Pushing on a string: US monetary policy is less powerful in recessions[*]

Silvana Tenreyro[a] and Gregory Thwaites[b]

[a] LSE, CEP, CfM; [b]LSE, CfM

April 2013

Abstract

We estimate the impulse response of key US macro series to the monetary policy shocks identified by Romer and Romer (2004), allowing the response to depend flexibly on the state of the business cycle. We find strong evidence that the effects of monetary policy on real and nominal variables are more powerful in expansions than in recessions. The magnitude of the difference is particularly large in durables expenditure and business investment. The effect is not attributable to differences in the response of fiscal variables or the external finance premium. We find some evidence that contractionary policy shocks have more powerful effects than expansionary shocks. But contractionary shocks have not been more common in booms, so this asymmetry cannot explain our main finding.

JEL classifications: E52, E32

Keywords: asymmetric effects of monetary policy, transmission mechanism, recession, durable goods, local projection methods.

1

1 Introduction

Is monetary policy more powerful in expansions or in recessions? A priori, the answer is unclear. Expenditure could be more or less sensitive to real interest rates at different points in the business cycle. Imperfections in the financial system might magnify or dampen the transmission of policy at different times. Prices might be more or less sticky. And the systematic component of monetary policy itself might behave differently. Previous work has addressed this question, and adjacent ones, both theoretically and empirically. The results have been mixed.

We investigate this question anew on US data, and find strong evidence that monetary policy shocks typically have much more powerful effects on output and inflation in an expansion than in a recession. We follow Auerbach and Gorodnichenko (2011) in adapting the local projection method of Jordà (2005) with the smooth transition regression method of Granger and Terasvirta (1994), to allow impulse response functions to depend on the state of the business cycle. We investigate the state-dependence of monetary policy impulse response functions in this framework, examining the response of a range of real and nominal variables to the estimate of monetary policy shocks introduced by Romer and Romer (2004).

The main result from our investigation is that shocks to the federal funds rate are more powerful in expansions than in recessions.[1] Nearly all of the effect we observe on average in the data is attributable to the effect in good times, and in particular to the response of durable consumption and business and household investment. In an expansion, output and then inflation fall in response to a negative monetary shock in the textbook fashion. Within this, and in line with previous findings, business investment and consumer expenditure on durable goods and housing are an order of magnitude more sensitive than other expenditures, whereas the responses of durables and nondurables prices are much closer together. In a recession, in contrast, the response of output and inflation to monetary policy interventions is negligible. These differences are not attributable to differences in the amplification afforded by

---

[1]Auerbach and Gorodnichenko (2012) find that fiscal policy is instead more powerful in recessions than in expansions; we find the opposite for monetary policy.

the response of credit prices or quantities, nor any systematic differences in fiscal policy across regimes. We find that contractionary shocks are more powerful than expansionary shocks, but given that they are equally common in both expansions and recessions, this cannot be the source of asymmetry across the business cycle.

These findings have important implications for theories of the business cycle, price setting and monetary policy transmission, and also for the design of stabilisation policy and the models used to analyse it. If changes in the policy rate have little impact in a recession, policymakers will possibly need to resort to other monetary policy measures to achieve the desired expansionary effect. The authorities may also need to rely more heavily on fiscal or financial policies to stabilize the economy in a deep or protracted slump.

The remainder of this paper is structured as follows. Section 2 reviews the literature. Section 3 explains the empirical method and the dataset. Section 4 sets out the main results and sensitivity analysis. Section 5 concludes with some thoughts for future research.

2 Literature

There is a small empirical literature on how the impact of monetary policy varies with the business cycle. This research has produced mixed results and, perhaps as a result, the mainstream monetary policy literature, both theoretical and empirical, has largely ignored the potential for asymmetries and their policy implications. See for example Christiano et al. (2005), Woodford (2002) and Gali (2008).

The most cited paper in this literature that is perhaps closest to ours in implementation is Weise (1999). Weise (1999) estimates regime-dependency with a smooth-transition technique (Granger and Terasvirta (1994)), as do we, but applies this to a VAR rather than a local projection model. The VAR contains industrial production, consumer prices and the M1, detrended in complicated piecewise fashion over 1960Q2-1995Q2. Monetary shocks are identified with Choleski orthogonalisation, putting money last. The regime is indicated by the first lag of quarterly GDP growth, such that high-frequency shifts in regime are possible. As with other VAR-

based regime-switching models (and in contrast to the local projection model we employ), the researcher must decide how to account for the possibility that a shock causes a shift in regime. In this case, impulse response functions are calculated as the difference between two stochastic simulations with different initial conditions for output.

Taken together, the results in this paper are difficult to interpret. In his linear model, a positive shock to the growth rate of M1 *reduces* output over a three-year horizon. The response of output in a high growth regime is similar to the linear model - i.e. a positive shock to money growth *reduces* output, whereas the response in a low-growth regime is almost nonexistent. The price level responds more positively in booms than in recessions. So the Phillips curve induced by this shock is approximately vertical in a low-growth regime, and actually negatively sloped in a high-growth regime.

Garcia (2002) study the response of quarterly industrial production growth to monetary policy in the US from 1955:2 to 1993:1 . They identify the business cycle with a two-state Markov switching regime and estimate

$$\Delta y_t - \mu_0 - S_t \mu_1 = \Sigma_{i=1}^{r} \phi_i \left(\Delta y_{t-i} - \mu_0 - S_{t-i} \mu_1 \right) + \beta_{iq} X_{t-i} + S_{t-i} \beta_{ip} X_{t-i} + \epsilon_t$$

where X_t is the interest rate in period t and $S_t = 1$ if the economy is in an expansion at time t. They strongly reject the null[2] that monetary policy, measured either as the simple level of Fed Funds rate or as Choleski innovations to a standard three-variable VAR, is equally powerful in both regimes, in favour of the alternative that they are more powerful in recessions. This method assumes, among other things, that the intrinsic persistence of GDP is the same in booms and recessions. If this assumption is violated in practice, the results will be biased.

