What Does Monetary Policy Do To Different People?

Pooyan Amir-Ahmadi, University of Illinois
Christian Matthes, Indiana University
Mu-Chun Wang, Deutsche Bundesbank *

September 3, 2021

Abstract

Does monetary policy affect people differently depending on their education level, their marital status, or their gender? To study this question, we use a Vector Autoregression where monetary policy effects are identified via an instrument to study how labor market outcomes differ across these groups after a monetary policy shock.
The response of the aggregate unemployment rate to a monetary policy shocks masks *massive* heterogeneity. We find that the magnitude of the response of the *difference* of unemployment rates across groups is often between 50 percent and 100 percent of the peak response of the *level* of the aggregate unemployment rate.

JEL Classification: E24, E50

Key Words: Monetary Policy, Heterogeneity

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Deutsche Bundesbank, the Eurosystem or their staff.

" In fact, the distributional effects of monetary policy are complex and uncertain."
— Ben Bernanke (Bernanke (2015))

1 Introduction

One of the central questions in macroeconomics is "What does monetary policy do?"[1] A common approach to answer this question for *aggregate* variables is to model the dynamics of a vector of variables via a vector autoregression (VAR) and then invoke an identifying assumption that maps forecast errors into estimates of structural shocks (see for example the work surveyed in Christiano et al. (1999)).

More recently, the economics profession has studied the effects of monetary policy changes on individuals. A key innovation that allows economists to quantitatively study these effects are heterogeneous-agent versions of equilibrium models that have long been used to analyze monetary policy (Kaplan et al., 2018; Gornemann et al., 2016; Auclert, 2019). While these equilibrium models feature rich heterogeneity, this heterogeneity is usually summarized by few state variables at the household level (e.g. various asset positions or the current income level). Empirically, Holm et al. (2021) use Norwegian micro data to empirically asses the impact of monetary policy shocks on individuals along the wealth distribution, while Andersen et al. (2020) use Danish micro data to study the effects of monetary policy on households along the income distribution. Both of those papers thus focus on measures of heterogeneity common in the aforementioned equilibrium models. In particular, these models, while tremendously useful, do not yet contain heterogeneity along many dimensions economists find useful to study, for example education, gender, and marital status, to name just the three dimensions we will focus on this paper. We study these specific groups to get a better sense of who is most exposed to monetary policy shocks among groups that are often featured in policy discussions.[2]

We go back to the VAR approach to monetary policy, but instead of focusing on aggregate variables alone, we use U.S. micro data to augment aggregate VARs with data on labor market outcomes of the socio-economic groups we are interested in. Our main focus in this paper is on the unemployment rate, but our VARs use both measures of intensive-margin adjustment in the labor market (hours) as well as changes in the extensive margin (the

[1]This is also the title of Leeper et al. (1996), which inspired our own choice of title.

[2]Our choices for slicing the micro data result in overlapping groups - individuals appear in different groups in our different exercises.

unemployment rate and labor-force participation).

To identify monetary policy shocks, we exploit variation in Federal Funds futures around Federal Open Market Committee (FOMC) meeting dates - we use the same monetary policy instrument as in Gertler and Karadi (2015). Our approach to incorporating instruments in VARs is borrowed from Mertens and Ravn (2013).

The issue of heterogeneous effects of monetary policy more broadly has been touched upon in various studies - besides the aforementioned studies, Bartscher et al. (2021) study the effects of monetary policy on the black/white unemployment gap, while Doepke and Schneider (2006) study the effects of inflation, partly controlled by monetary policy, on wealth inequality. Coibion (2012) use VARs to study the effects on inequality just like our paper, but these authors focus on broad summary measures of income and consumption inequality, whereas our focus is on differences in labor market outcomes across different socio-economic groups. Similar in focus to Coibion (2012), Lenza and Slacalek (2018) analyze the effects of monetary policy (and quantitative easing in particular) on wealth and income inequality in the Euro area.

The next section discusses our data. We then briefly introduce our VAR models before we show results from our procedure for aggregate and aggregated data as a benchmark. Section 5 contains our main results on the effects of monetary policy shocks across different socio-economic groups.

2 Data

Our VARs combine aggregate macro data with labor market data for various socio-economic groups that we aggregate from micro data. To construct these aggregated micro data, we use uniform extracts from the Current Population Survey (CPS) outgoing rotation group provided by the Center for Economic and Policy Research (`ceprdata.org`). We use micro data on hours, labor force participation, and unemployment status. Our sample mirrors that in Gertler and Karadi (2015): the data starts in July 1979 and ends in June 2012.[3] We only study prime age individuals throughout.[4] Unemployment rates and labor force participation rates are constructed as rates in the relevant groups: If, for example, everyone in one socio-economic group who is part of the labor force had a job, the unemployment rate in

[3] Our aggregate data as well as the data for the instrument are directly taken from the replication material for Mertens and Ravn (2019).

[4] There has been a major redesign of the CPS in 1994. This results in sudden breaks of hours for certain demographic groups. We follow the procedure described in Braun et al. (2018) to remove the breaks in those series. Our hours series is constructed to measure average hours for those who report positive hours.

that group would be 1 irrespective of developments in other groups. [5]

The macro series are the same as in Gertler and Karadi (2015): the one-year government bond rate, the excess bond premium (Gilchrist and Zakrajsek, 2012), the log of the Consumer Price Index (CPI) as well as the log of industrial production. All data series are seasonally adjusted.

