

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

Silvana Tenreyro^a and Gregory Thwaites^b

^a LSE, CEP, CfM; ^bLSE, CfM, Bank of England

This version: September 2015

First version: April 2013

Abstract

We estimate the impulse response of key US macro series to monetary policy shocks, allowing both the shocks, identified analogously to Romer and Romer (2004), and 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 asymmetry relates to how fast the economy is growing, rather than to the level of resource utilisation. There is some evidence that fiscal policy has counteracted monetary policy in recessions but reinforced it in booms. We also find 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, state-dependent impulse responses.

*Tenreyro acknowledges financial support from the European Research Council under the European Community's ERC starting grant agreement 240852 Research on Economic Fluctuations and Globalization. We thank Joshua Angrist, Francesco Caselli, Martin Eichenbaum, Wouter den Haan, Oscar Jorda, Robert Lieli, Valerie Ramey, Alan Taylor, Sarah Zubairy, the editor and our anonymous referees for helpful comments, Paola Ganum and Junyi Hou for superb research assistance, and Olivier Coibion for providing his data. This paper was written largely while Thwaites was at LSE and does not represent the views of the Bank of England.

1 Introduction

Is monetary policy effective in recessions? In recent years this perennial question took centre stage in the public policy debate, as central banks in the United States and Europe faced the deepest post-war crisis. A priori, whether monetary policy is more powerful in recessions or expansions 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 studied this question, and adjacent ones, finding mixed results.

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. In order to allow impulse response functions to depend on the state of the business cycle, we adapt the local projection method of Jordà (2005) and combine it with the smooth transition regression method of Granger and Terasvirta (1994).¹ 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 monetary policy shocks identified in the manner of 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. 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 substantially 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 and generally insignificantly different from zero. These differences are not attributable to differences in the amplification afforded by the response of credit prices or quantities. We find that contractionary shocks are more powerful than expansionary shocks - in line with Angrist, Jordà, and Kuersteiner (2013), who employ a different method. But given that they are equally common in both expansions and recessions, this cannot be the source of asymmetry across the business cycle. We study different indicators of the state of the economy and find that measures of the growth rate of activity such as GDP growth are the

¹Auerbach and Gorodnichenko (2011) and Ramey and Zubairy (2014) use a similar procedure to study the effect of fiscal policy, though the method has never been applied to the analysis of monetary policy.

most reliable determinant of monetary policy effectiveness, whereas measures of the level of resource utilization such as the output gap do not as clearly distinguish regimes. We find that fiscal policy seems to have counteracted monetary policy in recessions, but reinforced it in expansions, which provides one explanation for our results.

These findings are relevant for the design of stabilisation policy and the models used to analyse it. If changes in the policy rate have limited impact in a recession, central banks will be more likely to need to resort to other (unconventional) monetary policy measures to achieve the desired expansionary effect. Policymakers may also need to rely more heavily on fiscal or financial policies to stabilize the economy in a deep or protracted slump. On the modelling side, the findings call for macroeconomic models that generate a higher sensitivity in the response of the economy (and in particular, the durable-good sector) during expansions.

The remainder of this paper is structured as follows. Section 2 reviews the literature. Section 3 explains the empirical method and describes the dataset. Section 4 sets out the main results. We conduct sensitivity analysis in section 5. Section 6 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, mostly written a decade or more ago. Previous research 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, Eichenbaum, and Evans (2005), Woodford (2011) and Gali (2008). Our paper makes use of important subsequent methodological innovations in the estimation of impulse response functions in regime-switching environments.

The closest paper to ours in terms of 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 set of variables in the VAR is small: industrial production, consumer prices, and 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, against the weight of empirical evidence on this matter. 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 paper implies that monetary policy is virtually ineffective in a low-growth regime, and actually contractionary in a high-growth regime, a result that is hard to reconcile with the standard empirical result that, on average, monetary policy is expansionary.

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

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

where X_t is the interest rate in period t and $S_t = 1$ if the economy is in an expansion at time t . The procedure strongly rejects the null² 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 and other stochastic properties of GDP are the same in booms and recessions. There is substantial evidence that this assumption does not hold (see, for example, Acemoglu and Scott (1997) and references therein).

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

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

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

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

interest rate) to the AR(2)

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

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, it assumes that past monetary policy shocks can be aggregated across time in a linear model when the underlying environment may be nonlinear. Secondly, it assumes 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, and 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 a 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 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. In contrast to this and other papers discussed above—and importantly for understanding the transmission mechanism—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.

In summary, the general form of empirical model employed in the studies above is

$$(y_t - \bar{y}_t) = \alpha(L) (y_{t-1} - \bar{y}_{t-1}) + \beta x_t + \varepsilon_t$$

where x is the policy shock and y is the set of outcome variables. These studies typically allow only a proper subset of $\{\alpha(L), \beta, \bar{y}\}$ to depend on the state of the cycle. They must also take a stand on how the policy shock alters the transmission between regimes.

In contrast to the methods used previously, a 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 - the true model can take any form. Secondly, one can economize on parameters and, in some circumstances, increase the available degrees of freedom. In particular, one loses observations from the need to use 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.

Thirdly - and most importantly for the present study - with a regime-switching local projection model one does not need to take a stand on **how** the economy switches from one regime to another. More specifically, a regime-switching local projection model takes the form

$$y_{t+h} = F(z_t) (\beta_b^h x_t + \gamma_b' z_t) + (1 - F(z_t)) (\beta_r^h x_t + \gamma_r' z_t) + u_t$$

where $F(z_t)$ is an indicator of the regime. The coefficients 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. Alternatively, the transition from one regime to another effected by the policy shock must itself be modelled and simulated, involving a series of potentially erroneous and controversial modelling choices. Ramey and Zubairy (2014) finds that this can have an important bearing on the results when estimating the state-dependence of US fiscal policy. It may explain the difference between our findings and some of those in the previous literature on state-dependent monetary policy summarised above.

Overall the theoretical literature has not had much to say about the state-dependent impact of macroeconomic policy 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 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

In this section we first set out the specification of the econometric model used in this study. Then we explain our approach to statistical inference. Finally we describe our data sources, our state variables, and our identified policy shocks.

3.1 Specification

Our econometric model closely resembles the smooth transition - local projection model (STLPM) employed in Auerbach and Gorodnichenko (2011) and Ramey and Zubairy (2014) to analyze fiscal policy. The impulse response of variable y_t at horizon $h \in \{0, H\}$ in state $j \in \{b, r\}$ ³ to a shock ε_t is estimated as the coefficient β_h^j in the following regression

$$y_{t+h} = \tau t + (\alpha_h^b + \beta_h^b \varepsilon_t + \gamma^{b'} x_t) F(z_t) + (\alpha_h^r + \beta_h^r \varepsilon_t + \gamma^{r'} x_t) (1 - F(z_t)) + u_t \quad (1)$$

where τ is a linear time trend, α_h^j is a constant and x_t are controls.⁴ $F(z_t)$ 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

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

³ b denotes an expansion, r denotes a recession

⁴In the baseline specification, x_t contains one lag each of the dependent variable and Federal Funds rate

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.

In this paper, for each variable we estimate the $H + 1$ equations of the IRF at horizon $0, \dots, H$ 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 Inference

We employ two different approaches to conducting inference on our estimated impulse response functions. In order to conduct inference on cumulative impulse responses, moving averages and other functions of response variables at different horizons, each of these approaches needs to calculate the correlation of parameter estimates between equations. The first approach is to calculate standard errors analytically, allowing for the possibility of serially correlated residuals within equations and across equations. To capture this we follow Ramey and Zubairy (2014) and use the Driscoll and Kraay (1998) method to adjust standard errors for the possibility of correlation in the residuals across dates t and horizons h . This amounts to estimating the parameters of the equations separately, as above, and then averaging the moment conditions across horizons h when calculating Newey-West standard errors. Following Jordà (2005), we set the maximum autocorrelation lag $L = h + 1$ where h is the maximum horizon of the impulse response function.

The second approach is to bootstrap the key statistics of interest, namely the sign of $\beta_h^b - \beta_h^r$. This will allow not only for various forms of dependence among the residuals, but will also account for the fact that we scale the IRFs with estimated parameters (i.e. the impact effect of a policy shock on policy rates). Montiel Olea, Stock, and Watson (2012) argue that proper inference should take account of this, but also that standard large-sample 2SLS statistics can be misleading. To perform the bootstrap we construct 10,000 samples with replacement of size T and calculate the fraction of cases in which our null hypothesis does not hold. To account for the dependence between our observations, samples are constructed by aggregating contiguous blocks of observations of length H . We transform the resulting p-value into a t-statistic for comparability with the other measures.

Inference on the above families of t-statistics - $H + 1$ for each response variable - will generate a ‘multiple testing problem’: if we test n true null hypotheses at significance level

α , we will on average reject αn of them. Methods such as Holm (1979) exist to deal with this issue. However, in our setting there are no strong **a priori** grounds for specifying at what horizon the effects of monetary policy shocks depend on the state of the business cycle, rendering these methods inapplicable in the present study. It turns out that the t-statistics we present in section 4 are strongly correlated at adjacent horizons, alleviating the practical concern of this problem. But to deal further with this concern, we also calculate and conduct inference on cumulative impulse response functions at discrete horizons.

3.3 Data

We work predominantly with chain-linked US National Accounts data downloaded from the website of the Philadelphia Fed.⁵ 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 quarterly annualised log differences of implied deflators.

