Risk Premia, Subjective Beliefs, and Forward Guidance*

Anna Cieslak
Duke University Fuqua School of Business, NBER, and CEPR

Paymon Khorrami
Duke University Fuqua School of Business

March 20, 2025

Abstract

How can researchers identify monetary shocks when (i) agents may possess subjective beliefs and (ii) monetary authorities manage current and future interest rates (e.g., forward guidance)? In the presence of risk premia, 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. Through the lens of these models, we quantify a significant bias, especially for long-horizon forward guidance, from mistakenly assuming Long-Run 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 Ravi Bansal, Harjoat Bhamra (discussant), Nuno Clara, Marc Dordal i Carreras, Marco Grotteria (discussant), Ben Hébert, Matteo Leombroni, Shri Santosh, Dean Ryu (discussant), Jonathan Wright (discussant), and conference and seminar participants at ASU Sonoran Winter Finance Conference, the Macro-Finance Society, EFA, Boston College, Maryland, Loyola-Chicago, PUC-Rio, SAIF, HKUST, UChicago Blue Collar Working Group, the Duke Finance Brownbag, and the Fuqua Summer Series for helpful feedback.

1 Introduction

How does a monetary authority influence the economy through speeches, forward guidance, and other policies targeting the future? In this paper, we uncover and explore an identification challenge associated with this question. The basic idea is that forward guidance and related policies operate by manipulating investor beliefs, which are not directly observable. Our key question becomes whether or not investors' policy-induced forecast revisions can instead be obtained indirectly by looking at rich enough collections of asset markets. We first show that identification of the various monetary shocks and their impacts is possible, perhaps surprisingly, but it 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. In other words, our results suggest forward guidance shocks cannot be identified from asset markets. Our results also suggest that macroeconomics should entertain models where monetary policy has permanent impacts, which are uncommon in the current literature.

Communication as policy. To fix ideas here, suppose central banks can impact the economy through two types 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

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). Of course, forward-guidance shocks may also occur in isolation without any short-rate shock, through direct communication to the market. The question in this paper is how to recover the various policy shocks, focusing particularly on communication shocks like forward guidance.[1]

Challenges to identification. More specifically, we ask whether forward guidance shocks can 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), at many levels of contingency (e.g., isolating specific aspects of the probability

[1]For simplicity and to engage with more novel identification issues, we suppose that we can identify the short-rate shocks. 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).

distribution), and available at a high frequency. 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, contingency, or frequency.

To recover expected future interest rates, a first idea might be to use long-term yields to reveal expectations (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. 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. For us, isolating beliefs is critical.

The solution: Long-Run Neutrality. Identification is not hopeless. Our core set of theoretical results says that forward-guidance shocks can be identified from asset prices if and only if permanent risks and their risk prices are unaffected by monetary policy. That is, in a world with risk premia, identification requires long-run risk prices to be invariant to monetary policy. 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. We clarify that this is not a coincidence: identification requires a Long-Run Neutrality assumption.

Our theoretical 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. The return of the growth-optimal portfolio in excess of a long-maturity bond identifies the martingale in the pricing kernel (Alvarez and Jermann, 2005). Long-Run Neutrality implies this investment strategy has exactly a zero return on Fed announcements. To repeat, the return should not only be zero on average but zero on each and every announcement.

To implement this test, we proxy the growth-optimal portfolio and long-maturity bond with equities and Treasury bonds, respectively, and study their return dynamics near Fed announcements.

Our evidence suggests Long-Run Neutrality is violated. Indeed, equity and long-term bonds behave very differently near Fed announcements. Broadly speaking, Long-Run Non-Neutrality suggests central bank policies may be more powerful, especially through their risk premia effects, than previously understood. For the specific aims of the present paper, Long-Run Non-Neutrality implies that forward-guidance and other types of communication shocks cannot be identified from asset prices alone.

Of course, an existing literature on "announcement effects" (which we cite below) has also studied returns near Fed announcements. What is new here is the tighter connection between our data analysis and our theory. For one, the bonds we investigate are longer-maturity than those commonly studied, an important contribution given our test asks for an arbitrarily long-maturity bond. Second, the existing literature often studies equities and bonds separately and in differing sample periods; we provide a consistent sample to study them jointly. Third, the overwhelming majority of the literature studies average announcement returns, whereas our test requires us to examine other moments beyond the mean. One of our most striking and novel findings is the large announcement volatility of the long equity, short bond investment strategy.

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, there is an identification paradox: learning the effects of monetary policy requires precisely that monetary policy has no effects.

Literature review. The specific applications of our paper connect most directly to a literature seeking to measure the impact of forward guidance and central bank communication. This literature often either extracts a reduced-form "path" factor from long-term bond yields (Gürkaynak et al., 2005b; Swanson, 2021; Altavilla et al., 2019; Lunsford, 2020) or sidesteps the rate path by projecting outcomes on communication directly (Hansen and McMahon, 2016; Leombroni et al., 2021; Gómez-Cram and Grotteria, 2022). By contrast, we seek to identify the shock to investor beliefs about the interest rate path. And thus our approach diverges from this literature.

Methodologically, our paper is closely related to a literature on belief recovery theory (Ross, 2015; Borovička et al., 2016; Qin and Linetsky, 2016). That is, we focus on environments with both risk premia and potentially subjective beliefs and seek to disentangle them.[2] In such environments, extracting monetary policy shocks, which are belief shocks, requires a Long-Run Neutrality condition on the permanent component of the pricing kernel. As mentioned above, Backus et al. (2022) and Haddad et al. (2023) effectively impose Long-Run Neutrality as identification assumptions.

Following a related literature that decomposes the pricing kernel into permanent and stationary components (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. Our test zooms in on asset price changes around FOMC announcements, so we contribute 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 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.[3] Models of this type 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). In these environments, Long-Run Neutrality requires that monetary policy does not affect the probability distribution of future growth.

[2]A large literature documents belief distortions 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).

[3]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 and Matthies, 2022).

2 Theory: Identifying Monetary Shocks

We 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. For a simple VAR example that motivates why we would want to recover investor surprises about interest rates, see Appendix B. That example also walks the reader through the basic identification challenges and solutions in simpler mathematical terms.

2.1 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 objective probability. We work exclusively in the realm of investors' subjective beliefs, because our goal is to identify changes in interest rates relative to the market beliefs.

The econometrician does not know the investor beliefs $\mathbb{P}$. More specifically, the econometrician 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 a 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.[4]

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 state-

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

ments 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. For simplicity, assume the mean of ξ_t is equal to zero, so that "expected jumps" are implicitly reflected in the drift of X_t.

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

$$dX_t = \mu(X_{t-})dt + \sigma(X_{t-})dW_t + \xi_t dM_t. \tag{1}$$

The sequence of information sets $(\mathscr{F}_t)_{t\geq 0}$ available to investors is generated by histories of W, M, ξ, and X_0. In other words, investors observe $(X_t)_{t\geq 0}$. We assume the same information set for the econometrician.

Asset prices and the SDF. We assume there exists a real stochastic discount factor (SDF) process S whose increment is given by

$$\frac{dS_t}{S_{t-}} = -r(X_{t-})dt - \pi(X_{t-}) \cdot dW_t + \exp[\kappa(X_t, X_{t-})] - 1 - \chi_S(X_{t-})dt, \tag{2}$$

where $S_0 = 1$ and $\chi_S(x)dt$ is the jump compensator.[5] The variable $r(x)$ denotes the short-term real 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 real payoff $f(X_T)$ is

$$\mathbb{E}\Big[\frac{S_T}{S_t} f(X_T) \mid X_t\Big]. \tag{3}$$

That is, S represents the SDF from the perspective of investors and their beliefs.

In this environment, Hansen and Scheinkman (2009) show how Perron-Frobenius

[5]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 $\chi_S(x) := \int \big(\exp[\kappa(x', x)] - 1\big)\chi(dx' \mid x)$.

Theory can be leveraged to obtain a factorization of the SDF as

$$\frac{S_{t+T}}{S_t} = \exp(\eta T)\frac{e(X_t)}{e(X_{t+T})}\frac{H_{t+T}}{H_t}. \tag{4}$$

In (4), $\exp(\eta)$ 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 factorization (4). For the purposes of this section, we also assume the factorization is unique.[6]

Equation (4) 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 4. Using non-parametric methods, Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012) argue that H must play a significant role in pricing.

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 shocks are defined as follows.

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

$$z_\tau^0 := r(X_\tau) - \mathbb{E}[r(X_\tau) \mid X_{\tau-}]. \tag{5}$$

The shocks to the expected future short rates are given by

$$z_\tau^T := \mathbb{E}[r(X_{\tau+T}) \mid X_\tau] - \mathbb{E}[r(X_{\tau+T}) \mid X_{\tau-}], \quad T > 0. \tag{6}$$

[6]See Hansen and Scheinkman (2009) for sufficient conditions on the existence of such a factorization. In many cases, uniqueness will not hold. If there are multiple SDF factorizations satisfying (4), 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.

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

$$p_\tau^T(\mathsf{r}) := \mathbb{P}\{r(X_{\tau+T}) \leq \mathsf{r} \mid X_\tau\} - \mathbb{P}\{r(X_{\tau+T}) \leq \mathsf{r} \mid X_{\tau-}\}, \quad T > 0. \tag{7}$$

Remark 1 (Real versus nominal shocks)**.** *While the current environment allows inflation (i.e., it could be part of the state vector x), we choose to work directly monetary shocks to the* real interest rate $r(x)$ *to avoid going back and forth between nominal and real quantities. That said, all of our identification results below would continue to hold when applied to the nominal interest rate $i(x)$. Indeed, we will effectively show that identifying belief updates about* any *function of the economic state x hinges on our Long-Run Neutrality condition.*

Why do we care about these reduced-form monetary shocks? Motivated by Appendix B (see Procedure 2), regressing future outcomes on both z_τ^0 and z_τ^T, along with controls for the lagged state $X_{\tau-}$ can potentially identify the causal impact of short rates and forward guidance, jointly. If there are more dimensions to forward guidance (e.g., guidance at different horizons), one should include multiple horizons of the expected future short rate shocks $(z_\tau^T)_{T>0}$. If there is guidance about the distribution of future interest rates, one should consider including additional moments of the distributional shocks $(p_\tau^T)_{T>0}$. To run these procedures, we need to identify these reduced-form shocks.

