Monetary policy surprises and their transmission through term premia and expected interest rates*

Iryna Kaminska†, Haroon Mumtaz‡§ and Roman Šustek¶

June 29, 2021

Abstract

Monetary policy moves the yield curve. What is the economic interpretation of such moves and what are their macroeconomic consequences? Applying an affine term structure model to high-frequency yield curve movements around FOMC announcements, we shed new light on these questions. Estimation is subject to restrictions addressing estimation bias in previous studies. By imposing additional structure, expectations and term premia are decomposed into three components interpreted as monetary policy action, expected path and its uncertainty. In a local projections model, the shocks identified by the three components provide insights into monetary policy transmission in the context of existing theories.

JEL Classification Codes: E43, E52, E58, G12, C58.

Keywords: High-frequency data, affine term structure model, estimation bias, multidimensional policy shocks, monetary policy transmission.

*We thank Ambrogio Cesa-Bianchi, Refet Gürkaynak, Silvia Miranda-Agrippino, Juan Rubio-Ramírez, Eric Swanson, Stan Zin, an anonymous referee, and seminar and conference participants at the Bank of England, the University of Surrey and the 3rd RCEA Macro-Money-Finance conference for valuable comments and suggestions. The views expressed are those of the authors and not necessarily of the Bank of England or its Monetary Policy Committee.

†Bank of England; Iryna.Kaminska@bankofengland.co.uk.

‡Queen Mary University of London; h.mumtaz@qmul.ac.uk.

§Corresponding author: Haroon Mumtaz, address: School of Economics and Finance, Queen Mary University of London, Mile End Road, London, E1 4NS, UK, tel: +44 20 7882 8839, e-mail: h.mumtaz@qmul.ac.uk.

¶Queen Mary University of London and the Centre for Macroeconomics; r.sustek@qmul.ac.uk.

1 Introduction

A classic question in macroeconomics concerns the transmission of monetary policy surprises into the economy. The interest in this question stems from the notion that empirical impulse-responses can guide the development of theory (eg, Christiano, Eichenbaum, and Evans, 1999). This research strategy, however, rests on the assumption that one can identify the relevant impulses (shocks) in the data. The traditional approach to this identification problem relies on monthly or quarterly vector auto-regressions (VAR) combining macroeconomic data with a short-term nominal interest rate, taken as a proxy for a policy instrument. Various identification schemes have been proposed within this approach (see Ramey, 2016, for a review). What they have in common, however, is that the identified shocks at best reflect monetary policy surprises relative to the mathematical expectations of the regression model.[1] Furthermore, VAR-based identification is limiting once financial data are included. How does one invert the reduced-form VAR residuals to identify monetary policy shocks when, at monthly or quarterly frequency, financial markets react to monetary policy and policy makers partially base their decisions on information contained in asset prices? At the same time, ignoring financial data is inefficient, as asset prices may reveal expectations and uncertainty about future monetary policy and some sectors, for instance the housing market, are sensitive to asset prices (long-term interest rates).[2]

High-frequency (HF) data can ameliorate the identification problem (Bagliano and Favero, 1999; Kuttner, 2001; Cochrane and Piazzesi, 2002; Gürkaynak, Sack, and Swanson, 2005a, are early contributions). The idea is that, up to a measurement error, the announcement of the outcome of a policy meeting is the only (exogenous) event impacting on asset prices in a tight enough window around the announcement. Asset price movements in that window

[1] An alternative identification strategy, proposed by Romer and Romer (2004), is based on central bank narrative.

[2] Evans and Marshall (1998) is an early attempt to include long-term interest rates in a macro VAR model with monetary policy shocks identified in one of the traditional ways. Rudebusch (1998) questions VAR-based policy shocks. Woodford (2005) provides a theoretical discussion of the key role of long-term interest rates in the transmission of monetary policy.

can thus provide instruments for policy shocks.[3] The dynamic effects of the shocks identified by the HF instruments can then be studied in a standard empirical macroeconomic model. Gertler and Karadi (2015) carry out such an exercise and arrive at a stark conclusion: monetary policy transmits into the economy almost exclusively through changes in term premia, with expected future interest rates left almost unaffected.[4] This finding presents a challenge to quantitative-theoretical models used for monetary policy analysis. In most models, monetary policy transmits through changes in the conditional mean of the nominal pricing kernel, not its variance, the relevant part for movements in term premia (eg, Atkeson and Kehoe, 2009). Furthermore, in practice, communication aimed at managing expectations of future monetary policy is an integral part of modern central banking (eg, Woodford, 2005).

In this paper, we revisit the relevance of expected future interest rates vs. term premia in the monetary transmission mechanism. However, we go beyond this basic decomposition. By imposing additional structure on estimated expectations and term premia, we decompose the HF yield curve movements in terms of components that can be assigned economic interpretation. These structural components are then used to identify policy shocks in local projections and study their dynamic effects. Our focus is on the nominal yield curve in the period 1996-2007, characterized by conventional monetary policy. In more detail, the analysis proceeds as follows

First, we employ an estimated affine term structure model (ATSM) to decompose the HF movements in yields around Federal Open Market Committee (FOMC) announcements into expected future interest rates and term premia.[5] Importantly, the ATSM is estimated subject to restrictions (Joslin, Singleton, and Zhu, 2011), leading to more precise estimates of expected interest rates and term premia than those obtained from VARs, the framework

[3]An implicit assumption in this approach is that asset prices reflect all available public information up to the point of the announcement.

[4]Term premia reflect risk compensation for holding a long-term bond and can be estimated as a difference between the observed long-term interest rate of a given maturity and a forecast of the path of the short rate over that time horizon (ignoring technical details such as measurement errors and Jensen's inequality).

[5]ATSMs are the go-to models in empirical finance to study the term structure of interest rates. See Diebold, Piazzesi, and Rudebusch (2005), Piazzesi (2006), Duffee (2012), or Gürkaynak and Wright (2012) for an introduction.

used by Gertler and Karadi (2015).[6] The estimates from the restricted ATSM show that expected interest rates are as important as term premia in explaining yield curve movements, including those around FOMC announcements. For instance, at the 10-year maturity, the two components have about the same variance.[7] Second, we use principal components (PCs) of the estimated HF changes in expectations and term premia around FOMC announcements as basis to construct orthogonal instruments for monetary policy shocks. A particular rotation is applied to a subset of the PCs to obtain components with an economic interpretation: (i) *action*, taking the form of a change in the current policy rate; (ii) change in the *expected path* of future policy rates; and (iii) change in *uncertainty* about future monetary policy.[8] Finally, we use the instruments in a local projections (LP) macro model (Jordà, 2005) to trace out the dynamic effects of the policy shocks, identified by the instruments, on macro variables. Most of the estimated responses can be justified through the lenses of existing theories, although we also document some new patterns. The analysis delivers especially tight findings for the housing market, a sector which, through mortgage finance, is closely related to the term structure.[9]

We view our analysis as the natural next step in the line of research using HF data to identify monetary policy shocks. The first HF studies used a single asset, fed funds rate futures for the current month, to identify a single monetary policy shock—an action—capturing an unexpected change in the current policy rate (eg, Kuttner, 2001; Gürkaynak et al., 2005a; Beechey, 2007). Recognizing the complexity of monetary policy announcements, the work of Gürkaynak, Sack, and Swanson (2005b) extended the single-shock approach to two shocks:

[6]Due to a small sample bias (eg, Bauer, Rudebusch, and Wu, 2012), VARs substantially underestimate the responses of expected interest rates to current shocks, thus prescribing a bulk of the observed movements in long-term interest rates to term premia.

[7]In terms of the specific restrictions imposed on the ATSM, we follow two approaches proposed by Bauer et al. (2012) and Bauer (2018).

[8]In the literature, the term "target" is sometimes used for what we refer to as "action"; the terms "path", "statement" or "forward guidance" are used for what we call "expected path". As "forward guidance" is often used specifically in the context of the post-2008 zero-lower bound period, we prefer to avoid this term. We also prefer the term "expected path" to "path" or "statement" in order to stress that this component is extracted from the expectations part of the yield curve.

[9]The findings reported in the main text are based on the ATSM estimated on monthly data, which is the standard in the literature. In an Online Appendix we confirm that estimates based on daily data, which in terms of frequency are closer to the HF data, deliver similar properties of the ATSM model.

action and statement (see also Campbell, Evans, Fisher, and Justiniano, 2012). In this case, the shocks are identified from HF changes in a spectrum of fed funds rate futures with maturities up to a year. Under the assumption that term premia for such a short horizon are small, the fed funds rate futures reflect expectations of the policy rate for the coming year. In this approach, the statement does not affect the current rate but captures any changes in expectations for the policy rate one year ahead, not inferred from the action itself.[10] We extend this approach to information contained in the entire yield curve (up to 10-year maturity). This is possible due to the ATSM, which allows us to extract expectations separately from term premia, while avoiding the problems, in this task, inherent in a VAR. Two orthogonal instruments (action and expected path) are extracted from the expectations part of the yield curve. Unlike action, the expected path component is restricted not to affect the current short rate. The third orthogonal instrument (uncertainty) is obtained from term premia. This instrument affects neither the current short rate nor its expected future path and can be interpreted as any residual uncertainty surrounding future monetary policy not already inferred from the other two components.[11] Term premia and uncertainty in our framework are thus closely related. The three instruments have very different loadings on the HF changes in yields: action has a declining pattern across maturities, expected path has a tent-like pattern with a peak at the 2-year horizon, and uncertainty has an increasing pattern. To provide support for the economic interpretation of the components, we compare the first two components to those obtained by previous studies from fed funds rate futures (Gürkaynak et al., 2005b) and the third component to implied and estimated interest rate volatility.[12]

[10]For instance, the FOMC may surprise markets by a wording that makes bond traders revise their expectations about future monetary policy, even when there is no surprise in the action.

[11]As expectations and term premia in an ATSM can be correlated, the uncertainty instrument is obtained from the part of term premia orthogonalised with respect to the two expectations components.

[12]Swanson (2021) also uncovers three components of monetary policy surprises. However, in each subsample of his analysis only two components are operative: target and path (which incorporates both expectations and term premia) in the pre-2008 period and path and large scale asset purchases in the post-2008 period. Like here, his decomposition is based on the entire yield curve, but without separating expectations from term premia. Hanson and Stein (2015), Gertler and Karadi (2015) and Nakamura and Steinsson (2018), in contrast, summarise multidimensional monetary policy surprises by a single factor, based on either fed funds rate futures or one-year or two-year government bond rates.

The interpretation of the three instruments is derived solely from their HF effects on the yield curve. Further structural content of the shocks they identify is based on the responses of macro and financial variables in the LP model. The effects of the shock identified by action are consistent with a standard monetary policy shock in a New-Keynesian model, including its extensions with the financial accelerator (Bernanke, Gertler, and Gilchrist, 1999) and time-varying term premia (Rudebusch and Swanson, 2012). The shock identified by the expected path component is associated with a strong response of interest rate expectations and produces responses of other variables that are consistent with both the Fed information effect (Nakamura and Steinsson, 2018) and the Fed response to news channel (Bauer and Swanson, 2020).[13] Finally, the responses to the uncertainty component are a little less clear-cut to map into existing theories. We propose a hypothesis, based on the LP, that could be explored in future research. In the data, term premia and various measures of monetary policy uncertainty increase in response to the shock. The effect on output, however, is mixed and we ascribe it to a fall in excess bond premium (Gilchrist and Zakrajšek, 2012), a variable capturing tightness in the corporate credit market. Specifically, an increase in the term premium increases the 30-year mortgage rate. New home sales and demand for mortgages decline, thus possibly allowing more credit to flow to the corporate sector. This effect may be counteracting any negative effect of uncertainty on output.

For all three instruments, our analysis uncovers a particularly tight connection between monetary policy and the housing market. Regardless of the shock, an increase in the 10-year bond yield, no matter whether occurring due to expectations or term premia, is associated with a similar increase in the 30-year mortgage rate and a sharp contraction in the housing market (new home sales and house prices).

HF intra-day data have been increasingly used to study various phenomena. Besides the context most directly related to us, the literature can be divided into two mutually non-exclusive categories: yield curve decomposition (including real and inflation components)

[13] The local projections alone cannot discriminate between the two mechanisms. In an Online Appendix we show that the instrument extracted from a model that is subject to the small sample bias is unable to identify this shock.

and identification of shocks. The first category includes, for instance, Beechey (2007), Beechey and Wright (2009), Bauer (2015), Gertler and Karadi (2015), Hanson and Stein (2015), and Hördahl, Remolona, and Valente (2015). Daily data are sometimes also used (Abrahams, Adrian, Crump, Moench, and Yu, 2016). Some studies employ ATSMs, while others use regressions. The second category includes, eg, Bernanke and Kuttner (2005), Miranda-Agrippino and Ricco (2015), Nakamura and Steinsson (2018), Cieslak and Schrimpf (2019), Jarocinski and Karadi (2020), and Bauer and Swanson (2020).[14] In terms of the housing market, a subset of our findings is consistent with those reported by Hamilton (2008), who follows a different methodology.[15]

The paper proceeds as follows. Section 2 discusses the HF data, Section 3 introduces the ATSM and the necessary notation, Section 4 provides an overview of the estimation method and the restrictions imposed, Section 5 reviews the estimates, applies the model to the HF data, and carries out the LP analysis. Finally, Section 6 concludes. Robustness checks and technical details related to the estimation are included in an Online Appendix.

