Risk Premia, Subjective Beliefs, and Bundled Monetary Shocks[*]

Anna Cieslak
Duke University Fuqua School of Business

Paymon Khorrami
Duke University Fuqua School of Business

November 16, 2023

Abstract

We consider identification of monetary shocks and their causal impacts in quasi-linear environments where (i) agents may possess subjective beliefs and (ii) monetary authorities manage current and future interest rates (e.g., forward guidance). Assuming rational expectations or risk-neutrality trivially enables identification. Without those assumptions, identification of monetary shocks from asset prices hinges on a Long-Run Neutrality condition, roughly meaning policy does not affect the compensation for permanent risks. We construct a non-parametric test of the Long-Run Neutrality condition, related to the literature on FOMC announcement effects, and argue that it is violated in the data. Finally, we present some example models in which the Long-Run Neutrality condition is violated, illustrating how this condition is generally distinct from conventional notions of monetary neutrality.

JEL Codes: E44, E52, E7, G12

Keywords: monetary policy; subjective beliefs; asset pricing; forward guidance; macroeconomic announcement effects; belief recovery theory.

[*]Cieslak: anna.cieslak@duke.edu. Khorrami: paymon.khorrami@duke.edu. We are grateful to seminar participants at various institutions for helpful feedback.

1 Introduction

How does a monetary authority influence the economy? In this paper, we argue that the lack of consistent answers to this question stems from an identification problem. Our framework combines two key ingredients: unobservable investor beliefs and monetary policies affecting the future, such as forward guidance. Investor beliefs are a critical object, because the monetary shocks researchers typically consider are in fact surprises relative to those investor beliefs. Forward guidance, which operates by manipulating investor forecasts, becomes interesting and natural to explore jointly with beliefs. In this environment, we first show that identification of the various monetary surprises and their impacts requires a Long-Run Neutrality condition that differs from typical notions of monetary policy neutrality. Second, we propose a model-free test of this condition; the evidence suggests Long-Run Neutrality is violated.

Multiple monetary instruments. Even in a simplified world without any central bank information advantage and with only conventional tools (e.g., no long-term asset purchases/sales), monetary authorities impact the economy through several potential channels. Our starting point is that every FOMC meeting can contain up to three categories of shocks:

1. *Short-Rate Shocks:* unexpected changes to the short-term interest rate
2. *Forward-Guidance Shocks:* unexpected modifications to the expected future rate path
3. *Uncertainty Shocks:* unexpected changes in uncertainty about future rates

For example, when a central bank raises the short-term interest rate rate (Short-Rate Shock), investors must form beliefs about the persistence of the rate hike (Forward-Guidance Shock), and these investors furthermore must form beliefs about the risk that policy rates rise further in the future (Uncertainty Shock). Of course, Forward-Guidance Shocks and Uncertainty Shocks may also occur in isolation without any Short-Rate Shock. We do not impose a structure on how these three categories of shocks can be bundled together, and so we must actually identify each shock separately.

Non-identification. Given this framework, we ask two questions. First, what is the causal impact of short rates on the economy? The answer to this first question will typically require an identification of not just Short-Rate Shocks but also Forward-Guidance and Uncertainty Shocks, which leads to our second question: how can these various shocks be recovered? Let us address these two questions in more detail.

Can we can identify the causal impact of short rates on the economy? Imagine we are given a time series of the Short-Rate Shocks z_τ. A common approach examines high-frequency changes to Fed Funds futures prices on the FOMC meeting day (Krueger and Kuttner, 1996; Rudebusch, 1998; Kuttner, 2001; Rudebusch, 2002; Bernanke and Kuttner, 2005; Piazzesi and Swanson, 2008). While this procedure is not uncontroversial, we take knowledge of z_τ for granted to engage with more novel identification issues.

Given z_τ, suppose we run a regression of some future outcome $g_{\tau+t}$ on z_τ. One may think that, on average, other shocks wash out, so that a long enough sample at least identifies the average effect (ignoring nonlinearities). However, this logic fails whenever agents hold belief distortions, such that z_τ is a shock *relative to investors' beliefs*. If you want to identify the effect of z_τ alone, holding fixed $z_{\tau+1}$, $z_{\tau+2}$, etc., you need to control for these future surprises because they are correlated with z_τ under a large class of belief distortions entertained by the macro-finance literature. Without controlling for $z_{\tau+1}$, $z_{\tau+2}$, etc., we are always estimating the bundled impact of the surprise z_τ along with expected future surprises.

For example, if investors "overreact" to news as various survey forecasts suggest, then surprises will be negatively autocorrelated. Current approaches mistakenly assuming rational expectations will bundle the offsetting effects of the current and future surprises, masking the true impact of z_τ alone. In a calibrated example, we show that the "naïve local projection approach" described above (which effectively assumes rational expectations and omits $z_{\tau+1}$, $z_{\tau+2}$, etc.) leads to severe underestimation of the impacts of short-rate policy, precisely because investors appear to overreact.

If there was a single-factor structure to beliefs about short rates, it would be relatively simple to correct the bias discussed above. (We will explain this more in the paper, but the summary procedure is to include measures of z in the estimation of a larger vector autoregression.) However, if central banks also do forward guidance and other similar policies, there is a multi-factor structure to beliefs. Failing to control for these other factors further contaminates estimates of the causal effect of short rates. Unfortunately, we don't directly observe shocks to investor beliefs about future rates; some identification procedure is needed.

Of course, there are additional motivations to recover investor beliefs about future rates. Even if we knew the "causal effect of short rates"—say, because we knew that investor beliefs about short-term interest rates were approximately rational, justifying our regression of $g_{\tau+t}$ on z_τ—this effect assumes some baseline level of persistence. What if we want to know the impact of a more persistent or more transitory monetary action? Or what if we want to measure the effect of central bankers manipulating beliefs about future rates, even without changing the current rate (Gürkaynak et al., 2005b; Swanson, 2021)? How to answer such questions from the regression of $g_{\tau+t}$ on z_τ is unclear, particularly if permanent and transitory monetary shocks have very different effects (Uribe, 2022). Identifying our notion of Forward-Guidance and Uncertainty Shocks, which are primarily about future interest rates, helps with such questions.

This leads us to the second question in our paper: can Forward-Guidance and Uncertainty Shocks be recovered from asset price data? Asset markets are a natural arena to explore because of the richness in financial claims covering many horizons (e.g., very far into the future) and at many levels of contingency (e.g., isolating specific aspects of the probability distribution). We will discuss survey data as a viable alternative in various parts of the paper, but note for now that surveys do not possess the same richness in horizon or contingency.

Consider Forward-Guidance Shocks. A first thought might be to use long-term yields to reveal expectations about future short rates (Expectations Hypothesis). But the preponderance of evidence stands against the Expectations Hypothesis, because of the existence of bond risk premia, and more importantly the time-variation in these risk premia. Changes to the yield curve can only identify shocks to a risk-adjusted expectation of future short rates (for example, the risk-neutral expectation). Currently, no model-free mapping exists between these risk-adjusted expectations and investors' expectations. Assuming this gap away seems like wishful thinking, since unlike short-horizon Fed Funds futures, significant risk premia exist in long-term bonds.

Identification of Uncertainty Shocks faces an analogous issue, but in distribution space rather than mean space. Even given a full set of options on interest rates of all maturities, all we can identify is a risk-adjusted distribution of future interest rates. Given the existence of time-varying risk premia, it is nontrivial to map this risk-adjusted distribution into investors' subjective distribution. Some approaches have been proposed for separate identification of Short-Rate Shocks versus a second "path" factor encompassing all other monetary shocks (Gürkaynak et al., 2005b; Swanson, 2021). But interpreting this second factor is challenging without a model that allows us to separate the impacts of beliefs and risk premia. Our framework shows that isolating beliefs is critical.

Identification is not hopeless. Our core set of theoretical results says that, in quasi-linear environments, Forward-Guidance and Uncertainty Shocks can be identified from asset prices if and only if permanent risks and their risk prices are unaffected by monetary policy. That is, if rational expectations and risk-neutrality cannot be assumed to hold, and if we do not write down a fully-specified model of investor preferences and beliefs, then identification requires long-run risk prices to be invariant to monetary policy. For short, we refer to this monetary-invariance condition as *Long-Run Neutrality*. This exact condition is effectively imposed as an identification assumption in the recent papers of Backus et al. (2022) and Haddad et al. (2023), which had different but related goals. As we show, this is not a coincidence: identification requires such a Long-Run Neutrality assumption.

Our results connect closely to a broader issue in asset pricing, so-called recovery theory (Ross, 2015; Borovička et al., 2016). Investor beliefs are not revealed by asset prices, because beliefs are co-mingled with other permanent components of marginal utility. In the context of monetary policy, there is a nuance: we do not seek beliefs themselves, but rather shocks to beliefs. And this is why the key assumption of recovery theory, *absence* of a permanent component in marginal utility, is replaced by the *invariance* of such permanent component to monetary policy.

In our paper, we develop a non-parametric test of Long-Run Neutrality. A strategy that goes long the growth-optimal portfolio and short a long-maturity bond identifies the martingale in the pricing kernel (Alvarez and Jermann, 2005). Long-Run Neutrality implies this investment strategy behaves similarly on Fed announcement and non-announcement days. Supposing we can proxy the growth-optimal portfolio with equities and the long-maturity bond with 30-year Treasuries, a growing body of evidence rejects this prediction of Long-Run Neutrality. For example, across various time periods and methodologies, US equities earn about 2-3 times higher returns around FOMC announcements as long-term bonds. We cite this literature below and defer a discussion of the exact numbers and methodologies to the main text. This model-free test suggests that Forward-Guidance and Uncertainty Shocks cannot be identified from asset prices alone.

Finally, we discuss structural models in which economic growth and uncertainty are priced sources of risks. A leading example is Bansal and Yaron (2004). If we take these models seriously, our Long-Run Neutrality condition requires both growth and uncertainty to be invariant to monetary policy, *both in the short run and the long run*. Through the lens of these models, identifying the effects of monetary policy requires precisely that monetary policy has no effects.

Literature review. Our framework builds on two pillars. First, we entertain the possibility that market participants possess biased beliefs about future interest rates, asset returns, and monetary policy (Ball and Croushore, 2003; Hamilton et al., 2011; Chun, 2011; Giglio and Kelly, 2018; Cieslak, 2018; Crump et al., 2018; Kryvtsov and Petersen, 2019; d'Arienzo, 2020; Wang, 2021; Xu, 2019; Nagel and Xu, 2022; Bianchi et al., 2022b).

Second, we embrace environments in which monetary authorities can manage beliefs about future interest rates (Poole et al., 2002; Gürkaynak et al., 2005b; Campbell et al., 2012; Del Negro et al., 2012; Swanson, 2021), implying the presence of several distinct monetary effects, even ignoring any central bank information advantage (Romer and Romer, 2000; Melosi, 2017; Nakamura and Steinsson, 2018; Cieslak and Schrimpf, 2019; Miranda-Agrippino and Ricco, 2021).[1]

In this class of environments, estimating the causal impact of monetary policy requires a Long-Run Neutrality condition that is related to belief recovery theory (Ross, 2015; Borovička et al., 2016; Qin and Linetsky, 2016). Following a literature that decomposes the pricing kernel into permanent and stationary components (Kazemi, 1992; Alvarez and Jermann, 2005; Hansen and Scheinkman, 2009; Bakshi and Chabi-Yo, 2012; Qin and Linetsky, 2017; Corsetti et al., 2023), we develop a model-free test of the Long-Run Neutrality condition in the context of monetary policy.

Our test zooms in on asset price changes around FOMC announcements, which is related to the literature on "announcement effects." Several authors have argued that equity risk premia are strongly influenced by monetary policy announcements, including both actions and communications (Pearce and Roley, 1985; Rosa, 2011; Savor and Wilson, 2013, 2014; Lucca and Moench, 2015; Ai and Bansal, 2018; Cieslak et al., 2019; Cieslak and Pang, 2021; Bianchi et al., 2022a,c; Bauer et al., 2023). A related literature examines FOMC announcement effects in government bonds (Ederington and Lee, 1993; Gürkaynak et al., 2005a; Beber and Brandt, 2006; Faust et al., 2007; Hanson and Stein, 2015; Hillenbrand, 2021; Hanson et al., 2021). We connect these literatures by investigating the announcement effect of a particular long-short portfolio, which is the theoretically appropriate object for our purposes.

Finally, to illustrate how strong the Long-Run Neutrality condition can be, we consider a class of structural environments in which long-run growth and uncertainty become priced state variables. A large empirical literature suggests that certain longer-term prospects and news about these prospects matter (McQueen and Roley, 1993; Francis and Ramey, 2005; Beaudry and Portier, 2006; Barsky and Sims, 2011; Schmitt-Grohé and Uribe, 2012; Kurmann and Otrok, 2013; Barsky et al., 2015; Leduc and Liu, 2016; Nakamura et al., 2017; Basu and Bundick, 2017; Schorfheide et al., 2018; Berger et al., 2020; Liu

[1]A large theoretical literature has also considered the effects of monetary policy in environments with distorted beliefs. As we avoid putting too much structure on agents' beliefs, we do not directly engage with this literature, but some notable examples include Bernanke and Woodford (1997), Evans and Honkapohja (2003), Andolfatto and Gomme (2003), Schorfheide (2005), Milani (2008), Gasteiger (2014), Hommes et al. (2019), and Caballero and Simsek (2022). Caballero and Simsek (2022), in particular, argue that subjective belief dynamics themselves could be the *origin* of monetary policy shocks.

and Matthies, 2022). Models consistent with this evidence imply that persistent shocks to economic growth and uncertainty comprise the permanent component of the pricing kernel (Bansal and Yaron, 2004; Beaudry and Portier, 2004; Bloom, 2009; Bidder and Dew-Becker, 2016; Christiano et al., 2014; Fajgelbaum et al., 2017; Bianchi et al., 2018; Di Tella, 2017; Di Tella and Hall, 2022; Bianchi et al., 2023). A related literature also theorizes and documents the importance of monetary and other policy uncertainty (Baker et al., 2016; Creal and Wu, 2017; Husted et al., 2020; Pastor and Veronesi, 2012; Pástor and Veronesi, 2013; Kelly et al., 2016). If we accept these types of environments, Long-Run Neutrality implies that monetary policy does not affect the probability distribution of future growth.

2 A Simple VAR Example

We begin with the simplest possible example to illustrate the key results. This example can be made more complicated/realistic, and this will only worsen the identification problems we highlight. First, the presence of distorted beliefs complicates the identification of short-rate effects. Second, if a researcher is able to separately identify all types of monetary shocks, then these other shocks can be used as controls to then recover short-rate effects. Third, measuring the various monetary shocks is not possible in general without a model, combined with a slew of restrictions; we provide such conditions. Section 3 substantially generalizes the model and furthermore allows some types nonlinearities to permit the broadest statement of our identification results.

2.1 Model with two monetary actions

Consider a three-state model

$$X_t = \begin{bmatrix} g_t \\ r_t \\ f_t \end{bmatrix} = \begin{bmatrix} \text{(demeaned) growth rate} \\ \text{(demeaned) interest rate} \\ \text{forward guidance} \end{bmatrix}.$$

We presume that these states evolve dynamically according to

$$X_{t+1} = AX_t + B\Delta W_{t+1}, \tag{1}$$

where $\Delta W_{t+1} \sim \text{Normal}(0, I)$ is a 3-dimensional vector of Normal shocks. For the purpose of this example, we will think of $\Delta W_{t+1}^{(1)}$ as a "real shock" and $\Delta W_{t+1}^{(2)}$ and $\Delta W_{t+1}^{(3)}$ are "monetary shocks" that are completely governed by central bank actions. Equation (1) governs the *objective* state dynamics, which may differ from agents' *perceived* dynamics that we detail in Section 2.2.

