Functional Approximation of Impulse Responses☆

Regis Barnichon[a], Christian Matthes[b]

[a]*FRB San Francisco*
[b]*FRB Richmond*

Abstract

This paper proposes a new method, Functional Approximation of Impulse Responses (FAIR), to estimate the dynamic effects of structural shocks. FAIR approximates impulse responses with a few basis functions and then directly estimates the moving average representation of the data. FAIR can offer a number of benefits over earlier methods, including VAR and Local Projections: (i) parsimony and efficiency, (ii) ability to summarize the dynamic effects of shocks with a few key moments that can directly inform model building, (iii) ease of prior elicitation and structural identification, and (iv) flexibility in allowing for non-linear effects of shocks while preserving efficiency. We illustrate these benefits by summarizing the dynamic effects of monetary shocks, notably their asymmetric effects, with a few key statistics.

JEL classifications: C14, C32, C51, E32, E52
Keywords:

☆First draft: March 2014. We would like to thank Luca Benati, Francesco Bianchi, Christian Brownlees, Fabio Canova, Tim Cogley, Davide Debortoli, Wouter den Haan, Jordi Gali, Yuriy Gorodnichenko, Eleonora Granziera, Oscar Jorda, Thomas Lubik, Jim Nason, Kris Nimark, Mikkel Plagborg-Moller, Giorgio Primiceri, Ricardo Reis, Jean-Paul Renne, Barbara Rossi, Konstantinos Theodoridis, 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, SED 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. The views expressed here do not necessarily reflect those of the Federal Reserve Bank of Richmond or San Francisco or of the Federal Reserve System. Any errors are our own.

1. Introduction

The impulse response function (IRF) is a popular tool to describe the dynamic effects of shocks on macroeconomic time series, and this paper proposes a new method, Functional Approximation of Impulse Responses (FAIR), to estimate IRFs.

The FAIR methodology has two distinct features. First, instead of assuming the existence of a vector autoregression (VAR) representation, FAIR directly estimates the vector moving-average (VMA) representation of the data, i.e., FAIR directly estimates the IRFs. This confers a number of advantages, notably the ability to impose prior information and structural identifying restrictions in a flexible and transparent fashion, as well as the possibility to allow for a large class of non-linear effects. Second, FAIR approximates the impulse response functions with a few basis functions. The approximation serves as a dimension reduction tool, which makes the estimation of the VMA feasible. While different families of basis functions are possible, Gaussian basis functions can be of particular interest for macro applications: a one Gaussian function approximation can summarize a monotonic or hump-shaped IRF with only three parameters, each capturing a separate and interpretable feature of the IRF: (a) the magnitude of the peak effect of a shock, (b) the time to that peak effect, and (c) the persistence of that peak effect. This ability to summarize a high-dimensional IRF with a few key parameters amenable to statistical inference can make FAIR helpful to evaluate and guide the development of successful models.

After establishing the good finite sample properties of FAIR with Monte-Carlo simulations, we illustrate the benefits of FAIR by studying the effects of monetary shocks. FAIR can incorporate the different identification schemes found in the structural VAR literature, and we present results from three popular identification schemes: a recursive identification scheme, (ii) a set identification scheme based on sign restrictions, and (iii) a narrative identification scheme where a series of shocks has been previously identified (possibly with measurement error) from narrative accounts.

First, we focus on linear models, and we illustrate how influential stylized facts about the magnitude and dynamics of the effects of monetary shocks can be recast into statements about individual FAIR parameters. Second, we use FAIR to explore whether monetary shocks have asymmetric effects on unemployment, i.e., whether a contractionary monetary shock has a stronger effect (being akin to pulling on a string) than an expansionary shock (being akin to pushing on a string). While this question is central to the conduct of monetary policy, the evidence for asymmetric effects is relatively thin and inconclusive,[1] in part because estimating the asymmetric effect of a monetary shock is difficult within a VAR framework. Consistent with the string metaphor, we find evidence of strong asymmetries in the effects of monetary shocks. Regardless of our identification scheme, a contractionary shock has a strong adverse effect on unemployment, while an expansionary shock has little effect on unemployment. The response of inflation is also asymmetric but in the opposite direction of unemployment: prices react less when unemployment reacts more and vice-versa.

Our functional approximation of the IRFs has intellectual precedents in the distributed lags literature, notably Almon (1965), Jorgenson (1966), Hansen and Sargent (1981) and Ito and Quah (1989) approximation of the distributed lag function with polynomial or rational functions. More recently, Plagborg-Moller (PM, 2017) proposes a Bayesian method to directly estimate the structural VMA representation of the data. Different from our functional approximation approach, PM reduces estimation variance by using prior information about the shape and the smoothness of the impulse responses.

Section 2 presents our functional approximation of impulse responses and discusses the benefits of using Gaussian basis functions, Section 3 re-visits the linear effects of monetary shocks with FAIR; Section 4 extends FAIR to non-linear models; Section 5 studies the asymmetric effects of monetary shocks; Section 6 concludes.

[1] For instance, while Cover (1992), Angrist et al. (2016) and Tenreyro and Thwaites (2016) find evidence of asymmetric effects, Weise (1999), Ravn and Sola (2004) and Lo and Piger (2005) find nearly symmetric effects.

2. Functional Approximation of Impulse Responses (FAIR)

We first introduce FAIR in univariate setting and then generalize it to a multivariate setting.

2.1. Univariate setting

For a stationary univariate data generating process, the IRF corresponds to the coefficients of the moving-average model

$$y_t = \sum_{h=0}^{H} \psi(h)\varepsilon_{t-h} \tag{1}$$

where ε_t is an i.i.d. innovation with $E\varepsilon_t = 0$ and $E\varepsilon_t^2 = 1$, and H is the number of lags, which can be finite or infinite. The lag coefficients $\psi(h)$ is the impulse response of y_t at horizon h to innovation ε_t.

Since $\psi(h)$ is a large (possibly infinite) dimensional object, estimation of (1) can be difficult. Instead, one can invoke the Wold decomposition theorem to re-write (1) as an AR(∞)

$$A(L)y_t = \varepsilon_t \tag{2}$$

where $A(L) = \Psi(L)^{-1}$ with $\Psi(L) = \sum_{h=0}^{H} \psi(h)L^h$ and L the lag operator.

Since estimating an AR(∞) is not possible in finite sample, in practice researchers estimate a finite order AR(p) meant to approximate the AR(∞), and the IRF is recovered by inverting that AR(p). In effect, the AR(p) approximation serves as a dimension reduction tool that makes the estimation of (1) feasible. For instance, an AR(1) approximates the coefficients of the MA(∞) as an exponential function of a single parameter.

In this paper, we propose an alternative dimension-reducing tool (Functional Approximation of Impulse Responses or FAIR), which consists in representing the impulse response function as an expansion in basis functions. Instead of estimating an intermediate model –the AR(p)– and then recovering the impulse response, FAIR directly estimates the object of scientific interest; the impulse response function. Specifically, a functional approximation of ψ consists in decomposing ψ into a sum of basis functions

$$\psi(h) = \sum_{n=1}^{N} a_n g_n(h), \quad \forall h \geq 0 \tag{3}$$

with $g_n : \mathbb{R} \to \mathbb{R}$ the nth basis function, $n = 1, .., N$.

Different families of basis functions are possible, and in this paper we use Gaussian basis functions and posit

$$\psi(h) = \sum_{n=1}^{N} a_n e^{-(\frac{h-b_n}{c_n})^2}, \quad \forall h \geq 0 \tag{4}$$

with a_n, b_n, and c_n parameters to be estimated. Since model (4) uses N Gaussian basis functions, we refer to this model as a FAIR$_\mathcal{G}$ of order N. In the appendix, we prove that any mean-reverting impulse response function can be approximated by a sum of Gaussian basis functions.

A researcher can choose to use the basis function approximation (4) for $h \geq 0$ or for $h > 0$. As we will see in the identification section, there are benefits of excluding $h = 0$, i.e., $\psi(0)$, from the approximation, because it allows us to incorporate short-run restrictions into FAIR without overly restricting the IRF dynamics. Thus, for the rest of the paper, we will treat the elements of the impact coefficient $\psi(0)$ as a free parameter, although there is nothing in our approach that requires us to make this choice.

2.2. Multivariate setting

While the discussion has so far concentrated on a univariate process, we can easily generalize the FAIR approach to a multivariate setting. Consider the structural vector moving-average model of $\boldsymbol{y}_t$, a $(k \times 1)$ vector of stationary variables,

$$\boldsymbol{y}_t = \boldsymbol{\Psi}(L)\boldsymbol{\varepsilon}_t, \quad \text{where } \boldsymbol{\Psi}(L) = \sum_{h=0}^{H} \boldsymbol{\Psi}_h L^h \tag{5}$$

where $\boldsymbol{\varepsilon}_t$ is the vector of i.i.d. structural innovations with $E\boldsymbol{\varepsilon}_t = \mathbf{0}$ and $E\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t' = \mathbf{I}$, and H is the number of lags, which can be finite or infinite. The matrix of lag coefficients $\boldsymbol{\Psi}_h$ contains the impulse responses of $\boldsymbol{y}_t$ at horizon h to the structural shocks $\boldsymbol{\varepsilon}_t$.