2 High-frequency data

In order to study the HF yield curve reactions, we measure yields at various maturities in a narrow window around FOMC announcements. In doing so, we build on the literature studying monetary policy shocks within the HF approach. As noted in the Introduction, this literature focuses on short maturities, whereas we explore the reaction of the entire yield curve. Our HF data source is *Refinitiv Tick History*, except the 3-month T-bill rate, which has substantial gaps in the database at the required dates; the 3-month T-bill rate series was kindly provided to us by Refet Gürkaynak. As in the earlier literature, the changes in yields

[14]Swanson (2021) contains references for studies that, unlike us, focus on the post-2008 zero-lower bound period.

[15]A part of the literature, Kim and Orphanides (2012) being an early example, complements yield curve data with surveys of professional economists as a source of data for expected future interest rates. To keep the paper focused on the improvement of the estimation of the yield curve components relative to VARs, we confine ourselves only to yield curve data.

are measured in a 30-minute window starting 10 minutes before and ending 20 minutes after the announcement.

We focus on the period January 1996-August 2007, characterized by conventional monetary policy. In some cases, the data could be scarce, especially in the 1990s, with only a dozen of intra-day observations available. Therefore, for a few announcement dates our window has to be wider than 30 minutes. Despite this, the estimated changes in rates are similar to those reported by other studies (eg, Miranda-Agrippino, 2016). At the beginning of the sample, Treasury bonds with maturities longer than 10 years were traded relatively infrequently. Therefore, our longest maturity is based on the 10-year Treasury yield series. At medium-term maturities, Treasuries were not as frequently traded as LIBOR-based swaps (especially in the 1990s). Hence we faced a trade-off between having the same instrument but captured at different times due to relative illiquidity, or having all rates captured at the same time but taken from similar rather than the same instrument. We chose the latter and estimated the HF changes at 2-, 3-, and 5-year maturities from LIBOR-based swaps, which enabled us to create consistent narrow windows around the announcements. As noted above, at the short end, we use the change in the 3-month Treasury bill rate.[16]

The observed changes across the various maturities around the announcements are shown in Figure 1, which displays a consistent response pattern across all maturities. Table 1 presents basic statistics for the responses across maturities. Several observations follow. First, during the sample period, monetary policy surprises were slightly negative on average, with the shortest maturities affected the most and the impact declining with maturity. Second, all maturities display a strong reaction to the announcements, with the largest volatility occurring at the 2- and 3-year maturity.[17] Third, the yield curve tends to respond to the an-

[16]To analyse the behaviour of the yield curve around the announcements in a systematic way, we constructed a consistent yield curve across all maturities, adjusting for observed daily LIBOR spreads. We do this by estimating the spreads between LIBOR swap rates and the corresponding maturity yields observed at the close of business on the pre-announcement dates and then apply them to LIBOR rates around the announcements. The 3-month Treasury bill rate is left unadjusted, as the available data are already measured as a change in the rate.

[17]While the maximum response at the 2- and 3-year horizon persists across various splits of the sample, the relative volatility of the 3-month vs. 10-year maturity has changed towards the end of our sample. In the subsample 1996-2003, the standard deviation of the 3-month T-bill rate was 5.5 vs. 4.3 for the 10-year

nouncement in a consistent way, as indicated by the positive correlations between reactions across maturities, although the correlations are declining with maturity. Interestingly, the responses are highly correlated across medium and long maturities, with all the correlations between them being around 0.9.

3 The ATSM framework

The aim of this section is to provide a brief overview of the ATSM and introduce concepts and notation used in the rest of the paper. An underlying assumption behind an ATSM is the fundamental principle of finance, applied to default-free zero-coupon bonds of different maturities. Specifically,

$$E_t\left[M_{t+1}R_{t+1}^{(j)}\right] = 1, \tag{1}$$

where the expectation operator is with respect to information in period t, the scalar $M_{t+1} > 0$ is a kernel that prices all bonds and $R_{t+1}^{(j)}$ is a one-period gross return on a bond of any maturity j. That is, $R_{t+1}^{(j)} = P_{t+1}^{(j-1)}/P_t^{(j)}$, where $P_t^{(j)}$ is the price in period t of a bond of maturity j, which becomes a bond of maturity $j-1$ one period later. Of course, $P_t^{(0)} = 1$, as one dollar today has a value of one dollar.

ATSMs assume a specific functional form for the pricing kernel

$$-\log M_{t+1} = r_t + \frac{1}{2}\lambda_t^{'}\lambda_t + \lambda_t^{'}\varepsilon_{t+1}. \tag{2}$$

The popularity of this functional form lies in its practicality: when combined with the state space described below, it leads to a convenient affine solution for yields satisfying the no-arbitrage condition (1). Here, r_t is the continuously compounded short-term nominal interest rate, λ_t is a $N \times 1$ vector of risk prices for N underlying risk factors, and ε_{t+1} is a $N \times 1$ vector of innovations specified below. The N risk factors summarise the state space and are

bond. This has reversed to 2.0 vs. 4.26 in the period 2003-2007.

assumed to follow a first-order Gaussian VAR

$$X_t = \mu + \Phi X_{t-1} + \Sigma \varepsilon_t, \tag{3}$$

with $\varepsilon_t \sim N(0, I_N)$. This VAR process generates a '$\mathbb{P}$-measure' and the implied dynamics are referred to as the '$\mathbb{P}$-dynamics'.

Both the short rate and the risk prices are assumed to be related to the N factors through affine mappings

$$r_t = \delta_0 + \delta_1' X_t, \tag{4}$$

$$\lambda_t = \Sigma^{-1}(\lambda_0 + \lambda_1 X_t), \tag{5}$$

where δ_0 δ_1, Σ^{-1}, λ_0, and λ_1 are commensurate to the variables. In particular, λ_1 is $N \times N$. That is, the risk price of a particular factor can be affected by all factors. Observe that under risk neutrality (zero risk prices), the pricing kernel is simply $M_{t+1} = \exp(-r_t)$. That is, future cash flows are discounted with the short rate. Equations (1)-(5) summarize the ATSM.

Starting with $P_t^{(0)} = 1$, the model can be solved recursively for equilibrium bond prices (see, eg, Gürkaynak and Wright, 2012).[18] The vector of any J yields, $\hat{Y}_t$, can be written as

$$\hat{Y}_t = A + BX_t, \tag{6}$$

where $\hat{Y}_t$ is a $J \times 1$ vector. Equation (6) describes the model-implied yield curve—the cross-section of yields at a point in time that is consistent with no-arbitrage. (In an empirical implementation of the model, model-implied yields can potentially differ from observed yields due to measurement error and the lack of fit.) The arbitrage-free loadings A and B are

[18] Given the functional assumptions on the pricing kernel and the state space, the solution is an affine mapping from the factors to the logarithm of bond prices. Continuously compounded yields can then be inferred from the bond prices through standard discounting, $P_t^{(j)} = \exp(-jy_t^{(j)})$, which can be inverted to obtain yields as $y_t^{(j)} = (-1/j) \log P_t^{(j)}$. Yields are thus also affine in factors. For $j = 1$, we get the short rate: $y_t^1 = r_t$.

non-linear, recursive, functions of the model parameters δ_0, δ_1, λ_0, λ_1, μ, Φ, and Σ (see, eg, Gürkaynak and Wright, 2012). It can then be shown that the coefficients A and B in equation (6) are unaffected by switching to risk neutral pricing, $M_{t+1} = \exp(-r_t)$, and a risk-adjusted law of motion for the risk factors

$$X_t = \mu^{\mathbb{Q}} + \Phi^{\mathbb{Q}} X_{t-1} + \Sigma \varepsilon_t, \tag{7}$$

where

$$\mu^{\mathbb{Q}} = \mu - \lambda_0 \quad \text{and} \quad \Phi^{\mathbb{Q}} = \Phi - \lambda_1. \tag{8}$$

The VAR process (7) is referred to as the '$\mathbb{Q}$-measure', describing the '$\mathbb{Q}$-dynamics'. That is, dynamics under risk neutral pricing. Observe that under risk neutral pricing, the model is parameterised in terms of δ_0, δ_1, $\mu^{\mathbb{Q}}$, $\Phi^{\mathbb{Q}}$, and Σ. Thus, to derive the cross-sectional implications of the model summarized by equation (6), all that is required is the $\mathbb{Q}$-measure. The knowledge of the $\mathbb{P}$-measure and the risk prices λ_t is not required. To put it differently, the cross-section identifies the parameters of the $\mathbb{Q}$-measure, not the $\mathbb{P}$-measure.

Under the $\mathbb{Q}$-measure, the expected value of the short rate j periods ahead can be obtained from the short rate equation (4) and the VAR process (7). The effect of X_t on the expected value is given by $(\Phi^{\mathbb{Q}})^j$. The effect of X_t on the average expected short rate over the forecast horizon j under the $\mathbb{Q}$-measure is thus

$$B_j = \frac{1}{j} \delta_1' \left[I + \Phi^{\mathbb{Q}} + \cdots + (\Phi^{\mathbb{Q}})^{j-1} \right], \tag{9}$$

which is the jth row in the loading matrix B in equation (6). Under the $\mathbb{P}$-measure, the expected value of the short rate j periods ahead can be obtained from the short rate equation (4) and the VAR process (3). In this case, the effect of X_t on the expected value is given by Φ^j and the average expected short rate over the forecast horizon is given by

$$B_j^{\mathbb{P}} = \frac{1}{j} \delta_1' \left[I + \Phi + \cdots + \Phi^{j-1} \right]. \tag{10}$$

The difference $C_j \equiv B_j - B_j^{\mathbb{P}}$ is the effect of X_t on the term premium in yield $y_t^{(j)}$ and, as follows from the relationship (8), depends on λ_1. (Observe that $B = B^{\mathbb{P}} + C$.) Thus, while the knowledge of the $\mathbb{P}$-measure is not required for the cross-sectional implications of the model, it is necessary for deriving a decomposition between term premia and expected interest rates. Observe that the $\mathbb{P}$-measure can be identified either from the time-series of X_t and equation (3) or the cross-section of yields and the knowledge of λ_0 and λ_1 through the relationship (8).

4 Estimation of the ATSM

This section provides an overview of the estimation method and the restrictions imposed. All technical details are contained in the Online Appendix.

4.1 The importance of restrictions

In principle one could estimate a VAR on yields, and possibly macro variables, and then iterate it forward j times to obtain forecasts of the short rate between now and the jth period ahead, thus obtaining the expectations and term premium components for the jth maturity (eg, Gertler and Karadi, 2015). There are two problems with this approach. First, the VAR-based forecasts of future yields of different maturities may imply arbitrage opportunities. Second, nominal interest rates are highly persistent, which, in samples of the length typically observed, leads to both a downward bias in the persistence of the VAR process and high standard errors of its estimates. This problem arises because we do not observe frequent enough mean reversions of interest rates in the data to estimate the parameters of the driving process well.[19]

By construction, ATSMs resolve the first issue. ATSMs can also resolve the second issue, but only if they are estimated subject to restrictions. As ATSMs are estimated on both time

[19]See the classic results of Kendall (1954), Nicholls and Pope (1988), and Shaman and Stine (1988) and, for a discussion in the context of ATSMs, Bauer et al. (2012). As demonstrated by Pierse and Snell (1995), increasing the sampling frequency does not resolve the problem.

series and cross-sectional data, they use more information than a VAR. In particular, the cross section of yields at a point in time can *potentially* provide very precise information for the model dynamics.[20] However, the cross-section identifies only the parameters of the $\mathbb{Q}$-measure and to arrive at a decomposition into term premia and expectations, the knowledge of the parameters of either risk prices or the $\mathbb{P}$-measure is required. Further, Joslin et al. (2011) demonstrate that in a canonical ATSM—the maximally flexible model that is subject only to normalizing restrictions—the cross-sectional data convey no information for the estimation of the other parameters (see also Hamilton and Wu, 2012). As a result, the $\mathbb{P}$-dynamics are solely estimated from time series data and the estimates of expected interest rates (and thus term premia) are equivalent to those obtained from a simple VAR.[21] To improve the estimates relative to a VAR, the ATSM is estimated subject to restrictions to correct for the downward bias in the underlying VAR.

4.2 Model nomenclature: $\mathcal{M}_0$, $\mathcal{M}_1$ and $\mathcal{M}_2$

To ensure identification, we employ the normalising restrictions of Joslin et al. (2011), leading to their canonical representation. Under this representation, the N risk factors are defined as linear combinations of yields, $X_t = W\hat{Y}_t$, where W is a weighting matrix, and the model parameters are mapped into a set of unknowns $k^{\mathbb{Q}}$, $\phi^{\mathbb{Q}}$, μ, Φ, and Σ, which fully characterize the $\mathbb{P}$- and $\mathbb{Q}$-dynamics, (μ, Φ) and $(\mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}})$ respectively. Here, $k^{\mathbb{Q}}$ determines the mean of the short rate under the $\mathbb{Q}$-measure and $\phi^{\mathbb{Q}}$ is a $N \times 1$ vector that contains the eigenvalues of $\Phi^{\mathbb{Q}}$. Following Joslin et al. (2011), the risk factors X_t are calculated as the first $N < J$ PCs of the yields and W is the associated $N \times J$ loading matrix. Finally, the observed yields Y_t are assumed to be measured with error: $Y_t = \hat{Y}_t + e_t$. Under the assumption that X_t is observed in the estimation (ie, N linear combinations of yields using the weights W are

[20] To illustrate this, suppose investors were risk neutral (ie, prices of risk were equal to zero) and so observed yields were equal to expected future interest rates. Then one could simply read off expected future interest rates from the cross-section, thus avoiding the problematic time series data altogether.