Smets and Peersman (2001) study the response of quarterly industrial production growth to monetary policy in seven Euro-area countries. First, they identify the business cycle with a two-state Markov switching regime with fixed autoregressive coefficients but state-dependent means μ_{i,s_t} for each country i at time t in state s

---

[2] i.e. the hypothesis that $\Sigma_{i=1}^{r} \beta_{ip} = 0$ for $r = 4$

$$\Delta y_{i,t} - \mu_{i,s_t} = \phi_1 \left(\Delta y_{i,t-1} - \mu_{i,s_{t-1}} \right) + \phi_2 \left(\Delta y_{i,t-2} - \mu_{i,s_{t-2}} \right) + \epsilon_{i,t}$$

They then separately identify monetary policy shocks with a linear VAR and use the historical contribution to the time-t policy rate in this VAR as the measure of the shock. They add the first lag of monetary policy shocks (the contribution of historical shocks to the current interest rate) to the AR(2)

$$\Delta y_{i,t} - \mu_{i,s_t} = \phi_1 \left(\Delta y_{i,t-1} - \mu_{i,s_{t-1}} \right) + \phi_2 \left(\Delta y_{i,t-2} - \mu_{i,s_{t-2}} \right) + \beta_{s_{t-1}} MP_{t-1} + \epsilon_{i,t}$$

imposing that the state of the economy is the same across the countries in the sample. They find that β is more negative in recessions than in booms - essentially the opposite of our finding.

This method imposes strong assumptions on the dynamics of output. Firstly, that past monetary policy shocks can be aggregated across time in a linear model when the underlying environment may be nonlinear. Secondly, that the propagation of a given monetary shock (the ϕ coefficients) is the same in different regimes; in other words, all of the difference in the impact of monetary policy is apparent in the single β coefficient.

Lo and Piger (2005) estimate the following equation

$$\phi(L) y_t^T = \gamma_0(L) x_t + \gamma_0(L) x_t S_t + \epsilon_t$$

where y_t^T is the transitory component of log quarterly industrial production, x_t is a monetary policy shock identified from a three-variable structural VAR. S_t is a two-state Markov-switching process, in which the probabilities of transition from boom to recession is function of state variables z_t. The authors find that putting a constant and two lags of an NBER recession date indicator in z_t yields very strong evidence of asymmetry in the response of output to monetary policy. They calculate impulse response functions to a monetary policy shock in the four possible combinations of realisations of the state variable $\{S_t, S_{t+1}\}$ and find that monetary policy is most

5

powerful when the economy is in a recession either in period t or $t+1$. Accordingly, in calculating the impulse response, they do not allow the future state of the economy to change, either exogenously or in response to a monetary policy shock. Given that the aim of the exercise is to assess the impact of monetary policy on output - the state variable - this approach is difficult to defend.

In results, though not in method, our paper is closer to Thoma (1994), who estimates a non-linear VAR in output and monetary variables, allowing some of the coefficients to depend linearly on the deviation of output growth from trend. Like us, he finds that monetary shocks (especially contractionary ones) have more powerful effects in expansions than recessions. Unlike the approach we follow, however, his approach requires the researcher to make a number of discretionary decisions on the econometric specification. Differently from these papers—and importantly for understanding the transmission menchanism—our paper stresses the difference in the response during booms of durables and business investment on the one hand and non-durables on the other, a dimension glossed over in this literature.

Overall the literature has dismissed the empirical and theoretical relevance of policy asymmetries across the cycle. One notable exception is Vavra (2013), who in recent work argues that recessions are often characterized by high realized volatility, and thus frequent price changes, which leads to a steep Phillips curve and ineffective monetary policy. He estimates a New Keynesian Phillips Curve on US data and finds support for this hypothesis. Berger and Vavra (2012) simulate a model of durables expenditure in the presence of adjustment costs and show that durables purchases are less sensitive to subsidies when output is low. They also show that the conditional variance of an ARCH process describing durables expenditure is higher during booms than in recessions, suggesting that either aggregate shocks are larger in booms, or that durables expenditure is more sensitive to shocks of a given size. They supply additional evidence against the former possibility, suggesting that durables expenditure is more sensitive to aggregate shocks - including monetary shocks - during booms. Our findings, while silent about the specific mechanism, support the implication of Berger and Vavra (2012)'s model that monetary policy interventions are more effective during expansions and that most of the effect results from the

response of durables and business investment.

3   Econometric method

3.1   Specification

Our econometric model closely resembles the smooth transition - local projection model (STLPM) employed in Auerbach and Gorodnichenko (2011). The impulse response of variable y_t at horizon $h \in \{0, H\}$ in state $j \in \{b, r\}$ to a shock ε_t is estimated as the coefficient β_h^j in the following regression

$$y_{t+h} = \tau t + \left(\alpha_h^b + \beta_h^b \varepsilon_t + \gamma^{b\prime} x_t\right) F\left(z_t\right) + \left(\alpha_h^r + \beta_h^r \varepsilon_t + \gamma^{r\prime} x_t\right) \left(1 - F\left(z_t\right)\right) + u_t$$

where τ is a linear time trend, α_h^j is a constant and x_t are controls. $F\left(z_t\right)$ is a smooth increasing function of an indicator of the state of the economy z_t. Following Granger and Terasvirta (1994) we employ the logistic function

$$F\left(z_t\right) = \frac{\exp\left(-\theta \frac{(z_t - c)}{\sigma_z}\right)}{1 + \exp\left(-\theta \frac{(z_t - c)}{\sigma_z}\right)},$$

where c is a parameter that controls what proportion of the sample the economy spends in either state and σ_z is the standard deviation of the state variable z. The parameter θ determines how violently the economy switches from expansion to recession when z_t changes.