Setting aside the issue of identification, the common strategy to recover (5) is identical to the univariate case: rewrite (5) as a VAR(∞) and then estimate its VAR(p) approximation.

In this paper, we propose to use FAIR and directly estimate the impulse response functions by approximating each element of $\boldsymbol{\Psi}_h$, denoted by $\psi(h)$, as in (4). With k^2 IRFs to approximate, the parameter vector is of dimension $3 * N * k^2$. If the contemporaneous impact matrix $\boldsymbol{\Psi}_h$ is left unrestricted, this adds another k^2 terms to estimate.

2.3. FAIR benefits

We now argue that FAIR models with Gaussian basis functions (FAIR$_{\mathcal{G}}$) can be attractive in macro applications, and we will pay particular attention to the FAIR approximation with one Gaussian basis function, or FAIR$_{\mathcal{G}_1}$,

$$\psi(h) = a e^{-\frac{(h-b)^2}{c^2}} \tag{6}$$

illustrated in the top panel of Figure 1.

Parsimony of FAIR

A first advantage of using Gaussian basis functions is that a small number of Gaussians can already capture a large class of IRFs, that is Gaussian basis functions may offer an attractive dimension-reduction tool.

To illustrate this point, Figure 2 plots the IRFs of unemployment, inflation and the fed funds rate to a shock to the fed funds rate estimated from a standard VAR specification with a recursive ordering, along with the IRFs approximated with Gaussian basis functions.[2] For

[2] We describe the exact VAR specification in section 5. In Figure 2, the parameters of the Gaussian basis functions were set to minimize the discrepancy (sum of squared residuals) with the VAR-based impulse responses. Importantly, this is *not* how we estimate FAIR models.

unemployment and the fed funds rate we use a $\text{FAIR}_{\mathcal{G}_1}$. We can see that this simple model already does a good job at capturing the responses of unemployment and the fed funds rate implied by the VAR, while reducing the dimension of each IRF in Figure 2 from 25 to only 3 parameters. To capture the oscillating pattern of inflation following a monetary shock, two Gaussian basis functions are necessary, and we can see in Figure 2 that the IRF approximated with two Gaussian basis functions (a $\text{FAIR}_{\mathcal{G}_2}$) does a good job there as well.

With a small number of Gaussian basis functions per IRF, directly estimating the VMA representation of the data becomes feasible.[3] and this will allow us to estimate a VMA representation of the data.

Interpretability and portability of FAIR coefficients

A second advantage of using Gaussian basis functions is that the estimated coefficients can have a direct economic interpretation in terms of features of the IRFs. This stands in contrast to VARs where the IRFs are non-linear transformations of the VAR coefficients.

The ease of interpretation is most salient in a $\text{FAIR}_{\mathcal{G}_1}$ model (6) where the a, b and c coefficients have a direct economic interpretation, and in fact capture three separate characteristics of a hump-shaped impulse response. To see that graphically, the top panel of Figure 1 shows a hump-shape impulse response parametrized with a one Gaussian basis function: 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 $\tau = c\sqrt{\ln 2}$. These a-b-c coefficients are generally considered the most relevant characteristics of an impulse response function and

[3]For instance, for a trivariate model with three structural shocks, a $\text{FAIR}_{\mathcal{G}_1}$ only needs 27 parameters (9 impulse responses times 3 parameters per impulse response, ignoring intercepts) to capture the whole set of impulse responses $\{\mathbf{\Psi}_h\}_{h=1}^{H}$. In the example of Figure 2 where the impulse responses of inflation are approximated by a richer $\text{FAIR}_{\mathcal{G}_2}$, there would be 36 parameters ($3^2 * 2 + 6 * 3 = 36$). In comparison, a quarterly VAR with 3 variables and 4 lags has $4 * 3^2 = 36$ free parameters, and a monthly VAR with 12 lags has $12 * 3^2 + 6 = 108$ free parameters. Unlike VARs, the number of FAIR parameters to estimate need not increase with the frequency of the data.

the most discussed in the literature.

With more than one basis function, the ease of interpretation of the estimated parameters is no longer guaranteed. However, in some cases a FAIR$_{\mathcal{G}}$ model with multiple Gaussian-basis functions can retain its interpretability. For instance, consider an oscillating pattern as in the bottom panel of Figure 1, which is typical of the response of inflation to a contractionary monetary shock. In that case, the FAIR$_{\mathcal{G}}$ parameters can retain their interpretability if the first Gaussian basis function captures the first-round effect of the shock, a positive hump-shaped response of inflation usually refereed to as the price puzzle (Christiano et al. 1999), while the second Gaussian function captures the larger second-round effect, a negative hump-shaped effect. Going back to our VAR example from Figure 2, the two Gaussian basis functions used to match the impulse response of inflation show little overlap. In particular, one can summarize the 2^{nd}-round effect of the shock (the main object of scientific interest in the case of the inflation response to monetary shocks) with the *a*-*b*-*c* parameters of the second basis function: in that case, the a_2, b_2 and c_2 capture the peak effect, the time to peak effect, and the persistence of the 2^{nd}-round effect of the shock.[4]

Going beyond interpretation and ease of prior elicitation, the ability to summarize the dynamic effects of a shock makes FAIR$_{\mathcal{G}}$ particularly helpful to evaluate and guide the development of successful models. IRFs are high-dimensional objects whose main features can be hard to summarize and assess statistically. And while researchers traditionally resort to large panels of IRF plots to represent the responses of many variables to different shocks (e.g., Amir-Ahmadi et al., 2016), the overload of information can blur the key implications

[4]Formally, one can write a "no-overlap" condition as

$$\int_{h_2}^{\infty} g_1(h)dh \leq \epsilon \simeq 0 \quad \text{with } h_2 \text{ such that } \int_{h_2}^{\infty} g_2(h)dh = \alpha \simeq 1 \tag{7}$$

with g_n denoting the nth basis function. That condition is a restriction on the b_n and c_n coefficients that ensures that the two basis functions g_1 and g_2 have limited overlap. Intuitively, for say $\epsilon = .05$ and $\alpha = .9$, the condition states that the first Gaussian function (g_1) –the first-round effect of the shock– has a negligible (5 percent or less) overlap with 90 percent of the second Gaussian function (g_2) –the second-round effect of the shock–.

for theoretical modeling and make comparison of IRF estimates across different schemes, different model specifications and/or different sample periods difficult. By providing interpretable and portable moments amenable to statistical inference, FAIR can facilitate the emergence of a set of robust findings.

2.4. FAIR estimation

FAIR models can be estimated using maximum likelihood or Bayesian methods. The computational cost is not as trivial as OLS in the case of VARs, but the estimation is simple and relatively easy thanks to modern computational capabilities.

The key step in the estimation of FAIR models is the computation of the likelihood function of a VMA model. Since we assume that the variables in our model are stationary, we can truncate the moving average model at a large enough horizon H and the approximation error due to this truncation will be negligible. To initialize the computation, we set the first H structural innovations $\{\boldsymbol{\varepsilon}_j\}_{j=-H}^{0}$ to zero. Conditional on a set of parameters, we can then recursively back out the sequence of one step ahead forecast errors from the FAIR model (5) and then compute the likelihood function. We can then use either a Metropolis-Hastings algorithm to approximate the posterior distribution, or use an optimization algorithm to maximize the likelihood function and obtain Maximum Likelihood estimates of the parameters. More details are available in the Online Appendix.

A potentially important problem when estimating moving-average models is the issue of under-identification: In linear moving average models, 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 (e.g., Lippi and Reichlin, 1994). By constructing the likelihood recursively using past observations, our algorithm will effectively estimate the fundamental moving-average representation of the data.[5] In principle, one could

[5]If our estimation algorithm chose parameter estimates that implied non-invertibility, our $\{\boldsymbol{\varepsilon}_t\}$ estimates would ultimately grow very large, and this situation would lead to very low likelihoods, since we assume that

remedy this limitation and estimate non-invertible moving-average representations by using the procedure recently proposed by Plagborg-Moller (2017),[6] provided that one has enough prior information to favor one moving-average representation over the others.

