Gaussian Mixture Approximations of Impulse Responses and The Non-Linear Effects of Monetary Shocks*

Regis Barnichon
CREI, Universitat Pompeu Fabra, CEPR

Christian Matthes
Federal Reserve Bank of Richmond

January 2016
(first draft: March 2014)

Abstract

This paper proposes a new method to identify the (possibly non-linear) dynamic effects of structural shocks by using Gaussian basis functions to approximate impulse response functions. We apply our approach to the study of monetary policy and find that the effect of a monetary intervention depends strongly on the sign of the intervention and on the state of the labor market at the time of the intervention. A contractionary shock has a strong adverse effect on unemployment, larger than implied by linear estimates, but an expansionary shock has only a small effect. When the labor market is tight, an expansionary shock generates a burst of inflation and no significant change in unemployment. *JEL classifications: E24, E32*

*We would like to thank Luca Benati, Francesco Bianchi, Christian Brownlees, Fabio Canova, Tim Cogley, Davide Debortoli, Jordi Gali, Yuriy Gorodnichenko, Oscar Jorda, Thomas Lubik, Jim Nason, Kris Nimark, Giorgio Primiceri, Ricardo Reis, Barbara Rossi, Mark Watson, Yanos Zylberberg and seminar participants at the Barcelona GSE Summer Forum 2014, the 2014 NBER/Chicago Fed DSGE Workshop, William and Mary college, EUI Workshop on Time-Varying Coefficient Models, Oxford, Bank of England, NYU Alumni Conference, Society for Economic Dynamics Annual Meeting (Warsaw), Universitaet Bern, Econometric Society World Congress (Montreal), the Federal Reserve Board, the 2015 SciencesPo conference on Empirical Monetary Economics and the San Francisco Fed for helpful comments. Barnichon acknowledges financial support from the Spanish Ministerio de Economia y Competitividad (grant ECO2011-23188), the Generalitat de Catalunya (grant 2009SGR1157) and the Barcelona GSE Research Network. The views expressed here do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. Any errors are our own.

1 Introduction

There now exists a broad consensus on the average effect of monetary policy on economic activity, and it is widely accepted that a monetary contraction (expansion) leads to a persistent decline (increase) in output.

However, there is still little agreement about possible asymmetric or non-linear effects of monetary policy, and two questions at the core of monetary policy making are largely unsettled.[1] First, does monetary policy have asymmetric effects on economic activity? As captured by the string metaphor, does contractionary monetary policy have a much stronger effect –being akin to pulling on a string– than an expansionary shock –being akin to pushing on a string–? Second, does the effect of monetary policy vary with the state of the business cycle? For instance, does the central bank have more room to stimulate economic activity (without raising inflation) during recessions?

Providing answers to these questions has been difficult in part for one important technical reason: the standard approach to identify the dynamic effect of shocks relies on structural Vector-Autoregressions (VARs),[2] which are linear models. While VARs can accommodate certain types of non-linearities, some questions, such as the asymmetric effect of a monetary shock, cannot be easily answered within a VAR framework.

This paper proposes a new method to identify the (possibly non-linear) dynamic effects of structural shocks. Instead of assuming the existence of a VAR representation, our approach consists in working directly with the structural moving-average representation of the economy. Then, to make the estimation of the moving-average representation feasible, we approximate the impulse response functions with Gaussian basis functions.

Our approach builds on two premises: (i) any mean-reverting impulse response function can be approximated to any degree of accuracy by a mixture of Gaussian basis functions, and (ii),

[1]For instance, while Cover (1992) finds evidence of asymmetric effects, Ravn and Sola (1996, 2004) and Weise (1999) instead find nearly symmetric effects. And while Lo and Piger (2005) and Santoro et a. (2014) conclude that monetary policy has stronger effects during recessions, Tenreyro and Thwaites (2015) conclude the opposite.

[2]See e.g., Christiano, Eichenbaum, and Evans (1999) and Uhlig (2005).

in practice, only a very small number of Gaussian functions are needed to approximate a typical impulse response function. Intuitively, the impulse response functions of macroeconomic variables are often found to be monotonic or hump-shaped (e.g., Christiano, Eichenbaum, and Evans, 1999). In such cases, a single Gaussian function can already provide an excellent approximation of the impulse response function.

Thanks to the small number of free parameters allowed by a Gaussian Mixture Approximation (GMA), it is possible to directly estimate the impulse response functions from the data using Bayesian methods.[3] In turn, the parsimony of the approach allows us to estimate more general non-linear models.

We conduct a number of Monte-Carlo simulations to illustrate the performance of our approach in finite sample, first for linear models, then for non-linear models. In a linear model, we show that a GMA model can generate more accurate impulse response estimates (in a mean-squared error sense) than a well-specified VAR model. In a simulation with asymmetry and state-dependence, we find that a GMA model can accurately detect the presence of non-linearities and deliver good estimates of the magnitudes of the non-linearities.

We use our GMA approach to estimate the non-linear effects of monetary shocks identified with a recursive identification scheme.[4] Consistent with the string metaphor, our findings point towards the existence of strong asymmetries in the effects of monetary shocks. A contractionary shock has a strong adverse effect on output, larger than implied by linear estimates, but an expansionary shock has little effect on output. Interestingly, this asymmetry could be due the presence of downward price/wage rigidities. Although the evidence for inflation is more uncertain, we find that inflation displays a more marked price puzzle following a contractionary shock than following an expansionary shock. Finally, we also find that the effect of a monetary shock depends on the state of the business cycle at the time of the intervention:

[3] Another advantage of using Gaussian basis functions is that prior elicitation can be much easier than with Bayesian estimation of standard VARs, because the coefficients to be estimated are directly interpretable as features of impulse responses.

[4] While we introduce our GMA method in the context of a recursive identification scheme, our method is quite general and can also be applied to other population identification schemes, such as sign-restrictions (Uhlig, 2005) or long-run restrictions (Blanchard and Quah, 1989, Gali 1999).

An expansionary shock in a tight labor market generates no significant drop in unemployment but leads to a burst of inflation, consistent with a standard Keynesian narrative.

Although our use of Gaussian basis functions to model and estimate impulse response functions is new in the economics literature, our approach can be cast in the broader context of the machine (supervised) learning literature in that we project the function to be estimated onto the space spanned by a dictionary of basis functions (see Hastie, Tibshirani and Friedman, 2009). In basis functions methods, the number of basis functions is often too large for empirical purposes, and the complexity of the model is typically controlled through a combination of restriction, selection and/or regularization methods. Our approach, which consists in using a limited number of optimally chosen basis functions, uses both selection and restriction to control the complexity of the model.[5]

In economics, our parametrization of impulse responses relates to an older literature on distributed lag models and in particular the Almon (1965) lag specification, in which the successive weights, i.e., the impulse response function in our context, are given by a polynomial function.[6] Our use of basis functions of a Gaussian type relates to a large literature that relies on radial basis functions (of which Gaussian functions are one example) to approximate arbitrary multivariate functions (e.g., Buhmann, 2003) or to approximate arbitrary distributions using a mixture of Gaussian distributions (Alspach and Sorenson 1971, 1972, McLachlan and Peel, 2000). Although Gaussian basis functions provide a more natural and more parsimonious way than polynomials to approximate mean-reverting impulse response functions, our approach is general and other basis functions are possible. For instance, the inverse quadratic function, which is also a popular radial basis function, could be used to parametrize impulse response functions.[7] Finally, our approach shares with the non-parametric econometrics liter-

[5] It uses selection in the sense that our algorithm scans the dictionary of possible basis functions to find the basis functions that best fit the data, and it uses restriction in the sense that we restrict ourselves to the class of impulse response functions that can be generated by a few basis functions.

[6] Recently, Plagborg-Moller (2016) proposes a Bayesian method to directly estimate the structural moving-average representation of the data by using prior information about the shape and the smoothness of the impulse response.

[7] In fact, in a different context, Jorgenson (1966) suggested that ratios of polynomials, of which the inverse quadratic function is one example, could be used to parametrize distributed lag functions.

ature (e.g., Racine, 2008) the insight that mixtures of Gaussian kernels can approximate very general shapes, although we use that insight in a very different manner.

The economic literature has so far tackled the estimation of non-linear effects of shocks in two main ways.[8]

A first approach estimates non-linear effects by regressing a variable of interest on contemporaneous and lagged values of the structural shocks while allowing for possible non-linear effects. In the context of monetary policy, Cover (1992), DeLong and Summers (1988) and Morgan (1993) proxy shocks with unanticipated money innovations (obtained from a money supply process regression, following Barro, 1977) and test whether the impulse response function depends on the sign of the shock. This approach has been recently revived thanks to the use of narratively identified shocks (Romer and Romer, 2002) and thanks to the Local Projection method pioneered by Jorda (2005).[9] The narrative approach was precisely developed in order to identify exogenous monetary innovations, and Jorda's method can easily accommodate non-linearities in the response function.[10] However, the Local Projection method is limited by efficiency consideration. Indeed, while the Local Projection approach is intentionally model-free –not imposing any underlying dynamic system–, this can come at an efficiency cost (Ramey, 2012), which makes inferences on a rich set of non-linearities (e.g., sign- and state-dependence) difficult. In contrast, by positing that the response function can be approximated by one (or a few) Gaussian functions, our approach imposes strong dynamic restrictions between the parameters of the impulse response function, which in turn allow us to estimate a rich set of non-linearities.[11] Another advantage of our approach is that it can be used for

[8] A third non-linear approach was recently proposed by Angrist et al. (2013) who develop a semi-parametric estimator to evaluate the (possibly asymmetric) effects of monetary policy interventions. They find asymmetric effects of monetary shocks consistent with our findings.

[9] The combination of Jorda's method with narratively identified shocks was first introduced in the context of fiscal policy by Auerbach and Gorodnichenko (2013) in order to test for the existence of state dependence in the effects of fiscal policy.

[10] Tenreyro and Thwaites (2013) and Santoro et al. (2014) use the Jorda method to estimate the extent of state dependence in the effect of monetary policy.

[11] Naturally, this statement also implies that our results are valid under the assumption that response functions can be well approximated by a few Gaussian functions. In this respect, our approach is best seen as complementing the model-free approach of Jorda (2005).

model selection and model evaluation through marginal density comparisons.

A second strand in the literature has relied on regime-switching VAR models –notably threshold VARs (e.g., Hubrich and Terasvirta, 2013) and Markov-switching VARs (Hamilton, 1989)– to capture certain types of non-linearities.[12,13] However, such non-linear models are ill suited to identify how the impulse response to a structural shock depends on the value of that shock. The threshold variable in a threshold VAR applies to a switching variable, which is not the contemporaneous structural shock itself but instead a function of past shocks. In a Markov-switching model the switching variable is an unobserved state, but Markov-switching models are ill-suited to capture regime changes triggered by shocks. In contrast, our approach allows us to estimate how the impulse response function varies with the contemporaneous value of the shock. Our model with state dependence is a form of regime-switching model, in that we allow the effect of a shock to depend on the state of the business cycle. However, in our framework state dependence is a continuous function of an indicator variable, while regime-switching models display discrete switches between a finite number of regimes.

Section 2 describes how we approximate impulse responses using mixtures of Gaussians, Section 3 discusses the key steps of the estimation methodology; Section 4 generalizes our approach to non-linear models; Section 5 presents Monte Carlo simulations to evaluate the performance of our approach in finite sample, first for linear models, then for non-linear models; Section 6 applies GMA to the study of the non-linear effects of monetary shocks using US data; Section 7 concludes.

2 Gaussian Mixture Approximations

This section presents a new method to estimate impulse responses using Gaussian Mixture Approximations (GMA) of the structural moving-average representation of the economy. Al-

[12]For examples in the monetary policy literature, see Beaudry and Koop (1993), Thoma (1994), Potter (1995), Kandil (1995), Koop, Pesaran and Potter (1996), Koop and Potter, (1998), Ravn and Sola (1996, 2004), Weise (1999), Lo and Piger (2005).

[13]Another prominent class of non-linear VARs includes models with time-varying coefficients and/or time-varying volatilities (e.g., Primiceri, 2005).

though the use of GMAs was motivated in the introduction by the need to model and estimate certain types of non-linearities, the intuition and benefits of GMA models can be understood in a linear context, and this section introduces GMAs in a linear context. We postpone the modeling and estimation of non-linearities to Section 4.

2.1 A structural moving average representation

Our starting point is a structural moving-average description of the economy, in which the behavior of a system of macroeconomic variables is dictated by its response to past and present *structural* shocks. Specifically, denoting $\mathbf{Y}_t$ a vector of stationary macroeconomic variables, the economy is described by

$$\mathbf{Y}_t = \sum_{k=0}^{K} \mathbf{\Psi}_k \varepsilon_{t-k} \tag{1}$$

where boldface letters indicate vectors or matrices, ε_t is the vector of structural innovations with $E\varepsilon_t = \mathbf{0}$ and $E\varepsilon_t \varepsilon_t' = \mathbf{I}$, and K is the number of lags, which can be finite or infinite. The matrices $\{\mathbf{\Psi}_k\}$ are the coefficients of the impulse response functions to shocks. For now, the model is linear, and the $\{\mathbf{\Psi}_k\}$ matrices are fixed.