Our sample period (after the effects of the leads and lags described below are taken into account) runs from (shocks occurring in) 1969Q1-2002Q4, with the response variables measured up to five years later - i.e. the end of 2007. Our sample runs therefore over the four decades leading up to the beginning of the financial crisis but does not include the collapse of Lehman brothers or the ensuing major financial crisis, when the impact of monetary policy could have been different to a ‘normal’ recession.⁶

3.4 The state variable and the shocks

We define z_t as a seven-quarter moving average of real quarterly GDP growth. Following Ramey and Zubairy (2014), and in contrast to Auerbach and Gorodnichenko (2011), our moving average term z_t is a lagging rather than centred moving average, so that future values of response variables do not appear on the right-hand side of the regression. Higher values of θ 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 θ mean that more of the observations are taken to contain some information about behaviour in both regimes. We follow Auerbach and Gorodnichenko (2011) and calibrate rather than estimate the parameters of the smooth transition model, for the same reasons they cite - it is difficult in practice to identify the curvature and location of the transition function in the data -

⁵We use the latest vintage of the data rather than real-time estimates

⁶Using end-quarter data - i.e. the shock in the final month of the quarter - yielded qualitatively similar results to those below.

and given the need for distributional assumptions on the error term when estimating by maximum likelihood. We set $\theta = 3$ to give an intermediate degree of intensity to the regime switching, and also follow them in defining a recession as the worst 20 percent of the periods in our sample, setting c to make this so. The robustness of our results to each of these choices is investigated below.

3.5 Non-linear Romer regression

Romer and Romer (2004) identify monetary policy shocks as the residuals from an estimated reaction function. The premise of the current study is that the behaviour of the economy is characterised by important forms of non-linearity and state-dependence. It is conceivable that the reaction function of the Federal Reserve has also been state-dependent, such that estimating shocks with a standard linear framework would inject state-dependent measurement error, potentially generating an apparent asymmetry in the response of the economy at different points in the cycle. For example, if shocks are measured with greater error in a recession than a boom, response coefficients will be more attenuated, biasing our findings towards finding a difference when there may be none.

To obviate this possibility, we estimate a smooth transition analogue of the original Romer and Romer (2004) regression and use the resulting shocks in our baseline estimates. To be precise, if we write the original regression as

$$\Delta FFR_t = \beta X_t + \varepsilon_t \tag{2}$$

where X are the control variables employed by Romer and Romer (2004) and the estimated residuals $\hat{\varepsilon}_t$ are the identified monetary policy shocks, then our state-dependent identification scheme is

$$FFR_t = F(z_t) \beta^b X_t + (1 - F(z_t)) \beta^r X_t + \tilde{\varepsilon}_t \tag{3}$$

and $\tilde{\varepsilon}_t$ are our non-linearly identified shocks.

When estimated linearly over a common sample, we replicate the results exactly; when extending the sample the regression coefficients and hence the residuals change slightly, such that our extended series has a correlation of 0.987 with the original. The original Romer shocks, the shocks identified with the same method on a larger sample, and the non-linearly identified shocks are all shown in Figure I, which also shows our transformed state variable $F(z_t)$ at the baseline parameter values. All variables are aggregated to quarterly frequency. 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. The state-dependent policy shocks have a 0.902 correlation with the original series over a common sample. In Section 5 we examine the robustness of our findings to the use of linearly identified shocks.

4 Results

In this section of the paper, we first set out our baseline results. We then explore whether the asymmetry we find is due to a different pattern of shocks across the business cycle.

4.1 Baseline results

The first four columns of Figure II show the smoothed impulse responses of the volume of GDP, the level of the PCE deflator and the Federal Funds rate to an identified monetary-policy shock that generates an initial 1 percentage point rise in the Federal Funds Rate - i.e. h is on the x-axis, β^h is on the y-axis⁷. 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 constant across regimes). The second to fourth columns display central tendencies and 90 percent confidence intervals for the linear model, expansions and recessions respectively. The charts in the fifth column represent our two estimates of the t-statistic of the null hypothesis that $(\beta_h^b - \beta_h^r) = 0$, with the area between ± 1.65 shaded blue. So, for example, if the solid line in the fourth columns falls below the lower extreme of the blue area 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 10 percent significance level. The IRFs are scaled so that the shock results in a 1 percentage point increase in the Fed Funds rate in all three regimes.

Figure II 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 percent below baseline two to three years after the shock, before beginning to recover. The price level is initially sticky, but eventually falls by about 3 percent, flattening off 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 more strongly in an expansion than in a recession, with the

⁷In this case ‘smoothed’ means three-period centered moving averages of the IRFs, except at the endpoints of the function. The standard errors of these moving averages are calculated taking account of the covariance between the estimates at different points estimated above.

maximum fall about 1 percent in an expansion. The price level also falls much more sharply, by about 8 percent in an expansion against 4 percent in a recession. In a recession, the responses of output and prices are mostly statistically insignificantly different from zero. In an expansion, the nominal policy rate falls sharply below the conditional mean about two years after the shock, perhaps because of the systematic component of policy responding to the contraction the previous shock has created. 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 long-term nominal interest rates.

Table I cumulates the impulse response functions for the level of GDP and inflation, and shows two alternative estimates of how significant the difference between them is, as set out above - Driscoll-Kraay standard errors and bootstrapped significance tests. The cumulative effect of a monetary policy shock is significantly larger at standard levels, with the precise horizon depending on how the standard errors are calculated. Appendix A (not for publication) displays unsmoothed impulse responses of our headline variables.

Figure III plots the impulse response of the volumes of three expenditure aggregates to the same shock as before. In line with the response of aggregate output, all the volume indices respond much more in an expansion than in a recession, with the difference already significant one quarter after the shock. The top row - corresponding to an index of durable household expenditure - responds roughly an order of magnitude more than nondurable consumption, both in an expansion and in the linear model. In a recession, the response of all three kinds of expenditure is insignificant.

Figure IV plots the impulse responses of four other macroeconomic variables which may play a role 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. The first row shows that there is weak evidence that real government consumption responds positively on average to a tightening of monetary policy. Why this should be so is not clear. One possibility is that spending is set in nominal terms, such that real spending increases because the price level falls. However, the behaviour of government consumption in response to monetary policy at different points in the cycle does not support this explanation. In an expansion, when the disinflationary effects of policy are at their strongest, there is no evidence that real government consumption increases in response to a monetary policy tightening, whereas there is a significant increase in a recession. It could be that the larger fall in the price level we see in an expansion raises the real burden of public debt (Sterk and Tenreyro (2015)), such that spending cuts become necessary. Whatever the reason for this asymmetry, to the extent that increases in government consumption are expansionary, they will be offsetting the contractionary impulse provided by monetary policy: government

spending seems to have been ‘working against’ monetary policy during recessions but not during expansions.

There is weaker evidence for the same on the tax side. The second line of the figure shows that, after the first few quarters, the tax-GDP ratio seems to rise more sharply in response to a monetary tightening in an expansion than in a recession. This again may be due to the stronger response of the price level in an expansion, and its effect on the government’s intertemporal budget constraint. To the extent that tax rises are contractionary, this will reinforce the effect of monetary policy. However, much of the government debt that is being revalued by the disinflation is held by the US private sector, who will therefore enjoy a positive wealth shock that will offset much of the extra taxation needed to service the increase in debt. This could offset the contractionary effect of tax rises to some extent.

The third row of the table 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 strongly 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, there no evidence of an effect in this direction and by the end of the sample if anything 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. So the response of financing spreads cannot explain why policy is more powerful in a boom. 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.

The fourth and final row depicts the response of private nonfinancial credit volumes in relation to GDP. There is weak evidence that, on average, a monetary tightening reduces the credit-to-GDP ratio. The middle column shows that this evidence is stronger in a boom, with no consistent effect apparent in a recession. However, given the that both credit spreads and volumes appear to fall more following a monetary tightening in booms than in recessions, it is difficult to attribute the stronger response in the price and volume of activity to a stronger response in credit supply.

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

$$y_{t+h} = \tau t + \alpha_h^b + \beta_h^+ \max[0, \varepsilon_t] + \beta_h^- \min[0, \varepsilon_t] + \gamma^{b'} x_t + u_t$$

and again scale β so that the shock raises the policy rate by 1 percentage point on impact. The figure shows that positive shocks (i.e. monetary tightenings) have a much larger impact on output than negative shocks, although the estimates are only borderline significant at standard levels. This finding is consistent with those in Cover (1992), Long, Summers, Mankiw, and Romer (1988) and Angrist et al. (2013). The effects of positive and negative shocks on inflation are statistically much harder to distinguish, with the difference between them not significant at standard levels. However, the finding that contractionary shocks (monetary tightenings) appear to have a bigger impact on output, but not necessarily on inflation, than negative shocks is interesting in its own right.⁸

If positive shocks to the federal funds rate were more common in expansions than recessions, the results in Figure V 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 VI shows estimates of the pdf and the cdf of the shocks overall and depending on the state of the business cycle.⁹ There is little 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 VI, is that the distribution of shocks is more variable during recessions. If smaller shocks, which are more common in booms, are proportionally more powerful, this could also explain our finding of a larger

⁸We 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.

⁹The 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.

average impact of shocks. To check this we estimated the following equation

$$y_{t+h} = \tau t + \alpha_h^b + \beta_h^s \varepsilon_t + \beta_h^l \varepsilon_t^3 + \gamma^{b'} x_t + u_t$$

i.e. adding the cubed value of the policy shock as an additional explanatory variable. If the coefficient 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. The left-hand column of Figure VII plots the functions $\beta_h^l, h \in \{1, H\}$ with associated 90 percent confidence intervals, while the right-hand column shows t-statistics associated with the null hypothesis that $\beta_h^l = 0$ for each of the variables. The bottom row shows some evidence that larger shocks have tended to die out more quickly, which may explain why their negative effects on output and inflation are attenuated (top and middle rows). Table III shows estimates of the cumulative IRFs and significance as above.

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, and there is weak evidence that there are less persistent and have proportionally weaker effects on output and inflation. It is accordingly possible that differences across regimes in the distribution of the shocks, as opposed to differences across regimes in the response to a given shock, explain part of the weaker effect.

5 Sensitivity analysis

The following section examines the robustness of our findings to alternative choices of the policy shocks (subsection 5.1), lags and trends in the regression equation (subsection 5.2), the state variable z_t (subsection 5.3), the phase shift of the state variable (subsection 5.4), the intensity of regime-switching θ (subsection 5.5) and the proportion of the sample we call a recession c (subsection 5.6).

5.1 Alternative policy shocks

The monetary policy shocks used in this paper are identified as the residuals from an estimated reaction function, modified from Romer and Romer (2004) such that the reaction function is itself state-dependent. As an alternative to this, Figure VIII and Table IV display the results of using the original linear identification scheme. There is very little difference with our baseline results, with the differences we identify being slightly larger and statistically more significant.

