



BANK OF ENGLAND

Working Paper No. 401

Changes in the transmission of monetary policy: evidence from a time-varying factor-augmented VAR

Christiane Baumeister, Philip Liu and Haroon Mumtaz

October 2010



BANK OF ENGLAND

Working Paper No. 401

Changes in the transmission of monetary policy: evidence from a time-varying factor-augmented VAR

Christiane Baumeister,⁽¹⁾ Philip Liu⁽²⁾ and Haroon Mumtaz⁽³⁾

Abstract

This paper re-examines the evolution of the US monetary transmission mechanism using an empirical framework that incorporates substantially more information than the standard tri-variate VAR model used in most previous studies. In particular, we employ an extended version of a factor-augmented VAR, where we introduce time variation in the coefficients and stochastic volatilities in the variances of the shocks. Our formulation has two substantive advantages over earlier work: (i) the additional information summarised by the common factors that are extracted from a large panel of aggregate and disaggregate variables improves the identification of the monetary policy shocks since the factors capture more accurately the amount of information analysed by the monetary authority, (ii) we are able to estimate the time-varying effects of monetary policy surprises on macroeconomic aggregates and disaggregate prices and quantities of personal consumption expenditures. Our main results indicate that time variation is a dominant feature of key macroeconomic variables and their components. In analysing the temporal evolution of disaggregate dynamics, we uncover a considerable amount of heterogeneity in sectoral price responses which suggests that monetary policy actions exert an important, and potentially long-lasting, influence on relative prices in the US economy.

Key words: FAVAR, time-varying parameters, monetary transmission.

JEL classification: E30, E32.

(1) Bank of Canada. Email: cbaumeister@bank-banque-canada.ca

(2) International Monetary Fund. Email: pliu@imf.org

(3) Bank of England. Email: haroon.mumtaz@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England, the Bank of Canada and the International Monetary Fund. This paper has benefitted from helpful comments by Luca Benati, Özer Karagedikli, Lutz Kilian, Thomas Laubach, Gert Peersman, Simon Price, Paulet Sadler, Paolo Surico, Robert Tetlow, Sylvia Kaufmann, an anonymous referee and participants at the CEF 2009 conference, Banca d'Italia's second international conference in memory of Carlo Giannini and the Bundesbank seminar. This paper was finalised on 12 July 2010.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/publications/workingpapers/index.htm

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email mapublications@bankofengland.co.uk

Contents

Summary	3
1 Introduction	5
2 Methodology	8
2.1 Why factor-augmented VARs?	8
2.2 A time-varying FAVAR model of the US economy	10
2.3 Estimation	14
2.4 Data	15
2.5 Computation of impulse response functions	16
3 Results	17
3.1 Results using simulated data	17
3.2 Impulse responses to a monetary policy shock	19
3.3 Disaggregate price and quantity responses	22
4 Conclusions	31
Appendix A: Priors and estimation	33
Prior distributions and starting values	33
Simulating the posterior distributions	35
Appendix B: Results based on simulated data	38
Charts	39
References	48



Summary

Several recent studies have documented that the volatility of output and inflation in the United States showed a remarkable decline after the mid-1980s in common with the experience in many countries. In addition, there is evidence to suggest that the persistence of inflation also fell after this date. A growing empirical literature has examined this apparent change in the dynamics of the US economy. These papers usually employ empirical models that contain a limited amount of macroeconomic variables – typically using systems of equations known as vector autoregressions (VARs): a set of equations where the explanatory variables in each equation are the complete set of lagged variables in the system. GDP growth, inflation and the nominal interest rate are the typical variables included in simple VARs that describe the transmission mechanism of monetary policy. If, in reality, the central bank examines a wider set of variables when setting policy, estimates of the monetary policy shock derived from these small empirical models may be biased – ie not completely disentangled from non-policy shocks. As a consequence an accurate assessment of structural shifts may be hampered.

This paper therefore explores the dynamics of the US macroeconomy using a VAR model that incorporates a larger amount of economic information than a tri-variate model. In particular, we use an extended version of the ‘factor-augmented VAR’ (FAVAR) model recently proposed in the literature. The idea behind the FAVAR model is that the bias created by the difference in the information set of the researcher and the agents described in the model can be alleviated by augmenting the standard VAR with *common factors* that are extracted from a large set of macroeconomic indicators. These common factors summarise the relevant information in the macroeconomic indicators and therefore provide a proxy for the information set of agents in the model.

Our FAVAR model for the United States contains common factors extracted from data on real activity, inflation, money and credit and asset prices in addition to a short-term nominal interest rate. The innovation in our work is that we also allow the coefficients of the model and the variances of the shocks to vary over time. When this model is estimated on artificial data, it provides robust inference on changes in impulse response functions suggesting that the model is well suited to the task at hand.



The model is estimated over the period 1960 Q1 to 2008 Q3 (largely predating the recent recession). Our main results suggest that time variation is indeed a pervasive feature of key macroeconomic variables like output measures, price indices, money aggregates and asset prices. In this respect, we find important differences in the responses obtained from our FAVAR specification compared to low-dimensional systems. More specifically, in our data-rich environment we find that economic activity declines by less in more recent times after a restrictive monetary policy shock, whereas no time variation is detected in small-scale VARs.

We find no evidence of a ‘price puzzle’ (the common and counterintuitive finding that prices rise after a monetary contraction) for any of the aggregate price measures throughout the sample period. This may suggest that the extra information captured by the factors leads to more robust structural estimates in that it mimics the central bank’s practice of examining and reacting to a wide variety of data series. However at the disaggregate level, a considerable portion of sectoral price responses displays a significant price puzzle at short horizons during the 1970s which ameliorates from the early 1980s onwards. Our evidence therefore provides a case for the price puzzle not being a puzzle at disaggregate level, but rather a distinctive feature of sectoral dynamics. This should allow us to infer something about the price-setting behaviour of firms in reaction to monetary surprises.

Our results suggest that durable goods are most sensitive to interest rate innovations and show a considerable fall in consumption volumes and a decline in the price level. Durable goods also contribute the least to the dispersion of sectoral prices since individual impulse responses are closely aligned. Instead, non-durable goods and to some extent services account for a large share of cross-sectional heterogeneity, with price responses widely dispersed, covering a broad range of positive and negative values.

1 Introduction

In formulating policy decisions, central banks not only rely on information about the aggregate economy but also carefully monitor sectoral conditions by conducting business surveys that provide important information about the price-setting process of firms for example (Blinder (1991)). For monetary authorities it is crucial to know how their monetary actions affect the pricing decisions of firms since this determines the effectiveness of monetary policy in stabilising the economy. Over the past decades, macroeconomic developments such as increased monetary policy credibility, enhanced competition due to globalisation and technological advances might have contributed to alter the price-setting behaviour across sectors which ultimately changes the way monetary policy is transmitted to the economy as a whole.

An important empirical feature since the mid-1980s is that the volatility of output and inflation has declined considerably in the United States creating a more stable macroeconomic environment. In addition, the level and persistence of aggregate inflation have reached historical lows. Evidence supporting these changes can be found in a number of recent papers including Kim and Nelson (1999b), McConnell and Pérez-Quirós (2000), Cogley and Sargent (2005) and Benati and Mumtaz (2007). However, issues related to the causes and consequences of these changes have been more controversial. For example, the results by Cogley and Sargent (2002) and Clarida *et al* (2000) lend support to the idea that the change in US macroeconomic dynamics was linked to a change in the practice of monetary policy. In contrast, the evidence on US policy activism reported in Cogley and Sargent (2005) is less clear cut than the authors' earlier work. Similarly, results reported in Primiceri (2005) and Sims and Zha (2006) are more sympathetic to the idea that an absence of adverse non-policy shocks was the main driving force.

A strand of this literature has focused on the possibility of changes in the transmission of monetary policy shocks. Boivin and Giannoni (2006) estimate the responses of output and inflation to a monetary policy shock in the United States using a vector autoregression (VAR) model estimated on two subsamples: 1959-79 and 1980-2002. Their results suggest that the responses of output and inflation are smaller in the latter period. However, these results are at odds with those obtained by Primiceri (2005) and Sims and Zha (2006) using a more sophisticated approach to characterise time variation in the VAR parameters. All these papers find no significant change in the responses of inflation and output (or unemployment) across the

sample.

Most of these studies use small-scale VAR models extended to allow for time variation in VAR parameters and/or structural shocks. This methodology is undoubtedly powerful. However, one potential problem is the fact that the amount of information incorporated in these models is often quite limited. A canonical VAR model consists of three variables – a short-term interest rate, output growth and inflation. Although many studies have extended the canonical set-up to include other variables of interests, limited sample size and degrees of freedom problems still restricts the VAR specification to a limited number of variables. This feature has two potential consequences. First, missing variables could lead to biases in the reduced-form VAR coefficients. This may imply that reduced-form estimates of persistence and volatility are biased. Second, the omission of some variables could hinder the correct identification of structural shocks. One possible manifestation of these problems are impulse response functions that are at odds with economic theory. A number of recent studies have raised these points. For instance, Bernanke *et al* (2005) argue that if the information set used by the econometrician is smaller than that employed by the monetary authority, then structural shocks and their responses may be mis-measured because the empirical model excludes some variables that the central bank systematically responds to. Similarly, Castelnuovo and Surico (2009) and Benati and Surico (2009) building on Lubik and Schorfheide (2004), argue that during periods of indeterminacy, the dynamics of the economy are characterised by a latent variable. Therefore, (reduced-form and structural) estimates of the VAR model may be biased when estimation is carried out over these periods.

The purpose of this paper is to re-examine the evolution of the US monetary transmission mechanism using an empirical framework that incorporates substantially more information than the canonical three-variable model used in most previous studies. In particular, we employ an extended version of the factor-augmented VAR introduced in Bernanke *et al* (2005). This model includes information from a large number of macroeconomic indicators representing various dimensions of the economy. Our extensions include allowing for time variation in the coefficients and stochastic volatility in the variances of the shocks.

This formulation has two clear advantages over previous studies: (i) we identify the monetary policy shock using a model that incorporates around 600 macroeconomic and financial variables, hence making it less likely that our model suffers from the shortcomings discussed above, (ii) our



model allows us to estimate time-varying impulse responses for each of the variables contained in our panel. Therefore, we are able to derive results for the variation in responses of a wide variety of variables to a monetary policy surprise. In particular, this paper not only provides evidence on the possible change in responses of the main macroeconomic variables, but also on the time-varying responses of components of the consumption deflator and disaggregate consumption quantities. This latter feature is particularly relevant for the conduct of monetary policy since price-setting behaviour of firms plays a crucial role in the monetary transmission mechanism. Knowing which types of goods are more sensitive to monetary policy actions may improve our understanding of how monetary policy disturbances are propagated, enhancing policymakers' ability to stabilise the economy.

The main contribution of the paper is to analyse the temporal evolution of disaggregate dynamics, often hidden by aggregate measures, in order to inform policymakers about changes in the relative price effects of monetary actions. On the one hand, Lastrapes (2006) and Balke and Wynne (2007) demonstrate that money supply shocks have long-run effects on the distribution of relative commodity prices implying an important degree of monetary non-neutrality. On the other hand, Boivin *et al* (2009), in a recent empirical contribution, made the case that discrepancies between aggregate and sectoral measures of inflation derive from the fact that the bulk of fluctuations in individual prices is driven by sector-specific factors and that monetary shocks are of minor importance but induce sluggishness in price adjustment.¹

Our main results suggest that time variation is indeed a pervasive feature of key macroeconomic variables like output measures, price indices, money aggregates and asset prices. In this respect, we find important differences in the responses obtained from our FAVAR specification compared to low-dimensional systems. More specifically, in our data-rich environment we find that economic activity declines by less in more recent times after a restrictive monetary policy shock, whereas no time variation is detected in small-scale VARs. The latter specification also displays a substantial and persistent price puzzle which is absent in the FAVAR framework for all aggregate inflation measures throughout the sample.

Another salient aspect is that the propagation mechanism of monetary disturbances appears

¹ See also Altissimo *et al* (2009), Bils and Klenow (2004) and Clark (2006) for differences in inflation dynamics at disaggregate and aggregate level.

highly heterogeneous across components of personal consumption expenditures, suggesting that monetary policy actions exert an important influence on relative prices in the US economy. This heterogeneity across sectors might shed some light on the channels through which the transmission of monetary impulses occurs. We provide some evidence that at the disaggregate level the cost channel of monetary transmission seems to be active for several product categories.

The rest of the paper is organised as follows. The next section introduces the empirical methodology adopted in this study, outlines the estimation procedure and describes our data set. Section 3 presents and interprets the time-varying dynamics of selected macroeconomic aggregates and disaggregate prices and quantities in response to monetary policy shocks and discusses the implications for macroeconomic modelling and the conduct of monetary policy. Section 4 offers some concluding remarks.

