We consider the following specification:

$$\begin{aligned} y_{i,s,t^m} &= \beta^k \cdot \varepsilon_{t^q-k} \cdot S^{MW}_{s,t^q-k-1} + \gamma^k \cdot S^{MW}_{s,t^q-k-1} \\ &+ \text{controls} + \alpha_{t^m} + \alpha_s + u_{i,s,t+k} \end{aligned} \tag{5}$$

Here the unit of observation is the individual i in state s during month t^m. The outcome variable y_{i,s,t^m} is an indicator variable. In some cases, it is an indicator of labor market status, such as whether the individual is employed, unemployed, or out of the labor force. In other cases, we take this variable to be an indicator of a transition between two labor market states, such as the transition from non-employment to employment. t^q is the quarter associated with month t^m, so that the coefficients β^k and γ^k capture the impact of monetary shocks k quarters in the past and the married worker share $k+1$ quarters in the

past.[17] We restrict our attention to individuals 25-50 years old in order to abstract from entry and retirement. As before the time period is 1990-2008. Note that (5) is not a panel regression, but a repeated cross-section. The sample changes each month. To control for changes in the sample over time, we include various socio-demographic controls for age, race and education.[18] Otherwise, this specification is similar to our baseline specification. We run this regression separately for each horizon k. Standard errors are clustered at the state level.

Figure 5: Effect of an Increase in the Share of Married Women Working in the State on the Individual's Response to a Monetary Policy Shock

Using the relative measure of S^{MW}. Coefficients multiplied by 100. Dashed lines represent 95% confidence intervals. Errors clustered at the state level.

[17] We lag S^{MW} so that the timing in this regression is similar to the timing in the baseline regression.

[18] The omitted category (setting all controls to zero) is equivalent to a person age 35, white, non-hispanic, with only a high-school diploma.

Figure 5 shows the results for two groups (married women and all others) and three outcome variables (whether the respondent is employed, in the labor force, or unemployed). The results are largely consistent with the baseline regressions. In states with a higher share of married women working, both married women and all others are less likely to be employed or in the labor force and more likely to be unemployed following a monetary tightening. The effects on employment of married women are stronger and less precisely estimated. One possible reason for the latter is that there are simply many fewer married women in the labor force.

Figure 6: Effect of an Increase in the Share of Married Women Working in the State on the Individual's Response to a Monetary Policy Shock

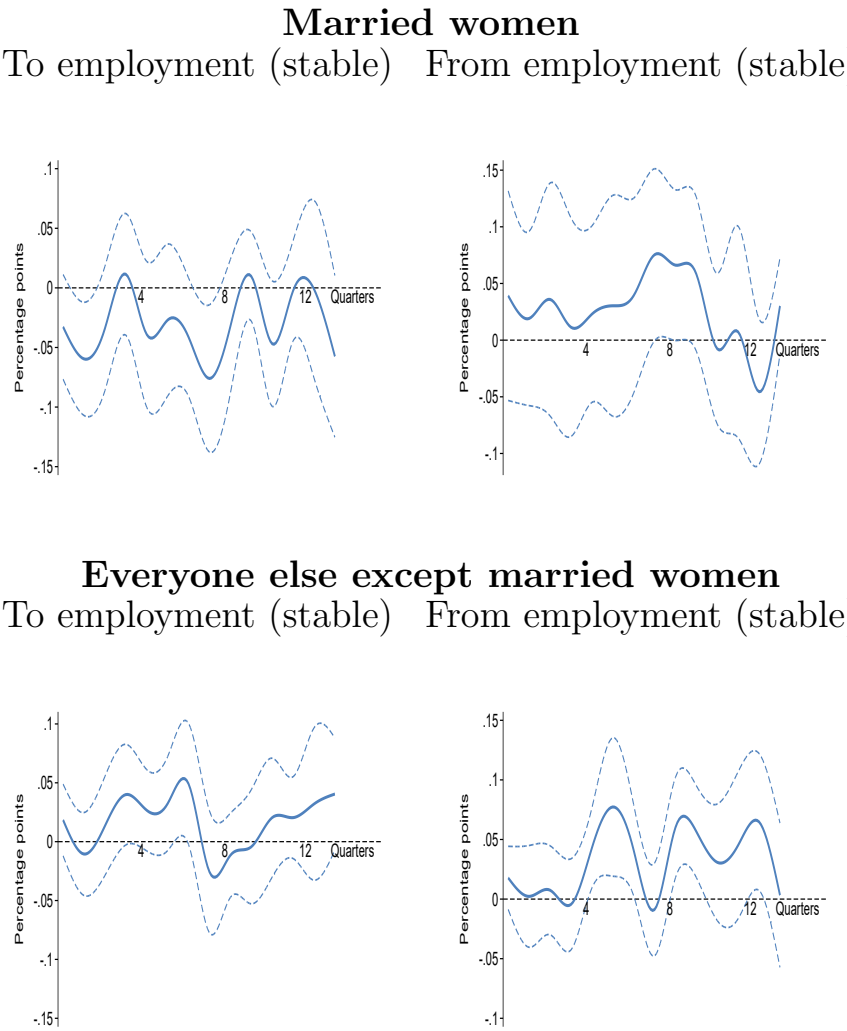

Using the relative measure of S^{MW}. Coefficients multiplied by 100. Errors clustered at the state level.

In Figure 6 we take the outcome variable y_{i,s,t^m} to be an indicator of whether an individual has changed employment status. As in Section 4 we use the panel dimension of the CPS to construct 'stable' measures of transitions from non-employment to employment

or from employment to non-employment where a person is observed for two months in one state and then for two months in the other. We construct these measures for both married women and for all others.

The responses in figure 6 are flatter than those in Figure 5. This is not surprising as Figure 6 graphs flows and whereas Figure 5 graphs stocks. One has to accumulate the flows into and out of employment to arrive at the stock of employment.

The main difference between the response of married women and all others is the response of flows into employment. In states in which more married women work, married women are less likely to enter employment following monetary tightening, whereas all others are more likely. Our interpretation of this difference is that it reflects the balance between two mechanisms. On the one hand, tight monetary policy has greater impact on the state, the more married women worked pre-shock. This is what we see in our baseline results. This greater impact appears to increase the flows into and out of employment. This is the effect that dominates in the case of all others. On the other hand, when married women do not work, they can use employment as an insurance mechanism for the family. This is the added worker effect. It appears to dominate in the sample of married women.

The flows out of unemployment in the figure are similar in magnitude. There is no obvious difference in the desire or ability to hold onto a job among the two groups.

Figure 7 takes y_{i,s,t^m} to be equal to the log of the respondent's wage, weekly earnings or hours. We restrict our attention to the sample of respondents that are working. The response of these variables to a monetary tightening for all others does not appear to differ significantly with the share of married women working. The response of these variables among married women is weaker the greater is the share of married women working. This weaker response of wages is consistent with the added worker effect. Fundamentally, the added worker effect is a labor supply response. Entry by married women increases labor supply in the markets in which they enter and reduce wages. As these marginal entrants are likely part time or low wage workers, earnings and hours also fall.

Overall, this individual behavior is consistent with an added worker effect that diminishes as married women become more attached to the labor market, as in Casella (2022). We see that married women are more likely to enter employment after a monetary tightening in states in which married women work less. We see that wages tend to fall more in these states. Both are consistent with the added worker effect being stronger in these states.

6 Conclusion

In this paper we provide empirical evidence that the added worker effect matters on the aggregate. As argued by Casella (2022), the added worker effect is a more potent insurance mechanism when married women are closer to the margin of working or not. In economies where married women are already firmly attached to the labor market, they are not close enough to the extensive margin of labor supply to be induced into countercyclical employment choices. As a result, this margin of insurance against shocks is shut down. Aggregate employment is more affected by shocks both because the direct countercyclical labor supply response of married women is missing and the lack of insurance motivates a stronger consumption (aggregate demand) response.