To choose N, the order of the FAIR model, the researcher can either choose to restrict himself to a class of functions if prior knowledge on the shape of the impulse response is available (for instance, using only one Gaussian basis function) or, similarly to a BIC criterion, use likelihood ratio tests/posterior odds ratios (assigning equal probability to any two models) to compare models with increasing number of basis functions. Finally, note that FAIR models can only be estimated for stationary series (so that the moving-average can be truncated). If the data are non-stationary, we can (i) allow for a deterministic trend in equation (5) and/or (ii) difference the data, and then proceed exactly as described above.[7]

2.5. Imposing structural identifying assumptions

To give a structural interpretation to the VMA innovations, researchers must use identifying assumptions. Given our later focus on the effects of monetary policy, we will discuss three popular approaches in the monetary literature: a recursive identification scheme, (ii) a narrative identification scheme where a series of shocks has been previously identified (possibly with measurement error) from narrative accounts, and (iii) a set identification scheme based on sign restrictions. We also mention how FAIR could open the door for more general identification schemes based on shape-restrictions.

Short-run restrictions and recursive ordering: Short-run restrictions in a fully

the $\{\boldsymbol{\varepsilon}_t\}$ are standard normal. As a result, our estimation procedure will effectively estimate the invertible representation.

[6]By using the Kalman filter with priors on the H initial values of the shocks $\{\boldsymbol{\varepsilon}_{-H}...\boldsymbol{\varepsilon}_0\}$, Plagborg-Moller's procedure can handle the estimation of both invertible and non-invertible representations and thus does not restrict the researcher to the invertible moving-average representation. However, unlike our proposed approach, that procedure would be difficult to implement in non-linear models.

[7]The presence of co-integration does not imply that a FAIR model in first-difference is misspecified. The reason is that a FAIR model directly works with the moving-average representation and does not require inversion of the moving-average, unlike VAR models. After estimation, one can even test for co-integration by testing whether the matrix sum of moving-average coefficients ($\sum_{h=1}^{H}\sum_{l=0}^{h}\boldsymbol{\Psi}_l$) is of reduced rank (Engle and Yoo, 1987).

identified model consists in imposing restrictions on the contemporaneous matrix $\mathbf{\Psi}_0$. For instance, a popular restriction based on a timing assumption is that a subset of the variables (ordered first in the vector $\boldsymbol{y}_t$) do not react on impact to that shock. As with VARs, this is ensured by setting the upper-right block of $\mathbf{\Psi}_0$ to zero. Note that since a FAIR approximation restricts the dynamics of the IRF, combining a FAIR approximation with a short-run restriction could overly restrict the dynamics of the IRF, as the short-run restriction at $h = 0$ would affect the whole path of the IRF. To avoid this implication, we can treat the elements of the initial impact matrix $\mathbf{\Psi}_0$ as free parameters and only use the FAIR approximation (4) for $h > 0$.

Narrative identification: In a narrative identification scheme, a series of shocks has been previously identified from narrative accounts. For that case, we can proceed as with the recursive identification, because the use of narratively identified shocks can be cast as a partial identification scheme. We order the narratively identified shocks series first in $\boldsymbol{y}_t$, and we assume that $\mathbf{\Psi}_0$ has its first row filled with 0 except for the diagonal coefficient, which implies that the narratively identified shock does not react contemporaneously to other shocks. In other words, we are assuming that the narrative shocks are contemporaneously correlated with the true monetary shocks and uncorrelated with other structural shocks.[8]

Sign restrictions: Set identification through sign restrictions consists in imposing sign-restrictions on the sign of the $\mathbf{\Psi}_h$ matrices, i.e., on the impulse response coefficients at different horizons. One can impose sign-restrictions on only the impact coefficients (captured by $\mathbf{\Psi}_0$, which could be left as a free parameter in this case) and/or sign restrictions on the impulse response. In a $\text{FAIR}_{\mathcal{G}_1}$ model, the sign restriction applies for the entire horizon of the impulse response. With oscillating pattern and a higher-order $\text{FAIR}_{\mathcal{G}}$ model, we can impose sign restrictions over a specific horizon by using priors on the location and the sign of the loading of the basis functions.

[8]Our procedure allows the narrative shocks to contain measurement error, as long as the measurement error is independent of structural shocks. This approach is similar to an external instrumental approach, in the sense of Stock (2008).

Identification restrictions through priors: When the FAIR parameters can be interpreted as "features" of the impulse responses, one can go beyond sign-restrictions and envision set identification through shape restrictions. Using the insights from Baumeister and Hamilton (2015), one could implement shape restrictions through informative priors on the *a*-*b*-*c* coefficients. For instance, one could posit priors on the location of the peak effect or posit priors on the persistence of the effect of the shock, among other possibilities.

2.6. Assessing FAIR performances from Monte-Carlo simulations

In this section, we summarize the findings of a set of Monte Carlo exercises that we use to asses the properties of our FAIR estimator in a linear context. Details can be found in the Online Appendix. We use two sets of data generating processes (DGP), one where the FAIR model is correctly specified and one where the true model is a VAR.

In the first set of simulations where the DGP is a trivariate $\text{FAIR}_{\mathcal{G}_1}$ model, we found that a correctly-specified FAIR model can generate substantially more accurate impulse response estimates (in a mean-squared error sense) than a VAR model, independently of whether the VAR includes many lags or is parsimoniously parametrized. To give a number, we found that, on average across Monte Carlo samples, the mean-squared error of the VAR is 150 percent higher than that of the FAIR model with flat priors. Intuitively, the VAR can only approximate the DGP with a large number of lags, and in that case the VAR parameters are imprecisely estimated. In those situations, a $\text{FAIR}_{\mathcal{G}}$ model can provide a useful alternative.

In the second set of simulations where the DGP is a VAR, we found that a misspecified FAIR performs just as well (or even slightly better in a mean-squared error sense) than a well-specified VAR model. The reason for the superior performances of the misspecified FAIR is the fact that the estimated VAR often shows counterfactual oscillation patterns. Indeed, because VAR-based IRFs are linear combinations of damped sine-cosine functions, the estimated VAR-based IRFs can display counter-factual oscillations even if the true data generating VAR does not feature these fluctuations. With its tighter parametrization, a $\text{FAIR}_{\mathcal{G}}$ with only a few basis functions avoids this problem. That being said, our goal is not to claim that FAIR models are always superior to VARs. Instead, the simulations are meant to convey that FAIR models can provide a useful alternative to VARs.

2.7. Relation to alternative IRF estimators

VARs have been the main approach to estimate IRFs since Sims (1980), but an increasing number of papers are now relying on Local Projections (LP, Jorda 2005) –themselves closely related to Autoregressive Distributed Lags (ADL, e.g., Hendry 1984)– to directly estimate impulse response functions.

FAIR aims to straddle between the parametric parsimony of VARs and the flexibility of LP. Indeed, while LP (or ADL in its naive form) is model-free –not imposing any underlying dynamic system–, this can come at an efficiency cost (Ramey, 2012), which can make inference difficult. In contrast, by positing that the response function can be approximated by one (or a few) Gaussian functions, FAIR imposes strong dynamic restrictions between the parameters of the impulse response function, which can improve efficiency. Moreover, FAIR alleviates another source of inefficiency in LP, namely the presence of serial correlation in the LP regression residuals. By modeling the behavior of a *system* of key macroeconomic variables (similarly to a VAR), a multivariate FAIR model is effectively modeling the serial correlation present in LP residuals, and this can further improve efficiency. Naturally, all these statement are only valid under the assumption that IRFs can be well approximated by a few Gaussian functions. In this respect, FAIR is best seen as complementing the model-free nature of LP.

Another benefit of FAIR over VAR and LP/ADL is the ease of prior elicitation and structural identification. In VARs, identification can be thorny and non-transparent (e.g., Baumeister and Hamilton, 2015), because the impulse-responses are non-linear transformations of the VAR coefficients. In LP/ADL, the scope for identification is more limited, because LP/ADL are univariate models, so that a series of previously identified shocks (or instruments) is typically required (e.g., Auerbach and Gorodnichenko, 2013).

3. A FAIR summary of the linear effects of monetary shocks

In this section, we illustrate the benefits of FAIR by summarizing stylized facts from the monetary literature with *a*-*b*-*c* parameters.

We consider a model of the US economy in the spirit of Primiceri (2005), where $\boldsymbol{y}_t$ includes the unemployment rate, the PCE inflation rate and the federal funds rate. We use one Gaussian basis function to parametrize the impulse responses of unemployment and the fed funds rate. For the response of inflation, we use two Gaussian functions to allow for the possibility of a price puzzle in which inflation displays an oscillating pattern.[9]

To put our results in the context of the literature, we identify monetary shocks using three different schemes: (i) a timing restriction whereby monetary policy affects macro variables with a one period lag (e.g., Christiano et al., 1999), (ii) a narrative approach based on Romer and Romer (2004) and extended until 2007 by Tenreyro and Thwaites (2016), and (iii) sign restrictions. For the latter scheme, we posit that positive monetary shocks are the only shocks that (a) raise the fed funds rate and (b) lower inflation roughly two years after the shock. Specifically, with a two Gaussian basis function specification for the response of inflation, we impose that the loading on the second basis function is negative ($a_{\pi,2} < 0$), while the first basis function (meant to capture a possible price puzzle) can load positively or negatively but is restricted to peak within a year ($b_{\pi,1} \leq 4$) with a "half-life" of at most a year ($c_{\pi,1}\sqrt{\ln 2} \leq 4$). In words, our sign-restriction is that the price puzzle cannot last for too long, so that the response of inflation must be negative after roughly two years. In the appendix, we plot the prior IRF of inflation implied by these priors. To ensure the same sample period across identification schemes, the data cover 1969Q1 to 2007Q4.