Let us specify the shock structure B and the persistence matrix A. We assume

$$B = \begin{bmatrix} \sigma_g & 0 & 0 \\ 0 & \sigma_r & 0 \\ 0 & 0 & \sigma_f \end{bmatrix}. \tag{2}$$

With B specified as such, there is no contemporaneous effect of monetary shocks on the growth rate. We do not need this assumption, but it will make the algebra more transparent. It also clarifies that our results are distinct from the classic timing issues in identifying monetary effects. Furthermore, short-rate and forward-guidance shocks are assumed orthogonal, for simplicity.

For the persistence matrix, we assume

$$A = \begin{bmatrix} a_{gg} & a_{gr} & a_{gf} \\ 0 & a_{rr} & a_{rf} \\ 0 & 0 & a_{ff} \end{bmatrix}. \tag{3}$$

and that A is stable. We make two remarks about this structure. First, the monetary block (r_t, f_t) evolves independently of g_t and the real shock $W_t^{(1)}$, a setup we have adopted for convenience. Second, the forward guidance variable f_t has the interpretation of a monetary action that affects future interest rates, which we will now show.

To understand the model, compute the impulse response functions (IRFs):

$$D_t^{(1)} := \frac{d}{dw}\mathbb{E}\left[X_{t+1} \mid X_0, \Delta W_1^{(1)} = w\right] = A^t B \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} a_{gg}^t \sigma_g \\ 0 \\ 0 \end{bmatrix}$$

$$D_t^{(2)} := \frac{d}{dw}\mathbb{E}\left[X_{t+1} \mid X_0, \Delta W_1^{(2)} = w\right] = A^t B \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ a_{rr}^t \sigma_r \\ 0 \end{bmatrix}$$

$$D_t^{(3)} := \frac{d}{dw}\mathbb{E}\left[X_{t+1} \mid X_0, \Delta W_1^{(3)} = w\right] = A^t B \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ a_{ff}^t \sigma_f \end{bmatrix} + (1_{t\geq 1}) \begin{bmatrix} 0 \\ \sigma_f a_{rf} \sum_{j=1}^t a_{rr}^{t-j} a_{ff}^{j-1} \\ 0 \end{bmatrix}$$

The IRFs confirm the labelling given above: $\Delta W_t^{(1)}$ is a pure growth-rate shock; $\Delta W_t^{(2)}$

is a pure short-rate shock, with persistence a_{rr}; and $\Delta W_t^{(3)}$ is a pure forward-guidance shock that raises future interest rates by $a_{rf}\sum_{j=1}^{t} a_{rr}^{t-j} a_{ff}^{j-1}$. Thus, we like to think of forward guidance as promises to modify the future path of short rates, while taking no action today.

By combining various amounts of $\Delta W^{(2)}$ and $\Delta W^{(3)}$, we engineer short-rate shocks of various persistences. Compute the expected future rate if $B\Delta W = (0, u, 1-u)'$:

$$\mathbb{E}\Big(r_{t+1} \mid B\Delta W_1 = \Big[\begin{smallmatrix} 0 \\ u \\ 1-u \end{smallmatrix}\Big]\Big) = u a_{rr}^t + (1-u) a_{rf} \sum_{j=1}^{t} a_{rr}^{t-j} a_{ff}^{j-1}, \quad t \geq 1.$$

Depending on the magnitudes of a_{rr}, a_{ff}, and a_{rf}, such combinations can alter the short rate persistence. The simple linear example here puts a very strong and unnecessary structure on how the short rate path is affected, but this is only for transparency.

2.2 Subjective beliefs

We allow agents' beliefs to potentially be non-rational. While we take no stand here, belief distortions could come from multiple sources: pure cognitive biases, imperfect information or attention, finite samples with imperfect priors, etc. To keep things simple, we consider a belief distortion that modifies the persistence of X_t. Under agents' subjective beliefs

$$X_{t+1} = \tilde{A} X_t + B\Delta\tilde{W}_{t+1}, \quad \Delta\tilde{W}_{t+1} \sim \text{Normal}(0, I). \tag{4}$$

We will assume $\tilde{A}$, like A, is a stable matrix. The notation $\tilde{\mathbb{E}}$ will stand for the subjective expectation operator for agents in our model, which may or may not coincide with the objective expectation $\mathbb{E}$.

Thus, agents perceive

$$\Delta\tilde{W}_{t+1} = \Delta W_{t+1} - L_t, \quad \text{where} \quad L_t := B^{-1}\big(\tilde{A} - A\big) X_t, \tag{5}$$

to be a standard Normal shock. One interpretation is that the vector L_t represents investors' time-varying degree of optimism.

If we observe the vector X_t for long enough, we can obtain A and B in (1) by linear regression. However, it will be difficult to obtain $\tilde{A}$ in (4) this way, because regressions take place under the objective measure.

Of particular interest are the perceived growth dynamics. To keep the belief distortion to the minimum level needed for our results, we assume

$$\tilde{A} = \begin{bmatrix} a_{gg} & a_{gr} & a_{gf} \\ 0 & \tilde{a}_{rr} & \tilde{a}_{rf} \\ 0 & 0 & \tilde{a}_{ff} \end{bmatrix}. \tag{6}$$

That is, agents hold no biases about the dynamics of g_t. Belief distortions still do affect growth forecasts, but only through biases about the dynamics of (r_t, f_t).

2.3 The causal effect of short rates

Following a large portion of the literature, we would like to answer the question "what is the causal effect of a short-rate shock on future growth?"

A *short-rate shock* at time τ defined as

$$z_\tau := r_\tau - \tilde{\mathbb{E}}[r_\tau \mid X_{\tau-1}].$$

In this paper, we take as given the ability to non-parametrically identify the short rate shock from data. In particular, assume the existence of a financial market (e.g., Fed Funds futures) whose price corresponds to $\tilde{\mathbb{E}}[r_\tau \mid X_{\tau-1}]$ at time $\tau - 1$. It is technically appropriate to use the investor expectation $\tilde{\mathbb{E}}$ here, although Section 2.6 treats the case investors hold unbiased beliefs for near-term short rates, despite being biased about forward guidance.

What is the causal effect of z_τ on $g_{\tau+t}$? To isolate the effect of z_τ, we would like to hold all other policies constant *under agents' beliefs*. An alternative calculation holds other policies constant under the objective measure, but in that alternative, a shock to z_τ becomes bundled with shocks to $z_{\tau+1}$, $z_{\tau+2}$, etc. We will explain this alternative calculation in more detail below. Holding policies constant under agents' beliefs means we seek the IRF to a perceived one-time pure short-rate shock: $B\Delta\tilde{W}_\tau = (0,1,0)'$ and $\Delta\tilde{W}_{\tau+s} = 0$ for all $s \geq 1$.[2] This IRF is

$$D_t^{g,z} := \frac{d}{dz}\mathbb{E}\Big[g_{\tau+t} \mid X_{\tau-1},\, z_\tau = z,\, (\Delta\tilde{W}_{\tau+s})_{s=1}^t = 0\Big] = \begin{bmatrix}1\\0\\0\end{bmatrix} \cdot \tilde{A}^t \begin{bmatrix}0\\1\\0\end{bmatrix}. \tag{7}$$

Suppose we want to measure $D_t^{g,z}$ empirically.

[2] For $z_\tau = 1$, without any perceived shocks to g_τ or f_τ, requires $B\Delta\tilde{W}_\tau = (0,1,0)'$.

Consider initially an incorrect but hopefully enlightening approach: run a regression of $g_{\tau+t}$ onto $z_\tau = (0,1,0)B\Delta\tilde{W}_\tau$, the lagged state $X_{\tau-1}$, and a constant. With infinite data, the slope coefficient on z_τ is

$$\beta_t^{g,z} := \frac{\mathrm{Cov}[g_{\tau+t}, z_\tau^*]}{\mathrm{Var}[z_\tau^*]} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \cdot A^t \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \tag{8}$$

where z_τ^* is the residual from a projection of z_τ onto $X_{\tau-1}$. Comparing (7) to (8), the discrepancy arises due to non-rational beliefs (i.e., $A \neq \tilde{A}$).

Understanding the statistical bias. The linear regression, which takes place under the objective probability measure, evidently estimates a rational-expectations IRF. But this deserves more discussion. There are three equivalent ways to understand what exactly the regression picks up:

$$\begin{aligned} \beta_t^{g,z} &= \frac{d}{dz}\mathbb{E}\Big[g_{\tau+t} \mid X_{\tau-1},\, z_\tau^{\text{rational}} = z\Big] \\ &= \frac{d}{dz}\mathbb{E}\Big[g_{\tau+t} \mid X_{\tau-1},\, z_\tau = z\Big] \\ &= \frac{d}{dz}\mathbb{E}\Big[g_{\tau+t} \mid X_{\tau-1},\, z_\tau = z,\, z_{\tau+s} = (0,1,0)'(A-\tilde{A})X_{\tau+s-1}\;\forall s \geq 1\Big], \end{aligned}$$

where $z_\tau^{\text{rational}} := r_\tau - \mathbb{E}[r_\tau \mid X_{\tau-1}]$ is the monetary shock relative to the objective belief. The three objects are all IRFs with different, but equivalent, interpretations.

The first object is an IRF in a rational world: it is the expected response of $g_{\tau+t}$ if agents held rational beliefs and received the surprise $z_\tau^{\text{rational}} = 1$. This is potentially an object of interest if central banks want to ignore the market's beliefs and instead care about responses to their objective perturbations. That said, $z_\tau^{\text{rational}} = 1$ never occurs in reality if market beliefs are distorted, and so this IRF may not possess a structural interpretation.

The second object measures the response when investors receive their actual surprise $z_\tau = 1$ and the economy proceeds afterward as expected by a rational observer. The equivalence to the first object holds because we are controlling for $X_{\tau-1}$, which completely governs the gap between the perceived and rational surprises:

$$z_\tau - z_\tau^{\text{rational}} = \underbrace{\mathbb{E}[r_\tau \mid X_{\tau-1}] - \tilde{\mathbb{E}}[r_\tau \mid X_{\tau-1}]}_{\text{belief distortion, function of } X_{\tau-1}}.$$

This second IRF sounds like a desirable quantity to measure, but it imposes a deceptive structure on future investor surprises. Indeed, the true expectation $\mathbb{E}$ is taken, which effectively holds future objective shocks to zero: $\Delta W_{\tau+s} = 0$ for $s \geq 1$. But if $\Delta W_{\tau+s} = 0$, then by (5) $\Delta \tilde{W}_{\tau+s} = -L_{\tau+s-1}$, which in particular pins down future short-rate surprises at $z_{\tau+s} = -(0,1,0)(\tilde{A} - A)X_{\tau+s-1}$.

The third and final object explicitly notes the structure imposed on future investor surprises. From our perspective, the equivalence of $\beta_t^{g,z}$ to this third object helps illuminate the deficiency in the regression. We are picking up the effect of z_τ bundled with future surprises $z_{\tau+s}$ of a particular relative magnitude. Our initial motivation in hoping to measure $D_t^{g,z}$ was to avoid this "double-counting."

Procedures to obtain the true IRF. What is the solution that identifies the IRF (7)? The first, and most obvious, approach is to measure beliefs directly. For example, suppose we have access to a survey that contains $\tilde{\mathbb{E}}_t[X_{t+j}]$. Regressing these beliefs onto X_t recovers $\tilde{A}^j$ for any j. After obtaining $\tilde{A}$, we can construct the IRF (7) algebraically. In a VAR(1), doing the estimation for $j = 1$ is obviously sufficient (although one may wish to use various survey horizons j to ameliorate the well-known issue that taking powers of $\tilde{A}$ exponentially blows up model misspecification errors).

Procedure 1. *Running a regression of $\tilde{\mathbb{E}}_t[X_{t+1}]$ onto X_t recovers $\tilde{A}$.*

An alternative to Procedure 1 is to obtain proxies for the perceived shocks $\Delta \tilde{W}_t$. A possible advantage of this approach is that perceived shocks might be easier to measure, e.g., from asset prices, than belief levels. For example, we will typically not have access to any survey that contains forecasts of future forward guidance. We explore and confirm, in Section 2.5 for this example and then more generally in Section 3, that the perceived shocks $\Delta \tilde{W}_t$ can sometimes be inferred from asset prices.

In such an approach, suppose we observe a time series of X_t and $\Delta \tilde{W}_t$. Running a regression of X_{t+1} on X_t and $\Delta \tilde{W}_{t+1}$ identifies $\tilde{A}$ as the coefficient on X_t. Even if B were not a diagonal matrix, these perceived shocks can always be conveniently represented in terms of observables by using

$$\begin{aligned}
z_t^g &:= g_t - \tilde{\mathbb{E}}[g_t \mid X_{t-1}] = (1,0,0)B\Delta \tilde{W}_t \\
z_t^r &:= r_t - \tilde{\mathbb{E}}[r_t \mid X_{t-1}] = (0,1,0)B\Delta \tilde{W}_t \\
z_t^f &:= f_t - \tilde{\mathbb{E}}[f_t \mid X_{t-1}] = (0,0,1)B\Delta \tilde{W}_t.
\end{aligned}$$

The object z^r is the same as what we had previously denoted z, but now we add the superscript to distinguish from the perceived shocks to g and f. In our particular environment, recall that agents hold no biases about the shock to g, so z^g need not be included as a control (but in general, one should include z^g). Therefore, we have established the validity of the following procedure:

Procedure 2. *Running a regression of X_{t+1} onto X_t, z^r_{t+1}, and z^f_{t+1} recovers $\tilde{A}$.*

Procedure 2 reveals the two sources of bias. First, even when forward guidance is irrelevant (i.e., even if $a_{gf} = a_{rf} = \tilde{a}_{rf} = 0$), distorted beliefs about r_t lead to biased IRFs. Second, forward guidance compounds the problem by introducing an additional belief distortion. These two sources of bias are eliminated in a one-period-ahead forecast via the controls z^r_{t+1} and z^f_{t+1}. Without these controls, a regression of X_{t+1} on X_t would estimate A; with the controls, we estimate $\tilde{A}$.

Why does Procedure 2 only estimate a one-period-ahead forecast? The reason is that a long-horizon forecast requires an onerous number of controls. For example, if we wanted to estimate the long-horizon IRF (7), we would need to include controls for $(z^r_{\tau+s})^t_{s=1}$ and $(z^f_{\tau+s})^t_{s=1}$, because IRF (7) conditions on $(\Delta\tilde{W}_{\tau+s})^t_{s=1} = 0$. Unless we include all proxies for future belief shocks, the regression coefficient of $g_{\tau+t}$ on z^r_τ would contain omitted variable bias: z^r_τ is correlated with $(z^r_{\tau+s}, z^f_{\tau+s})^t_{s=1}$, precisely because of the dynamical structure of beliefs.[3]

2.4 Numerical example

To get a sense for the consequences of misspecification, we provide a numerical example that is roughly calibrated to the evidence in Cieslak (2018). After calibrating, we illustrate how belief misspecification impacts the IRF of growth to a short-rate shock.

We perform this exercise in two ways. First, we shut down forward guidance (i.e., put $a_{gf} = a_{rf} = \tilde{a}_{rf} = 0$) to illustrate how biases can still arise in that simpler case; if this were the situation, it would be relatively easy to correct the bias by estimating $\tilde{A}$ from a regression of (g_{t+1}, r_{t+1}) on (g_t, r_t) and z^r_{t+1}—i.e., Procedure 2 without f and z^f. Second, we reintroduce forward guidance; not only do additional biases arise, but this is a case that is harder to correct because it requires observation of f and z^f.