3 Our VAR Model

We model all vectors of variables we are studying as follows:

$$\mathbf{y}_t = \mathbf{m} + \sum_{\ell=1}^{\mathcal{L}} \mathbf{A}_\ell \mathbf{y}_{t-\ell} + \mathbf{u}_t \tag{1}$$

where we set the lag length $\mathcal{L}$ to 12 since we use monthly data. To identify a monetary policy shock, we are looking to identify one column of a matrix $\mathbf{\Sigma}$ such that

$$\mathbf{u}_t = \mathbf{\Sigma}\mathbf{e}_t \tag{2}$$

We borrow our approach to identification from Mertens and Ravn (2013), who show how to use an instrument for a structural shock to identify said shock in a VAR. [6] In particular, we use an observed measure of a monetary policy shock $\mathbf{m}_t$ that has to satisfy the following two restrictions:

$$E(\mathbf{m}_t \mathbf{e^m}_t) = \Phi \qquad \text{[relevance condition]} \tag{3}$$

$$E(\mathbf{m}_t \mathbf{e^{r\prime}}_t) = \mathbf{0} \qquad \text{[exogeneity condition]} \tag{4}$$

where Φ is an unknown non-zero scalar, $\mathbf{e^m}_t$ is the monetary policy shock we want to identify, $\mathbf{e^r}_t$ are all other structural shocks such that $[\mathbf{e^m}_t \ \mathbf{e^{r\prime}}_t]' = \mathbf{e}_t$, and $\mathbf{0}$ is a conformable matrix of zeros. Our choice of instrument is the surprise in Federal Open Market Committee (FOMC) dates in the three month ahead monthly Fed Funds futures, which is also the benchmark

[5] We use a 3 point moving average filter to remove measurement error from the micro data. These filters are common in signal processing, where they are called finite impulse response filters. Our benchmark is a two-sided centered version of this filter, but we show in the appendix that our results are robust to using a one-sided (backward-looking) version as well.

[6] As highlighted by Jentsch and Lunsford (2019), the original bootstrap procedure used by Mertens and Ravn (2013) can be problematic. For inference, we instead use the delta method proposed by Montiel Olea et al. (2020) and used by Mertens and Ravn (2019) instead, but show in the appendix that our results are robust to using the parametric bootstrap of Montiel Olea et al. (2020) instead.

choice in Gertler and Karadi (2015).

One issue to keep in mind with this approach is that the impulse responses (IRFs) are only identified up to scale, as highlighted by Stock and Watson (2018). We scale all impulse responses so that the initial impact of a monetary policy shock is a 25 basis point increase in the short-term nominal interest rate. We run separate VARs for each socio-economic group. While these VARs allow us to get a sense of different responses across different socio-economic groups, they cannot help us to get a sense whether the differences are statistically significant since we need the joint distribution of the impulse responses across socio-economic groups for such statements. Thus, we also run VARs where we keep the aggregate variables as before, but now introduce the differences in labor market outcomes between any two socio-economic groups instead of the levels of the outcomes for one group alone. These additional VARs and their associated responses to a monetary policy shock allow us to asses statistical significance of the differences in impulse responses by studying the impulse responses of the differences in outcomes.

4 Aggregate Results

To assess whether our micro data is sensible, we do two things: (i) we aggregate our micro data and compare it to their aggregate counterpart, and (ii) we run a VAR with those aggregated data as well as our standard aggregate data to check if the impulse responses to a monetary shock look reasonable. This is especially important given the contributions by Ramey (2016) and Bu et al. (2020), who find that identification of monetary shocks could lead to counter-intuitive results if the sample is not informative enough about the effects of monetary shocks.

Figure 1 shows the aggregated data versus corresponding aggregate data[7] as well as selected results from the VAR with aggregated variables. Our aggregated unemployment calculated from CPS data tracks the aggregate unemployment rate closely. Concerning impulse responses for aggregate and aggregated variables, our results are standard: the short-term interest rate (labelled GS1 in Figure 1) increases on impact (this is due to our normalization), but stays positive for one year. The price level decreases (in particular, we do not find a price puzzle of the type that has often plagued this literature (Sims, 1992)), and industrial production (IP) decreases. As for the aggregated variables, the unemployment rate (UR) increases[8],and labor force participation (LFP) shows a persistent decrease after a short in-

[7]We focus here on unemployment since unemployment responses are the main focus of this paper.

[8]The response of the unemployment rate is similar to other results reported in the literature, see for example Figure 2 in Ramey (2016).

crease.

For the VARs with disaggregated data shown below, we display only responses of the relevant disaggregated data. The responses of the aggregate variables in those VARs are broadly in line with the results in Figure 1.

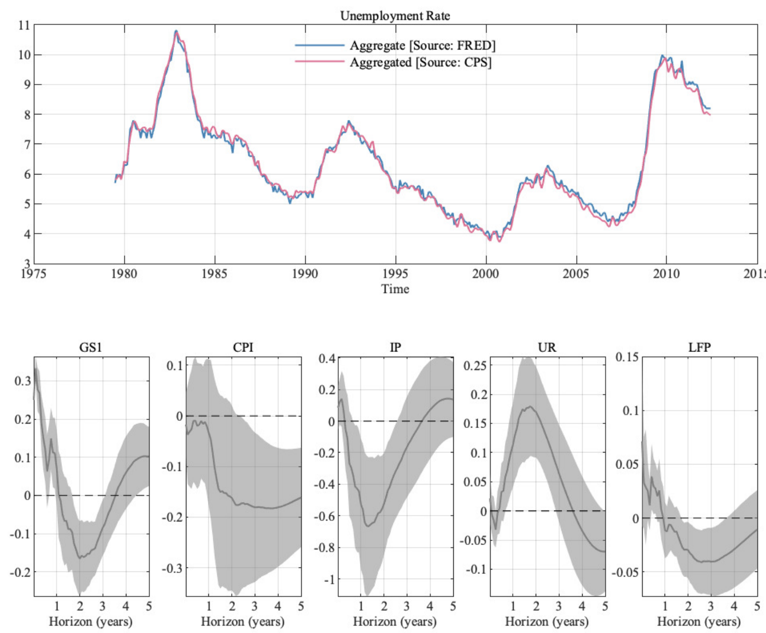

Figure 1: Impulse responses of the VAR with aggregate variables and aggregated versions of our micro data. Error bands are 68% significance bands computed using the delta method.

5 Disaggregated Results

We now turn to describing how the labor market outcomes of various socio-economic groups change after a monetary policy shock. The figures show the time series of the relevant labor market variable for the different subgroups in the top panel, the impulse responses for each subgroup in the middle row (with the response for the aggregated variable from the previous section also displayed in gray), and the bottom row shows the impulse response of the difference between two of the groups.[9]

[9]The impulse responses in the bottom row are not just the differences of the impulse responses in the middle row, as the VARs used to obtain the impulse responses in the bottom row use a different information set (lags of the differences across the groups feature as endogenous variables in those VARs).