The local projection model (Jordà (2005)) has a number of advantages relative to a VAR. First, it does not impose the dynamic restrictions implicit in a VAR. Secondly, one can economize on parameters and, in some circumstances, increase the available degrees of freedom. In particular, one loses H observations from the need to use up to H leads as dependent variables. But the number of variables on the right-hand side need only be enough to ensure that the shocks ε_t are exogenous; none are needed to describe the dynamics of the endogenous variable conditional on the shock. If the VAR representation involves a large number of variables and lags,

the net result will be an increase in the available degrees of freedom. And thirdly, in a regime-switching setting, one does not need to take a stand on how the economy switches from one regime to another. The coefficents β_h^j measure the average effect of a shock as a function of the state of the economy when the shock hits, and therefore encompasses the average effect of the shock on the future change in the economy's state. In contrast, when using a regime-switching VAR model, the impulse response of the VAR implicitly assumes no change in the state of the economy, an assumption that is difficult to defend when we are considering shocks with large real effects.

In this paper, for each variable we estimate the $H+1$ equations of the IRF at horizon $0, ..., H$ as a system of seemingly unrelated regression equations. By Kruskal's theorem, this yields the same point estimates of the regression coefficients as equation-by-equation OLS, because the explanatory variables are the same in each equation. But it enables us to calculate the distribution of functions of parameters at different horizons, such as the smoothed IRFs presented in the figures below.

3.2 Data

We work predominantly with chain-linked US National Accounts data downloaded from the website of the Philadelphia Fed. Where our aggregates do not correspond directly with published data, we construct our own approximations to the chain-linked aggregates with Tornqvist indices (Whelan (2000)). We work with log levels of volume indices, and log differences of implied deflators. Our monetary policy shocks ε_t are quarterly averages of the monetary policy shocks identified by Romer and Romer (2004), extended through 2008 by Coibion et al. (2012)[3]. Our sample period (after the effects of the leads and lags described below are taken into account) runs from (shocks occurring in) 1969Q1-2003Q2, with the response variables measured up to five years later. Our sample runs therefore over the four decades leading up to the collapse of Lehman brothers, but does not include the ensuing major financial crisis, when the impact of monetary policy could have been different to a 'normal'

---

[3]Using end-quarter data - i.e. the shock in the final month of the quarter - yielded qualitatively similar results to those below.

recession.

3.3   The state variable and the shocks

We follow Auerbach and Gorodnichenko (2011) and define z_t as a centred seven-quarter moving average of real quarterly GDP growth, and a recession as the worst 20 per cent of the periods in our sample, setting c to make this so. Higher values of γ mean that $F(z_t)$ spends more time close to the $\{0, 1\}$ bounds of the process, moving the model closer to a discrete regime-switching setup. Smaller values of γ mean that more of the observations are taken to contain some information about behaviour in both regimes. We set $\gamma = 3$ to give an intermediate degree of intensity to the regime switching. Chart 1 shows our transformed state variable $F(z_t)$ at these parameter values alongside the time series of the Romer monetary policy shocks. The figure shows, inter alia, that the monetary policy shocks associated with the early part of Paul Volcker's Chairmanship of the Federal Reserve - the period of greatest variability in the shocks - took place on the whole at a time of relatively weak economic activity. Appendix A examines the robustness of our findings to alternative choices of γ (the intensity of regime-switching), c (the proportion of the sample we call a recession), the length and phase-shift of the moving average state variable, and the identification scheme of the monetary policy shocks.

4   Results

4.1   Baseline results

The first four columns of Figure 2 show the smoothed impulse response of the volume of GDP, the inflation rate of the GDP deflator and the Federal Funds rate to an identified monetary-policy shock that generates an initial 1pp rise in the Federal Funds Rate - i.e. h is on the x-axis, β^h is on the y-axis. The first column displays the central estimate of the impulse response in expansions (dashed lines), recessions (dotted lines) and a linear model (solid lines, where we restrict the coefficient to be

9

constant across regimes). The second to fourth columns display central tendencies and 90% confidence intervals for the linear model, expansions and recessions respectively. The solid lines in the fifth column displays the t-statistic of the null hypothesis that $\left(\beta_h^b - \beta_h^r\right)$, with dotted lines at ± 1.65. So, for example, if the solid line in the fourth columns falls below the lower dotted line at some horizon h we can reject the null that the IRFs at that horizon are equal in favour of the alternative that they are more negative in expansions at a 5% significance level. The IRFs are scaled so that the shock results in a 1pp increase in the Fed Funds rate in all three regimes.

Figure 2 shows that the linear model delivers a familiar picture. Following a contractionary monetary policy shock, the level of output starts to fall, reaching a minimum of about half a percentage point below baseline two to three years after the shock, before beginning to recover. Inflation is initially sticky, but eventually falls by up to half a percentage point (at an annualised rate), and then begins to recover by the end of the horizon. The policy rate is persistent but reverts towards and eventually passes through the conditional mean.

The difference between expansions and recessions is seen most clearly in the left-hand column. Output responds almost an order of magnitude more strongly in an expansion than in a recession, with the maximum fall about 2% rather than 0.3%. Inflation also falls much more sharply. In a recession, the peak responses of output and inflation are statistically insignificantly different from zero, and the inflation response is more likely to be positive than negative. In an expansion, the nominal policy rate falls sharply below the conditional mean about two years after the shock, whereas it remains above it in a recession. It is therefore clear from the figures that the larger response of nominal and real variables in an expansion is not attributable to a bigger rise in real interest rates.

Figures 3 and 4 plot the impulse response of the volumes and prices of three expenditure aggregates to the same shock as before. Figure 3 shows that, in line with the response of aggregate output, all the volume indices respond much more in an expansion than in a recession. The top and bottom rows - indices of durable household and business expenditure - respond roughly an order of magnitude more than nondurable consumption, both in an expansion (7% vs 1%) and in the linear

10

model (2.5% vs 0.3%). In a recession, the response of all three kinds of expenditure is insignificant.

Figure 4 shows that the response of prices is also larger in an expansion. In an expansion, durables prices fall by more than nondurables, but the gap is much smaller than the gap in the responses of expenditure. This suggests that there is some substitutability of inputs between the two sectors, such that a contraction in one lowers cost pressures in both. Once again, the response of inflation in a recession is insignificant.