[3]For instance, compute the covariance

$$\begin{aligned}
\text{Cov}[z^r_{\tau+s}, z^r_\tau] &= (0,1,0)\text{Cov}[(X_{\tau+s} - \tilde{A}X_{\tau+s-1}), X'_\tau - X'_{\tau-1}\tilde{A}'](0,1,0)' \\
&= (0,1,0)(A-\tilde{A})A^s\text{Cov}[X_{\tau-1}, X'_{\tau-1}](A-\tilde{A})'(0,1,0)' \\
&= (0,1,0)A^{s-1}(A-\tilde{A})A^s\Sigma(A-\tilde{A})'(0,1,0)'
\end{aligned}$$

A similar calculation shows that $z^f_{\tau+s}$ is correlated with z^r_τ. Thus, omitting $z^r_{\tau+s}$ or $z^f_{\tau+s}$, for any s, will introduce omitted variable bias.

To calibrate, we mimic Cieslak (2018) and run the following two regressions in the model:[4]

$$r_{t+j} = \beta_0^{(j)} + \beta_g^{(j)} g_t + \beta_r^{(j)} r_t + \text{residual} \tag{9}$$

$$\tilde{\mathbb{E}}_t[r_{t+j}] = \tilde{\beta}_0^{(j)} + \tilde{\beta}_g^{(j)} g_t + \tilde{\beta}_r^{(j)} r_t + \text{residual} \tag{10}$$

where $\tilde{\mathbb{E}}_t[r_{t+j}]$ denotes subjective expectations of interest rates, obtained from survey data. Furthermore, to roughly match the evidence, we must in this section allow some impact of current growth on future rates (as might occur in a Taylor rule), so A will not be upper triangular here.

Excluding forward guidance. Without forward guidance, the calibration of A and $\tilde{A}$ that roughly matches the evidence in Cieslak (2018) is

$$A = \begin{bmatrix} 0.97 & -0.10 & 0 \\ 0.20 & 0.88 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad \tilde{A} = \begin{bmatrix} 0.97 & -0.10 & 0 \\ 0.07 & 0.96 & 0 \\ 0 & 0 & 0 \end{bmatrix}. \tag{11}$$

Notice that, relative to the econometric transition matrix, subjective beliefs perceive a smaller feedback of growth to future short rates, as well as a higher persistence of r. The comparison between our model and the data is presented in Table 1. Overall, the fit, while not perfect, is good enough to illustrate our main points here.

What are the consequences of subjective beliefs on the measurement of growth IRFs? Figure 1 compares the true IRF versus a mismeasured IRF that regresses future growth on the short-rate shock z_τ^r and economic state $X_{\tau-1}$.

The true causal impact of a short-rate shock is larger and more persistent than suggested by the regression-based procedure. Why? In our calibration, beliefs are extrapolative: the perceived persistence of r is higher than its econometric counterpart. An important property is that agents beliefs overreact to news, and mean-revert later. As discussed previously, the regression-based procedure effectively bundles all the shocks to beliefs, which have offsetting impacts on future growth. But the correct structural IRF holds future beliefs fixed, removing such offsetting. In terms of our notation, we can compare formulas for the true IRF (7) to the mismeasured IRF (8) to see that the bias is

[4] The regression coefficients are given by the following formulas:

$$\beta^{(j)} = \left[\Sigma_{(1:2,1:2)}\right]^{-1} \Sigma_{(1:2,1:3)} [A^j]'_{(2,:)} \quad \text{and} \quad \tilde{\beta}^{(j)} = \left[\Sigma_{(1:2,1:2)}\right]^{-1} \Sigma_{(1:2,1:3)} [\tilde{A}^j]'_{(2,:)}$$

where $\Sigma := (I - A)^{-1} BB' (I - A')^{-1}$ is the stationary variance-covariance matrix of X.

	A. Dependent variable r_{t+j}				B. Dependent variable $\tilde{\mathbb{E}}_t[r_{t+j}]$			
	$j=1$ quarter		$j=4$ quarters		$j=1$ quarter		$j=4$ quarters	
	$\beta_g^{(1)}$	$\beta_r^{(1)}$	$\beta_g^{(4)}$	$\beta_r^{(4)}$	$\tilde{\beta}_g^{(1)}$	$\tilde{\beta}_r^{(1)}$	$\tilde{\beta}_g^{(4)}$	$\tilde{\beta}_r^{(4)}$
Model (no f):	0.200	0.880	0.620	0.501	0.070	0.960	0.250	0.810
Model (yes f):	0.183	0.854	0.617	0.496	0.066	0.938	0.262	0.732
Data:	0.180	0.880	0.650	0.500	0.100	0.930	0.120	0.840

Table 1: Regressions of future short rates (Panel A) and survey-based expectations of future short rates (Panel B) on current growth and short rates (g_t, r_t). For the "Model (no f)" row, the calibrations of A and $\tilde{A}$ are given in (11). For the "Model (yes f)" row, the calibrations of A and $\tilde{A}$ are given in (12). In both models, the calibration of B is $\sigma_g = 0.0025/4$, $\sigma_r = 0.005/4$, and $\sigma_f = 0.010/4$. For the "Data" row, Cieslak (2018) proxies r_t by the Federal Funds Rate (with survey expectations obtained from the Blue Chip Financial Forecasts) and g_t by employment growth.

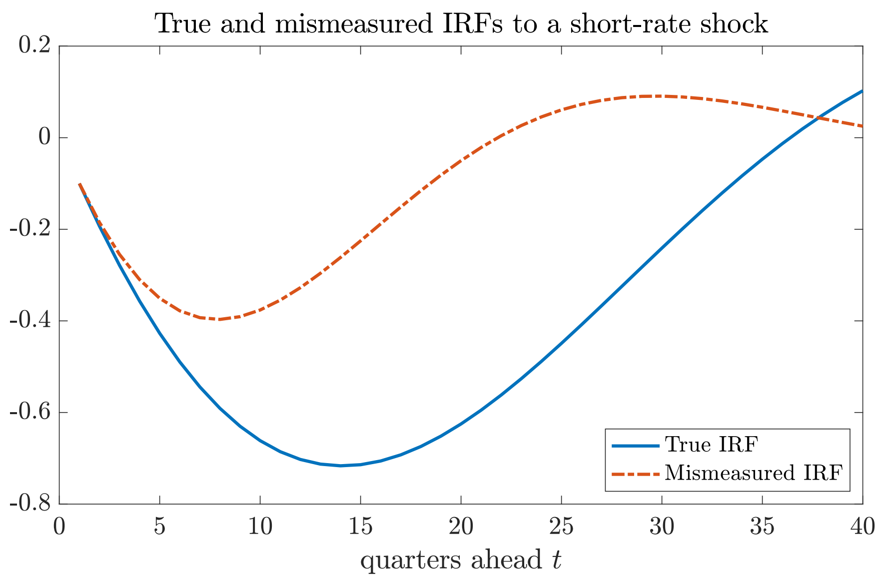

Figure 1: Impulse response functions of $g_{\tau+t}$ to a short-rate shock z_τ^r computed correctly as in (7) versus the mismeasured procedure in (8). The relevant model parameters are given in (11).

$\tilde{A}^t - A^t$, which grows with the horizon t. In Figure 2, this is reflected in the fact that the two IRFs are similar for 2-4 quarters, but diverge significantly in quarters 5-30. This exercise suggests that the impact of monetary policy may be even larger than estimated conventionally, especially in the 2-6 year range.

With forward guidance. With forward guidance reintroduced, we use the following calibration:

$$A = \begin{bmatrix} 0.97 & -0.10 & -0.10 \\ 0.20 & 0.88 & 0.05 \\ 0 & 0 & 0.70 \end{bmatrix} \quad \text{and} \quad \tilde{A} = \begin{bmatrix} 0.97 & -0.10 & -0.10 \\ 0.10 & 0.99 & 0.10 \\ 0 & 0 & 0.99 \end{bmatrix}. \tag{12}$$

The results from this calibration are also displayed in Table 1, row "Model (yes f)". Similar to the model that excludes forward guidance, we are again roughly able to match the evidence of Cieslak (2018). We then redo the comparison between the true growth IRF and its mismeasured counterpart, with the results in Figure 2. The discrepancy between the true and mismeasured IRFs is broadly similar to what we found without forward guidance, but slightly less extreme. Again, the reasoning is that subjective beliefs are extrapolative: $\tilde{A}$ contains larger persistences of r and f than the objective transition matrix A.

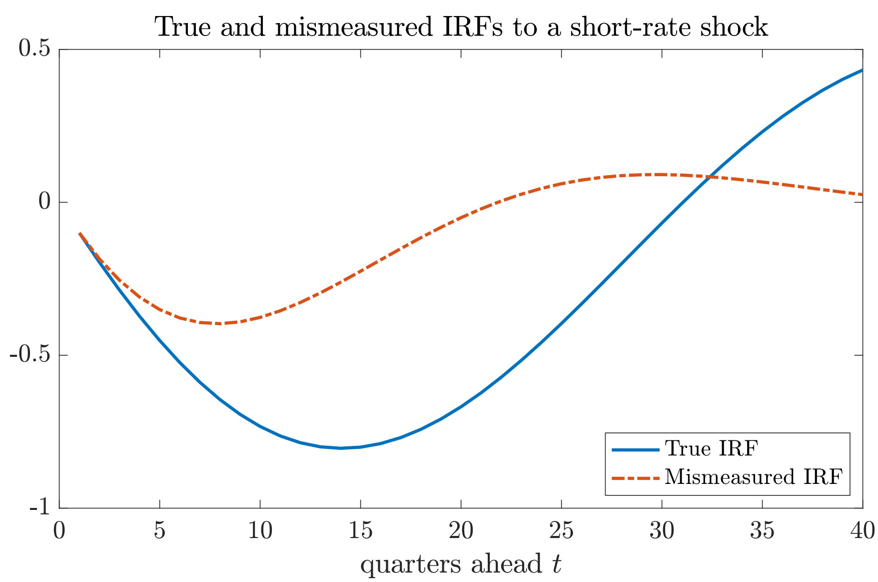

Figure 2: Impulse response functions of $g_{\tau+t}$ to a short-rate shock z^r_τ computed correctly as in (7) versus the mismeasured procedure in (8). The relevant model parameters are given in (12).

What does this calibration imply about the efficacy of forward guidance? We repeat the computation of true and mismeasured growth IRFs, but this time to the forward-guidance shock. The results are in Figure 3. The huge discrepancy, even larger than Figure 2, arises primarily because the subjective persistence of f is much larger than its objective persistence. Of course, the calibration of the dynamics of f and its interaction with r and g is done only using indirect evidence (in particular, we have not directly observed f, something we will try to correct in the next section). But if we take this numerical example as suggestive, the effects of forward guidance may be measured even less accurately than the effects of a short-rate shock.

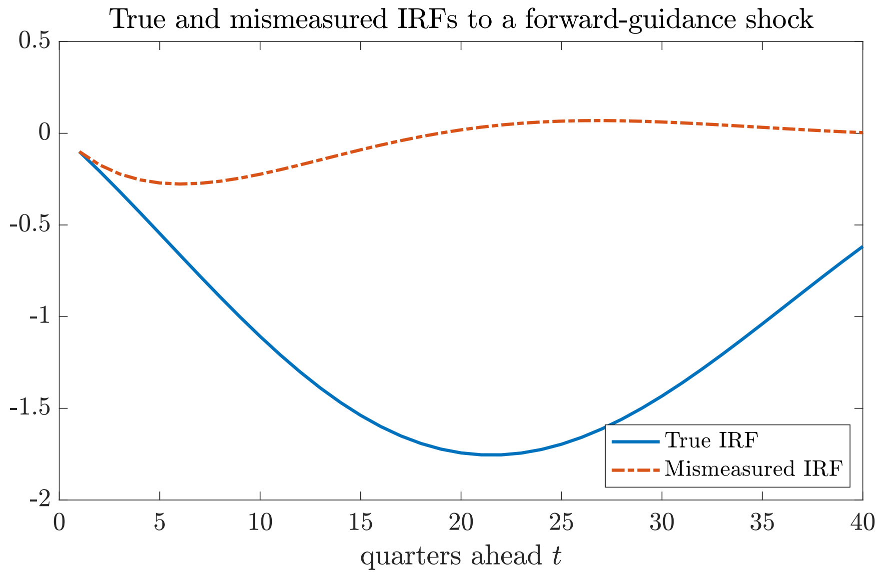

Figure 3: Impulse response functions of $g_{\tau+t}$ to a forward-guidance shock z_{τ}^{f} computed correctly as in (7) versus the mismeasured procedure in (8). The relevant model parameters are given in (12).

Discussion. Can the true IRF be recovered by adding controls to "clean up" the monetary shock? Some papers argue that adding enough appropriate controls for central bank information is required to get an appropriate monetary policy instrument (Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023a,b). Does a similar procedure help with distorted beliefs?

We clarify our connection to this literature by making two points. First, we have assumed throughout that the short-rate shock z_{τ}^{r} is measured perfectly. In other words, we sidestep the main problem discussed in this literature—effectively that enough $X_{\tau-1}$ variables have not been projected out when measuring z_{τ}^{r}. Furthermore, adding controls for beliefs does not "fix" the IRF here: in our environment the belief distortion $L_t = B^{-1}(\tilde{A} - A)X_t$ is a purely function of this state vector, which we are already controlling for in both the true and mismeasured IRFs.

Second, under distorted beliefs, it is generally true that the short-rate shock z_{τ}^{r} will be autocorrelated, a feature that the "clean up" procedure is designed to fix. The key point here is that such autocorrelation is a property of the underlying model, and should not be cleaned. By contrast, papers on the Fed information effect and response-to-news effect assume rational expectations, so autocorrelation should be purged.

2.5 Identifying forward guidance from asset prices

To implement Procedure 2, we somehow need to recover a time series for f, as well as its perceived shock z^f. Since f is primarily about future short-term interest rates, observation of the market's expectation of future short rates should suffice in place of f. A natural place to look for these market beliefs are Treasury bond markets. This leads us, in the context of our simple example, to a key identification challenge of the paper: can we infer expected future short rates from asset-price data?

The problem with using asset markets is that they do not directly reveal the market's expectation, but rather a risk-adjusted expectation. To address this discrepancy, we will first write down a standard affine term structure model (Duffie and Kan, 1996; Dai and Singleton, 2002; Duffee, 2002; Ang and Piazzesi, 2003) and then impose sufficient structure on the model. In doing so, we will be explicit about which conditions allow us to invert risk adjustments and recover the market expectation.

The asset-pricing model in this example features zero inflation for simplicity, and so the distinction between real and nominal is immaterial here. To the state dynamics in Section 2.1, we add the following one-period stochastic discount factor (SDF):

$$\frac{S_{t+1}}{S_t} = \exp\Big[-(\bar{r} + r_t) - \frac{1}{2}\pi_t'\pi_t - \pi_t'\Delta W_{t+1} \Big], \tag{13}$$

where $\pi_t = \pi_0 + \Pi X_t$ is the time-varying risk price vector. Notice that the SDF is specified under the objective probability measure. This is immaterial and only written this way to conform with the bond pricing literature.

We can also re-write the SDF in terms of the investor measure:

$$\begin{aligned} \frac{S_{t+1}}{S_t} &= \frac{\tilde{S}_{t+1}}{\tilde{S}_t} \exp\Big[L_t'\Delta W_{t+1} - \frac{1}{2}L_t'L_t \Big] \\ \text{where} \quad \frac{\tilde{S}_{t+1}}{\tilde{S}_t} &= \exp\Big[-(\bar{r} + r_t) - \frac{1}{2}(\pi_t + L_t)'(\pi_t + L_t) - (\pi_t + L_t)'\Delta\tilde{W}_{t+1} \Big]. \end{aligned} \tag{14}$$

The variable $\exp[L_t'\Delta W_{t+1} - \frac{1}{2}L_t'L_t]$ changes the probability measure from the objective one to investors' subjective one, while $\tilde{S}$ represents investor marginal utility. Notice that investors' perceived risk prices are $\pi_t + L_t$.