If (1) is invertible and admits a VAR representation, the model can be estimated from a VAR on $\mathbf{Y}_t$ (provided some structural identifying assumption, such as the recursive ordering of $\mathbf{\Psi}_0$). However, assuming the existence of a VAR representation can be restrictive. In particular, in a non-linear world where $\mathbf{\Psi}_k$ depends on the value of ε_{t-k} (for instance, when the impulse response function varies with the sign of the shock), the existence of a VAR is compromised, because (1) is unlikely to be invertible. Thus, in this paper, we propose an alternative method that side-steps the need to invert (1), i.e., that a method that side-steps the need for a VAR representation.

2.2 Gaussian Mixture Approximations of impulse response functions

Rather than looking for a VAR representation of the dynamic system (1), our aim is to directly estimate (1), the moving-average representation of the economy. Because the number of free parameters $\{\mathbf{\Psi}_k\}$ in (1) is very large or possibly infinite, our strategy consists in parameterizing the impulse response functions, and more precisely in using mixtures of Gaussian functions to approximate each impulse response function.

2.2.1 Theoretical background

Our parametrization of the impulse response functions builds on the following theorem, which states that any integrable function can approximated with a sum of Gaussian functions.

Theorem 1: Let f be a bounded continuous function on $\mathbb{R}$ that satisfies $\int_{-\infty}^{\infty} f(x)^2 dx < \infty$. There exists a function f_N defined by

$$f_N(x) = \sum_{n=1}^{N} a_n e^{-(\frac{x-b_n}{c_n})^2}$$

with a_n, b_n, $c_n \in \mathbb{R}$ for $n \in \mathbb{N}$, such that the sequence $\{f_N\}$ converges pointwise to f on every interval of $\mathbb{R}$.

Proof: See Appendix.

Motivated by Theorem 1, our approach will consist in approximating an impulse response function $\psi(.)$ with a sum of Gaussian functions, that is

$$\psi(k) \simeq \sum_{n=1}^{N} a_n e^{-(\frac{k-b_n}{c_n})^2} \tag{2}$$

with a_n, b_n, $c_n \in \mathbb{R}$ for k over some interval of $\mathbb{R}_+$.

Since our strategy consists in approximating impulse response functions with mixtures of Gaussians, we refer to this class of models as Gaussian Mixture Approximations (GMA), with a $GMA(N)$ denoting a GMA with N Gaussian basis functions.

2.2.2 Intuition and Motivation

Before describing the estimation of GMA models, it is instructive to first intuitively discuss the benefits of our approach over traditional VARs.

The advantage of our approach, and its use for studying the (possibly non-linear) effects of policy, will rest on the fact that, in practice, only a very small number of Gaussian basis functions are needed to approximate a typical impulse response function, allowing for efficiency gains and opening the door to estimating non-linearities.

Intuitively, impulse response functions of stationary variables are often found to be monotonic or hump-shaped (e.g., Christiano, Eichenbaum, and Evans, 1999).[14] And in such cases, a single Gaussian function can already provide a good approximate description of the impulse response. To illustrate this observation, Figure 1 plots the impulse response functions of unemployment, the price level and the fed funds rate to a monetary shock estimated from a standard VAR specification,[15] along with the corresponding $GMA(1)$, the Gaussian approximations with only *one* Gaussian function, i.e., using the approximation

$$\psi(k) \simeq ae^{-\frac{(k-b)^2}{c^2}}. \tag{3}$$

We can see that a $GMA(1)$ already does a good job at capturing the impulse responses implied by the VAR.[16] With a $GMA(2)$, the impulse responses are virtually on top on those of the VAR (Figure 1). For illustration, Figure 2 plots the Gaussian basis functions used for each impulse response in the GMA(2) case.

In both cases, the number of free parameters is manageable. For instance, in this 3 variables example, a GMA(1) only has 27 parameters (9 impulse responses times 3 parameters per

[14] This is also the case in theoretical models, e.g., New-Keynesian models, in which the impulse response functions are generally monotonic or hump-shaped (see e.g., Walsh, 2010).

[15] See Section 6 for the exact specification of the SVAR behind Figure 1. The VAR is specified with unemployment, PCE inflation and the fed funds rate. The impulse response for the price level is calculated from the response of inflation.

[16] In Figure 1, the parameters of the GMA (the a, b and c coefficients) were set to minimize the discrepancy (sum of squared residuals) between the two sets of impulse responses.

impulse response, ignoring intercepts) to capture the whole set of impulse responses $\{\mathbf{\Psi}_k\}_{k=1}^{K}$, while a GMA(2) has 48 free parameters ($9 * 3 * 2 = 48$).[17]

This relatively small number of free parameters has two main advantages. First, it allows us to directly estimate the impulse response functions from the vector moving-average representation (1), something that would otherwise be infeasible in finite sample (without additional assumptions). This point is important, because being able to directly work with the moving-average representation will allow us to estimate models in which shocks can have non-linear effects. Second, the parsimonious representation offered by GMA models may offer efficiency gains (relative to VARs) by tuning the bias-variance trade-off: GMA models aim to achieve lower variance by restricting the dimension of the parameter space, while tolerating more bias by restricting impulse response functions to belong to a certain class of functions.[18] These efficiency gains can be interesting not only in non-linear models but also in linear models.

To conclude this intuition section, we comment on a particularly interesting case: the $GMA(1)$ model, which has two additional advantages: (i) ease of interpretation, and (ii) ease of prior elicitation.

In a $GMA(1)$ model like (3), the a, b and c coefficients can be easily interpreted, because the impulse response function is summarized by three parameters –the peak effect, the time to peak effect, and the persistence of the impulse response–, which are generally considered the most relevant characteristics of an impulse response function.[19] As illustrated in Figure 3, parameter a is the height of the impulse-response, which corresponds to the maximum effect of a unit shock, parameter b is the timing of this maximum effect, and parameter c captures the persistence of the effect of the shock, as the amount of time τ required for the effect of a shock to be 50% of its maximum value is given by $\tau = c\sqrt{\ln 2.}$

[17]For comparison, a corresponding quarterly VAR with 3 variables and 4 lags has $4 * 3^2 + 6 = 42$ free parameters, and a monthly VAR with 12 lags has $12 * 3^2 + 6 = 114$ free parameters.

[18]Note that a GMA model will only be more biased than a VAR model if the true data generating process is actually a VAR. If the VAR is mis-specified or if the data generating process cannot be described by a VAR, a GMA model could display both lower variance *and* lower bias.

[19]For instance, when comparing the effects of monetary shocks across different specifications, Coibion (2012) focuses on the peak effect of the monetary shock, which in a GMA(1) model is simply parameter a.

Then, the ease of interpretation of the a, b and c parameters in turn makes prior elicitation easier than in standard VARs, in which the VAR coefficients have a less direct economic interpretation.

3 Bayesian estimation

To estimate our model, we use a Bayesian approach, which is particularly well suited for models that approximate the true DGP (Fernandez-Villaverde and Rubio-Ramirez, 2004). In particular, Bayes factors will allow us to evaluate GMA models against VAR models, even though the two classes of models are non-nested. Bayesian model comparison will also offer us a natural way to select the order of the GMA model, i.e., the number of Gaussian basis functions used in the approximation.

In this section, we describe the implementation and estimation of GMA models. We first discuss the structural identifying assumption, then describe how we construct the likelihood function by exploiting the prediction-error decomposition, discuss the estimation routine based on a multiple-block Metropolis-Hasting algorithm, and finally discuss prior elicitation, determination of the order of the GMA and identification issues related to fundamentalness. We conclude by discussing how to deal with non-stationary data.

3.1 Structural identifying assumption

Model (1) is under-identified without additional restrictions. As is common with structural VARs, we will assume that the model is just-identified and that restrictions on the contemporaneous impact matrix $\mathbf{\Psi}_0$ provide use with $\frac{N(N-1)}{2}$ additional restrictions. In fact, given our later focus on monetary policy, we will adopt a common recursive assumption, so that the contemporaneous impact matrix $\mathbf{\Psi}_0$ is assumed to be lower triangular with positive entries on the diagonal (a normalization).

3.2 Constructing the likelihood function

We now describe how to construct the likelihood function $p(y^T|\theta, Z^T)$ of a sample of size T for the moving-average model (1) with parameter vector θ where superscripts denote the sample of variables up to the date in the superscript.

To start, we use the prediction error decomposition to break up the density $p(y^T|\theta)$ as follows:[20]

$$p(\mathbf{Y}^T|\theta) = \prod_{t=1}^{T} p(\mathbf{Y}_t|\theta, \mathbf{Y}^{t-1}). \tag{4}$$

To calculate the one-step-ahead conditional likelihood function needed for the prediction error decomposition, we assume that all innovations $\{\varepsilon_t\}$ are Gaussian with mean zero and variance one,[21] and we note that the density $p(\mathbf{Y}_{t+1}|\mathbf{Y}^t, \theta)$ can be re-written as $p(\mathbf{Y}_{t+1}|\theta, \mathbf{Y}^t) = p(\mathbf{\Psi}_0\varepsilon_{t+1}|\theta, \mathbf{Y}^t)$ since

$$\mathbf{\Psi}_0\varepsilon_{t+1} = \mathbf{Y}_{t+1} - \sum_{k=0}^{K} \mathbf{\Psi}_k\varepsilon_{t-k}. \tag{5}$$

Since the contemporaneous impact matrix is a constant, $p(\mathbf{\Psi}_0\varepsilon_{t+1}|\theta, \mathbf{Y}^t)$ is a straightforward function of the density of ε_{t+1}.

To recursively construct ε_{t+1} as a function of θ and $\mathbf{Y}^t$, we need to uniquely pin down the value of the components of ε_{t+1}, that is we need that $\mathbf{\Psi}_0$ is invertible, which is guaranteed by the structural identification assumption discussed above ($\mathbf{\Psi}_0$ is lower triangular).[22]

Finally, to initialize the recursion, we set the first K innovations $\{\varepsilon_j\}_{j=-K}^{0}$ to zero.[23,24]

[20] To derive the conditional densities in decomposition (4), our parameter vector θ thus implicitly also includes the K initial values of the shocks: $\{\varepsilon_{-K}...\varepsilon_0\}$. We will keep those fixed throughout the estimation and discuss alternative initializations below.

[21] The estimation could easily be generalized to allow for non-normal innovations such as t-distributed errors.

[22] Importantly, we can see that our approach accommodates other structural identification schemes, as long as $\mathbf{\Psi}_0$ is invertible: short-run restrictions (restrictions on $\mathbf{\Psi}_0$), long-run restrictions (restrictions of the type $\sum_{k=0}^{K} \psi_{k,ij} = 0$ for some (i,j), Blanchard and Quah, 1989, Gali 1999) or sign-restrictions (restrictions on the signs of the coefficients of $\mathbf{\Psi}_0$, e.g., Uhlig, 2005). Although a detailed exploration of such possibilities is left for other applications, these identifying restrictions could be easily imposed in our MCMC estimation routine.

[23] Alternatively, we could use the first K values of the shocks recovered from a structural VAR.

[24] When K, the lag length of the moving average (1), is infinite, we truncate the model at some horizon K, large enough to ensure that the lag matrix coefficients $\mathbf{\Psi}_k$ are "close" to zero. Such a K exists since the variables are stationary.

3.3 Estimation routine

To estimate our model, we use a Metropolis-within-Gibbs algorithm (Robert & Casella 2004, Haario et al., 2001) with the blocks given by the different groups of parameters in our model (one block being composed of the a parameters, another composed of the b parameters and so on).

To initialize the Metropolis-Hastings algorithm in an area of the parameter space that has substantial posterior probability, we follow a two-step procedure: first, we estimate a standard VAR using OLS on our data set, calculate the moving-average representation, and we use the impulse response functions implied by the VAR as our starting point. More specifically, we calculate the parameters of our GMA model to best fit the VAR-based impulse response functions.[25] Second, we use these parameters as a starting point for a simplex maximization routine that then gives us a starting value for the Metropolis-Hastings algorithm.

3.4 Prior elicitation

We use (loose) Normal priors centered around the impulse response functions obtained from the benchmark (linear) VAR. Specifically, we put priors on the a, b and c coefficients that are centered on the values for a, b and c obtained by matching the impulse responses obtained from the VAR, as described in the previous paragraph.