Figure IX and Table V show 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 linear IRF is calculated using shocks from a linear VAR, whereas the non-linear IRFs employ shocks calculated with a non-linear VAR analogous to our non-linear specification of the Romer and Romer (2004) regression. The peak response of GDP is significantly larger in a boom, but the difference is generally nonmonotone, with a larger initial fall in GDP followed by a larger rebound in a boom. The difference in the inflation IRFs is much larger than the baseline case.

In summary, therefore, our baseline results are robust to identifying shocks with a linear regression in the manner of Romer and Romer (2004), or with a state-dependent VAR and Choleski identification scheme.

5.2 Trends and lags in the regression equation

The baseline regression equation (1) contains a log-linear trend and one lag of both the policy and endogenous variable. In this subsection we examine the robustness of our results to both choices.

Table VI and Chart X contain our standard IRFs calculated with a regression model identical to (1) except for the omission of a time trend. Qualitatively, the results retain the message that monetary policy is statistically and economically more contractionary in a boom than in a recession. But the anomalous shape of some of the IRFs suggests that a model without a log-linear trend would be misspecified. For example, the impulse response of GDP in a recession is uniformly positive and not hump-shaped, while the impulse response of interest rates in a boom becomes very large towards the end of the response horizon. For this reason, we retain a log-linear time trend in our baseline model.

Our baseline model contains one lag each of the dependent variable and the Federal Funds rate. This lag structure is optimal as indicated by the Schwartz Bayesian Criterion (SBC), given by

$$-2 \ln(\hat{L}) + k \ln(n) \tag{4}$$

where \hat{L} is the maximised value of the likelihood function, k is the number of regressors and n is the number of observations in the sample.

As a robustness test, we also calculated optimal lag length with the Akaike Information Criterion (AIC) given by

$$2k - 2 \ln(\hat{L}) \tag{5}$$

This criterion indicates that two lags are optimal for the dependent variable, and one lag is optimal for the policy variable. Table VI and Chart X contain our standard IRFs calculated with a regression model identical to the baseline but with the lag structure preferred by the AIC. The overall qualitative picture from the IRFs is unchanged. But there is statistically significant evidence that policy is somewhat more disinflationary in a recession than a boom in the short run, although the opposite is true in the long run. However, the GDP impulse response in a recession is once again not hump-shaped. So there is once again a question about whether the regression is misspecified when two lags are included. In the light of this evidence and that from our preferred SBC, our baseline lag structure is warranted.

5.3 The state variable

Our baseline results employ a measure of the economic cycle - a moving average of GDP growth - to which there are many reasonable alternatives. This subsection examines the sensitivity of our results to three of them.

Figure XII and Table VIII show the response of our headline variables when Z_t is a moving average of a $\{0, 1\}$ indicator of recession, defined as the proportion of the quarter in which the economy was in recession, as determined by the NBER. The difference in the response of output over the near term remains significant at standard levels when the standard errors are calculated asymptotically, but not when bootstrapped. The difference in cumulative inflation switches sign, and is significant at very short and long horizons when calculated asymptotically, but not at intermediate horizons. The cumulative response of the interest rate - evident in the bottom row of figure XII - is much weaker in a boom than in a recession, complicating any inference about the impact of a comparable monetary impulse.

Figure XIII and Table IX contain the results of the same test but define Z_t as an HP-filtered output gap. An HP filter is already essentially a centered moving average of the level of GDP, so no further filtering or phase shifting is undertaken. The charts and tables show that when output is low, policy contractions are more disinflationary. However, given that output detrended in this fashion contains information about future output, and that the evidence is statistically weak, we do not place much weight on this possibility.

To sum up, the balance of evidence suggests that when the economy is growing quickly, monetary policy shocks are more powerful with respect to output. There is some evidence that when the level of output is low, policy shocks are more disinflationary, but the interpretation of this evidence is complicated by the detrending procedure, which implicitly includes leads of the response variable.

5.4 Phase shift of state variable

Figure XIV and Table X show the baseline IRFs calculated when, following Auerbach and Gorodnichenko (2011) rather than Ramey and Zubairy (2014), z_t is a centered rather than lagging moving average of output. The gap between booms and recessions shrinks somewhat and appears earlier in the case of GDP growth, but the broad picture remains for both output and inflation, and remains statistically significant in the case of output. So our results do not appear to be an artefact of using a centered moving average to calculate the state of the economy.

5.5 Intensity of regime switching (θ)

Tables XI and XII are analogues of Table I but where we have set θ equal to 1 and 10 respectively. They show that the qualitative message of the earlier analysis is unchanged - our results are robust to reasonable changes in the intensity of regime switching.

5.6 Proportion of sample in a recession (c)

Figure XV and Table XIII show that our qualitative conclusions about the response of output to a monetary policy shock are robust to increasing to 50 percent the proportion of the sample judged to be more in a recession than in a boom. The response of inflation is now not statistically different across regimes, but similar in qualitative terms.

6 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. 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, but rather the outcome of differences in the economic effects of a given shock at different points in the business cycle. There is some evidence that fiscal policy offsets monetary policy more in recessions than in booms. We also find that monetary contractions are much more powerful than expansions. In other words, there is truth in the quote attributed to John Maynard Keynes that “you can’t push on a string” - when the economy is weak, monetary policy can do little about it.

Standard estimates in the literature that do not allow for state-dependent impulse responses have masked these differential effects. 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 literature has hitherto focused on linear, state-independent models of monetary policy transmission. In contrast, these 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 in the theoretical literature.

References

- Acemoglu, D. and A. Scott (1997). Asymmetric business cycles: Theory and time-series evidence. *Journal of Monetary Economics* 40(3), 501 – 533.
- Angrist, J. D., Ò. Jordà, and G. Kuersteiner (2013). Semiparametric estimates of monetary policy effects: String theory revisited. Working Paper 19355, National Bureau of Economic Research.
- Auerbach, A. J. and Y. Gorodnichenko (2011). Fiscal multipliers in recession and expansion. NBER Working Papers 17447, National Bureau of Economic Research, Inc.
- Berger, D. and J. Vavra (2012). Consumption dynamics during recessions. *mimeo*.
- Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy* 113(1), 1–45.
- Cover, J. P. (1992). Asymmetric effects of positive and negative money-supply shocks. *The Quarterly Journal of Economics* 107(4), pp. 1261–1282.
- Driscoll, J. C. and A. C. Kraay (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of economics and statistics* 80(4), 549–560.
- Gali, J. (2008). Introduction to monetary policy, inflation, and the business cycle: An introduction to the new keynesian framework. In *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 E. Zakrajsek (2012). Credit spreads and business cycle fluctuations. *American Economic Review* 102(4), 1692–1720.

- Granger, C. and T. Terasvirta (1994). Modelling nonlinear economic relationships. *International Journal of Forecasting* 10(1), 169–171.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, 65–70.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review* 95(1), 161–182.
- Lo, M. C. and J. Piger (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.
- Long, J. B. d., L. H. Summers, N. G. Mankiw, and C. D. Romer (1988). How does macroeconomic policy affect output? *Brookings Papers on Economic Activity* 1988(2), pp. 433–494.
- Montiel Olea, J. L., J. Stock, and M. W. Watson (2012). Inference in vars with external instruments.
- Ramey, V. A. and S. Zubairy (2014). Government spending multipliers in good times and in bad: Evidence from u.s. historical data.
- Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review* 94(4), 1055–1084.
- Smets, F. and G. Peersman (2001). Are the effects of monetary policy in the euro area greater in recessions than in booms? Working Paper Series 052, European Central Bank.
- Sterk, V. and S. Tenreyro (2015). The transmission of monetary policy through redistributions and durable purchases.
- 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. *FEDS Working Paper* 2000(35).

Table I: Cumulative impulse response of GDP and inflation: baseline specification

Cumulative impact on	At horizon h=	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0194	0.0109	0.0059	0.1233
	8	-0.0452	-0.0129	0.1319	0.2316
	12	-0.0751	-0.0240	0.0904	0.1100
	16	-0.0721	-0.0393	0.2379	0.2040
Inflation	4	0.0065	-0.0019	0.1722	0.7865
	8	-0.0080	0.0049	0.1920	0.4917
	12	-0.0481	0.0074	0.0041	0.1401
	16	-0.0778	-0.0153	0.0265	0.1958

Woodford, M. (2011). *Interest and prices: Foundations of a theory of monetary policy*. princeton university press.

Table II: Cumulative impulse response of GDP and inflation: positive and negative shocks

Cumulative impact on	At horizon h=	Regime		Significance level of difference	
		Positive	Negative	Driscoll-Kraay	Bootstrap
GDP	4	-0.0108	-0.0101	0.4833	0.1016
	8	-0.0415	-0.0153	0.1942	0.0284
	12	-0.0754	-0.0269	0.1137	0.0158
	16	-0.0837	-0.0288	0.1171	0.0287
Inflation	4	0.0018	-0.0098	0.2049	0.4487
	8	-0.0097	-0.0185	0.3846	0.2542
	12	-0.0254	-0.0039	0.2999	0.2597
	16	-0.0420	0.0141	0.1024	0.2742

Table III: Cumulative impulse response of GDP and inflation: small and large shocks

Cumulative impact on	At horizon h=	Cubed MP shock	Significance level of difference	
			Driscoll-Kraay	Bootstrap
GDP	4	0.0001	0.4686	0.4920
	8	0.0023	0.1405	0.7767
	12	0.0085	0.0014	0.8827
	16	0.0106	0.0001	0.9278
Inflation	4	-0.0013	0.1113	0.2387
	8	-0.0009	0.3678	0.3285
	12	0.0021	0.2817	0.5221
	16	0.0061	0.0472	0.8332

Table IV: Cumulative impulse response of GDP and inflation: linearly identified shocks