In this conditionally log-normal setting, the risk-neutral dynamics of the state vector are given by

$$\begin{aligned} X_{t+1} &= A_0^* + A^* X_t + B\Delta W_{t+1}^*, \\ \text{where} \quad A_0^* &:= -B\pi_0 \\ \text{and} \quad A^* &:= A - B\Pi, \end{aligned}$$

where ΔW_{t+1}^* is a Normal shock under the risk-neutral distribution. This framework can be used to solve for bond prices of all maturities. The equilibrium yield-to-maturity for an n-period risk-free zero-coupon bond is given by

$$\begin{aligned} y_t^{(n)} &= \frac{1}{n}\left[\mathcal{B}_0^{(n)} + \mathcal{B}^{(n)} X_t\right] \\ \text{where} \quad \mathcal{B}^{(n)} &= (1,0)(I - A^*)^{-1}(I - (A^*)^n), \\ \mathcal{B}_0^{(n)} &= n\bar{r} + \Big(\sum_{i=1}^{n-1} \mathcal{B}^{(i)}\Big) A_0^* - \frac{1}{2}\sum_{i=1}^{n-1} \mathcal{B}^{(i)} BB'(\mathcal{B}^{(i)})'. \end{aligned}$$

This solution is standard in the literature.

We use the model solution to identify f. In general, since bond yields are affine in the factors, we should be able to invert for the factor time series, given data on any three maturities. The setting here is even simpler because the growth rate and one-period yield (short rate) are presumed to be observable. So if we have a single n-period bond ($n > 1$), we can use its yield to obtain the state vector as

$$X_t = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ & \mathcal{B}^{(n)} & \end{bmatrix}^{-1} \begin{bmatrix} g_t \\ r_t \\ n y_t^{(n)} - \mathcal{B}_0^{(n)} \end{bmatrix}.$$

This expression suggests that f can be obtained via data on bond yields and an estimation of the asset-pricing model. Indeed, $n^{-1}\mathcal{B}_0^{(n)}$ is just the unconditional mean of $y_t^{(n)}$ and is easily obtained, while $\mathcal{B}^{(n)}$ can presumably be estimated using the joint dynamics over time of r_t and $y_t^{(n)}$.

This logic is more-or-less correct with some small caveats that do not concern us in this paper. The small caveat is that X_t itself is not recovered but a linear transformation $\hat{X}_t = UX_t$, for some matrix U. This is a well-known identification issue in affine term-structure models with latent state variables like our forward guidance variable.[5]

[5]The deep reasoning for non-identification is that, when some of the factors are latent and not directly observable, a continuum of true dynamics (A, B) are consistent with the same risk-neutral dynamics (A_0^*, A^*, B). However, only the risk-neutral dynamics matter for asset-pricing purposes, evident in the

For our purposes, observing a linear combination of the factors works just as well as observing the factors directly. Indeed, to implement Procedure 2, we only need enough additional variables to span the state X_t and the perceived shocks $\Delta\tilde{W}_{t+1}$. Since $\hat{X}_t$ spans the same space as X_t, this minor caveat causes no problems for us. Therefore, to avoid distractions and maintain our existing notation, we assume that the true state X_t is actually recovered and continue to write X_t rather than $\hat{X}_t$.

Having (essentially) recovered X_t, can we recover the perceived forward-guidance shock $z_t^f := f_t - \tilde{\mathbb{E}}[f_t \mid X_{t-1}] = (0,0,1)B\Delta\tilde{W}_t$? There are two fairly obvious situations under which this is possible:

(i) Investors have rational expectations;

(ii) Investors are risk-neutral.

Under rational expectations, recovery of X_t is sufficient to estimate the VAR (1) and thereby recover the shocks $B\Delta W_t$. In any case, shock recovery is not strictly speaking required, because $D_t^{g,z} = \beta_t^{g,z}$ under rational expectations. Under risk-neutrality, investor marginal utility features zero risk-pricing ($\pi_t + L_t = 0$), so the recovered risk-neutral dynamics actually correspond to investors' perceived dynamics:

$$A_0^* + A^* X_t = -B\pi_t + AX_t = BL_t + AX_t = \tilde{A}X_t.$$

So the risk-neutral shocks $B\Delta W_{t+1}^*$ actually coincide with the subjective shocks $B\Delta\tilde{W}_{t+1}$. Risk-neutrality allows recovery because it implies an "Expectations Hypothesis" but under the investor subjective belief.

A third, less obvious, condition also permits shock recovery. Assume

$$\Pi = -B^{-1}(\tilde{A} - A) \tag{15}$$

Condition (15) says that investor perceived risk prices $\pi_t + L_t$ are time-invariant (the coefficient on X_t is zero). While time-invariance may seem quite restrictive, it turns out that—if investors cannot be assumed rational or risk-neutral—it is both necessary and sufficient for identification of $B\Delta\tilde{W}$ in this model.

To understand sufficiency is fairly easy. Substitute (15) into the risk-neutral dynamics to obtain

$$A^* = A - B\Pi = A + BB^{-1}(\tilde{A} - A) = \tilde{A}.$$

expressions for $\mathcal{B}^{(n)}$ and $\mathcal{B}_0^{(n)}$. This identification issue is solely due to the latency of X_t and is a well-known feature of such models (Hamilton and Wu, 2012).

If we know A^*, then we know $\tilde{A}$, which allows us to obtain perceived shocks as $X_{t+1} - \tilde{A}X_t = B\Delta\tilde{W}_{t+1}$. Intuitively, if investor perceived risk prices are constant, then a version of the Expectations Hypothesis holds: long-term bond yields capture investor beliefs about future short-term yields, with a constant shifter.

Necessity is harder to see. The important fact is that condition (15) is required to make investor-perceived long-run risk prices constant. To see this, we follow the calculations in Backus et al. (2022) to compute the permanent component of the investor SDF $\tilde{S}$ in (14) as

$$\frac{H_{t+1}}{H_t} = \exp\Big[-\frac{1}{2}\|\pi_t + L_t - B'v\|^2 - (\pi_t + L_t - B'v)\cdot\Delta\tilde{W}_{t+1}\Big],$$

where $v := -(I - (\tilde{A} - B\Pi)')^{-1}(0,1,0)'$. In our general environment of Section 3, we explain this permanent component in more detail, and we show that perceived shock identification requires H_{t+1}/H_t to be independent of X_t (under investor beliefs). Taking that result as given, identification thus requires $\pi_t + L_t = \pi_0 + (\Pi + B^{-1}(\tilde{A} - A))X_t$ to be independent of X_t, which translates to condition (15).

As it turns out, our assumptions on the model structure in this example sneakily permit looser identification conditions than the more general case covered in Section 3. In particular, we have assumed a very strong structure where *monetary shocks are iid*: they happen every period and with the same time-invariant Normal distribution. As we will show in our more general model of Section 3, these type of iid monetary shocks are a knife-edge case without a reasonable intuition. In reality, we expect monetary policy actions to depend on the state of the economy. Outside of this knife-edge case of random monetary interventions, shock identification will require the even stronger condition that H_{t+1}/H_t be *invariant to policy shocks* (in our example here, that is $\tilde{W}^{(2)}_{t+1}$ and $\tilde{W}^{(3)}_{t+1}$).

2.6 Unbiased short-rate expectations

What if $z^r_\tau = z^{r,\text{rational}}_\tau$ but beliefs about future interest rates are not unbiased? This might be a reasonable position if near-term objects are easier to forecast than longer-term outcomes. In our framework, this is accommodated by setting $\tilde{a}_{rr} = a_{rr}$ and $\tilde{a}_{rf} = a_{rf}$. Then, the only bias in beliefs is about the persistence of the forward guidance process. (Although as mentioned earlier, the evidence in Cieslak (2018) is inconsistent with this view.)

The regression coefficient $\beta_t^{g,z}$ can now be interpreted as estimating the IRF

$$\beta_t^{g,z} = \frac{d}{dz}\mathbb{E}\Big[g_{\tau+t} \mid X_{\tau-1},\, z_\tau^r = z,\, z_{\tau+s}^r = 0\, \forall s \geq 1\Big],$$

which is much closer to what we want. This object measures the effect of a short-rate shock $z_\tau = 1$ holding all future short-rate shocks to zero. However, there is still a discrepancy between $\beta_t^{g,z}$ and the desired IRF in (7).

The problem: belief distortions about forward guidance imply our regression effectively sets future surprises to non-zero levels. If $\Delta W_{\tau+s} = 0$, then $B\Delta\tilde{W}_{\tau+s} = -L_{\tau+s-1}$, so the forward-guidance shock would be $z_{\tau+s}^f = (0,0,1)B\Delta\tilde{W}_{\tau+s} = (a_{ff} - \tilde{a}_{ff})f_{\tau+s-1}$. Thus, our regression can equivalently be interpreted as measuring

$$\beta_t^{g,z} = \frac{d}{dz}\mathbb{E}\Big[g_{\tau+t} \mid X_{\tau-1},\, z_\tau^r = z,\, z_{\tau+s}^r = 0,\, z_{\tau+s}^f = (a_{ff} - \tilde{a}_{ff})f_{\tau+s-1}\, \forall s \geq 1\Big].$$

We are picking up the impact of z_τ^r bundled with the expectation of future forward-guidance surprises. If forward guidance has a direct impact on growth, this bundled effect will differ from the impact of short rates in isolation.

3 Identifying Monetary Shocks in a Quasi-Linear World

We now consider the question of whether and how to recover investor surprises about future interest rates. Our setting, based on Hansen and Scheinkman (2009) and Borovička et al. (2016), is Markovian and has complete financial markets. We will work in continuous time, for several reasons. First, continuous time allows us to more naturally delineate between "typical shocks" that occur all the time and "monetary shocks" that occur only at specific dates. Second, we can obtain our results even allowing for some types of nonlinearities in continuous time, which is desirable if we would like to think not only about expected future interest rates but also rate uncertainty.

3.1 General setup

Beliefs. Let the probability measure $\mathbb{P}$ represent investor beliefs. Rational expectations is not assumed: $\mathbb{P}$ may or may not coincide with the true objective probability. We work exclusively in the realm of investors' subjective beliefs, because our goal when thinking about shock identification is to identify changes in interest rates relative to the market beliefs. Since we will never be referencing the objective probability measure, we will always use $(\mathbb{P}, \mathbb{E})$ for investor beliefs rather than the Section 2 notation $(\tilde{\mathbb{P}}, \tilde{\mathbb{E}})$.

The econometrician does not know the investor beliefs $\mathbb{P}$. To state the problem of the econometrician, he wants to learn monetary shocks—which will be policy surprises relative to $\mathbb{P}$—using only data on asset prices.

States, shocks, and information. There is stationary n-dimensional economic state X. The evolution of X is perturbed by two types of shocks. First, there are *non-monetary shocks* that occur continuously. Non-monetary shocks are modeled by the increments to W, which is an n-dimensional Brownian motion under $\mathbb{P}$. We could have included more of these shocks than state variables, but supposing they are the same number, as in most empirical applications, will streamline our arguments.[6]

Second, there are *monetary shocks* that occur only at specific times. To preserve a stationary and Markovian structure of our economy, we assume these times arrive according to a Poisson process with rate $\lambda(X_{t-})$, which can depend on the state. Whereas monetary announcement dates are deterministic and known in advance, one can think of randomness in these dates as capturing announcements during which some surprises actually occur. Furthermore, during some times of crisis, emergency actions and statements by the central bank can take place. We let M_t be the counter for announcements, so $dM_\tau = 1$ if and only if τ is an announcement date.

At these announcement dates, monetary shocks are modeled by the n-dimensional vector ξ_t. This random variable is independent of W and dictates the jump in the state variable: $X_t - X_{t-} = \xi_t dM_t$. Investors' perceived probability distribution of ξ_t is allowed to depend on the state X_{t-} just prior.

Subject to these two types of shocks, the state vector evolves as the jump-diffusion

$$dX_t = \mu(X_{t-})dt + \sigma(X_{t-})dW_t + \xi_t dM_t. \tag{16}$$

The sequence of information sets $(\mathscr{F}_t)_{t\geq 0}$ available to investors is generated by histories of W, M, ξ, and the initial condition X_0. In other words, investors observe $(X_t)_{t\geq 0}$. We will assume the same information set for the econometrician. (Thus, we sidestep the issue encountered in Section 2 whereby the econometrician would need to recover observations of X_t from asset prices plus an asset-pricing model.)

Asset prices and the SDF. We assume there exists a stochastic discount factor (SDF)

[6] Borovička et al. (2016) allow for $k > n$ shocks by adding more observables to the states in X.

process S whose increment is given by

$$\frac{dS_t}{S_{t-}} = -r(X_{t-})dt - \pi(X_{t-}) \cdot dW_t + \exp[\kappa(X_t, X_{t-})] - 1 - \chi_S(X_{t-})dt, \tag{17}$$

where $S_0 = 1$ and $\chi_S(x)dt$ is the jump compensator.[7] The variable $r(x)$ denotes the short-term interest rate, while $\pi(x)$ and $\kappa(x', x)$ denote risk prices associated to the non-monetary and monetary shocks (note that $\kappa(x, x) = 0$). Using the SDF, the date-t price of any payoff $f(X_T)$ is

$$\mathbb{E}\Big[\frac{S_T}{S_t} f(X_T) \mid X_t\Big]. \tag{18}$$

In this environment, Hansen and Scheinkman (2009) show how Perron-Frobenius Theory can be leveraged to obtain a decomposition of the SDF as

$$\frac{S_{t+T}}{S_t} = \exp(\eta T)\frac{e(X_t)}{e(X_{t+T})}\frac{H_{t+T}}{H_t}. \tag{19}$$

In (19), $\exp(\eta)$ is a positive eigenvalue of the instantaneous pricing operator; $e(\cdot)$ is the associated positive eigenfunction; and H_t is a martingale under $\mathbb{P}$. We assume existence of an SDF decomposition (19). For the purposes of this section, we also assume the decomposition is unique.[8]

Equation (19) decomposes the SDF into a deterministic component, a stationary component, and a permanent component. Some sources of H arising in structural representative-agent models are the permanent component of aggregate consumption or continuation value fluctuations in models with Epstein-Zin preferences and persistent growth or stochastic volatility. We flesh out some examples in Section 5. Using non-parametric methods, Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012) argue that H must play a significant role in pricing.

[7]More formally, let ν denote the random counting measure such that $\nu(\mathcal{B}, [0, t])$ gives the random number of jumps in time interval $[0, t]$ having size in the Borel set $\mathcal{B}$. We restrict attention to processes with a finite number of jumps in any finite time interval. Then, the compensator is the random measure $\chi(dx' \mid x)dt$ such that for any predictable function $g(x, t)$, the process $\int_0^t \int_{\mathbb{R}^n} g(x', s)\nu(dx', ds) - \int_0^t \int_{\mathbb{R}^n} g(x', s)\chi(dx' \mid X_{s-})ds$ is a martingale. With this notation, we define $\chi_S(x) := \int \big(\exp[\kappa(x', x)] - 1\big)\chi(dx' \mid x)$.