[21] Conceptually, the Gertler and Karadi (2015) results can thus be viewed as estimated from an unrestricted ATSM.

estimated exactly by the model), the $J - N$ independent measurement errors are normal with variance σ_e^2. We use $N = 4$, with the first three PCs resembling the standard level, slope and curvature factors. This choice is motivated by studies arguing that more than three factors are needed to properly capture the term structure (Cochrane and Piazzesi, 2008; Christensen, Diebold, and Rudebusch, 2009; Adrian, Crump, and Moench, 2013). To accommodate this viewpoint, while maintaining parsimony, we proceeded by testing $N = 3$ vs. $N = 4$. The RMSE criterion prefers $N = 4$ (eg, for the model $\mathcal{M}_1$, RMSE=4.39, as opposed to 7.35, when the fourth factor is dropped).

We estimate three versions of the model. Model $\mathcal{M}_0$ is the maximally flexible benchmark that is only subject to the Joslin et al. (2011) normalising restrictions. As a result, the estimates of the parameters of the $\mathbb{P}$-measure (μ, Φ) are based only on time-series data.[22] Model $\mathcal{M}_1$ places zero restrictions on λ_0 and λ_1.[23] To impose such restrictions, we use a stochastic search variable selection (SSVS) algorithm employed by Bauer (2018). It is clear from equation (8) that setting some risk prices to zero has the effect of 'pulling up' the VAR parameters μ and Φ towards $\mu^{\mathbb{Q}}$ and $\Phi^{\mathbb{Q}}$, thus ameliorating the small sample bias. Model $\mathcal{M}_2$ is based on the analysis of Bauer et al. (2012), who propose a statistical method to estimate and correct the small sample bias in μ and Φ. In this case, the model is estimated subject to the restriction that, assuming it is the data-generating process, it produces the same small sample bias as in the data. As a result, this procedure increases the persistence of the VAR under the $\mathbb{P}$-measure, relative to model $\mathcal{M}_0$.

4.3 Bayesian procedure

In the Joslin et al. (2011) canonical representation, the likelihood function factors into two components

$$f\left(Y_t|Y_{t-1}, \Theta\right) = f\left(Y_t|X_t, \phi^{\mathbb{Q}}, k^{\mathbb{Q}}, \Sigma, \sigma_e^2\right) \times f(X_t|X_{t-1}, \mu, \Phi, \Sigma), \tag{11}$$

[22] The parameters of risk prices are then obtained residually as $\lambda_0 = \mu - \mu^{\mathbb{Q}}$ and $\lambda_1 = \Phi - \Phi^{\mathbb{Q}}$.

[23] This strategy has been implemented, in various forms, by Cochrane and Piazzesi (2008), Duffee (2011), Joslin et al. (2011), Joslin, Priebsch, and Singleton (2014), and Bauer (2018).

where $\Theta = \left(\phi^{\mathbb{Q}}, k^{\mathbb{Q}}, \Sigma, \sigma_e^2, \mu, \Phi\right)$ denotes the parameters to be estimated. Note that the first term in this factorisation is the '$\mathbb{Q}$-likelihood', as it incorporates information from the cross-section of yields. In contrast, the second term is the '$\mathbb{P}$-likelihood', based on information derived from the time-series of the risk factors.[24]

We employ a Bayesian approach to estimate the three versions of the model, using the Gibbs sampling algorithm proposed by Bauer (2018). The Bayesian approach is particularly useful as it provides a systematic and efficient method to impose restrictions on μ and Φ (or equivalently on λ_0 and λ_1) in the likelihood function (11). This means that there is no need to carry out an explicit model comparison exercise that can involve estimation of a large number of restricted specifications. Moreover, maximisation of the likelihood of the ATSM is a non-trivial task that is made even more challenging by the small sample of the typical data set.[25]

4.4 Data for the ATSM estimation

The three versions of the model are estimated on monthly data for yields at maturities of 1, 3 and 6 months and 1 through 10 years. That is, thirteen maturities in total. The data at maturities of one year and above are obtained from the Federal Reserve Board database on the nominal yield curve (the Gürkaynak-Sack-Wright data set), with rates at shorter maturities taken from the FRED database. The sample runs from January 1990 to December 2008.[26]

[24] As Joslin et al. (2011) show, the fact that the two likelihoods share Σ does not affect the estimates of the other parameters.

[25] Bayesian estimation does not rely on maximisation of the likelihood function and, instead, aims to approximate the joint posterior distribution of the model parameters. MCMC methods make this task easy by working with the two conditional distributions associated with the joint posterior. Finally, as the Bayesian approach approximates the posterior distribution, error bands for parameter estimates are obtained directly. In contrast, frequentist approaches rely on asymptotic standard errors that may be inaccurate in small samples; bootstrap methods in the ATSM case have high computational costs.

[26] As noted in the Introduction, for robustness, the Online Appendix reports estimates obtained also on daily data.

5 Results

The results are presented in the following steps: (i) we inspect the impact of the restrictions on the estimated models (Section 5.1), (ii) extract and analyse three main components of monetary policy surprises from the HF data (Sections 5.2-5.4), and (iii) use the components as instruments in a local projections model (Section 5.5).

5.1 Inspecting the estimated ATSMs

All models display a good fit to the monthly data with root mean squared errors that are below five basis points.[27] Figure 2 shows the estimated posterior distributions of the largest eigenvalues of $\Phi^{\mathbb{Q}}$ and Φ. Under the $\mathbb{Q}$-measure, the three models have a very similar profile in terms of persistence. This, of course, is expected as the estimates are based on the same cross-sectional information and the partial likelihoods for the $\mathbb{Q}$-measure differ across the models only in terms of Σ. The results, however, are very different for the eigenvalues under the $\mathbb{P}$-measure. The maximally flexible model has the lowest median persistence out of the three models and the widest posterior distribution. Restrictions on risk prices or statistical bias correction thus lead to a substantial increase in persistence, as well as in the precision of the estimates.[28] As a result, at the median estimates, in $\mathcal{M}_1$ and $\mathcal{M}_2$ the volatility of expectations is about twice as high as in $\mathcal{M}_0$ and roughly at par with the volatility of term premia.

[27]The estimates of the $\mathbb{P}$ and $\mathbb{Q}$ parameters, and the implied λ's, are shown in the Online Appendix. There we also report results for the daily models, as well as additional results for the monthly models, including plots of the time series of expectations and term premia over time and their correlations with economic activity over the business cycle. Regarding the latter, here we only note that in $\mathcal{M}_1$ and $\mathcal{M}_2$ the 10-year term premium is counter-cyclical, whereas in $\mathcal{M}_0$ it is uncorrelated with the business cycle.

[28]To illustrate this, take the median values to the power of 120 to derive their effect on expected interest rates ten years ahead. This exercise results in 0.12, 0.41, and 0.63 percentage point increase in the nominal short rate in ten-years time for the three models respectively, for one percentage point increase in the current short rate.

5.2 High-frequency yield curve decomposition

The decomposition is based on the median estimates. Let $\Delta\tilde{X}_t$ denote the vector of changes in the risk factors in the 30-minute window around FOMC announcements. $\Delta\tilde{X}_t$ is obtained as the first four PCs of the changes in yields in that window. Recall that the HF data are for maturities of 3 months, and 2, 3, 5, and 10 years. The HF changes in term premia and expectations are then computed using $\Delta\tilde{X}_t$ and the estimated models.[29]

In terms of the notation of Section 3, the vector of changes in the expectations component, for the five maturities, is given by $\Delta\tilde{Y}_t^E = B^{\mathbb{P}}\Delta\tilde{X}_t$ and the vector of changes in term premia is given by $\Delta\tilde{Y}_t^{TP} = C\Delta\tilde{X}_t$, where $B^{\mathbb{P}}$ and C are derived from the parameters of the estimated models as described in Section 3. We also derive $\Delta\tilde{Y}_t = B\Delta\tilde{X}_t$, where $\Delta\tilde{Y}_t$ is a vector of changes in the fitted HF yields and $B = B^{\mathbb{P}} + C$. An implicit assumption in using the estimated ATSM for the HF decomposition is that a model estimated on monthly data is suitable to describe the yield curve at the HF. The fact that the models estimated on daily data (see the Online Appendix) have similar properties as the monthly models gives us confidence that this assumption, for our purposes, is reasonable.[30]

Figure 3 provides a summary of the movements of the yield curve around FOMC announcements explained by the three models. It plots the volatility curve of the HF changes in expectations and term premia across maturities (refer back to Table 1 for the volatility of the changes in the observed yields). The figure demonstrates that imposing restrictions on the estimated ATSM increases the reaction of expected future interest rates to FOMC announcements. While in the unrestricted model $\mathcal{M}_0$, term premia at the 10-year horizon, for instance, are significantly more volatile than expectations, the relative volatility is reversed in model $\mathcal{M}_1$ and in model $\mathcal{M}_2$ the variance of the two components is roughly at par.[31]

[29]Given that the set of maturities in the HF dataset is only a subset of the maturities used to estimate the models, one may wonder how different the estimated parameters of the ATSMs would be if only the maturities of the HF dataset were used in the estimation. It turned out that the estimates are almost identical. The maturities in the HF dataset thus seem to capture all of the main movements in the yield curve over time.

[30]The root mean squared error of the fit of the models at the HF is about three basis points across all models (monthly and daily), comparable to their fit at the monthly and daily frequencies.

[31]Recall that term premia and expectations can be correlated. The variances of the two components thus

Before moving on to the next stage, the HF reaction of the 3-month maturity (the shortest maturity at our disposal at the HF) deserves attention. In Figure 3, all three models display a standard deviation of term premia at the 3-month maturity of about one basis point. In the estimated ATSM, the shortest maturity is one month. At that maturity, risk premia are zero. There are, however, some nonzero elements in the C matrix at the 3-month maturity, especially for the fourth risk factor. The variation in risk premia at the 3-month maturity observed in Figure 3 occurs due to a few data points in the HF sample at which the typically unimportant fourth risk factor had an unusually large realisation. However, the restrictions imposed below effectively eliminate the effects of these sporadic events from the subsequent analysis.

5.3 Instruments for policy shocks

The instruments are obtained in three steps. First, we decompose the HF changes in expectations into PCs and select the most important PCs. Second, we orthogonalise term premia with respect to the selected PCs of expectations. The PCs of expectations are mutually orthogonal by definition. However, expectations and term premia (and thus their respective PCs) can be correlated. The second step addresses this correlation, leaving us with movements of term premia that are orthogonal to the PCs of expectations. We then carry out a PC decomposition of the part of term premia that is orthogonal to the PCs of expectations. Finally, in the third step, we apply a particular orthogonal rotation to the PCs of expectations and the PCs of the above part of term premia to assign them an economic interpretation.

Formally, recall that $\Delta\tilde{Y}_t^E = B^{\mathbb{P}}\Delta\tilde{X}_t$, where $\Delta\tilde{X}_t$ is the HF change in the N risk factors ($N = 4$), $\Delta\tilde{Y}_t^E$ has a dimension $\widetilde{J} \times 1$ ($\tilde{J} = 5$), and $B^{\mathbb{P}}$ is determined by the parameters of the estimated model. A PC decomposition of expectations returns: $\Delta\tilde{Y}_t^E = \Omega^E\mathcal{P}_t^E$. Here, Ω^E is a $\tilde{J} \times N$ loadings matrix and $\mathcal{P}_t^E$ are the corresponding PCs. The dimension of the

do not necessarily add up to the variance of the respective yield. The solid lines at the bottom of the charts in Figure 3 plot the correlation at a given maturity.

