Monetary Shock Asymmetry and Spatial Heterogeneity Along the Household Balance Sheet Channel

Thomas Decker*

UNC – Chapel Hill

Abstract

Using local projections, I find strong evidence for asymmetrical responses of employment, wages, and inflation to monetary shocks. In particular, I document that employment drops almost four times more subsequent to a contractionary shock than it rises following an expansionary shock. Further, I find that differences in the household balance sheet are partially responsible for state level heterogeneity, particularly in employment. Accounting for the household balance sheet limits the extremity in the states' employment responses and mildly homogenizes the states.

*I would like to acknowledge and express my gratitude for the generous support of the Matthew Guest Grant from the Economics Department at UNC. I would also like to thank my advisor, Dr. Christopher Roark, for phenomenal feedback throughout this process.

I. Introduction

The notion that contractions and expansions do not produce symmetric economic responses is not particularly novel. Nonetheless, it has not received much attention in the literature. Considering the responses of inflation, employment, and wages, the most straightforward driver of shock asymmetry is downward nominal wage rigidity. If nominal wages are sticky, they would not considerably fall following a contractionary shock; instead, employment would be likely to fall by more – workers would sooner be laid off than accept a reduction in the nominal wage. Indeed, Abritti and Fahr (2013) find evidence for sticky wages as a factor of monetary shock asymmetry. Developing a New Keynesian Dynamic Stochastic General Equilibrium (DSGE) model, they find that wage rigidity accounts for asymmetry amongst different business cycle shocks, not exclusively monetary shocks. Their findings are in line with others, such as Christiano, Eichenbaum, and Evans (2005), who find that nominal rigidities (specifically nominal wage rigidity) are a crucial component to adjustment nationally following a monetary shock.

While wage rigidity is the simplest explanation for shock asymmetries, it is certainly not exclusive. The household balance sheet channel (HHBSC), a means of monetary policy transmission, could also produce asymmetrical responses to monetary shocks. To flesh out a baseline theoretical framework for the HHBSC, I follow Mishkin (1996). He describes two views of the HHBSC, a credit-side view and liquidity-effect view. In the former lens, increases in the interest rate would decrease bank lending to consumers, in addition to decreasing household income. These declines in household financial position would then translate into declines in consumer durable and housing spending. In the latter view, an increase in the interest rate would decrease the value of financial assets held by consumers, thus increasing their perceived likelihood of suffering financial distress. As the risk of distress rises, consumers shift away from illiquid assets (i.e., durables, housing). Thus, aggregate demand would fall, affecting the macroeconomy (Mishkin 1996). Outside of the framework of asymmetry in monetary policy, this linkage between durables and declining aggregate demand has been studied as a complicating factor in both the Great Depression (Mishkin 1978) and the Great Recession (Mian and Sufi 2010; Mian and Sufi 2014). Since rigidities have been studied in multiple contexts related to monetary shocks, I instead focus on further exploring the household balance sheet.

Because illiquid assets, such as housing, are major sources of wealth, holding more liquid assets could have a problematic effect on the economy (Mishkin 1996). Since these assets have lower overall returns, consumers would then have a lower income after shifting to more liquid assets. In this way, a monetary shock leading to the holding of more liquid assets would drive a deeper cut in the macroeconomy, as the financial portion of the balance sheet would fall, augmenting general income declines. A later expansion would not see a symmetrical response because of the lost income during the contraction.

The household balance sheet should also vary over space; that is, heterogeneity in state level responses should at least partially be accounted for by the HHBSC. There is a modest literature researching the spatial heterogeneity of monetary policy, beginning with Carlino and Defina (1998). Using Bureau of Economic Analysis (BEA) regions as the level of study, they find evidence of asymmetry in the dispersion of monetary policy via the interest rate channel. The amount of manufacturing in the region is the most related to the degree to which monetary policy might impact a region differently, though the concentration of small firms—transmission through the bank lending channel—is weakly relevant, too. They extend their analysis to the state-level, finding that the bank lending channel's impact vanishes, but the credit channel's role is stronger, and the interest rate channel is still relevant (Carlino and Defina 1999). Owyang, Piger, and Wall (2005) separate monetary policy before and after the Volcker inflation, which yields differences in the regionality and severity of shocks across the nation. Francis, Owyang, and Sekhpoysan (2012) examine spatial effects at the city-level. Rather than the interest rate or the credit channels, they find that local fiscal policy and population density explain regional variations. Pedemonte and Herreño (2023) find a relationship between areas with high marginal propensity to consume (MPC), as well as poorer regions, and stronger effects of monetary policy. Auclert (2019) does investigate the household balance sheet as it relates to monetary policy transmission and proposes a redistribution channel for dispersion. He focuses, however, on the New Keynesian relationship between the HHBS and its ramifications on MPCs. Outside of a monetary policy context, regional heterogeneity has been explored at various levels (e.g., Canova and Pappa 2007; Mihov and Scott 2001). None of these studies, however, centralize regional differences in the household balance sheet. Asymmetries in information and access to the financial economy should exist geographically. That is, individuals in far-removed rural areas with limited financial sectors would not necessarily have the same interaction with the financial economy and therefore may not behave identically to individuals in more connected states.

To test whether these expectations of the household balance sheet channel are accurate, I use local projections, a form of direct forecasting first described by Jordà (2005), instead of the more traditional vector autoregression (VAR). A more in-depth discussion of the two is offered by Jordà (2005), but, briefly, local projections are more straightforward than VARs. Local projections have become a prominent alternative due to their relative ease of estimation, including in the arena of monetary shocks. Using data from the Quarterly Census of Wages and

Employment (QCEW) at the state-level from 1975 onwards, a novel state level inflation series developed by Hazell et al. (2022) and shocks from the Wieland and Yang (2020) extended Romer and Romer (2004) series and Bu, Rogers, and Wu (2021), I find the effect these shocks have on employment, wages, and inflation. I then check whether these are symmetrical via an interaction term before introducing the household balance sheet and extending the analysis to the states.

I find that there is no symmetry between contractionary and expansionary shocks for any of the variables. Specifically, employment troughs following a contraction are almost four times as deep as peaks are high following an expansion. While wage responses are asymmetrical, the effect is rather small, lending credence to the notion that wages are sticky. At the state level, I find evidence for heterogeneous responses across space, and the employment responses in particular are well-moderated by the household balance sheet channel. The household balance sheet can also explain some of the asymmetry between expansionary and contractionary shocks.

To the best of my knowledge, this research is the first to use local projections in investigating shock asymmetries, and to analyze the household balance sheet channel over space. Shock asymmetries have received little attention in the literature thus far. Cover (1992) does find that output is only impacted by negative money-supply shocks. Hanson, Hurst, and Park (2006) take a similar approach, though look at how expansionary and contractionary shocks affect state level growth rates. Their paper is similar to Tenreyo and Thwaites (2016), who look at how different starting points in the business cycle affect the economic response to monetary policy. Accordingly, this paper is a departure from these previous studies in several ways outside of the aforementioned two: I estimate the response of employment, wages, and inflation to two series of monetary shocks.

The papers most similar to mine are Pedemonte and Herreño (2023) and Abritti and Fahr (2013), though I make several key departures from each of these. While Pedemonte and Herreno (2023) also use local projections in their analysis, they focus on the metropolitan statistical area (MSA) as their aggregation level. This choice is fair considering the limitations of state-level time series data on inflation. However, I use a novel state level inflation series constructed by Hazell et al. (2022) to work around this barrier; this is, to the best of my knowledge, the first approach of its kind using their series. Moreover, Pedemonte and Herreño (2023) also centralize the role of New Keynesian heterogenous agents in their analysis, whereas I instead focus on more broad asymmetries in the effects of monetary shocks. Their research and this paper also both use the Wieland and Yang (2020)-extended Romer and Romer (2004) "narrative" monetary shock series, though, looking to extend my analysis after 2008, I also use the shock series created by Bu, Rogers, and Wu (2021). Abritti and Fahr (2013) do investigate asymmetries in the economic response to monetary shocks across a host of variables, but do so via a moments analysis, rather than a VAR or local projections. Accordingly, both theoretically and empirically, this research diverges.

Following the discussion of related literature, the remainder of this paper is divided into four sections. Section II presents the data. Section III contains the analysis. Section IV contains discussion and Section V concludes.

I.B Related Literature

In addition to the aforementioned spatial heterogeneity literature, this paper contributes to a burgeoning literature on general monetary heterogeneity. Coibion *et al*. (2018) and Auclert (2019) examine distributional implications of asymmetric consumption effects, finding evidence that contractionary monetary policy increases inequality and that redistribution itself is a channel for monetary policy. Wong (2018) finds an age effect of debt on the transmission process of monetary shocks, rooted in the tendency of individuals under 35 with mortgages to refinance their mortgages after monetary shocks, which she dubs the refinancing channel. Cloyne, Ferria, and Surico (2019) similarly investigate households with debt, focusing instead on the different responses to monetary policies between renters, homeowners with mortgages, and homeowners without mortgages in the US and UK. The variety in liquidity in the balance sheets of these categories of individuals, they find, drives consumption changes in response to monetary policies: mortgages take up much of their holder's balance sheets, with little other more liquid wealth being held. Beyond consumption, Bergman, Matsa, and Weber (2022) look at demographic-related effects of monetary policy, particularly on employment. They find that tighter US labor markets disproportionately benefit those with labor force attachment—Black individuals, women, and those without high school degrees. These sorts of demographic characteristics would, of course, show up across geographies.

More broadly, the household balance sheet channel has been shown to have contributed to both the Great Depression (Mishkin 1978) and the Great Recession (Mian and Sufi 2010; Mian and Sufi 2014). In addition to household balance sheet heterogeneity and monetary policy, a wide set of literature has examined the relationship between the household balance sheet and the whole economy generally. High household leverage has been shown to induce declining consumption—thus declining aggregate demand—in the lead up to the 2008 financial crisis (Sufi and Mian 2010). Mian and Sufi (2014) tie the persistent decline in employment following the Great Recession to the decrease in aggregate demand stemming from the household balance sheet shock. They dubbed this the aggregate demand channel of employment; the findings of my analysis vis-à-vis employment fit with this construction of aggregate demand and employment.

Similarly, Mian and Sufi (2014) also find a role for a housing net worth channel in driving unemployment in this time. Considering the household balance sheet from a different perspective, Chodorow-Reich, Nenov, and Simsek (2021) find a small but significant effect of stock market growth on local labor markets. As could be predicted, stock market growth induces a consumption wealth effect.

II. Data

For this analysis, I construct a panel of data for the US states. Employment and wage figures come directly from the Quarterly Census of Employment and Wages (QCEW), conducted by the BLS. Data starts with the earliest available year, 1975. Inflation data comes from the series constructed by Hazell et al. (2022). They scraped historical BLS Consumer Price Index (CPI) documents for most US states to construct a state-level inflation time series.[1] This data is balanced only after 1989, a drawback to using the metric.[2] The inflation series ends in 2017, thus limiting the analysis of prices by two years.