Specifically, denote $a^0_{ij,n}$, $b^0_{ij,n}$ and $c^0_{ij,n}$, $n \in \{1, N\}$ the values implied by fitting the GMA(N) to the VAR-based impulse response of variable i to shock j. The priors for $a_{ij,n}$, $b_{ij,n}$ and $c_{ij,n}$ are centered on $a^0_{ij,n}$, $b^0_{ij,n}$ and $c^0_{ij,n}$, and the corresponding standard-deviations are set as follows: $\sigma_{ij,a} = |a_{ij,0}|$, $\sigma_{ij,b} = K$ and $\sigma_{ij,c} = K^2$ (recall that K is the length of the moving-average). While there is clearly some arbitrariness in choosing the tightness of our priors, it is important to note that they are sufficiently loose to let us explore a large class of alternative specifications.

[25]Specifically, we set the parameters of our model (the a, b and c coefficients) to minimize the discrepancy (sum of squared residuals) between the two sets of impulse responses.

The use of informative priors is not critical for our approach, but we do this for a number of reason. First, since our current knowledge on the effect of monetary shocks is based to a large extent on VAR evidence of the kind reported in figure 1, it seems natural (and consistent with the Bayesian approach) to impose priors centered on our current state of knowledge. Second, given the inherent difficulty in estimating moving-average models, the priors help discipline the estimation by keeping the parameters in a reasonable set of the parameter space. Finally, and while we could have used improper uniform prior, the use of proper priors allows us to compute posterior odds ratio, which are important to select the order of the moving-average and to compare different GMA models.

3.5 Choosing N, the number of Gaussian basis functions

To choose N, the order of the GMA model, we use posterior odds ratios (assigning equal probability to any two model) to compare models with increasing number of mixtures. We select the model with the highest posterior odds ratio.[26]

3.6 Fundamentalness

In a linear moving average model, different representations (i.e., different sets of coefficients and innovation variances) can exhibit the same first two moments, so that with Gaussian-distributed innovations, the likelihood can display multiple peaks, and the moving average model is inherently underidentified. Since a GMA model works off directly with the moving-average representation, it cannot distinguish between invertible (also called "fundamental") and non-invertible representations. By using the VAR-based impulse responses as starting values, we implicitly focus on the invertible part of the parameter space.[27,28]

[26]This approach can be seen as analogous to the choice of the parameter lag in VAR models. While the Wold theorem shows that any covariance-stationary series can be written as a VAR(∞), one must select a finite lag order p that reasonable approximate the VAR(∞) (e.g., Canova, 2007). The usual approach is to use information criteria such as AIC and BIC, which is similar to our present approach.

[27]Since a VAR is obtained by inverting the fundamental moving-average representation, it automatically selects the fundamental representation (e.g., Lippi and Reichlin, 1994).

[28]An alternative estimation procedure to handle both invertible and non-invertible representations would be to use the Kalman filter with priors on the K initial values of the shocks $\{\varepsilon_{-K}...\varepsilon_0\}$, as recently proposed by

3.7 Dealing with non-stationary data

As can be seen from Theorem 1, GMA models can only capture impulse response functions that are bounded and integrable, which restricts our approach to stationary series. If the data are non-stationary, we can (i) allow for a deterministic trend in equation (1) and/or (ii) first-difference the data, and then proceed exactly as described above.

If a deterministic trend is suspected, we allow for a polynomial trend in each series, and we jointly estimate the parameters of the impulse responses (the $\boldsymbol{\Psi}_k$ coefficients) and the polynomial parameters.

If a stochastic trend is suspected, we can transform the data into stationary series by differencing the data. Importantly, and unlike with VARs, a GMA in first-difference is not misspecified if some variables are co-integrated.[29] After estimation, one can even test for co-integration by testing whether the matrix sum of moving-average coefficients ($\sum_{k=1}^{K} \sum_{l=0}^{k} \boldsymbol{\Psi}_l$) is of reduced rank (Engle and Yoo, 1987).

4 Gaussian Mixture Approximations of non-linear models

We now generalize the moving average model (1) by allowing for asymmetry and state-dependence, and we show how GMA models can easily accommodate such non-linearities.

Plagborg-Moller (2016). However, unlike our proposed approach, this procedure would be difficult to implement in non-linear models. Note also that the non-uniqueness of the moving average representation was proven for linear models (under Gaussian shocks). When we consider non-linearities, the non-uniqueness of the moving-average representation is not guaranteed anymore, and identification may be easier. In practice (and in Monte-Carlo simulations), the likelihood did not display multiple peaks when we allowed for asymmetry or state-dependence.

[29]The reason is that a GMA model directly works off of the moving-average representation and does not require inversion of the moving-average, unlike VAR models.

4.1 A non-linear moving-average model

In this section, we generalize model (1) by allowing the economy to respond non-linearly to shocks, and we consider the model

$$\mathbf{Y}_t = \sum_{k=0}^{\infty} \mathbf{\Psi}_k(\varepsilon_{t-k}, \mathbf{Z}_{t-k})\varepsilon_{t-k} \tag{6}$$

where ε_t is again the vector of structural innovations with $E\varepsilon_t = \mathbf{0}$ and $E\varepsilon_t\varepsilon_t' = \mathbf{I}$, and $\mathbf{Z}_t$ is a vector of macroeconomic variables that can be a function of past variables of $\mathbf{Y}$ or a function of variables exogenous to $\mathbf{Y}$.

Model (6) is a *non-linear* vector moving average representation of the economy, because in contrast to (1), the matrix of lag coefficients $\mathbf{\Psi}_k(\varepsilon_{t-k}, \mathbf{Z}_{t-k})$, i.e., the impulse response functions of the economy, are no longer constant. Instead, the coefficients of $\mathbf{\Psi}_k$ can depend on the values of the structural innovations ε_{t-k} and on the values of the macroeconomic variables $\mathbf{Z}_{t-k}$.

With $\mathbf{\Psi}_k$ a function of ε_{t-k}, the impulse response functions to a given structural shock depend on the value of the shock at the time of shock. For instance, a positive shock may trigger a different impulse response than a negative shock.

With $\mathbf{\Psi}_k$ a function of $\mathbf{Z}_{t-k}$, the impulse response functions to a structural shock depend on the value of the macroeconomic variables $\mathbf{Z}$ at the time of that shock. For instance, the response function may be different depending on the state of the business cycle (recession or expansion) at the time of the shock.

Because of its non-linear nature (6) does not admit a VAR representation, and the model cannot be recovered from a VAR. Instead, our GMA approach directly works off with the moving-average representation and can accommodate non-linearities. Moreover, the parametrization offered by Gaussian mixture approximations will ensure that the dimensionality of the problem remains reasonable. We now discuss in more details two cases of non-linear behavior that a GMA model can easily handle: (i) asymmetry and (ii) state-dependence.

4.1.1 Asymmetric effects of shocks

To allow for asymmetries, we let $\boldsymbol{\Psi}_k$ depend on the sign of the structural shock, i.e., we let $\boldsymbol{\Psi}_k$ take two possible values: $\boldsymbol{\Psi}_k^+$ or $\boldsymbol{\Psi}_k^-$. Specifically, a model that allows for asymmetric effect of shocks would write

$$\mathbf{Y}_t = \sum_{k=0}^{\infty} \left[\boldsymbol{\Psi}_k^+ \mathbf{1}_{\varepsilon_{\mathbf{t-k}} > \mathbf{0}} \varepsilon_{t-k} + \boldsymbol{\Psi}_k^- \mathbf{1}_{\varepsilon_{\mathbf{t-k}} < \mathbf{0}} \varepsilon_{t-k} \right] \tag{7}$$

with $\boldsymbol{\Psi}_k^+$ and $\boldsymbol{\Psi}_k^-$ the lag matrices of coefficients for, respectively, positive and negative shocks.

Denoting $\psi_{ij}^+(k)$, the i-row j-column coefficient of $\boldsymbol{\Psi}_k^+$ (that is, the impulse response of variable j to shock i), a GMA(N) model would then write

$$\psi_{ij}^+(k) = \sum_{n=1}^{N} a_{ij,n}^+ e^{-\left(\frac{k-b_{ij,n}^+}{c_{ij,n}^+}\right)^2}, \ \forall k > 0 \tag{8}$$

with a_{ij}^+, b_{ij}^+, c_{ij}^+ some constants to be estimated. A similar expression would hold for $\psi_{ij}^-(k)$.

4.1.2 Asymmetric and state-dependent effects of shocks

With asymmetry and state dependence, $\boldsymbol{\Psi}_k^+$ becomes $\boldsymbol{\Psi}_k^+(z_{t-k})$, i.e., the impulse response to a positive shock depends on the indicator variable z_t (and similarly for $\boldsymbol{\Psi}_k^-$).

Using a GMA(N) model, the impulse response function following a positive innovation (ψ_{ij}^+) can be parametrized as

$$\psi_{ij}^+(k) = (1 + \gamma_{ij}^+ z_{t-k}) \sum_{n=1}^{N} a_{ij,n}^+ e^{-\left(\frac{k-b_{ij,n}^+}{c_{ij,n}^+}\right)^2}, \ \forall k > 0, \, \varepsilon_{t-k} > 0 \tag{9}$$

with γ_{ij}^+, $a_{ij,n}^+$, $b_{ij,n}^+$ and $c_{ij,n}^+$ parameters to be estimated. An identical functional form holds for ψ_{ij}^-.

In this model, the amplitude of the impulse response depends on the state of the business cycle at the time of the shock. In (9), the amplitude of the impulse response is a function of the indicator variable z_t. Such a specification allows us to test whether, for instance, an expansionary policy has a stronger effect on output in a recession than in an expansion.

Note that in specification (9), the state of the cycle is allowed to stretch/contract the impulse response, but the shape of the impulse response is fixed (because a, b and c are all independent of z_t). While one could allow for a more general model in which all variables a, b and c depend on the indicator variable, specification (9) has two advantages. First, with limited sample size, it will typically be necessary to impose some structure on the data, and imposing a constant shape for the impulse response is a natural starting point.[30] Second, specification (9) generalizes trivially to GMAs of any order. The order of the GMA only determines the shape of the impulse response with higher order allowing for increasingly complex shapes. Then, for a given shape, the γ coefficient can stretch or expand the impulse response depending on the state of the cycle.[31]

4.2 Bayesian estimation of non-linear GMA models

We now discuss how estimating non-linear GMA models is a simple extension of the linear case. First, we present the construction of the likelihood, then we present the elicitation of priors.

4.2.1 Constructing the likelihood function

We discuss the construction of the likelihood in the more general case with non-linearities, i.e., when we consider the model $\mathbf{Y}_t = \sum_{k=0}^{\infty} \mathbf{\Psi}_k(\varepsilon_{t-k}, \mathbf{Z}_{t-k})\varepsilon_{t-k}$. The approach is identical to the approach described in section 3 but with two differences.

First, when the model allows for state dependence, the likelihood also depends on the

[30]Importantly, this assumption is easy to relax or to evaluate by model comparison using posterior odds ratios.

[31]Note the parallel and difference between (9) and a *varying coefficient model.* A varying coefficient model (e.g., Hastie and Tibshirani, 1993) is a (locally) linear model, whose coefficients are allowed to vary smoothly with some third variable z_t. In (9), the use of a *finite* sum of Gaussian basis functions (independent of z_t) plays a similar role to smoothness in varying coefficient models by restricting the shape of the impulse response and disciplining the estimates. Then, the effect of the third variable z_t is captured by letting the scale of the impulse response be a linear function of z_t.

value of the indicator vector $\mathbf{Z}_t$. Technically, constructing the likelihood of this specification is a straightforward extension of the linear case, when $\mathbf{Z}_t$ is a function of lagged values of $\mathbf{Y}_t$. To see that, note that we use the prediction-error decomposition to construct the likelihood function. Thus, we build a sequence of densities for $\mathbf{Z}_t$ that condition on past values of $\mathbf{Z}_t$, which includes past values of $\mathbf{Y}_t$. Thus, conditional on past values of $\mathbf{Y}_t$, $\mathbf{Z}_t$ is known.[32]

Second, with $\mathbf{\Psi}_k$ a function of ε_{t-k} (as in the case with asymmetry), one needs to make sure that the contemporaneous shocks are uniquely identified given a series of past shocks and given a set of model parameters. In other words, with the contemporaneous matrix depending on the value of the shock, it is important that there is a one to one mapping from structural to reduced form errors. Considering equation (5) in the non-linear case, we have

$$\mathbf{\Psi}_0^{\pm}(\varepsilon_{t+1})\varepsilon_{t+1} = \mathbf{Y}_{t+1} - \sum_{k=0}^{K} \mathbf{\Psi}_k^{\pm}\varepsilon_{t-k}, \tag{10}$$

and we need to ensure that $\mathbf{\Psi}_0^{\pm}(\varepsilon_{t+1})\varepsilon_{t+1}$ is only satisfied by a unique value ε_{t+1} given the model parameters. However, this is always the case when $\mathbf{\Psi}_0^{\pm}$ is lower triangular, i.e., with a recursive identification scheme. To see that, consider the bivariate case with $\varepsilon_{t+1} = (\varepsilon_{1,t+1}, \varepsilon_{2,t+1})'$ and denote $\psi_{0,ij}^{+}(\varepsilon_{t+1})$ the (i,j) elements of the contemporaneous impact matrix $\Psi_0^{+}(\varepsilon_{t+1})$. We have

$$\Psi_0^{\pm}(\varepsilon_{t+1}) = \begin{pmatrix} \psi_{0,11}^{\pm}(\varepsilon_{1,t+1}) & 0 \\ \psi_{0,21}^{\pm}(\varepsilon_{1,t+1}) & \psi_{0,22}^{\pm}(\varepsilon_{2,t+1}) \end{pmatrix}, \tag{11}$$

where the sign of $\psi_{0,ii}^{\pm}$ is fixed and independent of $\varepsilon_{i,t+1}$. Through a normalization, the elements on the diagonal are set to be always positive.