PCs is equal to N, as the changes in expectations are constructed from N risk factors.[32] We select the first $N_E \leq N$ most important PCs. Thus, $\Delta\tilde{Y}_t^E \approx \Omega_1^E \mathcal{P}_{1t}^E$, where the subscript "1" refers to the selected PCs and their corresponding loadings matrix, which is a partition of Ω^E. The approximation sign denotes the fact that we are not using all but only the most important PCs. To ensure orthogonality of the PCs of term premia, with respect to the selected PCs of expectations, we run the following regression for each $j = 1, ..., \tilde{J}$

$$\Delta\tilde{Y}_{jt}^{TP} = \alpha_j + \beta_j^\top \mathcal{P}_{1t}^E + \xi_{jt}, \tag{12}$$

where $\Delta\tilde{Y}_{jt}^{TP}$ is the jth element of $\Delta\tilde{Y}_t^{TP} = C\Delta\tilde{X}_t$, with C determined by the parameters of the estimated model. Let $\Delta\tilde{y}_t^{TP} = [\xi_{1t}, \ldots, \xi_{\tilde{J}t}]^\top$ collect the parts of term premia that are orthogonal to the selected PCs of expectations. We then carry out a PC decomposition of $\Delta\tilde{y}_t^{TP}$, retaining only the first $N_{TP} \leq N$ most important PCs, denoted by $\mathcal{P}_{1t}^{TP}$. Thus, $\Delta\tilde{y}_t^{TP} \approx \Omega_1^{TP}\mathcal{P}_{1t}^{TP}$. This procedure leaves us with a vector of mutually orthogonal components of the HF changes in expectations and term premia, $[\mathcal{P}_{1t}^E, \mathcal{P}_{1t}^{TP}]^\top$. Orthogonal matrixes Q^E and Q^{TP}, which have dimensions $N_E \times N_E$ and $N_{TP} \times N_{TP}$, respectively, are then applied to $\mathcal{P}_{1t}^E$ and $\mathcal{P}_{1t}^{TP}$, respectively, producing new components $\mathcal{P}_{1t}^{E*} \equiv Q^E\mathcal{P}_{1t}^E$ and $\mathcal{P}_{1t}^{TP*} \equiv Q^{TP}\mathcal{P}_{1t}^{TP}$. The associated loadings for $\Delta\tilde{Y}_t^E$ and $\Delta\tilde{y}_t^{TP}$ of these rotated components are, respectively: $\Omega_1^{E*} \equiv \Omega_1^E(Q^E)^{-1}$ and $\Omega_1^{TP*} \equiv \Omega_1^{TP}(Q^{TP})^{-1}$. That is, $\Delta\tilde{Y}_t^E \approx \Omega_1^{E*}\mathcal{P}_{1t}^{E*}$ and $\Delta\tilde{y}_t^{TP} \approx \Omega_1^{TP*}\mathcal{P}_{1t}^{TP*}$.

The rotated components $[\mathcal{P}_{1t}^{E*}, \mathcal{P}_{1t}^{TP*}]^\top$ are the instruments. By construction, they are orthogonal to each other. By imposing the rotation on the PCs of expectations and orthogonalised term premia, we are implicitly imposing a rotation on the underlying risk factors. Working with the PCs of expectations and term premia, however, is more intuitive and is closer to the practice in the literature.[33]

[32] A PC decomposition of $\tilde{J}$ time series returns $\tilde{J}$ PCs. However, as there are only $N < \tilde{J}$ risk factors, the remaining $\tilde{J} - N$ PCs have zero variance and can thus be ignored.

[33] The mapping between the two is as follows. Start with the fact that expectations can be expressed either in terms of their PCs or the risk factors. Thus, $\mathrm{var}(\Delta\tilde{Y}_t^E) = \Omega^E\Lambda^E(\Omega^E)^\top = B^{\mathbb{P}}\Lambda^X(B^{\mathbb{P}})^\top$, where $\Lambda^E = \mathcal{P}_t^E(\mathcal{P}_t^E)^\top$ is a diagonal covariance matrix and $\Lambda^X = \Delta\tilde{X}_t(\Delta\tilde{X}_t)^\top$ is also a diagonal covariance matrix,

5.4 Implementation and inspection of the instruments

We have shown that the restricted models $\mathcal{M}_1$ and $\mathcal{M}_2$ generate substantially stronger responses of expected interest rates to FOMC announcements than model $\mathcal{M}_0$. To economize on space, we therefore continue only with $\mathcal{M}_1$ (similar results for $\mathcal{M}_2$ are contained in the Online Appendix). Following the steps described above, the data suggest $N_E = 2$: the first two PCs of expectations account for 98.6% of the total variance of expectations across maturities, with the respective contributions of 87.6% and 11%. The orthogonalised term premia are explained by two PCs, which account, respectively, for 92% and 8% of their variance.[34]

The two PCs of expectations, $\mathcal{P}_{1t}^{E}$, and the two PCs of orthogonalised term premia, $\mathcal{P}_{1t}^{TP}$, are rotated to create $[\mathcal{P}_{1t}^{E*}, \mathcal{P}_{1t}^{TP*}]^\top$ such that Ω_1^{E*} and Ω_1^{TP*} have the following properties: (i) in Ω_1^{E*} the first element in the second column is equal to zero and (ii) in Ω_1^{TP*} the first element in the first column is equal to zero. This means that only the first component of $\mathcal{P}_{1t}^{E*}$ affects the 3-month T-bill rate; the second component of $\mathcal{P}_{1t}^{E*}$ does not. Also, the first component of $\mathcal{P}_{1t}^{TP*}$ does not affect the 3-month T-bill rate, while the second one does.[35]

The rotation of $\mathcal{P}_{1t}^{E}$ is based on Gürkaynak et al. (2005b). Given the above restrictions, the first component of $\mathcal{P}_{1t}^{E*}$ is interpreted as a surprise in *action*, while the second component is interpreted as a surprise in *expected path*. The expected path component captures any

as both $\mathcal{P}_t^E$ and $\Delta\tilde{X}_t$ are PCs. In general, however, $\Omega^E \neq B^{\mathbb{P}}$ and $\Lambda^E \neq \Lambda^X$ and $B^{\mathbb{P}}$ is not orthogonal ($B^{\mathbb{P}}$ is derived from the ATSM, not a PC decomposition). Nevertheless, there exists a $(N \times N)$ matrix H such that $\mathrm{var}(\Delta\tilde{Y}_t^E) = (B^{\mathbb{P}}H^{-1})H\Lambda^X H^\top (B^{\mathbb{P}}H^{-1})^\top$. We can thus relate $\mathcal{P}_t^E$ to $\Delta\tilde{X}_t$ as $\mathcal{P}_t^E = H\Delta\tilde{X}_t$, where $H = (\Omega^E)^\top B^{\mathbb{P}}$. The rotated PCs of expectations are thus related to the risk factors as $\mathcal{P}_t^{E*} = Q^E\mathcal{P}_t^E = Q^E H\Delta\tilde{X}_t$. When the rotation is applied only to a subset of $\mathcal{P}_t^E$, such as the N_E most important PCs, the relationship is $Q^E\mathcal{P}_{1t}^E = Q^E H_1\Delta\tilde{X}_t$, where Q^E is $N_E \times N_E$ and H_1 is the partition of H commensurate to the first N_E PCs. The same applies to the orthogonalised term premia.

[34] The first two PCs of raw term premia (ie, before they are orthogonalised with respect to the PCs of expectations) make up 78% and 21% of the total variance of raw term premia. Their correlations with the first two PCs of expectations are significantly different from zero only in the case of the second PC of expectations (around -0.45 for both PCs of term premia). These statistical relationships get picked up by the orthogonalisation regressions (12). The R^2s of the regressions for the five maturities are 0.30, 0.35, 0.32, 0.16, and 0.02, respectively. The regressions thus do not explain much of the term premia at the 5- and 10-year horizon. (The positive correlations between expectations and term premia at the 5- and 10-year horizon observed in Figure 3 are mainly due to the third PC of expectations.) As a result, term premia at the 5- and 10-year horizon are almost completely explained by the PCs of orthogonalised term premia.

[35] Before the rotation is applied, $[\mathcal{P}_{1t}^{E}, \mathcal{P}_{1t}^{TP}]^\top$ are normalised to have a unit standard deviation, a standard normalisation required for identification.

surprise in the FOMC announcement that affects expectations of future policy rates above and beyond the information already inferred from action. Following the same logic, we interpret the first component of $\mathcal{P}_{1t}^{TP*}$ (the one that does not affect the 3-month T-bill rate) as capturing any surprise in the FOMC announcement affecting perceived uncertainty surrounding the expected path of policy rates, not already inferred from the action and expected path components. We refer to it as a surprise in *uncertainty*.[36] A number of studies have established that central bank communication contains elements affecting uncertainty about future monetary policy, as perceived by financial markets. For instance, Swanson (2006) shows that improvements in Fed communication since the 1990s have substantially reduced policy rate uncertainty. Wright (2011) demonstrates that a decline in inflation uncertainty, achieved through advances in the monetary policy framework, has reduced term premia.[37]

There is no degree of freedom left to impose restrictions on the second component of $\mathcal{P}_{1t}^{TP*}$, whose presence in the analysis is a necessary consequence of the fact that there are four risk factors in the ATSM and term premia have been orthogonalised with respect to two PCs of expectations.[38] This "residual" component is thus free to affect the 3-month T-bill rate. Nevertheless, its contribution is visible only on a few occasions in the HF sample, as shown in the Online Appendix (these occasions are related to the unusually large realizations of the generally small fourth risk factor in the ATSM, see the discussion in Section 5.2). Its contribution to other maturities is equally small; at the 10-year horizon it is minuscule. We thus leave it out from the subsequent analysis and work with three instruments: action, expected path and uncertainty.

Table 2 reports the loadings of the five maturities in the HF dataset on the three in-

[36] Husted, Rogers, and Sun (2020) follow a similar orthogonalisation strategy with respect to action and expected path, although in a different framework, to isolate the marginal effect of an uncertainty factor.

[37] Blinder, Ehrmann, Fratzscher, de Haan, and Jansen (2008) provide a thorough review of an early literature on central bank communication, including its effects on monetary policy uncertainty; Tillmann (2020) contains a number of recent references on monetary policy uncertainty and term premia.

[38] The 2×2 rotation matrix Q^{TP} allows for only four restrictions. One is the orthogonality of the two components, other two impose normalised unit variance on the factors, and the fourth is the zero response of the 3-month T-bill rate to the first component.

struments (ie, it reports the first and second column of Ω_1^{E*} and the first column of Ω_1^{TP*}). The loadings are normalized relative to the loading at the two-year horizon to allow easy comparison with other studies. The patterns clearly differ across the three instruments. Action has a declining pattern across maturities, expected path has a tent-like pattern with a peak at the 2-year horizon, and uncertainty has an increasing pattern. Different components of policy announcements thus give rise to very different HF reactions of the term structure. Gürkaynak et al. (2005b) extract their two components, target and path from a spectrum of fed funds futures rates and regress the HF changes in longer maturities on the two components. Kuttner (2001) carries out the same exercise for a single target component, extracted from a single fed funds futures rate.[39] The declining pattern exhibited by our action component is consistent with the target component in both Kuttner (2001) and Gürkaynak et al. (2005b). Our expected path component has a similar pattern as the Gürkaynak et al. (2005b) path component.

The contributions of the three components to the HF changes in yields can be observed in Figure 4, for the 3-month, 5-year and 10-year maturities. By construction, only action affects the 3-month maturity. The contribution of this instrument declines with maturity. Expected path is important both at the 5- and 10-year maturity, while uncertainty has clearly the largest impact at the 10-year maturity. An interesting aspect of the figure is an apparent decline in the importance of uncertainty, and an increase in the importance of expected path, in contributing to the movements at the 10-year maturity from about 2001. This finding can be interpreted, at least partially, as being in line with the conclusions of Swanson (2006) and Wright (2011) that better Fed communication and transparency since the late-1990s have reduced monetary policy uncertainty.[40]

The bottom chart of Figure 4 shows that two (positive) realisations of the uncertainty

[39] There are some differences across the studies and ours in terms of the maturities and the period covered. Nevertheless, relative to other studies in the literature, the time span is quite similar, in the sense that it focuses on the pre-2008 period of conventional monetary policy.

[40] In principle, uncertainty could be about the underlying state of the economy, the transmission mechanism, and the response function of monetary policy with respect to the state of the economy. The evidence in Swanson (2006) and Wright (2011) concerns the last type of uncertainty.

component stand out: February 3, 1999 and January 3, 2001. The February 3, 1999, meeting was not accompanied by a statement (before May 1999, statements were not issued after every meeting). Based on the market commentary, a justification for the increase in the uncertainty component could be that the market was speculating if or when the Fed may embark on a tightening cycle, after the policy rate was cut on three occasions in the previous quarter in fear of a recession that did not materialise.[41] On January 3, 2001, FOMC cut the policy rate by 50 basis points, following a conference call, which came nearly four weeks ahead of the regularly-scheduled policy meeting. According to market commentary, this emergency meeting caught most investors off guard.[42]

To cross-check the economic interpretation of the instruments, Figure 5 compares the first two instruments with the target and path components derived from fed funds futures by Gürkaynak et al. (2005b), for the part of the sample where our and their sample overlap. The third instrument is compared with two popular proxies for monetary policy uncertainty: implied volatility from options on fed funds futures or swap rates (eg, Swanson, 2006; Wright, 2017) and estimated interest rate uncertainty (Jurado, Ludvigson, and Ng, 2015). These two proxies of uncertainty are for daily, rather than intra-day, changes bracketing the FOMC announcements. For implied volatility we use options on one-year swap rates; interest rate uncertainty is estimated as time-varying volatility of the forecast error in forecasts of the 3-month T-bill rate one year ahead.[43] Given that our instruments are derived from different data than any of the measures they are compared with, we would not expect perfect correspondence. Nevertheless, Figure 5 reports that in all four cases there is a statistically significant positive relationship, with the p-values in all but one case below 1% (below 5% in the remaining case).[44] As a caveat, the less then perfect correlation between our instruments

[41] Source: https://money.cnn.com/1999/02/03/economy/fed/.

[42] Source: https://money.cnn.com/2001/01/03/economy/fed/.

[43] Kaminska and Roberts-Sklar (2018) provide a list of various measures of monetary policy uncertainty proposed in the literature, including those based on computational linguistics and surveys. Most of these measures, however, are available only at monthly or lower frequency.