Figure 7: Effect of S^{MW} on the Individual's Response to a Monetary Policy Shock

Using the relative measure of S^{MW}. Coefficients multiplied by 100. Errors clustered at the state level. These are based on the 'earner study' sample.

We exploit cross-sectional variation in the response of US states to a monetary policy shock to provide evidence for this mechanism. We study how the impact of monetary policy varies with the share of married women who have worked prior to the shock. We find that the economy's response is more muted the lower the share of married women employed before the shock. We interpret this muted response as reflecting a stronger added worker effect. We provide three additional pieces of evidence consistent with this interpretation. First, married women are less likely to change labor market status in states in which a larger share of married women are employed. Moreover, monetary policy shocks have a stronger effect in such states. Second, following an increase in the federal funds rate married women themselves are comparatively more likely to be employed (and to enter employment) in states where the share of married women working is low. Third, in contrast to employment, wages of married women fall more in states where married women have worked less, consistent with a differential labor supply response to the shock.

References

Albanesi, S. (2019). Changing business cycles: The role of women's employment. *NBER Working Paper Series*, 1–64.

Barakchian, S. M. and C. Crowe (2013, 11). Monetary policy matters: Evidence from new shocks data. *Journal of Monetary Economics 60*, 950–966.

Bardoczy, B. (2022). Spousal insurance and the amplification of business cycles. *Working Paper*.

Bauer, M. D. and E. T. Swanson (2023a, 3). An alternative explanation for the "fed information effect". *American Economic Review 113*, 664–700.

Bauer, M. D. and E. T. Swanson (2023b, 1). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual 37*, 87–155.

Beraja, M., A. Fuster, E. Hurst, and J. Vavra (2019, 2). Regional heterogeneity and the refinancing channel of monetary policy*. *The Quarterly Journal of Economics 134*, 109–183.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016, 2). Consumption inequality and family labor supply. *American Economic Review 106*, 387–435.

Boone, C., A. Dube, L. Goodman, and E. Kaplan (2021, 5). Unemployment insurance generosity and aggregate employment. *American Economic Journal: Economic Policy 13*, 58–99.

Broer, T. ., J. V. Kramer, and K. Mitman (2025). The curious incidence of monetary policy across the income distribution. *Review of Economics and Statistics (Forthcoming)*.

Cajner, T., J. M. Coglianese, and J. Montes (2021, 7). The long-lived cyclicality of the labor force participation rate. *Finance and Economics Discussion Series 2021.0*, 1–48.

Carlino, G. and R. DeFina (1998, 11). The differential regional effects of monetary policy. *The Review of Economics and Statistics 80*, 572–587.

Casella, S. (2022). Women's labor force participation and the business cycle. *Working Paper*.

Cieslak, A. (2018, 9). Short-rate expectations and unexpected returns in treasury bonds. *The Review of Financial Studies 31*, 3265–3306.

Cravino, J., T. Lan, and A. A. Levchenko (2020, 4). Price stickiness along the income distribution and the effects of monetary policy. *Journal of Monetary Economics 110*, 19–32.

Doepke, M. and M. Tertilt (2016). *Families in Macroeconomics* (1 ed.), Volume 2. Elsevier B.V.

Ellieroth, K. (2019). Spousal insurance, precautionary labor supply, and the business cycle. *Working Paper*.

Faia, E., E. Shabalina, and D. Wiczer (2024). Heterogeneous effects of monetary policy across income and race: the labour mobility channel. *AEA Papers and Proceedings*.

Flood, S., M. King, R. Rodgers, S. Ruggles, R. J. Warren, D. Backman, A. Chen, G. Cooper, S. Richards, M. Schouweiler, and M. Westberry (2024). Ipums cps: Version 12.0 [dataset]. minneapolis, mn: Ipums.

Fukui, M., E. Nakamura, and J. Steinsson (2023). Women, wealth effects, and slow recoveries. *American Economic Journal: Macroeconomics 15*, 269–313.

Gertler, M. and P. Karadi (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics 7*, 44–76.

Graves, S., C. K. Huckfeldt, and E. T. Swanson (2023, 10). The labor demand and labor supply channels of monetary policy. *NBER Working Papers*.

Guner, N., Y. Kulikova, and A. Valladares-Esteban (2025). Does the added worker effect matter? *Review of Economic Dynamics*.

Gurkaynak, R. S., B. P. Sack, and E. T. Swanson (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *SSRN Electronic Journal*.

Imam, P. A. (2015, 3). Shock from graying: Is the demographic shift weakening monetary policy effectiveness. *International Journal of Finance and Economics 20*, 138–154.

Jarocinski, M. and P. Karadi (2020, 4). Deconstructing monetary policy surprises - the role of information shocks. *American Economic Journal: Macroeconomics 12*, 1–43.

Jorda, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*.

Juhn, C. and S. Potter (2007). Is there still an added worker effect? *NBER Working Paper*.

Kuttner, K. N. (2001, 6). Monetary policy surprises and interest rates: Evidence from the fed funds futures market. *Journal of Monetary Economics 47*, 523–544.

Leahy, J. V. and A. Thapar (2022). Age structure and the impact of monetary policy. *American Economic Journal: Macroeconomics*.

Lundberg, S. (1985, 1). The added worker effect. *Journal of Labor Economics 3*, 11–37.

Maggio, M. D., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017, 11). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review 107*, 3550–88.

Mankart, J. and R. Oikonomou (2016, 6). The rise of the added worker effect. *Economics Letters 143*, 48–51.

Mankart, J. and R. Oikonomou (2017, 10). Household search and the aggregate labour market. *Review of Economic Studies 84*, 1735–1788.

Mankart, J., R. Oikonomou, and F. Pascucci (2022). The rise of household insurance. *Working Paper*.

Mennuni, A. (2019, 7). The aggregate implications of changes in the labour force composition. *European Economic Review 116*, 83–106.

Miranda-Agrippino, S. and G. Ricco (2021, 7). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics 13*, 74–107.

Nakamura, E. and J. Steinsson (2018, 8). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics 133*, 1283–1330.

Olsson, J. (2025, 1). Singles, couples, and their labor supply: Long-run trends and short-run fluctuations. *American Economic Journal: Macroeconomics 17*, 1–34.

Ottonello, P. and T. Winberry (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica 88*, 2473–2502.

Parker, S. W. and E. Skoufias (2004, 8). The added worker effect over the business cycle: evidence from urban mexico. *Applied Economics Letters 11*, 625–630.

Ramey, V. A. (2016, 1). Macroeconomic shocks and their propagation. *Handbook of Macroeconomics 2*, 71–162.

Romer, C. D. and D. H. Romer (2004, 8). A new measure of monetary shocks: Derivation and implications. *American Economic Review 94*, 1055–1084.

Rudebusch, G. D. (1998, 11). Do measures of monetary policy in a var make sense? *International Economic Review 39*, 907.

Söderström, U. (2001, 4). Predicting monetary policy with federal funds futures prices. *Journal of Futures Markets 21*, 377–391.

Stephens, M. (2002, 7). Worker displacement and the added worker effect. *Journal of Labor Economics 20*, 504–537.

Thapar, A. (2008, 5). Using private forecasts to estimate the effects of monetary policy. *Journal of Monetary Economics 55*, 806–824.

Wang, H. (2019, 10). Intra-household risk sharing and job search over the business cycle. *Review of Economic Dynamics 34*, 165–182.

White, N., S. Bhattarai, A. Glover, N. Petrosky-Nadeau, H. Afrouzi, G. Hagemann, J. Ryngaert, and D. Koustas (2018). Gross worker flows, job loss, and monetary policy would like to thank olivier coibion. *Working Paper*.

Woytinsky, W. S. (1942). *Three aspects of labor dynamics*. Washington D.C: Social Science Research Council.

Wu, C. and D. Krueger (2021, 1). Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation. *American Economic Journal: Macroeconomics 13*, 79–113.