From (11) and (10), it is easy to show that the vector ε_{t+1} is uniquely identified. The first

[32]If we wanted to use an indicator function that was not part of the vector of endogenous variables $\mathbf{Y}_t$, this would also be possible by using a quasi-likelihood approach. That is, we would build a likelihood function that not only conditions on the parameters, but also the sequence of indicators $\mathbf{Z}_t$. This would in general not be efficient because the joint density of $\mathbf{Z}_t$ and $\mathbf{Y}_t$ could carry more information about the parameters in our model than the conditional density we advocate using. As long as $\mathbf{Z}_t$ is highly correlated with elements of (functions of) $\mathbf{Y}_t$, this loss in efficiency will likely be small.

component $\varepsilon_{1,t+1}$ is given by the first equation with

$$\varepsilon_{1,t+1} = \frac{\mathbf{Y}_{1,t}}{\psi^{\pm}_{0,11}(z_{t+1})}$$

which is uniquely identified since $\psi^{\pm}_{0,11}(\varepsilon_{1,t+1}) > 0$. Then, the second component $\varepsilon_{2,t+1}$ is given by

$$\varepsilon_{2,t+1} = \frac{\mathbf{Y}_{2,t} - \psi^{\pm}_{0,21}(\varepsilon_{1,t+1})\varepsilon_{1,t+1}}{\psi^{\pm}_{0,22}(\varepsilon_{2,t+1})}$$

which is uniquely identified since $\psi^{\pm}_{0,22}(z_{t+1}) > 0$. We can proceed similarly in a higher dimensional case.

Finally, to write down the one-step ahead forecast density $p(\mathbf{Y}_t|\theta, \mathbf{Y}^{t-1})$ as function of past observations and model parameters, we use the standard result (see e.g., Casella-Berger, 2002) that for $\mathbf{\Psi}_0$ a function of ε_{t+1}, we have

$$p(\mathbf{\Psi}_0(\varepsilon_{t+1})\varepsilon_{t+1}|\theta, y^t) = J_{t+1}p(\varepsilon_{t+1})$$

where J_{t+1} is the Jacobian of the mapping from ε_{t+1} to $\mathbf{\Psi}_0(\varepsilon_{t+1})\varepsilon_{t+1}$ and where $p(\varepsilon_{t+1})$ is the density of ε_{t+1}.[33]

4.2.2 Starting values and prior elicitation

As initial guesses, we set the parameters capturing asymmetry and state dependence to zero (i.e., no non-linearity).[34] This approach is consistent with the starting point (the null) of this paper: structural shocks have linear effects on the economy, and we are testing this null against the alternative that shocks have some non-linear effects. We then center the priors for these parameters at zero with loose priors.

[33] In our case with asymmetry, this Jacobian is simple to calculate, but the mapping is not differentiable at $\varepsilon = 0$. Since we will never exactly observe $\varepsilon = 0$ in a finite sample, we can implicitly assume that in a small neighborhood around 0, we replace the original mapping with a smooth function.

[34] An alternative would be to obtain initial estimates about possible non-linear effects. One option could be to combine Jorda's (2005) local projection method (which can accommodate non-linearities) with the structural shocks recovered from the VAR in order to get first estimates of the non-linear impulse responses.

5 Monte Carlo simulations

In this section, we conduct a number of Monte-Carlo simulations to illustrate the working of GMA models as well as to evaluate their performances in finite sample. We first evaluate the performances of GMA models in the linear case, and we then evaluate the ability of GMA models to detect (i) asymmetry alone and (ii) asymmetry and state-dependence.

Importantly, in all our Monte Carlo exercises, the estimated GMA models will be misspecified and only approximate the true Data Generating Process (DGP). We follow this strategy for two reasons. First, we want to be conservative and stack the odds against our proposed method. Second, this strategy is consistent with the idea that a GMA is meant to *approximate* the true DGP. By focusing on the approximate shape of the impulse response and thereby economizing on degrees of freedom, a GMA may (i) provide better estimates of the impulse responses in short sample, –a classical example of the bias-variance trade-off–, and (ii) be able to detect non-linearities. One goal of these simulation exercises is to evaluate whether this is indeed the case.

To simulate data, we proceed as follows. We first estimate a structural VAR on US data, invert it to obtain a set of impulse responses $\left\{\hat{\mathbf{\Psi}}_k\right\}_{k=0}^{\infty}$, and we modify these baseline impulse responses to introduce non-linearities, in particular asymmetry or state dependence. From these impulse responses, we generate simulated data from

$$\mathbf{Y}_t = \sum_{k=0}^{\infty} \hat{\mathbf{\Psi}}_k(\varepsilon_{t-k}, \mathbf{Z}_{t-k})\varepsilon_{t-k} \tag{12}$$

with ε_t Normally distributed, $E\varepsilon_t = 0$ and $E\varepsilon_t\varepsilon_t' = \mathbf{I}$.

In each scenario, we use 50 Monte-Carlo replications with a sample size $T = 200$, which roughly corresponds to the sample size available for the US.

5.1 Linear model

Our first simulation is meant to illustrate the working of Gaussian mixture approximations in the linear case. While we do not claim that GMAs are superior to VARs, we want to convey that GMAs can provide a useful alternative approach, especially in short samples.

The DGP is obtained from estimating the quarterly VAR(4) considered previously with the unemployment rate, the PCE inflation rate and the federal funds rate over 1959-2007. The impulse response functions to a monetary shock can be seen in Figure 1.

For each simulated dataset, we estimate (i) a GMA(2), and (ii) a VAR(4), and we evaluate the Mean-Square Error (MSE) of the estimated impulse response function over the horizons $H = 1...25$.[35] Importantly, we stack the odds in favor of the VAR and against the GMA model, because the estimated VAR is a well specified model.

The first row of Table 1 presents the average MSEs over the simulations. For unemployment and inflation, the GMA(2) is respectively 25 percent and 50 percent more accurate on average than the VAR. For the fed funds rate, the MSE is small in both cases, but again with a slight advantage for the GMA.[36] Table 1 also presents the average length and coverage rate of the confidence bands capturing the 95 percent posterior probability and compares it with the confidence bands implied by a Bayesian VAR with loose, but proper, Normal-Whishart priors. We report the average length and coverage rate at the time of the peak effect of the shock of the variable of interest. We can see that the average lengths are smaller for the GMA than for the VAR, while the coverage rate of the GMA remains good.

5.2 Non-linear models

We now evaluate the performances of GMA models in detecting non-linearities. For the DGP, we start from a VAR with (log) GDP, inflation and the fed funds rate, where we detrend

[35] Specifically, we report $MSE = \sum_{h=1}^{25}(\hat{\psi}_k - \psi_k)^2$ where $\hat{\psi}_k$ is the estimated impulse response function and ψ_k is the true function.

[36] Intuitively, the reason for the superior performances of GMA is the fact that the VAR often shows counter-factual oscillation patterns. In contrast, the GMA(2) is disciplined by its stricter parametrization.

GDP with a quadratic trend. Although we could have used the same VAR as previously, we preferred this one, because the price puzzle is more substantial in this specification (Figure 4), so that the Monte-Carlo exercise will be a more stringent test on a GMA(1) model that cannot capture the oscillating pattern in inflation. Again, the goal of the exercise is to assess whether a GMA model that only approximates the main feature of the impulse responses can still recover non-linearities.

Asymmetry

We first consider a DGP where the impulse response functions to monetary shocks depend on the sign of the shock. To introduce asymmetry, we modify the impulse responses $\left\{\hat{\boldsymbol{\Psi}}_k\right\}_{k=0}^{\infty}$ to make them depend on the sign of the monetary shock, and Figure 4 plots the asymmetric impulse response functions. For realism, the level of asymmetry that we simulate is chosen to roughly match the magnitude of the asymmetry we later find in US data. Note that we do not impose asymmetry for the response of the fed funds rate. This is done to test whether our procedure incorrectly reports the existence of asymmetry when there is none.

We estimate a GMA(1) with asymmetry on each set of simulated data, and Table 2 presents summary statistics for $a^+ - a^-$, which captures the amount of peak asymmetry for each one of the three variables in the model.

A number of results emerge. First, as shown by the frequency of rejection of zero coefficient for $a^+ - a^-$, the algorithm can detect asymmetry when it exists (case of output and inflation, first row of Table 2), even when the impulse response is not generated by one Gaussian, and even when, as with inflation, there is a strong oscillating pattern that cannot not captured by a one Gaussian approximation.[37] This is encouraging, because it supports our motivating idea that by approximating the most important feature of an impulse response, one can detect important non-linearities. Moreover, the algorithm does not detect asymmetry when there is none (case of the fed funds rate). Second, looking at the mean and standard-deviation of

[37]Specifically, the 90 percent posterior probability of $a^+ - a^-$ excludes zero for output and inflation respectively 94 and 90 percent of the time.

the estimates across Monte-Carlo replications (second row of Table 2), we can see that the algorithm under-estimates the amount of asymmetry (both for output and inflation). This indicates that in our empirical application on US data, our algorithm may under-estimate the magnitude of asymmetry present in the data. Third, the dispersion (third row) in the estimates across the Monte-Carlo replications is reasonably small, while the coverage rate of the posterior distribution – the frequency with which the true value lies within 90 percent of the posterior distribution–, is also good (fourth row).

Asymmetry and state dependence

We now consider a DGP where the impulse response functions to monetary shocks depend on the sign of the shock as well as the state of the business cycle. We introduce asymmetry exactly as in the previous exercise, but in addition, we posit that there is state dependence for output in response to a positive shock, i.e., $\gamma_Y^+ \neq 0$ in (9), where the indicator variable z_t is the US unemployment rate.[38] Again, the value of γ_Y^+ is chosen to be of the same order of magnitude as our later empirical findings with US data, and we set $\gamma_Y^+ = 1$.

We estimate a GMA(1) with asymmetry and state dependence on each set of simulated data, and Table 3 summarizes the results. A number of results emerge. First, the algorithm is very successful at detecting state dependence in output and the fact that $\gamma_Y^+ \neq \gamma_Y^-$ (first set of columns in Table 3). In the 50 Monte-Carlo replications, we detect $\gamma_Y^+ \neq \gamma_Y^-$ in all samples but one (first row). The algorithm also estimates the values of $\gamma_Y^+ - \gamma_Y^-$ without bias (second row), with reasonable dispersion (third row) and with good coverage (fourth row). Importantly, the algorithm detects no state dependence when there is none (case of inflation), as can be seen from the close to zero frequency of rejection of zero coefficient. Second, the algorithm can still pick up the existence of asymmetry for output and inflation ($\alpha^+ - \alpha^- \neq 0$, second set of columns). With a larger number of free parameters, estimation is more uncertain, but we can

[38]We could have used any indicator, but we wanted an indicator that has the same time series properties as the one we use on US data. We thus chose to use the US unemployment rate, which is the indicator we used in the application section.

still detect the existence of asymmetry in more than 80 percent of cases. Finally, looking at the estimates for γ_Y^+ and γ_Y^- separately, the algorithm estimates the value of γ_Y^+ –the magnitude of the non-linearity– with a downward bias, which seems to translate into an upward bias for γ_Y^-, although that bias is not significant over the 50 Monte-Carlo replications (last four columns of Table 3).

6 Application: the non-linear effects of monetary shocks

In this section, we apply our proposed GMA approach and study the non-linear effects of monetary shocks. We consider a small-scale model of the US economy in the spirit of Primiceri (2005), where $\mathbf{Y}_t$ includes the unemployment rate, the PCE inflation rate and the federal funds rate. As in Primiceri (2005), monetary policy affects the economy with a lag, and the matrix $\mathbf{\Psi}_0$ is assumed to be lower-triangular. The data cover 1959Q1 to 2007Q4, and we exclude the latest recession where the fed funds rate was constrained at zero and no longer captured variations in the stance of monetary policy.[39,40] When constructing the likelihood, we consider a moving-average model with $K = 45$, chosen to be large enough such that the lag matrix coefficients $\mathbf{\Psi}_k$ are close enough to zero for $k > K$.[41]

As a preliminary test, we start by checking that a linear GMA model performs well against a standard VAR model. Then, we present the non-linear impulse response functions obtained from a non-linear GMA with asymmetry alone first, and then with asymmetry and state dependence.