We first display our results in the usual way, and Figure 3 plots the impulse response

[9] As priors, we use very loose Normal priors on the *a*-*b*-*c* coefficients that are centered on the values obtained by matching the impulse responses obtained from the VAR and with standard-deviations $\sigma_a = 10$ ppt, $\sigma_b = 10$ quarters and $\sigma_c = 20$ quarters with the constraint $c > 0$ (in "half-life" units, this σ_c corresponds to a half-life of about 4 years, a very persistent IRF). To illustrate that these are very loose priors, in the appendix we show some corresponding prior IRFs.

functions of unemployment, inflation and the fed funds rate to a 100 basis point monetary innovation, as in Coibion (2012). Figure 4 presents the same results but through the lens of the FAIR parameter estimates: the blue error bars show the 90% posterior distributions of the *a*-*b*-*c* parameters summarizing the IRFs of unemployment and inflation.[10]

Peak effects of monetary shocks: the a parameter

In an influential paper, Coibion (2012) first drew attention to the fact that Romer and Romer (R&R, 2002) obtained much larger effects of monetary shocks than implied by structural VARs with a recursive ordering (e.g., Christiano et al., 2005).

Since impulse responses are high-dimensional and difficult to compare, Coibion (2012) made his point by summarizing the IRFs to monetary shocks with their peak effects. In a $\text{FAIR}_{\mathcal{G}_1}$ model, the peak effect is directly picked up by the a parameter, so that we can revisit Coibion's findings in lights of our a parameter estimates. A benefit of FAIR in this context is that FAIR delivers a posterior distribution for the a parameters, which allows us to qualify Coibion's findings with confidence intervals. Our estimated effects of monetary shocks are in line with Coibion (2012, Table 2), and the R&R shocks have a much larger effect than either recursively-identified or sign-identified shocks. For instance, a contractionary recursively-identified shock raises unemployment by $a_u^{rec} = .24_{[.19,.30]}$ ppt at its peak whereas a Romer and Romer shock raises unemployment by $a_u^{nar} = .56_{[.38,.79]}$, where the main entry denotes the median value and the subscript entry denotes the 90 percent credible interval.[11] Since the 90% credible interval for a_u^{nar} excludes the credible intervals for a_u^{rec} and a_u^{sgn}, we can conclude like Coibion (2012) that R&R shocks do have a larger effects on unemployment. A similar result holds for the response of inflation.

[10] For the impulse response of inflation, the *a*-*b*-*c* parameters of the second basis functions retain a useful interpretation, because the "no-overlap" condition (footnote 4) is satisfied by more than 99% of MCMC draws (taking with $\alpha = .9$ and $\epsilon = .05$). This can be seen graphically in the two median basis functions plotted in Figure 3.

[11] While Romer and Romer (2002) report a slightly larger baseline estimate for the peak response of unemployment (0.9 ppt), Coibion (2012) note that the R&R results are sensitive to the number of lags in their ADL model. Using instead a number of lags consistent with an AIC criterion, Coibion (2012) estimates a peak response of unemployment of 0.6ppt, in line with our results.

As Coibion noticed, an important caveat to this conclusion is that the monetary impetus is different across identification schemes, and indeed we also find that the R&R shocks generate a much larger response of the fed funds rate, peaking at close to 200 basis points instead of 100bp for the other identification schemes (figure 3). A simple approach to address this issue is to normalize the impulse responses by the peak response of the fed funds rate, as opposed to the impact response of the fed funds rate. After rescaling, the impulse responses of the fed funds rate all peak at 100 basis points, implying more similar monetary impetus across identification schemes. In a $\text{FAIR}_{\mathcal{G}}$ model, the normalization leaves the b and c parameters unchanged but rescales the a parameters capturing the peak effects. After constructing the posterior distributions of the rescaled a parameters (black error bars), we find that the point estimates for the peak effects of the R&R shocks are still slightly larger (especially for inflation), but the estimates are no longer significantly different: the error-bars for rescaled a show overlap across identification schemes. While Coibion (2012, figure 4) ultimately reached a similar conclusion, his conclusion was based on the IRF point estimates. In contrast, our conclusion is based on the posterior distribution of the peak effects of monetary shocks.[12]

Dynamic effects of monetary shocks: the b and c parameters

A second set of influential facts pertains to the dynamic effects of monetary shocks. In particular, three stylized facts that guided the development of New-Keynesian models (e.g., Mankiw and Reis, 2002, Gali, 2008) are (i) unemployment (or output) and inflation respond in a hump-shaped fashion, (ii) the peak response of inflation is delayed compared to the peak response of unemployment, and (iii) real and nominal variables show persistent responses.

These facts were based on visual inspections of the impulse response functions (e.g., Christiano et al. 2005, p8), but they refer precisely to the b and c coefficients of a $\text{FAIR}_{\mathcal{G}}$ model, and FAIR can make them more precise by providing confidence intervals around

[12] A natural question is then why the R&R shocks imply larger changes in the fed funds rate than other identification schemes. Coibion (2012) argues that the 1979-1982 period of non-FFR targeting, a period during which the identification of shocks is particularly thorny for the R&R approach, is behind the discrepancy.

them.

The IRF of unemployment is indeed hump-shaped with a peak occurring after nine-to-eleven quarters with $b_u^{rec} = 9.6_{[8.1,12.0]}$, $b_u^{nar} = 11.5_{[9.8,13.4]}$ and $b_u^{sgn} = 9.1_{[7.0,10.8]}$ quarters. The IRF of inflation is also hump-shaped (bar the initial price puzzle) but with a peak occurring two-to-four quarters later than the unemployment peak with $\Delta_{\pi,u} b^{rec} = 3.8_{[1.0,7.1]}$, $\Delta_{\pi,u} b^{nar} = 1.8_{[-.1,3.0]}$ and $\Delta_{\pi,u} b^{sgn} = 3.1_{[0.2,5.5]}$ where $\Delta_{\pi,u} b = b_\pi - b_u$. One can even properly test fact (ii) that $b_u < b_\pi$, figure 5 plots the joint posterior distribution of b_u (x-axis) and b_π (y-axis). The dashed red line denotes identical peak times, so that the figure can be seen as a test of no difference in peak times: a posterior density lying above or below the red line indicates statistical evidence for different peak times. Across the three identification schemes, more than 96, 88 and 92 percent of the posterior probability lies above the dashed-red line, confirming that the peak of the unemployment response does indeed occur significantly before that of inflation.

Regarding fact (iii) on persistence, for both inflation and unemployment the IRF returns to half of its peak value in four to six quarters with $c_u^{rec}\sqrt{\ln 2} = 4.7_{[3.7,5.7]}$, $c_u^{nar}\sqrt{\ln 2} = 4.9_{[3.8,6.0]}$ and $c_u^{sgn}\sqrt{\ln 2} = 5.0_{[4.2,5.8]}$. While a casual observation of Coibion (2012)'s IRFs does suggest that dynamics are roughly consistent across schemes, providing confidence intervals around such statements is more involved with estimates from a VAR or LP/ADL framework. Our results show that while the peak effects are different across identification schemes, the dynamic effects of monetary shocks are consistent, as the posterior distributions of c overlap across the identification schemes.

4. Non-linearities with FAIR: assessing the asymmetric effects of shocks

The parsimonious nature of FAIR makes it a good starting point to explore the presence of non-linearities while preserving degrees of freedom. Different non-linear effects of shocks are possible, and in this section, we focus on extending FAIR to estimate possible asymmetric effects of shocks, whereby a positive shock can trigger a different impulse response than a negative shock.

4.1. Introducing asymmetry

With asymmetric effects of shocks, the matrix of impulse responses $\mathbf{\Psi}_h$ depends on the sign of the structural shocks, i.e., we let $\mathbf{\Psi}_h$ take two possible values: $\mathbf{\Psi}_h^+$ or $\mathbf{\Psi}_h^-$, so that a model with asymmetric effects of shocks would be

$$\boldsymbol{y}_t = \sum_{h=0}^{H} \left[\mathbf{\Psi}_h^+ (\boldsymbol{\varepsilon}_{t-h} \odot \mathbf{1}_{\boldsymbol{\varepsilon}_{t-h}>0}) + \mathbf{\Psi}_h^- (\boldsymbol{\varepsilon}_{t-h} \odot \mathbf{1}_{\boldsymbol{\varepsilon}_{t-h}<0}) \right] \tag{8}$$

with $\mathbf{\Psi}_h^+$ and $\mathbf{\Psi}_h^-$ the lag matrices of coefficients for, respectively, positive and negative shocks and $\odot$ denoting element-wise multiplication.