2 Methodology

2.1 Why factor-augmented VARs?

Consider the following simple backward-looking model of the economy:

$$\pi_t = \beta\pi_{t-1} + \chi(y_{t-1} - y_{t-1}^*) + s_t \quad (1)$$

$$y_t = \alpha y_{t-1} + \varpi(R_{t-1} - \pi_{t-1}) + d_t \quad (2)$$

where the Phillips curve in equation (1) relates inflation (π_t) to the deviation of output (y_t) from potential (y^*) and a supply shock s_t . Equation (2) is a standard IS curve that describes the relationship between output and the real interest rate ($R_{t-1} - \pi_{t-1}$) and a demand shock d_t . Finally, the monetary authority sets interest rates according to a standard Taylor rule:

$$R_t = B\pi_{t-1} + \lambda(y_{t-1} - y_{t-1}^*) + v_t \quad (3)$$

where v_t is the monetary policy shock.

Bernanke *et al* (2005) argue that assumptions made about the information structure are crucial when deciding whether a standard VAR can describe such a model. In particular, if it is assumed that the specific data series included in the VAR correspond exactly to the model variables and

are observed by the central bank *and* the econometrician, then the VAR model provides an adequate description of the theoretical model. However, both these assumptions are difficult to justify. First, measurement error implies that measures of inflation and output are less than perfect proxies for model variables. Of course, this problem is much more acute for unobserved variables such as potential output. Furthermore, for broad concepts like economic activity and inflation there exists a multitude of observable indicators none of which will be able to match the theoretical construct precisely. Second, it is highly likely that the researcher only includes a subset of the variables monitored by the monetary authority in their models.

Measurement error and omitted variables can potentially affect VAR analysis of possible changes in the transmission of structural shocks. A crucial premise is that the structural shocks are identified correctly and the propagation mechanism of these shocks is estimated accurately. Both these assumptions are less likely to hold if important information is excluded from the VAR.

The obvious solution to this problem is to try and include more variables in the VAR. However, the degrees of freedom constraint becomes binding quite quickly in standard macro time-series data sets.² Bernanke *et al* (2005) suggest a practical solution. They propose a ‘Factor-Augmented’ VAR (FAVAR) model, where factors from a large cross-section of economic indicators are included as extra endogenous variables in a VAR. More formally, let $X_{i,t}$ be a $T \times N$ matrix of economic indicators in the central bank’s information set and let $Y_{j,t}$ denote a $T \times M$ matrix of observable economic variables assumed to have pervasive effects throughout the economy, then the FAVAR model can be written as:³

$$X_{i,t} = \Lambda F_t + \Psi Y_{j,t} + e_{i,t}, \quad (4)$$

$$\begin{pmatrix} F_t \\ Y_{j,t} \end{pmatrix} = \Theta \begin{pmatrix} F_{t-1} \\ Y_{j,t-1} \end{pmatrix} + v_t,$$

where $i = 1, 2 \dots N, j = 1, 2 \dots M,$

$$E(e'_{i,t} e_{i,t}) = R \quad (5)$$

$$E(v'_t v_t) = \Sigma$$

$$E(e'_{i,t} v_t) = 0$$

²This problem is even more acute in time-varying VARs as they usually impose a stability constraint (at each point in time) and this is less likely to be satisfied as the number of variables in the VAR increases.

³Ultimately, we are interested in uncovering the structural relationships among these variables.

and F_t is $T \times J$ matrix of common factors, Λ is an $N \times J$ matrix of factor loadings and Ψ is a $N \times M$ matrix of coefficients that relate $X_{i,t}$ to $Y_{i,t}$.

The first expression in (4) is the observation equation of the system and describes how the observed series are linked to the unobserved factors. The second expression (the transition equation) is a VAR(L) in F_t and Y_t (with a $((J + M) \times L) \times ((J + M) \times L)$ coefficient matrix Θ) and is used to describe the dynamics of the economy.

Two identification issues need to be dealt with in this extended VAR model. First, in order to identify the factors, restrictions need to be placed on either the observation or the transition equation. Bernanke *et al* (2005) leave the transition equation unrestricted and impose normalisation restrictions on the factor loadings. In particular, the top $J \times J$ block of Λ is assumed to be an identity matrix and the top $J \times M$ block of Ψ is assumed to be zero.⁴

The second identification issue concerns the identification of shocks to the transition equation. As in the standard VAR literature, this is carried out by imposing restrictions on the covariances of the VAR innovations, Σ , or by restricting the sign of the impulse response functions. Once the structural shocks are identified, impulse response functions can be constructed not only for F_t and $Y_{i,t}$ but for all the variables contained in $X_{i,t}$.

2.2 A time-varying FAVAR model of the US economy

Our FAVAR model for the US economy is closely related to the FAVAR model described above. There are, however, two crucial differences. First, we allow the dynamics of the system to be time-varying to capture changes in the propagation of structural shocks as a result of shifts in private sector behaviour and/or monetary policy preferences. Second, our specification incorporates heteroscedastic shocks which account for variations in the volatility of the underlying series.

⁴This normalisation solves the rotational indeterminacy problem inherent in dynamic factor models by ruling out linear combinations that lead to observationally equivalent models.

Consider first the observation equation:

$$\begin{pmatrix} X_{1,t} \\ \cdot \\ \cdot \\ X_{N,t} \\ R_t \end{pmatrix} = \begin{pmatrix} \Lambda^{11} & \Lambda^{21} & \Lambda^{31} & \Psi^{11} \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \Lambda^{NN} & \Psi^{1N} \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} F_t^1 \\ F_t^2 \\ F_t^3 \\ R_t \end{pmatrix} + \begin{pmatrix} e_{1,t} \\ \cdot \\ \cdot \\ e_{N,t} \\ 0 \end{pmatrix} \quad (6)$$

$X_{i,t}$ is a panel of variables that contains a large amount of information about the current state of the US economy along several dimensions which will be detailed in the data section below.

F_t^1 , F_t^2 and F_t^3 denote the latent factors which summarise the comovement among the underlying series at each date. In fact, we postulate that these three common factors capture the dynamics of the US economy.⁵ As in Bernanke *et al* (2005), we assume that the Federal funds rate R_t is the ‘observed factor’.

Λ and Ψ are the elements of the factor loading matrix. The structure of the loading matrix implies two things. First, some of the variables are allowed to have a contemporaneous relationship with the nominal interest rate, ie $\Psi \neq 0$ for data series that are expected to react promptly to monetary policy actions.⁶ Second, in contrast to Belviso and Milani (2006), we do not assign a structural interpretation to the factors, ie we do not impose that a factor only loads on a certain subset of data series that belong to a specific economic concept; instead, the dynamics of the variables included in $X_{i,t}$ are determined by a linear combination of all common factors. Since the aim of our study is to investigate possible heterogeneity in the reactions of individual prices and quantities across sectors, it would be unduly restrictive to force a proportionality constraint with a single factor upon the dynamics of disaggregate series.

As we describe below, time variation is introduced into the model by allowing for drift in the coefficients and the error covariance matrix of the transition equation. Note that an alternative way of modelling time variation is to allow the factor loadings (Λ and Ψ) to drift over time.⁷

⁵Stock and Watson (2002) and Bernanke *et al* (2005) have shown that a few factors are sufficient to summarise the common sources of variation in economic time series. The choice to set $k = 3$ was also motivated by the fact that estimation of the time-varying VAR becomes harder as the number of endogenous variables increases, as discussed by Del Negro (2003).

⁶Accounting for the contemporaneous relation between fast-moving variables and the interest rate directly in the observation equation amounts to removing the component of the factors that is contemporaneously affected by the funds rate. A classification of the data series according to their speed of adjustment to interest rate movements can be found in the data appendix.

⁷See Del Negro and Otrok (2008) for this kind of approach in a different context.

There are, however, two reasons why we do not adopt this alternative model. First, such a model implies that any time variation in the dynamics of each factor *and* the volatility of shocks to each factor is driven entirely by the drift in the associated factor loading. This assumption is quite restrictive, especially as it only allows changes in the mean and persistence of each factor to occur simultaneously with changes in the volatility of the shocks.⁸ Second, this model implies a much larger computational burden as the Kalman filter and smoother have to be employed for each underlying series. However, apart from the computational costs, this specification implies that the central bank will always react in the same way to the ‘state of the economy’ as captured by the latent factors which is difficult to justify given our sample period 1971 Q1 to 2008 Q3. Equally, allowing for time variation in both, the factor loadings and the coefficients of the transition equation, would entail serious identification problems since there would be three time-varying unobserved components, ie $\Gamma_t = [\Lambda_t, \Psi_t], \theta_t$ and F_t . However, substituting the transition equation (7) into the observation equation (6) imparts a restricted form of time variation also in the factor loadings. This interaction between the loadings and the time-varying coefficients of the factors should generate rich dynamics for the impulse response functions of the underlying series.

The transition equation of the system is a time-varying VAR model of the following form:

$$Z_t = \delta_t + \sum_{l=1}^L \phi_{l,t} Z_{t-l} + v_t \quad (7)$$

where $Z_t = \{F_t^1, F_t^2, F_t^3, R_t\}$ and L is fixed at 2.

Following Cogley and Sargent (2005) and Primiceri (2005) among others, we postulate the following law of motion for the coefficients $\theta_t = [\delta_t \ \phi_{l,t}]'$:

$$\theta_t = \theta_{t-1} + \eta_t \quad (8)$$

The time-varying covariance matrix of the VAR innovations v_t can be factored as

$$VAR(v_t) \equiv \Sigma_t = A_t^{-1} H_t (A_t^{-1})' \quad (9)$$

H_t is a diagonal matrix which contains the stochastic volatilities and A_t is a lower triangular

⁸This model implies that the dynamics of the observed factor are time invariant. In addition, the impact of the observed factor on the other variables in the transition equation is also assumed to be constant over time. Again, these assumptions are rather restrictive in a model designed to investigate the changing impact of monetary policy.

matrix that models the contemporaneous interactions among the endogenous variables:

$$H_t \equiv \begin{bmatrix} h_{1,t} & 0 & 0 & 0 \\ 0 & h_{2,t} & 0 & 0 \\ 0 & 0 & h_{3,t} & 0 \\ 0 & 0 & 0 & h_{4,t} \end{bmatrix} \quad A_t \equiv \begin{bmatrix} 1 & 0 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 & 0 \\ \alpha_{41,t} & \alpha_{42,t} & \alpha_{43,t} & 1 \end{bmatrix} \quad (10)$$

with the $h_{i,t}$ evolving as geometric random walks,

$$\ln h_{i,t} = \ln h_{i,t-1} + \varepsilon_t \quad (11)$$

Along the lines of Primiceri (2005), we postulate the non-zero and non-one elements of the matrix A_t to evolve as driftless random walks,

$$\alpha_t = \alpha_{t-1} + \tau_t, \quad (12)$$

and we assume the vector $[e'_t, v'_t, \eta'_t, \tau'_t, \varepsilon'_t]'$ to be distributed as

$$\begin{bmatrix} e_t \\ v_t \\ \eta_t \\ \tau_t \\ \varepsilon_t \end{bmatrix} \sim N(0, V), \text{ with } V = \begin{bmatrix} R & 0 & 0 & 0 & 0 \\ 0 & \Sigma_t & 0 & 0 & 0 \\ 0 & 0 & Q & 0 & 0 \\ 0 & 0 & 0 & S & 0 \\ 0 & 0 & 0 & 0 & G \end{bmatrix} \text{ and } G = \begin{bmatrix} \sigma_1^2 & 0 & 0 & 0 \\ 0 & \sigma_2^2 & 0 & 0 \\ 0 & 0 & \sigma_3^2 & 0 \\ 0 & 0 & 0 & \sigma_4^2 \end{bmatrix} \quad (13)$$

The model described by equations (6) to (13) incorporates a large amount of information about the US economy. In particular, if the factors in equation (6) contain relevant information not captured by a three-variable VAR used in studies such as Primiceri (2005), then one might expect policy shocks identified within the current framework to be more robust. Our flexible specification for the transition equation implies that the model accounts for the possibility of structural breaks in the dynamics that characterise the economy and allows the monetary authority to continuously update its knowledge about the macroeconomic environment.

2.2.1 Identification of monetary policy shocks

As in Bernanke *et al* (2005) we place the interest rate last in the transition equation (7) and use this recursive ordering to identify the monetary policy shock as the only shock that does not affect the latent factors in the system within the quarter. Note that the structure of observation

equation (6) implies that any contemporaneous correlation between ‘fast-moving’ variables (eg asset prices) and the policy rate is already accounted for by the factor loadings Ψ . In other words, this amounts to removing the component of the factors that are contemporaneously affected by the funds rate. This implies that a recursive ordering (with the interest rate placed last) can be employed in the transition equation.⁹

2.3 Estimation

The model described by equations (6) to (13) is estimated using the Bayesian methods described in Kim and Nelson (1999a). In particular, we employ a Gibbs sampling algorithm that approximates the joint posterior distribution. The algorithm exploits the fact that given observations on Z_t the model is a standard time-varying parameter model.