Can an econometrician identify the impacts of the central bank on current and future interest rates? In this paper, we take as given the ability to non-parametrically identify the short rate shock from data. 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.

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 (7).

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 factorization (4) 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

$$\hat{\mathbb{P}}(F) := \mathbb{E}[1_F H_t], \quad \forall F \in \mathscr{F}_t. \tag{8}$$

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

$$\hat{\mathbb{E}}[r(X_{\tau+T}) \mid X_\tau] - \hat{\mathbb{E}}[r(X_{\tau+T}) \mid X_{\tau-}] = \mathbb{E}[r(X_{\tau+T}) \mid X_\tau] - \mathbb{E}[r(X_{\tau+T}) \mid X_{\tau-}], \tag{9}$$

and similarly for other moments of $r(X_{\tau+T})$. If equality (9) 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, the 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 $d\log(H_t)$ *only depends on* monetary-insensitive economic states $X_t^\perp$,

where monetary-insensitivity is defined by $\mathbb{E}[f(X_{\tau+T}^\perp) \mid X_\tau] - \mathbb{E}[f(X_{\tau+T}^\perp) \mid X_{\tau-}] = 0$ *for all bounded functions* $f(\cdot)$*, all* $T > 0$*, and all monetary announcement dates* τ.

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. Importantly, Long-Run Neutrality does not imply constant or exogenous risk premia. Indeed, going back to the factorization (4), marginal utility contains both a permanent piece H_t and a stationary piece $e(X_t)$. Long-Run Neutrality leaves open the possibility that monetary policy impacts risk premia by moving the stationary piece $e(X_t)$.

For H to satisfy Definition 2, it must take the form

$$H_t = \exp\Big[-\frac{1}{2}\int_0^t \|\beta(X_s^{\perp})\|^2 ds - \int_0^t \beta(X_s^{\perp}) \cdot dW_s \Big] \tag{10}$$

for some function $\beta : \mathbb{R}^{\dim(X^{\perp})} \mapsto \mathbb{R}^n$, where $X^{\perp}$ denotes the sub-vector of monetary insensitive states. One can interpret β as the exogenous 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.

Remark 2 (Real SDF)**.** *While our environment allows inflation, we have decided to work with the* real SDF process *S. This is to make our analysis more interesting: whether or not Long-Run Neutrality holds for S is not obvious. For the nominal SDF $\frac{S}{P}$, a condition analogous to Long-Run Neutrality is much more likely to fail. Monetary policies, if they affect inflation and the price level P, are very likely to exert permanent impacts on the nominal SDF. For instance, even in a superneutrality world where the real economy is completely supply driven, the nominal SDF would necessarily move permanently in response to monetary shocks.*

2.2 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

$$\mu(x) = A_0 + Ax \tag{11}$$

$$\sigma(x) = B, \tag{12}$$

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:

$$r(x) = \rho_0 + \rho \cdot x, \tag{13}$$

for some constant ρ_0 and some vector ρ. Assuming (13) holds in a linear environment with (11)-(12) is tantamount to an assumption on the evolution of S (for example, the exponential-affine model of Section B.3 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 (11)-(13). 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 drift $B\beta(X_t^\perp)$:

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

While $\beta(\cdot)$ and A_0 are not separately identified, the long-term measure $\hat{\mathbb{P}}$ correctly identifies investors' perceived persistence A. The simplest case to see this is if there are no monetary-insensitive states, i.e., $\beta(x^\perp) \equiv \beta$ constant, in which case the drift state-dependence is AX_t under both $\mathbb{P}$ and $\hat{\mathbb{P}}$. This turns out to be the critical necessary object to compute investors' forecast revisions. By contrast, drift distortions like $B\beta(x^\perp)$ play no role in these forecast revisions, because investor forecasts just before and just after the monetary announcement are both distorted by the same amount. 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 $\mathbb{E}[X_{\tau+T} \mid X_\tau] - \mathbb{E}[X_{\tau+T} \mid X_{\tau-}]$.

Ultimately, Proposition 1 is just a generalization of what we observed in our example in Section B.3. 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.

2.3 Approximate identification with stochastic volatility

We continue to assume a linear drift (11) and a linear short rate function (13), but we dispense with homoskedasticity (12). 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 (10) characterizes the permanent component of the SDF, we assume there exists some function $\hat{\beta} : \mathbb{R}^{\dim(X^\perp)} \mapsto \mathbb{R}^n$ such that

$$\sigma(x)\beta(x^\perp) = \hat{\beta}(x^\perp) \quad \text{for all } x. \tag{14}$$

In other words, there is a zero in each element of β corresponding to a shock with non-constant volatility that also affects any monetary-sensitive state. Replacing homoskedasticity assumption (12) with the more general (14), 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 (11), (13), and (14). 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, (14) implies that the drift of X_t under the long-run measure $\hat{\mathbb{P}}$ is

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

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:

$$\sigma(x)\sigma(x)' = \varsigma_0\varsigma_0' + \sum_{i=1}^{n} \varsigma_i \mathrm{diag}(x_i)\varsigma_i', \tag{15}$$

where $\mathrm{diag}(x_i)$ is the diagonal matrix with x_i on the main diagonal. The affine approximation in (15) 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 (11), (13), (14), and (15). Define the variance surprises*

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

Then, on the event $\{\tau' > \tau + T\}$, $(v_\tau^T)_{T\in[0,\tau'-\tau)}$ 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 (14) that volatility shocks feature zero long-run risk prices.

2.4 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 (11)-(13) in Section 2.2. 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

$$\frac{dH_t}{H_{t-}} = \exp[\zeta \cdot (X_t - X_{t-})] - 1 - \chi_H(X_{t-})dt, \tag{16}$$

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[7]

$$\hat{\mu}(x) = \mu(x) + \int (x' - x)\Big(\exp[\zeta \cdot (x' - x)] - 1\Big)\chi(dx' \mid x), \tag{17}$$

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 $\int (x' - x)\big(\exp[\zeta \cdot (x' - x)] - 1\big)\chi(dx' \mid x)$ only depends on the monetary-insensitive states $x^{\perp}$. This 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

[7] Formally, under investor beliefs $\mathbb{P}$, the jumps $\xi_t dM_t$ have conditional mean $\int (x' - x)\chi(dx' \mid x)dt$. (Of course, we have assumed from the beginning that this conditional mean equals zero, but the argument applies more generally even when this is not the case.) 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 (17) in the text.

dM_t. In this case, rather than the log-normal form (10), H_t takes the form

$$H_t = \exp\Big[-\frac{1}{2}\int_0^t \|\beta(X_s)\| ds - \int_0^t \beta(X_s)\cdot dW_s \Big] \tag{18}$$

for some function $\beta(\cdot) : \mathbb{R}^n \mapsto \mathbb{R}^n$ that we suppose depends non-trivially on the monetary-sensitive states. The drift of X_t under the long-run measure $\hat{\mathbb{P}}$ is given by

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

Although $\hat{\mu}(x)$ is observable, we cannot separately distinguish between Ax and $B\beta(x)$. Previously when $\beta(\cdot)$ only depended on monetary-insensitive states $x^\perp$, we could identify A as the persistence of the monetary-sensitive states under $\hat{\mathbb{P}}$. Here instead, 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

$$\begin{aligned}
\mathbb{E}^{X_\tau}[X_T] - \mathbb{E}^{X_{\tau-}}[X_T] &= X_\tau - X_{\tau-} + \mathbb{E}^{X_\tau}[\int_0^T \mu(X_{t-})dt] - \mathbb{E}^{X_{\tau-}}[\int_0^T \mu(X_{t-})dt] \\
&= X_\tau - X_{\tau-} + \int_0^T A\Big(\mathbb{E}^{X_\tau}[X_t] - \mathbb{E}^{X_{\tau-}}[X_t]\Big)dt. \qquad (19)
\end{aligned}$$

Equation (19) is a recursive equation for $\mathbb{E}^{X_\tau}[X_T] - \mathbb{E}^{X_{\tau-}}[X_T]$, 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 (16) or (18) can be formalized. We have

Proposition 4. *Consider the linear environment defined by (11)-(13). Suppose either (i) dH_t features a contribution from dM_t; or (ii) the dynamics dH_t depend on X_t beyond the monetary-insensitive states $X_t^\perp$. 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. (The qualification "generically" comes from the fact that the proposition must rule out knife-edge cases like $\lambda(X_t)$ and ξ_t being state-independent.)

3 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. We evaluate the test with our own novel evidence, followed by a comparison to the existing literature.

3.1 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

$$R^{\infty}_{t,t+\Delta} := \lim_{T\to\infty} R^{T}_{t,t+\Delta} = \lim_{T\to\infty} \frac{\mathbb{E}[\frac{S_T}{S_{t+\Delta}} \mid X_{t+\Delta}]}{\mathbb{E}[\frac{S_T}{S_t} \mid X_t]} = \exp(-\eta)\frac{e(X_{t+\Delta})}{e(X_t)} \lim_{T\to\infty} \frac{\mathbb{E}[\frac{1}{e(X_T)}\frac{H_T}{H_{t+\Delta}} \mid X_{t+\Delta}]}{\mathbb{E}[\frac{1}{e(X_T)}\frac{H_T}{H_t} \mid X_t]}$$
$$= \exp(-\eta)\frac{e(X_{t+\Delta})}{e(X_t)}, \tag{20}$$

where the last equality holds if X_t is stochastically stable under the probability measure generated by H, which we implicitly assume (see footnote 6). The link between the long bond and the stationary component of the pricing kernel was originally discovered by Kazemi (1992). 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 $R^{*}_{t,t+\Delta} = \frac{S_t}{S_{t+\Delta}}$ (Bansal and Lehmann, 1997). In an incomplete market, where there are multiple possible SDFs, the growth-optimal portfolio will proxy the entropy-minimizing SDF. Putting these results together, and using the SDF factorization (4), the excess return of the growth-optimal portfolio relative to the infinite-horizon bond is

$$\log(R^{*}_{t,t+\Delta}) - \log(R^{\infty}_{t,t+\Delta}) = \log\left(\frac{H_t}{H_{t+\Delta}}\right) \tag{21}$$

over any horizon Δ. The long-short portfolio in (21) was originally studied by Alvarez and Jermann (2005). The result in (21) 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 $\log(R^{*}_{t,t+\Delta}) - \log(R^{\infty}_{t,t+\Delta})$ should be identically zero on monetary announcement days. Under condition (ii) of

Definition 2, the conditional risk premium $\mathbb{E}[\log(R^*_{t,t+\Delta}) - \log(R^\infty_{t,t+\Delta}) \mid X_t]$ should be time-invariant.