5.1 Education

Education has long been identified as a major determinant of economic outcomes. We therefore begin our study of potentially heterogeneous effects of monetary policy on labor market outcomes by studying education. We intentionally use a relatively coarse set of groups so that we can later interact it with other characteristics while keeping the (cross-sectional) sample sizes meaningful - less than high school (i.e. less than 12 years of schooling), less than a college degree (but a high school degree), and finally an undergraduate college degree or more (titled "AboveCollege"). As is well known, the unemployment rate of people with at least a college degree is not only substantially lower than for the other groups, but also less volatile, as can be seen from the top panel of figure 2. The effects on the levels of the unemployment rate differ substantially across these groups, as the unemployment rate of those with at least a college degree reacts substantially less than for the other groups or our aggregated unemployment rate (in gray in each panel in the middle row). The response of the aggregated unemployment rate in gray seems to provide an upper bound for the response across groups, a result that holds across most of the groups we study - we discuss this results in detail in Section 6. The bottom panel shows that differences between the most educated group and the others are long lived, with effects of monetary policy shocks being more severe the less educated a group is. The differences between the two less educated groups, on the other hand, are neither statistically nor economically significant.[10]

The magnitude of the response of the differences can amount to about 0.1 percentage points of the unemployment rate, which is roughly 50 percent of the peak response of the aggregate unemployment rate (in gray in the middle row). These differences are economically meaningful and, as we will see below, they can actually become larger as we dig deeper into heterogeneity among different socio-economic groups.

The fact that less skilled workers are more likely to experience unemployment during recessions (Heathcote et al., 2020) is thus also true for downturns driven by monetary policy.[11]

[10]This is one of the few instances where our choice of filter matters - with the one-sided filter, we see a short lived significant difference on impact in the difference between the two less educated groups. This can be seen in Figure A-2 in Appendix B.1.

[11]Further evidence for this hypothesis is provided in Jefferson (2005), who uses distributed lag models to compute cumulative multipliers of monetary policy shocks on relative unemployment rates. He does not use CPS data and is hence more limited in the sample length, which is just shy of 12 years, whereas we use over 30 years of data in our VAR.

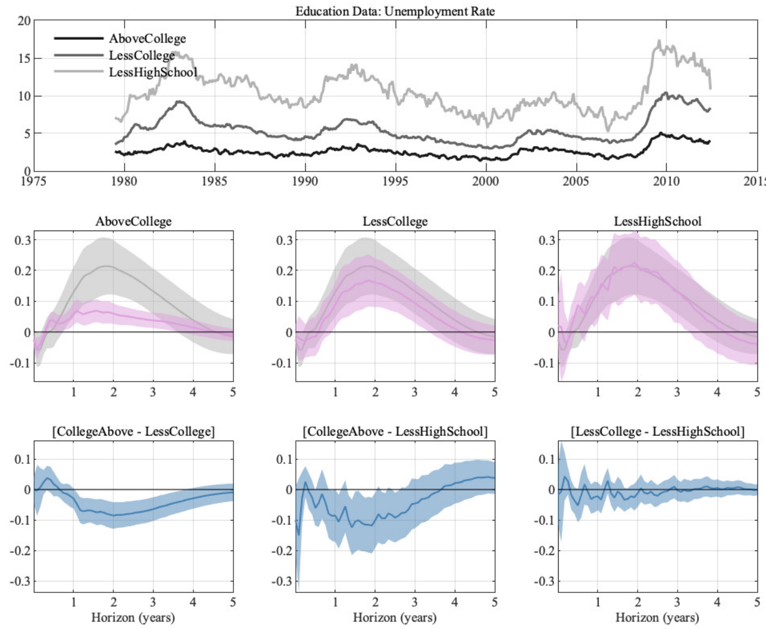

Figure 2: Results for different education levels. Error bands are 68% significance bands computed using the delta method.

5.2 Education and Gender

We now dig deeper into the micro data to assess whether the splitting of survey participants into education groups alone was hiding further heterogeneity. To do so, we use the same education grouping as before, but further split by gender.[12] While the male unemployment rate increases less the higher the level of education, for women the picture is not as clear: while the point estimate is smaller for more educated women, there is much more uncertainty surrounding the responses for less educated women, making those responses mostly statistically insignificant. Studying the IRFs of the differences in unemployment rates across these groups, we find that when comparing the same education level across gender, the unemployment rate increases significantly more for males. In terms of education the differences for a given gender between the most educated and the least educated groups are not only statistically significant, the maximum responses have approximately the same magnitude as the maximum response of the *level of the aggregate unemployment rate*.

[12]Differences in economic outcomes have long been a key issue studied by economists. For recent work, see for example Guvenen et al. (2020).

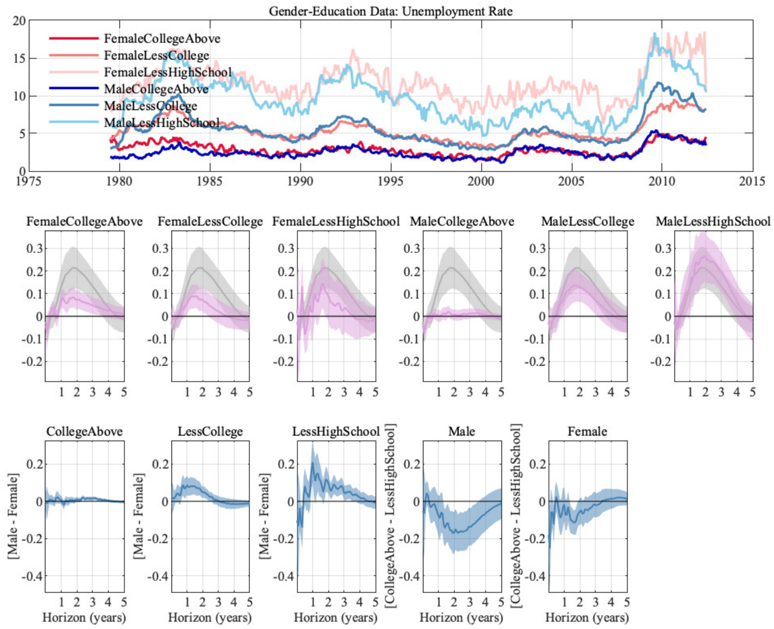

Figure 3: Results for different education and gender levels. Error bands are 68% significance bands computed using the delta method.