Cumulative impact on	At horizon h=	Regime			
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0243	0.0117	0.0050	0.0899
	8	-0.0565	-0.0138	0.1000	0.1853
	12	-0.0939	-0.0257	0.0620	0.0621
	16	-0.0901	-0.0420	0.1846	0.1271
Inflation	4	0.0081	-0.0020	0.1655	0.8062
	8	-0.0100	0.0053	0.1970	0.5020
	12	-0.0601	0.0079	0.0033	0.1081
	16	-0.0973	-0.0163	0.0173	0.1218

Table V: Cumulative impulse response of GDP and inflation: VAR shocks

Cumulative impact on	At horizon h=	Regime			
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0239	0.0361	0.0000	0.1299
	8	-0.0570	0.0487	0.0031	0.1252
	12	-0.0956	0.0859	0.0008	0.0950
	16	-0.0978	0.1185	0.0015	0.1001
Inflation	4	0.0003	-0.0022	0.4138	0.4365
	8	-0.0271	0.0070	0.0455	0.1869
	12	-0.0850	0.0107	0.0018	0.1547
	16	-0.1338	-0.0176	0.0125	0.1737

Table VI: Cumulative response of GDP and inflation: no trend in regression

Cumulative impact on	At horizon h=	Regime			
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0239	0.0361	0.0000	0.1299
	8	-0.0570	0.0487	0.0031	0.1252
	12	-0.0956	0.0859	0.0008	0.0950
	16	-0.0978	0.1185	0.0015	0.1001
Inflation	4	0.0003	-0.0022	0.4138	0.4365
	8	-0.0271	0.0070	0.0455	0.1869
	12	-0.0850	0.0107	0.0018	0.1547
	16	-0.1338	-0.0176	0.0125	0.1737

Table VII: Cumulative response of GDP and inflation: two lags of LHS variable

Cumulative impact on	At horizon h=	Regime			
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0217	0.0102	0.0029	0.1158
	8	-0.0487	-0.0135	0.1180	0.2225
	12	-0.0798	-0.0199	0.0500	0.0840
	16	-0.0761	-0.0364	0.1881	0.1813
Inflation	4	0.0112	-0.0085	0.0144	0.8536
	8	0.0002	-0.0073	0.3068	0.7026
	12	-0.0361	-0.0083	0.0843	0.2780
	16	-0.0652	-0.0329	0.1565	0.3069

Table VIII: Cumulative impulse response of GDP and inflation: NBER state variable

Cumulative impact on	At horizon h=	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0091	0.0170	0.0445	0.4960
	8	-0.0303	0.0267	0.0168	0.3464
	12	-0.0413	0.0047	0.0785	0.3120
	16	-0.0370	-0.0446	0.4060	0.4080
Inflation	4	0.0044	-0.0158	0.0435	0.8068
	8	-0.0033	-0.0211	0.2384	0.6857
	12	-0.0196	-0.0500	0.2120	0.5844
	16	-0.0296	-0.1109	0.0503	0.6805

Table IX: Cumulative impulse response of GDP and inflation: HP filtered output as state variable

Cumulative impact on	At horizon h=	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	0.0006	-0.0028	0.3891	0.4355
	8	-0.0099	-0.0248	0.2799	0.4860
	12	-0.0222	-0.0535	0.2681	0.5420
	16	-0.0301	-0.0773	0.2488	0.5646
Inflation	4	0.0027	0.0008	0.4482	0.5301
	8	0.0017	-0.0262	0.1295	0.6362
	12	-0.0056	-0.0919	0.0326	0.7933
	16	-0.0172	-0.1280	0.0658	0.7558

Table X: Cumulative impulse response of GDP and inflation: phase shift in state variable

Cumulative impact on	At horizon h=	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0141	0.0165	0.0339	0.1059
	8	-0.0573	0.0206	0.0051	0.1503
	12	-0.0954	0.0129	0.0094	0.1291
	16	-0.1014	-0.0142	0.0674	0.1410
Inflation	4	0.0218	-0.0108	0.0926	0.8080
	8	0.0233	-0.0106	0.1960	0.7775
	12	-0.0170	0.0010	0.3437	0.5515
	16	-0.0557	-0.0017	0.1699	0.3201

Table XI: Cumulative impulse response of GDP and inflation: $\theta = 1$

Cumulative impact on	At horizon h=	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0366	0.0349	0.0025	0.1432
	8	-0.0791	0.0190	0.0334	0.2199
	12	-0.1245	0.0194	0.0253	0.1087
	16	-0.1089	-0.0213	0.1657	0.1543
Inflation	4	0.0085	-0.0017	0.2968	0.7168
	8	-0.0170	0.0198	0.1283	0.4524
	12	-0.0902	0.0607	0.0006	0.1007
	16	-0.1333	0.0419	0.0063	0.1278

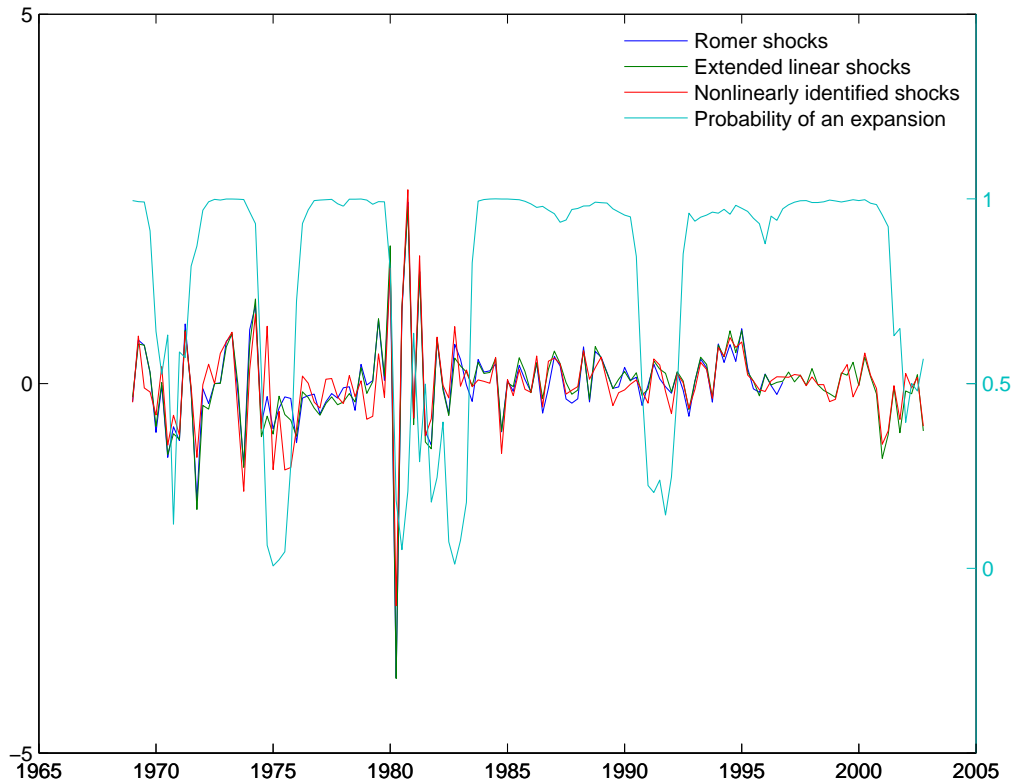
Table XII: Cumulative impulse response of GDP and inflation: $\theta = 10$

Cumulative impact on	At horizon h=	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0139	0.0028	0.0403	0.1027
	8	-0.0352	-0.0214	0.2798	0.2592
	12	-0.0662	-0.0323	0.1380	0.1013
	16	-0.0684	-0.0401	0.2154	0.1725
Inflation	4	0.0101	-0.0021	0.0389	0.8623
	8	0.0012	-0.0021	0.3982	0.7045
	12	-0.0340	-0.0101	0.0768	0.2682
	16	-0.0659	-0.0312	0.0888	0.2677

Table XIII: Cumulative response of GDP and inflation: $c = 0.5$

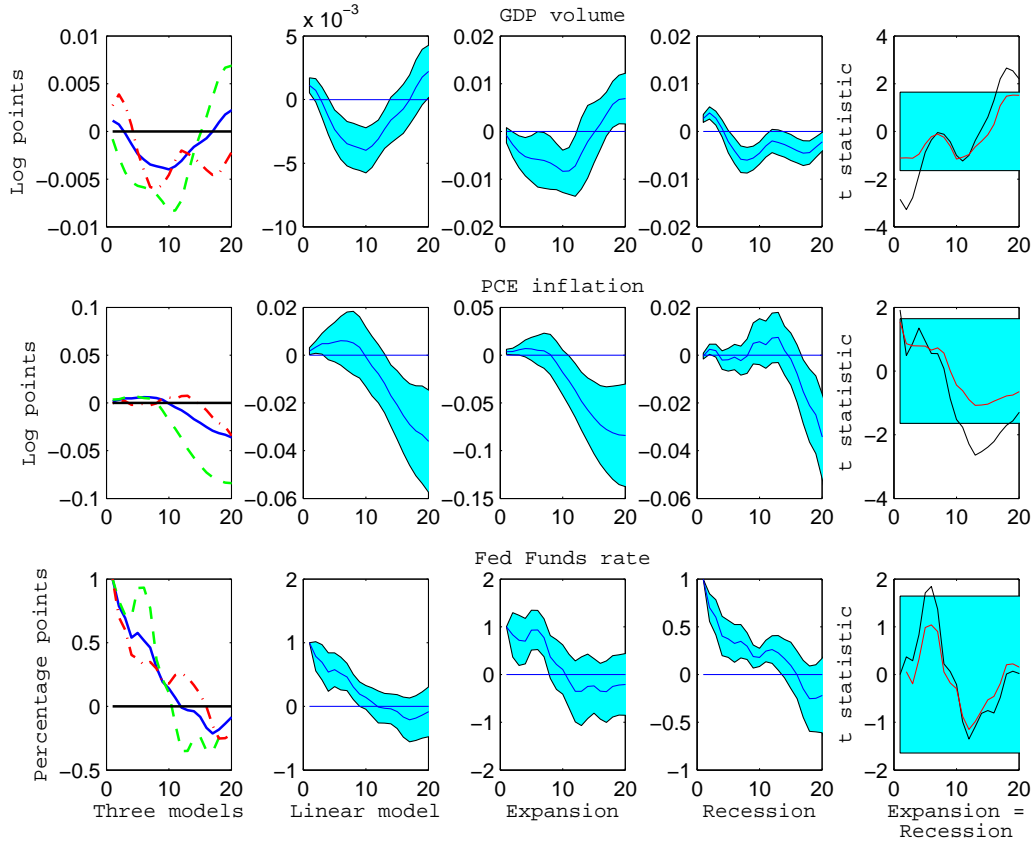
Cumulative impact on	At horizon h=	Regime		Significance level of difference	
		Expansion	Recession	Driscoll-Kraay	Bootstrap
GDP	4	-0.0141	0.0165	0.0339	0.1128
	8	-0.0573	0.0206	0.0051	0.1505
	12	-0.0954	0.0129	0.0094	0.1288
	16	-0.1014	-0.0142	0.0674	0.1446
Inflation	4	0.0218	-0.0108	0.0926	0.8050
	8	0.0233	-0.0106	0.1960	0.7684
	12	-0.0170	0.0010	0.3437	0.5485
	16	-0.0557	-0.0017	0.1699	0.3169