Equation (21) suggests a test: one can examine high-frequency changes in $R^*_{t,t+\Delta}$ and $R^\infty_{t,t+\Delta}$ 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 $\mathbb{E}[\log(R^*_{t,t+\Delta}) - \log(R^\infty_{t,t+\Delta}) \mid X_t]$, because it could be so under investor beliefs but not under the objective measure.)

The most important question is how to proxy $R^\infty_{t,t+\Delta}$ and $R^*_{t,t+\Delta}$. Following the literature, we will suppose $R^\infty_{t,t+\Delta}$ is well-approximated by returns on 30-year Treasuries. Then, we will assume $R^*_{t,t+\Delta}$ is well-approximated by aggregate stock market returns. Under these proxies, Long-Run Neutrality says that monetary policy impacts the stock market and 30-year Treasuries in an identical way. Note that we have been expositing Long-Run Neutrality as a condition for the *real* SDF, and so R^∞ should be a *real* bond. Therefore, we are effectively assuming that R^* is the inflation-hedged stock market return. Indeed, our empirical measures are a nominal bond and the nominal stock market, so inflation mechanically differences out in the calculation of the long-short return.

Of course, there will be two sources of error in our proxies. Long-term bonds are still not quite infinite-maturity bonds, and the stock market is not necessarily the investor-perceived growth-optimal portfolio. We view the imperfection in the bond proxy as slightly less troubling. For one, Nagel and Xu (2024) show in their appendix that average FOMC announcement day returns of long-term bonds (maturities 9, 15, 20, and 30 years) are statistically indistinguishable from each other. They show a similar result for these bonds' responses to short-rate shocks on FOMC days. Second, bonds of all maturities almost always respond to monetary announcements with the *same sign*. Assuming a hypothetical infinite-maturity bond would share this property with finite-maturity bonds, a violation of Long-Run Neutrality is suggested by the fact that the stock market and bond market returns have opposite signs on monetary announcement days over 55% of the time.[8]

On the other hand, our growth-optimal proxy is potentially more problematic, in that it stands in contrast to voluminous evidence against the CAPM and various asset-pricing anomalies. That said, we see two reasons why our results are likely to be robust. First, as shown by Savor and Wilson (2014) and Lucca and Moench (2015), stock returns are very

[8]We thank Jonathan Wright for pointing out this fact and explaining how it can be used to ameliorate our proxy issue.

well-explained by the CAPM *on FOMC announcement days*, unlike other days. Second, we seek the investor-perceived growth-optimal portfolio, rather than the objective one. There is some new evidence that investors and analysts *actually use the CAPM* in modeling (Dessaint et al., 2021; Décaire et al., 2023), which if true suggests they do perceive the market as the growth-optimal portfolio.

3.2 Existing announcement effect evidence

Given that we will proxy $R^{\infty}_{t,t+\Delta}$ and $R^{*}_{t,t+\Delta}$ by US Treasury and stock market returns, respectively, our empirical results relate to a large literature on FOMC announcement effects on stocks and bonds. Together, this collection of existing evidence suggests that H may respond to monetary policy.

First, there is some evidence that stock returns are higher than long-term bond returns around FOMC meeting days, on average. In particular, Lucca and Moench (2015) show (for 1994–2011) that the SPX return was on average 33 bps higher in the 24 hours before the FOMC announcements, relative to other days. By contrast, Hillenbrand (2021) shows (for 1989–2021) that 30-year Treasuries returns were approximately 13.8 bps to 18.6 bps higher per day in a 3-day window around the FOMC announcements, relative to other days.[9] This evidence suggests that the average returns on stocks and long-term bonds differ both in magnitudes and the timing around the FOMC announcements.

Although our focus is primarily on monetary announcements, it is worth reviewing evidence from a more comprehensive set of macroeconomic announcements. The broader announcement literature has argued that other macro announcements also induce asset price responses, suggesting that the mechanisms generating announcement premia around macroeconomic and policy announcements can be related. Using a much longer sample (1958-2009), Savor and Wilson (2013) report returns on stocks and 30-year Treasury bonds to be respectively 11.5 bps and 4.5 bps higher on announcement days. Additionally, 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.

A parallel strand of research examines asset responses to monetary policy surprises rather than average announcement returns. This literature relies on surprises identified from high-frequency short-term interest rate changes around FOMC announcements. For example, Gürkaynak, Sack and Swanson (2005b) study high-frequency responses

[9] We impute this range for the 30-year bond average return using evidence in Hillenbrand (2021) that 30-year Treasury yields decline between 0.46 bps to 0.62 bps more per day in the 3-day window surrounding FOMC meetings. We use the duration-approximation $\log(R^{T}_{t,t+\Delta}) \approx -T(y_{t+\Delta} - y_t)$, with $T = 30$.

to monetary surprises (during 1990–2004), finding that a 25 bp surprise rate cut leads to 1% SPX return but only a 0.32% 10-year Treasury return. Using an updated 1988–2019 sample, Bauer, Bernanke and Milstein (2023) find even stronger effects in stocks, with a 10 bps surprise cut associated with a 1% SPX return, although they do not study long-term Treasuries. The recent studies by Nagel and Xu (2024) and Boehm and Kroner (2023) perform detailed comparisons of stock and long-term bond returns during FOMC announcements, and we will compare and contrast our results to theirs below.

3.3 New evidence on Long-Run Neutrality

While suggestive, the evidence from the extant literature does not speak directly to the Long-Run Neutrality of the Fed-driven news. The samples vary substantially in length, choice of event windows, and are often inconsistent between equities and bonds, with little direct evidence on the behavior of long-short equity-bond portfolios. Importantly, most studies document average asset-specific announcement effects, ignoring higher moments. Below, we use a consistent sample for equities and long-duration bonds, compare results for various windows around FOMC announcements, and investigate other moments of equity-bond portfolios beyond the mean effect.

Following much of the literature, we consider scheduled FOMC decision announcements. The FOMC meets eight times per year on a pre-announced schedule.[10] Since 1994, the FOMC has published statements of the policy decisions, and since 2011, the statements have been followed by press conferences by the Fed Chair, initially every other announcement, and starting in 2019, after each announcement. Until 2011, statements were released at 14:15 ET. The time changed in 2011, alternating between 12:30 and 14:15 ET, depending on whether the meeting was followed by a press conference. Currently, statements are published at 14:00 ET, and the press conference is held at 14:30 ET.

We obtain price data at a one-minute frequency on the E-mini S&P 500 futures and the Treasury bond (T-bond) futures from TickData.com. Our main focus is on the Treasury futures with a 30-year T-bond as the underlying, the longest maturity available. We refer to this contract as the 30-year T-bond futures, recognizing that the actual delivery can take place in bonds with maturities of 15 years and above.[11] Our high-frequency sample

[10] In 2020, one scheduled meeting was canceled.

[11] In January 2009, the CME introduced a new 30-year Treasury bond futures contract, called "Ultra," which requires delivery of a bond with at least 25-year maturity. At that time, the range of eligible maturities for the original or "classic" 30-year T-bond futures was adjusted from a 15–30-year range to a 15–25-year range. Our current analysis focuses on the classic 30-year contract.

Panel A. Monetary policy decision announcements

	N	Mean	SE(mean)	SD	Skew	Kurtosis
(−24h,−15m)	210	26.0	8.2	118.2	1.8	18.6
(−15m,+15m)	210	-1.3	4.1	59.5	-1.5	11.3
(−15m,+24h)	210	-20.1	12.3	178.2	-0.8	6.0

Panel B. Press conferences

	N	Mean	SE(mean)	SD	Skew	Kurtosis
(−24h,−15m)	70	19.8	9.1	76.1	0.0	3.7
(−15m,+15m)	70	9.4	4.7	39.1	1.2	5.5
(−15m,+60m)	70	2.4	8.2	68.4	0.2	4.3
(+60m,+24h)	70	-43.5	19.2	160.5	-1.0	5.7

Table 1. Summary statistics. The table reports summary statistics for log returns on S&P500 E-mini futures minus log returns on 30-year Treasury bond futures, in basis points, in various windows around scheduled monetary FOMC decision announcements and press conferences. A futures "return" in window $(t, t+\Delta)$ is defined as $F_{t+\Delta}/F_t$, where F denotes the futures price. The sample covers FOMC meetings from 1997:09 through 2023:12, with press conferences introduced in 2011.

starts in September 1997 when the E-mini S&P 500 futures contract was introduced and runs through December 2023, covering 210 scheduled FOMC announcements and 70 press conferences.

We consider event windows from 24 hours before, narrowly around, and up to 24 hours after the FOMC decision announcements and press conferences. Table 1 summarizes the distribution of log returns of E-mini futures in excess of 30-year T-bond futures in different windows. The results in Panel A show that equities have done particularly well in the 24 hours before the FOMC announcement, earning on average 26 bps higher returns than bonds during the 1997–2023 sample. This result is consistent with, albeit somewhat weaker than, the original Lucca and Moench (2015) finding based on the 1994–2011 sample. While on average equities performed similarly to bonds in the narrow window of ±15 minutes around the announcements, they underperformed bonds by about 20 bps in the 24 hours after the announcement. Panel B summarizes returns around press conferences, showing a broadly similar pattern.