Denoting $\psi^+(h)$, an element of $\mathbf{\Psi}_h^+$ corresponding to a positive shock, a FAIR$_{\mathcal{G}}$ model with asymmetry would be

$$\psi^+(h) = \sum_{n=1}^{N} a_n^+ e^{-\left(\frac{h-b_n^+}{c_n^+}\right)^2}, \quad \forall h \in (0, H] \tag{9}$$

with a_n^+, b_n^+, c_n^+ parameters to be estimated. A similar expression would hold for $\psi^-(h)$.

4.2. Estimation and structural identification

The estimation of FAIR models with asymmetric impulse responses proceeds similarly to the linear case, but the construction of the likelihood involves one additional complication that we briefly mention here and describe in detail in the Appendix: one must make sure that the system $\mathbf{\Psi_0}(\boldsymbol{\varepsilon}_t)\boldsymbol{\varepsilon}_t = \mathbf{u}_t$ has a unique solution vector $\boldsymbol{\varepsilon}_t$ given a set of model parameters and given some vector $\mathbf{u}_t$. With the contemporaneous impact matrix $\mathbf{\Psi_0}$ a function of $\boldsymbol{\varepsilon}_t$, a unique solution is a priori not guaranteed. However, we show in the Appendix that there is a unique solution when we allow the identified shocks to have asymmetric effects in (i) the (full or partial) recursive identification scheme, (ii) the narrative identification scheme, and (iii) the sign-restriction identification scheme under the restriction that $sgn(\det \mathbf{\Psi_0^+}) = sgn(\det \mathbf{\Psi_0^-})$.

4.3. Monte-Carlo simulations

To assess how well a FAIR model can pick up on asymmetries in the data, we construct a moving average model with asymmetric responses to shocks by calculating the moving average representation of a trivariate VAR estimated on US data (details are in the Online Appendix) and then modifying this benchmark model to incorporate asymmetric responses to shocks for two of the three variables in our simulation study. Note that since the DGP is not based on Gaussian basis functions, the FAIR model is misspecified. We simulate 50 samples of length 200 periods each. For the two variables that feature asymmetric responses to shocks, we find that across Monte Carlo samples the FAIR algorithm detects the presence of asymmetric responses 94 percent and 90 percent of the time.[13] For the variable that reacts symmetrically to all shocks, the algorithm detects asymmetry in only 3 percent of the samples.

4.4. Relation to alternative non-linear IRF estimators

The economic literature has so far tackled the estimation of non-linear effects of shocks in two main ways: (i) LP (or ADL) combined with independently identified shocks or instruments, and (ii) Markov switching VARs.

LP can accommodate non-linearities in the response function, and a number of papers recently explored the asymmetric or state dependent effect of shocks using non-linear LP models (e.g., Auerbach and Gorodnichenko, 2013, Tenreyro and Thwaites, 2016). Relative to LP, the higher efficiency of FAIR can be of particular interest for non-linear models where degrees of freedom can decrease rapidly.[14]

Regime-switching VAR models, notably threshold VARs (e.g., Hubrich and Terasvirta, 2013) and Markov-switching VARs (Hamilton, 1989), can capture certain types of non-

[13] We consider that the FAIR estimation detects asymmetry in a Monte Carlo sample if the 90 percent posterior bands of the difference in the response to positive and negative shocks (centered at the median) excludes 0.

[14] For instance, when we allow for asymmetric effects of a shock, the number of parameters to capture an IRF until horizon 20 increases from 20 to 40 when using LP, but only from 3 to 6 when using a $\text{FAIR}_{\mathcal{G}_1}$.

linearities, notably state dependence (whereby the value of some state variable affects the impulse response functions). However, unlike FAIR, regime-switching VARs cannot easily capture asymmetric effects of shocks (whereby the impulse response to a structural shock depends on the sign of that shock).[15]

5. The asymmetric effects of monetary shocks

We now extend the empirical application of Section 3 by allowing monetary shocks to have asymmetric effects.

We can show our results in a standard fashion by displaying IRFs, and Figure 6 plots the IRFs obtained with a recursive ordering (results from the other identification schemes are left for the appendix). However, in the spirit of using *a*-*b*-*c* summary statistics, we will summarize our evidence for asymmetric effects with only one figure that focuses on the peak responses of inflation and unemployment to positive and negative monetary shocks.

Specifically, Figure 7 plots the posterior distribution of a_u and a_π for expansionary monetary shocks (a^-, x-axis) and contractionary shocks (a^+, y-axis) that trigger a peak change in the fed funds rate of a 100 basis points. Recall that a^- denotes the peak response to a decrease in the fed funds rate (an expansionary shock), while a^+ denotes the peak response to an increase in the fed funds rate (a contractionary shock). The dashed red line denotes identical peak responses, i.e., no asymmetry, so that the figure can be seen as a test for the existence of asymmetric effects: a posterior density lying above or below the red line indicates statistical evidence for asymmetric impulse responses. The three rows plot the posterior distributions of a_u and a_π for respectively the recursive identification, the narrative identification, and the sign-restrictions identification. To ease comparison we report $-a_\pi^+$

[15] We make this point formally in the appendix. Intuitively, with regime-switching VAR models, it is assumed that the economy can be in a finite number of regimes, and that each regime corresponds to a different set of VAR coefficients. However, if the true DGP features asymmetric impulse responses, a new set of VAR coefficients would be necessary each period, because the behavior of the economy at any point in time depends on all structural shocks up to that point. As a result, such asymmetric data generating processes cannot be easily captured by threshold VARs or Markov-switching models that only handle a finite (and typically small) number of state variables.

and $-a_\pi^-$, so that the peak effects of inflation and unemployment share the same sign.

Monetary shocks have asymmetric effects: For all three identification schemes, a contractionary monetary shock increases unemployment whereas an expansionary monetary shock has little on effect on unemployment (and non-significantly different from zero). Taking estimation uncertainty into account, the evidence in favor of asymmetry is strong: for the three identification schemes we estimate a .98, .99 and .98 posterior probability that the peak response of unemployment is larger following a contractionary shock than following an expansionary shock (i.e., that $a_u^+ > a_u^-$). In terms of magnitude, note that $a_u^{rec,+} = .22_{[.16,.28]}$ ppt whereas $a_u^{nar,+} = .35_{[.22,.43]}$, so that Coibion's (2012) finding that the R&R shocks have larger effects than recursively-identified shocks is also visible in the asymmetric impulse responses.

The response of inflation also displays an asymmetric pattern: the price level appears more sticky following a contractionary shock than following an expansionary shock. The evidence for asymmetry in the response of inflation is also good, although slightly less strong than with unemployment: the posterior probability that $a_\pi^+ < a_\pi^-$ is 0.93, 0.87 and >0.99 for the three identification schemes. In terms of magnitude, the sign-based identification points to a starker asymmetry in inflation that the other identification schemes. Notably, the response of inflation to a contractionary shock is estimated to be much more muted with sign restrictions ($a_\pi^{sgn,+} = -.03_{[-.04,-.02]}$) than with the other two schemes ($a_\pi^{nar,+} = -.15_{[-.34,.0]}$ and $a_\pi^{rec,+} = -.08_{[-.14,.0]}$).

As a final point, note that the asymmetry in inflation is the *mirror* image of the asymmetry in unemployment: looking at Figure 7, most of the posterior mass is *above* the 45 degree line for the peak response of unemployment, but most of the posterior mass is *below* the 45 degree line for the peak response of inflation. In other words, unemployment reacts less when prices react more, and vice-versa. Interestingly, this is the pattern that one would expect if (i) nominal rigidities were behind the real effects of monetary policy, and (ii) downward nominal rigidities were behind the asymmetric effects of monetary shocks on unemployment.

6. Conclusion

This paper proposes a new method to estimate the dynamic effects of structural shocks by using a functional approximation of the impulse response functions.

FAIR offers a number of benefits over other methods, including VAR and Local Projections: (i) parsimony and efficiency, (ii) ability to summarize the dynamic effects of shocks with a few key moments that can directly inform model building, (iii) ease of prior elicitation and structural identification, and (iv) flexibility in allowing for non-linearities while preserving efficiency. We illustrate these benefits by summarizing the dynamic effects of monetary shocks, notably their asymmetric effects, with a few key statistics.

Although this paper studies the effects of monetary shocks, Functional Approximation of Impulse Responses may be useful in many other contexts, notably when the sample size is small and/or the data are particularly noisy. FAIR could also be used to explore the non-linear effects of other important shocks; notably where the existence of non-linearities remains an important and resolved question, such as fiscal policy shocks (e.g., Auerbach and Gorodnichenko, 2013) or credit supply shocks (Gilchrist and Zakrajsek, 2012).