Figure 5 plots the impulse response of four other variables often implicated in the transmission of monetary policy shocks. The first two rows show the response of real government consumption and net tax revenues (as a share of GDP) respectively. If fiscal policy responds systematically to monetary policy, and this response differs systematically across the business cycle, then this could account for the difference in the response of output across the cycle. However, there is little support for this in the data. Government consumption does not respond significantly to the shock in either regime. If anything, it appears to be counteracting monetary policy more in an expansion than in an recession. Taxes appear to be supporting policy more in a recession. So a difference in fiscal policy does not appear to explain the main results.

The third row shows a measure of the external finance premium - the Gilchrist-Zakrajsek bond spread (Gilchrist and Zakrajsek (2012)). Monetary policy could be more powerful in a boom if the external finance premium is more increasing in interest rates in good times than in bad, such that the rates at which households and firms can borrow move by more than the policy rate suggests. However, the opposite appears to be the case: the external finance premium counteracts the effect of a monetary shock in an expansion. In a recession, the premium amplifies the shock. The difference in the response - which would tend to generate an opposite result to the one we find for the impact of monetary policy on expenditure and prices - is not quite significant at standard levels. So the response of financing spreads cannot explain why policy is more powerful in a boom.

The fourth row shows the response of the ratio of non-financial private debt to GDP to a policy tightening. Once again, the response appears to more negative

in an expansion than in a recession, although the significance of the estimate is marginal. There is accordingly no strong evidence that a greater non-price response of credit quantities in an expansion can explain the stronger response of the economy in expansions; if anything, the reverse is apparent in the data.

4.2 The distribution of shocks in expansions and recessions

One possible explanation for these findings is that the response of the economy to monetary policy shocks is indeed nonlinear, but is not directly a function of the state of the economy. Rather, it is possible that policy shocks of different kinds are more common at certain times, and it is this that generates the apparent dependence of the IRF on the state of the business cycle. If, say, large or positive shocks are proportionally more powerful than small or negative shocks, and if they are more common in expansions than recessions, then an empirical model like ours that is linear in the shocks, conditional on the regime, would misleadingly uncover a larger IRF in expansions than in recessions.

Figure 6 shows IRFs for state-independent model modified such that positive and negative shocks are allowed to have different effects. We plot $\left\{\beta_h^+, \beta_h^-\right\}, h \in \{0, H\}$ estimated from the following equation

$$y_{t+h} = \tau t + \alpha_h^b + \beta_h^+ \max[0, \varepsilon_t] + \beta_h^- \min[0, \varepsilon_t] + \gamma^{b\prime} x_t + u_t$$

and again scale β so that the shock raises the policy rate by 1 percentage point on impact. The figure shows that positive Romer shocks have a much larger impact on output than negative Romer shocks. Their effects on inflation are much more similar, with the difference between them not significant at standard levels. The bottom row shows that positive shocks are substantially more persistent than negative shocks, hampering any reliable inferences about the effect of a given interest rate. However, the finding that positive shocks (monetary tightenings) appear to have a bigger impact on output, but not necessarily on inflation, than negative shocks may be

interesting in its own right.[4]

If positive shocks to the federal funds rate were more common in expansions than recessions, the results in Figure 6 might account for the finding that policy tends to be more powerful in expansions than recessions. But no such regime-dependent pattern in the shocks exists. Figure 7 shows estimates of the pdf and the cdf of the shocks overall and depending on the state of the business cycle.[5] There is almost no difference between the central tendencies of the distributions of shocks in booms and recessions - positive shocks do not preponderate in booms.

The main difference between the two regimes apparent in Figure 7 is that there is more mass in the tails of the shock distribution during recessions. If smaller shocks, which are more common in booms, are proportionally more powerful, this could also explain our finding of a larger average impact of shocks. To check this we estimated the following equation

$$y_{t+h} = \tau t + \alpha_h^b + \beta_h^s \varepsilon_t + \beta_h^l \varepsilon_t^3 + \gamma^{b\prime} x_t + u_t$$

i.e. adding the cubed value of the Romer shock as an additional explanatory variable. If the coefficent β_h^l on this variable were significantly positive (negative), this would count as evidence that large shocks of either sign are more (less) powerful. Figure 8 plots the t-statistics associated with the null hypothesis that $\beta_h^l = 0$ for each of the variables featured above. There is some evidence that larger shocks are more persistent (a t-statistic close to +2), but no consistent evidence that this translates into a different effect on output or inflation. If anything, the positive t-statistic in the inflation panel suggests that bigger shocks have proportionally more powerful effects on inflation. Given that big shocks are more common in recessions but policy

---

[4]We estimated another equation in which the impact of policy was allowed depend both on the sign of the shock and on the state of the economy when it hit - i.e. to take on four values at any given horizon. We did not find any consistent statistically significant evidence of non-linearities by the sign of the shock, but the precision of our estimates was low given the loss of degrees of freedom inherent in this procedure.

[5]The linear estimate is the raw Romer shocks smoothed with a normally distributed kernel. The expansion and recession estimates are generated by weighting the kernel function with the $F(z_t)$ and $1 - F(z_t)$ respectively.

is more powerful in expansions, this would lend weight to the idea that a given policy shock is intrinsically more powerful in a boom than in a recession.

In summary, positive shocks appear to be more powerful than negative shocks, but they are not more common in expansions than recessions. Larger shocks are more common in recessions than expansions, but the effect of a shock does not clearly depend more or less than proportionally on its size. This suggests that differences across regimes in the distribution of the shocks, as opposed to differences across regimes in the response to a given shock, do not explain our key findings.

5   Concluding remarks

We have found statistically strong evidence that standard measures of US monetary policy shocks have had more powerful effects on expenditure quantities and prices during economic expansions than during recessions. Moreover, we find that virtually all of the response of activity during booms is due to the response of durables and business investment. These findings are robust to several variations in the empirical model. They do not appear to be an artefact of different patterns in the shocks themselves, and therefore must be due to differences in how a given shock affects the economy at different points in the business cycle.