[39] As an alternative, we could include data from the latest recession if we used the one-year or two-year government bond rate as the policy indicator (Gertler and Karadi, 2015). The one-year government bond rate remained positive until 2011 and it was argued that the zero lower bound was not a constraint on the Federal Reserve's ability to manipulate the two-year rate (Swanson and Williams, 2014). Results in terms of asymmetry and state dependence were similar.

[40] While we use quarterly data as in Primiceri (2005), we also conducted our estimation using monthly data. Results were very similar.

[41] As a robustness check, we consider a higher moving-average lag-length with $K = 55$. Results were identical.

6.1 The linear case: VAR versus GMA

First, we evaluate our GMA approach by doing a simple model comparison between a linear GMA(1) and a regular VAR with 4 lags.

Table 4 reports the (log) marginal densities for the GMA and the VAR, so that a model comparison can be readily obtained by computing the Bayes factor (obtained by taking the exponential of the difference in (log) marginal densities) after positing equal priors for the two competing models. Encouragingly for our approach, Bayesian model comparison favors the more parsimonious GMA(1) with a Bayes factor of about 400.

6.2 The asymmetric effects of monetary shocks

We now estimate an asymmetric GMA model in which the impulse responses to monetary shocks depend on the sign of the shock.[42,43] As detailed in the methodology section, to choose the appropriate order of the GMA model, we consider models with an increasing number of Gaussian basis functions. As shown in columns (3) to (5) of Table 4, Bayesian model comparison favors a GMA(2) , and from now on we will report and discuss the results obtained using a GMA(2).

We can see that Bayesian model comparison strongly favors a model with asymmetry in the impulse responses to monetary shocks: the (log) marginal density of an asymmetric GMA(2) is about 20 log-points larger than the linear GMA(1), which implies a Bayes factor of about 10^8. Against the VAR, the Bayes factor is 10^{11}.

Figure 5 plots the impulse responses (in percentage points) of unemployment, the price

[42] The impulse responses to the other shocks were modeled with linear (i.e., not asymmetric) GMAs. As a robustness check, we verified that our findings are not driven by this restriction by estimating a model with non-linearity in response to all shocks.

[43] As another robustness check, we evaluated the presence of asymmetry using monetary shocks identified through the narrative approach by Romer and Romer (2004) and extended until 2007 by Coibion et al. (2012). Encouragingly, despite very different identification methodologies, the correlation between the (median) shocks recovered with the recursively identified asymmetric GMA and the Romer and Romer shocks is at high at 0.62 over 1966q1-2007q4. We estimated an asymmetric GMA model with 4 variables included in the following order: the Romer and Romer shocks, unemployment, inflation and the fed funds rate, and we studied the asymmetric impulse responses to the first shock. The evidence for asymmetry was very similar to the one found with the recursive identification scheme.

level and the federal funds rate to a one standard-deviation monetary shock. The thick lines denote the impulse response functions implied by the posterior mode, and the error bands are the 5th and 95th posterior percentiles.

The evidence for asymmetry is striking: following a contractionary monetary shock, which represents a 70 basis points increase in the fed funds rate, unemployment increases by about 0.15 percentage points (ppt), whereas a (linear) VAR implies only a 0.10 ppt increase. In contrast, following an expansionary monetary shock (a 70 basis points decrease in the fed funds rate), the response of unemployment is small (a decline of 0.04 percentage points) and non-significantly different from zero. Figure 6 plots the posterior distribution of the difference in impulse responses between positive and negative shocks. This figure can be seen as a point-wise test of difference in impulse responses at different horizons. The 90 percent posterior interval of the difference in impulse responses of unemployment is substantially above zero for horizons 3 to 10, in line with the conclusion from the Bayes factors that the data support a model with asymmetric impulse responses to monetary shocks.[44]

Although the error bands are too large to be conclusive, the response of the price level also displays an interesting asymmetric pattern: the price level appears more sticky following a contractionary shock –displaying a larger price puzzle– than following an expansionary shock for which the price level drops on impact and displays no price puzzle. This is exactly the pattern one would expect if downward price (or wage) rigidity was responsible for the asymmetric response of unemployment.[45]

Finally, we also find asymmetry in the response of the fed funds rate to a monetary shock. A monetary shock generates a slightly more persistent increase in the fed funds rate than its expansionary counterpart. This can be seen in the bottom right panel where the response of the fed funds rate is more short-lived following an expansionary shock, or in Figure 6 where

[44] In the case of the GMA(1) model, an alternative test for asymmetry is a Wald-type test on $a^+ - a^-$. This test (not shown) gives a similar conclusion: for unemployment, the 90 percent posterior interval of $a^+ - a^-$ excludes zero.

[45] The existence of downward wage rigidity is supported empirically by the scarcity of nominal wage cuts relative to nominal wage increases (e.g., Card and Hyslop, 1997).

the posterior distribution of the differences in the responses of the fed funds rate exclude zero from horizons 1 to 3.[46]

6.3 The asymmetric and state-dependent effects of monetary shocks

In this section, we enrich our model by allowing the effect of monetary policy to depend on both the state of the business cycle and the sign of the shock. Intuitively, we would like to test whether monetary policy is more powerful at stimulating the economy in a period of economic slack, and whether an expansionary shock is more likely to generate inflation in a tight labor market.

We thus estimate model (9) with a GMA(2), and we use last period's unemployment rate as cyclical indicator (z_t).[47] Table 4 shows that Bayes model comparison strongly favors the model with asymmetry *and* state dependence over all the other models.

To visualize the effects of the state of the cycle on the impulse responses, Figure 8 shows how the peak effect of a monetary shock on unemployment or inflation depends on the state of the business cycle at the time of the shock, and to put results into perspective, Figure 7 plots the unemployment rate (i.e., the indicator variable z_t) along with the identified monetary shocks.

The first two rows plot the peak responses of respectively unemployment and inflation to contractionary and expansionary shocks, while the last row plots histograms of the distributions of respectively contractionary shocks and expansionary shocks over the business cycle. The

[46]One way to gauge how much of the asymmetric response of unemployment can be explained by the asymmetric response of the fed funds rate is to proceed as in the government spending multiplier literature (e.g., Ramey and Zubairy, 2014) and to compute the total change in unemployment relative to the total change in the fed funds rate, that is to compute $m = \sum_{k=0}^{K} \psi_k^U / \sum_{k=0}^{K} \psi_k^{ffr}$ for respectively positive and negative shocks. After "controlling" for the total change in the fed funds rate, the asymmetry is still present with $m^+ = .24 > m^- = .12$.

[47]This approach has the advantage of being self contained in that the unemployment rate is itself an endogenous variable whose behavior is described by (9). As an alternative, we also experienced with (i) the average growth rate of GDP over a one year period centered on the current quarter, and (ii) the unemployment rate detrended with an HP-filter ($\lambda = 10^5$). The latter specification was used to make sure that our results were not driven by slow moving trends (e.g., due to demographics) in the unemployment rate, which could make the unemployment rate a poor indicator of the amount of economic slack (see e.g. Barnichon and Mesters, 2015). Both specifications gave similar conclusions.

last row of Figure 8 has two purposes: (i) make sure that our results are not driven by an unusual distribution of shocks over the business cycle, say with more contractionary shocks in expansions than in recessions (which could happen in a short sample), (ii) get a sense of the range of unemployment rates over which we identify the coefficients capturing state dependence. Regarding (i), Figure 8 shows not marked difference in the distributions of positive and negative shocks over the business cycle.[48] Regarding (ii), most of the unemployment variations used to infer state dependence occur between 5 and 7 percent.

We first discuss the response of unemployment. The upper-left quadrant in Figure 8 depicts how the peak effect of a contractionary shock on unemployment varies as we move from a tight labor market (unemployment at 4 percent) to a slack labor market (unemployment at 8 percent). The thick dashed line represents the VAR estimate. Since the VAR is linear, that latter estimate is constant as the peak effect of monetary policy is independent of the state of the business cycle. The thick blue line depicts estimates from our non-linear framework. We can notice that the effect of a contractionary policy increases with the unemployment rate, being about 30 percent larger at a business cycle trough than at a business cycle peak.

For expansionary shocks (bottom left quadrant), the evidence is not as strong, but our estimates suggest some mild state dependence: the higher the unemployment rate, the larger the real effect of an expansionary policy. In fact, the 90th posterior probability bands start including the VAR estimate when the unemployment rate rises above 7 percent. That being said, the asymmetry between expansionary and contractionary interventions remains, and an expansionary policy is always considerably less potent than its contractionary counterpart.

We now turn to the response of inflation, depicted in the right-hand column of Figure 8. While we do not find any evidence of state dependence for contractionary shock, we find strong evidence that expansionary shocks generate a substantial rise in inflation when the unemployment rate is low. Interestingly, this finding is consistent with a standard Keynesian narrative, according to which a monetary authority trying to expand an economy already above

[48] A Kolmogorov-Smirnov test confirms this visual inspection, as we cannot reject the null that the two distributions are identical.

potential would only achieve higher inflation through increased price/wage pressures.

7 Conclusion

This paper proposes a new method to identify the (possibly non-linear) dynamic effects of structural shocks by using Gaussian basis functions to approximate impulse response functions. We apply our approach to the study of monetary policy and find that the effect of a monetary intervention depends strongly on the sign of the intervention. A contractionary shock has a strong adverse effect on output, larger than implied by linear estimates, but an expansionary shock has, on average, no significant effect on output. Interestingly, and while the evidence for inflation is relatively uncertain, the behavior of the inflation is consistent with asymmetry emerging (at least in part) out of downward price/wage rigidities: inflation displays a more marked price puzzle following a contractionary shock than following an expansionary shock. Finally, the effect of a monetary shock also depends on the state of the business cycle at the time of the intervention: An expansionary shock during a time a low unemployment generates not significant drop in unemployment but leads to a burst of inflation, consistent with a standard Keynesian narrative.

Although this paper studies non-linearities in the effect of monetary policy, Gaussian Mixture Approximations of the impulse responses may be useful in many other contexts. First, as a direct extension of the current paper, our method could be used to estimate the non-linear effects of other important shocks where the existence of asymmetry or state-dependence remains an important and unresolved question; notably fiscal policy shocks (Auerbach and Gorodnichenko, 2012, Ramey and Zubairy, 2014) or credit supply shocks (Gilchrist and Zakrajsek, 2012). Second, while we presented our method in the context of a recursive identification scheme, our method is quite general and can also be applied to other popular identification schemes, such as sign-restrictions (Uhlig, 2005) or long-run restrictions (Blanchard and Quah, 1989, Gali 1999). Finally, the parametrization offered by GMA models and the associated efficiency gains may be useful even for linear models, where the sample size is small and/or the data are particularly noisy.

References

[1] Almon, S. "The distributed lag between capital appropriations and expenditures," Econometrica, 33, January, 178-196, 1965

[2] Angrist, J., Jorda O and G. Kuersteiner. "Semiparametric estimates of monetary policy effects: string theory revisited," NBER Working Paper, 2013

[3] Alspach D. and H. Sorenson H. "Recursive Bayesian Estimation Using Gaussian Sums," Automatica, Vol 7, pp 465-479, 1971.

[4] Alspach D. and H. Sorenson H. "Nonlinear Bayesian Estimation Using Gaussian Sum Approximations," IEEE Transactions on Automatic Control, Vol 17-4, August 1972.

[5] Auerbach A and Y Gorodnichenko. "Measuring the Output Responses to Fiscal Policy," American Economic Journal: Economic Policy, vol. 4(2), pages 1-27, 2012

[6] Auerbach A and Y Gorodnichenko. "Fiscal Multipliers in Recession and Expansion." In Fiscal Policy After the Financial Crisis, edited by Alberto Alesian and Francesco Giavazzi, pp. 63–98. University of Chicago Press, 2013.

[7] Barnichon R, and G. Mesters, "On the Demographic Adjustment of Unemployment," Working Paper, 2015.

[8] Barro, Robert J., "Unanticipated Money Growth and Unemployment in the United States," American Economic Review, LXVII, 101-15, 1977.

[9] Beaudry, Paul, and Gary Koop. "Do Recessions Permanently Change Output?" Journal of Monetary Economics 31 (1993), 149 - 63.

[10] Blanchard, O. and D. Quah. "The Dynamic Effects of Aggregate Demand and Supply Disturbances," American Economic Review, 79(4), pages 655-73, September 1989.