[8]See Hansen and Scheinkman (2009) for sufficient conditions on the existence of such a decomposition. In many cases, uniqueness will not hold. If there are multiple SDF decompositions satisfying (19), we follow Proposition 1 of Borovička et al. (2016) in picking the unique one such that X is stationary and ergodic under the probability measure $\mathbb{P}^H$ induced by the martingale H (i.e., defined by $\mathbb{P}^H(A) := \mathbb{E}(1_A H_T)$ for all sets $A \in \mathscr{F}_T$, for any $T \geq 0$). For the purposes of this section, non-uniqueness will not be relevant to the monetary policy questions, which is why we sidestep these issues.

Reduced-form monetary shocks. Monetary actions perturb this environment at an announcement date τ through the shocks $\xi_\tau = X_\tau - X_{\tau-}$. The new state vector then feeds into current and future interest rates. But for our purposes, we will define monetary shocks directly in terms of their effect on interest rates. First, monetary policy can influence the short-term interest rate, which is given by $r(X_\tau)$. Second, policy can influence the sequence of future interest rates, namely $r(X_{\tau+T})$. Obviously, these future interest rates are random variables: altering future interest rates involves not only modifying the expected future rate path, but potentially also the entire probability distribution of future interest rates.

Our reduced-form monetary policy shocks are defined as follows.

Definition 1. *Suppose the central bank intervenes at time τ. The short rate shock is given by*

$$z_\tau^0 := r(X_\tau) - \mathbb{E}[r(X_\tau) \mid X_{\tau-}]. \tag{20}$$

The shocks to the expected future short rates are given by

$$z_\tau^T := \mathbb{E}[r(X_{\tau+T}) \mid X_\tau] - \mathbb{E}[r(X_{\tau+T}) \mid X_{\tau-}], \quad T > 0. \tag{21}$$

The shocks to the distribution of future short rates are given by

$$p_\tau^T(\mathsf{r}) := \mathbb{P}\{r(X_{\tau+T}) \le \mathsf{r} \mid X_\tau\} - \mathbb{P}\{r(X_{\tau+T}) \le \mathsf{r} \mid X_{\tau-}\}, \quad T > 0. \tag{22}$$

Can an econometrician identify the impacts of the central bank on current and future interest rates? First, as mentioned in the introduction and Section 2, let us take as given the ability to non-parametrically identify the short rate shock. In particular, we observe the value of $r(X_\tau)$ at time τ, and suppose we also observe $\mathbb{E}[r(X_\tau) \mid X_{\tau-}]$, presumably from a financial market (e.g., Fed Funds futures). Implicitly, this assumes risk prices are sufficiently small for short-horizon interest rates, such that the risk-neutral expectation $\mathbb{E}^*[r(X_\tau) \mid X_{\tau-}]$ coincides with the investor expectation. So let us assume z_τ^0 is observable. But following the logic of Section 2, we presume that identifying the causal impact of z_τ^0 requires controlling for surprises to the future interest rate path, via z_τ^T and p_τ^T.

Turning to z_τ^T and p_τ^T, we cannot use the same identification logic as with z_τ^0. The financial market still allows us to observe the risk-neutral expectations $\mathbb{E}^*[r(X_{\tau+T}) \mid X_\tau]$ and $\mathbb{E}^*[r(X_{\tau+T}) \mid X_{\tau-}]$, but the presence of risk premia embedded in longer-term interest rate futures implies $\mathbb{E}^* \neq \mathbb{E}$ when applied to future interest rates. Similarly, because $\mathbb{P}^* \neq \mathbb{P}$, we cannot expect the financial market to reveal the shock to the entire distribution of future short rates in (22).

What do financial markets reveal? Nevertheless, it turns out that z_τ^T and p_τ^T may sometimes be identified from financial market data. The basic idea, building on Borovička et al. (2016), is that the martingale H in the decomposition (19) may be used as a change-of-measure from investor beliefs $\mathbb{P}$ to the long-run risk-neutral measure $\hat{\mathbb{P}}$, defined by

$$\hat{\mathbb{P}}(F) := \mathbb{E}[1_F H_t], \quad \forall F \in \mathscr{F}_t. \tag{23}$$

It turns out that asset prices reveal this probability measure, as the next lemma verifies.

Lemma 1. *The econometrician observes* $\hat{\mathbb{P}}\{r(X_{\tau+T}) \leq \mathsf{r} \mid X_\tau\}$ *for every* r *and every* τ, T.

Except under the very particular degenerate situation $H \equiv 1$, investor beliefs $\mathbb{P}$ will not coincide with the recovered $\hat{\mathbb{P}}$, as explained by Borovička et al. (2016). But our goal is less ambitious. We do not seek $\mathbb{P}$ directly but rather investor surprises or *belief shocks*. As long as the gap, in some sense, between $\mathbb{P}$ and $\hat{\mathbb{P}}$ remains constant before and after monetary policy announcements, we may hope that

$$\hat{\mathbb{E}}[r(X_{\tau+T}) \mid X_\tau] - \hat{\mathbb{E}}[r(X_{\tau+T}) \mid X_{\tau-}] = \mathbb{E}[r(X_{\tau+T}) \mid X_\tau] - \mathbb{E}[r(X_{\tau+T}) \mid X_{\tau-}], \tag{24}$$

and similarly for other moments of $r(X_{\tau+T})$. If equality (24) were to hold, then we would be done: Lemma 1 proves that the left-hand-side is observable, so we will have inferred the belief shocks on the right-hand-side.

The key question is which conditions permit this procedure. As we will see, one critical condition is that policy cannot affect the permanent component of the SDF.

Definition 2. *We say that monetary policy possesses* Long-Run Neutrality *if*

(i) The evolution of $d\log(H_t)$ *is independent of the monetary shock* dM_t*;*

(ii) The evolution of $d\log(H_t)$ *is independent of the economic state* X_t.

Definition 2 specifies neutrality in terms of asset markets, via the martingale H. Condition (i) rules out direct effects of policy on the long-run SDF. Condition (ii) rules out indirect policy effects, which can be understood as follows. We are imagining a world in which monetary policy can have real effects and therefore generically affects the state vector X. In that case, if condition (ii) failed, then policy would indirectly affect H_{t+T} by moving X_t. For H to satisfy Definition 2, it must take the form

$$H_t = \exp\Big[-\frac{1}{2}\|\beta\|^2 t - \beta \cdot W_t\Big] \tag{25}$$

for some n-dimensional vector β. One can interpret β as the constant long-run risk price associated to non-monetary shocks. (Of course, by condition (i), there is a zero long-run risk price for monetary shocks.)

In the next two subsections, we illustrate how Long-Run Neutrality sometimes allows us to identify the shocks in Definition 1. After showing these positive identification results, we will explain how identification fails in some example environments without Long-Run Neutrality.

3.2 Exact identification: linear case

To start, we will make several assumptions such that the entire economy is linear. First, we assume the state dynamics are given by

$$\mu(x) = A_0 + Ax \tag{26}$$

$$\sigma(x) = B, \tag{27}$$

for some $n \times 1$ vector A_0, and $n \times n$ matrices A and B. Second, we assume that the short-term interest rate r is a linear function:

$$r(x) = \rho_0 + \rho \cdot x, \tag{28}$$

for some constant ρ_0 and some vector ρ. Assuming (28) holds in a linear environment with (26)-(27) is tantamount to an assumption on the evolution of S (for example, the exponential-affine model of Section 2.5 had affine bond yields). Alternatively, one could just think that the short rate r_t is one of the state variables in X_t.

With linear-Gaussian dynamics and a linear interest rate function, we no longer have to think about the uncertainties in future rates (captured by p_τ^T in Definition 1). The variance of the future state vector X_{t+T}, conditional on X_t, is a deterministic function of T. The same holds for all higher moments of X_{t+T}. Still, one wonders whether the forward-guidance shock z_τ^T is identified.

Proposition 1. *Suppose Long-Run Neutrality holds. Consider the linear environment defined by (26)-(28). Then, the forward-guidance shocks $(z_\tau^T)_{T\geq 0}$ are identified from asset price data alone.*

One may be skeptical that belief shocks can be identified at all. Going back to the fundamental identification issues raised by Harrison and Kreps (1979), asset prices do not directly reveal beliefs. More recently, Borovička et al. (2016) argued that beliefs are only revealed if $H_t \equiv 1$ is a degenerate martingale. In our context, how is identification is possible under seemingly weaker assumptions? The key simplification is that our environment is linear, whereas these previous papers have tried to argue non-parametrically. Imposing this stronger assumption on the economic dynamics allows us to weaken the conditions on H for shock identification.

However, even a linear environment is not enough. An additional simplification is that Proposition 1 does not seek beliefs directly, but rather surprises or *belief shocks*. It is easier to recover belief shocks, because they difference out any unobservable level effect in beliefs. Indeed, that is exactly what happens in the proof of Proposition 1.

Let us briefly elaborate on the method to identify z_τ^T. First, by solving an eigenvalue problem, one can use asset prices to recover the long-run risk-neutral probability measure $\hat{\mathbb{P}}$, as demonstrated by Lemma 1. (See also Ross (2015), Borovička et al. (2016), and Qin and Linetsky (2017) for this result more generally). The dynamics of X_t under $\hat{\mathbb{P}}$ are the same as those under $\mathbb{P}$, less the constant drift $B\beta$:

$$\hat{\mu}(x) = \mu(x) - B\beta = A_0 - B\beta + Ax$$

While β and A_0 are not separately identified, the long-term measure $\hat{\mathbb{P}}$ correctly identifies investors' perceived persistence A. This turns out to be the critical necessary object to compute investors' forecast revisions. By contrast, constant drift distortions like $B\beta$ play no role in these forecast revisions, because investor forecasts just before and just after the monetary announcement are both distorted by the same constant. In other words, the computable object $\hat{\mathbb{E}}[X_{\tau+T} \mid X_\tau] - \hat{\mathbb{E}}[X_{\tau+T} \mid X_{\tau-}]$ coincides with the desired investor forecast revision $\mathbb{E}[X_{\tau+T} \mid X_\tau] - \mathbb{E}[X_{\tau+T} \mid X_{\tau-}]$.

As a side note, we mention that using asset prices to recover $\hat{\mathbb{P}}$ and hence the perceived persistence matrix A allows us to circumvent performing Procedure 2. Instead, our proposed methodology in this section automatically recovers the IRFs to any monetary shock.

Ultimately, Proposition 1 is just a generalization of what we observed in our example in Section 2.5. But it is convenient that we can phrase the result in terms of the Long-Run Neutrality condition, which will be the center-piece of our emphasis going forward.

3.3 Approximate identification with stochastic volatility

We continue to assume a linear drift (26) and a linear short rate function (28), but we dispense with homoskedasticity (27). In such a world, the perceived probability distribution of $r(X_{\tau+T})$ becomes non-trivial (i.e., it is not fully characterized by its mean and the horizon T). And so we would ideally like to estimate the uncertainty shocks p_τ^T in addition to the forward-guidance shocks z_τ^T.

To proceed in this more general environment, we need an extra assumption. Roughly speaking, we need to assume that the sources of heteroskedasticity are not priced by the long-run risk-neutral measure. Supposing Long-Run Neutrality holds, so that equation (25) characterizes the permanent component of the SDF, we assume there exists some n-dimensional vector $\hat{\beta}$ such that

$$\sigma(x)\beta = \hat{\beta} \quad \text{for all } x. \tag{29}$$

In other words, there is a zero in each element of β corresponding to a shock with non-constant volatility. Replacing homoskedasticity assumption (27) with the more general (29), we are still able to identify the forward-guidance shocks but not the uncertainty shocks. Formally, we have the following generalization of Proposition 1.

Proposition 2. *Suppose Long-Run Neutrality holds. Consider the quasi-linear environment defined by (26), (28), and (29). Then, the forward-guidance shocks $(z_\tau^T)_{T\geq 0}$ are identified from asset price data alone.*

The key intuition for Proposition 2 is the same as Proposition 1. Indeed, (29) implies that the drift of X_t under the long-run measure $\hat{\mathbb{P}}$ is

$$\hat{\mu}(x) = \mu(x) - \sigma(x)\beta = A_0 - \hat{\beta} + Ax.$$

As in Proposition 1, investors' perceived persistence A can be inferred from financial data, which is the critical necessary object to compute investors' forecast revisions.

Unfortunately, in the environment considered by Proposition 2, the uncertainty shocks p_τ^T are non-trivial and non-identified. In some applications, we may have a priori reasons to care less about p_τ^T. But in situations where uncertainty matters, we will want to recover p_τ^T.

To make partial progress, we make the following linearity assumption about the form of the state diffusion:

$$\sigma(x)\sigma(x)' = \varsigma_0\varsigma_0' + \sum_{i=1}^{n} \varsigma_i \mathrm{diag}(x_i)\varsigma_i', \tag{30}$$

where $\mathrm{diag}(x_i)$ is the diagonal matrix with x_i on the main diagonal. The affine approximation in (30) is consistent with standard stochastic volatility models having "square-root dynamics." With this structure, we can at least identify shocks to *investors' perceived variance* of future interest rates, even if we cannot recover the entire probability distribution of $r(X_{\tau+T})$. (Indeed, one can verify that the same method of proof used in Proposition 3 does not work for third and higher moments.)

Proposition 3. *Suppose Long-Run Neutrality holds. Consider the quasi-linear environment defined by (26), (28), (29), and (30). Define the variance surprises*

$$v_\tau^T := Var[r(X_{\tau+T}) \mid X_\tau] - Var[r(X_{\tau+T}) \mid X_{\tau-}].$$

Then, $(v_\tau^T)_{T\in[0,\tau'-\tau)}$ are identified from asset price data alone, where τ' is the subsequent monetary announcement date after τ.

Together, Propositions 2-3 demonstrate that a forward-guidance shocks and some aspects of uncertainty shocks, at least those pertaining to variances, can be obtained from asset-market data. We require assumptions both on the dynamic evolutions and on the underlying economic model. The key assumption on the dynamics is quasi-linearity, with variance dynamics taking a "square-root" form. The critical economic assumption in all cases is Long-Run Neutrality, along with assumption (29) that volatility shocks feature zero long-run risk prices.

3.4 Non-identification without Long-Run Neutrality

We now provide some examples to illustrate why shock recovery requires Long-Run Neutrality. To provide the best possible chance at achieving identification, let us specialize to the linear setup defined by (26)-(28) in Section 3.2. First, we consider a world where monetary policy affects H directly (violating condition (i) of Definition 2). Second, we consider a world where monetary policy indirectly affects H through its impact on the state vector X (violating condition (ii) of Definition 2). In either environment, monetary policy shocks are generally not identified from asset prices alone.

Direct monetary effects. Consider what happens if dH_t is directly impacted by the monetary shock dM_t. For simplicity, suppose the non-monetary shocks W_t do not impact H_t at all. We will furthermore assume that the jumps in H are log-linear in the jumps in X. The evolution of H_t then takes the form

$$\frac{dH_t}{H_{t-}} = \exp[\zeta \cdot (X_t - X_{t-})] - 1 - \chi_H(X_{t-})dt, \tag{31}$$

for some vector ζ that encodes the long-run risk price of monetary shocks, and where $\chi_H(x)dt$ is the jump compensator that makes H a martingale.

The crux of the identification issue is that monetary interventions that shift X_t are, except in a knife-edge case, dependent on the economic state. In our Markov environment, H_t inherits the shocks to X_t, so state-dependence in monetary shocks translates into state-dependence in H-shocks, which obfuscates the recovery of belief shocks.

To see the problem, use H to again define the long-run probability measure $\hat{\mathbb{P}}$. The relation between the drift of X_t under measures $\hat{\mathbb{P}}$ and $\mathbb{P}$ is[9]

$$\hat{\mu}(x) = \mu(x) + \int (x' - x)\Big(\exp[\zeta \cdot (x' - x)] - 1\Big)\chi(dx' \mid x), \tag{32}$$

where χ is the compensator of the jumps in X_t. The object that is observed from financial data is $\hat{\mu}(x)$. But we would like to recover the persistence matrix A from $\mu(x) = A_0 + Ax$. Such recovery is only possible if the distortion $\int (x' - x)\big(\exp[\zeta \cdot (x' - x)] - 1\big)\chi(dx' \mid x)$ is a constant independent of x. This constant case arises if and only if both the arrival rate $\lambda(x)$ and size of monetary surprises ξ_t are state-independent. Such a knife-edge case essentially means monetary policy acts randomly.