A detailed description of the prior distributions and the sampling method is given in Appendix A. Here we summarise the basic algorithm which involves the following steps:

1. Given initial values for the factors, simulate the VAR parameters and hyperparameters.
 - The VAR coefficients θ_t and the off-diagonal elements of the covariance matrix α_t are simulated by using the methods described in Carter and Kohn (2004).
 - The volatilities of the reduced-form shocks H_t are drawn using the date-by-date blocking scheme introduced by Jacquier *et al* (1994).
 - The hyperparameters Q and S are drawn from an inverse Wishart distribution, while the elements of G are simulated from an inverse gamma distribution.
2. Given starting values for the factors, draw the factor loadings (Λ and Ψ) and the covariance matrix R .
 - Given data on Z_t and $X_{i,t}$, standard results for regression models can be used and the coefficients and the variances are simulated from a normal and inverse gamma distribution.
3. Simulate the factors conditional on all the other parameters.

⁹A classification of the data series according to their speed of adjustment to interest rate movements can be found in the data appendix.

- This is done in a straightforward way by employing the methods described in Bernanke *et al* (2005) and Kim and Nelson (1999a).

4. Go to step 1.

We use 20,000 Gibbs sampling replications and discard the first 19,000 as burn-in. To assess convergence we compare posterior moments computed using different subsets of the retained draws. The results of this exercise (which are available upon request) show little variation across the retained draws providing some evidence of convergence to the ergodic distribution.

2.4 Data

The data set consists of a balanced panel of quarterly observations on 138 US macroeconomic and financial time series spanning the period from 1960 Q1 to 2008 Q3¹⁰ which cover a broad range of measures of real activity and income, employment, asset prices, interest rates and spreads, exchange rates, price indices and money aggregates. We also include a set of forward-looking variables like consumer expectations, commodity prices, orders and inventories that should capture signals of the future course of the economy as well as inflationary pressures to which the monetary authority might react pre-emptively.

All the series have been obtained from the Bureau of Economic Analysis (BEA), the Bureau of Labor Statistics, the US Bureau of the Census and the FRED database. They provide a comprehensive description of the state of the economy containing indicators that are commonly analysed by a central bank in the monetary policy decision process.¹¹ This macroeconomic information set has been augmented by a large panel of disaggregate price and quantity series for a wide range of consumer expenditure categories obtained from the National Income and Product Accounts published by the BEA. We collect data at the highest level of disaggregation and only if observations for one category were missing for the time span we consider, we moved to the next level of aggregation and hence, excluded the sub-categories to avoid double-counting. The

¹⁰ However, the first ten years are used as a training sample to calibrate our priors.

¹¹ Ideally, the assessment of central bank behaviour would rely only on information that was available at the time of policymaking ie real-time data as opposed to fully revised ones. However, Bernanke and Boivin (2003) provide compelling evidence that this distinction makes little difference given the latent nature of the factors; what matters most, is the variety of data included in the information set.

remaining sectoral price and volume series are the same as in Boivin *et al* (2009) amounting to 190 disaggregated deflator series for personal consumption expenditures (PCE) and the corresponding series on real consumption to which we add price and quantity data for overall PCE, durable and non-durable goods, and services.

In total, our data set includes the effective Federal funds rate as the monetary policy instrument and 592 aggregate and disaggregate time series from which we extract the common factors. Data that are available on a monthly basis have been converted to quarterly frequency by taking monthly averages. The variables have been appropriately transformed to induce stationarity and have been demeaned and standardised before estimation.

2.5 Computation of impulse response functions

We calculate the impulse responses Δ_t of F_t^1 , F_t^2 , F_t^3 and R_t to the monetary policy shock for each quarter, where we normalise the shock such that it increases the Federal funds rate by 100 basis points at each date in the sample to make the responses comparable over time. With these in hand, the time-varying impulse responses of each underlying variable can be easily obtained using the observation equation (6) of the model. That is, the impulse responses of $X_{1,t}, \dots, X_{N,t}$ are computed as:

$$\begin{pmatrix} \Lambda^{11} & \Lambda^{21} & \Lambda^{31} & \Psi^{11} \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \Lambda^{NN} & \Psi^{1N} \\ 0 & 0 & 0 & 1 \end{pmatrix} \times \begin{pmatrix} \Delta_t^{F_t^1} \\ \Delta_t^{F_t^2} \\ \Delta_t^{F_t^3} \\ \Delta_t^{R_t} \end{pmatrix} \quad (14)$$

Given the presence of time-varying parameters in the transition equation, computation of impulse response functions has to take into account the possibility of parameter drift over the impulse response horizon. Therefore, following Koop *et al* (1996), we define the impulse response functions at each date as Δ_t

$$\Delta_{t+k} = E(Z_{t+k} | \Xi_{t+k}, \mu_{MP}) - E(Z_{t+k} | \Xi_{t+k}) \quad (15)$$

where Ξ denotes all the parameters and hyperparameters of the VAR and k is the horizon under consideration. Equation (15) states that the impulse response functions are calculated as the difference between two conditional expectations. The first term in equation (15) denotes a

forecast of the endogenous variables conditioned on a monetary policy shock μ_{MP} . The second term is the baseline forecast, ie conditioned on the scenario where the monetary policy shock equals zero. Therefore, in effect, equation (15) integrates out future uncertainty in the VAR parameters. The conditional expectations in (15) are computed via Monte Carlo integration for 500 replications of the Gibbs sampler. Details on the Monte Carlo integration procedure can be found in Koop *et al* (1996).

3 Results

3.1 Results using simulated data

To test the effectiveness of our empirical strategy we consider the following experiment. We generate artificial data using a medium-scale dynamic stochastic general equilibrium (DSGE) model (to be described below) that incorporates a shift in structural parameters. We use the simulated data in two ways. First, we estimate a time-varying VAR using the minimal set of variables necessary to identify a monetary policy shock (inflation, output and the interest rate) and use a recursive identification scheme to calculate the time-varying responses to monetary policy shocks. Second, we use the full set of model variables in the time-varying FAVAR model described above and calculate the time-varying response of output and inflation to policy shocks using a recursive identification scheme. Given the discussion above, one would expect the time-varying FAVAR to outperform the time-varying VAR in recovering the pattern of responses over time as the latter contains less information than the former.

As a data-generating process we use the open economy DSGE model developed in Justiniano and Preston (2009) which is a generalisation of the model in Monacelli (2005). In particular, Justiniano and Preston (2009) extend the model by introducing incomplete asset markets, habit formation and indexation of prices to past inflation. We refer the reader to Justiniano and Preston (2009) for a detailed derivation and description of the model. The model is relatively large and describes the behaviour of ten key macroeconomic variables including consumption, output, inflation, import prices, the real and the nominal exchange rate, the domestic interest rate and a set of foreign variables. The behaviour of the monetary authority is described by a policy rule where the domestic interest rate is set based on output, inflation and exchange rate fluctuations:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) [\lambda_1 \pi_t + \lambda_2 y_t + \lambda_3 \Delta e_t] + \eta_{m,t}$$

where π_t denotes inflation, y_t is the output gap and Δe_t is the change in the nominal exchange rate. We generate data from the model where the model parameters are calibrated at values estimated in Justiniano and Preston (2009) (*the full list of calibrated parameters and a description of the model is available from the authors on request*). However, we assume that the value of λ_1 changes over the sample period in consideration, from 1.8 in the first 125 generated observations to 0.9 in the subsequent 125 observations. This change in behaviour of the monetary authority implies a change in the impulse responses to monetary policy shocks in the model. Our experiment aims to gauge how well a time-varying VAR and a time-varying FAVAR perform in picking up this change in the true impulse responses. Note that in order to solve this regime-switching DSGE model we use the solution algorithm described in Farmer *et al* (2008). Their algorithm incorporates the possibility that model agents form rational expectations about the possibility of changes in regime.¹²

The solution of the switching DSGE model delivers a VAR (in the state variables) with regime-dependent impact and lag coefficients. In each regime we draw the shocks from a normal distribution and generate 500 observations. The first 375 observations are discarded to reduce the influence of initial conditions. The data in the two regimes is concatenated to form the artificial sample. This is repeated 500 times and the average across the 500 replications is used to estimate the two models. Note that computational concerns imply that a full Monte Carlo analysis is not feasible. The required estimation time for the time-varying FAVAR implies that several weeks of computer time is required for this task

The time-varying VAR is estimated using model data on output, inflation and the domestic interest rate. The time-varying FAVAR uses the ten macroeconomic variables described by the model (in particular: consumption, output, inflation, import price inflation, the real exchange rate, the domestic interest rate, terms of trade, foreign output, foreign inflation and foreign interest rates). Both models use the first 50 observations to calibrate the priors (which are identical across models) and therefore the effective sample is 200. In both cases a recursive identification scheme (with the interest rate ordered last) is used to identify the monetary policy shock.

The top panel of the chart in Appendix B plots the model impulse responses to a unitary

¹²In other words, we assume the presence of two regimes in this Markov-switching DSGE model with transition probabilities set equal to 0.95. Note that existence and uniqueness of the solution depends on the values of λ_1 in both regimes and on the transition probabilities.

monetary policy shock in a set of three dimensional graphs. The x-axis of the figures represent each time period in the simulated data. We incorporate a regime shift in the model at the 100th observation, with λ_1 falling from 1.8 to 0.9 after this point. This is reflected in the impulse response functions by an increase in the magnitude of both the output and inflation cumulated responses at this point. The second panel of the chart presents the median estimate of the impulse response to a monetary policy shock (calibrated to increase the interest rate by 100 basis points) from the small-scale time-varying VAR. Two results are immediately apparent: first, while the initial response of both output and inflation is negative, the cumulated response becomes positive at later horizons. Second, the model fails to pick up the shift in shock transmission and the impulse responses display no change across the sample. The results are different when we consider the median impulse response to this shock from our proposed time-varying FAVAR model. In particular, both the output gap and inflation responses are negative and persistent as in the model. Notably, the estimates display a shift around the 100th observation. The output response becomes stronger after this date. The inflation response displays a similar change, although the onset of the shift is more gradual. In general, the FAVAR tracks the changes in the impulse responses quite well. This provides strong evidence that this empirical model is well suited to the task of measuring changes in the transmission mechanism.

3.2 Impulse responses to a monetary policy shock

In this section, we examine the responses of key US macroeconomic variables to a monetary policy shock. Chart 1 displays the estimated (cumulated) impulse responses of selected real activity measures to a monetary policy contraction. The left panels of the chart show the median responses in each quarter. The two middle panels compare the responses in 1975 Q1 and 2008 Q1 as two representative dates of the sample period. The last column considers the relative importance of time variation in impulse responses. In particular, it plots the joint posterior distribution of the cumulated responses at the one-year horizon with values for 1975 Q1 plotted on the x-axis and those for 2008 Q1 on the y-axis. Shifts of the distribution away from the 45-degree line indicate a systematic change across time.

Chart 1 illustrates that a 1% increase in the Federal funds rate reduces the level of real GDP by around 0.5% at a horizon of two years and this magnitude is about half than that estimated in the earlier part of the sample. Comparable results are obtained for the other real activity measures

with the consumption response showing the smallest decline after the mid-1980s. The reaction of gross investment is strongest at all times but gets less sensitive in the more recent past according to the median estimate. These results are in line with Boivin and Giannoni (2006) who report a fall in the impact of policy shocks after 1980. However, our estimates differ from those obtained by Primiceri (2005) who finds no change in the response of economic activity across this sample period using a three-variable time-varying VAR model. The last column of Chart 1 provides evidence that the milder reaction of all real activity indicators in more recent times is a pervasive feature since for all measures at least 75% of the joint distribution lies above the 45-degree line.

Chart 2 plots the time-varying cumulated responses of several inflation indicators. The top row of Chart 2 displays the responses of CPI. A 1% increase in the policy rate reduces the price level by around 0.4% three years after the shock during the 1970s, while it falls by 0.6% in the 2000s. Boivin and Giannoni (2006) instead find that the CPI response dampens in the Great Moderation period. The responses of all price variables set in with a delay of a couple of quarters but then gradually decline before stabilising at a lower level. Due to this initial period of sluggishness, the evidence of systematic shifts in the responses of inflation measures to an exogenous monetary policy contraction at the one-year horizon is mixed. With about 55% of the estimated distribution in the current period larger than in 1975, the change for producer prices is the least significant. The strongest shift occurs for the GDP deflator with the majority of points (around 69%) lying below the 45-degree line. At later horizons (not reported here) the change in responses becomes more significant pointing towards a uniformly larger price decline in the more recent past.

The most interesting result is however, that none of the aggregate price measures exhibits a price puzzle. These results support the analysis of Castelnuovo and Surico (2009) who argue that the price puzzle in structural VARs may be a symptom of omitted variable bias that may arise when the Taylor principle is violated. In particular, they show that when the economy is operating under indeterminacy, an additional unobserved variable characterises the dynamics of the economy. The factors included in our model summarise a large amount of information that may proxy this latent variable. The fact that the price puzzle is absent throughout the sample lends support to this idea. This becomes even more apparent when we compare our FAVAR results for real activity and inflation with estimates obtained from a tri-variate time-varying VAR.