[11] Buhmann, Martin D, Radial Basis Functions: Theory and Implementations, Cambridge University Press, 2003.

[12] Canova, F. "Methods for Applied Macroeconomic Research," Princeton University Press, 2007.

[13] Card, D. and D. Hyslop, "Does Inflation Greases the Wheels of the Labor Market?" in Reducing Inflation: Motivation and Strategy, C. Romer and D. Romer, eds., University of Chicago Press, 1997.

[14] Casella, G. and R. L. Berger, Statistical Inference, Duxbury, 2002.

[15] Cover, J. "Asymmetric Effects of Positive and Negative Money-Supply Shocks," The Quarterly Journal of Economics, Vol. 107, No. 4, pp. 1261-1282, 1992.

[16] Coibion, O. "Are the Effects of Monetary Policy Shocks Big or Small?," American Economic Journal: Macroeconomics, vol. 4(2), pages 1-32, April, 2012.

[17] Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia, "Innocent Bystanders? Monetary Policy and Inequality in the US," NBER Working Papers 18170, National Bureau of Economic Research, Inc, 2012.

[18] Christiano, L., M. Eichenbaum, and C. Evans. "Monetary policy shocks: What have we learned and to what end?," Handbook of Macroeconomics, volume 1, chapter 2, pages 65-148, 1999.

[19] DeLong B and L. Summers, "How Does Macroeconomic Policy Affect Output?," Brookings Papers on Economic Activity, Economic Studies Program, The Brookings Institution, vol. 19(2), pages 433-494, 1988.

[20] Engle, R. and B. Yoo, "Forecasting and testing in co-integrated systems," Journal of Econometrics, vol. 35(1), pages 143-159, May 1987.

[21] Fernandez-Villaverde, J. and J. Rubio-Ramirez, “Comparing dynamic equilibrium models to data: a Bayesian approach,” Journal of Econometrics, vol. 123(1), pages 153-187, November, 2004.

[22] Gali, J. ”Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?,” American Economic Review, vol. 89(1), pages 249-271, March 1999.

[23] Gertler M and P Karadi, ”Monetary Policy Surprises, Credit Costs, and Economic Activity,” American Economic Journal: Macroeconomics, vol. 7(1), pages 44-76, January 2015.

[24] Gilchrist, S and E Zakrajsek, ”Credit Spreads and Business Cycle Fluctuations,” American Economic Review, American Economic Association, vol. 102(4), pages 1692-1720, June 2012.

[25] Haario, H., E. Saksman, and J. Tamminen, ”An adaptive Metropolis algorithm,” Bernoulli 7, no. 2, 223–242, 2001.

[26] Hamilton, J. “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” Econometrica 57, 357-384, 1989.

[27] Hastie, T. and Tibshirani, R. “Varying-coefficient models,” J. Roy. Statist. Soc. Ser. B 55, 757–796, 1993.

[28] Hastie, T., R. Tibshirani and J. Friedman, *The Elements of Statistical Learning*, Springer 2009.

[29] Hubrich, K. and T. Terasvirta. ”Thresholds and smooth transitions in vector autoregressive models,” Advances in Econometrics “VAR Models in Macroeconomics, Financial Econometrics, and Forecasting”, Vol. 31, 2013.

[30] Jorda O., "Estimation and Inference of Impulse Responses by Local Projections," American Economic Review, pages 161-182, March 2005.

[31] Koop G, M. Pesaran and S. Potter "Impulse response analysis in nonliner multivariate models," Journal of Econometrics, 74 119-147, 1996.

[32] Koop G and S. Potter "Dynamic asymmetries in US unemployment," Journal of Business & Economic Statistics Volume 17, Issue 3, 1999.

[33] Koreyaar J. *Mathematical Methods*, Vol. 1, pp 330-333. Academic Press, New York, 1968.

[34] Lippi, M. & Reichlin, L. "VAR analysis, nonfundamental representations, Blaschke matrices," Journal of Econometrics 63(1), 307–325, 1994.

[35] Lo, M. C., and J. Piger "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, 2005.

[36] McLachlan, G. and D. Peel. *Finite Mixture Models*. Wiley Series in Probability and Statistics, 2000.

[37] Morgan, D. "Asymmetric Effects of Monetary Policy." Federal Reserve Bank of Kansas City Economic Review 78, 21-33, 1993.

[38] Plagborg-Moller, P. "Bayesian Inference on Structural Impulse Response Functions," Working Paper, 2016.

[39] Potter, S. "A non-linear approach to US GNP," Journal of Applied Econometrics," Vol 10 109-125, 1995

[40] Primiceri, G. "Time Varying Structural Vector Autoregressions and Monetary Policy," Review of Economic Studies, vol. 72(3), pages 821-852, 2005.

[41] Racine, J. "Nonparametric Econometrics: A Primer," Foundations and Trends in Econometrics, now publishers, vol. 3(1), pages 1-88, March 2008.

[42] Ravn M and M Sola. "A Reconsideration of the Empirical Evidence on the Asymmetric Effects of Money-supply shocks: Positive vs. Negative or Big vs. Small," Archive Discussion Papers 9606, Birkbeck, 1996.

[43] Ravn M. and M. Sola. "Asymmetric effects of monetary policy in the United States," Review, Federal Reserve Bank of St. Louis, issue Sep, pages 41-60, 2004.

[44] Ramey V. "Comment on "Roads to Prosperity or Bridges to Nowhere? Theory and Evidence on the Impact of Public Infrastructure Investment"," NBER Chapters, in: NBER Macroeconomics Annual 2012, Volume 27, pages 147-153.

[45] Ramey V. and S. Zubairy. "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data ," Working Paper, 2014.

[46] Robert, Christian P. and George Casella "Monte Carlo Satistical Methods" Springer, 2004.

[47] Romer, C., and D. Romer. "A New Measure of Monetary Shocks: Derivation and Implications," American Economic Review 94 (4): 1055–84, 2004

[48] Santoro, E, I. Petrella, D. Pfajfar and E. Gaffeo "Loss Aversion and the Asymmetric Transmission of Monetary Policy," Journal of Monetary Economics, 68 19-35, 2014.

[49] Swanson, E. and J. Williams. "Measuring the Effect of the Zero Lower Bound on Medium– and Longer–Term Interest Rates." American Economic Review 104 (10): 3154–85, 2014.

[50] Tenreyro, S., and G. Thwaites (2015): "Pushing on a string: US monetary policy is less powerful in recessions," Working Paper.

[51] Thoma, M. "Subsample Instability and Asymmetries in Money-Income Causality." Journal of Econometrics 64, 279-306, 1994

[52] Uhlig, H. "What are the effects of monetary policy on output? Results from an agnostic identification procedure," Journal of Monetary Economics, vol. 52(2), pages 381-419, March 2005.

[53] Walsh C. Monetary Theory and Policy, 3nd. ed., The MIT Press, 2010.

[54] Weise, C. "The Asymmetric Effects of Monetary Policy: A Nonlinear Vector Autoregression Approach," Journal of Money, Credit and Banking, vol. 31(1), pages 85-108, February 1999.

Appendix: Proof of Theorem 1

Following Alspach and Sorenson (1971, 1972) in the context of approximating distributions, the problem of approximating a function f can be considered within the context of delta families of positive types.

Delta families are families of functions which converge to a delta function as a parameter characterizing the family converges to a limit value.

Let $\{\delta_\lambda\}$ be a family of functions on the interval $]-\infty, +\infty[$ which are integrable over every interval. $\{\delta_\lambda\}$ forms a delta family of positive type if the following conditions are satisfied:

1. For every constant $\gamma > 0$, δ_λ tends to zero uniformly for $\gamma \leq |x| \leq \infty$ as $\lambda \to \lambda_0$
2. There exist s in $\mathbb{R}$ so that $\int_{-s}^{s} \delta_\lambda(x)dx \longrightarrow 1$ as λ tends to some limit value λ_0
3. $\delta_\lambda(x) \geq 0$ for all x and λ

Defining

$$\delta_\lambda(x) \equiv G_\lambda(x) = \frac{1}{\sqrt{2\pi\lambda^2}} e^{-\frac{x^2}{\lambda^2}}, \tag{13}$$

it is easy to see that the Gaussian functions $\{G_\lambda\}$ form a delta family of positive type as $\lambda \to 0$ (i.e., $\lambda_0 = 0$). That is, the Gaussian function tends to the delta function as the variance tends to zero.[49]

We can then make use of the following theorem.

Theorem: The sequence $\{f_\lambda\}$ which is formed by the convolution of δ_λ and f

$$f_\lambda(x) = \int_{-\infty}^{+\infty} \delta_\lambda(x-u) f(u) du \tag{14}$$

converges uniformly to f as $\lambda \to \lambda_0$ for x on every interval $[x_0, x_1]$ of $\mathbb{R}$.

Proof: see Korevaar (1968).

[49] Note that this proof can be easily applied to other functions (such as the inverse quadratic function $x \to \frac{1}{1+(\frac{x}{\lambda})^2}$) that form a delta family of a positive type, so that our approach is not restricted to Gaussian functions.

Using (13) in (14), the function f_λ given by

$$f_\lambda(x) = \int_{-\infty}^{+\infty} G_\lambda(x-u)f(u)du \tag{15}$$

converges uniformly to f as $\lambda \to 0$ for x in some arbitrary interval $[x_0, x_1]$ of $\mathbb{R}$.

Next, we want to approximate (15) with a Riemann sum. To do so, first rewrite f_λ as

$$f_\lambda(x) = \underbrace{\int_{-\infty}^{-s} G_\lambda(x-u)f(u)du}_{=A(\lambda,x)} + \int_{-s}^{+s} G_\lambda(x-u)f(u)du + \underbrace{\int_{s}^{+\infty} G_\lambda(x-u)f(u)du}_{=B(\lambda,x)} \tag{16}$$

for $s > 1$.

Note that for any $s > 1$, we have

$$\begin{aligned} 0 &\leq \int_s^{+\infty} G_\lambda(u)du \\ &\leq \frac{1}{\sqrt{2\pi\lambda^2}} \int_s^{+\infty} e^{-\frac{u}{\lambda^2}} du \text{ since } u^2 > u \text{ for any } u \text{ in } [s, +\infty[,\ s > 1 \\ &\leq \left[\frac{-\lambda^2}{\sqrt{2\pi\lambda^2}} e^{-\frac{u}{\lambda^2}} \right]_s^{+\infty} = \frac{|\lambda|}{\sqrt{2\pi}} e^{-\frac{s}{\lambda^2}} \underset{\lambda \to 0}{\longrightarrow} 0 \end{aligned}$$

which shows that $\forall s > 1$, $\lim\limits_{\lambda -> 0} \int_s^{+\infty} G_\lambda(u)du = 0$. Symmetrically, we can show $\lim\limits_{\lambda -> 0} \int_{-\infty}^{-s} G_\lambda(u)du = 0$.

Going back to (16), we have

$$0 \leq |B(\lambda, x)| \leq M \int_{-\infty}^{x-s} G_\lambda(t)dt$$

where $M = \sup\limits_{x \in \mathbb{R}} |f(x)|$. Since $x \in [x_0, x_1]$, we can choose an $s > 1$ such that $x - s < -1$, so that we can apply the previous result and get

$$\lim_{\lambda \to 0} |B(\lambda, x)| = 0. \tag{17}$$

Proceeding symmetrically, we have $\lim_{\lambda\to 0} |A(\lambda, x)| = 0$.

Finally, since the function $u \mapsto G_\lambda(x-u)f(u)$ is continuous over $[-s, s]$, we can approximate $\int_{-s}^{+s} G_\lambda(x-u)f(u)du$ with a Riemann sum. Denoting

$$f_{\lambda,N}(x) = \sum_{n=1}^{N} G_\lambda(x-\xi_n)f(\xi_n)\,(\xi_n - \xi_{n-1})$$

where $\xi_n = -s + n\frac{2s}{N}$, we get that

$$\lim_{N\to\infty} f_{\lambda,N}(x) = \int_{-s}^{+s} G_\lambda(x-u)f(u)du. \tag{18}$$

Denoting $a_n = f(\xi_n)\,(\xi_n - \xi_{n-1})$, $b_n = \xi_n$ and $c_n = \lambda$, using (18), (17) in (16) and combining with (15), we get that

$$\lim_{\lambda\to 0}\left(\lim_{N\to\infty} f_{\lambda,N}(x)\right) = f(x)$$

which completes the proof .