The findings question the common wisdom that cuts in policy rates can stop or mitigate recessions, calling for the analysis of alternative policy measures during contractions. On the modelling side, the findings call for monetary models that generate a higher sensitivity in the response of durable goods during expansions, an asymmetry that has been largely glossed over by the theoretical literature.[6]

A   Appendix: sensitivity analysis

The following appendix examines the robustness of our findings to alternative choices of γ (the intensity of regime-switching), c (the proportion of the sample we call a

---

[6]As noted, a recent exception is Berger and Vavra (2012),.who propose a promising and plausible mechanism that could account for our empirical findings.

recession), the state variable z_t, and source of monetary policy shocks ϵ_t. The bottom line is that our qualitative results are robust to reasonable alternatives along each of these margins.

A.1   Intensity of regime switching (γ)

Figures 9 and 10 are analogues of figure 2 but where we have set γ equal to 1 and 10 respectively. They show that the qualitative message of the earlier analysis is unchanged.

A.2   Proportion of sample in a recession (c)

Figure 11 shows that the main qualitative conclusions are robust to increasing to 50% the proportion of the sample judged to be more in a recession than in a boom.

A.3   State variable

Figure 12 shows the baseline IRFs calculated when z_t is a lagging rather than centred moving average of output. The gap between booms and recessions shrinks, and become statistically less significant, but the broad picture remains.

A.4   Policy shocks

Figure 13 shows the baseline IRFs calculated when ϵ_t are the structural shocks recovered from a VAR in the log-levels of GDP, the GDP deflator and the Federal Funds rate, with a Choleski identification scheme in which monetary policy is ordered last. The message from the figures is very similar to our baseline case.

References

Auerbach, A. J. and Gorodnichenko, Y. (2011). Fiscal multipliers in recession and expansion, *NBER Working Papers 17447*, National Bureau of Economic Research, Inc.

Berger, D. and Vavra, J. (2012). Consumption dynamics during recessions, *mimeo* .

Christiano, L. J., Eichenbaum, M. and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy, *Journal of Political Economy* **113**(1): 1–45.

Coibion, O., Gorodnichenko, Y., Kueng, L. and Silvia, J. (2012). Innocent by-standers? monetary policy and inequality in the u.s, *NBER Working Papers 18170*, National Bureau of Economic Research, Inc.

Gali, J. (2008). Introduction to monetary policy, inflation, and the business cycle: An introduction to the new keynesian framework, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.

Garcia, R. (2002). Are the effects of monetary policy asymmetric?, *Economic Inquiry* **40**(1): 102–119.

Gilchrist, S. and Zakrajsek, E. (2012). Credit spreads and business cycle fluctuations, *American Economic Review* **102**(4): 1692–1720.

Granger, C. and Terasvirta, T. (1994). Modelling nonlinear economic relationships, *International Journal of Forecasting* **10**(1): 169–171.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections, *American Economic Review* **95**(1): 161–182.

Lo, M. C. and Piger, J. (2005). Is the response of output to monetary policy asymmetric? evidence from a regime-switching coefficients model, *Journal of Money, Credit and Banking* **37**(5): 865–86.

Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications, *American Economic Review* **94**(4): 1055–1084.

Smets, F. and Peersman, G. (2001). Are the effects of monetary policy in the euro area greater in recessions than in booms?, *Working Paper Series 052*, European Central Bank.

Thoma, M. A. (1994). Subsample instability and asymmetries in money-income causality, *Journal of Econometrics* **64**(1-2): 279–306.

Vavra, J. (2013). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation, *mimeo* .

Weise, C. L. (1999). The asymmetric effects of monetary policy: A nonlinear vector autoregression approach, *Journal of Money, Credit and Banking* **31**(1): 85–108.

Whelan, K. (2000). A guide to the use of chain aggregated nipa data, *Finance and Economics Discussion Series* .

Woodford, M. (2002). *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton.