[1] Angrist, J., Jorda O and G. Kuersteiner. "Semiparametric estimates of monetary policy effects: string theory revisited," Journal of Business and Economic Statistics, forthcoming, 2016

[2] Amir-Ahmadi P., C. Matthes, and M. Wang. "Drifts and Volatilities under Measurement Error: Assessing Monetary Policy Shocks over the Last Century," Quantitative Economics, vol. 7(2), pages 591-611, July 2016

[3] Auclert, A, "Monetary Policy and the Redistribution Channel," Manuscript, May 2017

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

[5] Baumeister, C. and J. Hamilton, "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information," Econometrica, 83(5), 1963-1999, 2015

[6] Bunn P., J. Le Roux, K. Reinold and P. Surico "The Consumption Response to Positive and Negative Income Shocks," CEPR Working Paper DP No. 11829, 2017

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

[8] Christelis D., D Georgarakos, T. Jappelli, L. Pistaferri and M. van Rooij, "Asymmetric Consumption Effects of Transitory Income Shocks," CSEF Working Papers 467, Centre for Studies in Economics and Finance (CSEF), 2017

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

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

[11] 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

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

[13] Gali, J. Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework, Princeton University Press, 2008

[14] 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

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

[16] Hansen, L. and Sargent, T. Exact linear rational expectations models: Specification and estimation. Federal Reserve Bank of Minneapolis Staff Report No. 71, 1981

[17] Hendry, D. Dynamic econometrics. Oxford University Press on Demand, 1995.

[18] 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

[19] Ito, T. and Quah, D. "Hypothesis Testing with Restricted Spectral Density Matrices, with an Application to Uncovered Interest Parity," International Economic Review 30(1), 203-215, 1989.

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

[21] Jorgenson D., "Rational Distributed Lag Functions," Econometrica, 34 134-49, 1966

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

[23] 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

[24] Mankiw G. and R. Reis, "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," The Quarterly Journal of Economics, Oxford University Press, vol. 117(4), pages 1295-1328, 2002

[25] Mertens K. and M. Ravn. "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," American Economic Review, American Economic Association, vol. 103(4), pages 1212-1247, June 2013

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

[27] Plagborg-Moller, P. "Bayesian Inference on Structural Impulse Response Functions," Working Paper, 2017

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

[29] 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

[30] 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

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

[32] Sims, C., “Macroeconomics and Reality,” Econometrica 48(1), 1–48, 1980

[33] Stock, J., What’s New in Econometrics: Time Series, Lecture 7. Short course lectures, NBER Summer Institute, at http://www.nber.org/minicourse_2008.html, 2008

[34] Tenreyro, S. and G. Thwaites.”Pushing on a String: US Monetary Policy Is Less Powerful in Recessions.” American Economic Journal: Macroeconomics, 8(4): 43-74, 2016

[35] 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

[36] 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

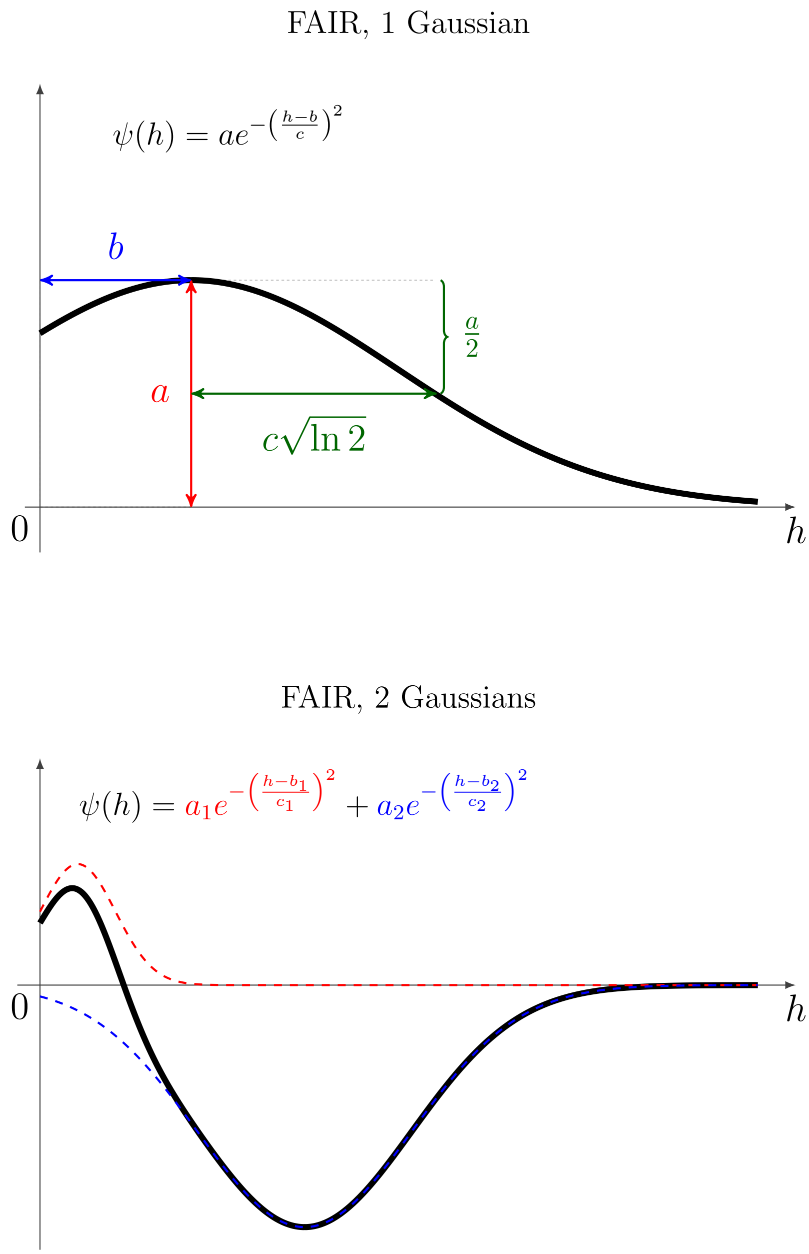

Figure 1: **Functional Approximation of Impulse Responses (FAIR)** with one Gaussian basis function (top panel) or two Gaussian basis functions (bottom panel).

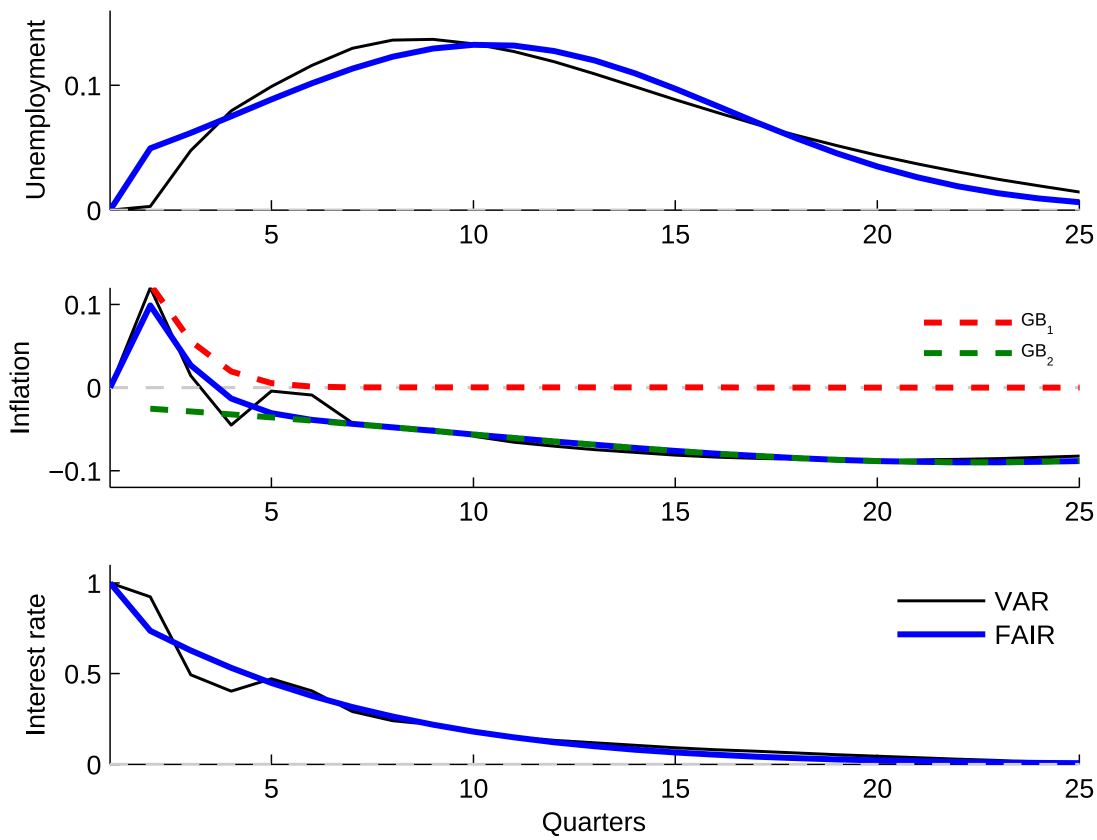

Figure 2: **Approximating IRFs with Gaussian basis functions**: IRFs of the unemployment rate (in ppt), annualized PCE inflation (in ppt) and the federal funds rate (in ppt) to a 100bp monetary shock, as estimated from a VAR or approximated using one Gaussian basis function (top and bottom panels) or two Gaussian basis functions (middle panel). The two basis functions in the middle panel (dashed-green and dashed-red lines) are appropriately weighted so that their sum gives the functional approximation of the impulse response function (thick blue line). Estimation using data covering 1959-2007.