For monetary shocks, I employ the Romer and Romer (2004) shock series, which begins in 1969, extended to the end of 2007 by Wieland and Yang (2020). For a comparative check with the advent of monetary policy via large scale asset purchases (LSAP, or more familiarly, quantitative easing), I also use the series constructed by Bu, Rogers, and Wu (2021), which spans from the start of 1994 through the third quarter of 2019. The overlap in the series is not highly correlated (R = 0.21 at the monthly level), so combining the series is not logical. Moreover, the

[1] Several small states are omitted from the time series, as well as Arizona due to barriers to entering the BLS to access the data during the COVID-19 Pandemic. The full list of excluded states is: Arizona, Delaware, Kentucky, Iowa, Idaho, Montana, Maine, New Hampshire, New Mexico, Nebraska, Nevada, North Dakota, South Dakota, Wyoming, West Virginia, Vermont, Rhode Island.

[2] The data begin at varied times, depending on the state. For uniformity, I begin in 1989. Note that Alaska (2015 Q4) is not in the data.

two are constructed quite distinctly. Bu, Rogers, and Wu (2021) use a Fama-MacBeth multistep technique to derive their shocks, first estimating, via an instrumental variable, the sensitivity of asset yields at different maturities to monetary policy announcements. Then, they use the sensitivity as the explanatory variable in cross-sectional regressions where the interest rate at different times along the 30-year yield curve is the dependent variable; the coefficient in this regression is the shock. Romer and Romer (2004), dissimilarly, use a "narrative" approach by which they determine an intended interest rate, based on what was discussed during FOMC meetings as opposed to what is contained in forecasts. With this "intended series," they find the change between the projected and discussed Federal Funds rates for a given meeting and regress it on the Greenbook forecasts of different variables. Their shocks are the residuals of this regression, the change in the interest rate not determined by their forecasts.

As mentioned, the overlap of the two series has a low correlation, and the two also differ in several other ways. By magnitude, the Romer and Romer (2004) data are much larger, with a minimum of -0.56 and a maximum of 0.75. In comparison, the Bu, Rogers, Wu (2021) shocks have a minimum of -0.11 and a maximum of 0.09. The mean, standard deviation, and median are all also quite different. The ramifications of these differences on this analysis will be discussed below, though it is reasonable to expect, at the very least, that they might produce different shocks.

To investigate the household balance sheet channel, I use quarterly dividend, interest, and rent income (DIR) reported by the BEA to assess the relative financial position of each state. Using annual population data from the BEA, I then convert this data into per capita terms. Throughout the analysis, I consider two ways of measuring DIR. In the aggregate, I control for the several quarters of lagged DIR, taking logs to approximate the percent change between quarters. At the state-level, I instead consider *relative* DIR. That is, the DIR of each state at a given time relative to the average across all states. The full rationale for the change is discussed later, but, in short, the national average for each period is 1 (and thus stationary). This might cause issues in panel regressions, but change in the level of DIR would not, hence its use in the aggregate.

For the QCEW data, I take the total state level employment and mean weekly wages for all types of employers.[3] Since employment is reported by month, I create a simple average of the three data points for quarterly employment. Data for Alaska and Washington, DC, are missing from 1977-1980; accordingly, I linearly interpolate estimates for these years to obtain a full panel. Because the QCEW is maintained as the results of each individual survey year, the data are initially seasonal. To highlight this, Figure A.I shows wages and employment in Alabama, as an example, before and after adjustments are made. I smooth the employment data using a double exponential time series adjustment. A simple moving average smoothing did not effectively smooth employment for Alaska, Wyoming, and Idaho, among a few other states, hence the decision for the double exponential process. The only drawback to the double exponential adjustment is that the smoothed data are lower than the original time series (see Figure A.II for a visual comparison of the two methods). For wages, however, I opt for a simple moving average, as the process sufficed in smoothing wages. After adjusting, I then take logs and find the percent change between quarters for both variables.

III. Aggregate Analysis

---

[3] Data from 1975-1989 use the SIC industry classification system, whereas data after 1990 use the NAICS; this does not implicate on this analysis.

Through this analysis, I use local projection regressions as the main form of empirical analysis. Local projections were introduced by Jordà (2005) and have since become a prominent alternative to the more traditional VAR since they are direct forecasting. They perform as well, if not better than VARs, even if the VAR is optimal for the data structure and are less vulnerable to errors of misspecification (Plagborg-Moeller and Montiel Olea 2021). To begin, I estimate local projections across all locations as baselines for the state level analysis. These can be thought of as a national average, or equivalently as the effects for the average state. Then, I use the results to generate cumulative impulse response functions. For inflation, I estimate:

$$\pi_{i,t+h} = \alpha_i^h + \beta^h s_t + \sum_{j=1}^{J} \gamma_{t-j}\pi_{i,t-j} + \varepsilon_{i,t+h} \; \forall \; h \in [0, H]$$

where *i* indexes the state, *h* is the number of quarters after the shock, and *t* indexes time. β^h represents the effect of the monetary shock, s_t, in time t on inflation, $\pi_{i,t+h}$, at horizon h. α_i^h is the state level fixed effect, $\varepsilon_{i,t+h}$ is the error term, and the summed γ_{t-j} is the combined coefficient for lagged inflation. J is the maximum number of lags.

Analogously, for change in employment, I estimate:

$$e_{i,t+h} = \alpha_i^h + \beta^h s_t + \sum_{j=1}^{J} \gamma_{t-j}e_{i,t-j} + \varepsilon_{i,t+h} \; \forall \; h \in [0, H]$$

with the only difference being that changes in employment from the previous quarter are represented by e. Similarly, for wages, I estimate:

$$w_{i,t+h} = \alpha_i^h + \beta^h s_t + \sum_{j=1}^{J} \gamma_{t-j}w_{i,t-j} + \varepsilon_{i,t+h} \; \forall \; h \in [0, H]$$

with *w* representing the change in the nominal wage from the previous quarter.

It is worthwhile to note that, as I estimate cumulative impulse response functions, the dependent variable in equations 2 and 3 is difference between the logs of wages and employment from time t+h to time t-1. Since these variables are then percent changes, the lagged percent change (as the difference in logs between quarters) is then the corresponding control. As inflation is already a percent change, the left-hand side variable is the difference between the forecasted period and time t-1, and the lagged control is the simple inflation rate.

As mentioned, I use two variants for the shock data, the extended Romer and Romer (RR) series (Romer and Romer 2004; Wieland and Yang 2020) and the Bu, Rodgers, Wu (BRW) series (2021). Impulse response functions at the 95% confidence level are shown in Figure I, with the RR shocks presented on the left and the BRW shocks on the right. I use H = 24 and J = 4, with the use of 4 lags for consistency with the ensuing state level analysis, for which more lags would infringe on the degrees of freedom.

Graphs can be interpreted as the cumulative percent change of the dependent variable (inflation, employment, or wages) to an unanticipated one percentage point monetary shock. The black line tracks the dynamic multipliers at each horizon (number of quarters ahead), and the area banded by the dashed lines is the confidence interval.

Employment and wages have noticeably different responses to the RR shocks than they do from the BRW shocks. In shape, these results are partially akin to the findings of Pedemonte and Herreño (2023) for employment, which are in turn alike the original findings of Romer and Romer (2004). Their research is, however, at the Metropolitan Statistical Area (MSA) level, so differences between the analysis here and previous results are to be expected. Moreover, they do

Figure I – Aggregate Impulse Response Functions

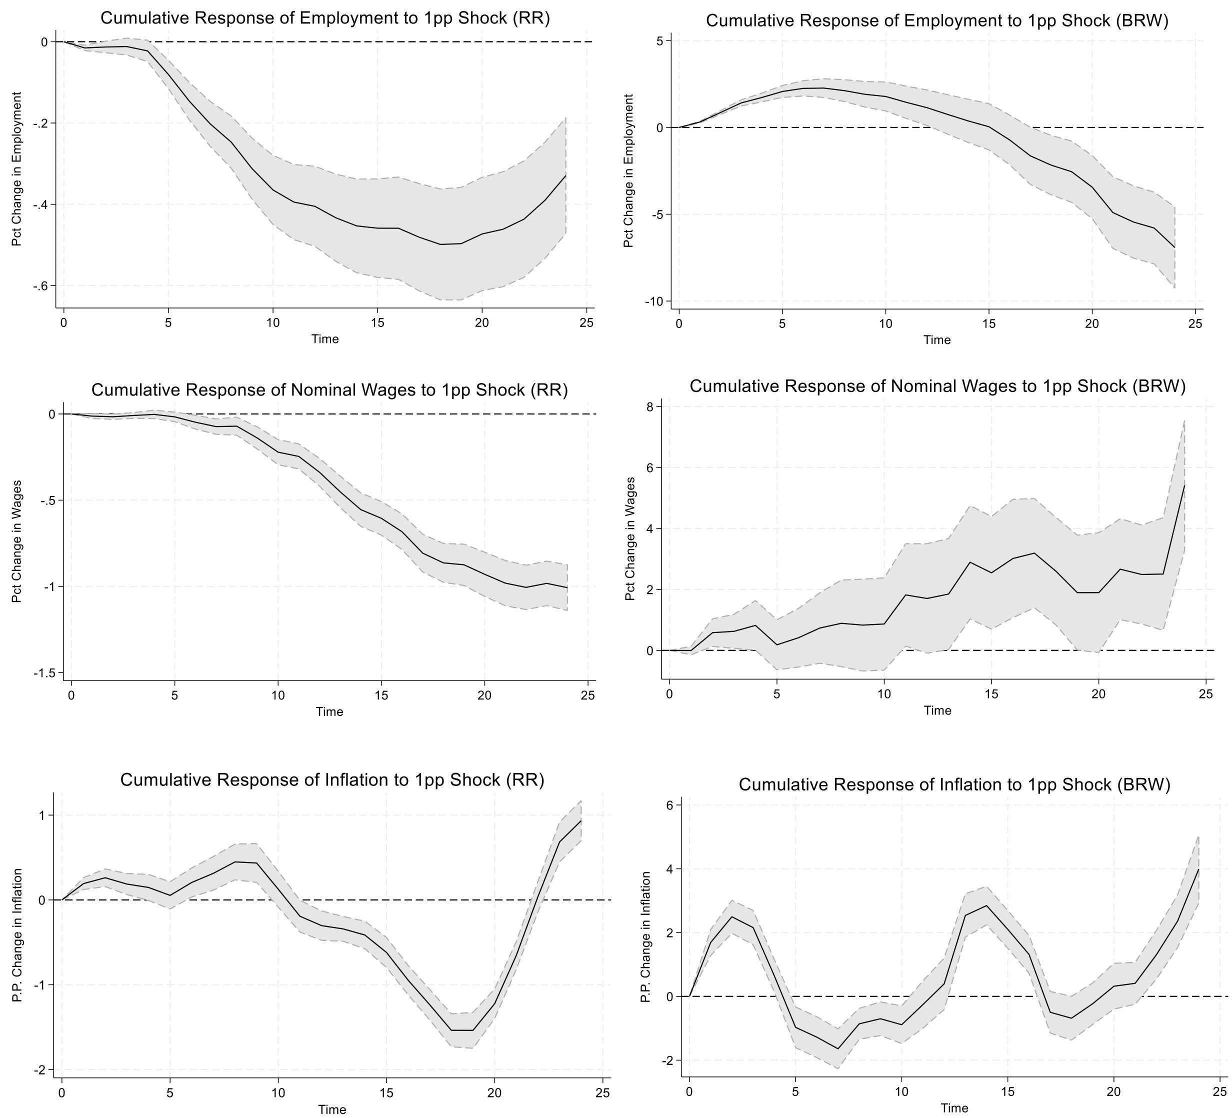

Note: Cumulative impulse responses of adjusted change in employment, adjusted change in nominal wages, and inflation due to the RR shocks (left) and BRW shocks (right). The gray region in each graph represents a 95% confidence interval for the estimate, which is the solid black line. Standard errors are clustered at the state level.

not note seasonally adjusting the data, and their analysis uses 8 lags, rather than the 4 used here. Interestingly, both the employment and wage responses between the two shock series are somewhat mirrored. The RR shock to employment is exclusively negative after initially being insignificant, but the BRW-induced response is at first positive before declining. There is not a highly meaningful response of wages to the BRW shocks—the undulating nature of the changes between zero and a positive suggest as much—but they do trend positive over time, whereas the RR shocks drive a decrease in wages after a few periods. In both cases, the magnitude of the responses between the series are different by a factor of about ten. The explanation likely lies in the aforementioned difference in the magnitudes between the series, though it could also be an effect of the time period; I investigate this possibility next.