Indirect monetary effects. Next, consider what happens if dH_t depends on X_t but not dM_t. In this case, rather than the log-normal form (25), H_t takes the form

$$H_t = \exp\Big[-\frac{1}{2}\int_0^t \|\beta(X_s)\|ds - \int_0^t \beta(X_s) \cdot dW_s\Big] \tag{33}$$

for some non-constant function $\beta(\cdot) : \mathbb{R}^n \mapsto \mathbb{R}^n$. The drift of X_t under the long-run measure $\hat{\mathbb{P}}$ is given by

$$\hat{\mu}(x) = A_0 + Ax - \beta(x)B.$$

[9]Formally, under investor beliefs $\mathbb{P}$, the jumps $\xi_t dM_t$ have conditional mean $\int (x' - x)\chi(dx' \mid x)dt$. By contrast, under the probability distribution $\hat{\mathbb{P}}$ induced by H, the jumps $\xi_t dM_t$ have conditional mean $\int (x' - x)\exp[\zeta \cdot (x' - x)]\chi(dx' \mid x)dt$ (c.f., Kunita and Watanabe, 1967, Theorem 6.2). Combining these points leads to formula (32) in the text.

Although $\hat{\mu}(x)$ is observable, we cannot separately distinguish between Ax and $\beta(x)B$. When $\beta(\cdot)$ was a constant, we could identify A as the persistence of X_t under $\hat{\mathbb{P}}$. Here, the persistence of X_t differs under $\hat{\mathbb{P}}$ and $\mathbb{P}$, complicating matters.

Investors' perceived persistence A is the critical determinant of forecast revisions. Indeed, we have

$$\begin{aligned}\mathbb{E}^{X_\tau}[X_T] - \mathbb{E}^{X_{\tau-}}[X_T] &= X_\tau - X_{\tau-} + \mathbb{E}^{X_\tau}[\int_0^T \mu(X_{t-})dt] - \mathbb{E}^{X_{\tau-}}[\int_0^T \mu(X_{t-})dt] \\ &= X_\tau - X_{\tau-} + \int_0^T A\Big(\mathbb{E}^{X_\tau}[X_t] - \mathbb{E}^{X_{\tau-}}[X_t]\Big)dt. \end{aligned} \tag{34}$$

Equation (34) is a recursive equation for $\mathbb{E}^{X_\tau}[X_T] - \hat{\mathbb{E}}^{X_{\tau-}}[X_T]$, but we can only solve it if we know the value of A. Since we cannot infer A, we cannot solve for these forecast revisions.

The above suggestive analysis of violating Long-Run Neutrality by either (31) or (33) can be formalized. We have

Proposition 4. *Consider the linear environment defined by (26)-(28). Suppose either (i) dH_t features a contribution from dM_t; or (ii) the dynamics dH_t depend on X_t. Then, generically, z_τ^T cannot be identified from asset price data.*

Whereas Propositions 1-2 demonstrated the sufficiency of Long-Run Neutrality for identifying forward-guidance shocks in quasi-linear environments, Proposition 4 demonstrates the corresponding necessity result.

4 Testing Long-Run Neutrality

In this section, we construct a simple non-parametric test of Long-Run Neutrality. This test builds on insights by Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012) in proxying the permanent and transitory components of the SDF. But our analysis focuses on the monetary announcement returns and premia rather than unconditional moments. We then collect some numbers from existing studies to evaluate the non-parametric test.

4.1 A non-parametric test

Roughly speaking, Long-Run Neutrality means that long-run risk premia are invariant to monetary policy. To formalize this, consider two portfolios: (i) an infinite-maturity bond with return $R^{\infty}_{t,t+\Delta}$; and (ii) the growth-optimal portfolio with return $R^{*}_{t,t+\Delta}$. The holding period return on the infinite-maturity bond is given by

$$R_{t,t+\Delta}^{\infty} := \lim_{T\to\infty} R_{t,t+\Delta}^{T} = \lim_{T\to\infty} \frac{\mathbb{E}[\frac{S_T}{S_{t+\Delta}} \mid X_{t+\Delta}]}{\mathbb{E}[\frac{S_T}{S_t} \mid X_t]} = \exp(-\eta)\frac{e(X_{t+\Delta})}{e(X_t)} \lim_{T\to\infty} \frac{\mathbb{E}[\frac{1}{e(X_T)}\frac{H_T}{H_{t+\Delta}} \mid X_{t+\Delta}]}{\mathbb{E}[\frac{1}{e(X_T)}\frac{H_T}{H_t} \mid X_t]}$$
$$= \exp(-\eta)\frac{e(X_{t+\Delta})}{e(X_t)}, \tag{35}$$

where the last equality holds if X_t is stochastically stable under the probability measure generated by H, which we implicitly assume (see footnote 8). On the other hand, the growth-optimal portfolio return $R_{t,t+\Delta}^{*}$ is defined as investors' expectation of the maximal log return: it is the time-$(t+\Delta)$ measurable return R that maximizes $\mathbb{E}[\log(R) \mid X_t]$ subject to $\mathbb{E}[\frac{S_{t+\Delta}}{S_t}R \mid X_t] = 1$, the solution of which is $R_{t,t+\Delta}^{*} = \frac{S_t}{S_{t+\Delta}}$. Putting these results together, and using the SDF decomposition (19), the excess return of the growth-optimal portfolio relative to the infinite-horizon bond is

$$\log(R_{t,t+\Delta}^{*}) - \log(R_{t,t+\Delta}^{\infty}) = \log\left(\frac{H_t}{H_{t+\Delta}}\right) \tag{36}$$

over any horizon Δ. The result in (36) holds in even more general environments than the one considered here—for instance, in non-Markovian environments (Qin and Linetsky, 2017). Under condition (i) of Definition 2, the excess return $\log(R_{t,t+\Delta}^{*}) - \log(R_{t,t+\Delta}^{\infty})$ should be identically zero on monetary announcement days. Under condition (ii) of Definition 2, the conditional risk premium $\mathbb{E}[\log(R_{t,t+\Delta}^{*}) - \log(R_{t,t+\Delta}^{\infty}) \mid X_t]$ should be time-invariant.

Equation (36) suggests a test: one can examine high-frequency changes in $R_{t,t+\Delta}^{*}$ and $R_{t,t+\Delta}^{\infty}$ around monetary announcements to detect the policy impact on H. As long as investor beliefs are not singular with respect to the objective probability, invariance of H to policy under investor beliefs is equivalent to invariance under the objective measure, justifying this test. (By contrast, it is harder to test the time-invariance of $\mathbb{E}[\log(R_{t,t+\Delta}^{*}) - \log(R_{t,t+\Delta}^{\infty}) \mid X_t]$, because it could be so under investor beliefs but not under the objective measure.) For example, if we suppose $R_{t,t+\Delta}^{\infty}$ is well-approximated by returns on 30-year Treasuries, and $R_{t,t+\Delta}^{*}$ is well-approximated by stock market returns, then Long-Run Neutrality says that monetary policy impacts the stock market and 30-year Treasuries in an identical way.

4.2 Existing announcement effect evidence

Following Alvarez and Jermann (2005), suppose $R^{\infty}_{t,t+\Delta}$ is well-approximated by returns on long-term US Treasuries, and $R^{*}_{t,t+\Delta}$ is well-approximated by US stock market returns. Then, by piecing together existing evidence from various sources, we can shed light on the portfolio in (36). Together, this collection of evidence suggests that H responds to monetary policy.

First, stock returns are significantly higher than long-term bond returns around FOMC meeting days. In particular, Lucca and Moench (2015) shows (for 1994-2011) that the SPX return was 0.330% higher on FOMC meeting days, relative to non-meeting days. By contrast, Hillenbrand (2021) shows (for 1989-2021) that 30-year Treasuries returns were approximately $[0.138\%, 0.186\%]$ higher per day in 3-day window around FOMC, relative to other days.[10] Not only are the return magnitudes different, their timing does not match: the stock return is almost entirely earned pre-announcement, where majority of the bond returns are earned post-announcement.

Although our paper is primarily about monetary announcements, we think it is worthwhile to mention some studies on a more complete set of macroeconomic announcements. This may be relevant, because the mechanisms generating announcement premia are often related, and the papers in the broader announcement effects literature have argued effect sizes are similar between FOMC announcements and non-FOMC announcements. Using a much longer sample (1958-2009), Savor and Wilson (2013) show that stock returns are approximately 11.5bp higher on announcement days, while 30-year Treasuries earned 4.5bp higher returns on announcement days. Savor and Wilson (2014) show in a similar sample (1964-2011) that the CAPM beta of 30-year Treasuries is 0.14 on announcement days, whereas Long-Run Neutrality predicts it should be 1.

There is also a literature that does not take averages across all announcements but rather examines asset responses to monetary policy shocks. For example, Gürkaynak, Sack and Swanson (2005b) study high-frequency responses to monetary shocks (during 1990-2004), finding that a 25bp surprise rate cut leads to 1% SPX return but only a 0.32% 10-year Treasury return. When extended further (1988-2019), Bauer, Bernanke and Milstein (2023) found even stronger effects in stocks (a 10bp surprise cut leads to a 1% SPX return), although they did not study long-term riskless bonds.

[10]This is constructed using his evidence that 30-year Treasury yields decline between 0.46bps to 0.62bps more per day in the 3-day window surrounding FOMC meetings. Then, we use the duration-approximation $\log(R^{T}_{t,t+\Delta}) \approx -T(y_{t+\Delta} - y_t)$.

5 Examples of H: Interpreting Long-Run Neutrality

We present some example economies in which the SDF S features a permanent component H. In each example, we discuss what is meant, economically, by the Long-Run Neutrality statement "monetary policy does not affect H." Thus, we can evaluate the stringency of conditions that allow identification of monetary policy shocks. The examples in this section are based on Bansal and Yaron (2004), with related analysis in Hansen and Scheinkman (2009) and Borovička et al. (2016). Generalizing these economies to explicitly include monetary policy is an interesting avenue for future research.

5.1 Long-run risk model

Suppose aggregate consumption has the trend-stationary dynamics

$$\log C_{t+1} = \log C_t + \alpha \cdot (X_{t+1} - X_t),$$

where the state vector X_t follows a stationary VAR(1):

$$X_{t+1} = A_0 + AX_t + B\Delta W_{t+1}, \quad \Delta W_{t+1} \sim \text{Normal}(0, I).$$

If investors have subjective beliefs, the same form of these equations hold in reality and under investor beliefs, but with alternative values of A_0.

Suppose the representative investor has recursive preferences as in Kreps and Porteus (1978) and Epstein and Zin (1989), and unitary elasticity of intertemporal substitution (EIS). The investor's continuation value satisfies the following recursion:

$$V_t = (1-\beta)\log C_t + \beta \frac{\log \mathbb{E}_t[\exp((1-\gamma)V_{t+1})]}{1-\gamma},$$

where $\gamma > 1$ denotes the investor's coefficient of relative risk aversion and β is the subjective discount factor. Guess that the solution is $V_t = A_0 t + v_0 + v \cdot X_t$, for some constant v_0 and vector v. In that case, one can show that the solution is

$$v = (I - A')^{-1}\alpha \quad \text{and} \quad v_0 = (1-\beta)^{-1}[\beta A_0 + \frac{1}{2}(1-\gamma)v'BB'v].$$

In this model, the SDF is given by

$$\frac{S_{t+1}}{S_t} = \beta \frac{C_t}{C_{t+1}} \frac{\exp((1-\gamma)v \cdot X_{t+1})}{\mathbb{E}[\exp((1-\gamma)v \cdot X_{t+1}) \mid X_t]}$$

Given that consumption is a trend-stationary process, the permanent component of the SDF is clearly given by the third piece, i.e.,

$$\frac{H_{t+1}}{H_t} = \frac{\exp((1-\gamma)v \cdot X_{t+1})}{\mathbb{E}[\exp((1-\gamma)v \cdot X_{t+1}) \mid X_t]} = \exp\Big[-\frac{1}{2}(1-\gamma)^2 v' BB' v + (1-\gamma) v' B \Delta W_{t+1}\Big].$$

In this model, if monetary policy shocks do not affect H, then there are two possibilities. One trivial possibility is that $\gamma = 1$ corresponding to log utility, which rules out priced growth-rate shocks. In that case, Long-Run Neutrality corresponds to the conventional wisdom that monetary policy does not have a permanent effect on the consumption level, which is hard-wired in this example with trend-stationary consumption.

Alternatively, supposed growth-rate shocks are priced. Then, letting $B^{(i)}$ denote the ith column of B, Long-Run Neutrality requires $v'B^{(i)} = 0$ for every shock $\Delta W^{(i)}_{t+1}$ that can be impacted by monetary policy. For example, if $\Delta W^{(1)}_{t+1}$ is the short-rate shock, and $\Delta W^{(2)}_{t+1}$ is a shock corresponding to forward guidance, then a requirement for identification is $v'B^{(1)} = v'B^{(2)} = 0$. But since the elements of $v = (I - A')^{-1}\alpha$ are generically non-zero, the requirement implies that $B^{(i)} = 0$. In words, identification requires that *growth, in both the short and long run, is invariant to monetary policy*.

5.2 Stochastic-volatility model

Suppose consumption features the following dynamics, with stochastic volatility:

$$\begin{aligned} \log C_{t+1} &= \log C_t + g + \sqrt{X_t}\Delta W^{(1)}_{t+1} \\ X_{t+1} &= \mu + a(X_t - \mu) + \sigma\sqrt{X_t}\Delta W^{(2)}_{t+1}, \quad \Delta W_{t+1} \sim \text{Normal}(0, I), \end{aligned}$$

where $a < 1$. As above, the representative investor has Epstein-Zin utility with unitary EIS. In this model, that means that the SDF takes the form:

$$\frac{S_{t+1}}{S_t} = \beta \exp\Big[-g - \gamma X_t - \frac{1}{2} X_t \big|\big(\begin{smallmatrix}\gamma \\ (\gamma-1)\sigma\kappa\end{smallmatrix}\big)\big|^2 - \sqrt{X_t}\big(\begin{smallmatrix}\gamma \\ (\gamma-1)\sigma\kappa\end{smallmatrix}\big) \cdot \Delta W_{t+1}\Big],$$

where $\kappa < 0$ is the larger root of the quadratic equation $\frac{1-\gamma}{2}\sigma^2\kappa^2 + \log(\beta)\kappa - \frac{1}{2}\gamma = 0$.

In this environment, it is easy to verify that the stationary component of the SDF is characterized by the eigenfunction $e(x) = \exp(vx)$, where v is a root of the quadratic equation $0 = \frac{1}{2}\sigma^2 v^2 - [(\gamma-1)\sigma^2\kappa + (1-a)]v - \gamma$ (the choice of the root is to ensure the dynamics of X_t are stable under the measure induced by the resulting H_t). Consequently, the permanent component of this SDF is

$$\frac{H_{t+1}}{H_t} = \exp\Big[-\frac{1}{2}\Big|\Big(\begin{smallmatrix}\gamma \\ (\gamma-1)\sigma\kappa-\sigma v\end{smallmatrix}\Big)\Big|^2 X_t - \sqrt{X_t}\Big(\begin{smallmatrix}\gamma \\ (\gamma-1)\sigma\kappa-\sigma v\end{smallmatrix}\Big)\cdot\Delta W_{t+1}\Big],$$

Imagine we are not in the knife-edge case where $(\gamma-1)\kappa = v$. Then, identification of monetary shocks requires that *monetary policy does not affect uncertainty*, since uncertainty affects H. For example, identification requires output growth volatility and stock market volatility be invariant to monetary actions.