[44] Interestingly, the R^2 in the regression of the expected path component on the Gürkaynak et al. (2005b) path component raises from under 0.1 before 2000 to 0.35 in 2004, while the slope coefficient raises from 0.2 in 1996 to 0.65 in 2004 (the estimates are based on time-varying coefficient regression, using the Gibbs sampling algorithm proposed by Cogley and Sargent, 2002). It appears that as the Fed communication has

and the variables used to cross-check their interpretation suggests that the instruments may be picking up some additional mechanisms than those proposed by their interpretation. In particular, the uncertainty instrument may be picking up some other factors determining term premia movements around FOMC meetings, such as liquidity or the demand effects of yield oriented investors stressed by Hanson and Stein (2015).

5.5 Local projections

To estimate the dynamic impact of policy shocks on macroeconomic and financial variables of interest, we use Bayesian local projections, introduced by Miranda-Agrippino and Ricco (2015). The Bayesian approach addresses concerns regarding efficiency of standard LP estimates.[45] As in Jordà (2005), the model is

$$Z_{t+h} = c^{(h)} + B_1^{(h)} Z_t + \sum_{j=1}^{P} b_j^{(h)} Z_{t-j} + v_{t+h},$$

where Z_t is a vector of the M variables of interest, h is the impulse-response horizon and v_{t+h} denotes residuals. The impulse-responses for the shocks of interest at horizon h can be calculated as $B_1^{(h)} A_0$, where A_0 denotes the contemporaneous impact matrix. The contemporaneous impulse-responses in a LP are equivalent to those in a VAR (see Miranda-Agrippino and Ricco, 2015). A column of the A_0 matrix corresponding to a given shock can thus be estimated from residuals of a VAR (in Z_t) and a HF instrument using the method of Mertens and Ravn (2013). The three HF instruments identify three different contemporaneous responses at the monthly frequency (columns of A_0), $A_{0,k}$, $k = 1, 2, 3$. The dynamic impulse-responses in the LP model are then computed as $B_1^{(h)} A_{0,k}$.

The LP model is estimated on monthly data for 1990-2007, a period typical for studies that focus on conventional monetary policy, using twelve lags as controls.[46] The benchmark

improved over time, the information content about the expected future path of policy rates obtained from different markets got more aligned.

[45] Technical details and sensitivity analysis are contained in the Online Appendix.

[46] As in Miranda-Agrippino and Ricco (2015), the prior distributions are set using a training sample, which

model has the following variables: log of industrial production, the CPI inflation rate, the Gilchrist and Zakrajšek (2012) excess bond premium (EBP), and the first two PCs of yields that were used as risk factors in the ATSM. The first three variables are standard in the empirical macro literature.[47] The first two PCs are included as summary statistics for the responses of the yield curve (they account for 99% of the total variation in yields across maturities at the monthly frequency). The responses of the short rate and the 10-year yield are then obtained by multiplying the responses of the two PCs with the PC loadings corresponding to these two interest rates (adding the third and fourth PCs did not affect the LP results in any substantial way). Then additional variables are added one by one, including their twelve lags as controls:[48] the 30-year mortgage rate, implied volatility used in Section 5.4, the Husted et al. (2020) monetary policy uncertainty index (MPU)[49], the log of S&P 500, the log of real house prices, the log of new single-family home sales, and the estimated 10-year expectation and term premium components.[50]

Figures 6-8 report the findings. The responses to the shock identified by the action instrument (Figure 6) appear to be broadly consistent with responses to a standard Taylor rule shock in a New-Keynesian (NK) model. Industrial production declines and inflation also exhibits a declining tendency. As in a version of the NK model with a financial accelerator (eg, Bernanke et al., 1999), EBP rises. The 10-year bond yield at the monthly frequency initially declines. A number of NK models in which the implicit inflation target is not constant have this property. In, eg, Gürkaynak et al. (2005a) the decline occurs due to expectations, whereas in Rudebusch and Swanson (2012) it is due to term premia. The

spans the period 1982-1989.

[47] The excess bond premium is the component of the spread between an index of rates of return on corporate securities and a similar maturity government bond rate that is left after the component due to default risk is removed. It is typically interpreted as a measure of tightness in the credit market for non-farm business sector.

[48] In principle, this can change the responses of the original variables, but in practice the responses remained similar. The alternative is to have a larger set of variables from the outset, but this is difficult from a computational view point due to the relatively small sample size.

[49] This is a broader measure of monetary policy uncertainty than implied volatility, derived from media analysis using computational linguistics.

[50] Except the excess bond premium, implied volatility, the MPU index, and the yield curve data, the data come from either FRED or Haver.

responses in Figure 6 give more support to the latter. The observed decline in the term premium is accompanied also by initial reductions in monetary policy uncertainty, exhibited by both proxies used. The S&P500 falls, which is consistent with the standard discount factor channel (Bernanke and Kuttner, 2005).[51] Finally, the 30-year mortgage rate essentially mimics the 10-year bond yield. The housing market variables (house prices and new home sales) in turn mimic the mortgage rate, but with a negative sign.[52]

The responses to the shock identified by the expected path instrument (Figure 7) are markedly different from the responses to the shock identified by action. Specifically, industrial production and inflation increase. Also the 10-year bond yield increases and the increase is mainly due to an increase in the expectations component. The S&P500 rises too and there is not much change in the two measures of monetary policy uncertainty. These responses are suggestive of either the Fed information effect (Nakamura and Steinsson, 2018) or the Fed response to news channel (Bauer and Swanson, 2020).[53] If the Fed information effect is present, the instrument identifies a revelation, by the FOMC announcement, of positive news about the future state of the economy, which was not in the public domain before the FOMC meeting. If the Fed response to news channel is present, the instrument instead identifies a change in the market's assessment of the Fed's future path of monetary policy due to a revision in the market's estimate of the Fed's responsiveness to the economy. On the basis of the LP alone, it is not possible to discriminate between the two theories (see Bauer and Swanson, 2020, for how to discriminate between the two theories). Although both theories are based on a positive underlying news, the housing market contracts, as the 30-year mortgage rate increases in line with the 10-year bond yield.[54]

Finally, Figure 8 contains responses to a shock identified by the uncertainty instrument.

[51] See also the 'monetary policy shock' in Jarocinski and Karadi (2020).

[52] The negative relationship between the mortgage rate and the housing market variables is in line with the price effect of monetary policy in Garriga, Kydland, and Šustek (2017).

[53] See also the 'central bank information shock' in Jarocinski and Karadi (2020).

[54] In the Online Appendix we show that the expected path instrument extracted from the unconstrained model $\mathcal{M}_0$ is unable to identify the shock. This is because the small sample bias in model $\mathcal{M}_0$ implies that the expectations component at long horizons is relatively unimportant. This reduces drastically the relevance of the expected path instrument.

Supporting the uncertainty interpretation of the instrument, the two proxies of monetary policy uncertainty rise on impact. The term premium at the monthly frequency also increases, accompanied by an increase in the expectations component, leading to an increase in the 10-year yield. Interestingly, the responses of industrial production and inflation do not conform to macroeconomic effects of uncertainty shocks (Bloom, 2009) and the S&P500 remains broadly flat. However, the shocks explored by the uncertainty shocks literature are not about future monetary policy and therefore it is not guaranteed that the conclusions carry over to the present context. The responses in Figure 8 suggest a mechanism that could be explored in future research. The increase in the 10-year bond yield is followed by a similar increase in the 30-year mortgage rate and a contraction in the housing market. The resulting decline in demand for mortgages may free up loanable funds for the corporate sector, leading to the observed decline in the EBP. The easier access to credit by firms (Bernanke et al., 1999) may in turn counteract any negative effects of uncertainty on industrial production.

6 Conclusions

HF changes in the yield curve around FOMC announcements are used to advance our understanding of monetary policy surprises and their effects on the macroeconomy. To this end, we adopt a three-stage procedure. First, we decompose high-frequency movements in the yield curve around FOMC meetings into expectations and term premia. Unlike existing work on the topic, we carry out this decomposition using term structure models (and we also correct for a small sample bias in the estimates of the two components).

Second, we decompose the HF reaction of expected interest rates and term premia across maturities into their respective PCs and use these to construct orthogonal instruments to identify monetary policy shocks. An orthogonal rotation of the PCs provides an economic interpretation of the instruments as a monetary policy action, expected path and its uncertainty. The instruments extend the proxies for monetary policy shocks employed in previous studies, which were typically based either on a single maturity or extracted only from the short-end of the yield curve.

Third, impulse-responses provide further structural interpretation. Responses to the shock identified by the action instrument are consistent with a standard monetary policy shock in a New Keynesian model with financial frictions. The expected path instrument appears to identify a shock that induces responses that are consistent with a Fed information effect or the Fed response to news channel. The shock identified by the uncertainty instrument is associated with an increase in term premia and monetary policy uncertainty. However, the excess bond premium, measuring tightness in corporate credit market, declines in response to the shock, mitigating the impact of a rise in uncertainty on output. All three shocks have a pronounced effect on the housing market, whereby an increase in long-term interest rates is associated with a decline in new home sales and house prices.

Our analysis has been carried out on the sample preceding the 2008 global financial crisis and the subsequent zero lower bound and unconventional monetary policies. The findings thus characterise the transmission mechanism in a conventional setting. Following the approach of Swanson (2021), the analysis could be extended to the subsequent period. However, to adequately account for the zero lower bound, the term structure model would need to depart from the convenient affine representation, as, for example, in Wu and Xia (2016). We see such extensions as a promising avenue for future research.

References

Abrahams, M., Adrian, T., Crump, R. K., Moench, E., Yu, R., 2016. Decomposing real and nominal yield curves. Journal of Monetary Economics 84, 182–200.

Adrian, T., Crump, R. K., Moench, E., 2013. Pricing the term structure with linear regressions. Journal of Financial Economics 110, 110–38.

Atkeson, A., Kehoe, P. J., 2009. On the need for a new approach to analyzing monetary policy. In: NBER Macroeconomics Annual, Volume 23. National Bureau of Economic Research, Inc.

Bagliano, F., Favero, C., 1999. Information from financial markets and VAR measures of monetary policy. European Economic Review 43, 825–837.

Bauer, M. D., 2015. Inflation expectations and the news. International Journal of Central Banking 11, 1–40.

Bauer, M. D., 2018. Restrictions on risk prices in dynamic term structure models. Journal of Business and Economic Statistics 36:2, 196–211.

Bauer, M. D., Rudebusch, G. D., Wu, J. C., 2012. Correcting estimation bias in dynamic term structure models. Journal of Business and Economic Statistics 30, 454–467.

Bauer, M. D., Swanson, E. T., 2020. The Fed's response to economic news explains the "Fed information effect". Working Paper Series 2020-06, Federal Reserve Bank of San Francisco.

Beechey, M., 2007. A closer look at the sensitivity puzzle: The sensitivity of expected future short rates and term premia to macroeconomic news. Finance and Economics Discussion Series 2007-06, Federal Reserve Board.

Beechey, M. J., Wright, J. H., 2009. The high-frequency impact of news on long-term yields and forward rates: Is it real? Journal of Monetary Economics 56, 535–44.

Bernanke, B., Gertler, M., Gilchrist, S., 1999. The financial accelerator in a quantitative business cycle framework. In: Taylor, J., Woodford, M. (Eds.), Handbook of Macroeconomics. Amsterdam: North Holland.

Bernanke, B. S., Kuttner, K., 2005. What explains the stock market's reaction to Federal Reserve policy? Journal of Finance 60, 1221–57.

Blinder, A. S., Ehrmann, M., Fratzscher, M., de Haan, J., Jansen, D. J., 2008. Central bank communication and monetary policy: A survey of theory and evidence. Journal of Economic Literature 46, 910–45.

Bloom, N., 2009. The impact of uncertainty shocks. Econometrica 77, 623–85.

Campbell, J. R., Evans, C. L., Fisher, J. D. M., Justiniano, A., 2012. Macroeconomic effects of federal reserve forward guidance. Brookings Papers on Economic Activity Spring, 1–80.

Christensen, J. H. E., Diebold, F. X., Rudebusch, G. D., 2009. An arbitragefree generalized NelsonSiegel term structure model. The Econometrics Journal 12, C33–C64.

Christiano, L. J., Eichenbaum, M., Evans, C. L., 1999. Monetary policy shocks: What have we learned and to what end? In: Taylor, J., Woodford, M. (Eds.), Handbook of Macroeconomics. Amsterdam: North Holland.

Cieslak, A., Schrimpf, A., 2019. Non-monetary news in central bank communication. Journal of International Economics 118, 293–315.

Cochrane, J. H., Piazzesi, M., 2002. The Fed and interest rates—A high-frequency identification. American Economic Review—Papers and Proceedings 92, 90–95.

Cochrane, J. H., Piazzesi, M., 2008. Decomposing the yield curve. Mimeo.

Cogley, T., Sargent, T., 2002. Evolving post-World War II U.S. inflation dynamics. In: NBER Macroeconomics Annual 2001, Volume 16. National Bureau of Economic Research, Inc, pp. 331–388.

Diebold, F. X., Piazzesi, M., Rudebusch, G. D., 2005. Modelling bond yields in finance and macroeconomics. American Economic Review: Papers and Proceedings 95, 415–420.

Duffee, G. R., 2011. Information in (and not in) the term structure. Review of Financial Studies 29, 2895–934.

Duffee, G. R., 2012. Forecasting interest rates. In: Timmermann, A., Elliot, G. (Eds.), Handbook of Economic Forecasting. Elsevier, Amsterdam, the Netherlands.