Inflation for the Romer and Romer series shares some similarity with the response of inflation estimated by Pedemonte and Herreño (2023) in shape, though the magnitude of the response is lower here. The timing of the major decline in inflation is the same between the two (about 10 quarters). Similarities should be expected, but the divergence is fine, given the differences in the time frame studied and in the underlying inflation data. As before, the BRW shocks are greater in magnitude than the RR series. The shapes of the response are quite different, though: a one percentage point shock to monetary policy is generally inflationary for the BRW series, but mostly negative for the RR shocks.

Aggregate Asymmetries

To test for asymmetries, I estimate equations 1-3 with an added interaction term for the sign of the shock in the series. In this way, the multiplier from the shock is augmented for contractionary shocks. Formally, I estimate:

$$\pi_{i,t+h} = \alpha_i^h + \beta^h s_t + \delta_+^h s_t a_+ + \sum_{j=1}^{J} \gamma_{t-j}\pi_{i,t-j} + \varepsilon_{i,t+h} \;\forall\; h \in [0,H]$$

with δ, the interaction coefficient, subscripted with a "+" to emphasize this is the interaction for contractionary (positive) shocks, and the *a*, an indicator distinguishing data points greater than 0. For wages and employment, I add the same terms into equations to 2 and 3, formally:

$$e_{i,t+h} = \alpha_i^h + \beta^h s_t + \delta_+^h s_t a_+ + \sum_{j=1}^{J} \gamma_{t-j} e_{i,t-j} + \varepsilon_{i,t+h} \ \forall\ h \in [0, H]$$

$$w_{i,t+h} = \alpha_i^h + \beta^h s_t + \delta_+^h s_t a_+ + \sum_{j=1}^{J} \gamma_{t-j} w_{i,t-j} + \varepsilon_{i,t+h} \ \forall\ h \in [0, H]$$

The results of the estimation of these regressions are shown in Figure II. As in Figure I, the graphs show the cumulative impulse responses of wages, employment, and inflation to the BRW and RR shocks. Graphs can be interpreted as showing the cumulative effect on the response variable of a one percentage point unanticipated positive (or negative, depending on the case) monetary shock. To be clear, the graphs showing the effect of a contractionary shock show the interaction coefficient plus the shock coefficient, whereas the graphs depicting the response due to an expansionary shock are only displaying the latter.

Across the graphs, there is no evidence for symmetry between positive and negative monetary shocks. While the response shapes may be mirrored, in some cases, the difference between troughs, peaks, and in significance is stark. For the RR shocks, employment following a tightening initially declines, reaching a low of -1.45%, before the shock dissipates. Conversely, a loosening shock peaks at 0.38%, though generally insignificantly impacts employment. Wages are mostly unaffected by a contraction, though are slightly negatively affected by an expansion. This appears to stem from the fall in prices after an expansionary shock, which itself is unexpected. Since wages here are nominal, a fall in the price level would pull down the nominal wage rate, though the magnitude suggests that real wages may have risen after the expansion. As was the case after an expansion, inflation behaves surprisingly following a contraction as well. Rather than the expected decrease, prices initially increase, but do eventually reach a trough after

18 quarters (the shock dissipates shortly thereafter). Although they are surprising, the inflation responses are also not suggestive of symmetry, and do share some similarities with the loosening/tightening shock estimations done by Tenreyo and Thwaites (2016). Expansionary shocks minimally affect inflation, whereas contractionary shocks negatively affect it. This is imperfectly akin to what I find here: the expansionary shocks are inconsistent and even negative, and the contractionary shocks are slightly negative over time (though their predictions do not rebound).

As was the case for the RR series, the BRW-induced responses do not offer evidence of symmetry. The BRW shocks, however, once again yield responses of a much greater magnitude than the RR series, but at least for employment and wages predict the expected response of the variables. That is, following a contractionary shock, employment and wages decline substantially; following an expansionary shock, they increase substantially. Interestingly, the impulse response functions do not suggest a meaningful dissolution of the shock for these four settings, except for an expansionary shock to wages. Despite the lack of a prediction towards steady state, the responses are still asymmetric. Employment's lowest point following a contraction is about double its highest point subsequent an expansion. Wages are closer to symmetrical, but a quick glance at the graphs reveals that a tightening does decrease the wage rate more than a loosening. The responses of inflation also do not suggest symmetry in the responses, but, akin to the wage and employment responses, are generally in the expected direction. Following an unexpected increase in the interest rate, prices decline slightly, before briefly rising before declining again more substantially after 16 quarters. After an unexpected decrease in the interest rate, the response is generally inflationary.

Time Period Differences

Figure II.A – Asymmetric Romer and Romer Shocks

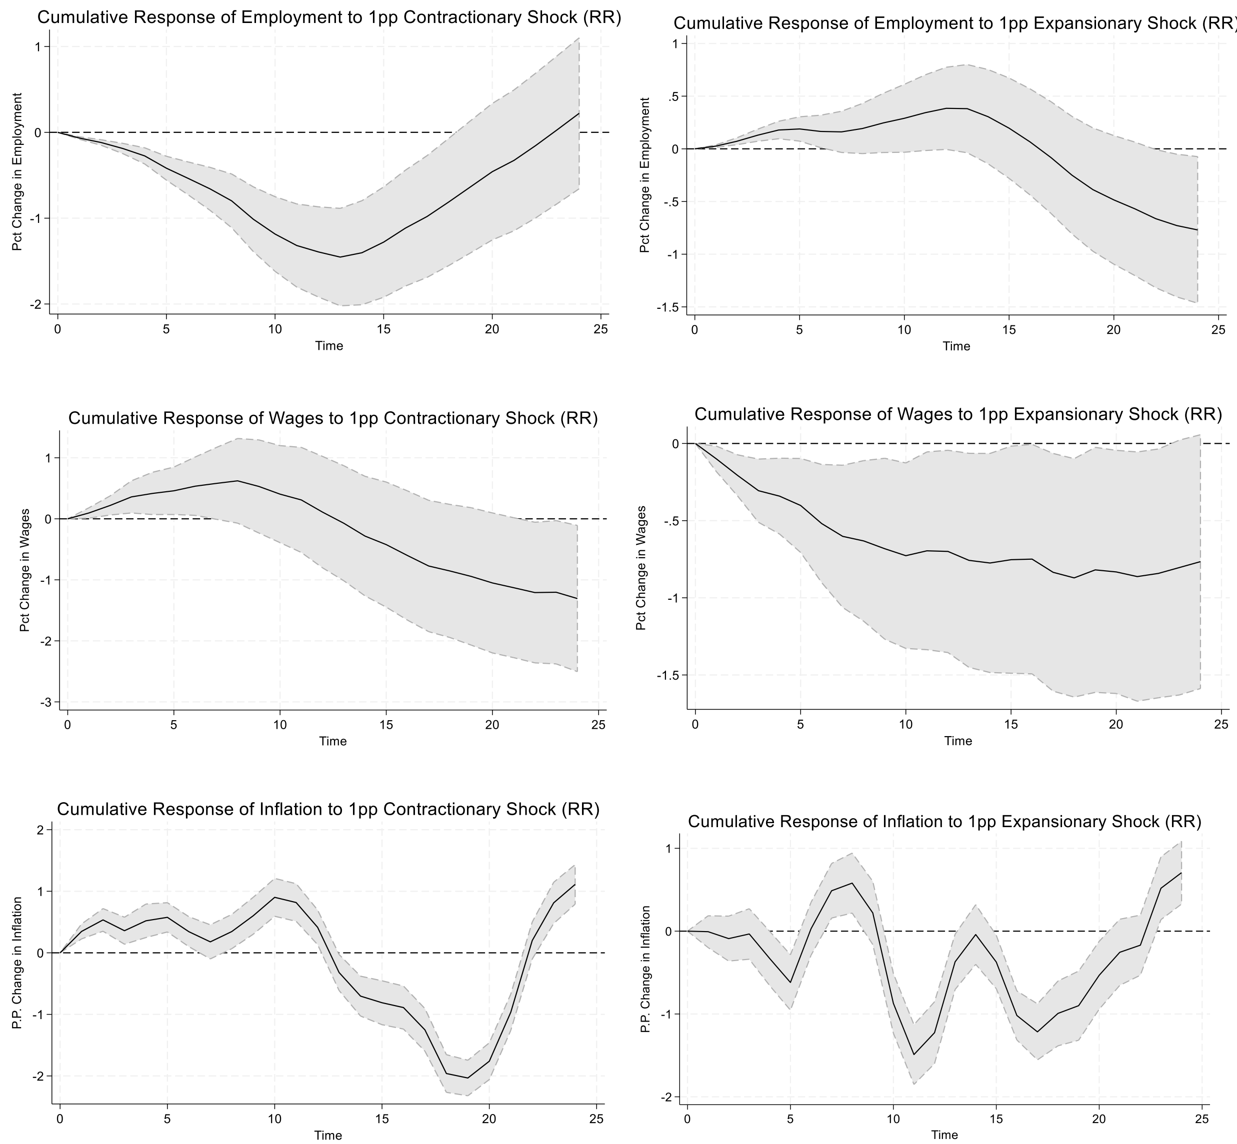

Note: Cumulative impulse response functions for wages, employment, and inflation. Positive interactions are on the left side; negative interactions are on the right side. IRFs for the RR series. The gray area is the 95% confidence interval for the estimate, the thick black line. Standard errors are robust and clustered at the state-level.