Table 2: Effect of an Increase in the Relative Share of Married Women Working on the Response to a MP shock

Lag	Employment				Payroll			
	b	se	p	N	b	se	p	N
0	-0.0738	0.0462	0.1150	2274	-0.3330	0.1090	0.0031	2324
1	-0.1800	0.0826	0.0324	2274	-0.5040	0.1440	0.0008	2324
2	-0.3160	0.1090	0.0048	2274	-0.7160	0.1750	0.0001	2324
3	-0.4850	0.1410	0.0010	2274	-1.0300	0.2340	0.0000	2324
4	-0.6230	0.1820	0.0010	2274	-1.1000	0.2680	0.0001	2324
5	-0.7430	0.2330	0.0021	2274	-1.2900	0.3290	0.0002	2324
6	-0.8270	0.2750	0.0036	2274	-1.4800	0.3610	0.0001	2324
7	-0.9030	0.3060	0.0043	2274	-1.5600	0.4190	0.0004	2324
8	-0.9160	0.3320	0.0073	2274	-1.6400	0.4050	0.0001	2324
9	-0.9310	0.3610	0.0119	2274	-1.7100	0.4850	0.0007	2324
10	-0.9350	0.3680	0.0132	2274	-1.8200	0.4930	0.0004	2324
11	-0.9290	0.3860	0.0185	2274	-1.7600	0.4910	0.0006	2324
12	-0.9250	0.4010	0.0239	2274	-1.6000	0.5270	0.0034	2324

Based on Driscoll-Kraay standard errors.

A Appendix

A.1 Baseline Coefficient Estimates

Table 2 presents the coefficient estimates behind Figure 1.

A.2 Potentially Confounding Variables

We first check that our main conclusion is robust to only using within-region variation across states, by running a regression of the form

$$y_{s,t,k} = \beta^k \cdot \epsilon_t \cdot S^{MW}_{s,t-1} + \gamma^k \cdot S^{MW}_{s,t-1} + \sum_r \theta^k_r \cdot \epsilon_t \cdot I_{s:r} + \alpha_t + \alpha_s + u_{s,t,k}$$

where $y_{s,t+l}$ stands for a cumulative change in log of employment between $t-1$ and $t+k$, ϵ_t is a monetary policy shock (with higher values representing a tighter monetary policy, i.e. a negative shock). $S^{MW}_{s,t-1}$ is a measure of how much married women already work relative to the whole population (using population 16 and above). $I_{s:r}$ is an indicator that takes the value 1 if state s is in BEA region r.

As in the baseline analysis, β^k is the vector of coefficients of interest: if $\beta^k < 0$, states in which married women are more firmly attached to the labor market compared to the rest of their region respond more severely to a monetary policy shock. γ^k_r measures the

Figure 8: Effect of an Increase in the Share of Married Women Working (relative to the average) on the Economy's Response to a Monetary Policy Shock

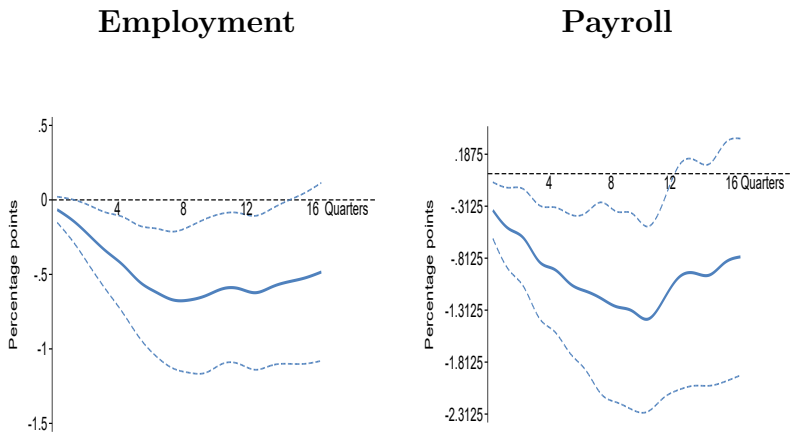

Each figure depicts the effects over time of a one percentage point increase in S^{MW} on the responsiveness of the relevant variable (state personal employment or payroll) to a monetary policy shock, including region fixed effects interacted with the shock. Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

differential sensitivity of different regions.[19] The results (depicted in figure 8) are robust, with the measured effects being slightly smaller and more noisy.

Next, we investigate the role of potential confounders–other state-level descriptors that might correlate with the share of married women working that might affect the responsiveness of the state to monetary policy. Table 3 presents correlations of the share of married women working with other measures of labor force participation.

[19]The regions are Far West (excluded category), Great Lakes, Mideast, New England, Plains, Rocky Mountains, Southeast and Southwest.

Table 3: Correlation between share of married women working and other state-level variables

		S^{MW}	S^{MW}_{abs}
Relative share of married women working, 16+	S^{MW}	1.00	
Share of married women 16+ working	S^{MW}_{abs}	0.62	1.00
Share of women 16+ working		0.33	0.93
Share of men 16+ working		-0.14	0.66
Share of 16+ working		0.12	0.85
LFP of married women 16+		0.62	0.99
LFP of women 16+		0.33	0.92
LFP of men 16+		-0.18	0.60
LFP of 16+		0.10	0.83
Share of population 25-40		-0.33	-0.15
Share of population 40-65		0.41	0.20
Share of population 65+		0.34	-0.08
Personal income per capita		0.24	0.27
Manufacturing share of emplpoyment		0.01	-0.05
Construction share of emplpoyment		-0.33	-0.12
Service share of emplpoyment		0.21	0.10
Financial services share		-0.11	0.07
Average max UI weekly benefit		0.20	0.12
Max # weeks of UI benefits		-0.02	0.00
Max UI weeks times amount		0.19	0.12

Now, we run regressions of the form These regressions take the form:

$$\Delta \ln y_{s,t,k} = \beta^k \cdot \epsilon_t \cdot S^{MW}_{s,t-1} + \gamma^k \cdot S^{MW}_{s,t-1} + \theta^k \cdot \epsilon_t \cdot X_{s,t-1} + \delta^k \cdot X_{s,t-1} + \alpha_t + \alpha_s + u_{s,t+k}. \tag{6}$$

where $y_{s,t,k}$ stands for a cumulative change in log of employment between $t-1$ and $t+k$, ϵ_t is a monetary policy shock (with higher values representing a tighter monetary policy, i.e. a negative shock). $S^{MW}_{s,t-1}$ is a measure of how much married women already work relative to the whole population (using population 16 and above). $X_{s,t-1}$ is the confounding variable which captures some state-level characteristic that might be influencing the impact of monetary policy.

As in the baseline analysis, β^k is the vector of coefficients of interest: if $\beta^k < 0$, states in which married women are more firmly attached to the labor market respond more severely to a monetary policy shock. Our focus is on how the inclusion of the confounding variables $X_{s,t-1}$ impact our baseline estimates of the β^k. Note that we include the $X_{s,t-1}$ both in levels and interacted with the monetary policy shock ϵ_t.

Figure 11 shows that using the absolute measure of the share of married women working in the baseline regression leads to the same conclusion.