Evans, C. L., Marshall, D. A., 1998. Monetary policy and the term structure of nominal interest rates: Evidence and theory. Carnegie-Rochester Conference Series on Public Policy 49, 53–111.

Garriga, C., Kydland, F. E., Šustek, R., 2017. Mortgages and monetary policy. Review of Financial Studies 30, 3337–75.

Gertler, M., Karadi, P., 2015. Monetary policy surprises, credit costs, and economic activity. American Economic Journal: Macroeconomics 7, 44–76.

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

Gürkaynak, R., Sack, B., Swanson, E., 2005a. The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. American Economic Review 95, 425–36.

Gürkaynak, R., Sack, B., Swanson, E., 2005b. Do actions speak louder that words? The response of asset prices to monetary policy actions and statements. International Journal of Central Banking 1, 55–93.

Gürkaynak, R., Wright, J. H., 2012. Macroeconomics and the term structure. Journal of Economic Literature 50, 331–67.

Hamilton, J., Wu, J. C., 2012. Identification and estimation of gaussian affine term structure models. Journal of Econometrics 168, 315–31.

Hamilton, J. D., 2008. Daily monetary policy shocks and new home sales. Journal of Monetary Economics 55, 1171–1190.

Hanson, S. G., Stein, J. C., 2015. Monetary policy and long-term real rates. Journal of Monetary Economics 115, 429–48.

Hördahl, P., Remolona, E. M., Valente, G., 2015. Expectations and risk premia at 8:30AM: Macroeconomic announcements and the yield curve. Working Paper 527, Bank for International Settlements.

Husted, L., Rogers, J., Sun, B., 2020. Monetary policy uncertainty. Journal of Monetary Economics 115, 20–36.

Jarocinski, M., Karadi, P., 2020. Deconstructing monetary policy surprises: The role of information shocks. American Economic Journal: Macroeconomics 12, 1–43.

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

Joslin, S., Priebsch, M., Singleton, K. J., 2014. Risk premiums in dynamic term structure models with unspanned macro factors. Journal of Finance LXIX, 1197–1233.

Joslin, S., Singleton, K. J., Zhu, H., 2011. A new perspective on gaussian dynamic term structure models. Review of Financial Studies 24, 926–970.

Jurado, K., Ludvigson, S. C., Ng, S., 2015. Measuring uncertainty. American Economic Review 105, 1177–1216.

Kaminska, I., Roberts-Sklar, M., 2018. Volatility in equity markets and monetary policy rate uncertainty. Journal of Empirical Finance 45, 68–83.

Kendall, M. G., 1954. A note on bias in estimation of autocorrelation. Biometrica 41, 403–404.

Kim, D. H., Orphanides, A., 2012. Term structure estimation with survey data on interest rate forecasts. Journal of Financial and Quantitative Analysis 47, 241–272.

Kuttner, K. N., 2001. Monetary policy surprises and interest rates: Evidence from the fed funds futures market. Journal of Monetary Economics 47, 523–44.

Mertens, K., Ravn, M. O., 2013. The dynamic effects of personal and corporate income tax changes in the United States. American Economic Review 103, 1212–47.

Miranda-Agrippino, S., 2016. Unsurprising Shocks: Information, Premia, and the Monetary Transmission. Working Paper 626, Bank of England.

Miranda-Agrippino, S., Ricco, G., Sep. 2015. The Transmission of Monetary Policy Shocks. Discussion Papers 1711, Centre for Macroeconomics.

Nakamura, E., Steinsson, J., 2018. High-frequency identification of monetary non-neutrality: The information effect. Quarterly Journal of Economics 133, 1283–1330.

Nicholls, D. F., Pope, A. L., 1988. Bias in the estimation of multivariate autoregressions. Australian and New Zealand Journal of Statistics 30, 296–309.

Piazzesi, M., 2006. Affine term structure models. In: Ait-Sahalia, Y., Hansen, L. P. (Eds.), Handbook of Financial Econometrics. Elsevier, Amsterdam.

Pierse, R. G., Snell, A. J., 1995. Temporal aggregation and the power of tests for a unit root. Journal of Econometrics 65, 333–45.

Ramey, V. A., 2016. Macroeconomic shocks and their propagation. In: Taylor, J. B., Uhlig, H. (Eds.), Handbook of Macroeconomics. Elsevier, Amsterdam.

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

Rudebusch, G. D., 1998. Do measures of monetary policy in a Var make sense? International Economic Review 39, 907–31.

Rudebusch, G. D., Swanson, E. T., 2012. The bond premium in a DSGE model with long-run real and nominal risks. American Economic Journal: Macroeconomics 4, 105–143.

Shaman, P., Stine, R. A., 1988. The bias of autoregressive coefficient estimators. Journal of the American Statistical Association 83, 842–48.

Swanson, E., 2006. Have increases in Federal Reserve transparency improved private sector interest rate forecasts? Journal of Money, Credit, and Banking 38, 791–819.

Swanson, E., 2021. Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets. Journal of Monetary Economics 118, 32–53.

Tillmann, P., 2020. Monetary policy uncertainty and the response of the yiled curve to policy shocks. Journal of Money, Credit, and Banking 52, 803–833.

Woodford, M., 2005. Central bank communication and policy effectiveness. NBER Working Paper 11898.

Wright, J. H., 2011. Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. American Economic Review 101, 1514–34.

Wright, J. H., 2017. Forward-looking estimates of interest rate distributions. Annual Review of Financial Economics 9, 333–51.

Wu, J. C., Xia, F. D., 2016. Measuring the macroeconomic impact of monetary policy at the zero lower bound. Journal of Money, Credit and Banking 48, 253–91.

Table 1: Effect of FOMC announcements on yields across maturities

	3-month	2-year	3-year	5-year	10-year
Average response, bps	-1.4	-1	-1.1	-0.5	-0.3
Minimum, bps	-23	-22	-23	-16	-16
Maximum, bps	9	19	21	19	13
St. Deviation	4.6	5.9	6.1	5.2	4.3
Correlations					
3-month	1	0.57	0.49	0.41	0.35
2-year		1	0.92	0.93	0.86
3-year			1	0.91	0.85
5-year				1	0.90

Note: The sample is from January 1996 to August 2007.

Table 2: Loadings on the components of policy surprises

	3-M	2-YR	3-YR	5-YR	10-YR
ROTATED PCS $\mathcal{P}_{1t}^{E*}$ AND $\mathcal{P}_{1t}^{TP*}$					
Expectations					
Action	1.48	1.00	0.84	0.68	0.51
Expected path	0	1.00	0.97	0.84	0.64
Term premia					
Uncertainty	0	1.00	1.05	1.40	2.23
GÜRKAYNAK ET AL. (2005B)					
Target	2.07	1.00	n/a	0.57	0.27
Path	0	1.00	n/a	0.90	0.69
KUTTNER (2001)					
Target	1.29	1.00	n/a	0.78	0.51

Note: The loadings for action are the first column of Ω_1^{E*}; the loadings for expected path are the second column of Ω_1^{E*}; the loadings for uncertainty are the first column of Ω_1^{TP*}. For ease of comparison across studies, the loadings are normalised to be equal to one at the 2-year maturity. Our sample is January 1996-August 2007. Gürkaynak et al. (2005b): Table 5, sample July 1991-December 2004. Kuttner (2001): Table 3, sample June 1989-February 2000, daily changes.

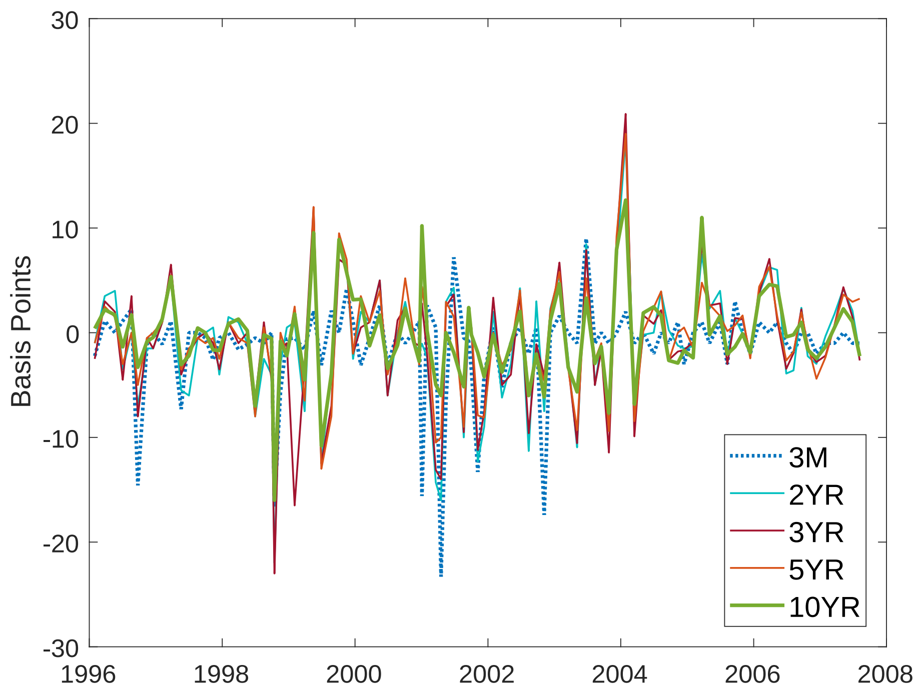

Figure 1: Yield changes around FOMC announcements across maturities.

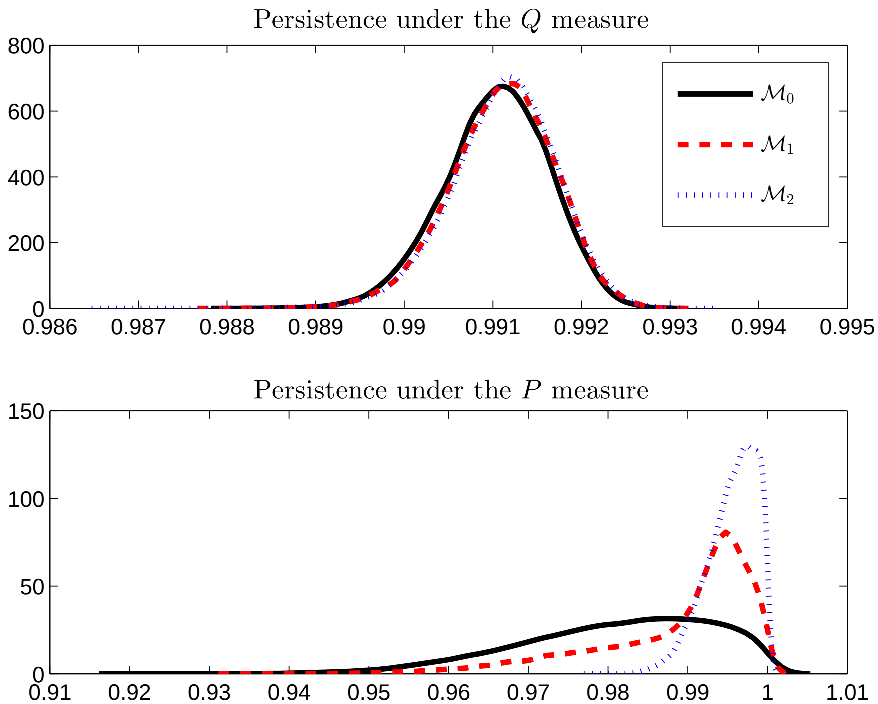

Figure 2: The effects of estimating the ATSM subject to restrictions: posterior distribution of persistence (largest eigenvalue).