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



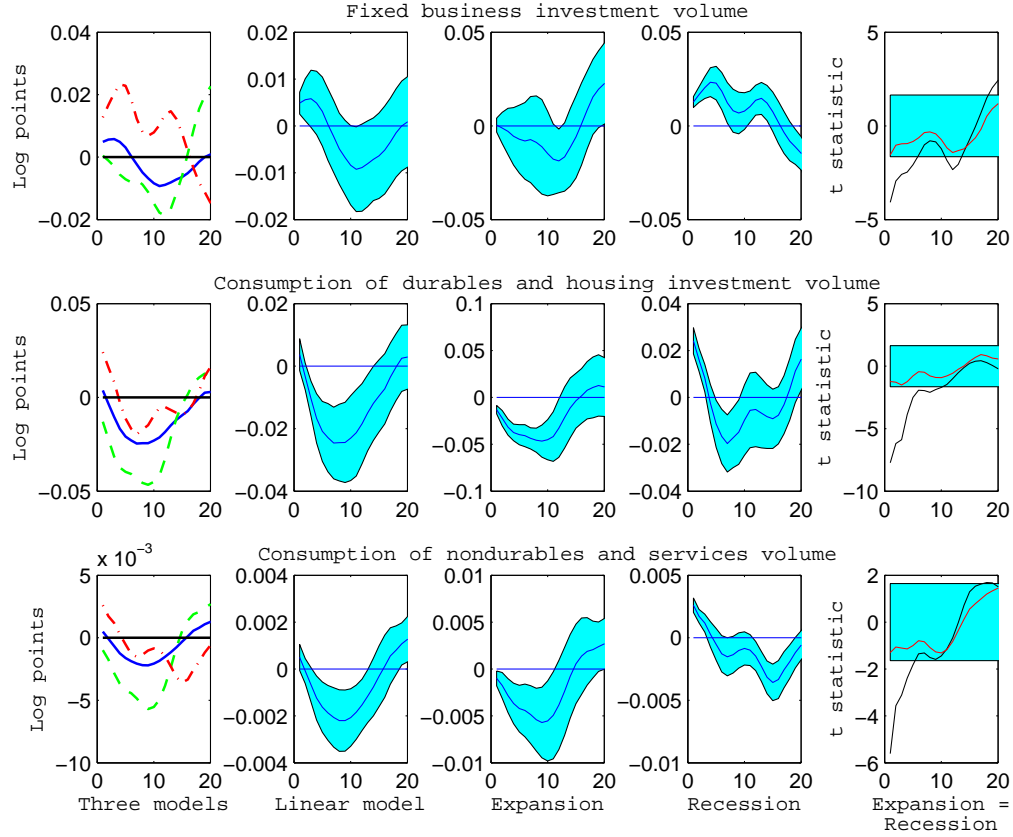
The blue line is the series of monetary policy in Romer and Romer (2004). The green line is constructed in an identical fashion but over a longer sample. The red line is constructed over the same longer sample but with a state-dependent regression model. The turquoise line is the value of the cdf of our state variable $F(z_t)$. See main text for details.

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



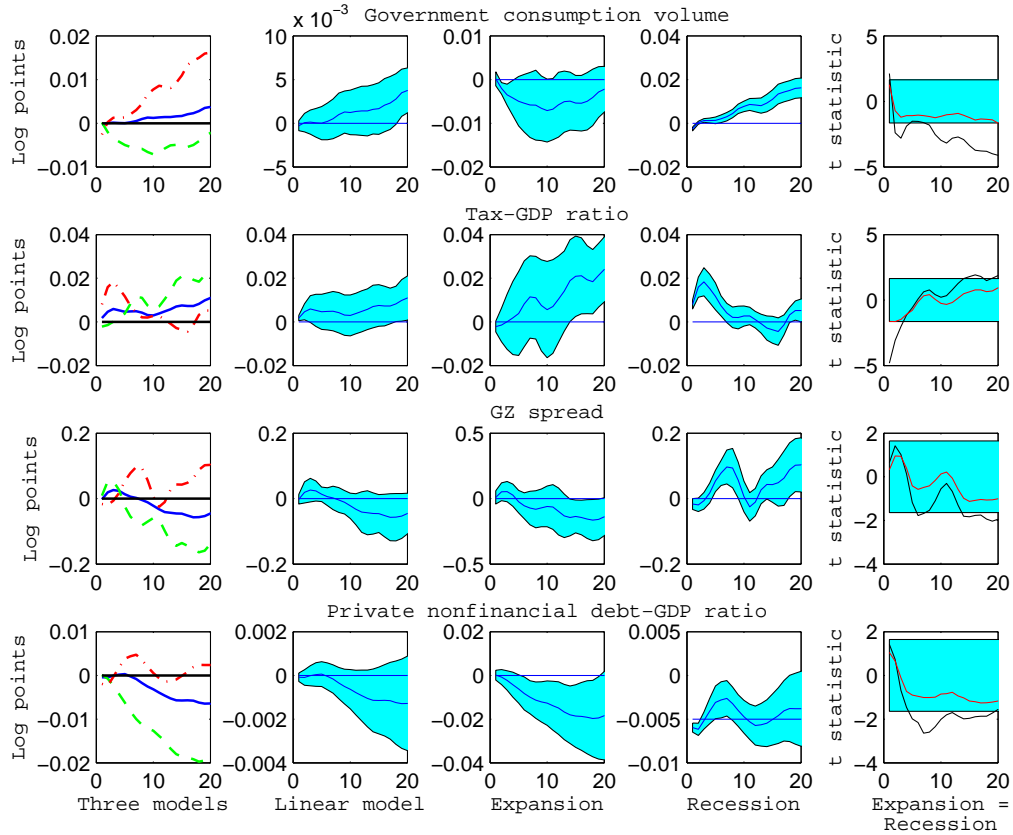
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure III: Impulse response of expenditure volumes



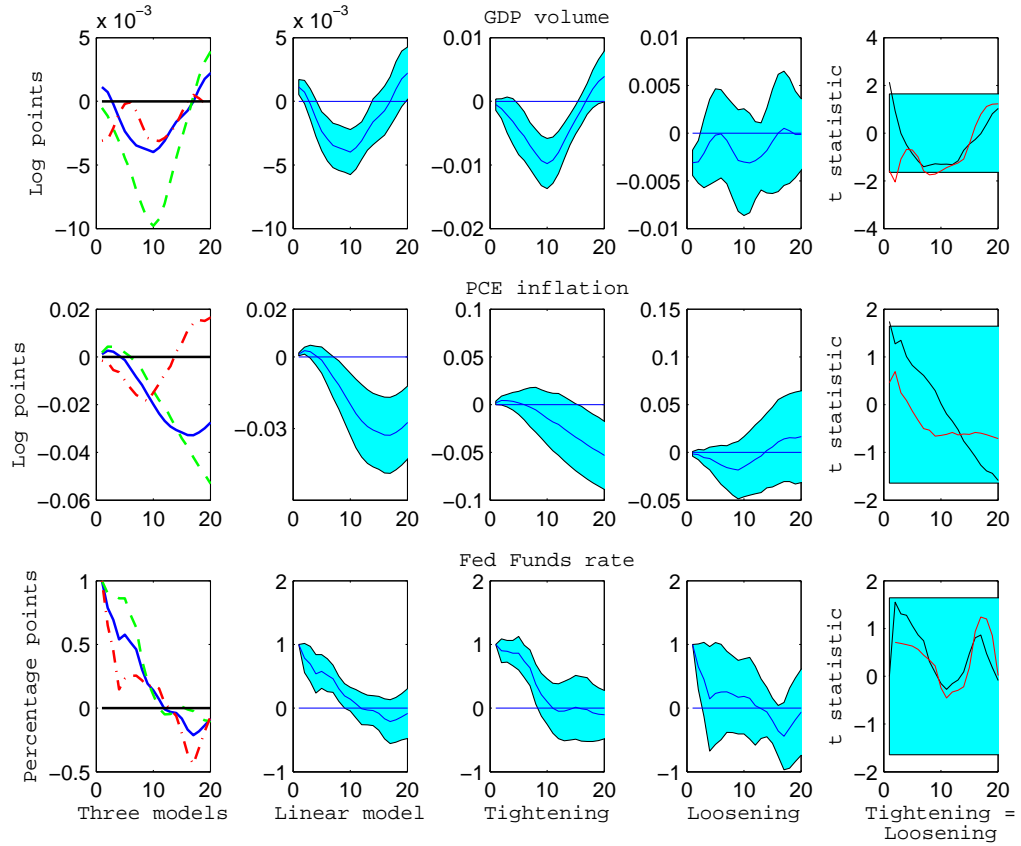
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of an index of real durables consumption and housing investment, the second row an index of real consumption of nondurable goods and services, and the third row an index of real fixed business investment.