Figure 9: Comparing baseline estimates with including $X_{s,t-1}$ and $\epsilon_t X_{s,t-1}$ controls

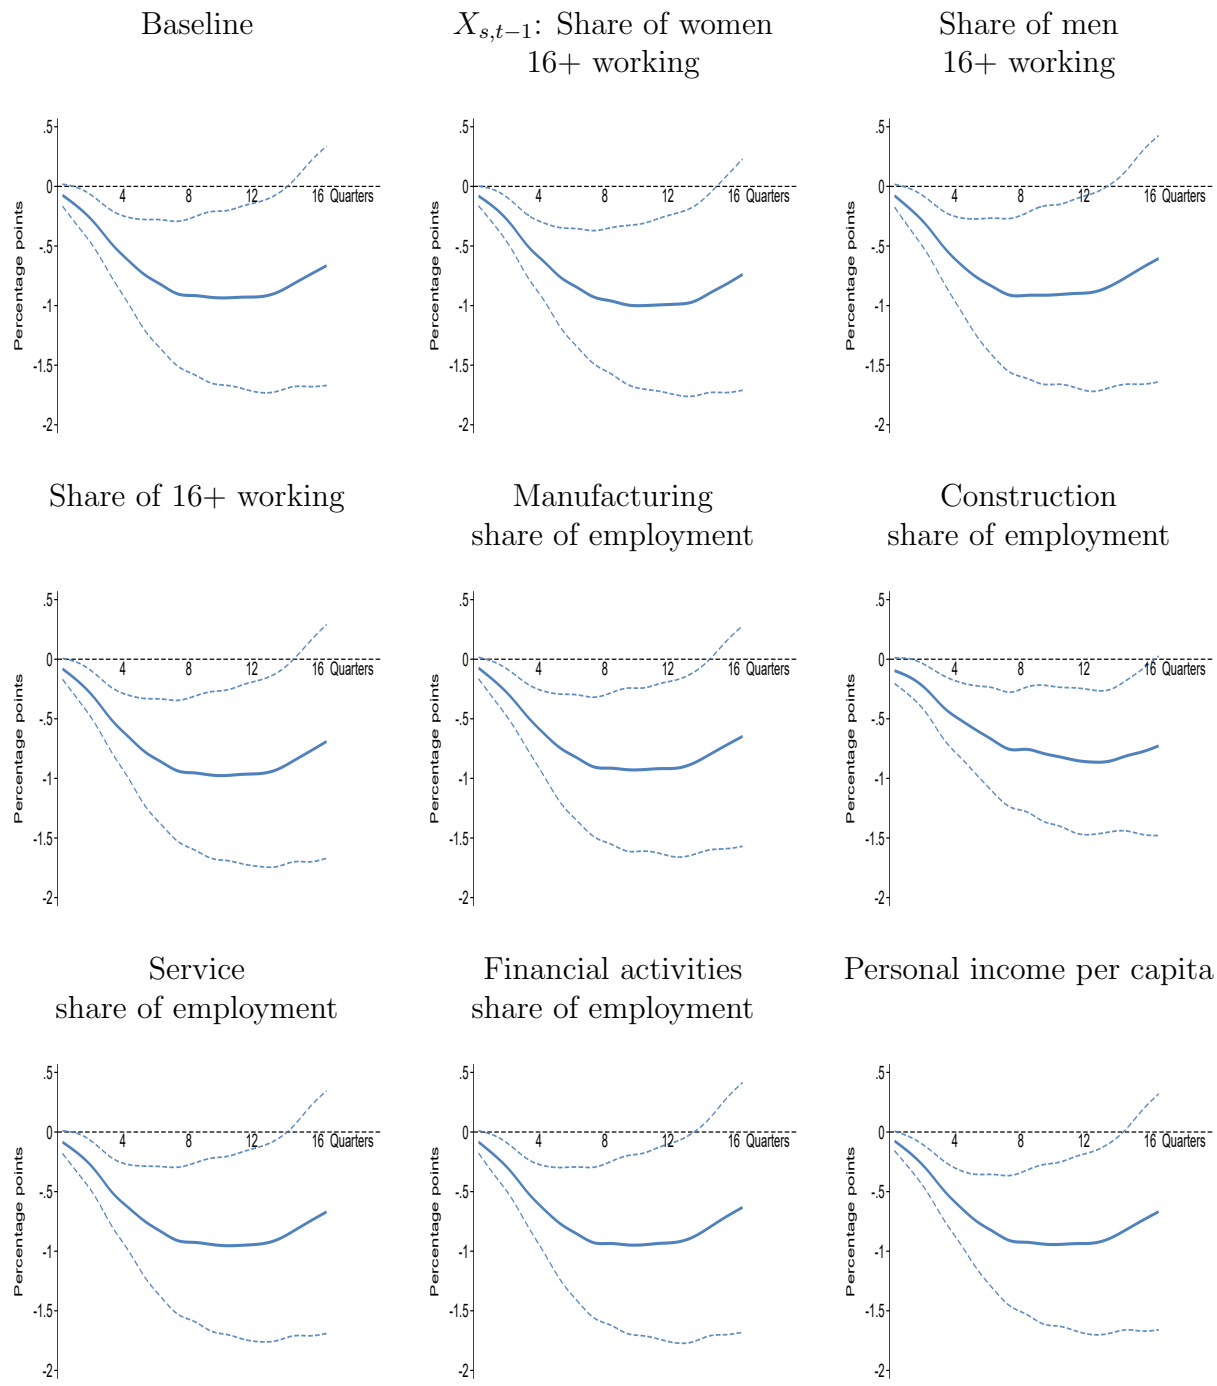

$X_{s,t-1}$: Share of 16+ working is based on CPS IPUMS with quantities smoothed as a 4 quarter moving average. Industry shares and personal income per capital taken from Leahy and Thapar (2022).

Figure 10: Comparing baseline estimates with including $X_{s,t-1}$ and $\epsilon_{t-1}X_{s,t-1}$ controls

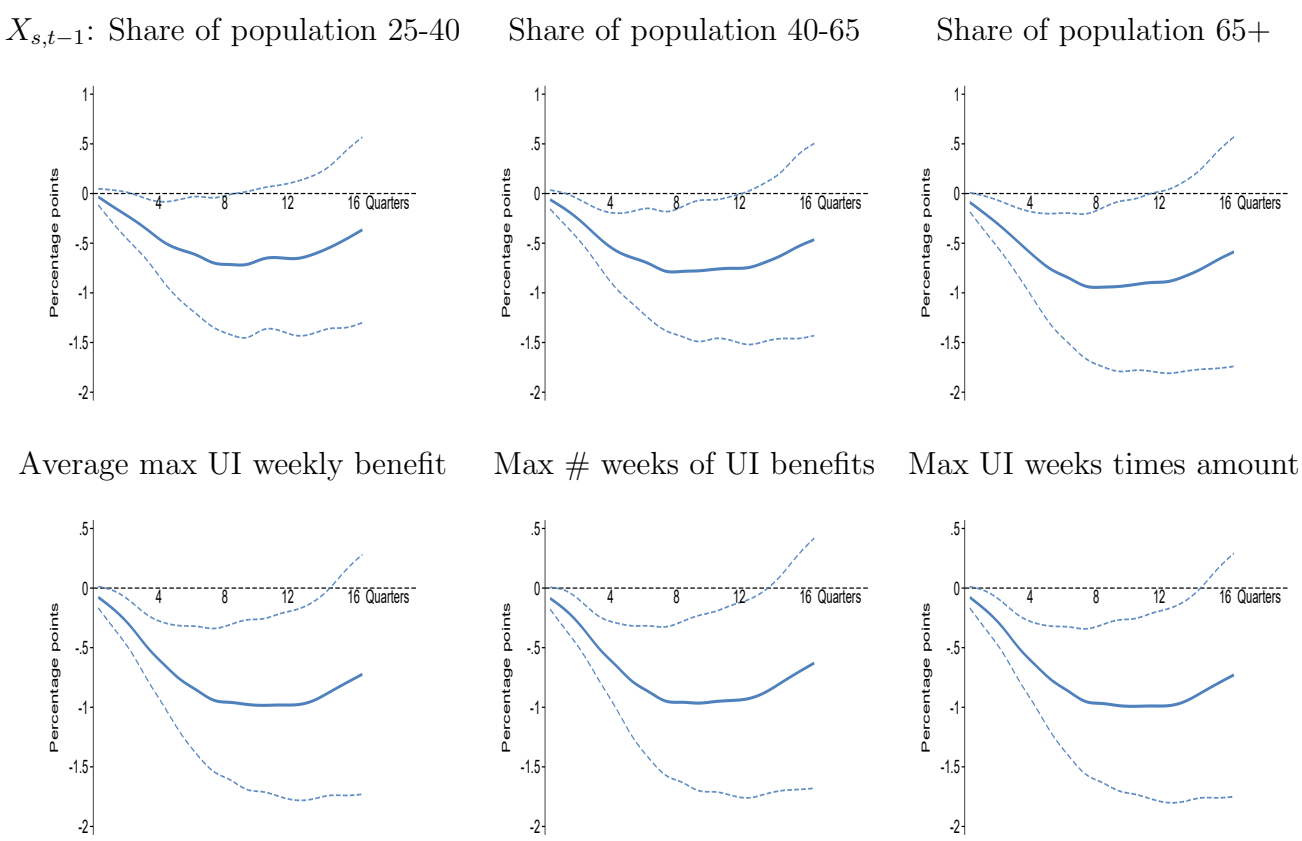

$X_{s,t-1}$: Share of population by age based on CPS IPUMS, quantities smoothed as a 4 quarter moving average. UI generosity variables taken from Boone et al. (2021).