The large differences in the size of the responses between the two series begs the question as to whether that is a feature of their construction or a time period effect. While I have discussed the former in the preceding sections, the latter requires further analysis. To investigate, I restrict

Figure II.B – Asymmetric BRW Shocks

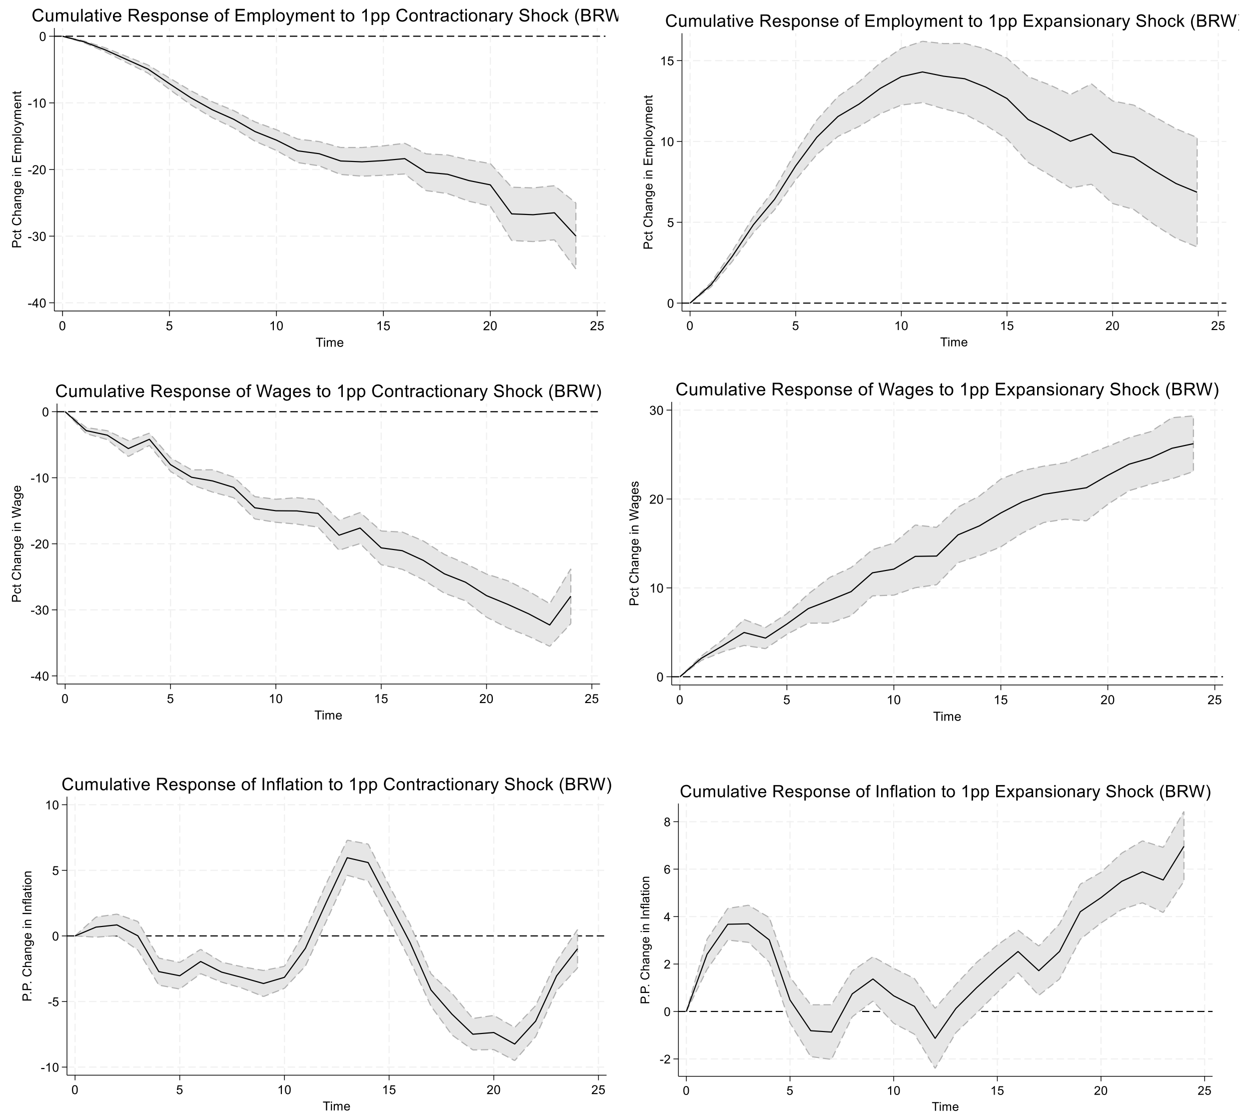

Note: Cumulative impulse response functions for wages, employment, and inflation. Positive interactions are on the left side; negative interactions are on the right side. IRFs for the BRW series. The gray area is the 95% confidence interval for the estimate, the thick black line. Standard errors are robust and clustered at the state-level.

the local projections for the BRW series to before and after 2008 and re-estimate equations 1-3. In the interest of space, the impulse response functions for employment, wages, and prices are shown in Appendix B in Figure B.I.

As is evident, there are certainly differences before and after 2008, though the differences are not related to the size of the responses but rather their shape. Accordingly, the difference in magnitudes between the responses generated by the two series appears to be a feature of their construction. Given the mean for the BRW shocks is about a tenth of the mean of the RR shocks, this makes sense that the coefficients would be much larger. Moreover, given the low correlation between the two series, there does not appear to be a good reason that they *should* be expected to generate similar responses. That they do not do so is thus to be anticipated. This does not mean that one is necessarily better, simply that there are fundamental differences between the two which manifest here.

Household Balance Sheet

Before analyzing the role of the HHBSC in driving heterogeneity, I first estimate its prominence at the aggregate level. Specifically, I estimate the following for prices:

$$\pi_{i,t+h} = \alpha_i^h + \beta^h s_t + \sum_{j=1}^{J} \varphi_{t-j} f_{i,t-j} + \sum_{j=1}^{J} \gamma_{t-j} \pi_{i,t-j} + \varepsilon_{i,t+h} \; \forall \; h \in [0, H]$$

which is equal to equation (1) but with an added control for the household balance sheet, the coefficient for which is φ and the variable for which is f. The metric I use to estimate the household balance sheet is the change in each state's log dividend, interest, and rent income from the previous period. I then multiply it by a hundred so that it is the percent change, rather than a decimal change. I estimate analogous local projections for employment and wages, shown below, with H = 24 and J = 4, the same specifications I had previously used.

$$e_{i,t+h} = \alpha_i^h + \beta^h s_t + \sum_{j=1}^{J} \varphi_{t-j} f_{i,t-j} + \sum_{j=1}^{J} \gamma_{t-j} e_{i,t-j} + \varepsilon_{i,t+h} \; \forall \; h \in [0, H]$$

$$w_{i,t+h} = \alpha_i^h + \beta^h s_t + \sum_{j=1}^{J} \varphi_{t-j} f_{i,t-j} + \sum_{j=1}^{J} \gamma_{t-j} w_{i,t-j} + \varepsilon_{i,t+h} \; \forall \; h \in [0, H]$$

The cumulative impulse response functions are graphed in Figure III. As before, each graph may be interpreted the response of the respective dependent variable to a one percentage point unanticipated shock to the interest rate.

After controlling for the household balance sheet, the impact of the shock is marginally different, though essentially the same. For example, the employment response to the RR shocks is slightly deeper after controlling for the HHBS (though not significantly). In the aggregate, this suggests the role of the household balance sheet is minimal at best. It does not, however, mean that it cannot extend as an explanation to the states. Indeed, it makes sense that a channel which may be spatially heterogeneous but not nationally would not massively impact the effect of monetary policy. Thus, while we see a marginal role for the HHBSC in the aggregate, this does not preclude its role regionally.

Spatial Heterogeneity

To analyze the behavior of employment, wages, and prices spatially, I estimate local projections for each state. Essentially, I estimate equations 1-3 for each state separately, using robust standard errors and controlling for four lags of the dependent variable. The cumulative responses are rather noisy, likely in part due to the limitations posed on the sample size due to the start of the QCEW and inflation data, respectively, and the starts and ends of the two shock series: the highest degrees of freedom (which is for the RR shocks) is only 121. Accordingly, the graphs for the impulse responses are omitted. Nevertheless, each state exhibits individualized responses, some of which are in-line with aggregate trends, though others are not. To this end, I

Figure III – HHBSC-controlled Impulse Response Functions

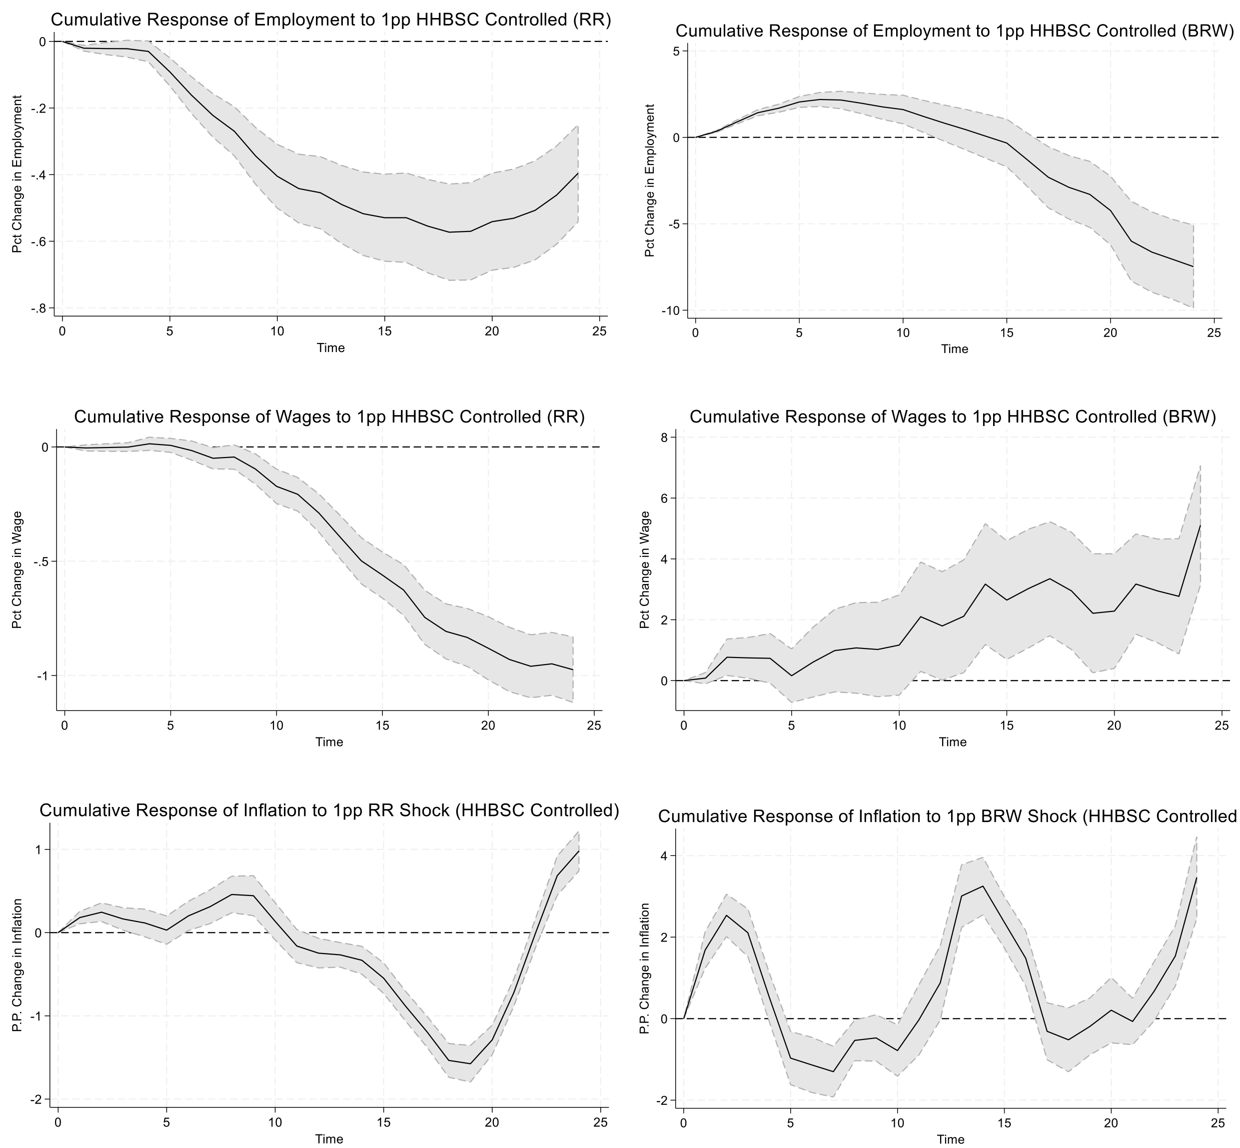

Note: cumulative impulse responses of wages, employment, and CPI to the combined coefficient for the shocks (RR shocks on the left, BRW shocks on the right) and the interaction with relative dividend, interest, and rent income per capita. The gray area bound by the dashed lines represents the 95% confidence interval for the coefficient, shown along the solid black line.

graph histograms of each set of variables at the most extreme point in the aggregate responses. Thus, the RR employment responses at the state level are shown at the 18th quarter, RR wages at the 21st, RR inflation at the 21st, BRW employment at the 24th, BRW wages at the 18th, and inflation CPI at the 24th. These are graphed in Figure IV.