5.3 Gender and Marital Status

To understand the differential responses of males and females deeper, we now split our micro data both across gender and marital status.[13]

Differences between married and single women are small and short-lived, as can be seen both from comparing the relevant plots in the middle row of Figure 4, as well as in the second entry of the last row of the same figure. For men, the picture is quite different: single males are hit much harder by a contractionary monetary shock (but also benefit much more from an expansionary shock) than their married counterparts. Furthermore, the responses of single men are much larger than those of female singles.

The magnitudes here are astounding: The estimated peak effect for single males is 50 percent higher than the aggregate response (last panel of the second row of Figure 4). The responses of the differences involving single males are of the same magnitude as the peak effect of the level response of aggregate unemployment.

[13]To keep the samples sizes for the micro data reasonable, we drop education level in this section.

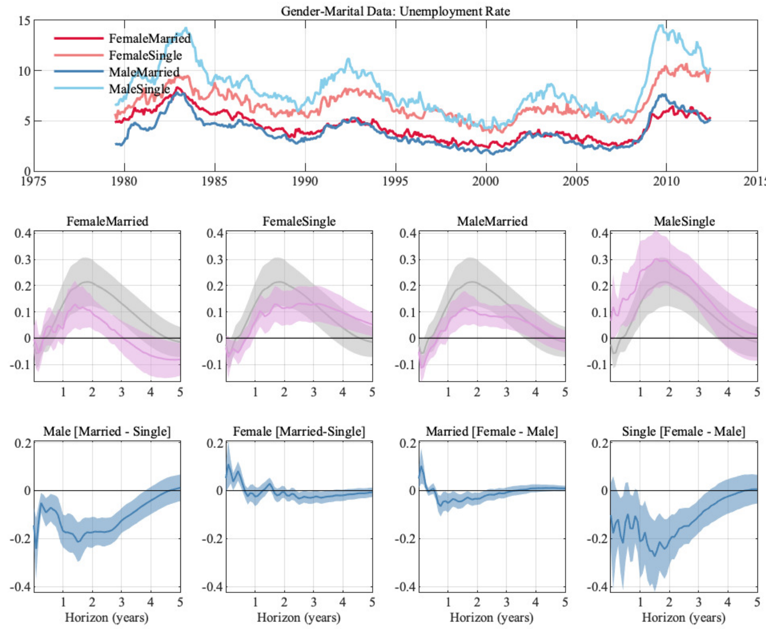

Figure 4: Results across marital status and gender. Error bands are 68% significance bands computed using the delta method.

5.4 Labor Force Participation

While the main focus of our study is on the unemployment rate, we find it useful to also consider how labor force participation (LFP) varies across these groups, highlighting how our results are related to other work in macroeconomics.

We focus in Figure 5 on the responses of the level of LFP for the different groups (as well as for our aggregated version of LFP).

As with the unemployment rate, various different groups have much larger swings in labor force participation than those we see for the aggregated variable (in gray in the last panel of the second row). However, another pattern emerges: Some women change their labor force participation in the *opposite* direction relative to the aggregate response. In particular, women with less than a high school degree and those that are married increase their labor force participation rate in response to a contractionary monetary policy shock.[14]

There is by now a growing literature highlighting how women's labor market participation has changed and how a woman's labor market status may change to smooth shocks faced by the

[14] There is obviously a overlap between those groups - at the beginning of our sample 24 percent of married women had not completed high school, whereas at the end of our sample it was 6 percent. The majority of women without high school degree are married in our sample - 80 percent at the beginning of our sample and 59 percent at the end.

entire family (Albanesi, 2019; Doepke and Tertilt, 2016; Gorbachev, 2016; Ellieroth, 2019). Our results show that over our sample married women entered the labor force, possibly to counteract the prospect of job loss by their spouse. The aggregate negative response of LFP seems to be mainly driven by both male and female singles.

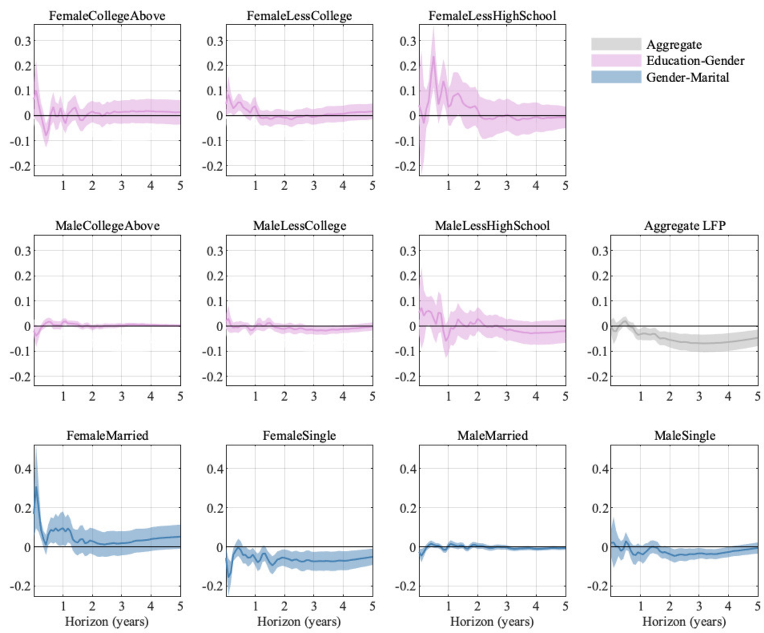

Figure 5: Impulse responses of the labor force participation rate across groups. Error bands are 68% significance bands computed using the delta method.