The first column of Chart 3 presents the responses of unemployment, the consumer price level

and the Federal funds rate to a monetary policy shock over time derived from a standard three-variable time-varying VAR that uses the same prior distributions (for VAR coefficients, variances and hyperparameters) as our FAVAR benchmark. The results from this small-scale system are in stark contrast to the FAVAR estimates as can be seen from the snapshots in 1975 Q1 and 2008 Q1 where we have added the median responses generated by our data-rich model. The unemployment response shows no time variation with impulse response functions being essentially identical across the sample period. The response of the price level has an anomalous positive sign throughout the sample with little change in magnitude and persistence over time. This comparison highlights the importance of the additional information contained in the factors and provides some evidence that the identification of the monetary policy shock may be more robust in our time-varying FAVAR framework than in its low-dimensional counterpart.

Chart 4 displays the responses of money and credit measures. The estimates indicate important time variation. The first half of the sample is characterised by an unexpected positive response of M1 to an increase in the funds rate which however turns negative from the mid-1980s onwards. M2 instead declines in response to the monetary policy contraction throughout the sample but significantly so only in the period after the mid-1980s. In contrast to the increasingly larger responses of the money indicators, the median response of consumer credit is more muted in the current period. The fourth column of Chart 4 shows compelling evidence of a systematic increase in the responses of both monetary aggregates to a restrictive policy shock with respectively 88% and 72% of the joint distribution lying below the 45-degree line. Similarly, a considerable share (around 75%) of the distribution of the credit reaction is located above the threshold indicating a significant decline in the response across the representative dates.

Chart 5 shows the responses of selected asset prices to the monetary policy shock. The Dow Jones industrial average falls only slightly in response to a 100 basis point increase in the interest rate at all times and starts rising as soon as the funds rate response reverts back to baseline. The points of the joint distribution are almost equally spread out across the dividing line which indicates no significant change in the stock index reaction over time. Oil prices respond negatively to the contractionary policy shock. However, evidence for significant variation in their responses is at best weak with pairs clustered near the 45-degree line pointing to a marginally milder response in current times at the one-year horizon. This only reflects slight changes in the speed of adjustment as the magnitude of the responses is the same in the long run. The largest

change appears to occur in the response of the yen-US dollar exchange rate. In the period before the mid-1980s the response of the exchange rate peaks after six quarters and reverts back to zero over the considered horizon, whereas the depreciation of the dollar is more persistent in the latter half of the sample. This finding is supported by the lower right panel of Chart 5 which suggests a rise in the effect of the policy shock on the exchange rate in 2008 compared to 1975 with 62% of the probability mass lying above the 45-degree line.

3.3 Disaggregate price and quantity responses

In this section, we attempt to shed some light on the evolution of disaggregate price and quantity responses over time since movements in relative prices determine the extent to which monetary policy impulses have real effects. If all individual prices were to adjust rapidly and by similar amounts to monetary disturbances, then policy actions would only have moderate and short-lived effects on real economic activity. Knowing how price dynamics differ across goods and services that are part of household consumption expenditures helps improve our understanding of the monetary transmission mechanism at the disaggregate level. In Section 3.3.1 we study the impact of monetary shocks on the cross-sectional distribution of individual responses and how it has changed over time, and in Section 3.3.2 we discuss the implications of our findings for macroeconomic models and monetary policy.

3.3.1 Time pattern of sectoral responses

Impulse responses

Chart 6, panel A displays the median impulse response functions of the individual components of the personal consumption expenditure deflator and the corresponding real quantities after a contractionary monetary policy shock of 100 basis points at our two representative points in time: 1975 Q1 and 2008 Q1. As a reference, we have also plotted the median responses of the aggregate price and real consumption measures along with the 16th and 84th percentiles of the posterior distribution (red lines) as well as the unweighted average of the disaggregate price and quantity responses (black stars). While the mean of the sectoral price and quantity responses closely tracks the path of the aggregate indices, lying on or within the error bands at all horizons,

there is widespread heterogeneity in their adjustment dynamics at the disaggregate level in terms of degree of responsiveness and direction.

In contrast to 1975 where almost half of the individual prices rise permanently following an unexpected monetary tightening, in 2008 we observe that the cross-sectional distribution of price reactions is compressed and shifts downward with the majority of items responding negatively. Instead, the heterogeneous dynamics of the quantity responses do not exhibit such a noticeable change with regard to the range of responses across time in the long run, but the distribution moves upward implying that the reduction of consumption volumes is more short-lived. In fact, the aggregate PCE consumption as well as most of the sectoral quantities of goods and services show a substantial and more persistent fall in 1975, while the reactions are more muted in 2008. Only a few categories of goods and services experience an increase in demand despite the interest rate rise.

To get a better sense of changes in the pattern of the individual price and quantity responses over time, Chart 6, panel B depicts the evolution of the disaggregate, aggregate and mean reactions over the whole sample period four and twelve quarters after the restrictive monetary policy action. A considerable fraction of price responses displays a price puzzle during the 1970s, especially at short horizons, which attenuates as time progresses and from the mid-1980s onwards many of the sectoral prices decline in response to a negative monetary shock. In comparison to the price responses, there seems to be less evidence for time variation in the reaction of real quantities, with a large share of individual goods' and services' consumption volumes falling after an unexpected increase in the Fed funds rate. However, the slightly stronger reduction in consumption quantities in the early part of the sample might be linked to higher individual prices.

We follow Boivin *et al* (2009) to get an idea of the interaction between sectoral quantity and price responses conditional on a monetary policy shock over time. Chart 7 shows two scatter plots of price-volume combinations for all PCE items one year after the monetary innovation, together with the cross-sectional regression line for 1975 Q1 and 2008 Q1. In line with Boivin *et al* (2009), we find that sectors where prices react the most, quantities adjust the least. Over time two things happen. First, the cloud of pairs shifts to the lower left quadrant implying that more disaggregate price and consumption responses are negative in the more recent past. Second, the

regression line flattens indicating that consumption volumes get less responsive relative to prices. In the lower left graph of Chart 7, we plotted the estimated slope coefficients which show how the price-volume relationship evolves over time. The gradual decline of the slope coefficient over the sample period reveals that adjustments to monetary shocks take place more and more by prices than by quantities. Thus, in more recent times, sectors on average absorb contractionary monetary policy actions by adapting sales prices rather than production volumes. This finding at the disaggregate level is in line with the decline in real effects of monetary surprises observed for economy-wide activity measures and the greater sensitivity of aggregate price responses.

Cross-sectional distribution of prices

Another way of summarising the information of the effects of monetary policy shocks on the spread of disaggregate price responses is by looking at their entire distribution across items. Chart 8, panel A shows the estimated smoothed densities of sectoral prices for the years 1975 and 2008 at selected horizons in the upper part, and the distributions in four different years 4 and 16 quarters after the monetary innovation in the lower part. The evolution of the cross-sectional distribution in the quarters following the shock illustrates how a monetary policy disturbance propagates through the individual prices of goods and services inducing changes in relative prices which in turn affect the real economy.

We observe a gradual increase in the dispersion over the first years which highlights differences in speed and magnitude of price adjustments. While the distribution of disaggregate price responses in 1975 widens around zero as the horizon lengthens, it gradually shifts to the right of the origin in 2008 which implies that now more individual prices decline after a contractionary shock albeit with some delay. Also the size of the shift is greater in more recent times confirming the evidence at the aggregate level of a greater responsiveness of prices to monetary impulses. As is evident from a comparison of the evolution of the cross-sectional densities at horizons 4 and 16 at different points in time, there is a gradual transition from mainly positive to more negative responses both over the horizon and the sample.

A more concise way of capturing the evolutionary pattern of the cross-sectional distribution for all horizons over the whole sample period is by describing the densities by means of their moments which are depicted in Chart 8, panel B. During the first years of the sample, the median

of the distribution falls only slightly below zero over the horizon confirming our evidence that the price responses of single PCE components are spread around the origin with almost half of them still being in the positive range 20 quarters after the shock. Across time and horizons the median gets more and more negative pointing to a considerable leftward shift of the sectoral density. The standard deviation and the interquartile range (the difference between the 25th and 75th percentiles) capture the dispersion of relative price movements induced by the contractionary monetary policy shock. Both measures rise steeply over the first ten quarters and then stabilise at this level as the horizon extends further. The widening of the price dispersion is symptomatic of the fact that price reactions are not uniform across PCE categories.

The spread of the density is remarkably stable at around 0.75% during the 1970s and 1980s but decreases steadily from the early 1990s onwards. A similar evolutionary pattern emerges for the interquartile range. Both moments impart that the relative price dispersion is very persistent reaching a permanently higher level over the horizon considered here, but moderates slightly along the time line. We also observe that the distribution of sectoral price responses is skewed to the right in the initial periods after the contractionary monetary policy shock indicating that there are large upward deviations from the mean. This initial positive skewness is somewhat larger in the current period. The skew turns negative after about six quarters and continues to decline until the end of the horizon. Thus, at longer horizons the distribution is strongly skewed to the left throughout the sample which means that a couple of individual price responses are far below the mean whereas the majority of responses are close to the mean or exceed it by a small amount.

The kurtosis tells us how many individual price responses are located in the tails of the sectoral densities. In response to restrictive monetary disturbances, we observe an accumulation of responses in the tails of the cross-sectional distribution which diminishes in subsequent quarters. However, at longer horizons the degree of leptokurtosis experiences a substantial increase and remains at this higher level implying that extreme price reactions are a permanent feature of the cross-sectional distribution. While the long-run size of the tails does not differ much over time, changes in the fat-tailedness of the distribution across horizons get more pronounced in the latter part of the sample.

3.3.2 Implications of sectoral price responses

Two aspects of our results stand out so far: first, the existence of a price puzzle in the short run at a high level of disaggregation and second, considerable variation in price responses both across the panel of individual goods and services, and over time. In what follows, we will analyse the potential implications of these two findings for the monetary transmission mechanism and macroeconomic modelling.

Price puzzle

While a price puzzle is absent for the aggregate PCE deflator, the finding of a price puzzle at disaggregate level over the short horizon warrants further investigation. Despite the fact that the price puzzle attenuates considerably over time, it does not vanish for all categories of goods and services contained in the personal consumption basket. The conventional explanation of omission of relevant variables from the information set of the policymaker does not seem plausible in our data-rich environment (see Bernanke *et al* (2005)) ie mis-specification is highly unlikely to account for the positive responses of some disaggregate price classes since we arguably have included a large amount of macroeconomic indicators and sectoral conditions which a central bank might take into consideration when making policy decisions. Hence, the widely held view that the price puzzle is the result of the Fed possessing more accurate information regarding incipient inflation does not apply here.

By the same token, missing or latent variables have the potential to induce a bias in recovering structural monetary policy shocks. In fact, Bils *et al* (2003) who also find anomalous reactions of individual prices and quantities to a contractionary monetary policy shock attribute this finding to shocks derived from small-scale systems as not being truly exogenous. However, since we filter out endogenous policy reactions that might impair the correct identification of genuine structural monetary policy shocks, the risk of mis-identification should be minimised compared to low-dimensional VARs as shown above. It is also worth stressing that in contrast to most previous work (eg Balke and Wynne (2007) and Lastrapes (2006)), the disaggregated series are an integral part of our empirical model and hence, not subject to the criticism that sectoral data are merely appended to a macro VAR with potentially controversial implications.¹³ Instead,

¹³One of these implications is that aggregate price measures react to a monetary policy shock with a lag, whereas all individual price

macroeconomic and sectoral developments are modelled in a unifying framework establishing a direct link between macro and micro dynamics. Furthermore, we allow for time variation thereby taking account of the observed instability in macroeconomic time series due to changes in the economic environment as well as improvements in the conduct of monetary policy ie different monetary policy regimes, and variations in the volatility of shocks (as documented in the literature on the Great Moderation).

Consequently, if mis-specification and other biases can be excluded as an explanation for the sectoral price puzzles, our evidence provides a case for the price puzzle not being a puzzle at disaggregate level but rather a distinctive feature of sectoral dynamics which should allow us to infer something about the price-setting behaviour of firms in reaction to monetary surprises. In fact, the finding that the prices of a non-negligible fraction of individual consumption goods and services still respond positively at the short horizon after the mid-1980s could indicate that pricing strategies play an important role at the firm level which are hidden in (the response of) aggregate price measures.

There are various reasons why a firm might opt for raising the price of its products when confronted with an unexpected monetary policy tightening leading to a supply-side channel of monetary transmission. Barth and Ramey (2002) argue that in view of financial market frictions firms experience an unanticipated increase in the Federal funds rate as a cost-push shock and cope with it by passing the increased production costs on to consumers, at least in the short run. In a similar vein, Stiglitz (1992) suggests that in an imperfectly competitive environment firms tend to raise their prices following a monetary contraction in order to increase their cash flows momentarily before sales recede, at the expense of facing higher costs of their behaviour in terms of greater demand reductions in the future. Both pricing strategies depend on the cost structure and balance sheet situation of individual firms and hence, are the result of financial constraints.

Other factors that might influence a firm's price-setting behaviour in such circumstances is the low demand elasticity for its products where the price can be raised without incurring too great a loss in terms of volumes purchased as well as the degree of competition where more market power facilitates passing on higher costs. However, supply-side related propagation mechanisms of monetary policy shocks became weaker and more short-lived over time which is reflected in

components can respond contemporaneously.

the smaller share of individual responses displaying a price puzzle. Also Barth and Ramey (2002) show by means of a sample split that this cost-side channel of monetary transmission dominates in many industries in the period before 1980 and weakens thereafter, which is consistent with our disaggregate evidence. Factors that could have contributed to this weakening might be sought in changes in the financial structure such as deregulation and financial innovations which mitigate frictions, and globalisation resulting in greater competition (so that firms do not just take the domestic situation into account for their pricing decisions but also the international context).