The narrow event windows are particularly informative given that the FOMC-driven news is plausibly the main source of variation in asset prices at those times. The equity-bond portfolio returns show a volatility of nearly 60 bps in the ±15 minutes window of the decision announcement and 40 bps in the ±15 around the start of the press conference, which cover the opening remarks by the Chair. For comparison, the volatility is

27 bps in the ±15-minutes window around 14:00 ET on all other days in the 1997:09–2023:12 sample, i.e., less than half of that around the FOMC decision announcements. The insignificant average ±15-minutes announcement return in Table 1 thus masks a sizeable time variation in returns around announcements. To illustrate that variation, Figure 1 plots the cumulative equity-bond portfolio returns obtained by summing the ±15 window decision announcement returns across announcements. The cumulative returns show a persistent downward trend in the first half of the sample, reaching -1004 bps at the November 2009 meeting, and from then onward a persistent upward trend through the end of our sample. Thus, equities underperformed long-term bonds on FOMC announcements from the late 1990s through the first rounds of the quantitative easing implemented after the global financial crisis but outperformed long-term bonds afterward through the most recent period.

While speculative, let us mention a theory that highlights time-varying Fed *credibility* as a possibility for the reversal shown in Figure 1. A long literature beginning with Barro and Gordon (1983) explores the role of a central bank's reputation in modulating the efficacy of its policy. In a recent contribution, Bhattarai et al. (2023) argue that quantitative easing programs can increase central bank commitment power to keep interest rates low, which amplifies the power of forward guidance. This mechanism was particularly relevant after 2008, when the binding zero lower bound made the combination of quantitative easing and forward guidance the relevant policy tools.

If the Long-Run Neutrality condition holds, the FOMC-driven news should move equity and T-bond returns one for one, implying that a regression of T-bond returns on equity returns should have a slope coefficient (beta) and an R-squared both equal to one. Figure 2 shows that these predictions are rejected in the data. Bond-equity betas are statistically different from one across various windows around decision announcements and press conferences. While the ±15-minute betas are positive, they are significantly below one, reaching 0.32 for decision announcements and 0.22 for press conferences, with R-squareds of 9.9% and 20.1%, respectively. Outside narrow windows, the evidence against Long-Run Neutrality strengthens further, with even lower R-squared and betas becoming negative. This evidence echoes Boehm and Kroner (2023), who document that yield curve responses to FOMC announcements contain low explanatory power for stock market responses to those same announcements.

To assess how often the Long-Run Neutrality condition could potentially hold in our sample, we compute high-frequency betas and R-sqaureds using realized covariances and variances of one-minute equity and T-bond returns in the ±15 window around decision announcements. The histograms in Figure 3 indicate that 90% of announcements

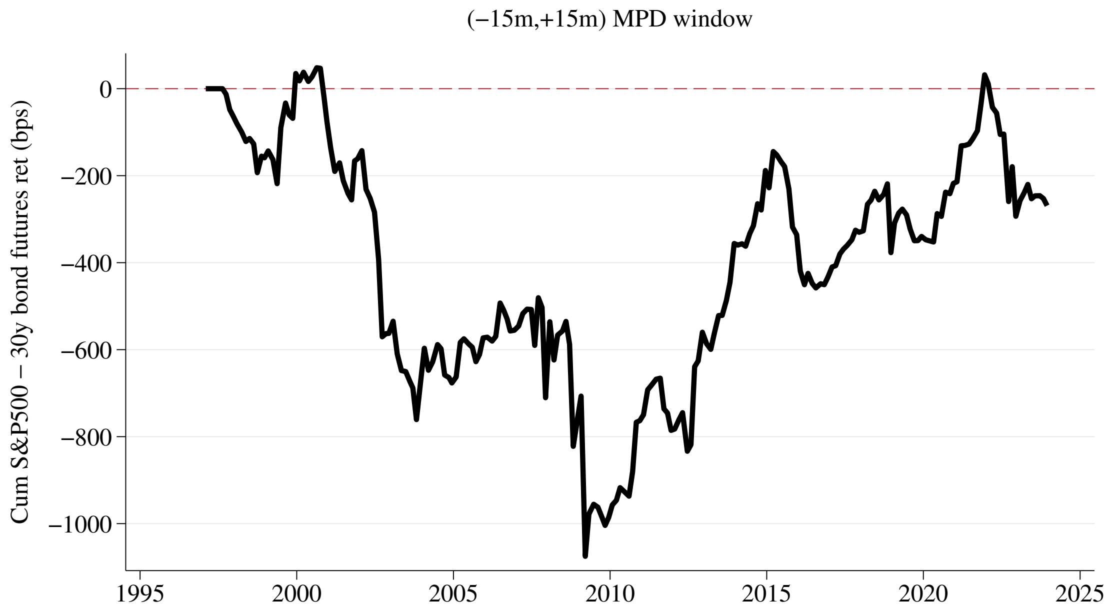

Figure 1. Cumulative returns in a narrow window around FOMC policy announcements. The figure plots cumulative log returns on a long-short portfolio of S&P 500 E-mini futures in excess of 30-year Treasury bond futures. The returns are calculated from −15 to +15 minutes around scheduled FOMC announcements. The sample period covers 210 meetings from 1997:09 through 2023:12.

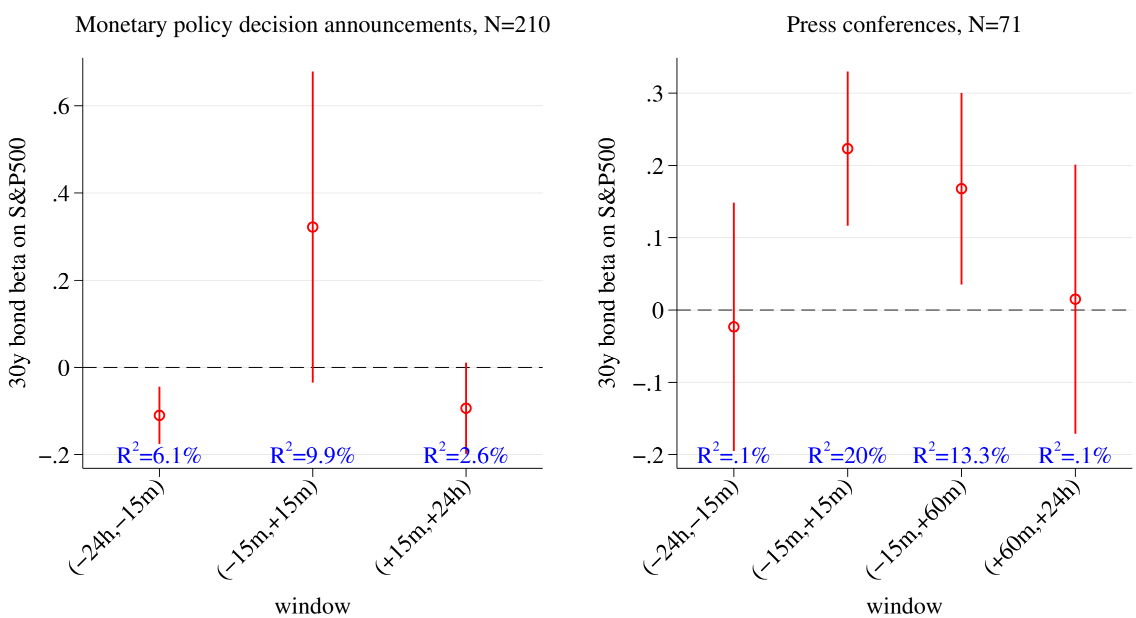

Figure 2. Bond betas. The figure presents betas of 30-year Treasury bond futures returns on S&P 500 E-mini futures returns in different windows around scheduled FOMC announcements and press conferences. Robust 95% confidence intervals are included.

feature a beta below 0.5 and an R-squared below 0.42, again suggesting that the Long-Run Neutrality most of the time remains violated.

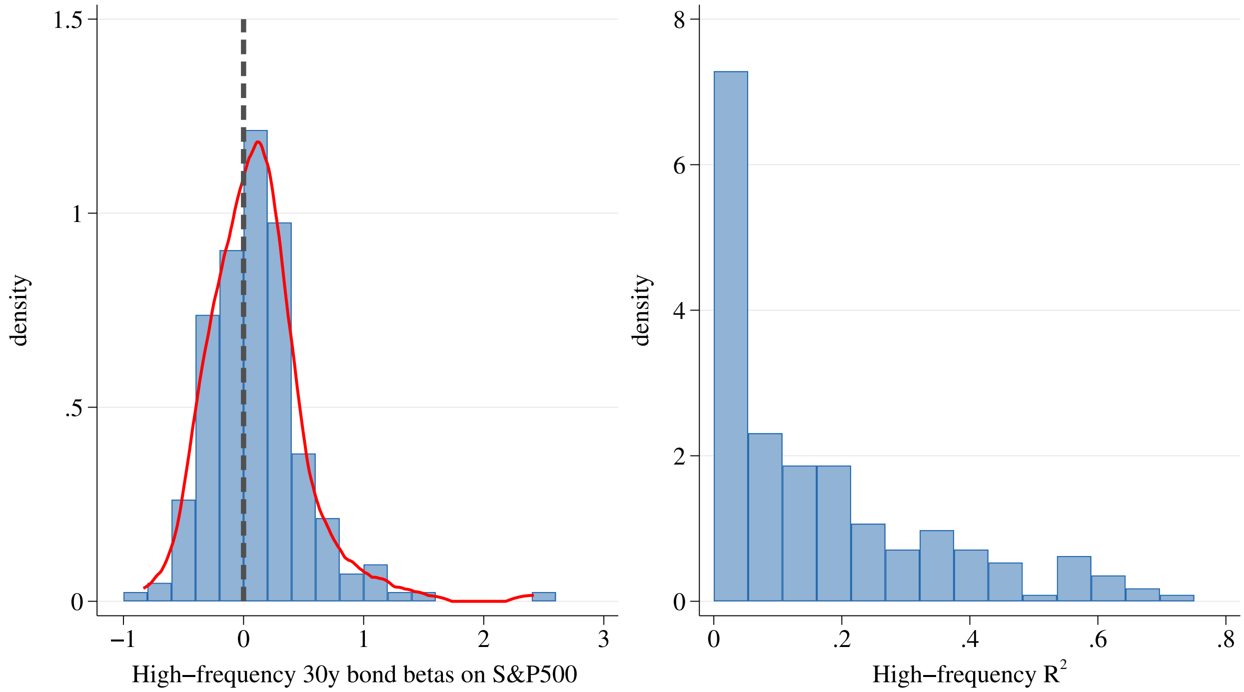

Figure 3. High-frequency bond betas. The figure presents the distribution of high-frequency realized betas of the 30-year Treasury bond futures returns on S&P 500 E-mini futures returns. For each ±15-minutes window around scheduled FOMC decision announcements, the beta and R-squared are calculated from realized covariances and variances of one-minute returns.