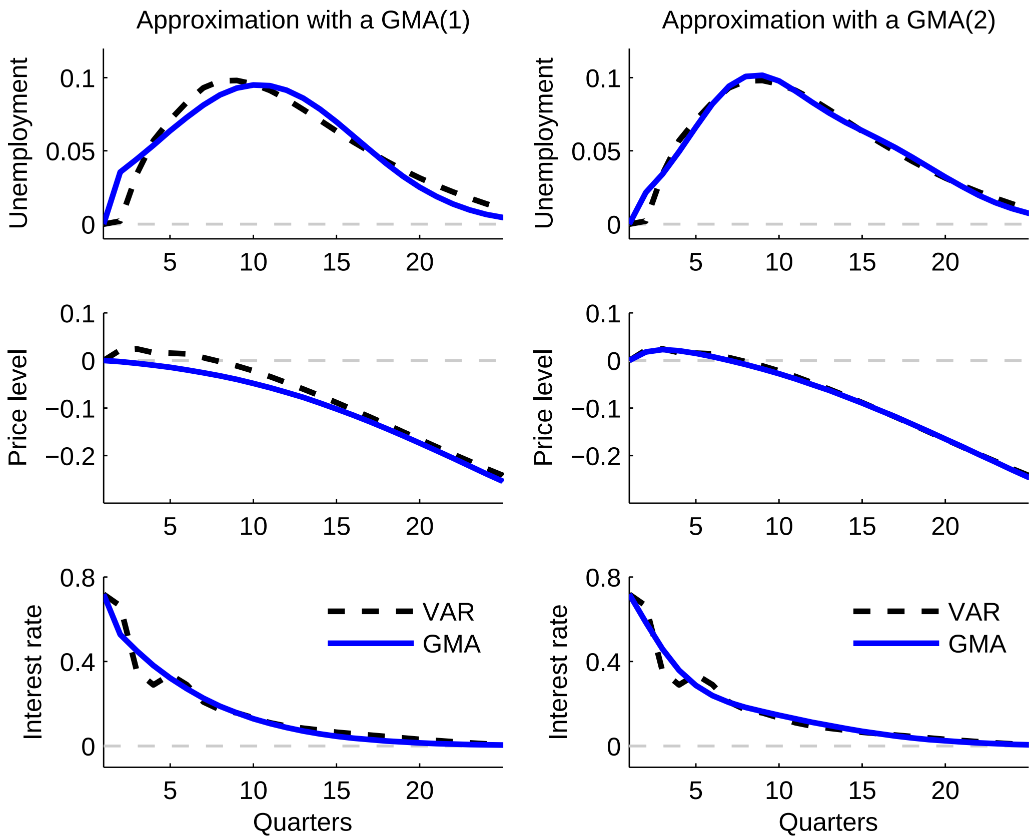

Figure 1: Impulse response functions (in ppt) of the unemployment rate, the (log) price level and the federal funds rate to a one standard-deviation monetary shock. Impulse responses estimated with a VAR (dashed-line) or approximated using one Gaussian basis function (GMA(1), left-panel, thick line) or two Gaussian basis functions (GMA(2), right panel thick line). Estimation using data covering 1959-2007.

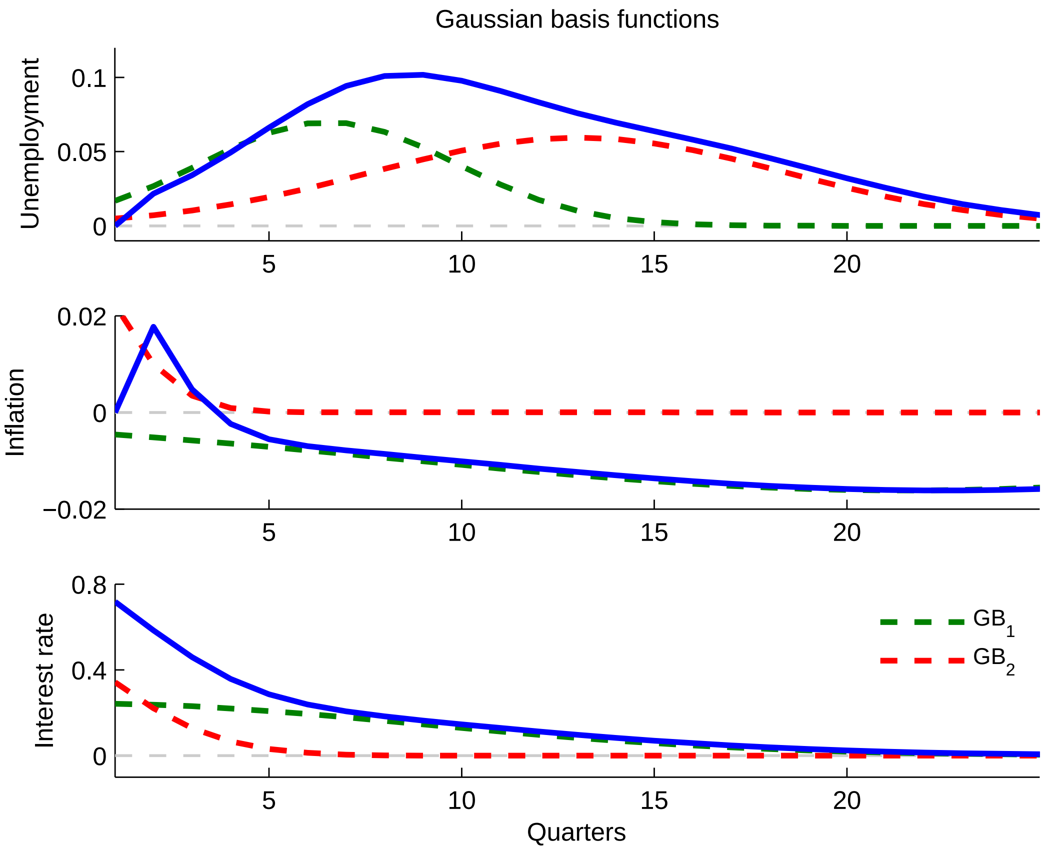

Figure 2: Gaussian basis functions (dashed lines) used by a GMA(2) to approximate the responses of unemployment, inflation and the fed funds rate to a monetary shock. The basis functions are appropriately weighted so that their sum gives the GMA(2) parametrization of the impulse response functions (solid lines).

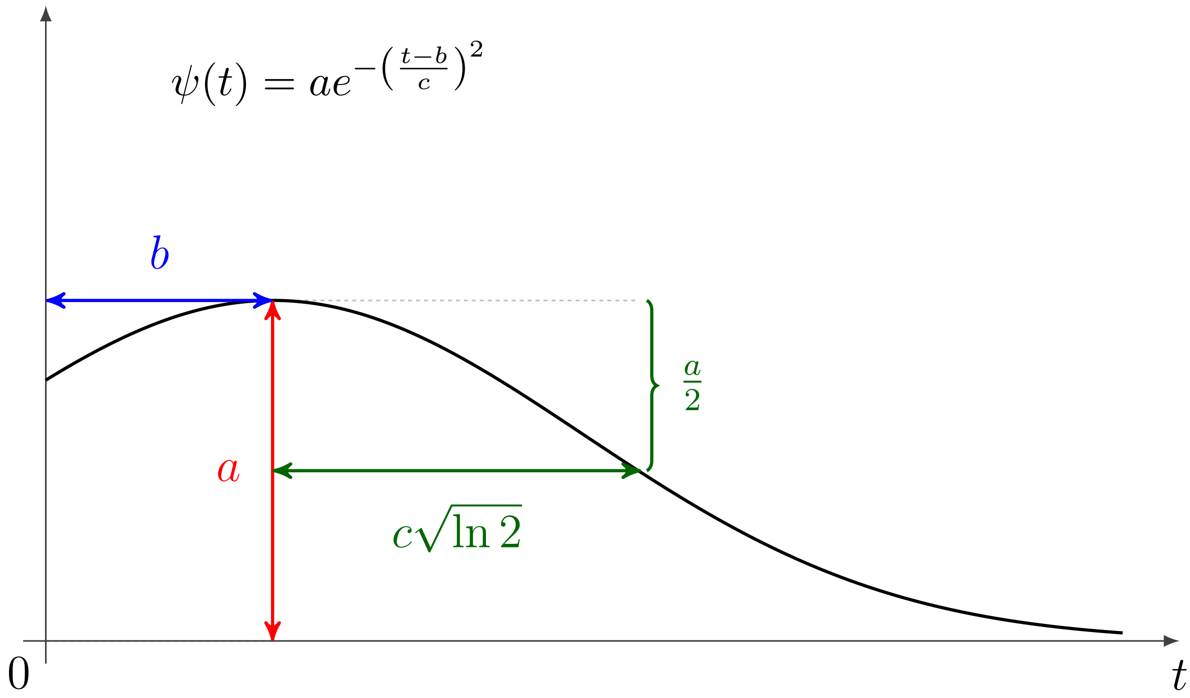

Figure 3: Interpreting an impulse response function with a GMA(1) model.

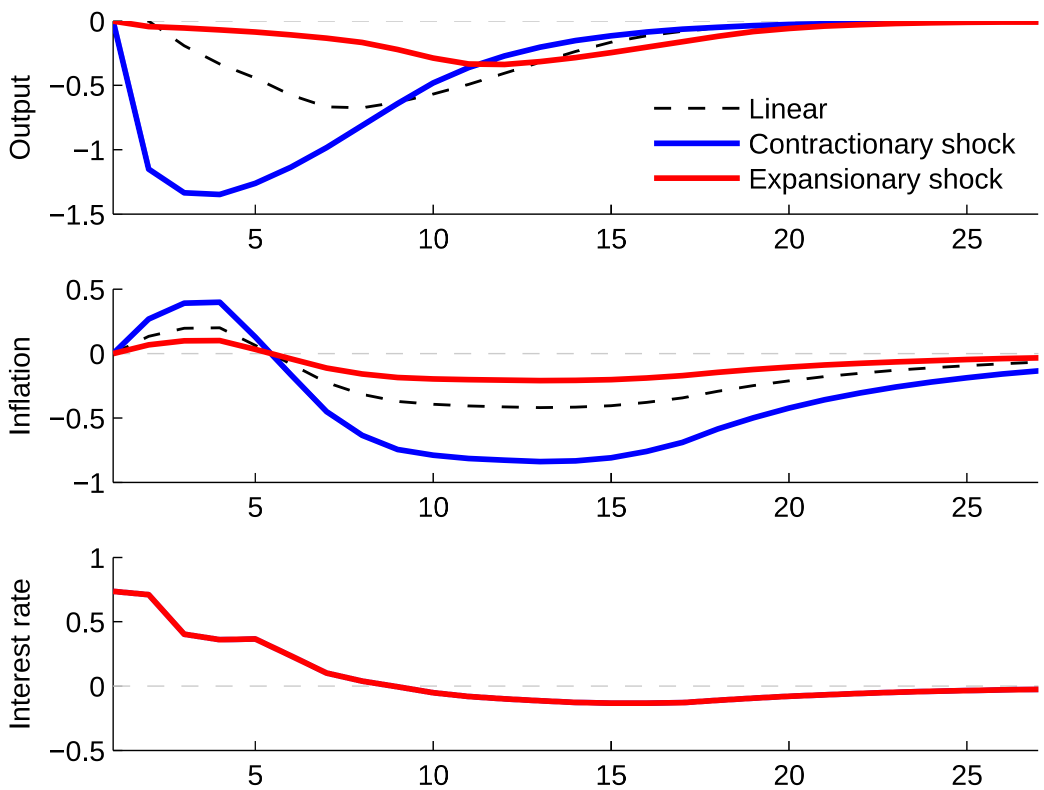

Figure 4: Asymmetric impulse response functions to a monetary shock. The thick green line reports the impulse response to a positive shock, and the thick blue line reports the impulse response to a negative shock (with sign flipped for clarity of exposition). The dashed lines are the impulse responses estimated from a VAR.