6 Disaggregation - Some Analytics

Finally, as we mentioned before, throughout our examples the aggregate impulse response seems to provide an approximate upper bound to the impulse responses of most socio-economic groups (the one group that has a substantially larger response are single males in the previous section). To understand possible sources of this result, consider the following stylized environment: We have at our disposal data on a unit-variance economic shock ε_t (for simplicity, we abstract here from mismeasurement that we allow for in our empirical work) as well as two (mean zero) data series $x_{1,t}$ and $x_{2,t}$, which we can aggregate into $x_t \equiv ax_{1,t} + bx_{1,t}$, where a and b are non-negative weights used for aggregation. To further simplify this example, we (i) assume all data are independently and identically distributed (iid), and (ii) focus on population regressions, i.e. the hypothetical scenario where we have access to an infinite amount of data. In such a situation, we could recover the impact effect of the monetary shock on these variables, i.e. the population regression coefficients of three regressions of interest:

$$
\begin{aligned}
x_t \quad &= \alpha \ \varepsilon_t + u_t && (5)\\
x_{1,t} \quad &= \alpha_1 \ \varepsilon_t + u_t^1 && (6)\\
x_{2,t} \quad &= \alpha_2 \ \varepsilon_t + u_t^2 && (7)
\end{aligned}
$$

u_t, u_t^1, u_t^2 are the regression residuals. The question we want to know is how the weights a and b in the data construction are linked to a weight w defined via:

$$\alpha = w\alpha_1 + (1 - w)\alpha_2 \tag{8}$$

This weight summarizes how heavily the impact response of $x_{1,t}$ influences the impact response of aggregated data x_t. Using the standard OLS formula and the fact that the variance of ε_t is 1, we find that

$$w = \frac{a * cov(y_{1,t}, \varepsilon_t) + (b - 1) * (cov(y_{2,t}, \varepsilon_t))}{cov(y_{1,t}, \varepsilon_t) - cov(y_{2,t}, \varepsilon_t)} \tag{9}$$

We can clearly see that the weights depend on the weights used in aggregation and the larger a, the larger the weight w. Furthermore, we can see from the data series in our earlier plots that those groups that have IRFs closest to the IRF of the aggregated series have more volatile unemployment series. If part of that volatility is driven by monetary shocks then the covariance of that group's outcome with the monetary shock is larger than for other groups, leading to an even larger weight w, as is evident from equation (9). Thus, it is not surprising that the responses of larger groups in the population that also have more volatile unemployment rates such as single males have a large influence on the IRF of the aggregated unemployment series.

7 Conclusion

We have used standard identification assumptions in macroeconomics in conjunction with well-known micro data on labor market outcomes to shed light on how the unemployment rate (and labor force participation) for different groups react to monetary policy shocks.
We find substantial heterogeneity across individuals in the US when it comes to the sensitivity

to monetary policy shocks. The impulse response of the differences across groups often has a peak that is as large in magnitude as the peak of the response of the aggregate unemployment. Aggregate responses thus mask a large amount of heterogeneity across groups: less educated individuals show a substantially larger sensitivity to monetary policy shocks, as do single males (with there certainly being substantial overlap between these two groups).

Dynamic equilibrium models (Auclert, 2019; Kaplan et al., 2018; Gornemann et al., 2016) highlight channels that can lead to substantial heterogeneity in individuals' responses to monetary policy shocks. An interesting question for future research is whether the frictions already present in these models are enough to account for our findings or if additional frictions related to the groups we studied are needed. While we normalized our impulse responses to show a contractionary monetary policy shock, this also means that these more sensitive groups benefit more from expansionary policy than what one would expect form looking at the aggregate impulse responses alone.

References

Albanesi, S. (2019, March). Changing business cycles: The role of women's employment. Working Paper 25655, National Bureau of Economic Research.

Andersen, A. L., N. Johannesen, M. Jørgensen, and J.-L. Peydró (2020, December). Monetary Policy and Inequality. Working Papers 1227, Barcelona Graduate School of Economics.

Auclert, A. (2019, June). Monetary Policy and the Redistribution Channel. *American Economic Review 109*(6), 2333–2367.

Bartscher, A. K., M. Kuhn, M. Schularick, and P. Wachtel (2021, January). Monetary Policy and Racial Inequality. Staff Reports 959, Federal Reserve Bank of New York.

Bernanke, B. S. (2015, Jun). Monetary policy and inequality.

Braun, C., F. Kydland, and P. Rupert (2018, February). Quality Hours: Measuring Labor Input. Working paper, UC Santa Barbara.

Bu, C., J. Rogers, and W. Wu (2020, February). Forward-Looking Monetary Policy and the Transmission of Conventional Monetary Policy Shocks. Finance and Economics Discussion Series 2020-014, Board of Governors of the Federal Reserve System (U.S.).

Christiano, L., M. Eichenbaum, and C. Evans (1999). Monetary policy shocks: What have we learned and to what end? In J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, Volume 1A, pp. 65–148. Elsevier.

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

Doepke, M. and M. Schneider (2006, December). Inflation and the Redistribution of Nominal Wealth. *Journal of Political Economy 114*(6), 1069–1097.

Doepke, M. and M. Tertilt (2016, December). Families in Macroeconomics. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2 of *Handbook of Macroeconomics*, Chapter 0, pp. 1789–1891. Elsevier.

Ellieroth, K. (2019). Spousal Insurance, Precautionary Labor Supply, and the Business Cycle - A Quantitative Analysis. 2019 Meeting Papers 1134, Society for Economic Dynamics.

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

Gilchrist, S. and E. Zakrajsek (2012, June). Credit Spreads and Business Cycle Fluctuations. *American Economic Review 102*(4), 1692–1720.

Gorbachev, O. (2016, May). Has the increased attachment of women to the labor market changed a family's ability to smooth income shocks? *American Economic Review 106*(5), 247–51.

Gornemann, N., K. Kuester, and M. Nakajima (2016, May). Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy. International Finance Discussion Papers 1167, Board of Governors of the Federal Reserve System (U.S.).

Guvenen, F., G. Kaplan, and J. Song (2020, April). *The Glass Ceiling and the Paper Floor: Changing Gender Composition of Top Earners since the 1980s*, pp. 309–373. University of Chicago Press.

Heathcote, J., F. Perri, and G. Violante (2020, August). The Rise of US Earnings Inequality: Does the Cycle Drive the Trend? *Review of Economic Dynamics 37*, 181–204.

Holm, M., P. Paul, and A. Tischbirek (2021). The Transmission of Monetary Policy under the Microscope. *Journal of Political Economy, forthcoming*.