It is worth comparing our results to the recent study by Nagel and Xu (2024), who revisit the seminal analysis of Bernanke and Kuttner (2005) and argue that stock returns on FOMC announcement days, in response to short-rate shocks, are entirely attributable to a change in the term premium. More specifically, they maturity-match equity dividend strips to Treasury bonds and show that the returns of these two portfolios, on average in response to a short-rate shock, are equal. Our analysis differs in several respects. First, the theoretically appropriate object for us is not a Treasury whose duration matches that of the stock market but rather one whose duration is as high as possible. Second, with tools such as forward guidance and quantitative easing, monetary policy does much more than change the short rate. By looking at returns during the FOMC announcements, we are effectively pooling responses to all monetary shocks, as we should, whereas Nagel and Xu (2024) look only at short-rate shocks. Third, the effects *on average* mask tremendous time-variation: in our data, looking at 30-minute windows around FOMC announcements, equities and long-term bonds do earn similar *average* returns (Table 1, row 2), but their returns are only weakly correlated (Figures 2-3). Again, this weak correlation echoes the results in Boehm and Kroner (2023).

One concern with the results so far is that the bond underlying the 30-year T-bond futures contract is a poor proxy for the R^{∞} return, given that the effective duration of the underlying is about 15 years on average. Therefore, we perform additional analysis using daily zero-coupon yield curve data regularly updated by the Federal Reserve (Gurkaynak et al., 2007). The longest reported zero-yield maturity is 30 years, which is also the maximum maturity issued by the US Treasury. To proxy for the market portfolio, we use the daily CRSP market return from Ken French's website. We construct the equity-bond portfolio as the log return on the CRSP market return in excess of the log return on the 30-year zero coupon bond.

Figure 4, left panel, displays the cumulative return on the equity-bond portfolio in the three-day window (days $-1, 0, +1$) around scheduled FOMC announcements. The sample starts in 1994, when the FOMC began releasing public statements. The choice of the window is motivated by the finding in Hillenbrand (2021) that Treasury bonds earn essentially all returns in the three days surrounding the announcement. The right panel in Figure 4 presents cumulative returns disaggregated by day $-1, 0, +1$ around the announcement. To assess the magnitudes, we juxtapose the cumulative FOMC window returns against the cumulative returns on all other days and scale each by the total number of days in the respective sample. There are 717 days falling in the three-day FOMC window and 7108 days falling outside. The last observation along each trajectory represents the sample average return and is reported on the graph.

The return trajectories indicate that, in economic terms, the equity-bond portfolio has earned larger (in absolute value) returns in the FOMC window than on all other days, consistent with the idea that FOMC-driven news is associated with deviations from the Long-Run Neutrality condition. At the same time, the disaggregated results in the right panel of Figure 4 reveal a complex interpretation of the directional effect of the FOMC news in three-day announcement windows, suggesting that the equity-bond portfolio undergoes regime-like shifts and can switch sign over time and across specific days. In particular, most of the deviations from neutrality appear in years 2008-2012, but oppositely on announcement day versus surrounding days. The evidence suggests that the permanent component of marginal utility falls substantially on announcement days but rises substantially the day before and after.

The balance of our evidence on the log excess return $\log(R^{*}_{t,t+\Delta}) - \log(R^{\infty}_{t,t+\Delta})$ suggests that monetary Long-Run Neutrality is violated. This violation is visible in the mean return, in line with the existing literature, but the stronger violations appear in higher moments: the long-short portfolio displays significant time-variation, and its two legs co-move weakly in windows surrounding FOMC announcements. Beyond documenting these higher moments, our analysis also contributes by studying various time windows

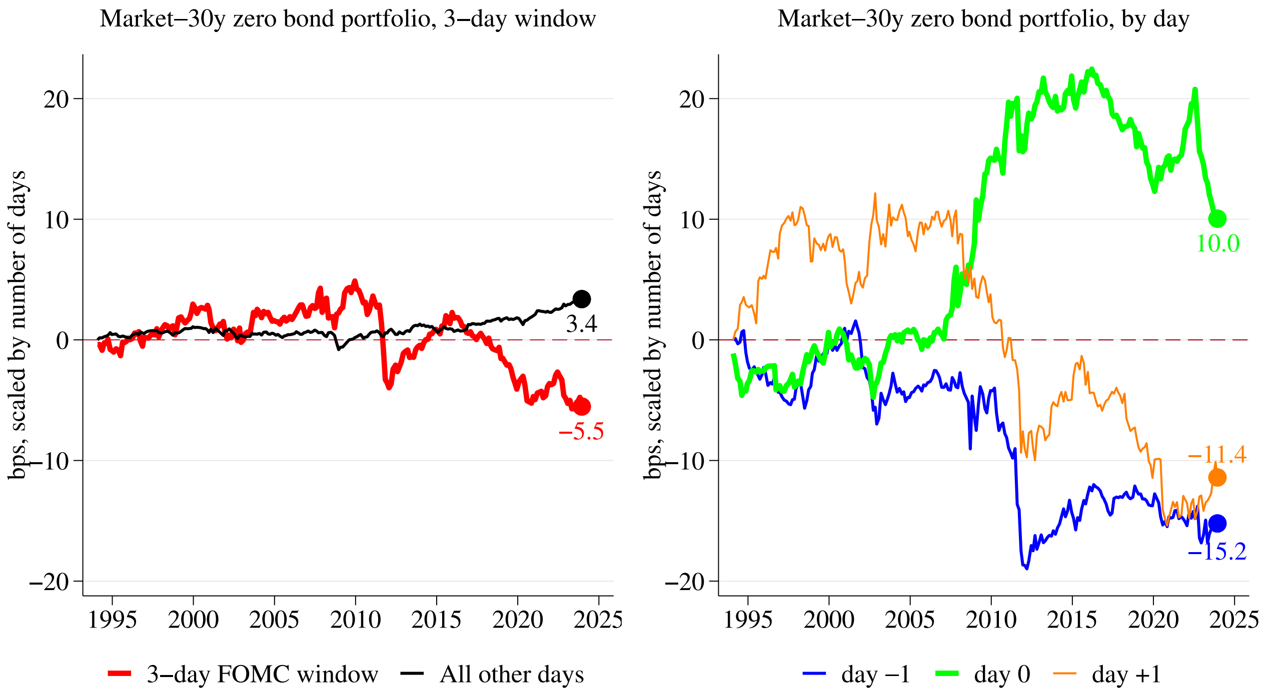

Figure 4. Market-bond portfolio returns, daily frequency. The figure presents cumulative daily close-to-close log returns on the market portfolio in excess of the 30-year zero-coupon Treasury bond. The three-day FOMC window comprises days -1, 0, and $+1$ around the scheduled FOMC announcements. Cumulative returns are expressed in basis points and scaled by the number of days in a given sample. The sample runs from 1994:01 through 2023:12. There are 717 days in the three-day FOMC window (239 on each of days $-1, 0, +1$) and 7108 on all other days outside the three-day FOMC window.

and return frequencies. We uncover several nuanced patterns deserving of further investigation, such as the reversal around 2009 of the long-short portfolio's performance in narrow windows around FOMC, and the opposite performance on announcement day versus the surrounding days.

4 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.

4.1 Long-run risk model

Suppose aggregate consumption has the following trend-stationary subjective 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:

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

where $\gamma > 1$ denotes the investor's coefficient of relative risk aversion and β is the subjective discount factor. Guess that the solution is $U_t = A_0 t + u_0 + u \cdot X_t$, for some constant u_0 and vector u. In that case, one can show that the solution is

$$u = (I - A')^{-1}\alpha \quad \text{and} \quad u_0 = (1-\beta)^{-1}[\beta A_0 + \frac{1}{2}(1-\gamma)u'BB'u].$$

In this model, the investor marginal utility (investor SDF) is given by

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

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

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

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 $u'B^{(i)} = 0$ 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 Long-Run Neutrality requires $u'B^{(1)} = u'B^{(2)} = 0$. But since the elements of $u = (I - A')^{-1}\alpha$ are generically non-zero, the requirement implies that $B^{(i)} = 0$. In words, Long-Run Neutrality here requires that *growth, in both the short and long run, is invariant to monetary policy*.

4.2 Stochastic-volatility model

Consumption has the following perceived 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 and $\gamma > 1$. In this model, that means that the investor SDF takes the form:[12]

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

where $u < 0$ is the larger root $(\gamma - 1)\sigma^2 u^2 + 2(\beta^{-1} - a)u + (\gamma - 1) = 0$.

In this environment, it is easy to verify that the stationary component of the SDF is characterized by the eigenfunction $e(x) = \exp(\kappa x)$, where κ is a root of the quadratic equation $0 = \frac{1}{2}\sigma^2\kappa^2 - [(\gamma - 1)\sigma^2\kappa + (1 - a)]\kappa + \frac{1}{2}(2\gamma - 1)$ (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

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

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

[12] See Appendix C for these calculations in a more general model where g is also stochastic.

4.3 Discussion and takeaways

The preceding models clarify that Long-Run Neutrality may be a severe assumption. Through the lens of models where news shocks are priced, monetary authorities who influence the probability distribution of future real outcomes will often violate Long-Run Neutrality. Even more broadly, the dynamics in the models above are specified in terms of investor beliefs. So even if the monetary authority merely conveys information about growth to the public without impacting it in a true causal sense, Long-Run Neutrality will be violated.