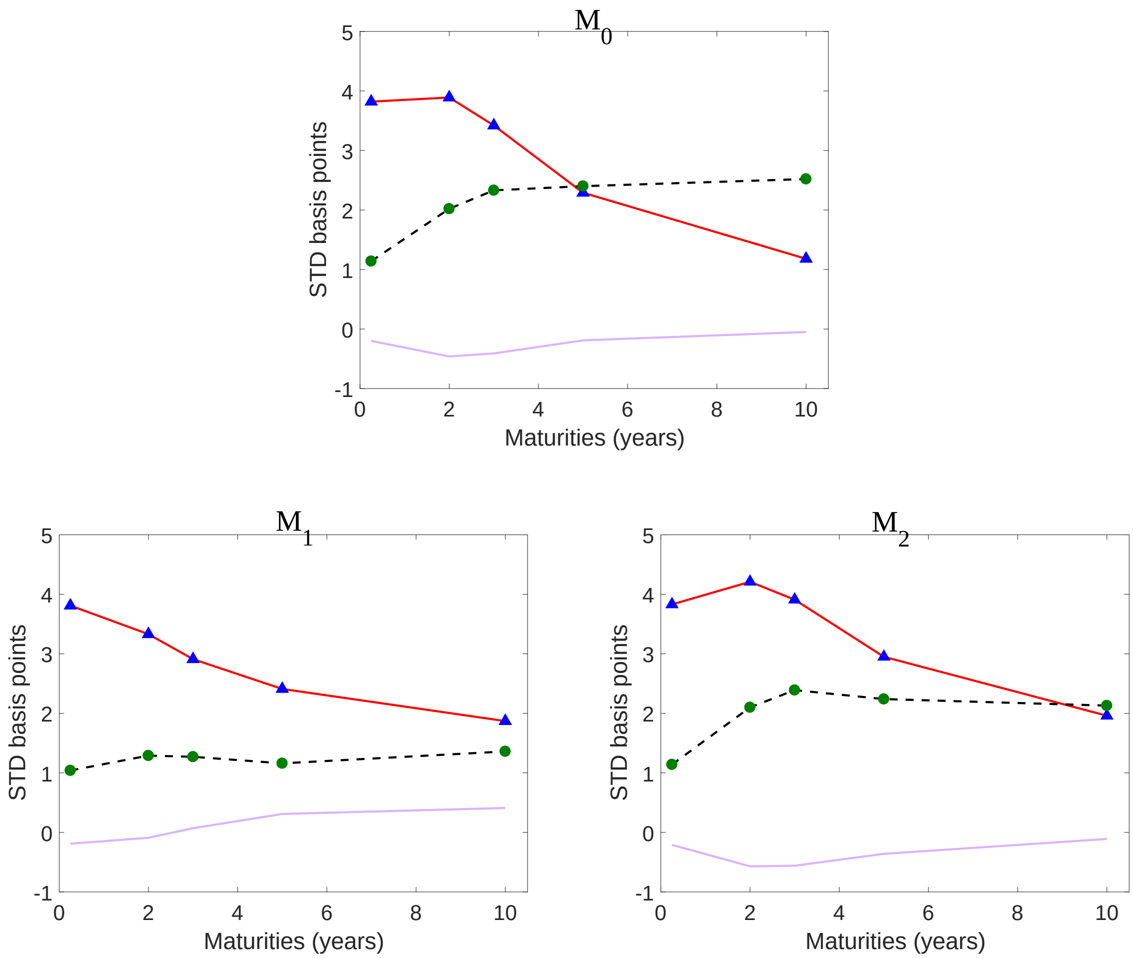

Figure 3: Volatility curve of the change in expectations and term premia around FOMC announcements. Solid line with markers: expectations. Dashed line: term premia. Solid line without markers: correlation between the two components. Markers denote the available maturities at the high frequency. The shortest maturity is three months.

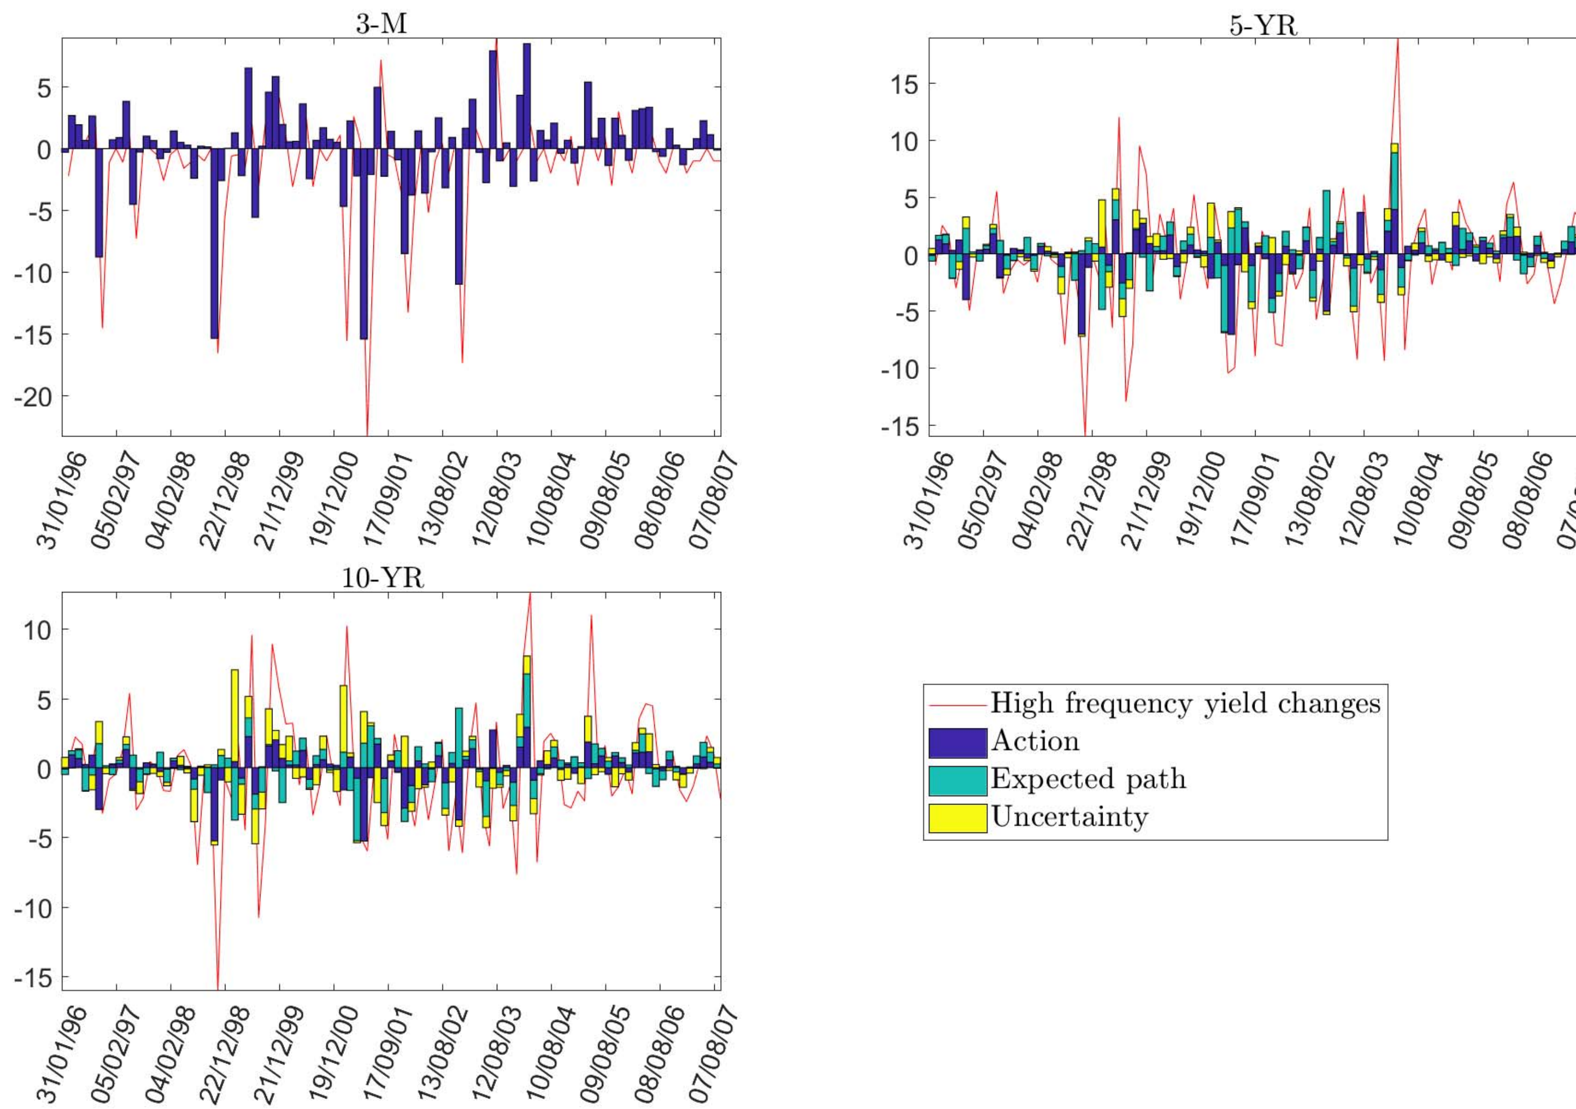

Figure 4: Contribution of the three policy instruments to the high-frequency changes in yields around FOMC announcements. The units are basis points.

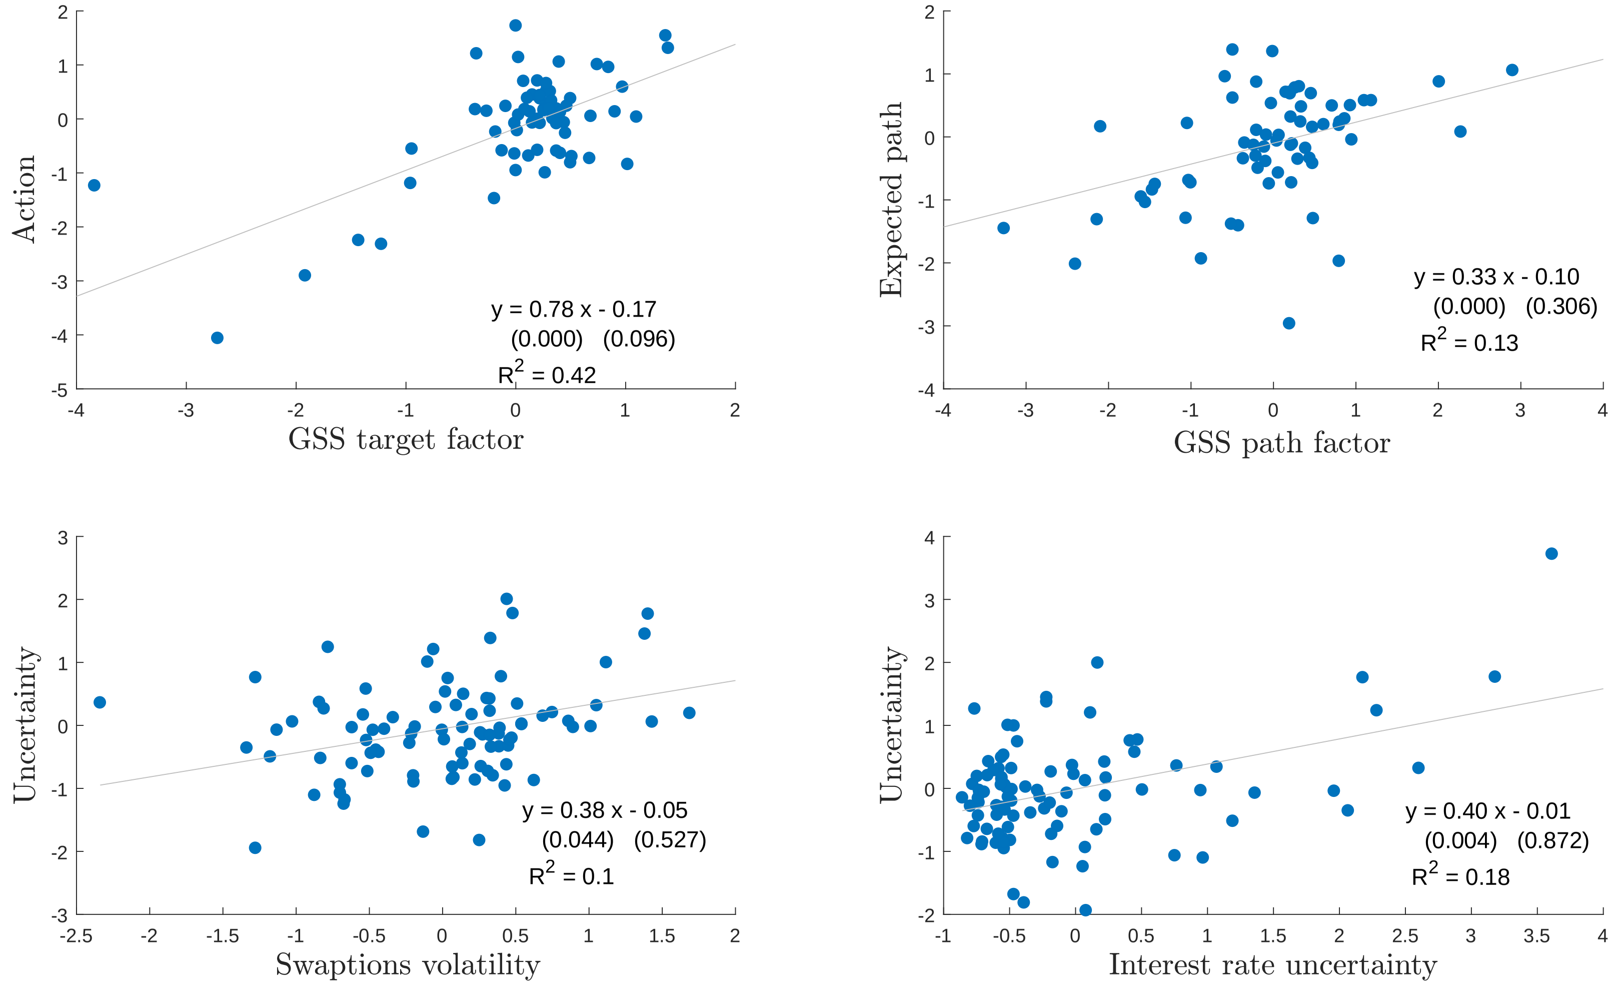

Figure 5: Cross-validation of the instruments. The numbers in the parentheses are p-values.