Figure IV – Distribution of the State-Level Coefficients of Interest

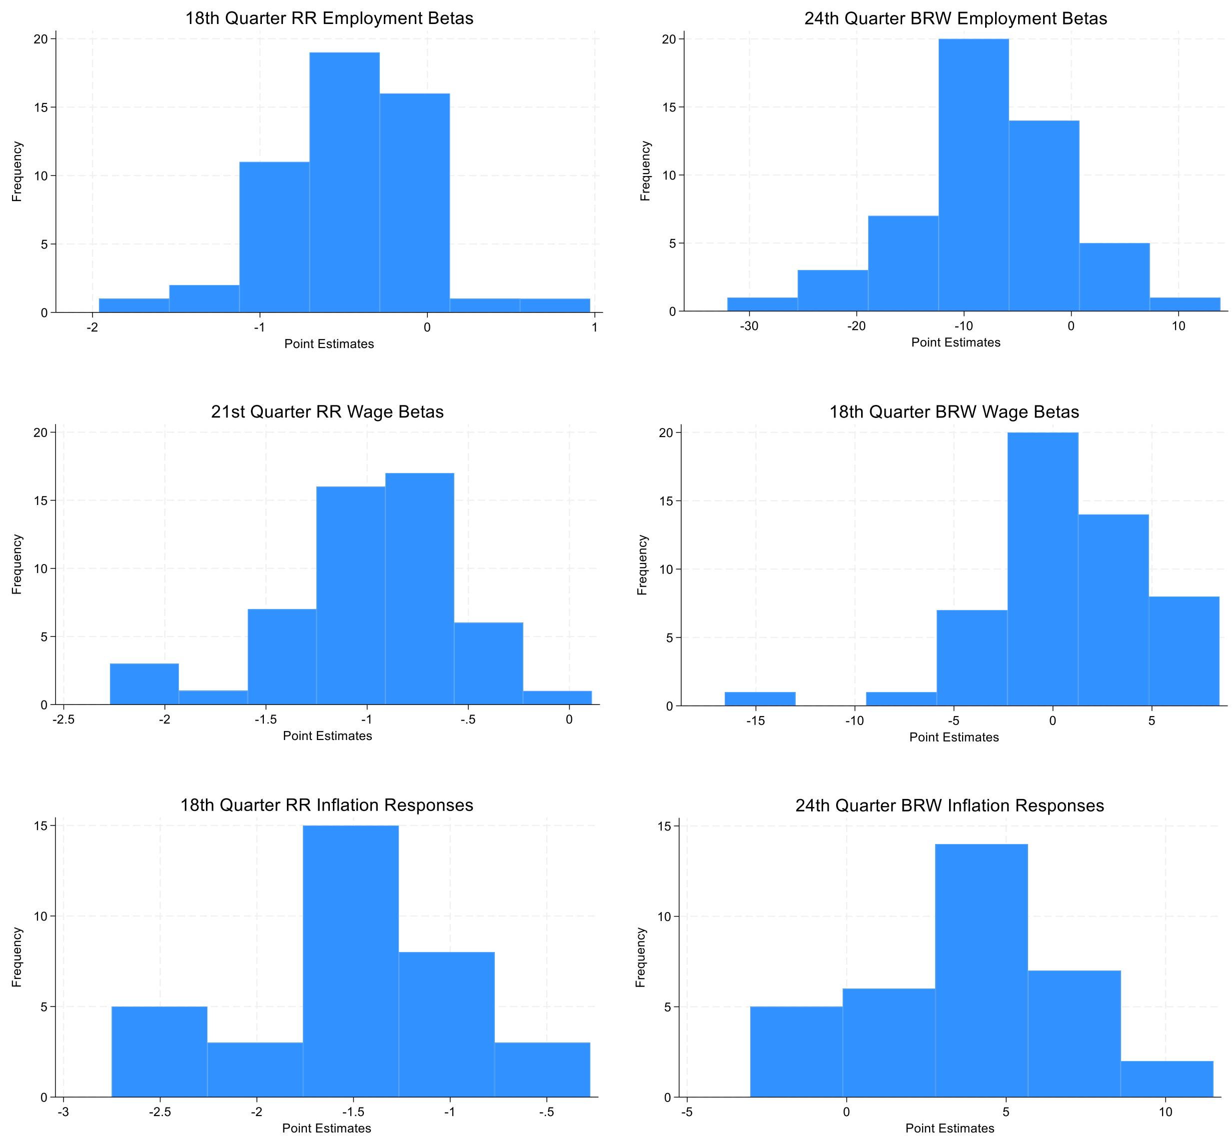

As is evident in the histograms in Figure IV, most states are clustered close together, but several states display extreme responses. While they are nonetheless close, the existence of any heterogeneity is noteworthy. As previously noted, several channels may account for this heterogeneity, such as the interest rate or credit channels or local fiscal policy and population density (Carlino and Defina 1998; Francis, Owyang, and Sekhpoysan 2012). Given this channel has been well-researched, however, I save my analysis including the interest rate channel for

Appendix C[4]. Instead, I focus on the less-investigated household balance sheet channel's role in impacting heterogeneity at the state level.

Sources of Heterogeneity and Asymmetry

If the HHBSC indeed produces heterogeneity across the states, then controlling for it in state-level local projections should drive state-level responses towards homogeneity. To test this expectation, I again re-estimate equations 1-3 for each state, rather than as a panel; standard errors are again robust. The difference between this analysis and the antecedent is that I add four quarters of lagged relative dividend, interest, and rent income per capita (DIR) as a control. The "relative" nature of this data is that each state's level at time t is a fraction of the state's total relative to the (unweighted) average of all states at time t. Whereas using this in the panel specification earlier may have caused issues as the average every quarter is 1, estimating each state separately allows its usage without concern. Further, while a change in the DIR may be relevant to individuals at a given time, a greater relative balance sheet suggests that there is a better connectedness to the financial economy. Accordingly, even though the specific levels may fluctuate, if a state is relatively better connected to the financial economy, individuals may be more likely to adjust their expenditures, and thus impact the real economy, in line with Mishkin (1996).

Comparing histograms visually would not be particularly illusory. Instead, I take the kurtosis of the 51 (or 34, for inflation) coefficient estimates from the above local projections at the previous horizons, comparing this to their kurtosis with the HHBSC controlled for. The specific values are contained in Table I.A. In addition to kurtosis, I also look at standard

[4] It is worth noting here that the effects of the HHBSC are different when the interest rate channel is accounted for. See Appendix C for a more careful discussion.

deviation, the values of which are contained in table I.B. The choice of kurtosis and standard deviation is because they reveal whether the HHBSC would help to alleviate extreme values in the former case, and if it would homogenize the entire distribution in the latter case.

As is evident, kurtosis of the shock coefficients decreased exclusively for the employment responses and for the RR-induced inflation responses, but not for other variables. The kurtosis for wages increased quite considerably, suggesting that the household balance sheet deepens spatial heterogeneity here, rather than homogenizing it. Employment, as mentioned, does become less extreme. This finding is underscored by the standard deviation for the distribution, which decreases after controlling for the HHBSC for employment. This suggests that the HHBSC both moderates the response of the state level employment to monetary shocks in addition to homogenizing it. The spread of the other four distributions increases, again furthering the notion that the household balance sheet drives the state-level responses of wages, and at best only modestly moderates inflation.

To investigate if the effect of the HHBSC extends beyond heterogeneity to asymmetry, I estimate the asymmetry equations (4-6) for each state. I omit any IRFs or descriptive charts for this estimation, and instead proceed to analyze the source of any differences in responses. In short, I attempt to determine if the HHBSC is not only a factor in spatial heterogeneity, but also in the asymmetry of expansionary and contractionary shocks.

Following the approach Carlino and Defina (1998) in analyzing spatial heterogeneity, I take the coefficients for each state and regress them on each state's average relative DIR for the time frame. Due to the reduced number of states in the inflation data, I limit this analysis to only employment and wage responses. The point estimates from these regressions can be seen as the

Table I.A – Kurtosis of the Coefficient Distributions

	HHBSC Not Controlled	HHBSC Controlled
RR Employment Betas	4.98	4.37
RR Wage Betas	3.80	4.81
RR Inflation Betas	2.86	2.66
BRW Employment Betas	4.31	3.01
BRW Wage Betas	7.71	9.87
BRW Inflation Betas	2.83	2.94

Table I.B Standard Deviation of the Coefficient Distributions

	HHBSC Not Controlled	HHBSC Controlled
RR Employment Betas	0.472	0.470
RR Wage Betas	0.467	0.616
RR Inflation Betas	0.60	0.66
BRW Employment Betas	7.87	7.47
BRW Wage Betas	4.08	4.13
BRW Inflation Betas	3.26	3.55

amount of change created in the response of the states to the monetary shocks driven by the household balance sheet channel. In this way, any significance in these coefficients is evidence that the household balance sheet channel is involved in creating asymmetry between expansionary and contractionary shocks. The results of the estimation of these regressions are shown in Figure V, with the positive (contractionary) shocks shown in Figure V.A and their expansionary counterparts in Figure V.B. I plot the point estimates (red dots) with the 95% confidence interval at each horizon, the upper and lower bounds of which are denoted by the blue dots.

Between expansions and contractions, the role of the household balance sheet is mixed. Heterogeneity amongst the RR-induced employment responses is partially explained by the household balance sheet for both expansionary and contractionary shocks. The only other responses to be significantly impacted by the HHBSC are the expansionary RR-driven wages and the contractionary BRW-driven wages. The BRW-induced employment responses are unaffected by the household balance sheet. In this way, this reaffirms the general heterogeneity findings—that the HHBSC affects differences in employment—but does not do so by a huge amount, nor does it highly impact wages.

In addition to the general spatial heterogeneity analysis, Appendix C also adds in the interest rate channel to this analysis by controlling for the percent of employment in manufacturing, following Carlino and Defina (1998), plus the percent of employment in construction (often leveraged and additionally cyclical). The graphs are shown in Appendix figure C.II. They tell the same story about the household balance sheet channel as the graphs in figure V do, that it mostly affects the employment response (from the RR-series) and inconsistently explains asymmetry in wages. Interestingly, the RR-wages are slightly explained for both expansionary and contractionary shocks, suggesting the interest rate channel is a useful control.