An additional takeaway from both of these models is that Long-Run Neutrality and conventional notions of monetary neutrality can differ. In the first model, consumption is trend-stationary, so monetary policy can never impact long-term consumption, even though it could impact H. In the second model, volatility is stationary, so long-horizon consumption is fully determined by the level shock; thus, monetary policy could affect H through volatility without affecting long-term consumption. These two example models are hard-wired to exhibit conventional neutrality even though Long-Run Neutrality fails.[13]

On the other hand, it is possible to come up with models where both conventional neutrality fails but Long-Run Neutrality holds. For instance, consider a conventional long-run risk setting where consumption grows over time and features an autoregressive process for the conditional expected growth rate (Bansal and Yaron, 2004). If monetary policy exhibits a persistent impact on growth rates, conventional neutrality automatically fails. However, Long-Run Neutrality will hold in the absence of recursive utility preferences. The broader point is that our Long-Run Neutrality simply differs from conventional notions. (This is why, for example, Backus et al. (2022) use both conventional neutrality and Long-Run Neutrality as distinct identification assumptions in their term structure model.)

Finally, it is possible to write examples where an FOMC announcement premium is consistent with Long-Run Neutrality. For instance, consider the heterogeneous-agent model of Kekre and Lenel (2022). The state variable X_t in this class of models captures the (stationary) wealth distribution across agents with different risk preferences.[14] Mon-

[13]Models with similar properties to our example economies include Kung (2015) and Bianchi et al. (2022a). In Kung (2015), monetary policy can affect R&D investment, hence growth rates. In Bianchi et al. (2022a), monetary policy occasionally modifies its reaction function, which induces stochastic volatility.

[14]For example, suppose there are two agent types, both with CRRA preferences. The SDF can be given by the marginal utility of the more risk-tolerant agent, i.e., $\frac{S_{t+1}}{S_t} = \beta\left(\frac{C_{t+1}}{C_t}\frac{X_{t+1}}{X_t}\right)^{-\gamma}$, where C is aggregate consumption, X is the consumption share of the more risk-tolerant agent, and γ is her risk aversion. With appropriate assumptions (e.g., OLG, taxation), the wealth share X is a stationary state variable.

etary policy causes redistribution, which impacts asset prices and risk premia. But because redistribution is fundamentally mean-reverting (i.e., X_t is stationary), there need not be any permanent impact on any agent's marginal utility. Thus, monetary policy can impact risk premia, perhaps substantially, even though Long-Run Neutrality holds. (Without heterogeneity, and with rational expectations, the existence of an FOMC announcement premium automatically induces a violation of Long-Run Neutrality (Ai and Bansal, 2018; Ai et al., 2023).)

4.4 Quantifying the bias in a model

So far, we have seen that Long-Run Neutrality may not hold empirically. But the fact that it fails literally may be of minuscule quantitative importance for extracting monetary shocks. In this section, we seek to *quantify* the size of the bias in identification, using a framework related to the long-run risks models above. The model features exogenous consumption growth with time-varying expected growth and time-varying growth variance. As before, the dynamics of these states are purely under the representative investor's subjective belief. More details on this model are in Appendix C.

The main result from this analysis is presented in Figure 5. Suppose you asked the investor her beliefs about the expected path of interest rates after a monetary shock. Her answer, which is the object we seek to identify and is plotted in solid blue, depends on what the monetary shock affects (growth versus variance) and the forecasting horizon. By contrast, we also plot what a researcher would extract by assuming the Expectation Hypothesis (dotted red) or our Long-Run Neutrality condition (dashed yellow).

Interestingly, if the monetary shock is a growth shock, assuming either the Expectation Hypothesis or Long-Run Neutrality would produce zero bias in this environment. The reason is related to the IID nature of shocks in this model, which is the knife-edge case that Proposition 4 omits. In that knife-edge case, what matters for identification is something weaker than Long-Run Neutrality: that a particular shock induces zero time-variation in "long-run risk prices".[15]

But the bias is substantial if monetary policy affects uncertainty. The top panel of Figure 5 shows that, while biases are negligible for shocks at horizons up to two years, extracting longer-horizon forward guidance is more problematic. One takeaway is that

[15] Here, "long-run risk prices" refers to the loading of $\log H_{t+1}/H_t$ on the shock ΔW_{t+1}. In the IID shock case here, since "long-run risk prices" are independent of the current growth state, there is no bias in simply assuming Long-Run Neutrality. This independence turns to hold given the linear structure of the growth state dynamics and the unitary EIS assumption on preferences. For a similar reason, namely that "short-run risk prices" are independent of the growth state, the Expectation Hypothesis assumption also produces zero bias.

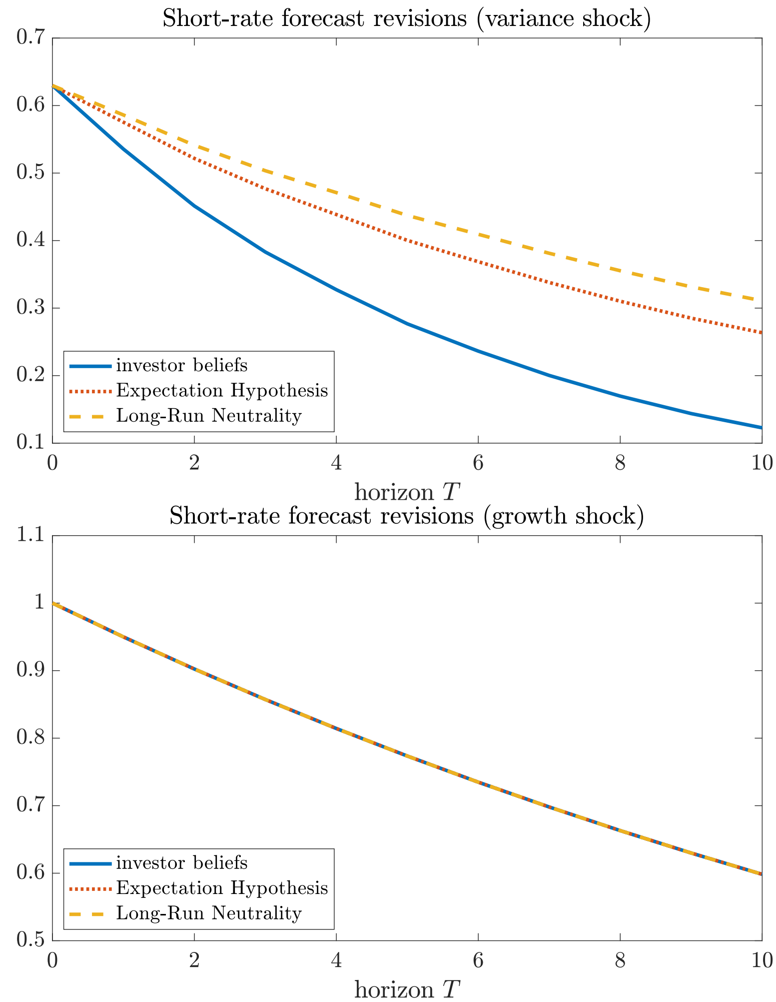

Figure 5. Short-rate forecast revisions. The figure presents forecast revisions regarding short-term interest rates, at various forecast horizons, in the calibrated long-run risk model. Forecasts are computed under the investor belief, under the risk-neutral measure (i.e., labeled "Expectation Hypothesis"), and under the long-run risk-neutral measure (i.e., labeled "Long-Run Neutrality").

identification procedures (e.g., assuming the Expectation Hypothesis or Long-Run Neutrality) will tend to *overstate* forward-guidance shocks. Another takeaway is that Long-Run Neutrality is actually worse as an identification assumption than the Expectation Hypothesis. The reason is simple: the model considered here features a prominent permanent component H that is actually even more volatile than the SDF S itself.

5 Final remarks

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 says this problem has no general solution, except in the degenerate case where the pricing kernel features no permanent shocks. But if a researcher more humbly seeks 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, evidence from monetary announcements 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.

Three questions come to mind. First, if monetary policy does not possess Long-Run Neutrality, how should researchers proceed? For instance, the strong impacts of monetary announcements on stock markets may not be evidence of significant long-term real impacts but rather an impact on long-term risk pricing. Writing and estimating structural models that feature a built-in violation of Long-Run Neutrality seems like a productive direction for future monetary research.

Second, how can researchers identify the effects of forward guidance and other such policy promises when asset prices do not? One idea is to combine such structural models with survey data on future interest rates, inflation, and the like. Analysis of survey evidence has been a fruitful 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. One must, however, overcome the low-frequency and short-horizon nature of existing surveys. A potentially helpful higher-frequency methodology is that proposed by Acosta (2023).

Third, one can generalize our model to many types of policies that either make promises or operate through beliefs about the future. To identify the effects of these interventions, one needs an analogous long-run neutrality condition, and this is testable by investigating the log excess return $\log(R^*_{t,t+\Delta}) - \log(R^\infty_{t,t+\Delta})$ at times of intervention.

References

Acosta, Miguel, "The perceived causes of monetary policy surprises," 2023. Unpublished working paper.

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

__ , __ , Hongye Guo, and Amir Yaron, "Identifying preference for early resolution from asset prices," 2023. NBER Working Paper.

Altavilla, Carlo, Luca Brugnolini, Refet S Gürkaynak, Roberto Motto, and Giuseppe Ragusa, "Measuring euro area monetary policy," *Journal of Monetary Economics*, 2019, *108*, 162–179.

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.

Andersen, Torben G and Luca Benzoni, "Do bonds span volatility risk in the US Treasury market? A specification test for affine term structure models," *The Journal of Finance*, 2010, *65* (2), 603–653.

Andrade, Philippe and Filippo Ferroni, "Delphic and odyssean monetary policy shocks: Evidence from the euro area," *Journal of Monetary Economics*, 2021, *117*, 816–832.

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.

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.

__ and Bruce N Lehmann, "Growth-optimal portfolio restrictions on asset pricing models," *Macroeconomic dynamics*, 1997, *1* (2), 333–354.

Barro, Robert J and David B Gordon, "Rules, discretion and reputation in a model of monetary policy," *Journal of Monetary Economics*, 1983, *12* (1), 101–121.

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.

Bhattarai, Saroj, Gauti B Eggertsson, and Bulat Gafarov, "Time consistency and duration of government debt: A model of quantitative easing," *The Review of Economic Studies*, 2023, *90* (4), 1759–1799.

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.

Boehm, Christoph E and T Niklas Kroner, "Monetary Policy without Moving Interest Rates: The Fed Non-Yield Shock," 2023. Unpublished working paper.