Heterogeneity of price responses

A second striking feature is the dispersion across responses of individual prices to monetary policy impulses which might reinforce the idea that various channels of monetary transmission are at work in different sectors that differ in strength and importance, ie industries respond differently to monetary shocks depending on which channel they are most sensitive to. Hence, it might be of interest to explore which sectors contribute most to the dispersion and are more prone to displaying a price puzzle by grouping the price and quantity responses into different categories.¹⁴

We first organise the responses into three major sub-categories – durables, non-durables and services – which are depicted in Chart 9, panel A, for the entire sample period. We chose to report the dynamic effects eight quarters after the monetary policy innovation since this strikes a good balance between the short run, ie the price puzzle dying out, and the long run, ie the widening of the dispersion. As emerges from the graphs, durables are most sensitive to interest rate innovations and react in the expected way showing a considerable fall in consumption volumes and a decline in the price level after the 1970 period with the exception of two items. Durable goods also contribute the least to the dispersion of sectoral prices since individual impulse responses are closely aligned.

Instead, non-durable goods and to some extent services account for a large share of cross-sectional heterogeneity since price responses are widely dispersed covering a broad range of positive and negative values. Supply-side effects appear to play an important role in the

¹⁴Boivin *et al* (2009) explain the observed dispersion of producer prices by the varying degree of market power and the volatility and persistence of the idiosyncratic components.

propagation of monetary impulses to service categories like transportation, household operation and recreational activities, as can be seen from Chart 9, panel B, where we group disaggregate responses according to more detailed product classes. However, the positive price responses in these sectors are accompanied by a relatively large decrease in quantities purchased providing some evidence for the cash-flow argument advanced by Stiglitz (1992). Responses belonging to the product groups vehicles, clothes as well as food and furniture for most of their components rather comply with the demand channel of monetary transmission since the early 1980s.

Relative price effects

Another consequence of this heterogeneity of disaggregate responses is that monetary policy disturbances exert a considerable effect on relative prices. Standard macro models that try to account for relative price movements in response to monetary shocks and hence, for monetary non-neutralities are based on the assumption that: (1) prices are flexible but different categories of goods and services are affected in different proportions as a result of misperceptions about the source of price fluctuations, or (2) frictions that arise from the existence of menu costs, informational stickiness (eg Mankiw and Reis (2002)) or rational inattention (eg Sims (2003)) allow only a subset of firms to change their prices immediately. However, at least two features of our empirical results cannot be reconciled with the sources of relative price effects suggested by these models.

First, the effects on dispersion appear to be relatively long-lasting which is suggestive of the fact that monetary policy shocks lead to important non-neutralities even at high levels of disaggregation.¹⁵ If price dispersion were just related to timing lags and informational delays, then the dispersion should widen initially but converge to a new general price equilibrium (new steady state) as time passes, but we observe that the dispersion persists, at least for the horizon we consider. Bils *et al* (2003) also emphasise that this persistency of movements in relative prices and quantities consumed runs counter the premise that monetary non-neutrality derives from differences in price flexibility across consumption categories. As a consequence, there are also permanent effects on the composition of output which is reflected in the dispersion of real consumption responses and hence, the efficiency of resource allocation across sectors which

¹⁵Lastrapes (2006) using VARs identified with long-run restrictions also finds that money supply shocks have persistent effects on the cross-sectional distribution of relative commodity prices.

conflicts as well with conventional models for price determination. However, Carvalho (2006) shows that by introducing heterogeneity in price-setting behaviour into macro models, the real effects of monetary shocks are amplified and more persistent than in standard models.

Second, common to all these standard models of price determination is that frictions imply changes in the same direction but of different magnitude and speed of adjustment. Thus, the source of relative price effects lies in differences in the frequency of price adjustments either due to timing or informational constraints. Hence, what existing micro-founded models cannot account for are moves of sectoral prices in different directions, ie falling and rising individual prices following a monetary policy contraction which seems a key feature of the disaggregated data as documented above. Balke and Wynne (2007) who also find positive and negative price responses after an unexpected interest rate increase propose to model financial market frictions in view of the cost-of-capital channel of monetary transmission.

Hence, based on our empirical findings, endogenising pricing decisions of firms to account for strategic behaviour in combination with sectoral heterogeneity in financial constraints could yield richer dynamics that better comply with the microeconomic evidence on price-setting behaviour.

Policy implications

Given the key role of price-setting behaviour in understanding the monetary transmission mechanism, knowing how monetary policy shocks affect sectoral prices and quantities provides the monetary authority with useful information on how to interpret sectoral signals in order to devise an optimal policy reaction. Measures of aggregate inflation may hide disaggregate dynamics and pricing behaviour at the microeconomic level. In fact, Aoki (2001) stresses the importance of selecting the appropriate goal variables for the conduct of monetary policy given that the behaviour of sectoral prices can differ substantially from aggregate indices.

In view of this, Carlstrom *et al* (2006) propose a two-sector model which includes an interest rate rule according to which the central bank can react to sectoral inflation rates with different intensities. It might indeed be optimal to put different weights on sub-indices of aggregate price measures in policy analysis depending on the underlying pricing behaviour of sectors to monetary impulses. Our findings should make policymakers more aware about the importance of

heterogeneity and the potentially important compositional effects of monetary policy actions.

4 Conclusions

In this paper, we have re-examined the evolution of the US monetary transmission mechanism over time using an extended version of the factor-augmented VAR model introduced in Bernanke *et al* (2005). Our empirical framework incorporates information from almost 600 macroeconomic and financial indicators representing various sectors of the economy. By allowing for time variation in the coefficients and stochastic volatility in the variances of the shocks, we investigate the time-varying dynamic responses to a contractionary monetary policy shock for macroeconomic aggregates and disaggregate prices and quantities of personal consumption expenditures. This is important not only to get a better understanding of the behaviour of sectoral prices in response to monetary disturbances, but also to assess the role of price-setting behaviour in the propagation mechanism and how this might have altered in response to changes in the macroeconomic environment.

We find no evidence of a price puzzle for any of the aggregate price measures throughout the sample period. This may suggest that the extra information captured by the factors leads to more robust structural estimates in that it mimics the central bank's practice of examining and reacting to a wide variety of data series. Likewise, accounting for time-varying dynamics might matter for these findings since it allows the central bank to continuously learn about the state of the economy and adapt its policy behaviour accordingly.

Instead, at the disaggregate level, a considerable portion of sectoral price responses displays a significant price puzzle at short horizons during the 1970s which ameliorates from the early 1980s onwards. On the other hand, there seems to be less evidence for time variation in the reaction of real quantities at both the aggregate and sectoral level following an unexpected increase in the Federal funds rate. Interestingly, we find a shift in the relationship between the responses of prices relative to consumption volumes which tends to suggest that over time nominal adjustment of the economy following a restrictive monetary policy shock increasingly takes place by prices rather than by quantities.

In addition to significant time variation in the median responses of prices at the disaggregate



level, we also observe that the cross-sectional distribution of responses has undergone substantial changes over the sample period. While the price dispersion diminishes slightly over time, it widens considerably over the forecast horizon and is very persistent in the long run. This indicates that the transmission mechanism of monetary innovations is highly heterogeneous across components of personal consumption expenditures suggesting that monetary policy actions exert an important, and potentially long-lasting, influence on relative prices in the US economy. Put differently, this cross-sectional heterogeneity might be a sign that various sectors employ different pricing strategies to deal with an unanticipated increase in the funds rate.

The period around the mid-1980s for which we have identified important changes in disaggregate price responses to unexpected monetary policy actions corresponds to the time for which many studies have documented an increase in macroeconomic stability as well as changes in the Fed's attitude towards inflation stabilisation. Future research could therefore be directed towards understanding the interactions between monetary policy objectives and the heterogeneity of disaggregate price responses to macroeconomic disturbances.

Appendix A: Priors and estimation

Consider the time-varying FAVAR model given by equations (6) and (7). As shown by Bernanke *et al* (2005) identification requires some restrictions on the factor loading matrix. Following Bernanke *et al* (2005) the top $J \times J$ block of Λ is assumed to be an identity matrix and the top $J \times M$ block of Ψ is assumed to be zero.

Prior distributions and starting values

Factors and factor loadings

Following Bernanke *et al* (2005) we centre our prior on the factors (and obtain starting values) by using a principal components (PC) estimator applied to each $X_{i,t}$. In order to reflect the uncertainty surrounding the choice of starting values, a large prior covariance of the states ($P_{0|0}$) is assumed by setting it equal to the identity matrix.

Starting values for the factor loadings are also obtained from the PC estimator (with the restrictions given above imposed). The prior on the diagonal elements of R is assumed to be inverse gamma:

$$R_{ii} \sim IG(5, 0.01)$$

In choosing this diffuse prior we closely follow Bernanke *et al* (2005), but employ a slightly higher scale parameter in order to reflect the high volatility of some of the series in the panel.

VAR coefficients

The prior for the VAR coefficients is obtained via a fixed-coefficient VAR model estimated over the sample 1960 Q2 to 1970 Q2 using the PC estimates of the factors. θ_0 is therefore set equal to

$$\theta_0 \sim N(\hat{\theta}^{OLS}, V)$$

where V equals the OLS estimates of $var(\hat{\theta}^{OLS})$ on the main diagonal.

Elements of H_t

Let \hat{v}^{ols} denote the OLS estimate of the VAR covariance matrix estimated on the pre-sample data described above. The prior for the diagonal elements of the VAR covariance matrix (see **(10)**) is as follows:

$$\ln h_0 \sim N(\ln \mu_0, 10 \times I_N)$$

where μ_0 are the diagonal elements of \hat{v}^{ols} .

Elements of A_t

The prior for the off-diagonal elements A_t is

$$A_0 \sim N(\hat{a}^{ols}, V(\hat{a}^{ols}))$$

where \hat{a}^{ols} are the off-diagonal elements of the Choleski decomposition of \hat{v}^{ols} , with each row scaled by the corresponding element on the diagonal. $V(\hat{a}^{ols})$ is assumed to be diagonal with the diagonal elements set equal to ten times the absolute value of the corresponding element of \hat{a}^{ols} .

Hyperparameters

The prior on Q is assumed to be inverse Wishart

$$Q_0 \sim IW(\bar{Q}_0, T_0)$$

where \bar{Q}_0 is assumed to be $var(\hat{\theta}^{OLS}) \times 10^{-4} \times T_0$ and T_0 is the length of the sample used for calibration.

The prior distribution for the blocks of S is inverse Wishart:

$$S_{i,0} \sim IW(\bar{S}_i, K_i)$$

where $i = 1, 2, 3$ indexes the blocks of S . \bar{S}_i is calibrated using \hat{a}^{ols} . Specifically, \bar{S}_i is a diagonal matrix with the relevant elements of \hat{a}^{ols} multiplied by 10^{-3} and K_i are the degrees of freedom which are set to $i + 1$ to obtain a proper prior as in Primiceri (2005).

Following Cogley and Sargent (2005), we postulate an inverse gamma distribution for the

elements of G :

$$\sigma_i^2 \sim IG\left(\frac{10^{-4}}{2}, \frac{1}{2}\right)$$

Simulating the posterior distributions

Factors and factor loadings

This closely follows Bernanke *et al* (2005). Details can also be found in Kim and Nelson (1999a).

Factors

The distribution of the factors F_t is linear and Gaussian:

$$F_T | X_{i,t}, R_t, \Xi \sim N(F_{T|T}, P_{T|T})$$

$$F_t | F_{t+1}, X_{i,t}, R_t, \Xi \sim N(F_{t|t+1, F_{t+1}}, P_{t|t+1, F_{t+1}})$$

where $t = T - 1, \dots, 1$, Ξ denotes a vector that holds all the other FAVAR parameters and:

$$F_{T|T} = E(F_T | X_{i,t}, R_t, \Xi)$$

$$P_{T|T} = Cov(F_T | X_{i,t}, R_t, \Xi)$$

$$F_{t|t+1, F_{t+1}} = E(F_t | X_{i,t}, R_t, \Xi, F_{t+1})$$

$$P_{t|t+1, F_{t+1}} = Cov(F_t | X_{i,t}, R_t, \Xi, F_{t+1})$$

As shown by Carter and Kohn (2004), the simulation proceeds as follows. First, we use the Kalman filter to draw $F_{T|T}$ and $P_{T|T}$ and then, proceed backwards in time using:

$$F_{t|t+1} = F_{t|t} + P_{t|t} P_{t+1|t}^{-1} (F_{t+1} - F_t)$$

$$P_{t|t+1} = P_{t|t} - P_{t|t} P_{t+1|t}^{-1} P_{t|t}$$

If more than one lag of the factors appears in the VAR model, this procedure has to be modified to take account of the fact that the covariance matrix of the shocks to the transition equation (used in the filtering procedure described above) is singular. For details see Kim and Nelson (1999a).