Figure V.A – Expansionary Shocks

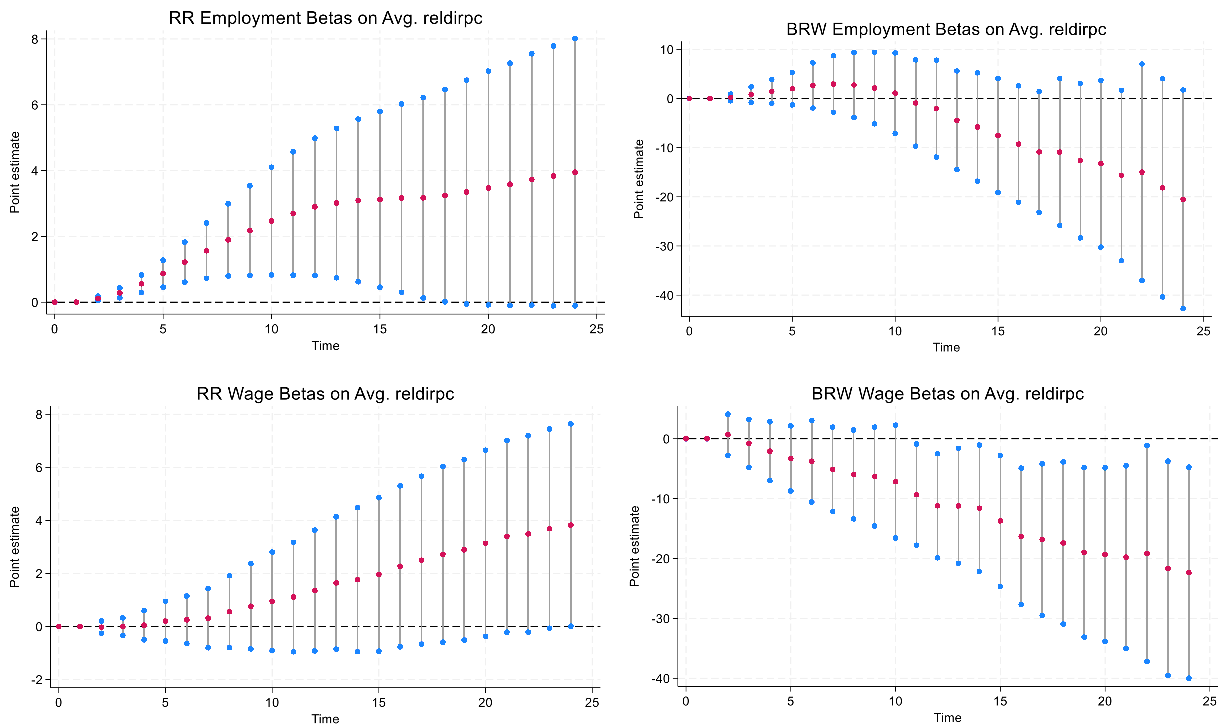

Note: impact of the HHBSC on the coefficients of the state-level regressions. The bands are the black line, the point estimate of the impact variable, are the 95% confidence interval of the estimators. All responses are to contractionary shocks.

Figure V.B – Contractionary Shocks

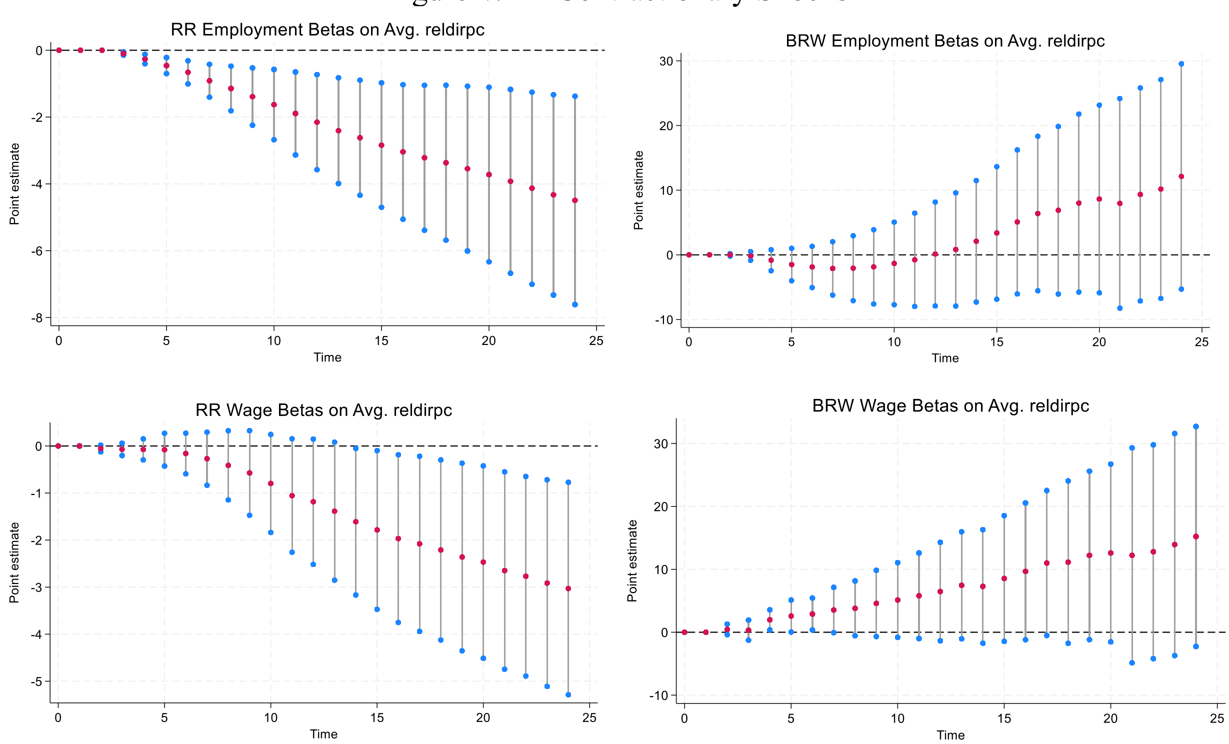

Note: impact of the HHBSC on the coefficients of the state-level regressions. The bands are the black line, the point estimate of the impact variable, are the 95% confidence interval of the estimators. All responses are to expansionary shocks.

Discussion

The household balance sheet has been shown in many previous works to affect recessions, aggregate demand, and other facets of the macroeconomy (see, for example, Mishkin [1978] or Mian and Sufi [2010; 2014]). To the best of my knowledge, this paper is the first to document the household balance sheet as a means of moderating state-level employment responses to monetary shocks. It is interesting that this explanatory power of the household balance sheet is limited to employment, and the opposite holds for wages and inflation. It is possible that the change in the increased heterogeneity in the nominal wage stems from inflation, as the real wage for each state is unobserved (no sufficient state-level deflators exist for the time series). Thus, as the household balance sheet ultimately factors into aggregate demand through consumption of durables (Mishkin 1996), employment can feasibly be growing more homogenous while the nominal wage and inflation are growing apart. Because the household balance sheet channel captures investment, changes in capital deriving from these investments would vary by state. Since aggregate demand is affected by investment which affects employment, then employment could be the only of the three impacted by the household balance sheet. Generally, this finding is in accordance with the aggregate demand channel of employment documented by Mian and Sufi (2014), where employment growth is negatively predicted by the household financial positions. Their findings on wages are different than those found here, but they do find the employment impact to be most noteworthy, as is the case here.

The findings that the household balance sheet channel plays into asymmetry as well as heterogeneity are a new flavor of existing findings in the literature. Mishkin (1978) notes that a declined household balance sheet dragged down aggregate demand during the Great Depression, and Mian and Sufi (2010; 2014) found a similar relationship for the Great Recession. Following the same logic as to why the household balance sheet only had visible impacts on employment responses to monetary shocks across the states, that effect of the household balance sheet on driving asymmetrical employment responses is then in line with these findings in the literature. That is, this relationship between employment responses and the household balance sheet is to be expected. Nonetheless, the evidence that the household balance sheet is partially responsible for asymmetry itself is novel, as it had not, to my knowledge, been considered asymmetrically, but rather as a mechanism that might deepen recessions.

Another potential source of asymmetry, seen in the literature but unaccounted for in this analysis, is rigidity in the nominal wage. If workers resist reductions in the nominal wage, then the employment decline stemming from a contractionary policy would be relatively greater than that stemming from an expansionary policy. Wage rigidity has been found to be a driving force of business cycle asymmetries in previous research (Abritti and Fahr 2013). Although not tested here, the asymmetry found in this research lends itself to their results. Moreover, if downwards wage rigidity is indeed a major source of asymmetric response, then the results (or lack thereof) from the wage and employment analyses lend further credence to the role of wage stickiness.

IV. Conclusion

Using local projections to develop impulse response functions for the effect of monetary shocks on employment, nominal wages, and inflation, I find several novel results. Foremost, I confirm that expansionary and contractionary shocks are indeed asymmetrical. Considering the difference between peaks and troughs, employment is deeper after a contractionary shock than it is higher after an expansionary shock by nearly four-fold for the RR series, and more than double for the BRW series. Despite the responses of inflation and wages being less easily quantifiable, asymmetry is still present in all cases. Second, I find a homogenizing role for the household balance sheet in the different responses of employment across the US states. While the household balance sheet does appear to drive the responses of inflation to be more heterogeneous, the effect on the nominal wage is necessarily uncertain. Third, as could be reasonably predicted by past empirical findings (Mian and Sufi 2010; Mian and Sufi 2014; Mishkin 1978), I confirm that, especially for employment, the household balance sheet functions to serve as a source of heterogeneity across state level responses to monetary shocks.

The biggest implication of this paper on policy and future research is the need to attend to both expansionary and contractionary shocks. Because regression techniques largely impose symmetry, ignoring the enhanced effects of a contraction relative to an expansion could have adverse effects for employment, which is obviously crucial to individuals' livelihoods. Furthermore, the documentation of the household balance sheet as a source of heterogeneity for employment can better equip leaders to act proactively subsequent monetary shocks.

References

Abbritti, Mirko, and Stephan Fahr. 2013. "Downward Wage Rigidity and Business Cycle Asymmetries." *Journal of Monetary Economics* 60 (7): 871–86. https://doi.org/10.1016/j.jmoneco.2013.08.001.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6): 2333–67. https://doi.org/10.1257/aer.20160137.

Bergman, Nittai, David A. Matsa, and Michael Weber. 2022. "Inclusive Monetary Policy: How Tight Labor Markets Facilitate Broad-Based Employment Growth." SSRN Scholarly Paper. Rochester, NY. https://doi.org/10.2139/ssrn.4007639.

Bu, Chunya, John Rogers, and Wenbin Wu. 2021. "A Unified Measure of Fed Monetary Policy Shocks." *Journal of Monetary Economics* 118 (March): 331–49. https://doi.org/10.1016/j.jmoneco.2020.11.002.

Canova, Fabio, and Evi Pappa. 2007. "Price Differentials in Monetary Unions: The Role of Fiscal Shocks." *The Economic Journal* 117 (520): 713–37. https://doi.org/10.1111/j.1468-0297.2007.02047.x.