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

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

Cochrane, John H and Monika Piazzesi, "Bond risk premia," *American Economic Review*, 2005, *95* (1), 138–160.

Collin-Dufresne, Pierre and Robert S Goldstein, "Do bonds span the fixed income markets? Theory and evidence for unspanned stochastic volatility," *The Journal of Finance*, 2002, *57* (4), 1685–1730.

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

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.

Décaire, Paul H, Denis Sosyura, and Michael D Wittry, "Model Input Choice and Valuation Disagreement," 2023. Unpublished Working Paper.

Dessaint, Olivier, Jacques Olivier, Clemens A Otto, and David Thesmar, "CAPM-based company (mis) valuations," *The Review of Financial Studies*, 2021, *34* (1), 1–66.

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

__, "Information in (and not in) the term structure," *The Review of Financial Studies*, 2011, *24* (9), 2895–2934.

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.

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.

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

Gómez-Cram, Roberto and Marco Grotteria, "Real-time price discovery via verbal communication: Method and application to Fedspeak," *Journal of Financial Economics*, 2022, *143* (3), 993–1025.

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.

Gurkaynak, Refet S., Brian Sack, and Jonathan H. Wright, "The U.S. Treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, 2007, *54*, 2291–2304.

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.

Hansen, Stephen and Michael McMahon, "Shocking language: Understanding the macroeconomic effects of central bank communication," *Journal of International Economics*, 2016, *99*, S114–S133.

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.

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

Joslin, Scott, Marcel Priebsch, and Kenneth J Singleton, "Risk premiums in dynamic term structure models with unspanned macro risks," *The Journal of Finance*, 2014, *69* (3), 1197–1233.

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.

Kekre, Rohan and Moritz Lenel, "Monetary policy, redistribution, and risk premia," *Econometrica*, 2022, *90* (5), 2249–2282.

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.

Kung, Howard, "Macroeconomic linkages between monetary policy and the term structure of interest rates," *Journal of Financial Economics*, 2015, *115* (1), 42–57.

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.

Leombroni, Matteo, Andrea Vedolin, Gyuri Venter, and Paul Whelan, "Central bank communication and the yield curve," *Journal of Financial Economics*, 2021, *141* (3), 860–880.

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.

Lunsford, Kurt G, "Policy language and information effects in the early days of Federal Reserve forward guidance," *American Economic Review*, 2020, *110* (9), 2899–2934.

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.

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.

— **and** — , "Movements in Yields, not the Equity Premium: Bernanke-Kuttner Redux," 2024. 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.

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.

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, 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.

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 Forward Guidance

Anna Cieslak and Paymon Khorrami

March 20, 2025

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 (3) 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.[16]

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 factorization (4) 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 (8) 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 (12) implies the condition (14). □

[16] 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 (11), 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 $\sum_{k=0}^\infty \frac{t^k}{k!}A^k$.)

On the other hand, using assumption (14), we have that the drift of X_t under $\hat{\mathbb{P}}$ is $A_0 + AX_t - \hat{\beta}(X_t^\perp)$, where $X^\perp$ are the monetary-insensitive states. Compute similarly an equation for $\hat{m}_t^x := \hat{\mathbb{E}}^x[X_t]$:

$$\begin{aligned}\hat{m}_t^{X_\tau} - \hat{m}_t^{X_{\tau-}} &= A\int_0^t \Big(\hat{\mathbb{E}}^{X_\tau}[X_{\tau+s}] - \hat{\mathbb{E}}^{X_{\tau-}}[X_{\tau+s}]\Big)ds - \int_0^t \Big(\hat{\mathbb{E}}^{X_\tau}[\hat{\beta}(X_{\tau+s}^\perp)] - \hat{\mathbb{E}}^{X_{\tau-}}[\hat{\beta}(X_{\tau+s}^\perp)]\Big)ds \\ &= A\int_0^t \Big(\hat{m}_s^{X_\tau} - \hat{m}_s^{X_{\tau-}}\Big)ds,\end{aligned}$$

where the last line uses the monetary-insensitivity of $X^\perp$ and the fact that $\mathbb{P}$ and $\hat{\mathbb{P}}$ are mutually absolutely continuous. The solution to this integral equation is

$$\hat{m}_t^{X_\tau} - \hat{m}_t^{X_{\tau-}} = \exp(At)(X_\tau - X_{\tau-}).$$

Since the latter quantity is observable (by Lemma 1), we have that $m_t^{X_\tau} - m_t^{X_{\tau-}}$ is observable for all t. By assumption (13), 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 (11) and the square-root assumption on the diffusion (15), 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

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

Again, this solution holds for any time t prior to the next monetary surprise (i.e., on the event $\{t < \tau'\}$ as the proposition states). We may then compute

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

By Proposition 2, the object $m_s^{X_\tau} - m_s^{X_{\tau-}}$ is observable. In addition, examining the proof of Proposition 2, we actually have A observable as well. Hence, $V_T^{X_\tau} - V_T^{X_{\tau-}}$ is observable. This is enough, since by assumption (13), 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 (8). Indeed, Proposition 2 of Borovička et al. (2016) says that the observable asset prices can be obtained by formula (3) 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. Recall $r(x) = \rho_0 + \rho \cdot x$. Without loss of generality, let us assume that ρ has non-zeros in all of its entries. It must therefore generically be the case that

$$\mathbb{E}\big[H_T X_T \mid X_0 = x\big] = \mathbb{E}\big[X_T \mid X_0 = x\big] + \epsilon(T), \tag{A.8}$$

where $\epsilon(T)$ is an n-dimensional vector 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-}, including at least some states that are not monetary-insensitive. (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-}.) We may summarize both cases by writing the drift of X_t under $\hat{\mathbb{P}}$ as

$$\hat{\mu}(x) = \underbrace{A_0 + Ax}_{\mu(x)} - \Gamma(x), \tag{A.9}$$

where $\Gamma(x)$ is a non-constant function of x that depends non-trivially on at least some monetary-sensitive states. Using (A.9), compute

$$\begin{aligned} \mathbb{E}^x\left[H_T X_T\right] - \mathbb{E}^x\left[X_T\right] &= \hat{\mathbb{E}}^x\left[X_T\right] - \mathbb{E}^x\left[X_T\right] \\ &= A\int_0^T \left(\hat{\mathbb{E}}^x\left[X_t\right] - \mathbb{E}^x\left[X_t\right]\right)dt - \int_0^T \hat{\mathbb{E}}^x\left[\Gamma(X_t)\right]dt \end{aligned} \tag{A.10}$$

Guess that (A.8) holds, and replace in (A.10) to get

$$\epsilon(T) = A\int_0^T \epsilon(t)dt - \int_0^T \hat{\mathbb{E}}^x\left[\Gamma(X_t)\right]dt. \tag{A.11}$$

Clearly, the last term in (A.11) generically depends on x, contradicting the guess that $\epsilon(T)$ is independent of x. Thus, z_τ^T is not identified. □

B Motivating example: A Simple VAR

We begin with a simple example to illuminate the key issues. The point of this model is motivational: it allows us to set up our questions, explain the crux of the identification challenges, and then demonstrate conditions under which some types of identification are possible. In particular, we discuss why it is desirable to obtain investor forecast revisions about the interest rate path, as well as why obtaining these forecast revisions is highly non-trivial. Section 2 substantially generalizes the model and furthermore allows some types nonlinearities to permit the broadest statement of our identification results.

B.1 Model with two monetary actions and subjective beliefs

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 assume that g_t and r_t are observable, whereas f_t is a latent factor. Suppose the state evolves dynamically according to

$$X_{t+1} = A^o X_t + B\Delta W^o_{t+1}, \tag{B.1}$$

where $\Delta W^o_{t+1} \sim \text{Normal}(0, I)$ is a 3-dimensional vector of Normal shocks. For the purpose of this example, think of the shocks ΔW^o_{t+1} as containing a "real shock" and two "monetary shocks" that are completely governed by central bank actions—if B were a diagonal matrix, we could assign these labels to the individual shocks, but that is not necessary for our purposes. Equation (B.1) governs the *objective* state dynamics (hence the "o" superscript), which may differ from agents' *perceived* dynamics that we detail below.

Let us specify the persistence matrix A^o. Assume

$$A^o = \begin{bmatrix} a^o_{gg} & a^o_{gr} & a^o_{gf} \\ a^o_{rg} & a^o_{rr} & a^o_{rf} \\ 0 & 0 & a^o_{ff} \end{bmatrix}. \tag{B.2}$$

and that A^o is stable. Note that a^o_{gr} and a^o_{gf} capture the effects of short rates and forward guidance on growth. By contrast, a^o_{rg} captures the "feedback effect" of growth into the interest rate rule. And a^o_{rf} captures the transmission from forward guidance into the short-rate. In fact, the entire purpose of including f in this system is to add an additional factor governing future short rates. Finally, forward guidance f_t evolves as a univariate AR(1) independently of (g_t, r_t), a setup that is not necessary to any result but is transparent. We make no assumptions about B.

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

$$X_{t+1} = AX_t + B\Delta W_{t+1}, \quad \Delta W_{t+1} \sim \text{Normal}(0, I). \tag{B.3}$$

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

Thus, agents perceive

$$\Delta W_{t+1} = \Delta W^o_{t+1} - L_t, \quad \text{where} \quad L_t := B^{-1}\big(A - A^o\big) X_t, \tag{B.4}$$

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

B.2 The causal effects of monetary policy

Consider the question "what are the causal effects of monetary policy on future growth?" Here, there are two components of policy: short rates and forward guidance.

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

$$z^r_\tau := r_\tau - \mathbb{E}[r_\tau \mid X_{\tau-1}]. \tag{B.5}$$

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 $\mathbb{E}[r_\tau \mid X_{\tau-1}]$ at time $\tau - 1$. It is technically appropriate to use the investor expectation $\mathbb{E}$ here, because financial markets reflect investor beliefs. A *forward-guidance shock* at time τ is defined as

$$z^f_\tau := f_\tau - \mathbb{E}[f_\tau \mid X_{\tau-1}]. \tag{B.6}$$

Unlike the short-rate shock, we do not assume the forward-guidance shock is non-parameterically identified. This will be a key challenge. For completeness, also define the *growth shock* $z^g_\tau := g_\tau - \mathbb{E}[g_\tau \mid X_{\tau-1}]$. Stack the shocks into

$$z_\tau := \left(z^g_\tau, z^r_\tau, z^f_\tau\right)' = B\Delta W_\tau.$$

These shocks are reduced-form in nature, but that is not the key issue. The complication, instead, comes from the distorted beliefs: the reduced-form shocks

$$z_\tau = B\Delta W^o_\tau - (A - A^o)X_{\tau-1} \tag{B.7}$$

are related to the true structural shocks ΔW^o, with a bias that is a function of lagged X.