Figure 11: Comparing the relative and absolute measures of married woman's participation

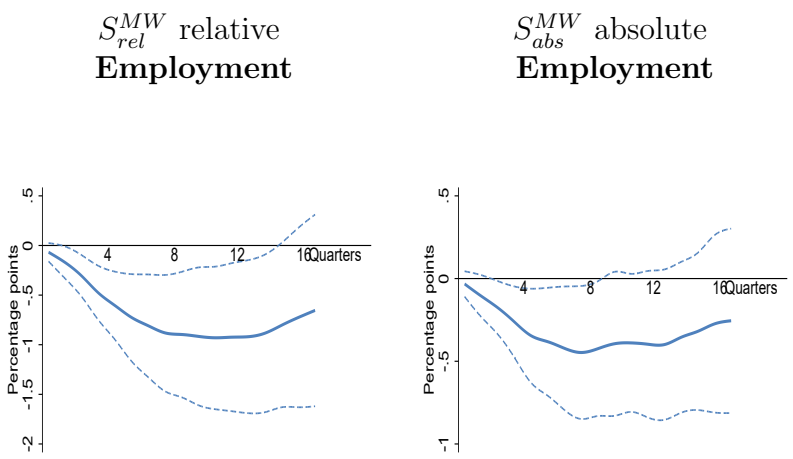

Each figure depicts the effects over time of a one percentage point increase in S^{MW} on the responsiveness of employment to a monetary policy shock. Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

A.3 Other robustness tests

Here we discuss other robustness exercises. In figure 12, we check that our baseline results are not driven by small states only. We rank states by the average of total employment within our sample period. In the first figure, we drop the 6 smallest states (Alaska, North and South Dakota, Montana, Vermont and Wyoming). In the second figure, states in the panel analysis are weighted by the average total employment over the sample period.

Figure 12: Effect of an Increase in the Share of Married Women Working (relative to the average) on the Economy's Response to a Monetary Policy Shock

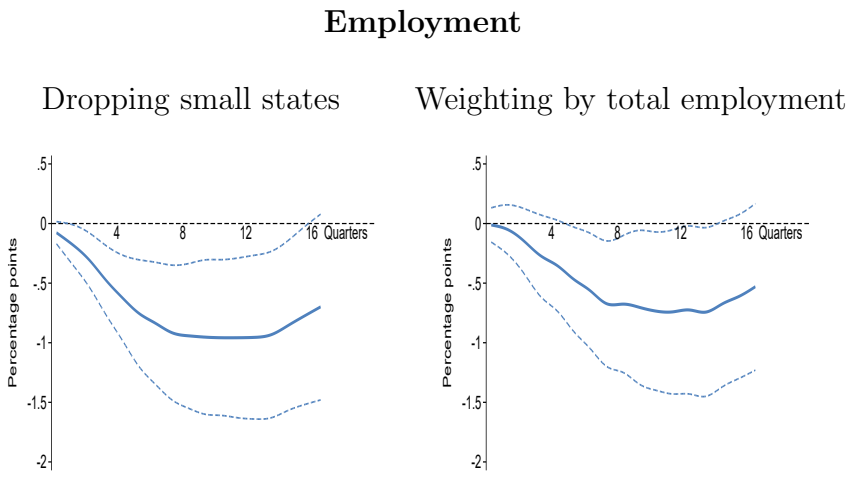

Each figure depicts the effects over time of a one percentage point increase in S^{MW} on the responsiveness of employment to a monetary policy shock. Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

Next, we discuss robustness to using alternative monetary policy shocks.

Figure 13: Comparison with using the Romer and Romer (2004) narrative shocks.

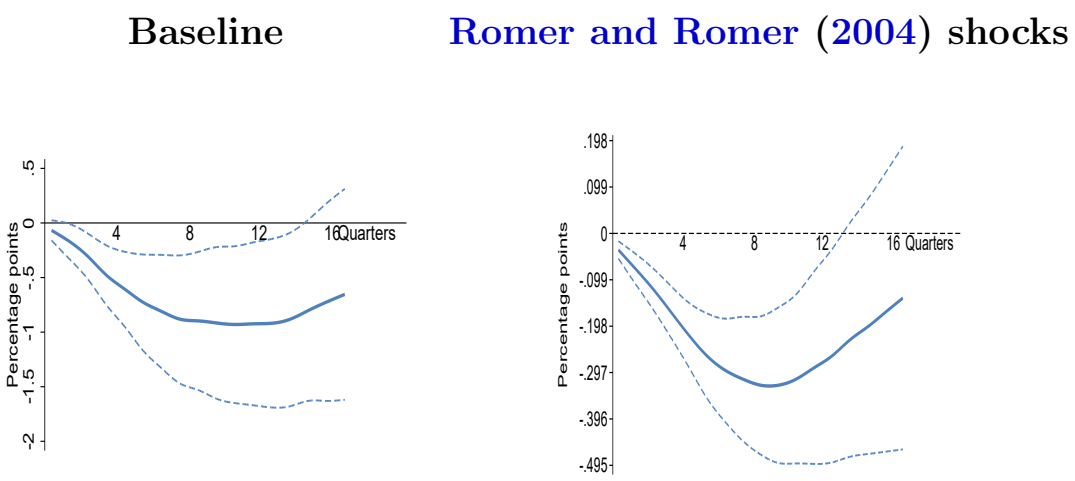

Romer and Romer (2004) shock series, extended by Johannes Wieland and Max Breitenlechner

The strategy of using high-frequency surprises as monetary policy shocks has come under some scrutiny since its inception. Several papers have documented that even this very targeted measurement method leads to a series of shocks that likely measure more than the change in monetary policy orthogonal to the current economic conditions. Specifically,

Figure 14: Comparison between baseline and Miranda-Agrippino and Ricco (2021) shocks.

the shocks can still be predictable and autocorrelated (Ramey (2016), Cieslak (2018)) and can correlate with central banks? private macroeconomic forecasts (Barakchian and Crowe (2013), Gertler and Karadi (2015), Ramey (2016)). This can be either because the Fed is revealing some of its private information through their policy action (as in Nakamura and Steinsson (2018)) or that the Fed is reacting with a surprising intensity compared to what the markets expect (as argued by Bauer and Swanson (2023a) and labeled the 'Fed response to news' channel). To deal with these potential issues several authors have proposed strategies to orthogonalize the series of high-frequency surprises in asset prices around FOMC meetings with respect to the information effect. Jarocinski and Karadi (2020)'s approach uses the information-processing power of the markets and identifies central bank information shocks from the high-frequency co-movement of interest rate and stock market surprises (based on the intuition that if a monetary policy tightening also has a signalling content about an improving economy, stock prices would increase, despite the contractionary effects of the tightening itself). Bauer and Swanson (2023b) project the standard high-frequency surprises in the federal funds rate on publicly available macroeconomic and financial news. Miranda-Agrippino and Ricco (2021) orthogonalize the surprises with respect to central bank?s economic projections as well as with past market surprises (purging the serial correlation as well).