Jefferson, P. N. (2005, May). Does monetary policy affect relative educational unemployment rates? *American Economic Review 95*(2), 76–82.

Jentsch, C. and K. G. Lunsford (2019, July). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States: Comment. *American Economic Review 109*(7), 2655–2678.

Kaplan, G., B. Moll, and G. L. Violante (2018, March). Monetary Policy According to HANK. *American Economic Review 108*(3), 697–743.

Leeper, E. M., C. A. Sims, and T. Zha (1996). What Does Monetary Policy Do? *Brookings Papers on Economic Activity 27*(2), 1–78.

Lenza, M. and J. Slacalek (2018, October). How does monetary policy affect income and wealth inequality? Evidence from quantitative easing in the euro area. Working Paper Series 2190, European Central Bank.

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

Mertens, K. and M. O. Ravn (2019, July). The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States: Reply. *American Economic Review 109*(7), 2679–2691.

Montiel Olea, J. L., J. H. Stock, and M. W. Watson (2020). Inference in structural vector autoregressions identified with an external instrument. *Journal of Econometrics*.

Ramey, V. (2016). Macroeconomic Shocks and Their Propagation. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2 of *Handbook of Macroeconomics*, Chapter 0, pp. 71–162. Elsevier.

Sims, C. A. (1992, June). Interpreting the macroeconomic time series facts : The effects of monetary policy. *European Economic Review 36*(5), 975–1000.

Stock, J. H. and M. W. Watson (2018, May). Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *Economic Journal 128*(610), 917–948.

Appendix For "What Does Monetary Policy Do To Different People?"

A Noise Filtering

To remove noise from the raw micro data we employ a 3-point moving average filter. To highlight that our findings are robust, we show results with a one-sided filter in this appendix (the results for the two-sided filter are in the main text). Let us denote the noisy time t raw data by x_t. For the two-sided filter we implement

$$y_t = a_1 x_{t+1} + a_2 x_t + a_3 x_{t-1} \tag{A-1}$$

with y_t denoting the corresponding time t noise filtered data and corresponding weights $a_1 = .25$, $a_2 = .5$ and $a_3 = .25$. For the one-sided filter we implement

$$y_t = a_1 x_t + a_2 x_{t-1} + a_3 x_{t-2} \tag{A-2}$$

with corresponding weights $a_1 = .5$, $a_2 = .25$ and $a_3 = .25$. In Figure A-1 we show the comparison of raw, one-sided and two-sided filtered time series of our aggregated unemployment data.

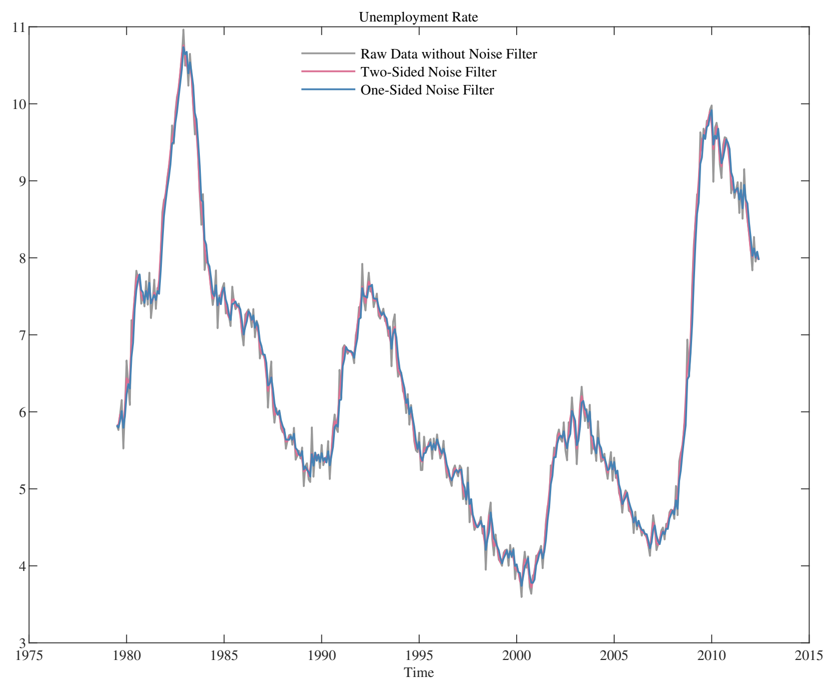

Figure A-1: Raw data versus one-sided and two-sided filter applied to our aggregated unemployment data.

B Additional Impulse Response Graphs

B.0.1 Results with One-sided Filter

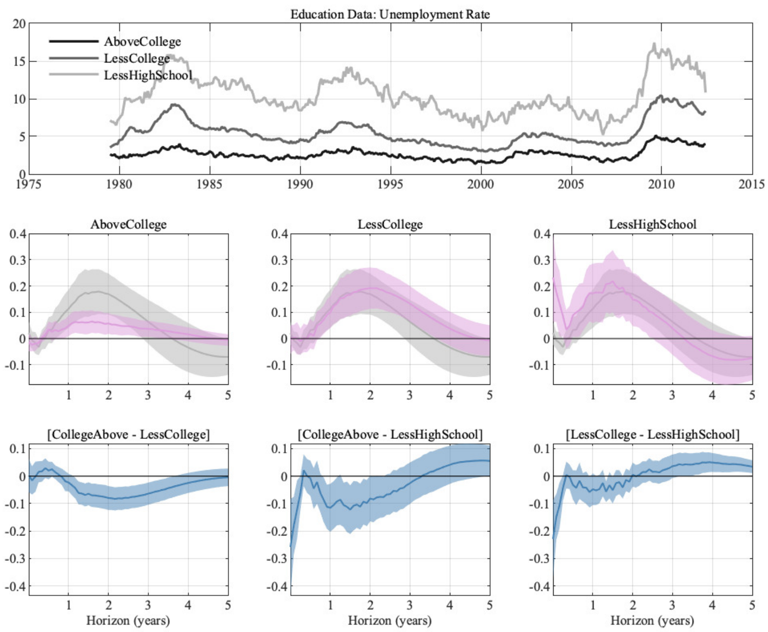

Figure A-2: Results for different education levels. Error bands are 68% significance bands computed using the delta method.