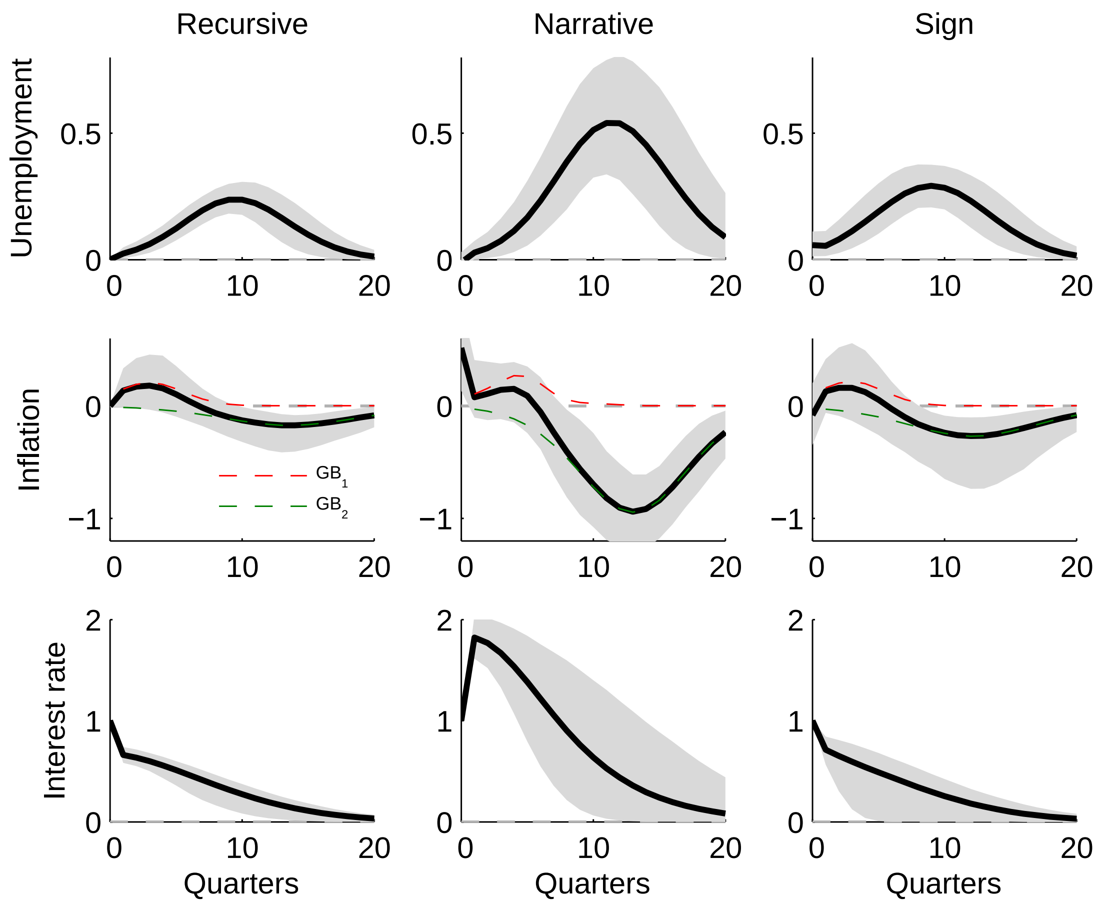

Figure 3: **FAIR Impulse Response Functions**: IRFs of the unemployment rate (in ppt), inflation (in ppt) and the federal funds rate (in ppt) to a 100bp monetary shock identified from a recursive ordering (left column), a narrative approach (middle column), and sign-restrictions (right column). Shaded bands denote the 5th and 95th posterior percentiles. Sample 1969-2007.

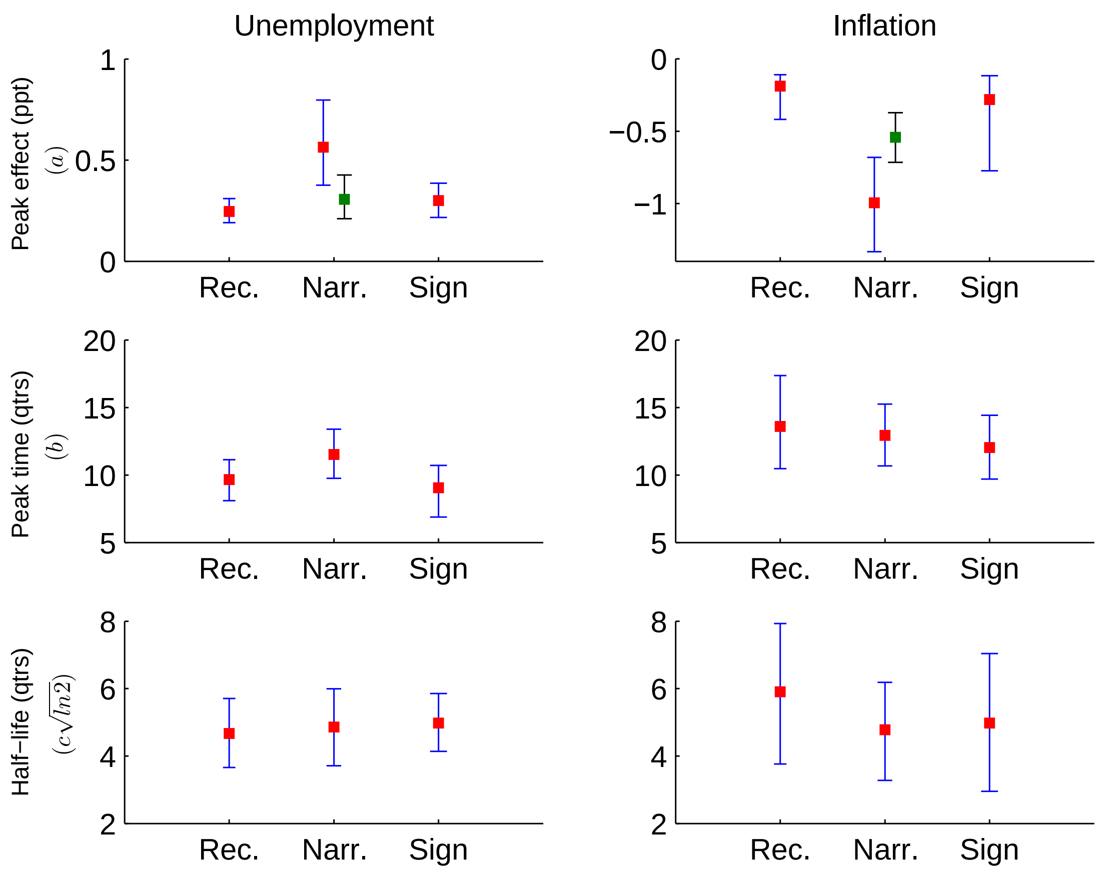

Figure 4: **A FAIR summary of the effects of monetary shocks**: 90th posterior range (blue error bars) of the a-b-c parameters for the IRFs of unemployment and inflation to a 100 bp monetary shocks identified from a recursive ordering ("Rec."), a narrative approach ("Narr."), and sign-restrictions ("Sign"). The red square marks the median value. The black error bars (with green square mark) denote the a parameter estimates rescaled by the peak response of the fed funds rate. For the recursive and sign identification schemes, the blue and black error bars coincide. Sample 1969-2007.

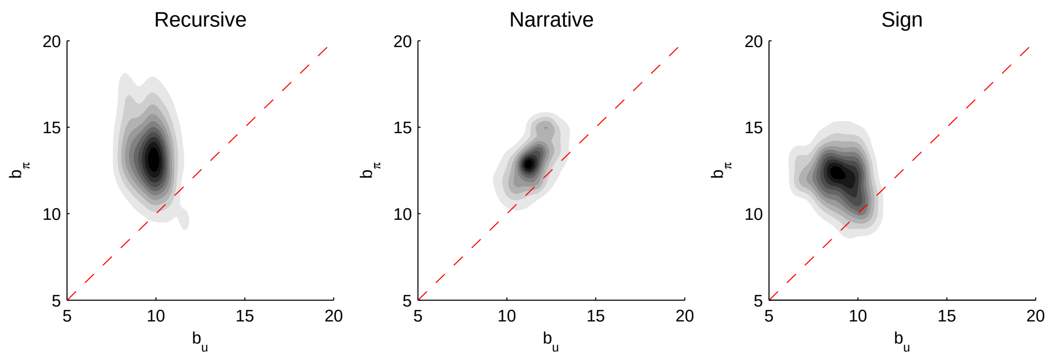

Figure 5: **FAIR estimates of the time to peak effect**: 90th joint posterior density of b_u (x-axis) and b_π (y-axis), the time to peak effects of respectively unemployment and inflation following a monetary shock identified from a recursive ordering, a narrative approach, and sign-restrictions. The 45° dashed-red line denotes identical times to peak response. Sample 1969-2007.