Figure IV: Impulse response functions of fiscal and credit variables



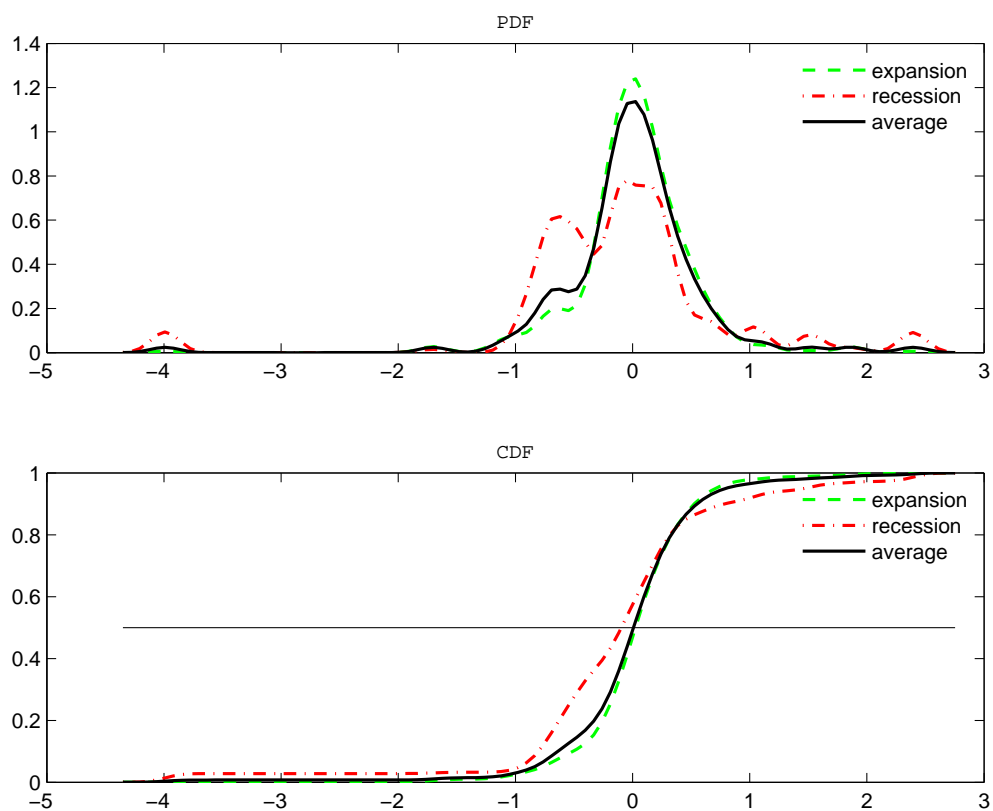
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of an index of real government consumption, the second row the level of the net tax-GDP ratio, the third row the Gilchrist-Zakrajsek bond spread index, and the fourth row the level of private nonfinancial debt to GDP.

Figure V: Impulse response to positive and negative monetary policy shocks



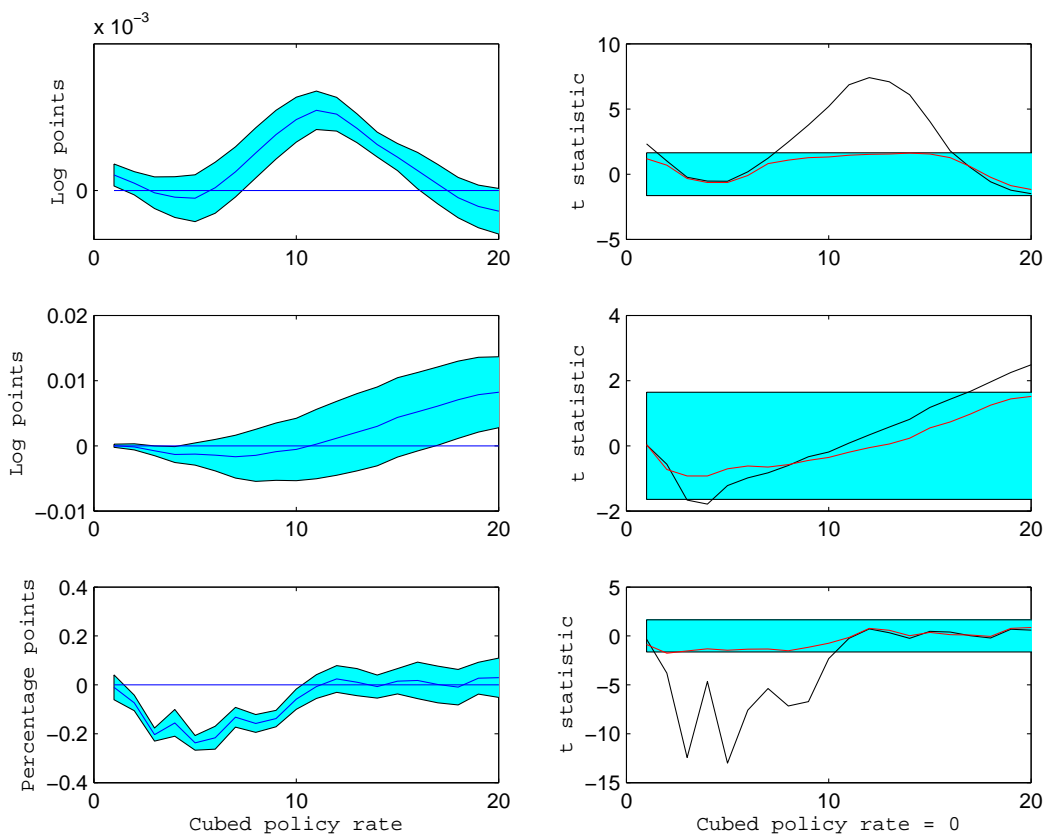
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response to a positive shock, and the red dotted line the response to a negative shock (normalised to have the same sign). The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response to a positive shock, and the fourth column the interval around the response to a negative shock. The fifth column shows t-statistics testing the hypothesis that the difference between the responses to positive and negative shocks is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

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



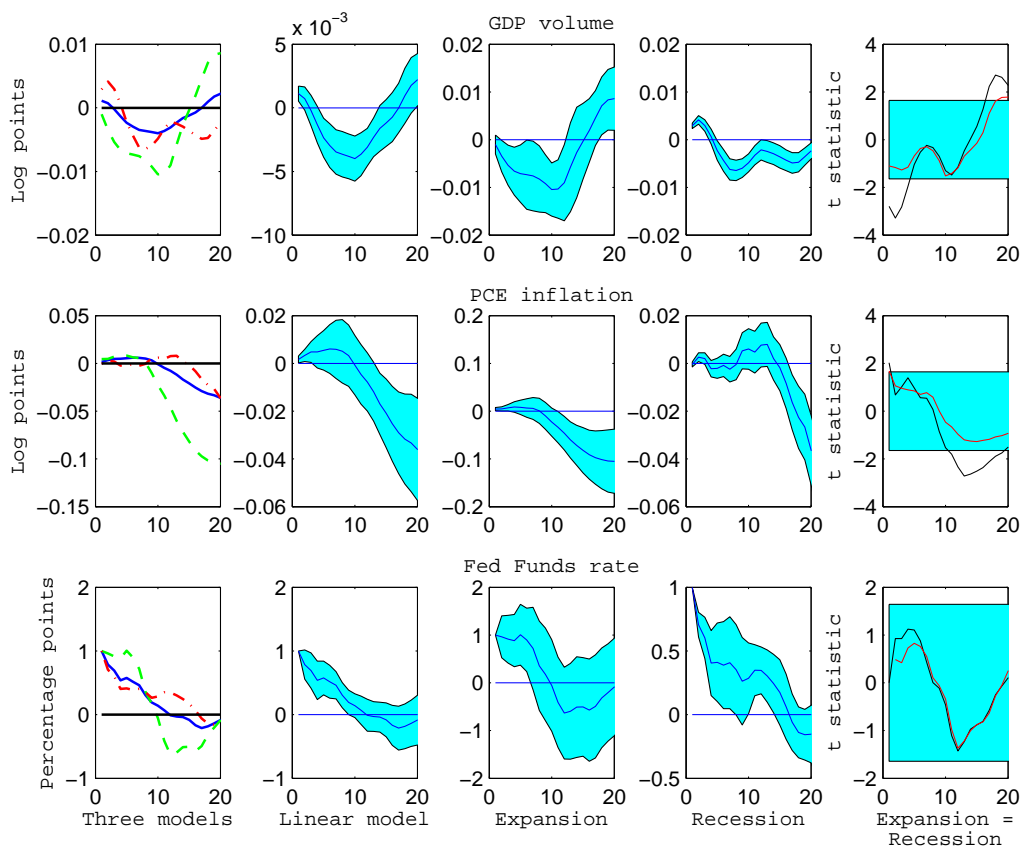
The top pane shows the pdf of the shocks in the different regimes. The bottom pane shows the cdf. The green lines show the distribution during an expansion, the red lines in a recession, and the black line the average of the two regimes.