Figure 1: Monetary policy shocks and the state of the economy

Figure 2: Impulse response of headline variables to a monetary policy shock

19

Figure 3: Impulse response of expenditure volumes
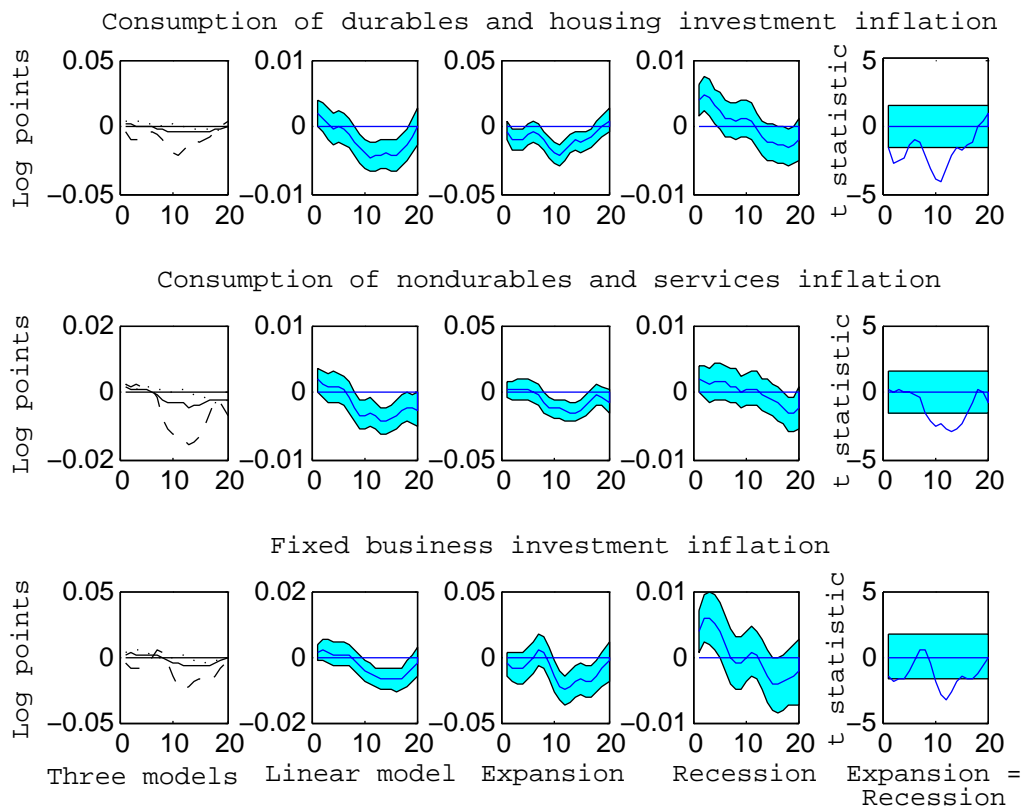
Figure 4: Impulse response of expenditure prices
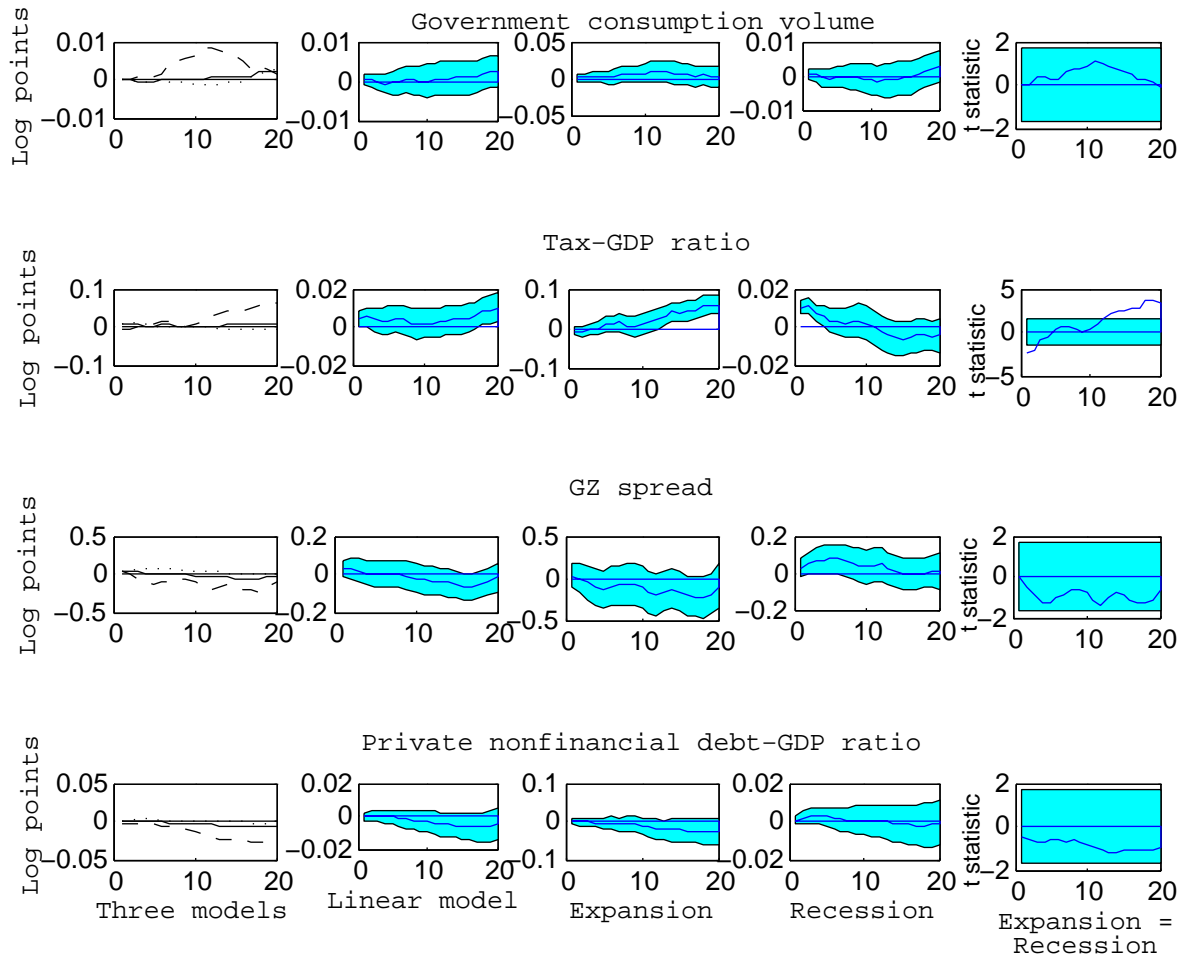
Figure 5: Impulse response functions of fiscal and credit variables
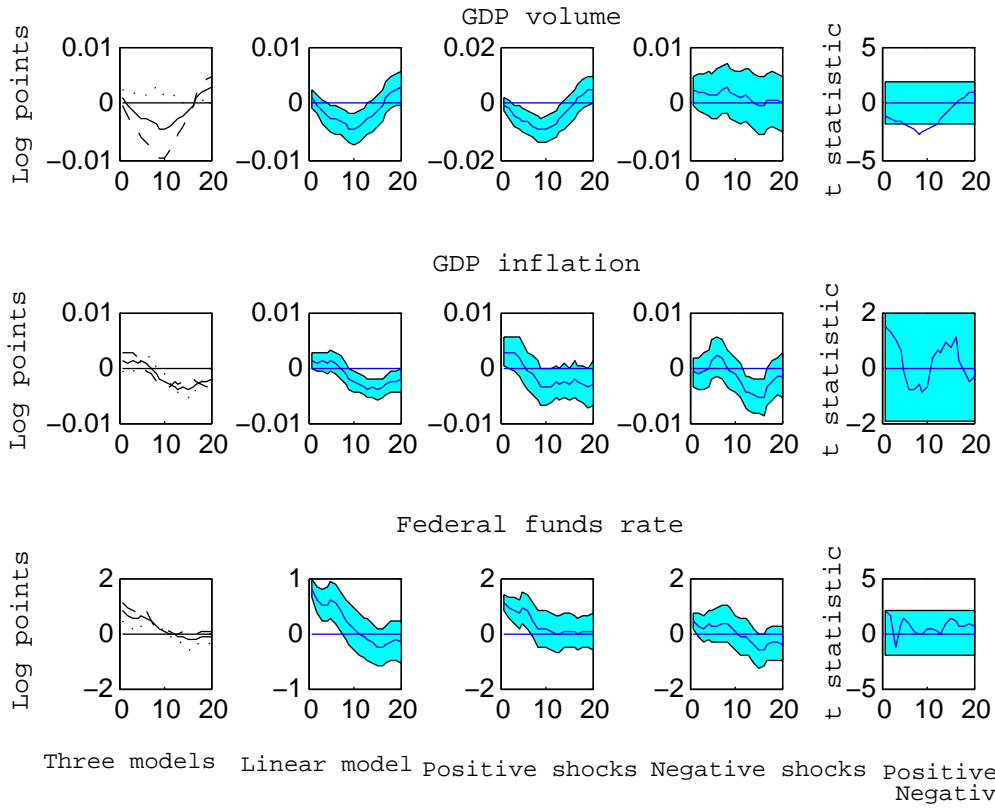
Figure 6: Impulse response to positive and negative monetary policy shocks