6 Conclusion

If researchers do not want to impose rational expectations or risk-neutrality, how can they use asset prices to recover beliefs of investors? The current frontier of knowledge suggests this problem has no general solution, except in the degenerate case of long-run risk-neutrality. However, if a researcher more humbly seeks only to identify *shocks to investor beliefs*, then identification is possible under weaker conditions.

We explore such shock identification in the context of monetary policy that can affect current and future interest rates. In quasi-linear environments (either linear or with stochastic volatility of the "square-root" form), shock identification is possible provided a *Long-Run Neutrality* condition holds: policy must not affect variables that permanently shift the pricing kernel.

Unfortunately, existing empirical evidence on the monetary announcement effect suggests Long-Run Neutrality is violated. Furthermore, in some popular structural models featuring priced news about growth and uncertainty, Long-Run Neutrality is equivalent to saying monetary policy does not affect the real economy. Through the lens of these models, identification of monetary policy effects relies paradoxically on monetary policy having no effects.

We see three important outstanding questions. First, quantitatively, how much bias do existing methods introduce by neglecting to identify investor beliefs about forward guidance and future rate uncertainty? Our simple linear example in Section 2 demonstrates the amount of bias in the "causal effect of short rate policy" is related to $\tilde{A} - A$, which represents investors' *misunderstanding of persistences* (e.g., extrapolative beliefs). This bias grows with the forecast horizon, suggesting that we may know very little about the effects of monetary policy beyond the very short horizon. In an example calibration reflecting this extrapolation, we have shown that discretionary short-rate shocks may in fact have stronger impacts than existing methods estimate. Richer models and calculations of this type would be informative.

Second, how can researchers identify investor forecast revisions when asset prices do not? For this question, we so no obvious alternative than to incorporate survey data on future interest rates, inflation, and the like. Analysis of survey evidence has been a fruitful and growing area of research, and we see promise in connecting these survey data with more structural models of monetary policy. Our framework sheds light on how such survey data should be included in regressions with short-rate shocks to estimate the effects of monetary policy jointly (e.g., Procedures 1-2). At the same time, our paper effectively assumes a representative agent (and so homogeneous beliefs, or at least that the relevant set of beliefs for any economic outcome are those of the marginal investor). Seriously thinking about how belief heterogeneity impacts the identification of monetary shocks and their effects seems like a promising area for future research in this direction.

Third, and perhaps most importantly, what really are monetary shocks? Our approach follows most of the empirical literature in measuring shocks as interest rate forecast errors. Several potential theoretical explanations exist for these forecast errors: e.g., true randomization in rate setting; uncertain and time-varying interest rate rules; signalling future prospects via policy; belief disagreements among central bankers and investors. Further developing these monetary shock microfoundations seems like an important research direction, to understand whether and how monetary policy has long-term impacts or not (e.g., what do these stories imply about H), as well as the appropriate procedure for estimating monetary IRFs.

References

Ai, Hengjie and Ravi Bansal, "Risk preferences and the macroeconomic announcement premium," *Econometrica*, 2018, *86* (4), 1383–1430.

Alvarez, Fernando and Urban J Jermann, "Using asset prices to measure the persistence of the marginal utility of wealth," *Econometrica*, 2005, *73* (6), 1977–2016.

Andolfatto, David and Paul Gomme, "Monetary policy regimes and beliefs," *International Economic Review*, 2003, *44* (1), 1–30.

Ang, Andrew and Monika Piazzesi, "A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables," *Journal of Monetary Economics*, 2003, *50* (4), 745–787.

Backus, David K, Mikhail Chernov, Stanley E Zin, and Irina Zviadadze, "Monetary Policy Risk: Rules versus Discretion," *The Review of Financial Studies*, 2022, *35* (5), 2308–2344.

Baker, Scott R, Nicholas Bloom, and Steven J Davis, "Measuring economic policy uncertainty," *The Quarterly Journal of Economics*, 2016, *131* (4), 1593–1636.

Bakshi, Gurdip and Fousseni Chabi-Yo, "Variance bounds on the permanent and transitory components of stochastic discount factors," *Journal of Financial Economics*, 2012, *105* (1), 191–208.

Ball, Laurence and Dean Croushore, "Expectations and the Effects of Monetary Policy," *Journal of Money, Credit and Banking*, 2003, pp. 473–484.

Bansal, Ravi and Amir Yaron, "Risks for the long run: A potential resolution of asset pricing puzzles," *The Journal of Finance*, 2004, *59* (4), 1481–1509.

Barsky, Robert B and Eric R Sims, "News shocks and business cycles," *Journal of Monetary Economics*, 2011, *58* (3), 273–289.

__ , Susanto Basu, and Keyoung Lee, "Whither news shocks?," *NBER Macroeconomics Annual*, 2015, *29* (1), 225–264.

Basu, Susanto and Brent Bundick, "Uncertainty shocks in a model of effective demand," *Econometrica*, 2017, *85* (3), 937–958.

Bauer, Michael D and Eric T Swanson, "An alternative explanation for the "Fed information effect"," *American Economic Review*, 2023, *113* (3), 664–700.

__ and __ , "A reassessment of monetary policy surprises and high-frequency identification," *NBER Macroeconomics Annual*, 2023, *37* (1), 87–155.

__ , Ben S Bernanke, and Eric Milstein, "Risk appetite and the risk-taking channel of monetary policy," *Journal of Economic Perspectives*, 2023, *37* (1), 77–100.

Beaudry, Paul and Franck Portier, "An exploration into Pigou's theory of cycles," *Journal of Monetary Economics*, 2004, *51* (6), 1183–1216.

— and — , “Stock prices, news, and economic fluctuations,” *American Economic Review*, 2006, *96* (4), 1293–1307.

Beber, Alessandro and Michael W Brandt, “The effect of macroeconomic news on beliefs and preferences: Evidence from the options market,” *Journal of Monetary Economics*, 2006, *53* (8), 1997–2039.

Berger, David, Ian Dew-Becker, and Stefano Giglio, “Uncertainty shocks as second-moment news shocks,” *The Review of Economic Studies*, 2020, *87* (1), 40–76.

Bernanke, Ben S and Kenneth N Kuttner, “What explains the stock market’s reaction to Federal Reserve policy?,” *The Journal of Finance*, 2005, *60* (3), 1221–1257.

— **and Michael Woodford**, “Inflation Forecasts and Monetary Policy,” *Journal of Money, Credit, and Banking*, 1997, pp. 653–684.

Bianchi, Francesco, Cosmin L Ilut, and Martin Schneider, “Uncertainty shocks, asset supply and pricing over the business cycle,” *The Review of Economic Studies*, 2018, *85* (2), 810–854.

— **, Howard Kung, and Mikhail Tirskikh**, “The origins and effects of macroeconomic uncertainty,” *Quantitative Economics*, 2023, *14* (3), 855–896.

— **, Martin Lettau, and Sydney C Ludvigson**, “Monetary policy and asset valuation,” *The Journal of Finance*, 2022, *77* (2), 967–1017.

— **, Sydney C Ludvigson, and Sai Ma**, “Belief distortions and macroeconomic fluctuations,” *American Economic Review*, 2022, *112* (7), 2269–2315.

— **,** — **, and** — , “Monetary-based asset pricing: A mixed-frequency structural approach,” Technical Report, National Bureau of Economic Research 2022.

Bidder, Rhys and Ian Dew-Becker, “Long-run risk is the worst-case scenario,” *American Economic Review*, 2016, *106* (9), 2494–2527.

Bloom, Nicholas, “The impact of uncertainty shocks,” *Econometrica*, 2009, *77* (3), 623–685.

Borovička, Jaroslav, Lars Peter Hansen, and José A Scheinkman, “Misspecified recovery,” *The Journal of Finance*, 2016, *71* (6), 2493–2544.

Caballero, Ricardo J and Alp Simsek, “Monetary policy with opinionated markets,” *American Economic Review*, 2022, *112* (7), 2353–2392.

Campbell, Jeffrey R, Charles L Evans, Jonas DM Fisher, and Alejandro Justiniano, “Macroeconomic effects of federal reserve forward guidance,” *Brookings papers on economic activity*, 2012, pp. 1–80.

Christiano, Lawrence J, Roberto Motto, and Massimo Rostagno, “Risk shocks,” *American Economic Review*, 2014, *104* (1), 27–65.

Chun, Albert Lee, “Expectations, bond yields, and monetary policy,” *The Review of Financial Studies*, 2011, *24* (1), 208–247.

Cieslak, Anna, "Short-rate expectations and unexpected returns in treasury bonds," *The Review of Financial Studies*, 2018, *31* (9), 3265–3306.

_ , Adair Morse, and Annette Vissing-Jorgensen, "Stock returns over the FOMC cycle," *The Journal of Finance*, 2019, *74* (5), 2201–2248.

_ and Andreas Schrimpf, "Non-monetary news in central bank communication," *Journal of International Economics*, 2019, *118*, 293–315.

_ and Hao Pang, "Common shocks in stocks and bonds," *Journal of Financial Economics*, 2021, *142* (2), 880–904.

Corsetti, Giancarlo, Simon Lloyd, Emile Marin, and Daniel Ostry, "US Risk and Treasury Convenience," 2023. Unpublished working paper.

Creal, Drew D and Jing Cynthia Wu, "Monetary policy uncertainty and economic fluctuations," *International Economic Review*, 2017, *58* (4), 1317–1354.

Crump, Richard K, Stefano Eusepi, and Emanuel Moench, "The term structure of expectations and bond yields," 2018. Unpublished working paper.

Dai, Qiang and Kenneth J Singleton, "Expectation puzzles, time-varying risk premia, and affine models of the term structure," *Journal of Financial Economics*, 2002, *63* (3), 415–441.

d'Arienzo, Daniele, "Maturity increasing overreaction and bond market puzzles," 2020. Unpublished working paper.

Duffee, Gregory R, "Term premia and interest rate forecasts in affine models," *The Journal of Finance*, 2002, *57* (1), 405–443.

Duffie, Darrell and Rui Kan, "A yield-factor model of interest rates," *Mathematical Finance*, 1996, *6* (4), 379–406.

Ederington, Louis H and Jae Ha Lee, "How markets process information: News releases and volatility," *The Journal of Finance*, 1993, *48* (4), 1161–1191.

Epstein, Larry G and Stanley E Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 1989, pp. 937–969.

Evans, George W and Seppo Honkapohja, "Adaptive learning and monetary policy design," *Journal of Money, Credit and Banking*, 2003, pp. 1045–1072.

Fajgelbaum, Pablo D, Edouard Schaal, and Mathieu Taschereau-Dumouchel, "Uncertainty traps," *The Quarterly Journal of Economics*, 2017, *132* (4), 1641–1692.

Faust, Jon, John H Rogers, Shing-Yi B Wang, and Jonathan H Wright, "The high-frequency response of exchange rates and interest rates to macroeconomic announcements," *Journal of Monetary Economics*, 2007, *54* (4), 1051–1068.

Francis, Neville and Valerie A Ramey, "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," *Journal of Monetary Economics*, 2005, *52* (8), 1379–1399.

Gasteiger, Emanuel, "Heterogeneous expectations, optimal monetary policy, and the merit of policy inertia," *Journal of Money, Credit and Banking*, 2014, *46* (7), 1535–1554.

Giglio, Stefano and Bryan Kelly, "Excess volatility: Beyond discount rates," *The Quarterly Journal of Economics*, 2018, *133* (1), 71–127.

Gürkaynak, Refet S, Brian Sack, and Eric Swanson, "The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models," *American Economic Review*, 2005, *95* (1), 425–436.

— , — , and Eric T Swanson, "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 2005, *1* (1), 55–93.

Haddad, Valentin, Alan Moreira, and Tyler Muir, "Whatever it takes? The impact of conditional policy promises," Technical Report, National Bureau of Economic Research 2023.

Hamilton, James D and Jing Cynthia Wu, "Identification and estimation of Gaussian affine term structure models," *Journal of Econometrics*, 2012, *168* (2), 315–331.

— , Seth Pruitt, and Scott Borger, "Estimating the market-perceived monetary policy rule," *American Economic Journal: Macroeconomics*, 2011, *3* (3), 1–28.

Hansen, Lars Peter and José A Scheinkman, "Long-term risk: An operator approach," *Econometrica*, 2009, *77* (1), 177–234.

Hanson, Samuel G and Jeremy C Stein, "Monetary policy and long-term real rates," *Journal of Financial Economics*, 2015, *115* (3), 429–448.

— , David O Lucca, and Jonathan H Wright, "Rate-amplifying demand and the excess sensitivity of long-term rates," *The Quarterly Journal of Economics*, 2021, *136* (3), 1719–1781.

Harrison, J Michael and David M Kreps, "Martingales and arbitrage in multiperiod securities markets," *Journal of Economic Theory*, 1979, *20* (3), 381–408.

Hillenbrand, Sebastian, "The Fed and the secular decline in interest rates," *Available at SSRN 3550593*, 2021.

Hommes, Cars, Domenico Massaro, and Matthias Weber, "Monetary policy under behavioral expectations: Theory and experiment," *European Economic Review*, 2019, *118*, 193–212.

Husted, Lucas, John Rogers, and Bo Sun, "Monetary policy uncertainty," *Journal of Monetary Economics*, 2020, *115*, 20–36.

Kazemi, Hossein B, "An intemporal model of asset prices in a Markov economy with a limiting stationary distribution," *The Review of Financial Studies*, 1992, *5* (1), 85–104.

Kelly, Bryan, L'uboš Pástor, and Pietro Veronesi, "The price of political uncertainty: Theory and evidence from the option market," *The Journal of Finance*, 2016, *71* (5), 2417–2480.

Kreps, David M and Evan L Porteus, "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica*, 1978, pp. 185–200.

Krueger, Joel T and Kenneth N Kuttner, "The Fed funds futures rate as a predictor of Federal Reserve policy," *Journal of Futures Markets*, 1996, *16* (8), 865–879.

Kryvtsov, Oleksiy and Luba Petersen, "Expectations and monetary policy: experimental evidence," *Available at SSRN 4343328*, 2019.

Kunita, Hiroshi and Shinzo Watanabe, "On square integrable martingales," *Nagoya Mathematical Journal*, 1967, *30*, 209–245.

Kurmann, André and Christopher Otrok, "News shocks and the slope of the term structure of interest rates," *American Economic Review*, 2013, *103* (6), 2612–2632.

Kuttner, Kenneth N, "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market," *Journal of Monetary Economics*, 2001, *47* (3), 523–544.

Leduc, Sylvain and Zheng Liu, "Uncertainty shocks are aggregate demand shocks," *Journal of Monetary Economics*, 2016, *82*, 20–35.

Liu, Yukun and Ben Matthies, "Long-Run Risk: Is It There?," *The Journal of Finance*, 2022, *77* (3), 1587–1633.

Lucca, David O and Emanuel Moench, "The pre-FOMC announcement drift," *The Journal of Finance*, 2015, *70* (1), 329–371.

McQueen, Grant and V Vance Roley, "Stock prices, news, and business conditions," *The Review of Financial Studies*, 1993, *6* (3), 683–707.

Melosi, Leonardo, "Signalling effects of monetary policy," *The Review of Economic Studies*, 2017, *84* (2), 853–884.

Milani, Fabio, "Learning, monetary policy rules, and macroeconomic stability," *Journal of Economic Dynamics and Control*, 2008, *32* (10), 3148–3165.