Figure VII: Impulse response to cubed monetary policy shocks



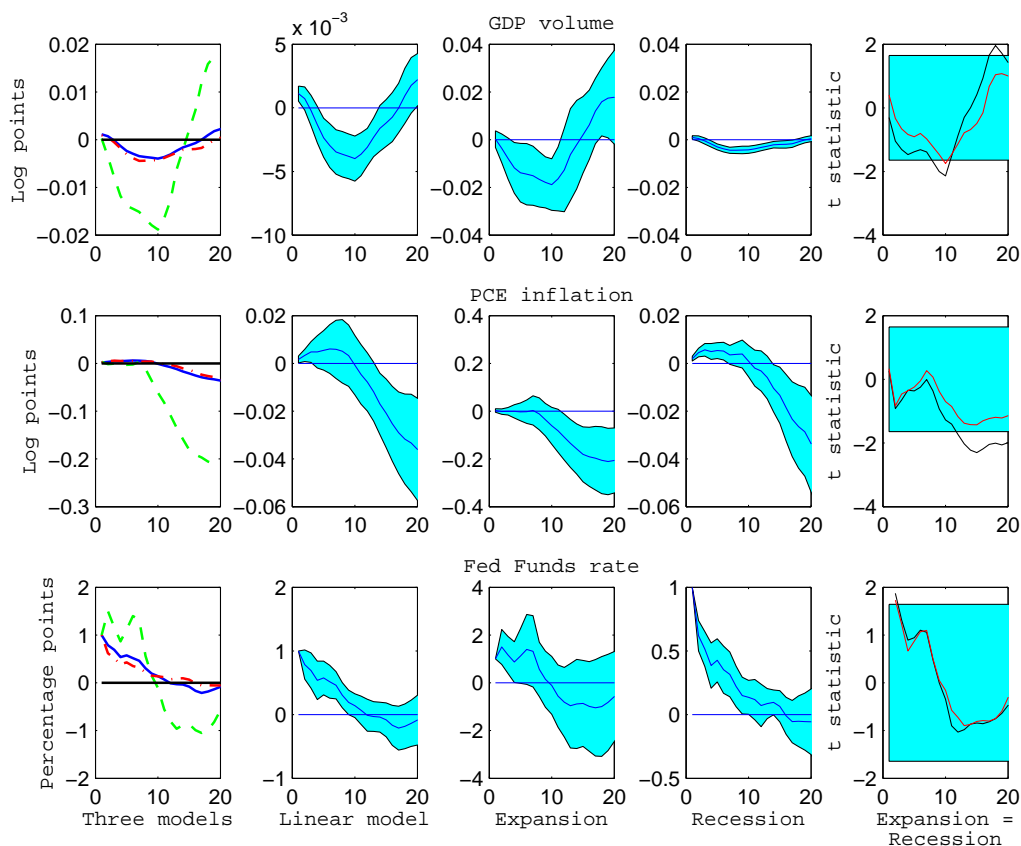
The left hand column shows point estimates and a 90 percent confidence interval for the impulse response on cubed monetary policy shocks ε_t^3 , i.e. β_h^l in the equation $y_{t+h} = \tau t + \alpha_h^b + \beta_h^s \varepsilon_t + \beta_h^l \varepsilon_t^3 + \gamma^{bl} x_t + u_t$. The right-hand column shows three estimates of the t-statistic testing the hypothesis that $\beta_h^l = 0$. The black line is calculated using a modified Newey-West method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The light blue shaded area is ± 1.65 .

Figure VIII: Impulse response of headline variables to monetary policy shocks identified linearly



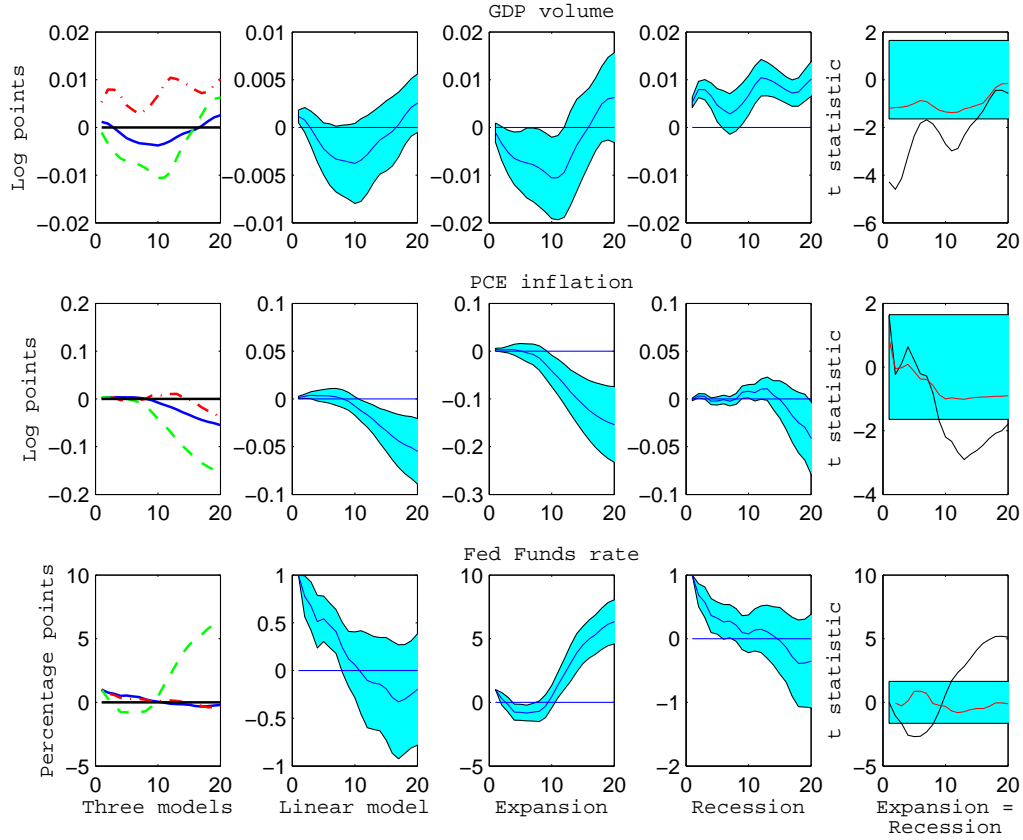
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure IX: Impulse response of headline variables to monetary policy shocks identified with a VAR



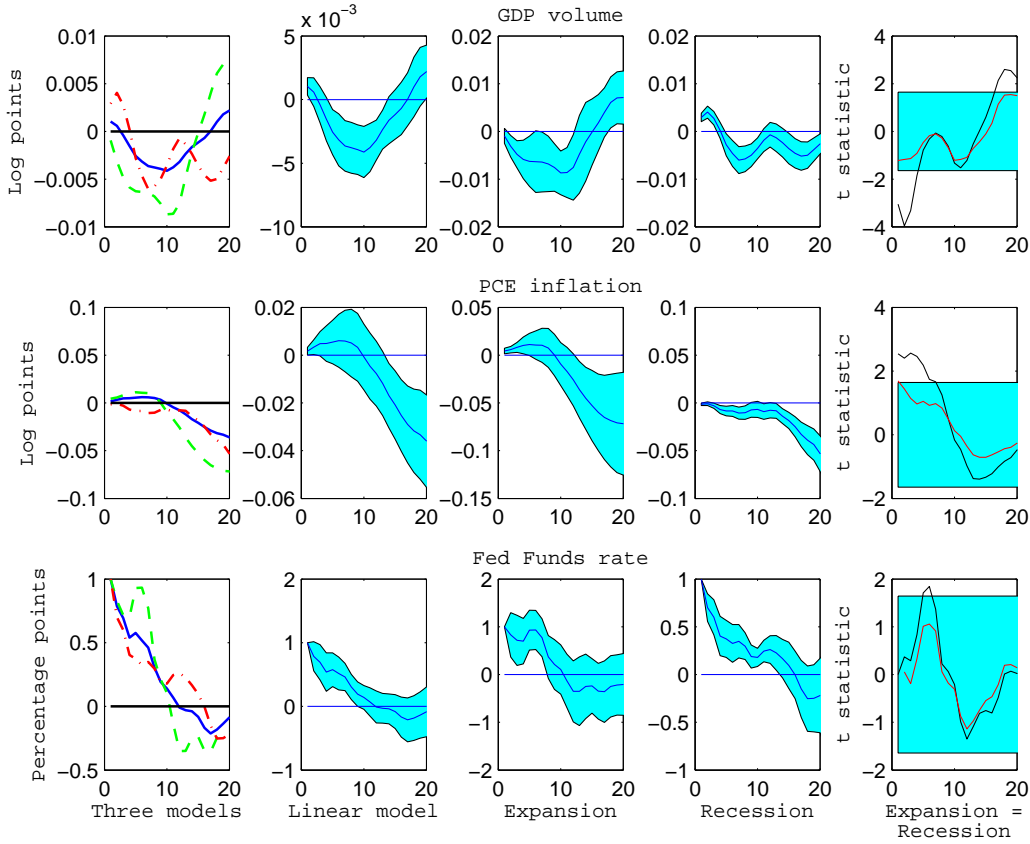
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure X: IRFs with no trend in regression equation



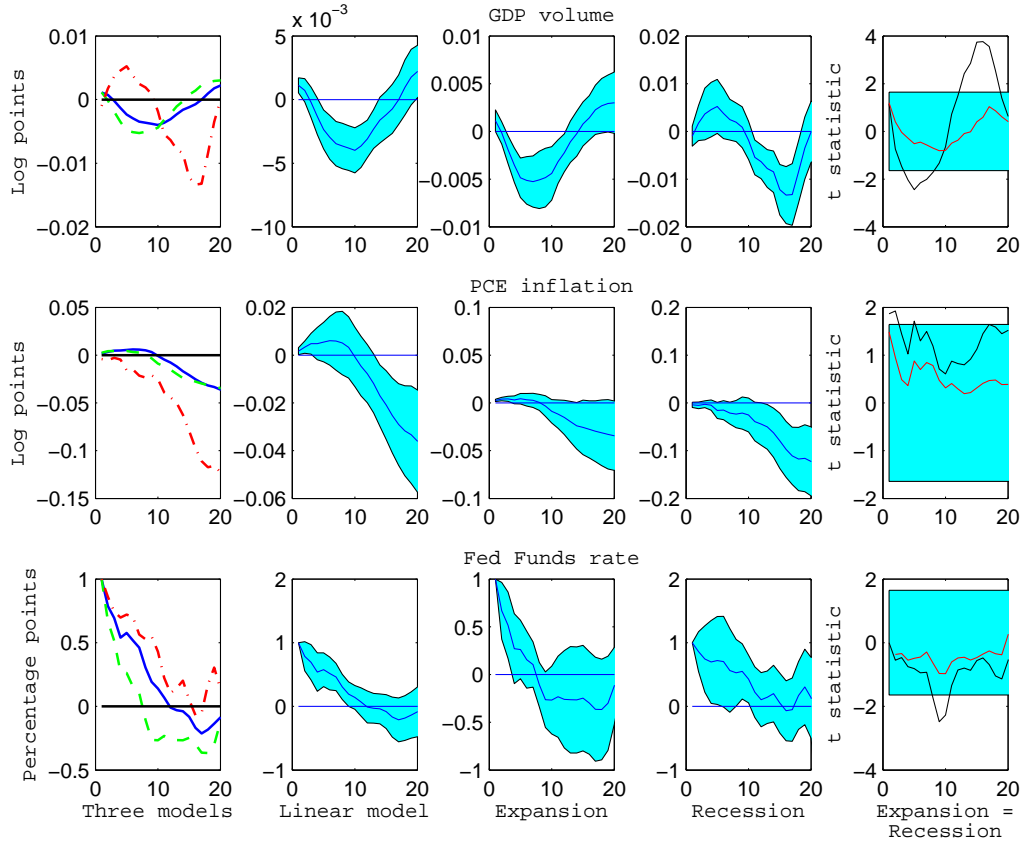
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure XI: IRFs with two lags in regression equation