Figure 15 shows the comparison between the baseline result with using the monetary policy shocks series proposed by Miranda-Agrippino and Ricco (2021). The first line shows that using this alternative series of shocks does not lead to the same conclusion. However, the second line reveals that the descrepency is driven almost entirely by how the methods interpret the monetary policy actions of the Fed during the Great Recession. Specifically, if we replace the value of the shock for both our baseline series and the Miranda-Agrippino and Ricco (2021) series during the quarters of the great recession with its pre-great-recession mean (completely neutralizing the shocks during this team), the results are very comparable across the two series. This conclusion is the same when comparing to other shocks aiming to correct for the 'information effect', namely Bauer

and Swanson (2023a) and Jarocinski and Karadi (2020). Figure 14 shows why, plotting the baseline, Miranda-Agrippino and Ricco (2021) and Bauer and Swanson (2023a) series. While the high-frequency surprises interpret the cutting of interest rates during the Great Recession as an easing of monetary policy, the other series 'overcorrect' and interpret it as tightening. When we treat these actions as neutral, variation from the rest of the sample supports our baseline conclusion. Figure 13 shows that our result is also robust to using the narrative approach by Romer and Romer (2004) to identifying monetary policy shocks.

Figure 15: Comparison with using the Miranda-Agrippino and Ricco (2021) shocks.

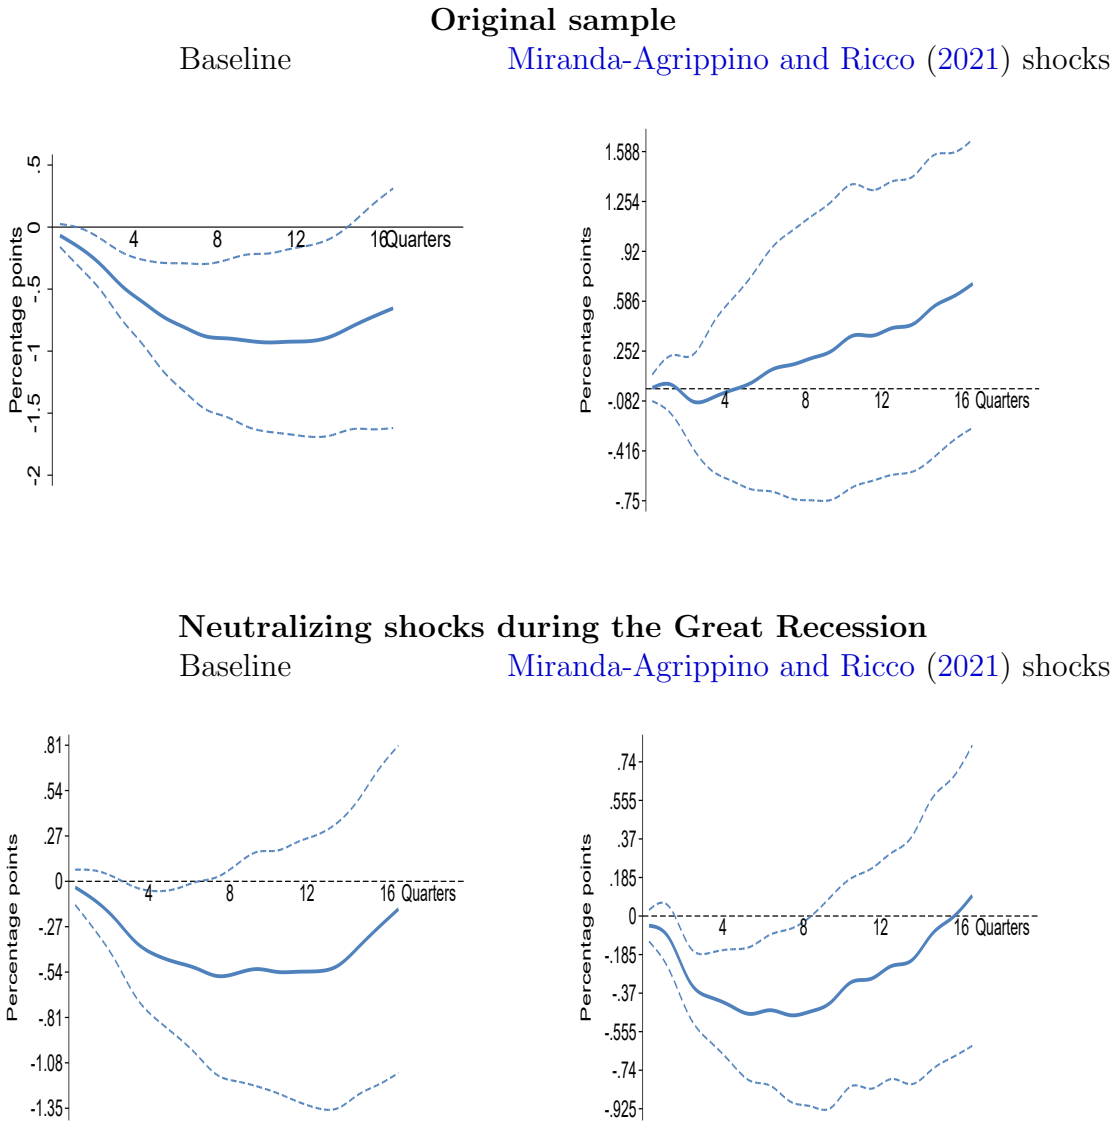

Comparing the baseline results with using the monetary policy shocks series proposed by Miranda-Agrippino and Ricco (2021). In the second line, we replaced the values of the shock in quarters 2007q4-2008q4 with its sample mean in the period before the Great Recession (neutralizing the values of the shocks during the Great Recession).

A.4 Robustness checks for the marginality measures

In this section we include robustness exercises regarding the marginality measures. Figure 16 shows that measuring marginality of married women with the restricted sample of those who are 20-50 years old yields the same conclusion as the baseline marginality measures. This shows that the differential response of employment to monetary shocks is not driven by moves related to retirement.

Figure 16: Effect of an Increase in the 'Marginality' of Married Women 25-50 on the Economy's Response to a Monetary Policy Shock

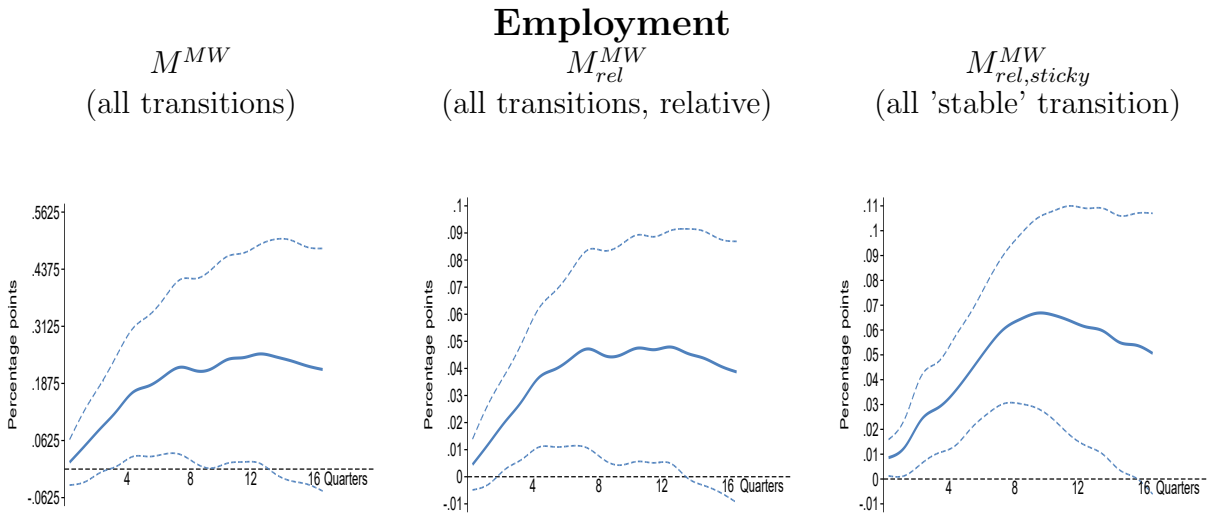

Marginality measures defined on the population 25-50.