Elements of R

As in Bernanke *et al* (2005), R is a diagonal matrix. The diagonal elements R_{ii} are drawn from the following inverse gamma distribution:

$$R_{ii} \sim IG(\bar{R}_{ii}, T + 0.01)$$

where

$$\bar{R}_{ii} = 5 + \hat{e}_i' \hat{e}_i + \Gamma_i' \left[\bar{M}_0^{-1} + (F_{i,t}' F_{i,t})^{-1} \right]^{-1} \Gamma_i$$

where $\Gamma_i = \Lambda_i$ or if appropriate $\Gamma_i = [\Lambda_i, \Psi_i]$ and \hat{e}_i denotes the OLS estimate the i^{th} element of R . As in Bernanke *et al* (2005), we set $M_0 = I$.

Elements of Λ and Ψ

Letting $\Gamma_i = \Lambda_i$ or $\Gamma_i = [\Lambda_i, \Psi_i]$ for the appropriate equation, the factor loadings are sampled from

$$\Gamma_i \sim N(\bar{\Gamma}_i, R_{ii} \bar{M}_i^{-1})$$

where $\bar{\Gamma}_i = \bar{M}_i^{-1} (F_{i,t}' F_{i,t}) \hat{\Gamma}_i$, $\bar{M}_i = \bar{M}_0 + (F_{i,t}' F_{i,t})$ and $\hat{\Gamma}_i$ represents an OLS estimate.

Time-varying VAR

Given an estimate for the factors, the model becomes a VAR model with drifting coefficients and covariances. This model has become fairly standard in the literature and details on the posterior distributions can be found in a number of papers including Cogley and Sargent (2005), Cogley *et al* (2005) and Primiceri (2005). Here, we describe the algorithm briefly. Details can be found in the papers mentioned above.

VAR coefficients θ_t

As in the case of the unobserved factors, the time-varying VAR coefficients are drawn using the methods described in Carter and Kohn (2004).

Elements of H_t

Following Cogley and Sargent (2005), the diagonal elements of the VAR covariance matrix are sampled using the methods described in Jacquier *et al* (1994).

Elements of A_t

Given a draw for θ_t , the VAR model can be written as


$$A_t' (\tilde{Z}_t) = u_t$$

where $\tilde{Z}_t = Z_t - \delta_t - \sum_{l=1}^L \phi_{l,t} Z_{t-l} = v_t$ and $\text{var}(u_t) = H_t$. This is a system of equations with time-varying coefficients and given a block diagonal form for $\text{var}(\tau_t)$, the standard methods for state-space models described in Carter and Kohn (2004) can be applied.

VAR hyperparameters

Conditional on Z_t , θ_t , H_t , and A_t , the innovations to θ_t , H_t , and A_t are observable, which allows us to draw the hyperparameters – the elements of Q , S , and the σ_i^2 – from their respective distributions.

Appendix B: Results based on simulated data



Charts

Chart 1: Time-varying median impulse response functions of real economic activity indicators at each point in time (first column) and in 1975 Q1 and 2008 Q1 (second and third columns) with 16th and 84th percentiles (shaded areas) to a 1% increase in the funds rate and joint distribution of the cumulated responses one year after the monetary policy shock in 1975 Q1 and 2008 Q1 (fourth column).

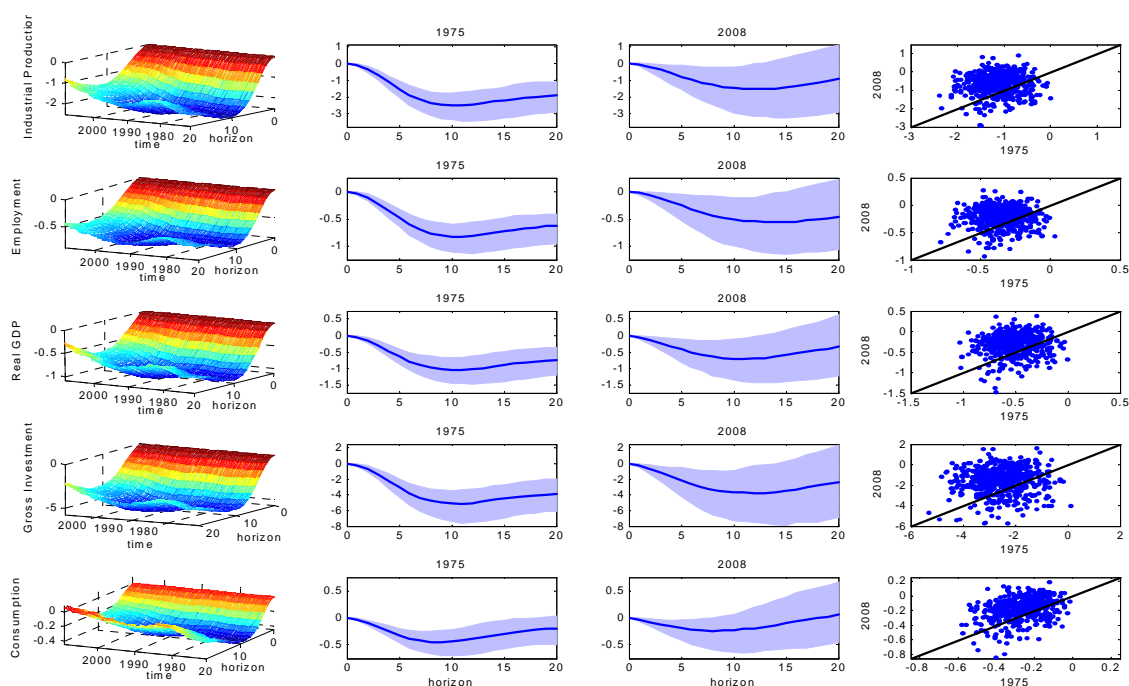


Chart 2: Time-varying median impulse response functions of inflation indicators at each point in time (first column) and in 1975 Q1 and 2008 Q1 (second and third columns) with 16th and 84th percentiles (shaded areas) to a 1% increase in the funds rate and joint distribution of the cumulated responses one year after the monetary policy shock in 1975 Q1 and 2008 Q1 (fourth column).

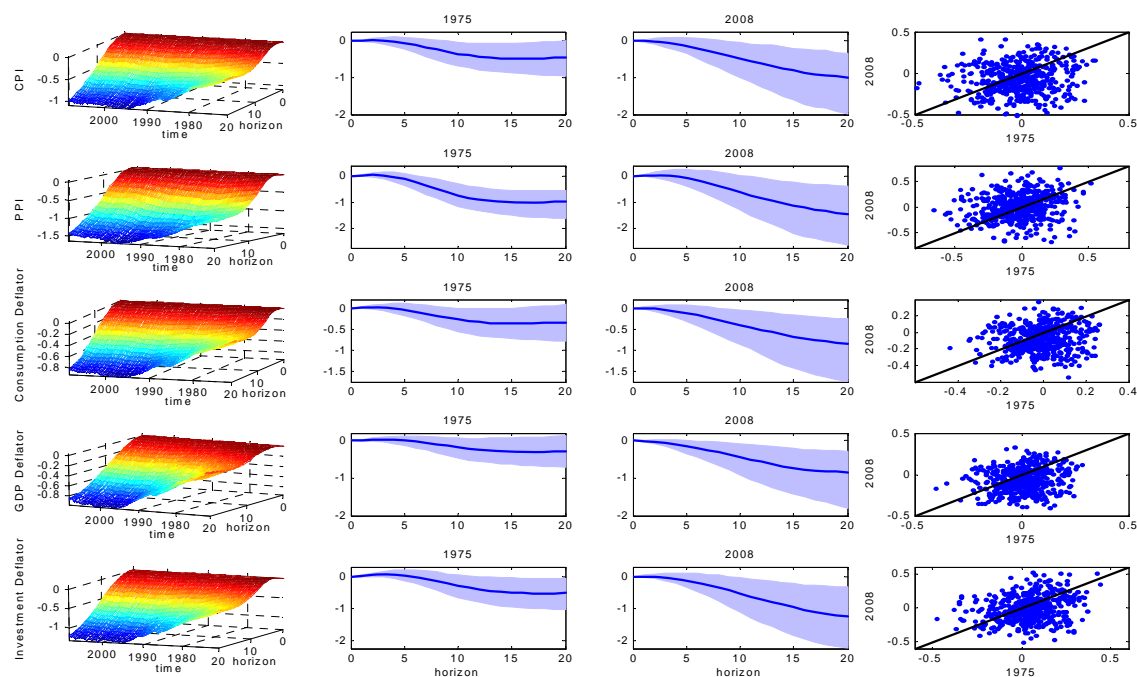


Chart 3: Time-varying median impulse response functions of unemployment, consumer prices and the funds rate obtained from a tri-variate VAR at each point in time (first column) and in 1975 Q1 and 2008 Q1 (blue lines in second and third columns) with 16th and 84th percentiles (shaded areas) to 1% increase in the funds rate. The red lines are the corresponding responses from the FAVAR specification.

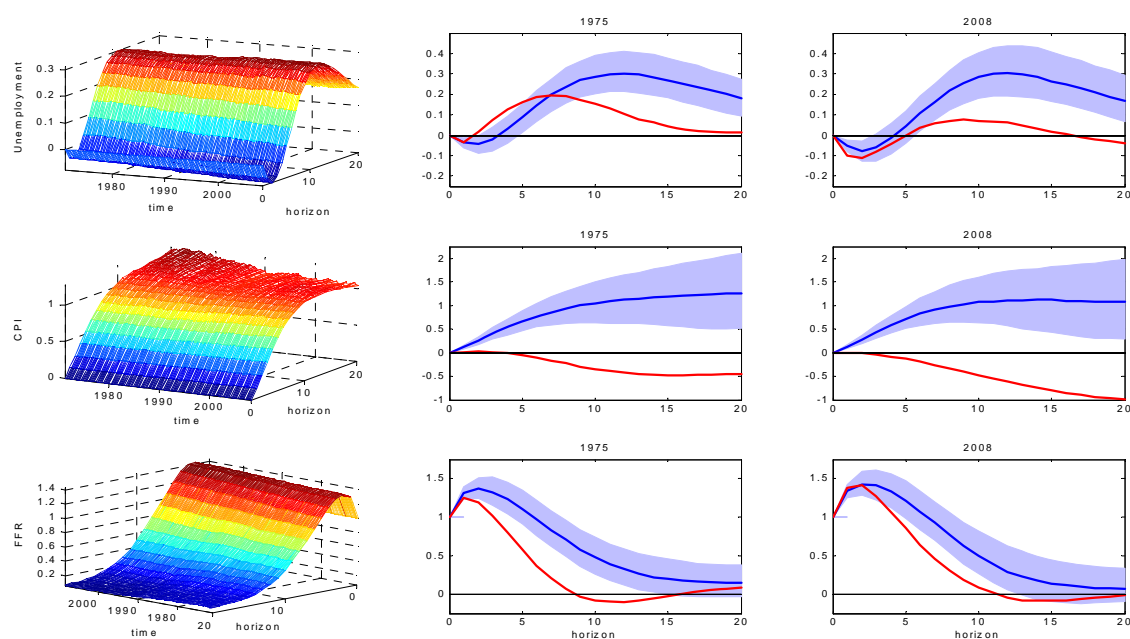


Chart 4: Time-varying median impulse response functions of money and credit measures at each point in time (first column) and in 1975 Q1 and 2008 Q1 (second and third columns) with 16th and 84th percentiles (shaded areas) to a 1% increase in the funds rate and joint distribution of the cumulated responses one year after the monetary policy shock in 1975 Q1 and 2008 Q1 (fourth column).

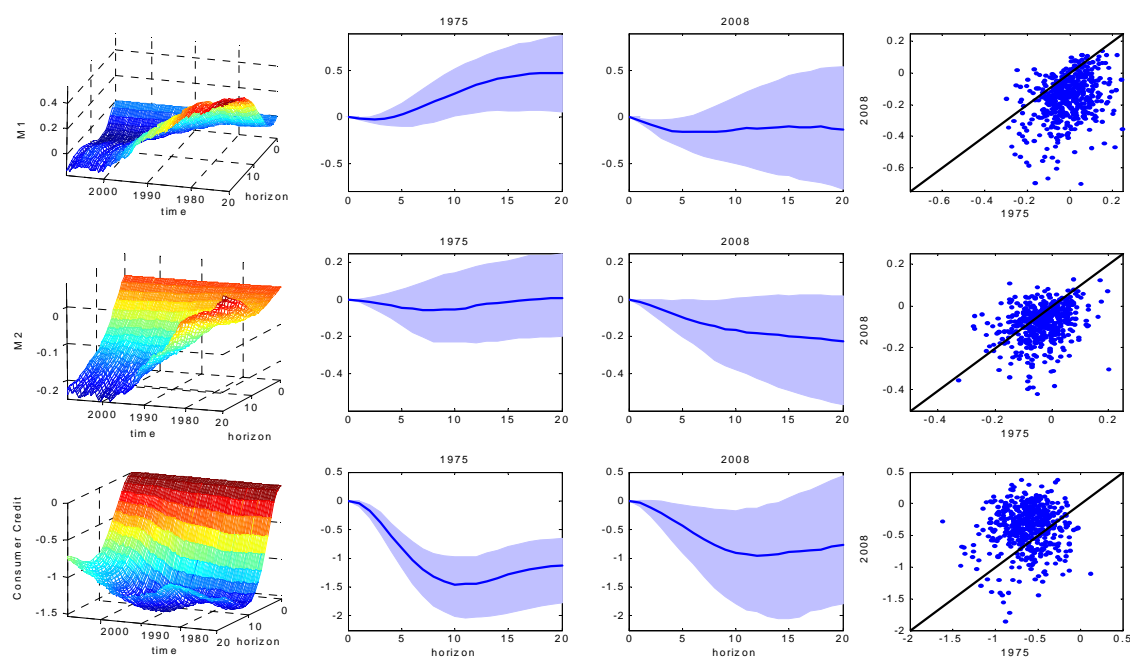


Chart 5: Time-varying median impulse response functions of several asset prices at each point in time (first column) and in 1975 Q1 and 2008 Q1 (second and third columns) with 16th and 84th percentiles (shaded areas) to a 1% increase in the funds rate and joint distribution of the cumulated responses one year after the monetary policy shock in 1975 Q1 and 2008 Q1 (fourth column).