Miranda-Agrippino, Silvia and Giovanni Ricco, "The transmission of monetary policy shocks," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 74–107.

Nagel, Stefan and Zhengyang Xu, "Dynamics of subjective risk premia," 2022. Unpublished working paper.

Nakamura, Emi and Jón Steinsson, "High-frequency identification of monetary non-neutrality: the information effect," *The Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.

—, Dmitriy Sergeyev, and Jón Steinsson, "Growth-rate and uncertainty shocks in consumption: Cross-country evidence," *American Economic Journal: Macroeconomics*, 2017, *9* (1), 1–39.

Negro, Marco Del, Marc P Giannoni, and Christina Patterson, "The forward guidance puzzle," *FRB of New York Staff Report*, 2012, (574).

Pastor, Lubos and Pietro Veronesi, "Uncertainty about government policy and stock prices," *The Journal of Finance*, 2012, *67* (4), 1219–1264.

Pástor, L'uboš and Pietro Veronesi, "Political uncertainty and risk premia," *Journal of Financial Economics*, 2013, *110* (3), 520–545.

Pearce, Douglas K and V Vance Roley, "Stock Prices and Economic News," *The Journal of Business*, 1985, *58* (1), 49–67.

Piazzesi, Monika and Eric T Swanson, "Futures prices as risk-adjusted forecasts of monetary policy," *Journal of Monetary Economics*, 2008, *55* (4), 677–691.

Poole, William, Robert H Rasche, Daniel L Thornton et al., "Market anticipations of monetary policy actions," *Review-Federal Reserve Bank of Saint Louis*, 2002, *84* (4), 65–94.

Qin, Likuan and Vadim Linetsky, "Positive eigenfunctions of Markovian pricing operators: Hansen-Scheinkman factorization, Ross recovery, and long-term pricing," *Operations Research*, 2016, *64* (1), 99–117.

— **and** — , "Long-term risk: A martingale approach," *Econometrica*, 2017, *85* (1), 299–312.

Romer, Christina D and David H Romer, "Federal Reserve information and the behavior of interest rates," *American Economic Review*, 2000, *90* (3), 429–457.

Rosa, Carlo, "Words that shake traders: The stock market's reaction to central bank communication in real time," *Journal of Empirical Finance*, 2011, *18* (5), 915–934.

Ross, Steve, "The recovery theorem," *The Journal of Finance*, 2015, *70* (2), 615–648.

Rudebusch, Glenn D, "Do measures of monetary policy in a VAR make sense?," *International Economic Review*, 1998, pp. 907–931.

— , "Term structure evidence on interest rate smoothing and monetary policy inertia," *Journal of Monetary Economics*, 2002, *49* (6), 1161–1187.

Särkkä, Simo and Arno Solin, *Applied stochastic differential equations*, Vol. 10, Cambridge University Press, 2019.

Savor, Pavel and Mungo Wilson, "How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements," *Journal of Financial and Quantitative Analysis*, 2013, *48* (2), 343–375.

— **and** — , "Asset pricing: A tale of two days," *Journal of Financial Economics*, 2014, *113* (2), 171–201.

Schmitt-Grohé, Stephanie and Martin Uribe, "What's news in business cycles," *Econometrica*, 2012, *80* (6), 2733–2764.

Schorfheide, Frank, "Learning and monetary policy shifts," *Review of Economic Dynamics*, 2005, *8* (2), 392–419.

— **, Dongho Song, and Amir Yaron**, "Identifying long-run risks: A Bayesian mixed-frequency approach," *Econometrica*, 2018, *86* (2), 617–654.

Swanson, Eric T, "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets," *Journal of Monetary Economics*, 2021, *118*, 32–53.

Tella, Sebastian Di, "Uncertainty shocks and balance sheet recessions," *Journal of Political Economy*, 2017, *125* (6), 2038–2081.

— **and Robert Hall**, "Risk premium shocks can create inefficient recessions," *The Review of Economic Studies*, 2022, *89* (3), 1335–1369.

Uribe, Martín, "The Neo-Fisher effect: Econometric evidence from empirical and optimizing models," *American Economic Journal: Macroeconomics*, 2022, *14* (3), 133–162.

Wang, Chen, "Under-and overreaction in yield curve expectations," 2021. Unpublished working paper.

Xu, Zhengyang, "Expectation Formation in the Treasury Bond Market," 2019. Unpublished working paper.

Appendix:

Risk Premia, Subjective Beliefs, and Bundled Monetary Shocks

Anna Cieslak and Paymon Khorrami

November 16, 2023

A Proofs

We first prove Lemma 1 regarding what is recovered from asset price data. Then, we prove Propositions 1-4.

Proof of Lemma 1. Given the Markovian environment, the asset prices in (18) can be represented by a family of pricing operators $(\mathcal{Q}_t)_{t\geq 0}$ as

$$[\mathcal{Q}_t f](x) = \mathbb{E}[S_t f(X_t) \mid X_0 = x]. \tag{A.1}$$

The operator $\mathcal{Q}_t$ is the t-period pricing operator for any claim f that is a function of the Markov state. In all that follows, we assume we observe $\mathcal{Q}_t$ (i.e., this is what is meant by "asset price data" in a complete market environment).

Now, solve the eigenvalue problem

$$[\mathcal{Q}_t e](x) = \exp(\eta t) e(x). \tag{A.2}$$

By the Perron-Frobenius theory, $\exp(\eta) > 0$ is a positive eigenvalue of $\lim_{t\to 0} t^{-1}\mathcal{Q}_t$, and its associated eigenfunction e is strictly positive. Given $\mathcal{Q}_t$ is observable, we thus can infer e and η from data.[11]

After recovering these objects, we may construct

$$H_t := \exp(-\eta t) S_t \frac{e(X_t)}{e(X_0)}. \tag{A.3}$$

Of course, S is not directly observable in data, and so neither is H, but the important point is that the H in (A.3) is the same one in the decomposition (19) by construction. Note that H_t is a strictly positive martingale since

$$\mathbb{E}[H_T \mid \mathscr{F}_t] = \frac{\exp(-\eta T)}{e(X_0)} \mathbb{E}[S_T e(X_T) \mid \mathscr{F}_t] = \frac{\exp(-\eta T)}{e(X_0)} \exp(\eta(T-t)) e(X_t) S_t = H_t,$$

by (A.2). Although the construction of $\hat{\mathbb{P}}$ in (23) depends on the unobservable H, note that

$$\begin{aligned}
\hat{\mathbb{P}}\{r(X_{\tau+T}) \leq \mathsf{r} \mid X_\tau\} = \mathbb{E}\Big[\frac{H_{\tau+T}}{H_\tau} 1_{\{r(X_{\tau+T})\leq \mathsf{r}\}} \mid X_\tau\Big] &= \mathbb{E}\Big[\exp(-\eta T) \frac{S_{\tau+T}}{S_\tau} \frac{e(X_{\tau+T})}{e(X_\tau)} 1_{\{r(X_{\tau+T})\leq \mathsf{r}\}} \mid X_\tau\Big] \\
&= \frac{\exp(-\eta T)}{e(X_\tau)} [\mathcal{Q}_T \hat{e}](X_\tau).
\end{aligned}$$

Note that $\hat{e}(x) := e(x) 1_{\{r(x)\leq \mathsf{r}\}}$ is a computable payoff as a function of x. Since η, e, and $\mathcal{Q}_T$ are all also observable, we can observe $\hat{\mathbb{P}}\{r(X_{\tau+T}) \leq \mathsf{r} \mid X_\tau\}$ from asset price data. □

Proof of Proposition 1. This proposition is implied by Proposition 2, since the constant diffusion condition (27) implies the condition (29). □

[11] In a discrete-time model, it would suffice to study the instantaneous pricing operator $\mathcal{Q}_1$, since the law of iterated expectations allows us to apply $\mathcal{Q}_1$ in succession t times in order to obtain $\mathcal{Q}_t$. In continuous time, the analogous operator is the instantaneous pricing operator $\lim_{t\to 0} \mathcal{Q}_t / t$.

Proof of Proposition 2. Let $m_t^x := \mathbb{E}^x[X_t]$ denote the (investor-perceived) conditional mean of X_t, starting from point x. By applying Itô's formula to X_t, we have that m_t^x solves the differential equation (since the compensated monetary shock has zero mean)

$$\frac{d}{dt} m_t^x = \mathbb{E}^x[\mu(X_t)]$$

subject to the initial condition $m_0^x = x$. Specializing to the linear drift from (26), the ODE becomes

$$\frac{d}{dt} m_t^x = \mathbb{E}^x[A_0 + AX_t] = A_0 + Am_t^x$$

This ODE is affine, and the solution takes the well-known form

$$m_t^x = \exp(At)\Big[x + \int_0^t \exp(-As)A_0 ds\Big].$$

We may then compute

$$m_t^{X_\tau} - m_t^{X_{\tau-}} = \exp(At)(X_\tau - X_{\tau-}).$$

(The interpretation of $\exp(At)$ is as the Taylor series $\sum_{k=0}^{\infty} \frac{t^k}{k!} A^k$.)

Finally, using assumption (29), we have that the drift of X_t under $\hat{\mathbb{P}}$ is $\hat{A}_0 + \hat{A}X_t$, where $\hat{A}_0 = A_0 - \hat{\beta}$ and $\hat{A} = A$. Since $\hat{A}$ is observable (by Lemma 1), we have that $A = \hat{A}$ is also observable. Therefore, $m_t^{X_\tau} - m_t^{X_{\tau-}}$ is observable for all t. By assumption (28), we have obtained $z_\tau^T = \rho \cdot (m_T^{X_\tau} - m_T^{X_{\tau-}})$. □

Proof of Proposition 3. Let $m_t^x := \mathbb{E}^x[X_t]$ and $V_t^x := \mathbb{E}^x[(X_t - m_t^x)(X_t - m_t^x)']$ denote the conditional mean and variance of X_t, starting from point x. A standard result on SDEs (e.g., Chapter 5.5 of Särkkä and Solin, 2019) is that

$$\frac{d}{dt} V_t^x = \mathbb{E}^x[\mu(X_t)(X_t - m_t^x)'] + \mathbb{E}^x[(X_t - m_t^x)\mu(X_t)'] + \mathbb{E}^x[\sigma(X_t)\sigma(X_t)'].$$

This equation holds at times t that are non-announcement dates. Specializing to the linear drift from (26) and the square-root assumption on the diffusion (30), we obtain

$$\begin{aligned}
\frac{d}{dt} V_t^x &= \mathbb{E}^x[(A_0 + AX_t)(X_t - m_t^x)'] + \mathbb{E}^x[(X_t - m_t^x)(A_0 + AX_t)'] + \mathbb{E}^x[\varsigma_0\varsigma_0' + \sum_{i=1}^n \varsigma_i \mathrm{diag}(u^{(i)} \cdot x)\varsigma_i'] \\
&= A\mathbb{E}^x[X_t(X_t - m_t^x)'] + \mathbb{E}^x[(X_t - m_t^x)X_t']A' + \varsigma_0\varsigma_0' + \sum_{i=1}^n \varsigma_i \mathrm{diag}(u^{(i)} \cdot m_t^x)\varsigma_i' \\
&= AV_t^x + V_t^x A' + \varsigma_0\varsigma_0' + \sum_{i=1}^n \varsigma_i \mathrm{diag}(u^{(i)} \cdot m_t^x)\varsigma_i',
\end{aligned}$$

where $u^{(i)}$ is the ith elementary vector. Subject to the initial condition $V_0^x = [0]_{n\times n}$, this ODE for V_t^x is a Riccati equation, for which the solution has the well-known form

$$V_t^x = \int_0^t \exp(A(t-s))\Big[\varsigma_0\varsigma_0' + \sum_{i=1}^n \varsigma_i \mathrm{diag}(u^{(i)} \cdot m_t^x)\varsigma_i'\Big] \exp(A'(t-s))ds.$$

Again, this solution holds for any time t prior to the next monetary surprise. We may then compute

$$V_t^{X_\tau} - V_t^{X_{\tau-}} = \int_0^t \exp(A(t-s)) \sum_{i=1}^n \varsigma_i \mathrm{diag}\big[u^{(i)} \cdot (m_s^{X_\tau} - m_s^{X_{\tau-}})\big] \varsigma_i' \exp(A'(t-s)) ds.$$

By Proposition 2, the object $m_s^{X_\tau} - m_s^{X_{\tau-}}$ is observable. In addition, under assumption (29), we have $A = \hat{A}$ observable. Hence, $V_T^{X_\tau} - V_T^{X_{\tau-}}$ is observable. This is enough, since by assumption (28), we have $v_\tau^T = \rho'(V_T^{X_\tau} - V_T^{X_{\tau-}})\rho$. □

Proof of Proposition 4. Suppose, leading to contradiction, that z_τ^T is identified by asset price data. The same procedure also identifies

$$\hat{z}_\tau^T := \hat{\mathbb{E}}\big[r(X_{\tau+T}) \mid X_\tau\big] - \hat{\mathbb{E}}\big[r(X_{\tau+T}) \mid X_{\tau-}\big], \quad T > 0, \tag{A.4}$$

where the probability measure $\hat{\mathbb{P}}$ is defined in (23). Indeed, Proposition 2 of Borovička et al. (2016) says that the observable asset prices can be obtained by formula (18) using either (i) probability measure $\mathbb{P}$ and SDF S, or (ii) probability measure $\hat{\mathbb{P}}$ and SDF $\hat{S}$, where

$$\hat{S}_t := S_t \frac{H_0}{H_t}. \tag{A.5}$$

Therefore, the same asset price data that identify z_τ^T also identify $\hat{z}_\tau^T$.

Since both z_τ^T and $\hat{z}_\tau^T$ are identified by the same procedure on asset prices, their values must be identical:

$$\begin{aligned}\mathbb{E}\Big[\frac{H_{\tau+T}}{H_\tau} r(X_{\tau+T}) \mid X_\tau\Big] - \mathbb{E}\Big[\frac{H_{\tau+T}}{H_{\tau-}} r(X_{\tau+T}) \mid X_{\tau-}\Big] \\ = \mathbb{E}\big[r(X_{\tau+T})) \mid X_\tau\big] - \mathbb{E}\big[r(X_{\tau+T})) \mid X_{\tau-}\big].\end{aligned} \tag{A.6}$$

Using the fact that (A.6) holds for all X_τ and $X_{\tau-}$, we must have

$$\mathbb{E}\big[H_T r(X_T) \mid X_0 = x\big] = \mathbb{E}\big[r(X_T) \mid X_0 = x\big] + \alpha(T), \tag{A.7}$$

where $\alpha(T)$ may depend on T but is independent of x.

Now, under both hypotheses (i) and (ii) of the Proposition, it must generically be the case that dH_t depends on X_{t-}. (Generically, because hypothesis (i) can be consistent with $dH_t \perp X_{t-}$ in the knife-edge case that the announcement arrival rate $\lambda(x) \equiv \lambda$ is constant and the probability distribution of monetary shocks ξ_t is independent of X_{t-}.) As a result, the probability distribution of H_T generically depends on X_0. Since X_T also depends on X_0, we have that H_T and $r(X_T)$ are generally non-orthogonal.

Based on this discussion, both of

$$\begin{aligned}\hat{R}_T(x) &:= \mathbb{E}\big[H_T r(X_T) \mid X_0 = x\big] \\ \text{and} \quad R_T(x) &:= \mathbb{E}\big[r(X_T) \mid X_0 = x\big]\end{aligned}$$

are non-constant functions of x for any time horizon $T > 0$. Furthermore, $\hat{R}_T(x) - R_T(x)$ depends on x. This contradicts the fact that $\alpha(T)$ is independent of x. Thus, x_τ^T is not identified. □