Figures 17 and 18 show a horse-race between the share of married women working and marginality measures. Overall, while the results are consistent with the idea that both are correlated proxies for the same mechanism, the share of married women working show a more robust effect. The effect of M^{MW}_{rel} disappears when share of married women working is included. $M^{MW}_{rel,sticky}$ survives but diminishes in magnitude.

Figure 17: Horse-race between share of married women working and marginality measure

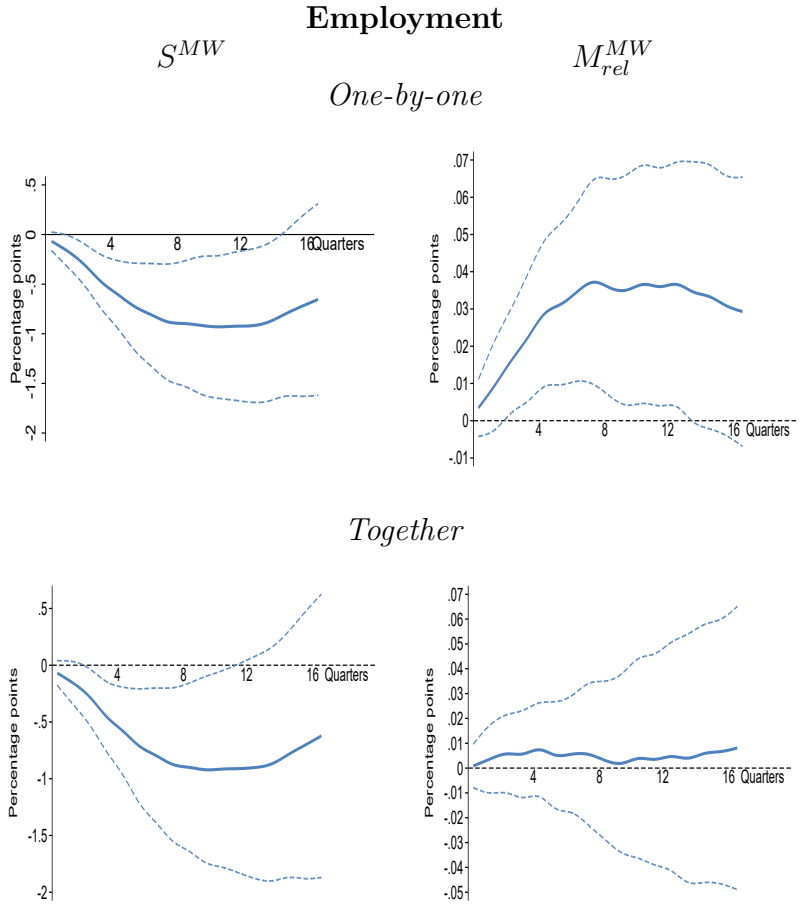

Each figure depicts the effects over time of a one percentage point increase in S^{MW} and M^{MW}_{rel} on the responsiveness of employment to a monetary policy shock. The first row includes one per regression, as in the paper. The second includes both in a horse-race. Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

Figure 18: Horse-race between share of married women working and marginality measure

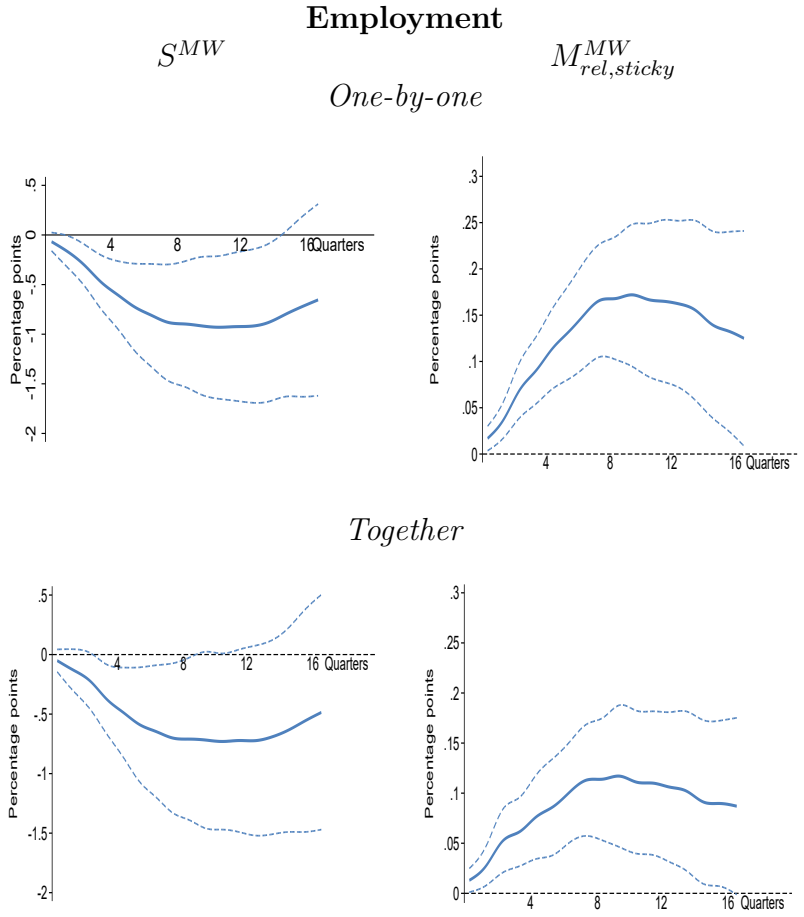

Each figure depicts the effects over time of a one percentage point increase in S^{MW} and $M^{MW}_{rel,sticky}$ on the responsiveness of employment to a monetary policy shock, as in the paper. The first row includes one per regression. The second includes both in a horse-race. Solid lines represent point estimates. Dashed lines represent 95 percent confidence intervals based on Driscoll-Kraay standard errors.

Figure 19: Effect of an Increase in the 'Marginality' of Married Women 25-50 on the Economy's Response to a Monetary Policy Shock

Employment

M_{rel}^{MW} M_{rel}^{MW}

All transitions de-NUN-ified Transitions to and from the labor force

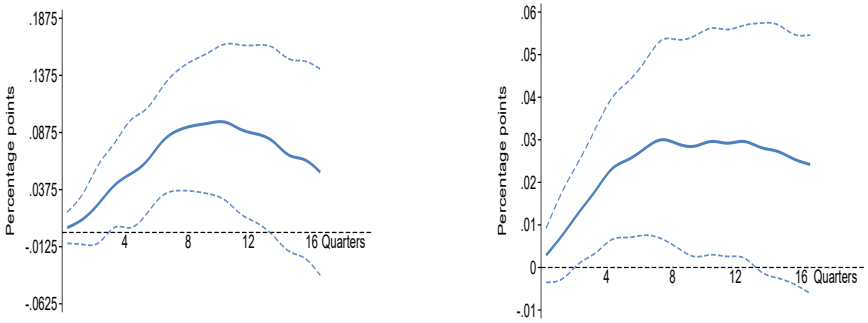

Alternative marginality measures defined on the population 25-50.

A.5 Supplemental results

In the analysis if transitions to and from employment we choose to report changes in the probability of making a 'stable' transition. Since individuals in the CPS are interviewed for four months in a row, we construct a measure of a 'stable' transitions from non-employment to employment as equal to 1 if a person is observed for two months non-employed and then for two months employed. It turns out most transitions in the CPS microdata are incredibly short lived: only 39.8% of all raw transitions to employment in our sample are 'stable' by our definition (and only 30.3% of transitions to labor force are 'stable'). Given two months is itself a very short period of time, we find it unlikely that many of the raw transitions can be interpreted as true changes in an optimized labor supply decision.