23

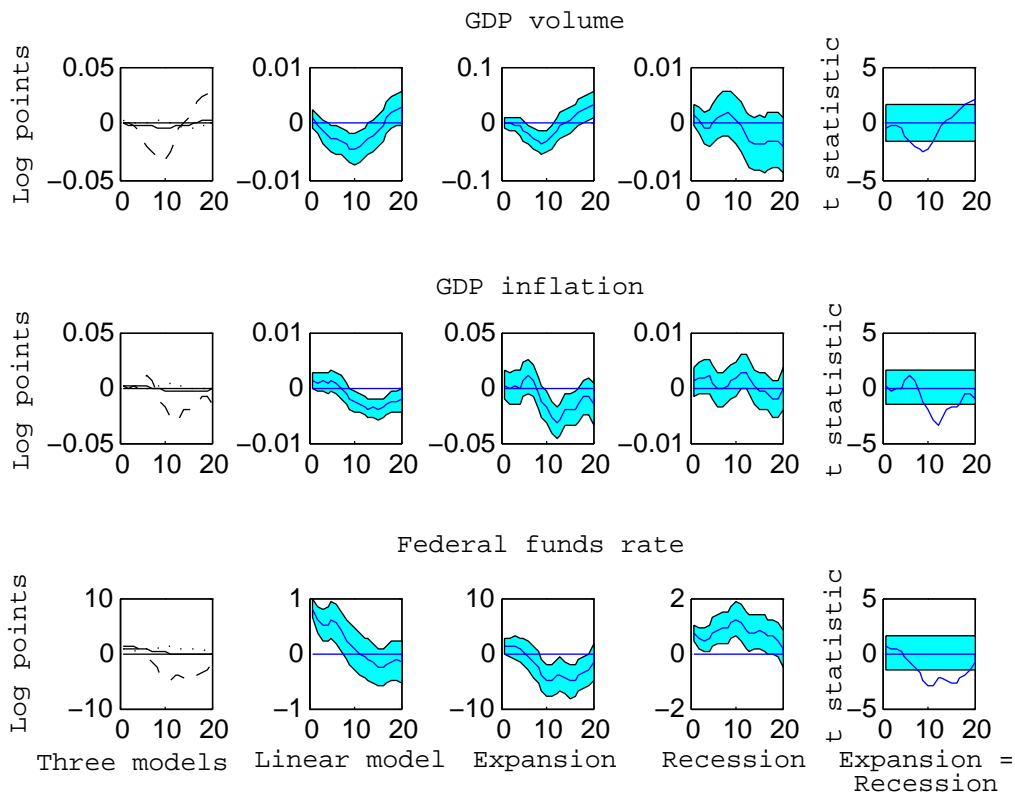Figure 7: Pdfs and cdfs of the regime-specific shocks

Figure 8: t statistics on cubed monetary policy shocks

Figure 9: Impulse response of headline variables to monetary policy shock, $\gamma = 1$

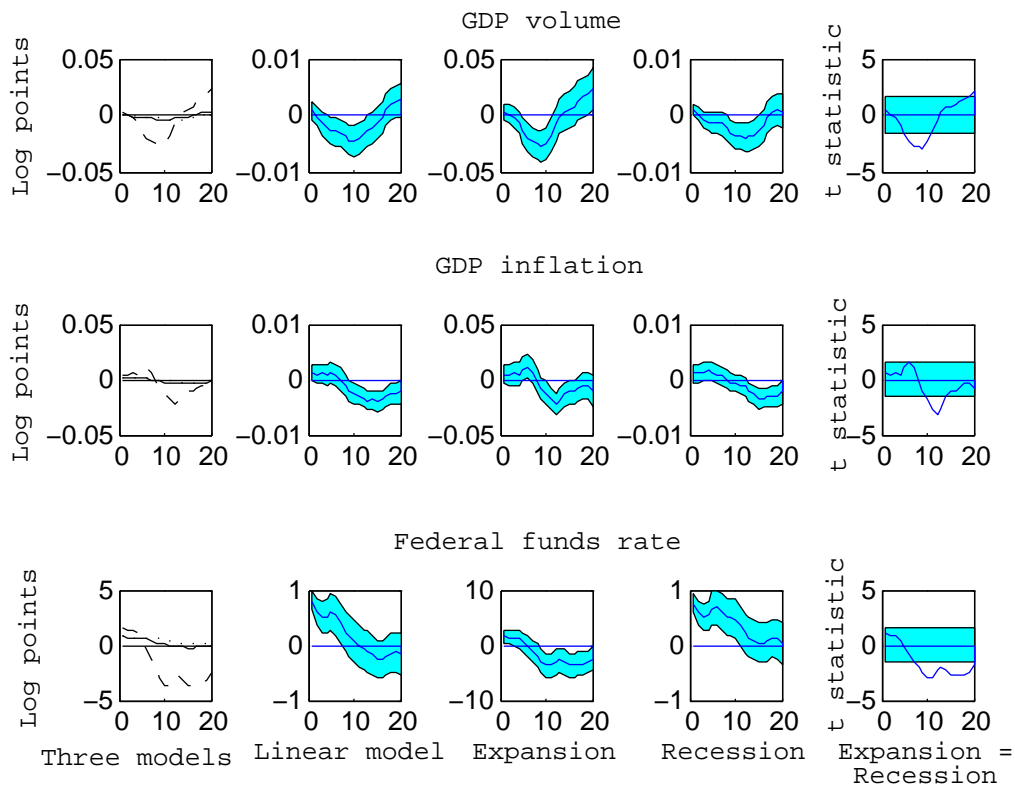Figure 10: Impulse response of headline variables to monetary policy shock, $\gamma = 10$