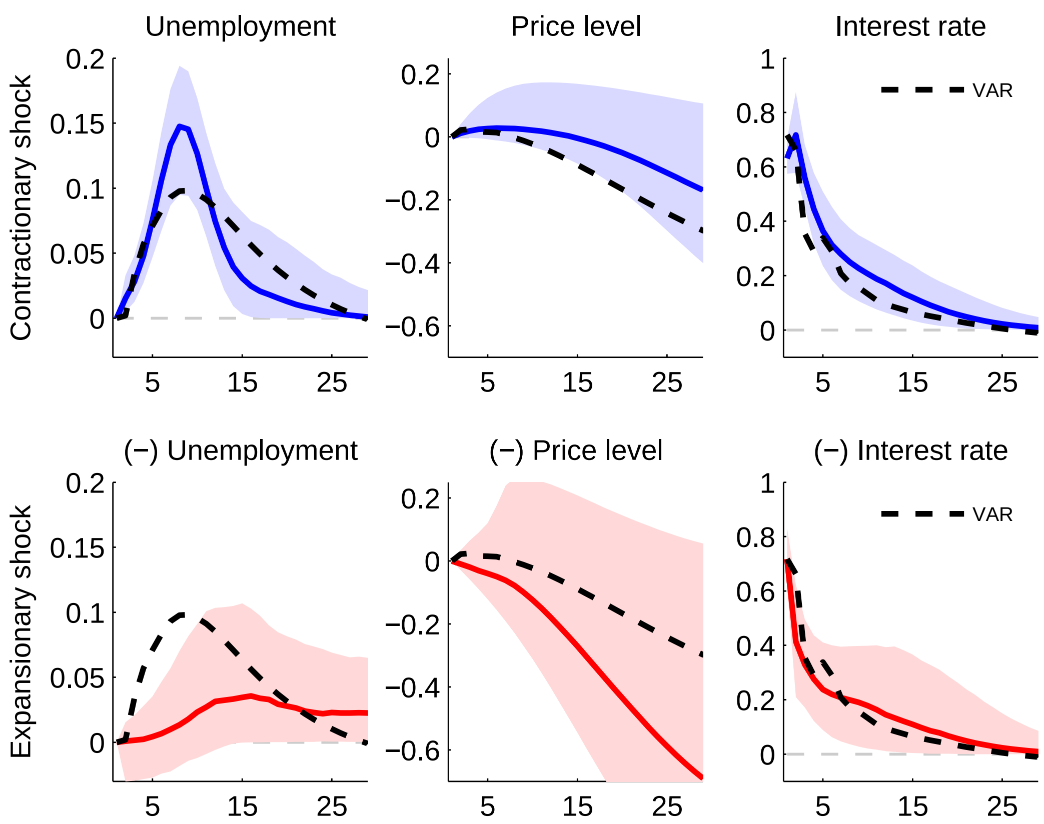

Figure 5: Impulse response functions (in ppt) of the unemployment rate, the (log) price level and the federal funds rate to a one standard-deviation monetary shock. Estimation from a VAR (dashed-line) or from a GMA(2) with asymmetry (plain line). Shaded bands denote the 5th and 95th posterior percentiles. For ease of comparison, responses to the expansionary shock are multiplied by -1. Estimation using data covering 1959-2007.

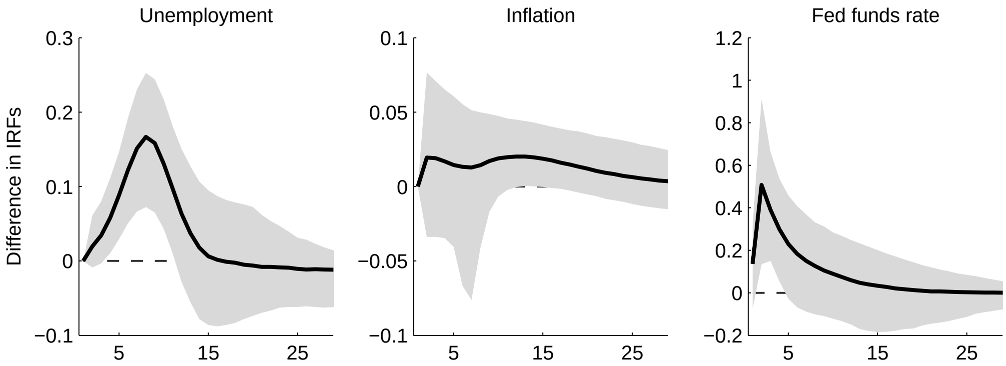

Figure 6: Differences in impulse response functions (in ppt) of the unemployment rate, the (log) price level and the federal funds rate to a one standard-deviation monetary shock. Shaded bands denote the 5th and 95th posterior percentiles. Estimation using data covering 1959-2007.

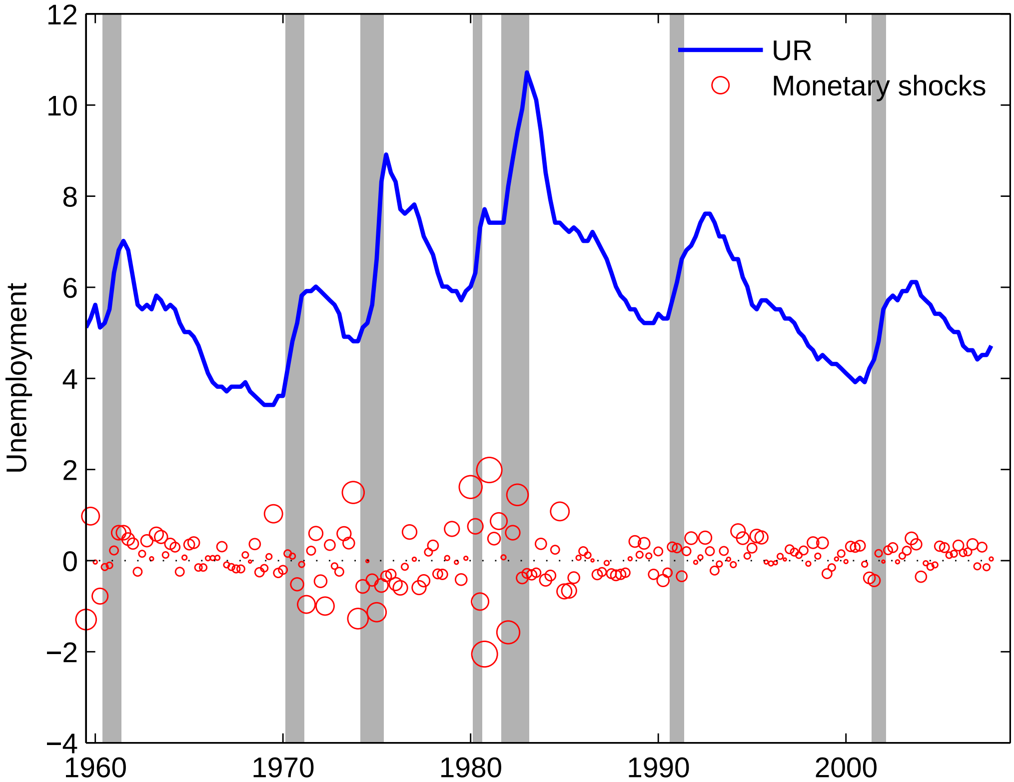

Figure 7: Unemployment rate –the business cycle indicator (solid line, left scale)–, and estimated monetary shocks (circles, right scale) with larger circles indicating larger shocks.

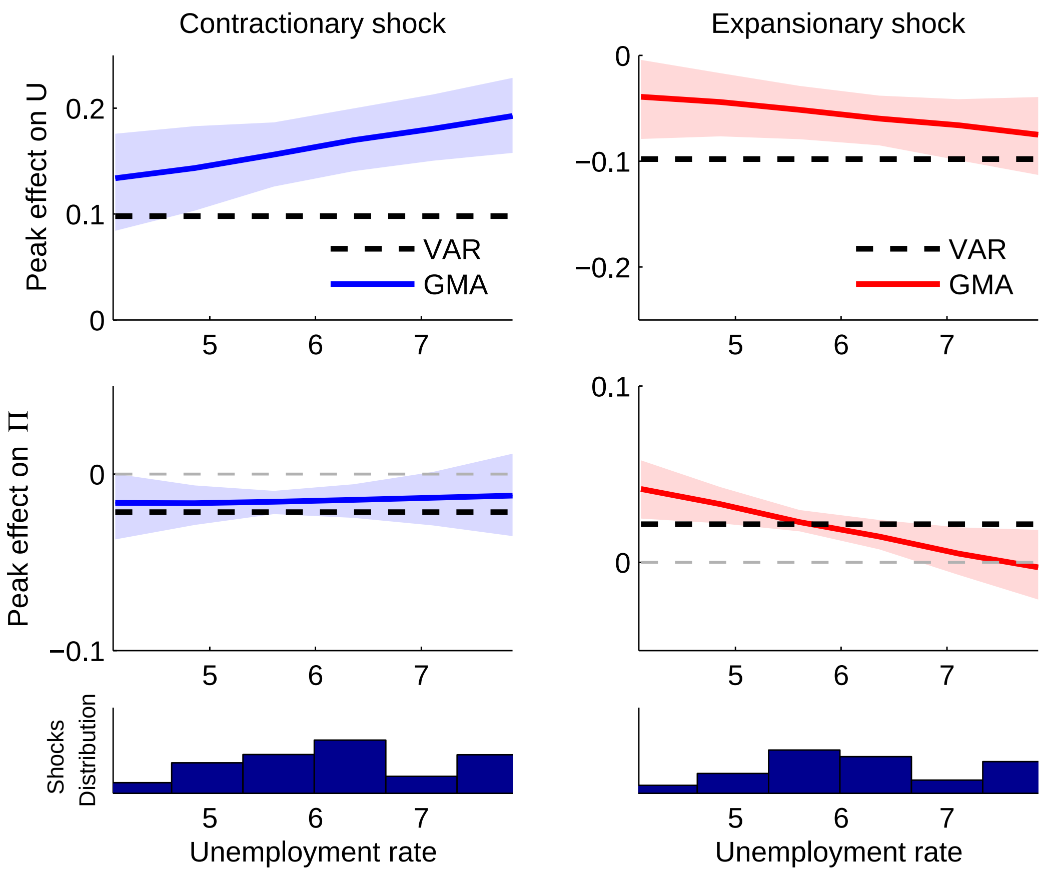

Figure 8: Peak effect of monetary policy on unemployment and inflation (in percentage points) as a function of the state of the business cycle (measured with the unemployment rate) for one standard deviation contractionary monetary shocks (left panel) and expansionary monetary shocks (right panel). The dashed lines represent the 5th and 95th posterior percentiles. The thick-dashed line is the linear VAR estimate. The bottom panel plots the distribution of (respectively) contractionary shocks and expansionary shocks over the business cycle. Estimation using data covering 1959-2007.

Table 1: Summary statistics for Monte Carlo simulation with a linear model

	U		π		ffr	
	VAR	**GMA**	**VAR**	**GMA**	**VAR**	**GMA**
MSE	0.057	0.043	0.077	0.041	0.003	0.002
Avg length (at peak effect)	0.16	0.13	0.27	0.11	0.05	0.03
Coverage rate (at peak effect)	0.94	0.83	1	0.78	0.94	0.93

Note: Summary statistics over 50 Monte-Carlo replications. MSE is the mean-squared error of the estimated impulse response function over horizons 1 to 25. Avg length is the average distance between the lower (2.5%) and upper (97.5%) confidence bands at the time of peak effect of the monetary shock. The coverage rate is the frequency with which the true value lays within 95 percent of the posterior distribution. The VAR estimates and confidence bands are obtained from a Bayesian VAR with Normal-Whishart priors. U, **π** and ffr denote respectively unemployment, inflation and the fed funds rate.

Table 2: Summary statistics for Monte Carlo simulation with asymmetry

	a^{+}-a^{-}		
	y	**π**	**ffr**
Frequency of rejection of zero coefficient	0.94	0.90	0.08
Mean (true value)	-0.82 (-1.00)	-0.50 (-0.60)	0.03 (0.00)
Std-dev	0.28	0.17	0.12
Coverage rate	0.82	0.86	0.88

Note: Summary statistics over 50 Monte-Carlo replications. For each coefficient of interest, "Frequency of rejection of zero coefficient" is the frequency that 0 lies outside 90 percent of the posterior distribution, and "Coverage rate" is the frequency with which the true value lies within 90 percent of the posterior distribution. y, π and ffr denote respectively output, inflation and the fed funds rate.

Table 3: Summary statistics for Monte Carlo simulation with asymmetry and state dependence

	γ^{+}-γ^{-}		α^{+}-α^{-}		γ^{+}		γ^{-}	
	y	**π**	**y**	**π**	**y**	**π**	**y**	**π**
Frequency of rejection of zero coefficient	0.96	0.03	0.82	0.80	0.87	0.06	0.20	0.05
Mean (true value)	0.96 (1.00)	0.02 (0.00)	-0.78 (-1.00)	-0.48 (-0.60)	0.71 (1.00)	0.00 (0.00)	-0.21 (0.00)	-0.00 (0.00)
Std-dev	0.26	0.17	0.37	0.23	0.31	0.19	0.23	0.19
Coverage rate	0.84	0.92	0.71	0.70	0.68	0.92	0.65	0.90

Note: Summary statistics over 50 Monte-Carlo replications. For each coefficient of interest, "Frequency of rejection of zero coefficient" is the frequency that 0 lies outside 90 percent of the posterior distribution, and "Coverage rate" is the frequency with which the true value lies within 90 percent of the posterior distribution. y and π denote respectively output and inflation.

Table 4: Marginal densities, BIC and AIC

	VAR	**GMA(1)**	**GMA(1) Asymmetry**	**GMA(2) Asymmetry**	**GMA(3) Asymmetry**	**GMA(1) Asymmetry State dep.**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
(log) marginal density	112	118	127	138	107	158

Note: Trivariate models with unemployment, PCE inflation and the fed funds rate estimated over 1959-2007. The VAR estimates and confidence bands are obtained from a Bayesian VAR with Normal-Whishart priors.