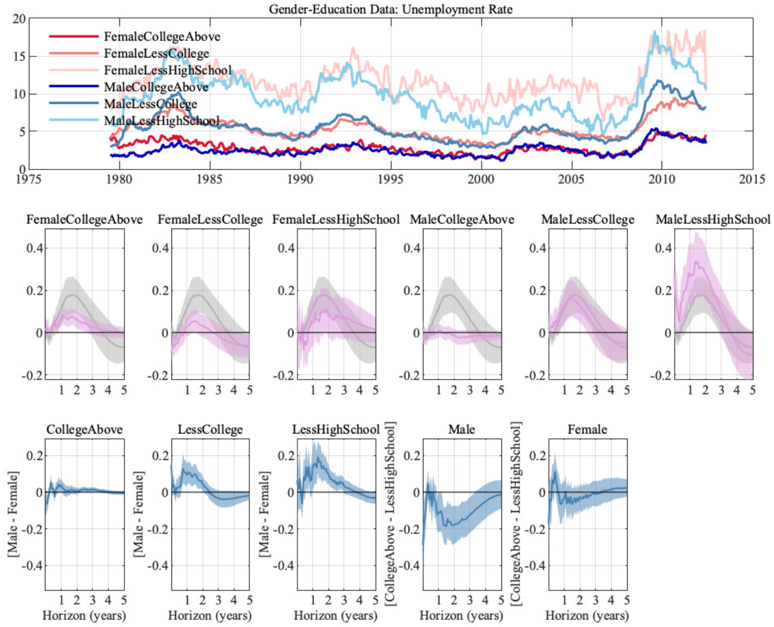

Figure A-3: Results for different education and gender levels. Error bands are 68% significance bands computed using the delta method.

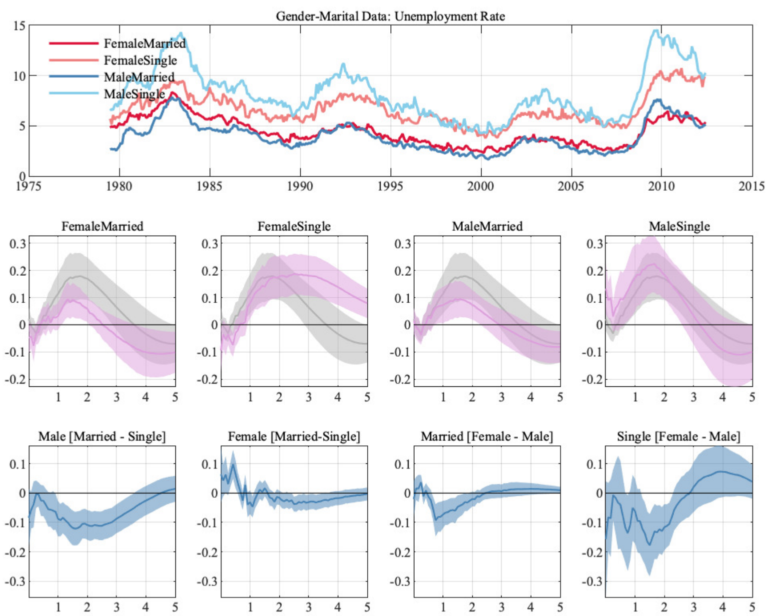

Figure A-4: Results across marital status and gender. Error bands are 68% significance bands computed using the delta method.

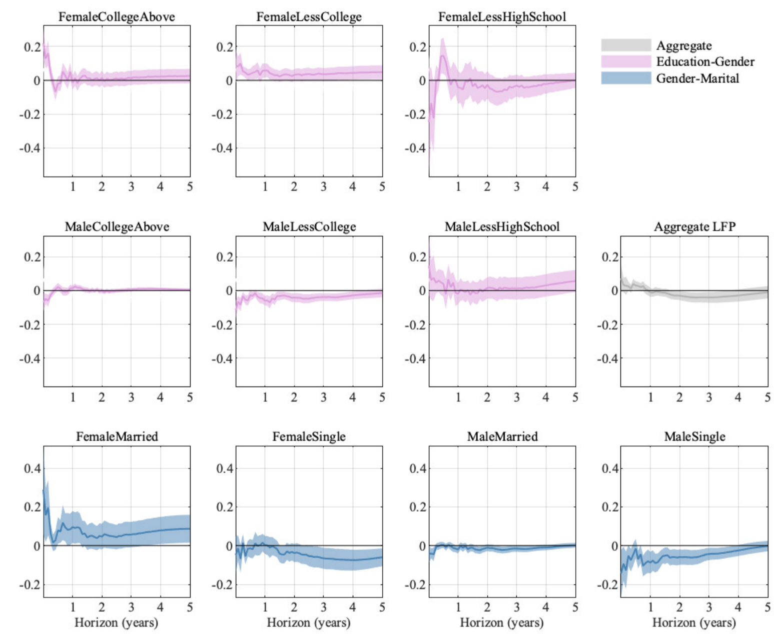

Figure A-5: Impulse responses of the labor force participation rate across groups. Error bands are 68% significance bands computed using the delta method.

B.1 Alternative Computation of Confidence Intervals

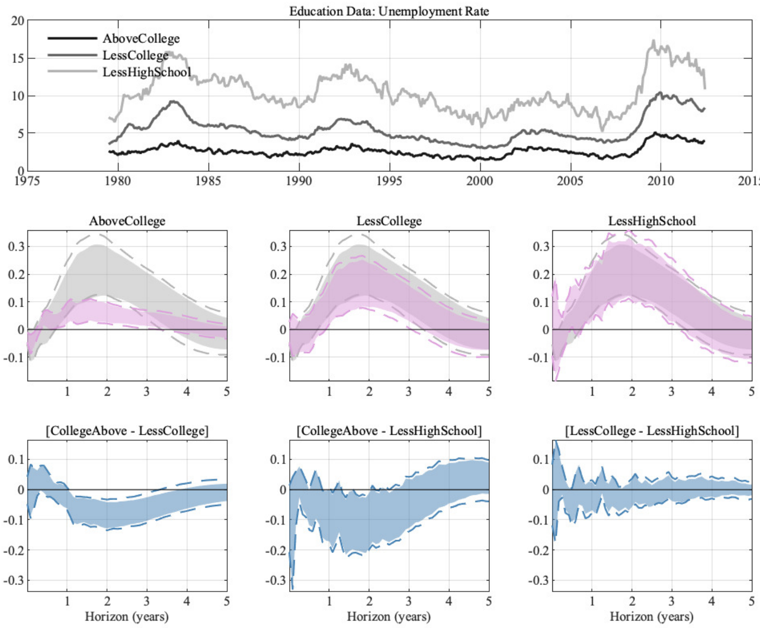

Figure A-6: Results for different education levels. Dashed error bands are 68% significance bands computed using the Montiel Olea et al. (2020) parametric bootstrap, shaded error bands use the delta method.