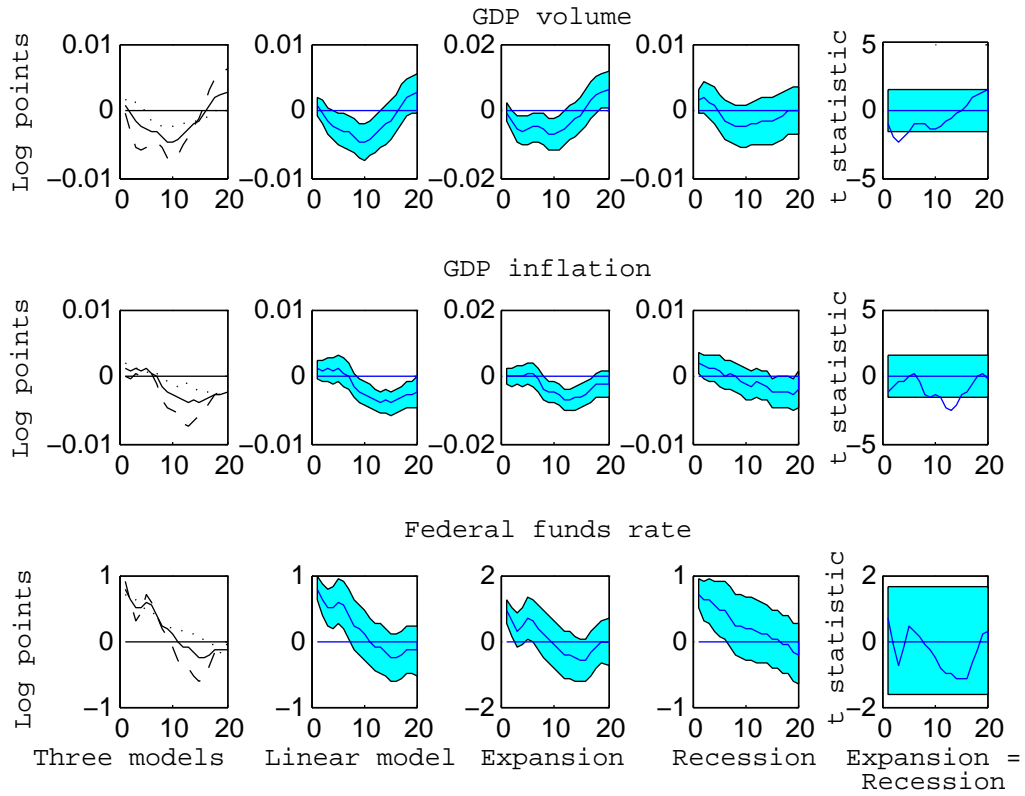Figure 11: Impulse response of headline variables to monetary policy shock, $c = 50$

Figure 12: Impulse response of headline variables to monetary policy shock, z_t lagging
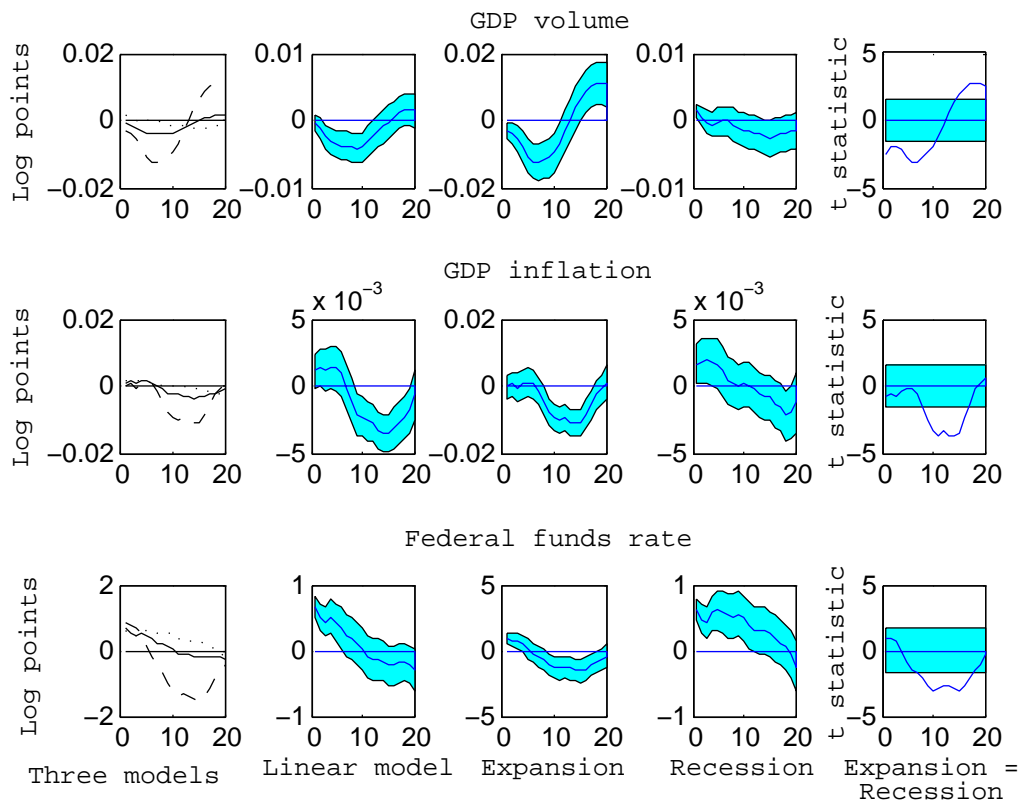
Figure 13: Impulse response of headline variables to monetary policy shocks identified with a VAR