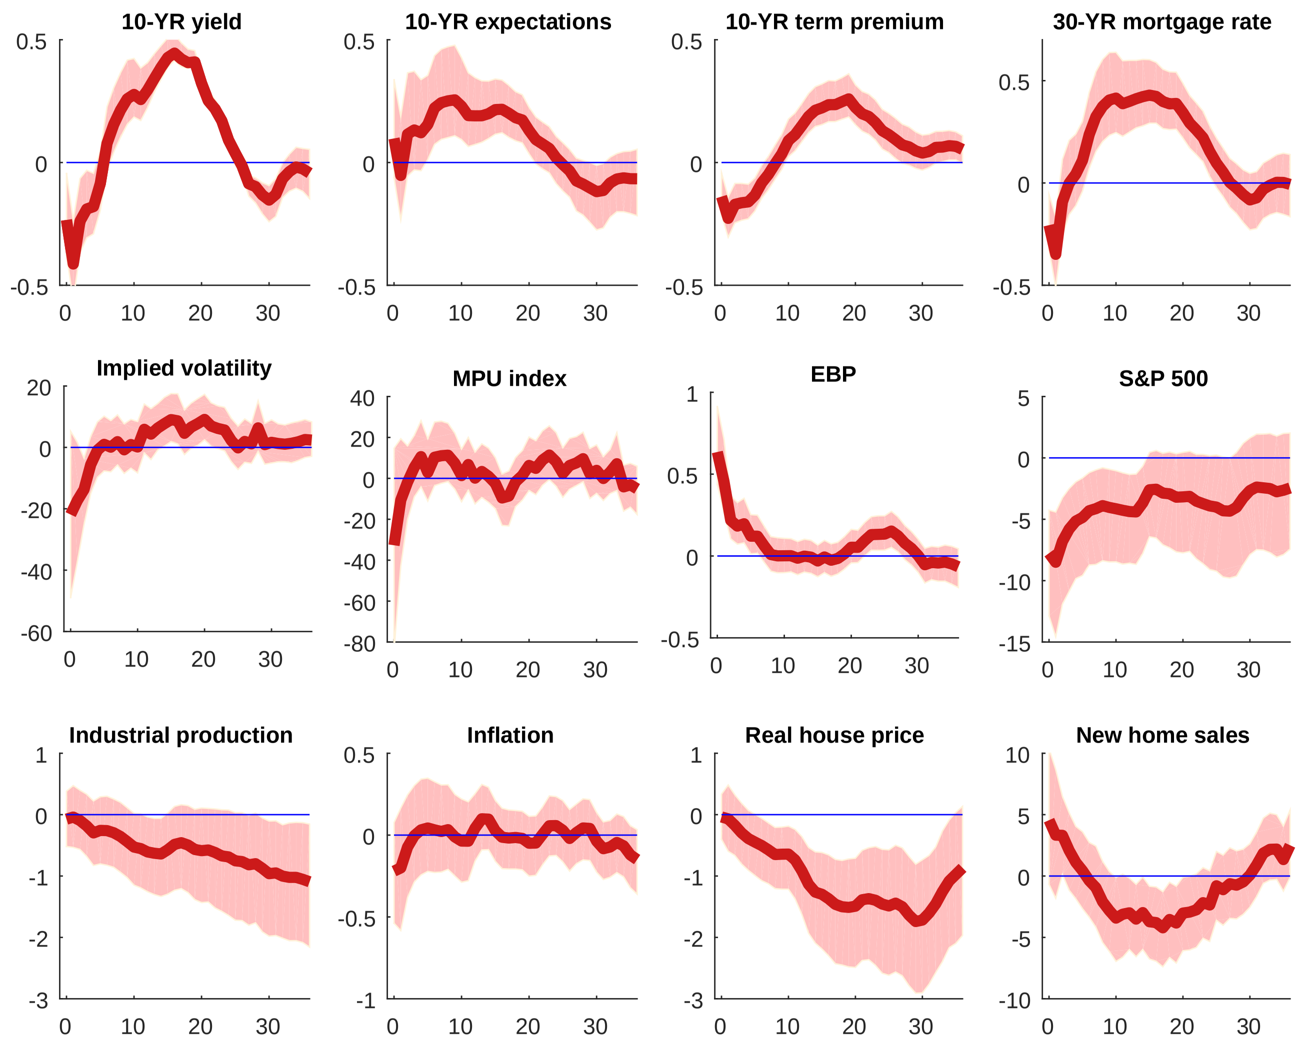

Figure 6: Local projections: action. The horizon for the responses is 36 months. Responses of interest and inflation rates are in percentage points, other responses are in percent. The shock is normalised to increase the slope factor by one percentage point. The charts plot the median response and the 90% error band.

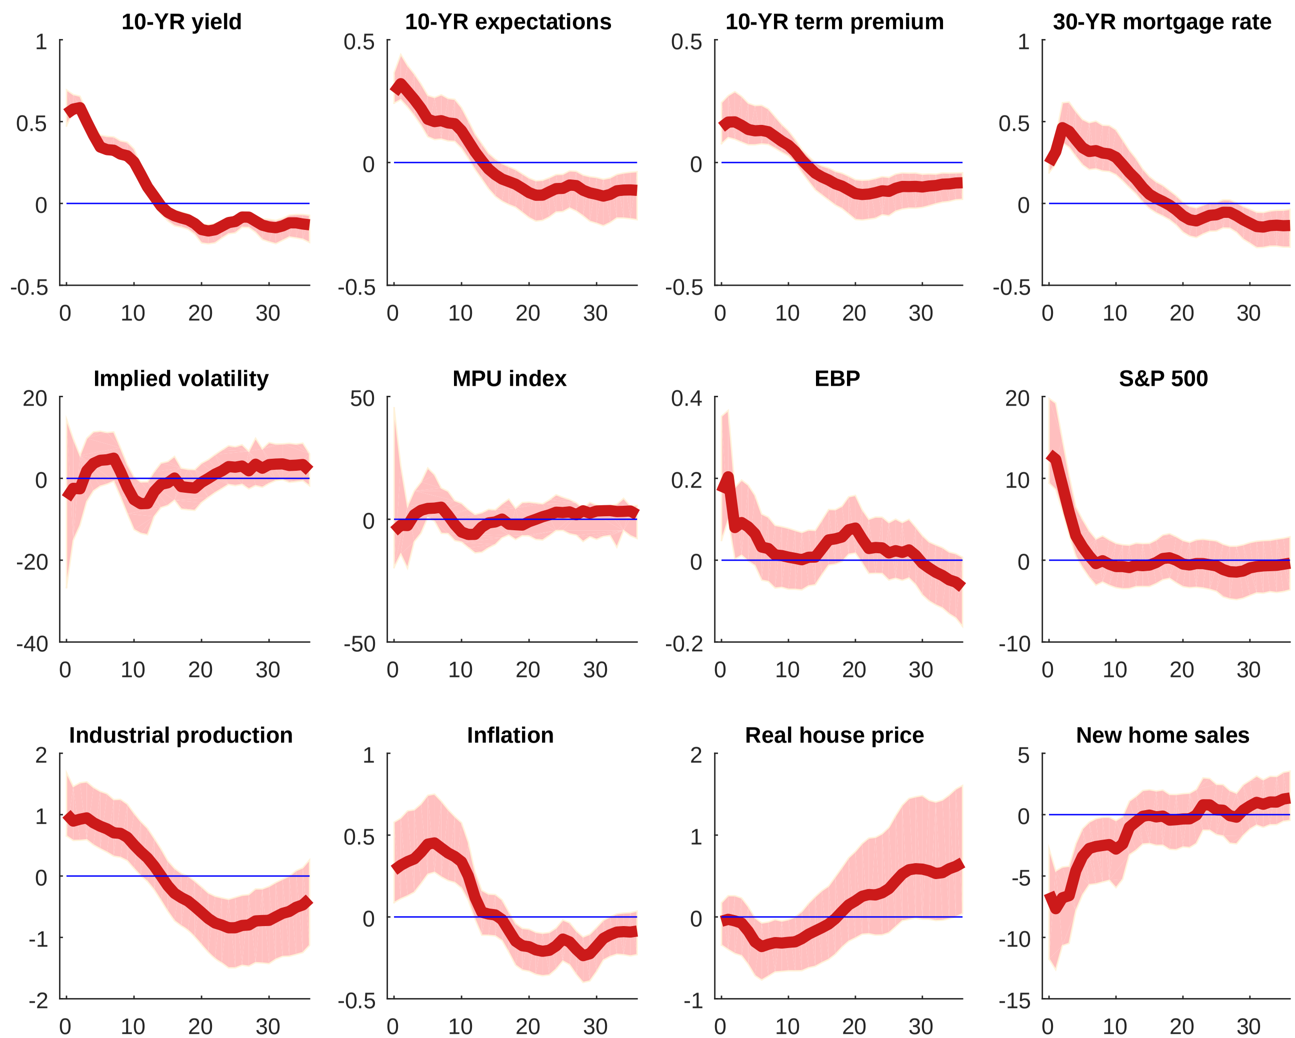

Figure 7: Local projections: expected path. The horizon for the responses is 36 months. Responses of interest and inflation rates are in percentage points, other responses are in percent. The shock is normalised to increase the level factor by one percentage point. The charts plot the median response and the 90% error band.

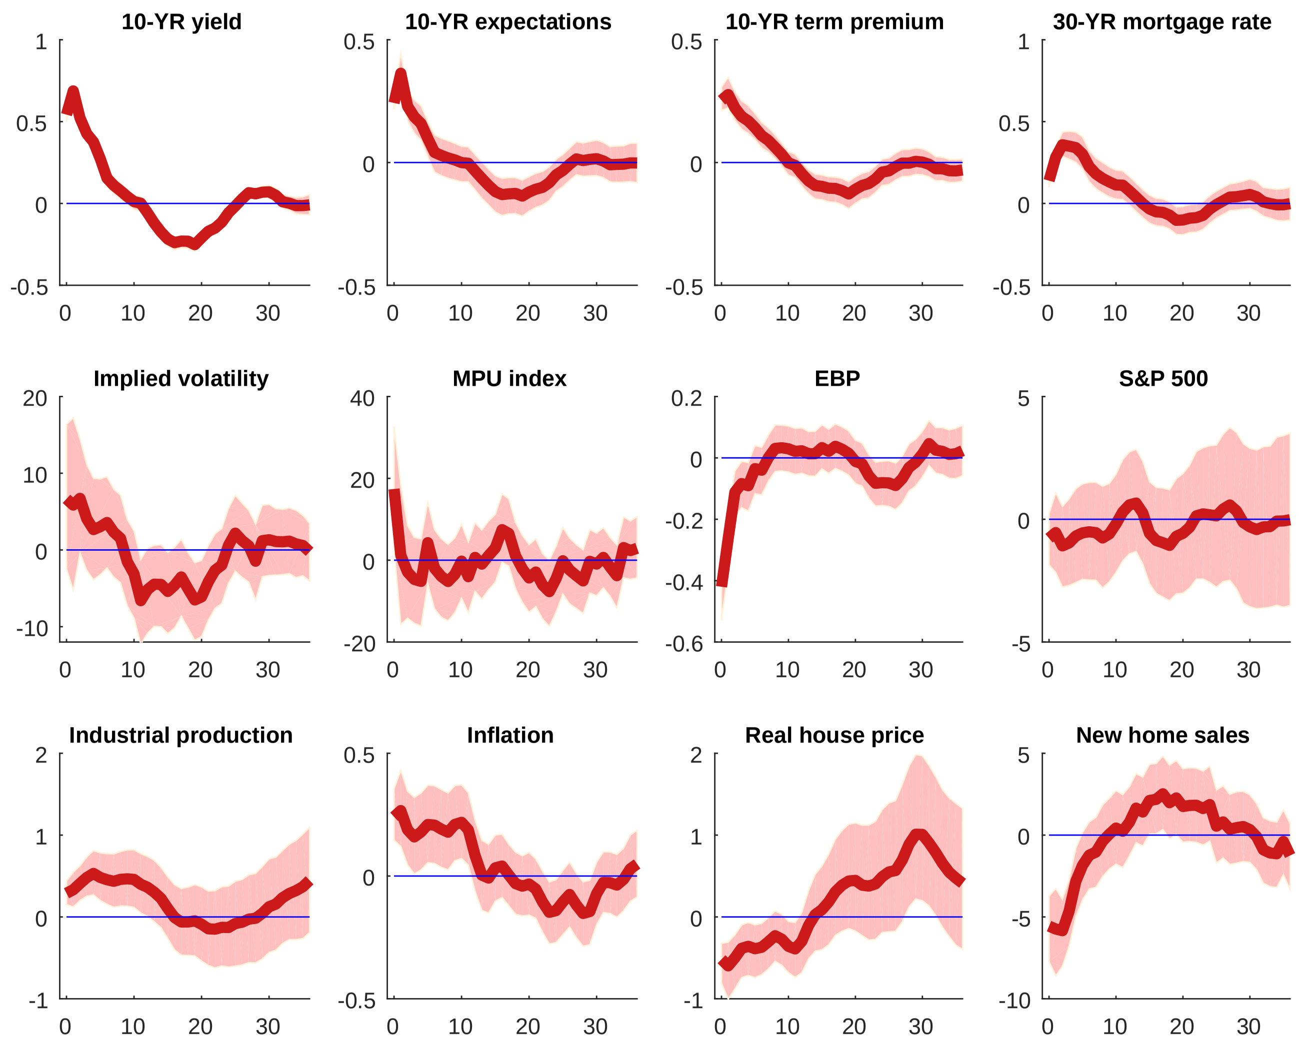

Figure 8: Local projections: uncertainty. The horizon for the responses is 36 months. Responses of interest and inflation rates are in percentage points, other responses are in percent. The shock is normalised to increase the level factor by one percentage point. The charts plot the median response and the 90% error band.