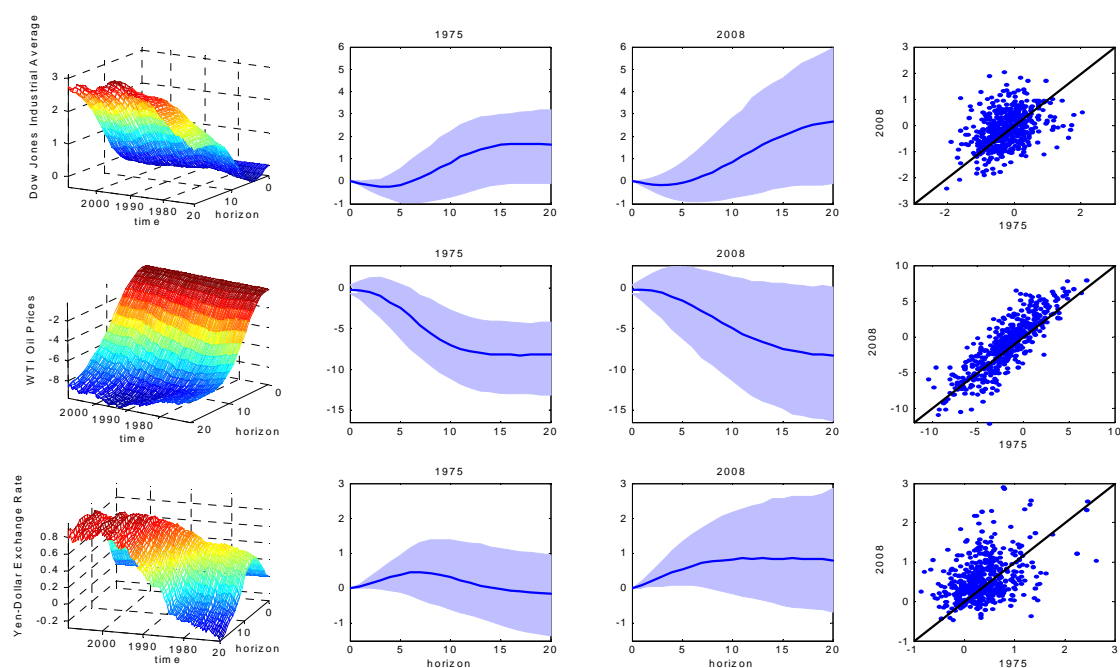


Chart 6: Median impulse responses to a 1% increase in the funds rate of disaggregated prices and quantities (dotted blue lines), aggregate PCE deflator and real consumption (solid red lines) with 16th and 84th percentiles (dashed red lines), unweighted average of individual responses (black stars). Panel A: At two points in time – 1975 Q1 and 2008 Q1. Panel B: At each point in time – four and twelve quarters after the monetary policy shock.

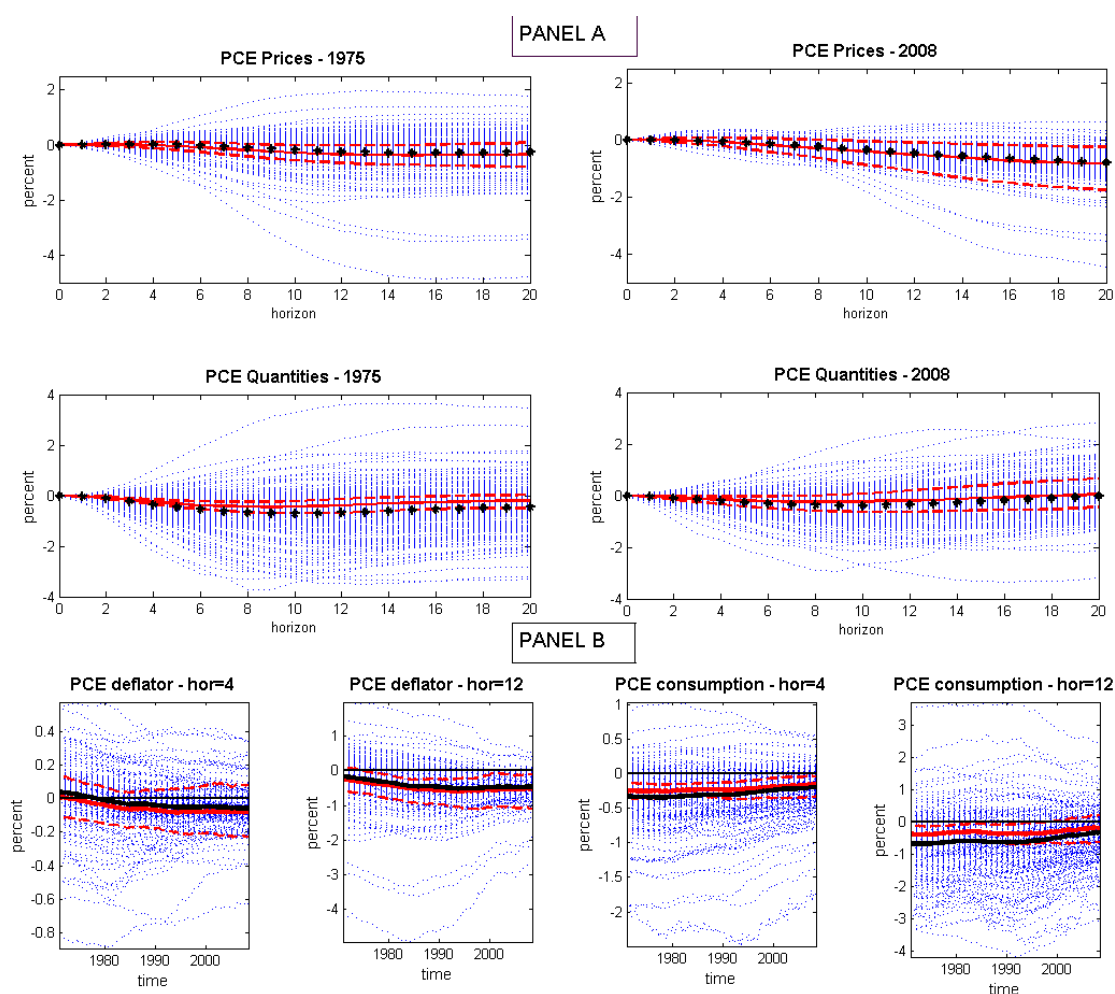


Chart 7: Relation between sectoral price and quantity responses after four quarters conditional on a monetary policy shock: (a) scatter plot and cross-sectional regression line for 1975 Q1 and 2008 Q1 (b) evolution of estimated slope coefficients over time.

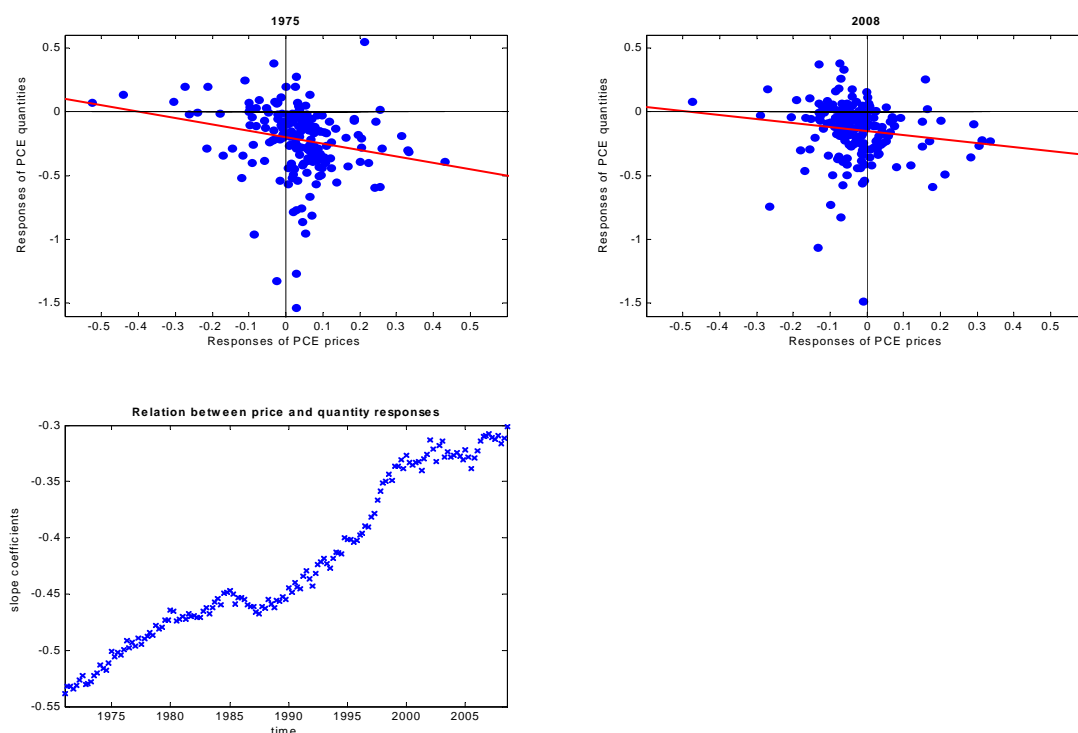


Chart 8: Panel A: Smoothed densities of cross-sectional price responses to a 1% increase in the funds rate at selected horizons and different points in time. Panel B: Moments of the cross-sectional distribution of individual price responses over time.