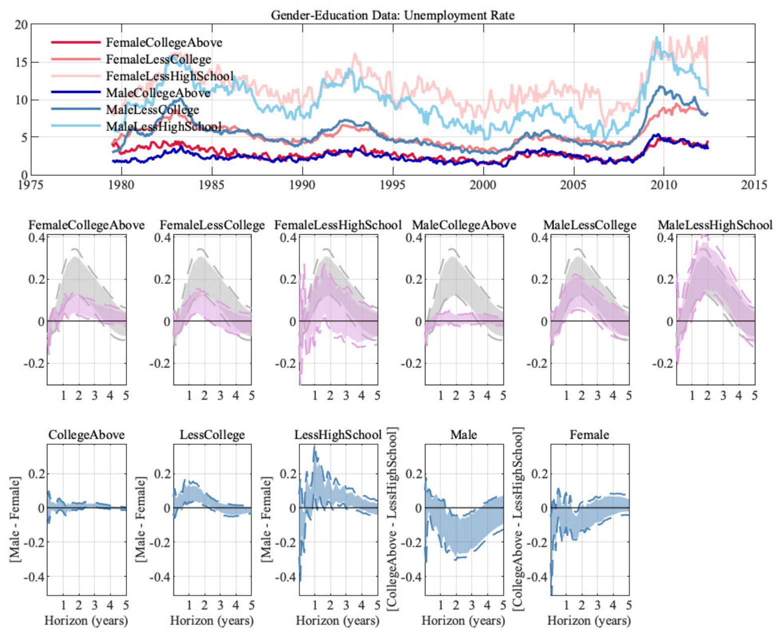

Figure A-7: Results for different education and gender levels. Dashed error bands are 68% significance bands computed using the Montiel Olea et al. (2020) parametric bootstrap, shaded error bands use the delta method.

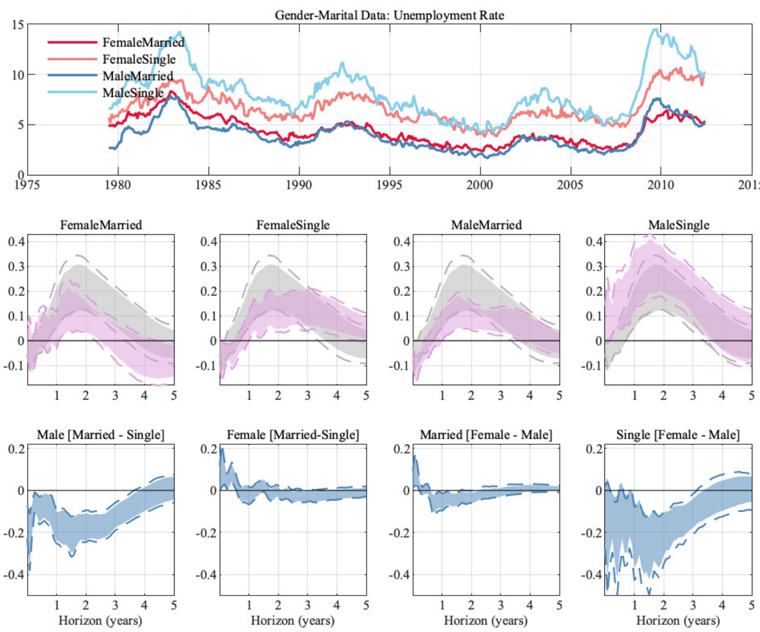

Figure A-8: Results across marital status and gender. Dashed error bands are 68% significance bands computed using the Montiel Olea et al. (2020) parametric bootstrap, shaded error bands use the delta method.

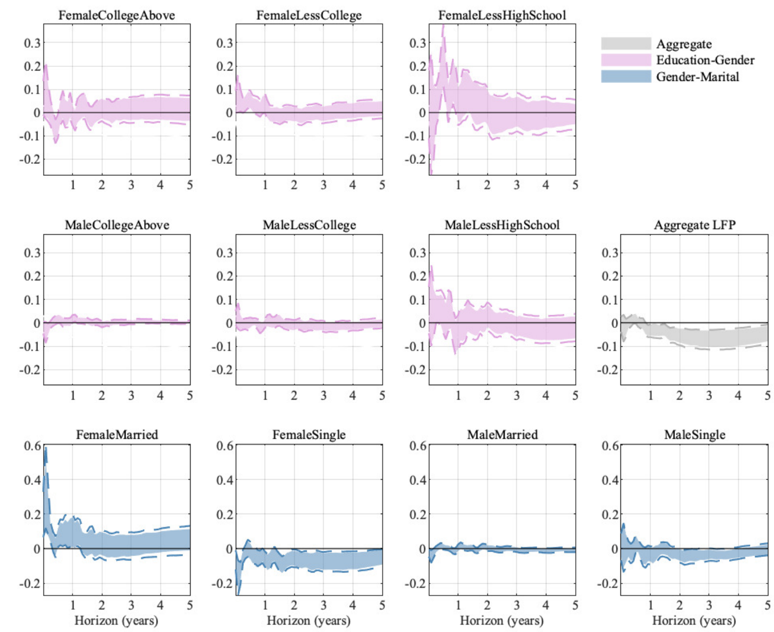

Figure A-9: Impulse responses of the labor force participation rate across groups. Dashed error bands are 68% significance bands computed using the Montiel Olea et al. (2020) parametric bootstrap, shaded error bands use the delta method.