Rather they probably reflect being employed in a very temporary jobs or classification errors. The CPS documentation defines being employed as follows: "In the CPS, individuals' employment status was determined on the basis of answers to a series of questions relating to their activities during the preceding week. Those who reported doing any work at all for pay or profit, or working at least fifteen hours without pay in a family business or farm, were classified as "at work." Those who did not work during the previous week but who acknowledged having a job or business from which they were temporarily absent (e.g., due to illness, vacation, bad weather, or labor dispute) were also classified as employed, under the heading "has job, not at work last week. ... Individuals were coded as unemployed if they ... (were not employed) and either reported looking for work as their major activity during the previous week (for 1962 through 1993) or answered yes to a question about whether they had been looking for work in the past four weeks." One can easily imagine someone who did not work a week before the survey one month but did in the next month, and they are not formally employed, because the nature of their work

does not require a regular schedule. This person would be classified as out-of-labor-force one month and employed the next, even though nothing has fundamentally changed for them.

Figure 20 shows the results of comparing the responses of stable transitions with the raw measure (of counting all changes from non-employment to employment, or from employment to non-employment). Using the raw measures is much noisier (see the difference in the axis scale). Overall, the result that in more equal states married women are less likely to enter and more likely to exit in response to a monetary tightening is not robust to using the raw measure. For the rest of the population, figure 21 shows that the difference is less striking. In more equal states, there are more transition both to and from employment after a monetary tightening.

Figure 20: Effect of an Increase in the Share of Married Women Working in the State on the Individual's Response to a Monetary Policy Shock

Using the relative measure of S^{MW}. Coefficients multiplied by 100. Errors clustered at the state level.

For completion, we also repeat the important individual level results by first aggregating outcomes by state and running them as a state-quarter-level panel, similar to the baseline analysis. Specifically we follow Cajner, Coglianese, and Montes (2021), to first adjust the inidivual level outcome for socio-demographic factors, then aggregate them to counts per state per quarter, seasonally adjust them, and compute aggregate shares. Bellow we do this for being employed (the share of people working), doing a stable transition to employment (the share of people who were out for at least two months entering employment for at least two months) and doing a stable transition from employment. We then run the main specification 1 replacing the left hand side with these shares to mimick as much as possible the individual-level regressions in 5. The results are presented in figure 22 and mimick the individual-level conclusions from section 5.

Figure 21: Effect of an Increase in the Share of Married Women Working in the State on the Individual's Response to a Monetary Policy Shock

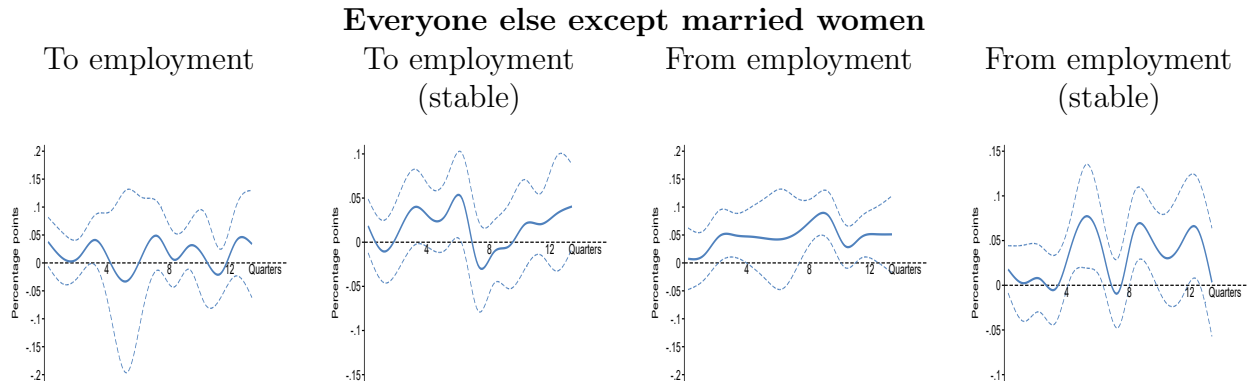

Using the relative measure of S^{MW}. Coefficients multiplied by 100. Errors clustered at the state level.

Figure 22: Effect of an Increase in the Share of Married Women Working in the State on the Aggregated Response to a Monetary Policy Shock

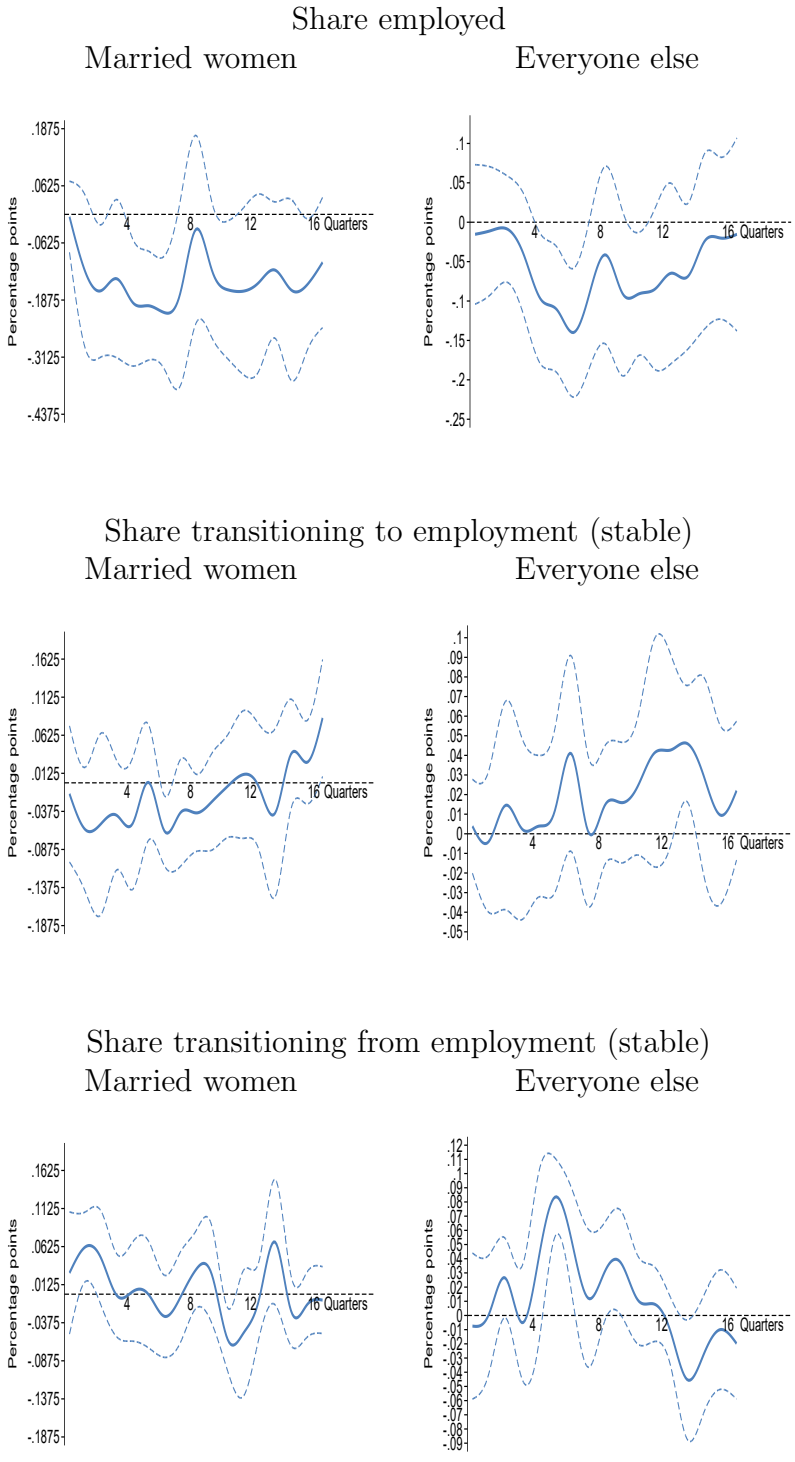

Using the relative measure of S_{rel}^{MW}. Coefficients multiplied by 100. Dashed lines represent 95% confidence intervals based on Driscoll-Kraay standard errors.