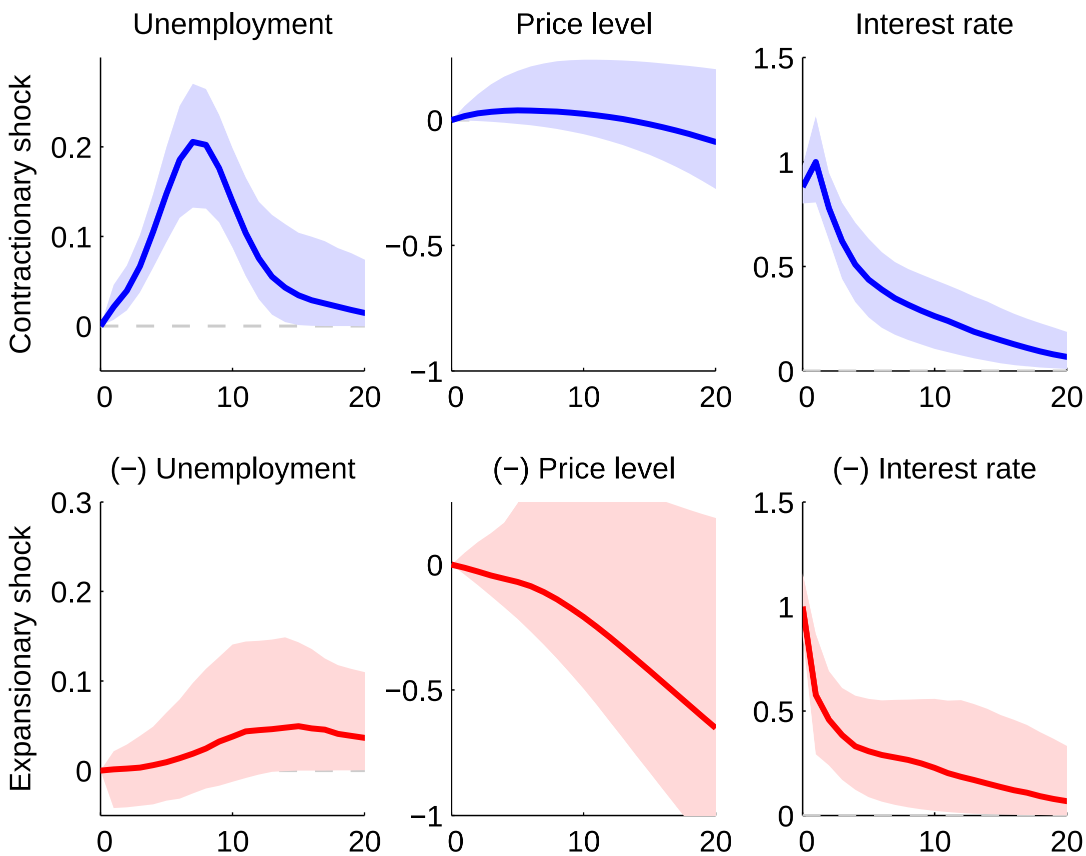

Figure 6: **Asymmetric IRFs, recursive identification**: FAIR estimates of the IRFs of unemployment (in ppt), the (log) price level (in percent) and the federal funds rate (in ppt) to a 100bp monetary shock identified from a recursive ordering. Shaded bands denote the 5th and 95th posterior percentiles. For ease of comparison, responses to the expansionary shock are multiplied by -1. Sample 1959-2007.

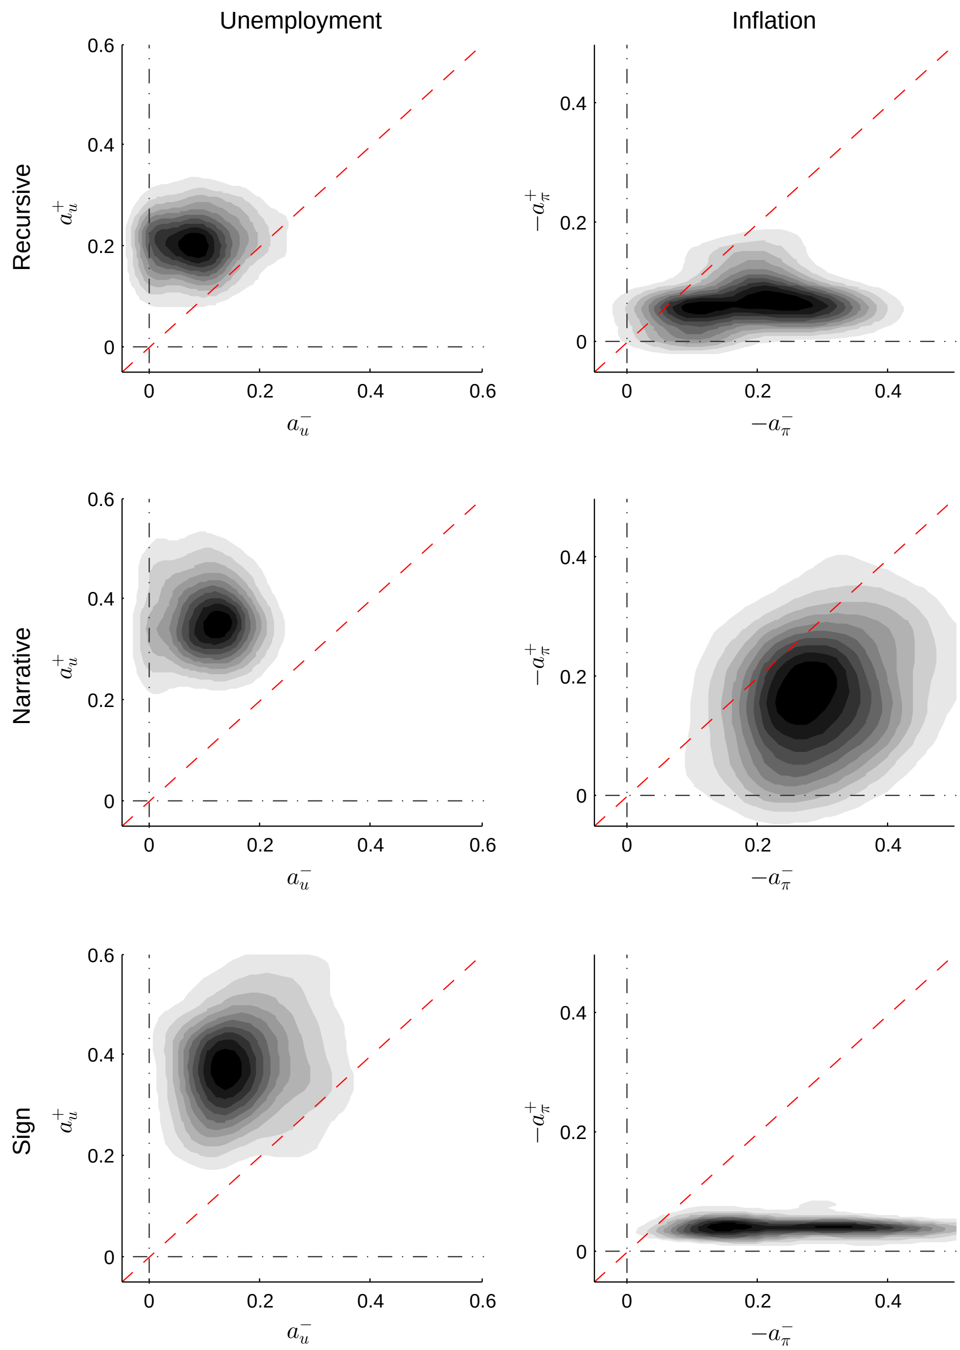

Figure 7: **A FAIR summary of the asymmetric effects of monetary shocks**: posterior distribution of the peak responses of unemployment (a_u, left panel) and negative inflation ($-a_\pi$, right panel) to a 100 bp monetary shock. a^+ denotes the peak response to a contractionary shock (a +100bp shock to the fed funds rate) and a^- denotes the peak response to an expansionary shock (a −100bp shock). The dashed red line denotes symmetric peak responses. Results from a recursive identification scheme over 1959-2007 ("Recursive", top row), a narrative identification scheme over 1966-2007 ("Narrative", middle row), and a set identification scheme with sign restrictions over 1959-2007 ("Sign", bottom row).