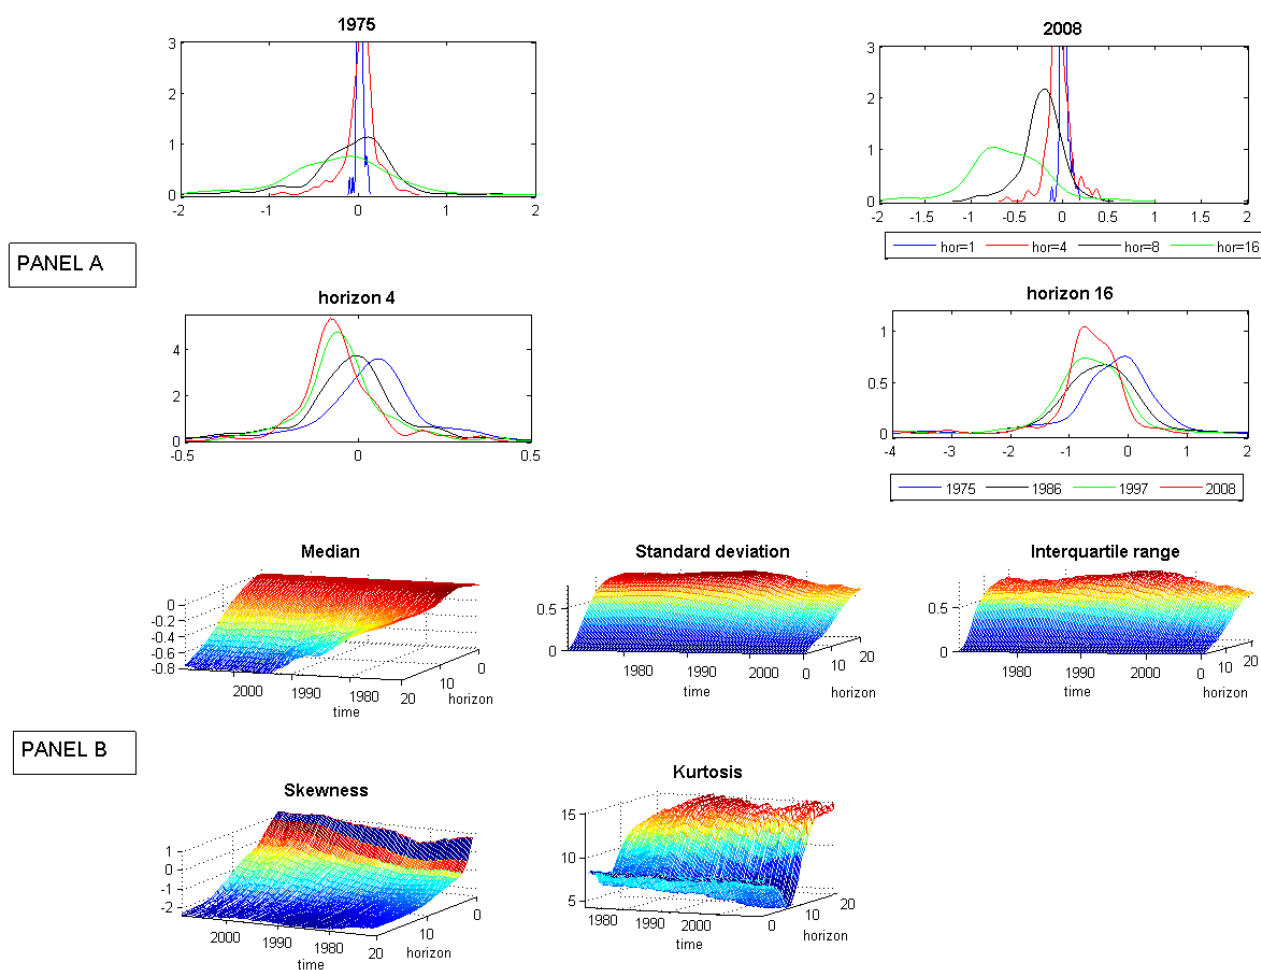
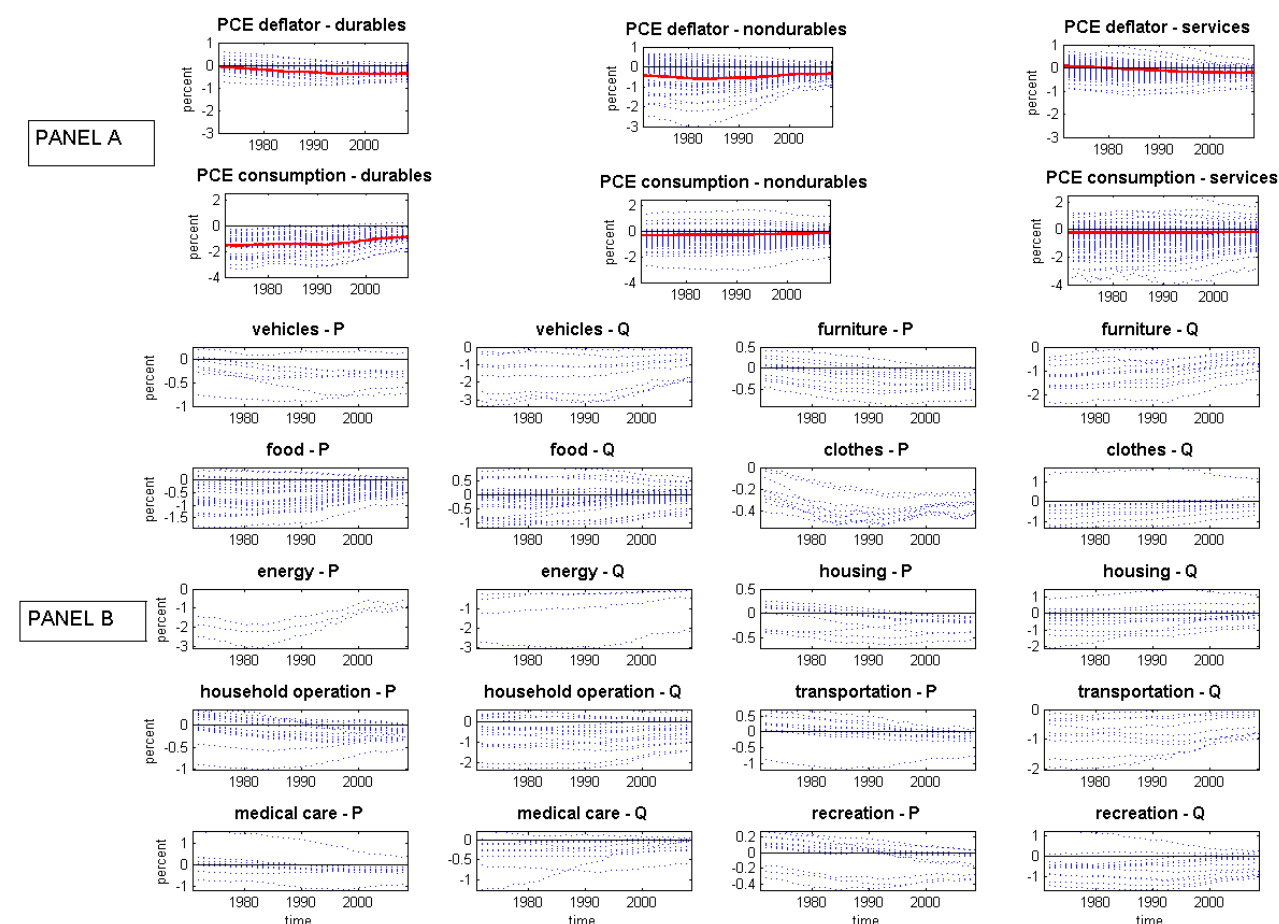


Chart 9: Time-varying median responses of prices and quantities eight quarters after the shock (dotted blue lines). Panel A: Major PCE components along with response of aggregate sub-category (solid red lines). Panel B: Grouped according to different product categories.



References

- Altissimo, F, Mojon, B and Zaffaroni, P (2009)**, ‘Can aggregation explain the persistence of inflation?’, *Journal of Monetary Economics*, Vol. 56, pages 231-41.
- Aoki, K (2001)**, ‘Optimal monetary policy responses to relative-price changes’, *Journal of Monetary Economics*, Vol. 48, pages 55-80.
- Balke, N and Wynne, M (2007)**, ‘The relative price effects of monetary shocks’, *Journal of Macroeconomics*, Vol. 29, pages 19-36.
- Barth, M and Ramey, V (2002)**, ‘The cost channel of monetary transmission’, in Bernanke, B and Rogoff, K (eds), *NBER Macroeconomics Annual 2001*, Cambridge, Massachusetts: MIT press, pages 199-239.
- Belviso, F and Milani, F (2006)**, ‘Structural factor-augmented VARs (SFAVARs) and the effects of monetary policy’, *The BE Journal of Macroeconomics*, Vol. 6, No. 3 (topics).
- Benati, L and Mumtaz, H (2007)**, ‘US evolving macroeconomic dynamics: a structural investigation’, *ECB Working Paper No. 476*.
- Benati, L and Surico, P (2009)**, ‘VAR analysis and the Great Moderation’, *American Economic Review*, Vol. 99, pages 1,636-52.
- Bernanke, B and Boivin, J (2003)**, ‘Monetary policy in a data-rich environment’, *Journal of Monetary Economics*, Vol. 50, pages 525-46.
- Bernanke, B, Boivin, J and Elias, P (2005)**, ‘Measuring the effects of monetary policy: a factor-augmented vector autoregressive (FAVAR) approach’, *Quarterly Journal of Economics*, Vol. 120, pages 387-422.
- Bils, M and Klenow, P (2004)**, ‘Some evidence on the importance of sticky prices’, *Journal of Political Economy*, Vol. 112, pages 947-85.
- Bils, M, Klenow, P and Kryvtsov, O (2003)**, ‘Sticky prices and monetary policy shocks’, *Federal Reserve Bank of Minneapolis Quarterly Review*, Vol. 27, pages 2-9.
- Blinder, A (1991)**, ‘Why are prices sticky? Preliminary results from an interview study’, *American Economic Review*, Vol. 81, pages 89-96.
- Boivin, J and Giannoni, M (2006)**, ‘Has monetary policy become more effective?’, *Review of Economics and Statistics*, Vol. 88, pages 445-62.
- Boivin, J, Giannoni, M and Mihov, I (2009)**, ‘Sticky prices and monetary policy: evidence from disaggregated US data’, *American Economic Review*, Vol. 99, pages 350-84.



- Carlstrom, C, Fuerst, T, Ghironi, F and Hernandez, K (2006)**, ‘Relative price dynamics and the aggregate economy’, Boston College, *mimeo*.
- Carter, C and Kohn, R (2004)**, ‘On Gibbs sampling for state space models’, *Biometrika*, Vol. 81, pages 541-53.
- Carvalho, C (2006)**, ‘Heterogeneity in price stickiness and the real effects of monetary shocks’, *Frontiers of Macroeconomics*, Vol. 2, pages 1-56.
- Castelnuovo, E and Surico, P (2009)**, ‘Monetary policy, inflation expectations and the price puzzle’, *Economic Journal*, forthcoming.
- Christiano, L, Eichenbaum, M and Evans, C (2005)**, ‘Nominal rigidities and the dynamic effects of a shock to monetary policy’, *Journal of Political Economy*, Vol. 113, pages 1-45.
- Clarida, R, Galí, J and Gertler, M (2000)**, ‘Monetary policy rules and macroeconomic stability: evidence and some theory’, *Quarterly Journal of Economics*, Vol. 115, pages 147-80.
- Clark, T (2006)**, ‘Disaggregate evidence on the persistence of consumer price inflation’, *Journal of Applied Econometrics*, Vol. 21, pages 563-87.
- Cogley, T, Morozov, S and Sargent, T (2005)**, ‘Bayesian fan charts for UK inflation: forecasting and sources of uncertainty in an evolving monetary system’, *Journal of Economic Dynamics and Control*, Vol. 29, No. 11, pages 1,891-925.
- Cogley, T and Sargent, T (2002)**, ‘Evolving post-WWII US inflation dynamics’, in Bernanke, B and Rogoff, K (eds), *NBER Macroeconomics Annual 2001*, Cambridge, Massachusetts: MIT Press, pages 331-47.
- Cogley, T and Sargent, T (2005)**, ‘Drift and volatilities: monetary policies and outcomes in the post WWII U.S.’, *Review of Economic Dynamics*, Vol. 8, pages 262-302.
- Del Negro, M (2003)**, ‘Comment on ‘Drifts and volatilities’ by Cogley and Sargent’, Federal Reserve Bank of Atlanta, *mimeo*.
- Del Negro, M and Otrok, C (2008)**, ‘Dynamic factor models with time-varying parameters: measuring changes in international business cycles’, *Staff Reports 326*, Federal Reserve Bank of New York.
- Farmer, R, Waggoner, D and Zha, T (2008)**, ‘Minimal state variable solutions to Markov-switching rational expectations models’, *Working Paper No. 2008-23*, Federal Reserve Bank of Atlanta.
- Jacquier, E, Polson, N and Rossi, P (1994)**, ‘Bayesian analysis of stochastic volatility models’, *Journal of Business and Economic Statistics*, Vol. 12, pages 371-418.
- Justiniano, A and Preston, B (2009)**, ‘Monetary policy and uncertainty in an empirical small open economy model’, *Journal of Applied Econometrics*, forthcoming.

Kim, C-J and Nelson, C R (1999a), *State-space models with regime switching*, Cambridge, Massachusetts: MIT Press.

Kim, C-J and Nelson, C R (1999b), ‘Has the US economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle’, *Review of Economics and Statistics*, Vol. 81, pages 608-16.

Koop, G, Pesaran, H and Potter, S (1996), ‘Impulse response analysis in nonlinear multivariate models’, *Journal of Econometrics*, Vol. 74, No.1, pages 119-47.

Lastrapes, W (2006), ‘Inflation and the distribution of relative prices: the role of productivity and money supply shocks’, *Journal of Money, Credit, and Banking*, Vol. 38, pages 2,159-98.

Lubik, T and Schorfheide, F (2004), ‘Testing for indeterminacy: an application to US monetary policy’, *American Economic Review*, Vol. 94, pages 190-217.

Mankiw, N and Reis, R (2002), ‘Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve’, *Quarterly Journal of Economics*, Vol. 117, pages 1,295-328.

McConnell, M and Pérez-Quirós, G (2000), ‘Output fluctuations in the United States: what has changed since the early 1980’s?’, *American Economic Review*, Vol. 90, pages 1464-1476.

Monacelli, T (2005), ‘Monetary policy in a low pass-through environment’, *Journal of Money, Credit, and Banking*, Vol. 37, pages 1,047-66.

Primiceri, G (2005), ‘Time varying structural vector autoregressions and monetary policy’, *Review of Economic Studies*, Vol. 72, No. 3, pages 821-52.

Sims, C (2003), ‘Implications of rational inattention’, *Journal of Monetary Economics*, Vol. 50, pages 665-90.

Sims, C and Zha, T (2006), ‘Were there regime switches in US monetary policy?’, *American Economic Review*, Vol. 96, pages 54-82.

Stiglitz, J (1992), ‘Capital markets and economic fluctuations in capitalist economies’, *European Economic Review*, Vol. 36, pages 269-306.

Stock, J and Watson, M (2002), ‘Macroeconomic forecasting using diffusion indexes’, *Journal of Business and Economic Statistics*, Vol. 20, pages 147-62.

Stock, J and Watson, M (2005), ‘Implications of dynamic factor models for VAR analysis’, *NBER Working Paper No. 11467*.