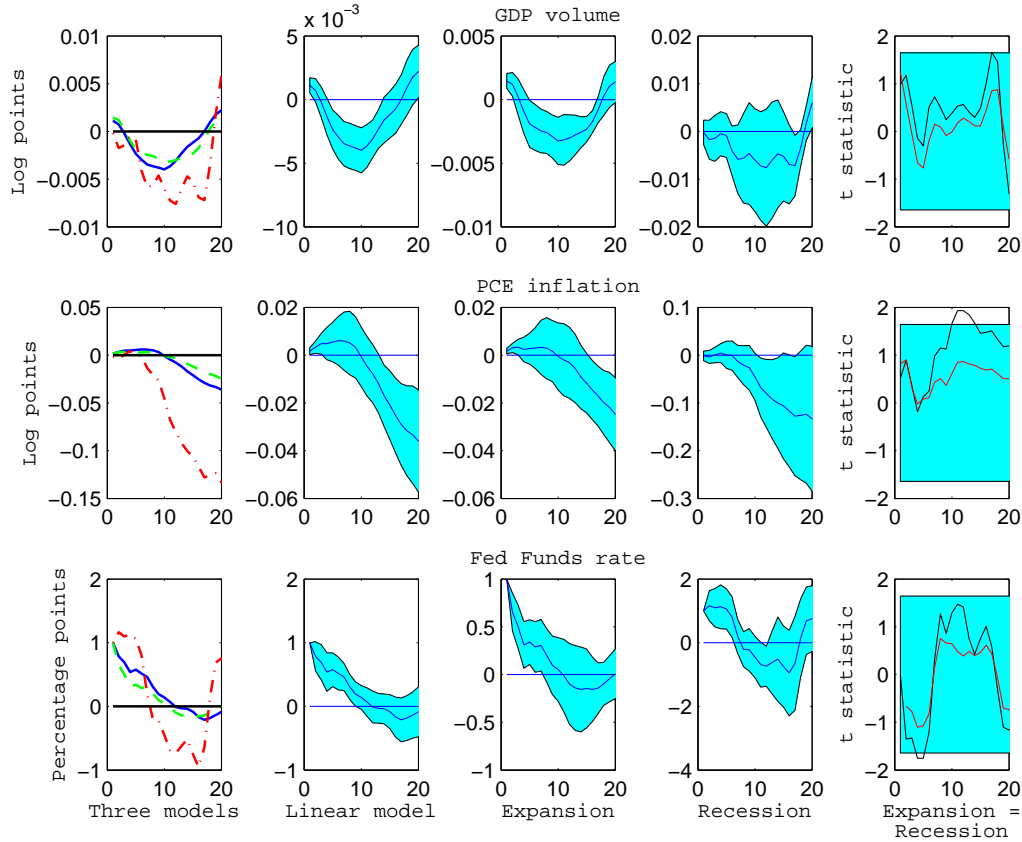
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure XII: IRFs with NBER recession state variable



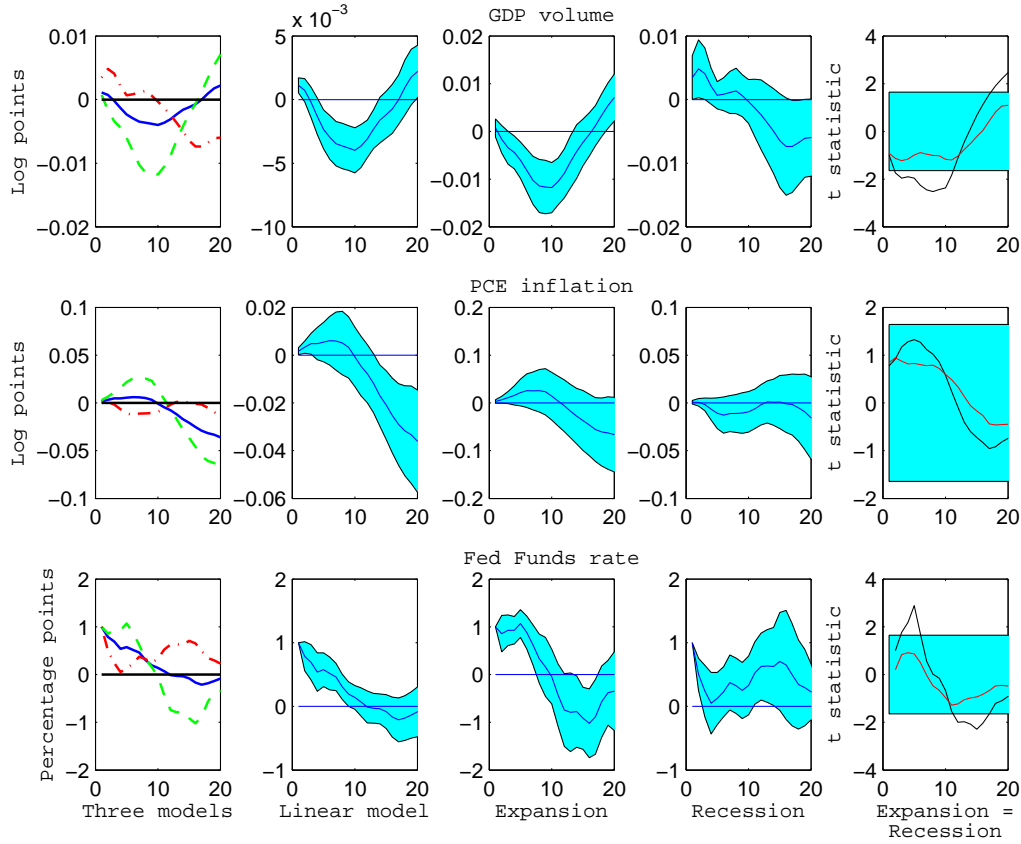
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure XIII: IRFs with HP-filtered output gap as state variable



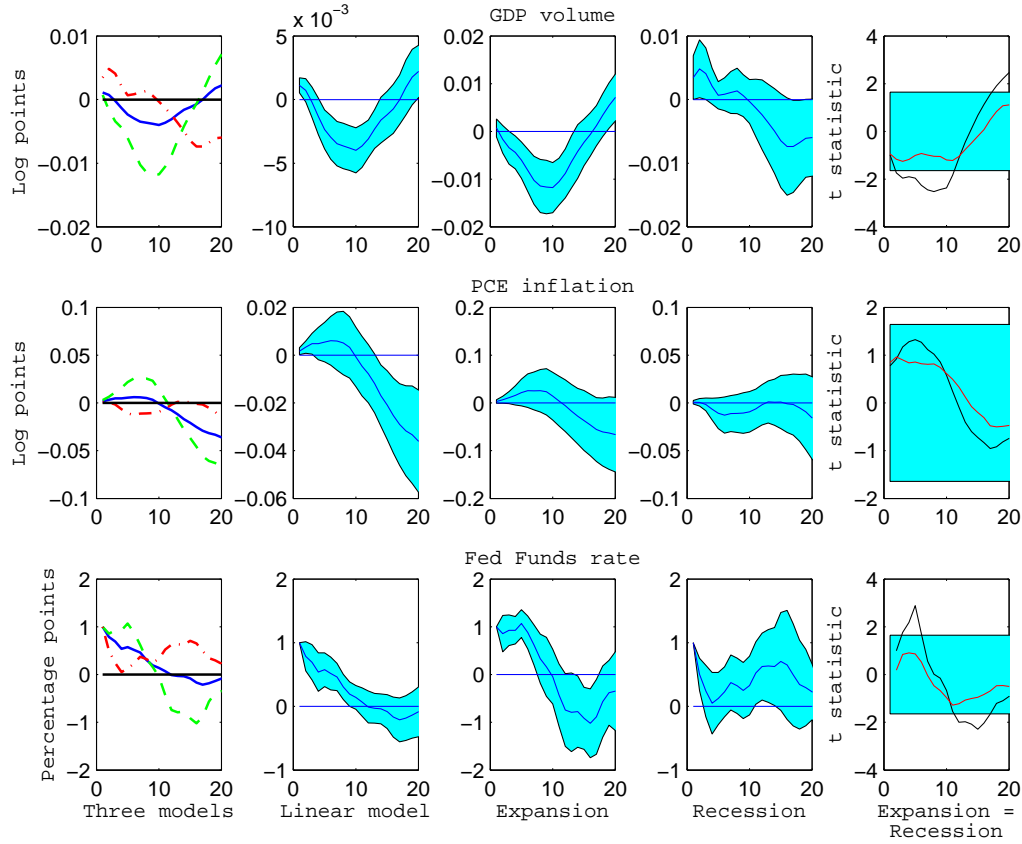
The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

Figure XIV: IRFs with centered state variable



The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.

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



The first four columns show the impulse response to a monetary policy shock that increases the Federal Funds rate by 1 percentage point on impact. In the first column, the solid blue line shows the response in a linear, state-independent model, the green dashed line shows the response in an expansion, and the red dotted line the response in a recession. The second column shows a 90 percent confidence interval around the state-independent response, the third column the same interval around the response in an expansion, and the fourth column the interval around the response in a recession. The fifth column shows t-statistics testing the hypothesis that the difference between the coefficients in an expansion and a recession is zero. The black line is calculated using the Driscoll-Kraay method, and the red line using a bootstrap approach (see main text for details). The light blue shaded area is ± 1.65 . The first row is the log-level of real GDP, the second row is the quarterly annualised inflation rate of the GDP deflator, and the third row is the level of the Federal Funds rate.