Carlino, Gerald, and Robert DeFina. 1998. "The Differential Regional Effects of Monetary Policy." *Review of Economics and Statistics* 80 (4): 572–87. https://doi.org/10.1162/003465398557843.

Carlino, Gerald and Robert DeFina. 1999. "Do states respond differently to changes in monetary policy?" Business Review, Federal Reserve Bank of Philadelphia, issue Jul, 17-27.

Chodorow-Reich, Gabriel, Plamen T. Nenov, and Alp Simsek. 2021. "Stock Market Wealth and the Real Economy: A Local Labor Market Approach." *American Economic Review* 111 (5): 1613–57. https://doi.org/10.1257/aer.20200208.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1): 1–45. https://doi.org/10.1086/426038.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism." *The Review of Economic Studies* 87 (1): 102–29. https://doi.org/10.1093/restud/rdy074.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary Policy and Inequality." *Journal of Monetary Economics* 88 (June): 70–89. https://doi.org/10.1016/j.jmoneco.2017.05.005.

Cover, J. P. 1992. "Asymmetric Effects of Positive and Negative Money-Supply Shocks." *The Quarterly Journal of Economics* 107 (4): 1261–82. https://doi.org/10.2307/2118388.

Francis, Neville, Michael Owyang, and Tatevik Sekhposyan. 2012. "The Local Effects of Monetary Policy." *The B.E. Journal of Macroeconomics* 12 (2): 1–38. https://econpapers.repec.org/article/bpjbejmac/v_3a12_3ay_3a2012_3ai_3a2_3an_3a3.htm.

Hanson, Michael S., Erik Hurst, and Ki Young Park. 2006. "Does Monetary Policy Help Least Those Who Need It Most?" *Wesleyan Economics Working Papers*, Wesleyan Economics Working Papers. https://ideas.repec.org//p/wes/weswpa/2006-006.html.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson. 2022. "The Slope of the Phillips Curve: Evidence from U.S. States." *The Quarterly Journal of Economics* 137 (3): 1299–1344. https://doi.org/10.1093/qje/qjac010.

Mian, Atif, and Amir Sufi. 2010. "Household Leverage and the Recession of 2007–09." *IMF Economic Review* 58 (1): 74–117. https://www.jstor.org/stable/25762071.

Mian, Atif, and Amir Sufi. 2014. "What Explains the 2007-2009 Drop in Employment?" *Econometrica* 82 (6): 2197–2223. https://www.jstor.org/stable/43616911.

Mihov, Ilian, and Andrew Scott. 2001. "Monetary Policy Implementation and Transmission in the European Monetary Union." *Economic Policy* 16 (33): 371–406. https://www.jstor.org/stable/1344646.

Mishkin, Frederic S. 1978. "The Household Balance Sheet and the Great Depression." *The Journal of Economic History* 38 (4): 918–37. https://www.jstor.org/stable/2118664.

Mishkin, Frederic S. 1996. "The Channels of Monetary Transmission: Lessons for Monetary Policy." Working Paper. Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w5464.

Montiel Olea, José Luis, and Mikkel Plagborg-Møller. 2021. "Local Projection Inference Is Simpler and More Robust Than You Think." *Econometrica* 89 (4): 1789–1823. https://doi.org/10.3982/ECTA18756.

Owyang, Michael T., Jeremy Piger, and Howard J. Wall. 2005. "Business Cycle Phases in U.S. States." *The Review of Economics and Statistics* 87 (4): 604–16. https://www.jstor.org/stable/40042879.

Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94 (4): 1055–84. https://doi.org/10.1257/0002828042002651.

Tenreyro, Silvana, and Gregory Thwaites. 2016. "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions." *American Economic Journal: Macroeconomics* 8 (4): 43–74. https://www.jstor.org/stable/26156440.

U.S. Bureau of Economic Analysis. 2023. "SQINC4 Personal income by major component."

U.S. Bureau of Labor Statistics. 2023. "Quarterly Census of Wages and Employment."

Wieland, Johannes F., and Mu-Jeung Yang. 2020. "Financial Dampening." *Journal of Money, Credit and Banking* 52 (1): 79–113. https://doi.org/10.1111/jmcb.12681.

Wong, Arlene. 2021. "Refinancing and The Transmission of Monetary Policy to Consumption." *Princeton University Economics Department,* Working Papers 2021-57. https://ideas.repec.org//p/pri/econom/2021-57.html.

Appendix A – Supplementary Graphs

Figure A.I – Baseline and Smoothed Employment and Wages

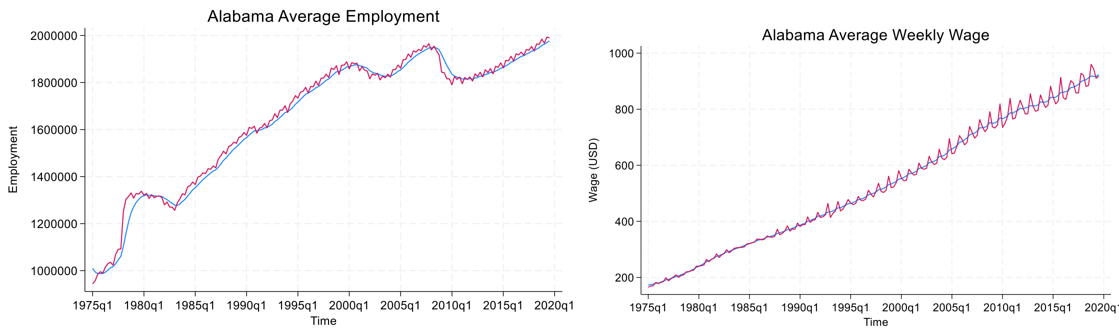

Note: Adjusted data is shown in blue, original data in red. Average quarterly weekly wage and employment data are shown over the time frame.

Figure A.II – Baseline and Comparative Smoothing Methods for Select States

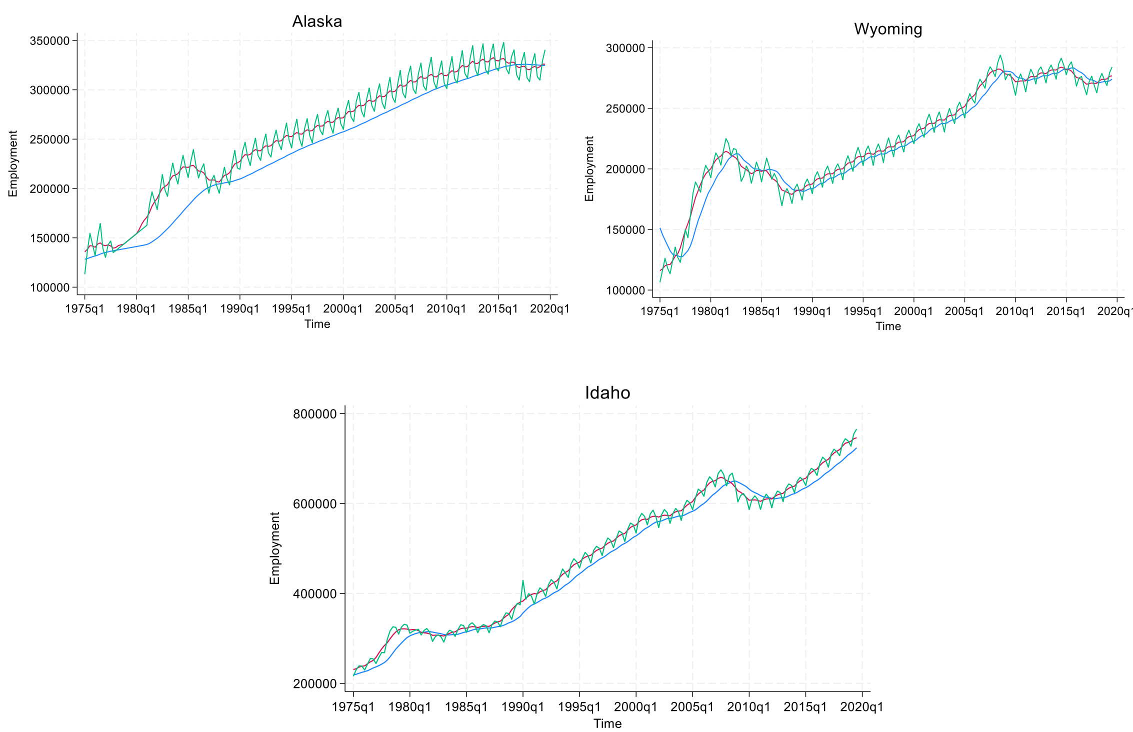

Note: the red line is the moving average; the blue line is the double exponential smoothing; the green line is the interpolated (for Alaska) or original data shape. Average quarterly employment is shown for each state over the time frame across these three metrics.

Appendix B – Time Period Analysis Functions

Figure B.I – Pre-Post 2008 Differences for BRW Shocks

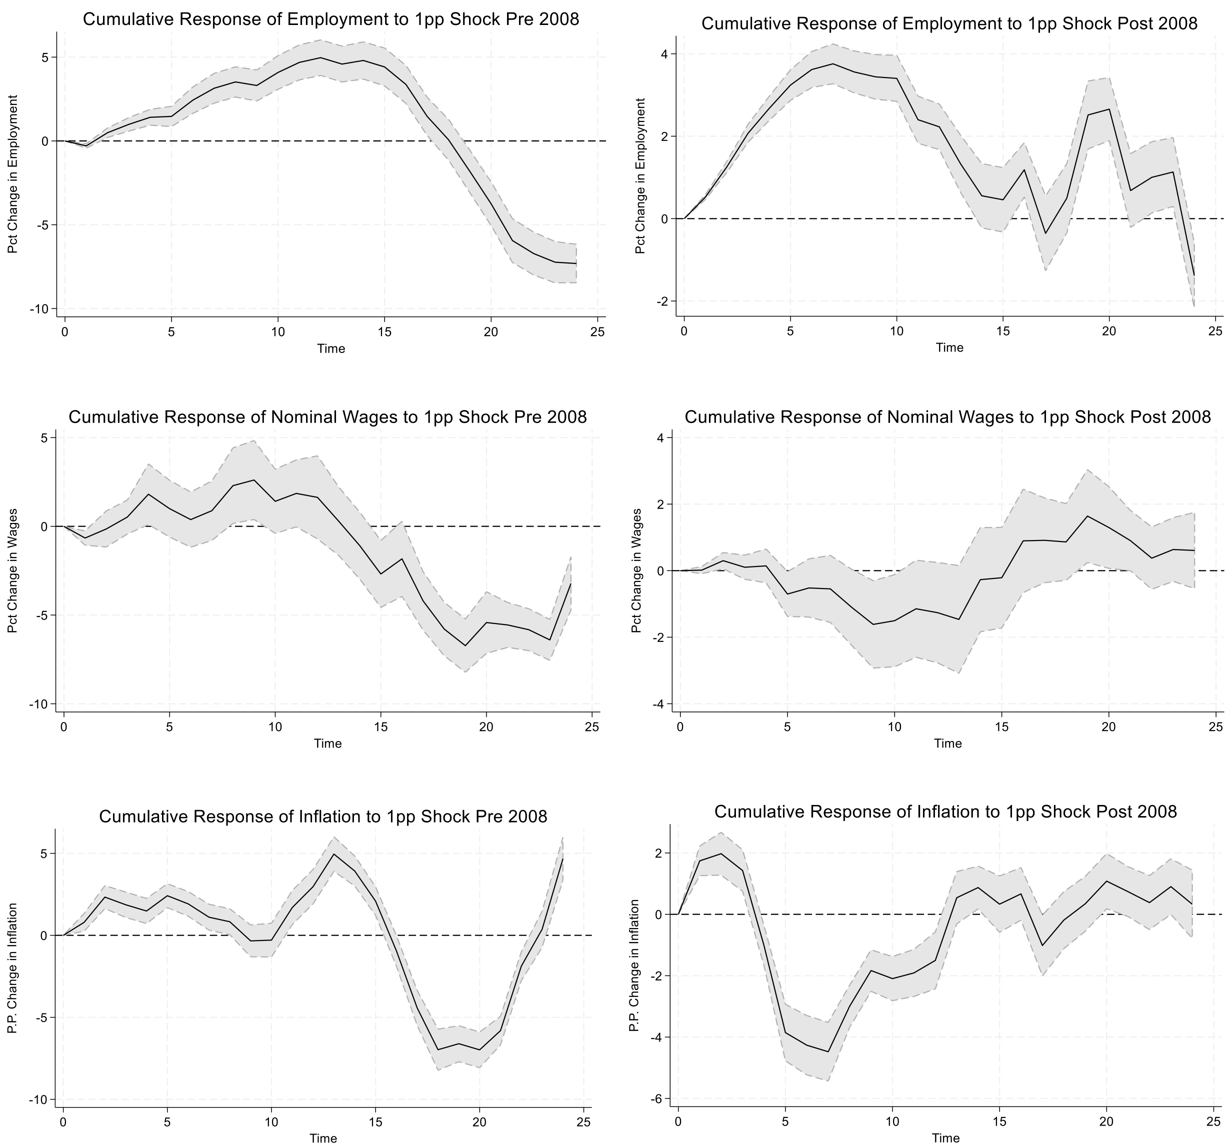

Note: graphs show the cumulative responses of employment (top), wages (middle), and CPI (bottom), separated by whether the BRW shock data is before or after 2008, when the RR shocks end. The black line is the estimate of the response at each horizon, with the surrounding area its 95% confidence interval. Standard errors are robust and clustered at the state level.

Appendix C – Interest Rate Channel Extensions

In Figure C.I, I recreate the histograms of figure IV, but with 4 quarters of lagged relative DIR income and of the construction and manufacturing percentage. I then recreate Table I in Table C.I to show how the interest rate channel (IRC) affects the kurtosis and standard deviation at these horizons.

A few things are surprising when comparing Figure C.I to figure IV. First, the Romer shocks are driven further apart after accounting for the household balance sheet and the interest rate channel, and the BRW shocks more homogenous. As is visible in Table C.I, the increase in homogeneity is considerable for all the BRW-induced effects. The rationale for why these responses have manifested in this way is unclear, and my analysis is not such that I can provide a full explanation. One possibility is, however, that the period selection is inapt for the states. Looking more granularly, though, would not be concise, and I do not believe another approach would be unarbitrary. That the heterogeneity is not solved by balance sheet and interest rate channel is consistent with Carlino and Defina (1998), as they find only so much explanatory power stemming from their selected channels. For the BRW shocks, however, the increased homogeneity suggests a strong role for the interest rate channel.

In figure C.II, I recreate figure V, but add the average time period percent of construction and manufacturing per state as an additional control. Figure C.II.A show the point estimates with 95% confidence intervals for the listed combination of shocks and dependent variables regressed at each horizon on each states average relative DIR income and average percent of construction and manufacturing employment. The responses are similar to figure V, but we do see slightly more asymmetry explained for wages (the RR-induced expansionary results are predicted by average relative DIR income in later quarters when the interest rate channel is controlled for.

Otherwise, the differences are rather unremarkable, though that the explanatory power does not vanish suggests a role for both channels.

Figure C.I – Distribution of the State-level Coefficient, Interest Rate Channel Controlled

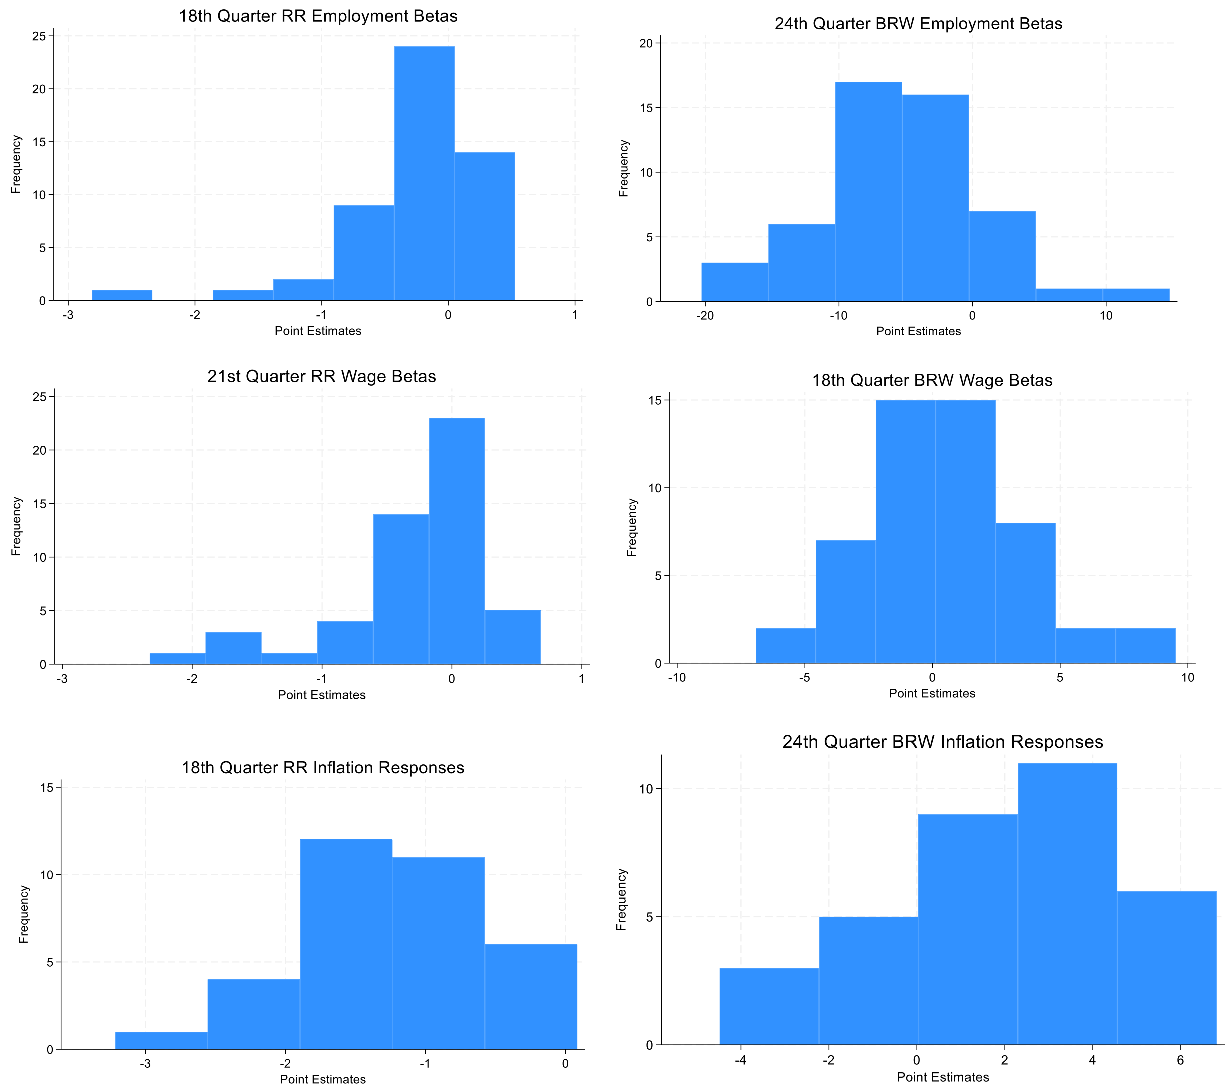

Table C.I –Distribution Statistics with Interest Rate Channel Controlled

Coefficients	Kurtosis	Standard Deviation
RR - Employment	9.85	0.57
RR - Wages	5.47	0.58
RR – Inflation	3.41	0.70
BRW - Employment	4.02	6.41

BRW - Wages	3.71	3.12
BRW - Inflation	2.42	2.91

Figure C.II.A – Expansionary Shocks (IRC Controlled)

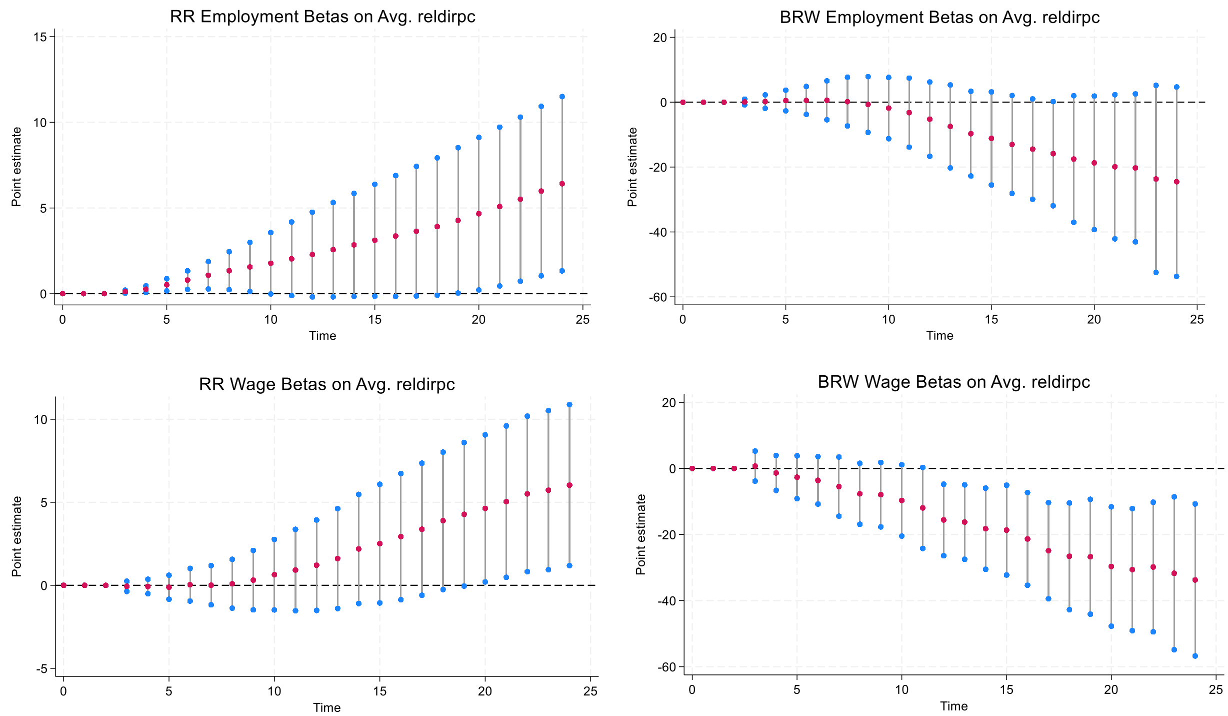

Figure C.II.B – Contractionary Shocks (IRC Controlled)