What are the causal effects of monetary policy on $g_{\tau+t}$? Because we do not have a full structural model of monetary action, we *define* all causal effects by the following IRFs

$$D_h := \frac{\partial}{\partial w'}\mathbb{E}^o\Big[g_{\tau+h} \mid B\Delta W^o_\tau = w\Big] = \left[\begin{smallmatrix}1\\0\\0\end{smallmatrix}\right] \cdot (A^o)^h \tag{B.8}$$

The second and third entries of this IRF tell us the effects of the short-rate and forward-guidance shocks. Suppose we seek D_h.

A first convenient result is that, to measure D_h empirically, we may run a regression of future growth $g_{\tau+h}$ onto z_τ, the lagged state $X_{\tau-1}$, and a constant. In other words, we can employ the *subjective shocks* and need not obtain proxies for the objective shocks; this simplification is beneficial because some shocks like z^r will be identified from financial market data and therefore reflect investors' beliefs. In principle, this "local projection" approach, highlighted by Jordà (2005), flexibly estimates the effect of shocks on future g at various horizons. The Jorda approach is considered to be more robust to model misspecification than the alternative that estimates the VAR(1) and computes $(A^o)^h$. For emphasis, we give this procedure a name:

Procedure 1. *The estimate of b_h from the regression*

$$g_{\tau+h} = a_h + b_h' z_\tau + c_h' X_{\tau-1} + \epsilon_{\tau+h} \tag{B.9}$$

is a consistent estimator of D_h.

Why does Procedure 1 work? An important, and underappreciated, feature is the lagged state control $X_{\tau-1}$: in a distorted-belief world, this control is necessary as it captures the discrepancy between the objective and subjective shocks—see equation (B.7). Approaches in the literature, instead, have mostly hoped to find objective shocks to monetary policy. If a researcher succeeds in measuring an objective shock, there is no need to control for drivers of belief distortions. But if instead one merely measures a subjective surprise, omitting the lagged states $X_{\tau-1}$ effectively, and perhaps wrongly, assumes rational expectations holds.

A simple numerical example highlights the importance of controlling for $X_{\tau-1}$. We calibrate A^o, A, and B to roughly match the evidence in Cieslak (2018), who compares the objective dynamics of employment growth and short rates to the subjective dynamics of these same variables from surveys. Details of this calibration are provided in Appendix B.4.1. The results for the calibrated A^o and A imply agents perceive a smaller feedback of growth to future short rates, as well as higher persistences of r and especially of f.

After calibrating, we ask: what is the impact of wrongly assuming that z_τ^r and z_τ^f are objective shocks, instead of perceived investor surprises? Assuming rational expectations corresponds to implementing Procedure 1 without controlling for the lagged state $X_{\tau-1}$. On the other hand, controlling for $X_{\tau-1}$ corrects for belief distortions and therefore yields the correct effect of forward guidance. Figure B.1 plots the outcomes of these two regressions. The lines labelled "True IRF" come from correctly implementing Procedure 1, while the lines labelled "Wrong IRF" come from forgetting to control for $X_{\tau-1}$ (and thereby implicitly assuming rational expectation). One notices a large discrepancy between these measures, so much that even the signs of the effect can be opposite in the two approaches.

As an aside, we note that the sign-flip from our numerical example may provide an alternative mechanism to generate what looks like a "Fed information effect." For example, some studies argue that monetary tightening can be associated with higher output, investment, or stock prices if the central bank reveals positive information through its policy decision (Romer and Romer, 2000; Melosi, 2017; Nakamura and Steinsson, 2018; Miranda-Agrippino and Ricco, 2021). Through the lens of our framework, the positive response to monetary tightening is evidence of misspecification. The true structural effect of monetary tightening re-emerges with the proper controls for investors' belief distortion.

To summarize, the obstacles to implementing Procedure 1 are twofold. First, as mentioned, we need to control for $X_{\tau-1}$ in our estimation. But since X includes the unobservable forward-guidance variable f, we will need a method to recover f. What is this method? Second, we require the shocks z_τ. We will continue to assume that the short-rate shock z_τ^r is observed from financial markets. But how can we recover z_τ^f? Given the forward-looking nature of financial markets, let us broach the possibility that asset prices can inform us about both of f and z^f.[17]

[17] We do not discuss the growth shock z^g further because its inclusion is already emphasized by the existing literature on the "information effect" of monetary policy. In particular, one should include z^g in Procedure 1, because it may be correlated to monetary surprises z^r and z^f (and so omission of z^g results in an omitted variable bias). For example, if the central bank responds to high-frequency growth information, one needs to control for this growth information to get an appropriate monetary policy instrument (Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023a,b; Andrade and Ferroni, 2021).

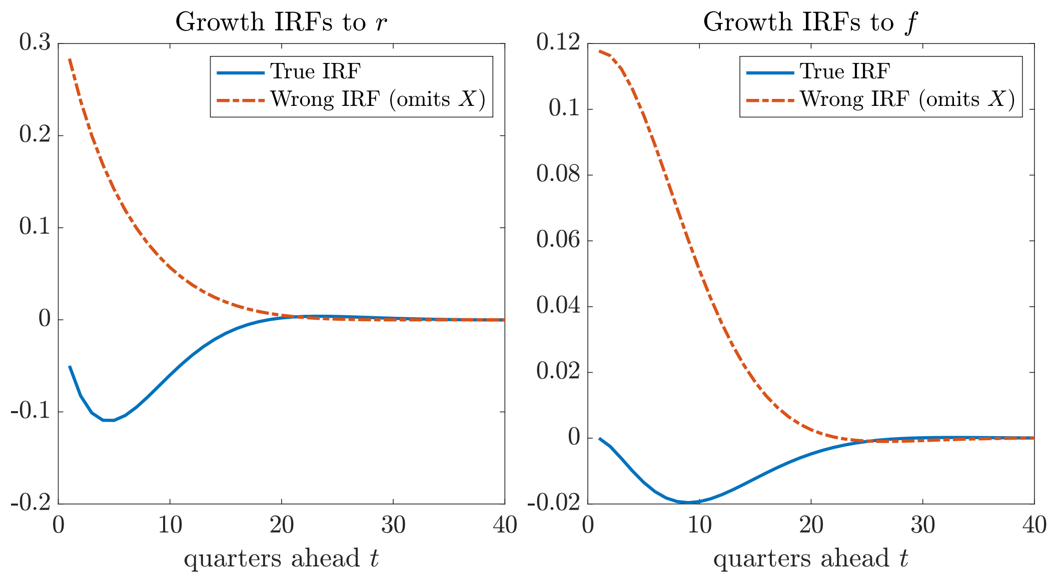

Figure B.1. Growth IRFs to forward guidance. Regressions of $g_{\tau+t}$ on the forward-guidance shock z_{τ}^{f} with and without the lagged state $X_{\tau-1}$ as a control. The solid curve controls for $X_{\tau-1}$ and coincides with the "True IRF." The dashed curve omits $X_{\tau-1}$ which yields the "Wrong IRF" and corresponds to assuming rational expectations. The calibrations of A^{o} and A are given in Appendix B.4.1, equation (B.18). The calibration of B is a diagonal matrix with $\sigma_g = 0.0025/4$, $\sigma_r = 0.010/4$, and $\sigma_f = 0.050/4$.

B.3 Identifying forward guidance from asset prices

We 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.

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.

Pricing model. 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 dynamics in Section B.1, we add the following one-period stochastic discount factor (SDF):

$$\frac{S_{t+1}^{o}}{S_{t}^{o}} = \exp\Big[-(\bar{r} + r_t) - \frac{1}{2}\pi_t'\pi_t - \pi_t'\Delta W_{t+1}^{o} \Big], \tag{B.10}$$

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:

$$\frac{S_{t+1}}{S_{t}} = \exp\Big[-(\bar{r} + r_t) - \frac{1}{2}(\pi_t + L_t)'(\pi_t + L_t) - (\pi_t + L_t)'\Delta W_{t+1} \Big]. \tag{B.11}$$

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

In this setting, the risk-neutral dynamics of the state are given by

$$\begin{aligned} X_{t+1} &= A_0^* + A^* X_t + B\Delta W^*_{t+1}, \\ \text{where} \quad A_0^* &:= -B\pi_0 \\ \text{and} \quad A^* &:= A^o - B\Pi, \end{aligned} \tag{B.12}$$

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

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

This solution is standard in the literature.

The challenge: extracting forward guidance. Can we use the model solution to identify f? 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 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.

Unfortunately, this identification logic is incorrect in general, as Appendix B.4.2 explains in more detail. Bond yields themselves do not reveal *particular latent factors*. This is a well-known identification issue in affine term-structure models with latent state variables like our forward guidance variable (Hamilton and Wu, 2012). Without specific knowledge of A, B, or Π, we cannot decide whether the underlying state vector is X_t or some rotation $\hat{X}_t$ of the state, as X and $\hat{X}$ lead to the same pricing implications. This issue goes beyond our particular example. Indeed, the example above represents the "best-case scenario" where we know the true model is a three-factor model with two of the factors observable. Even then, the sole latent factor cannot be identified. The problem intensifies if the environment is more complex, e.g., if there are additional latent factors driving bond yields or bond risk premia. Therefore, bringing in bond yields generically does not help us identify the causal effects of forward guidance. (As Appendix B.4.2 discusses, only in a knife-edge case—where forward-guidance is the unique latent variable and where its correlations with all observable states are known—can f be identified from bond yields.)

Surprises to future short rates: A resolution? Given the impossibility of identifying f, hence z^f, let us now ask a humbler question: can our asset-pricing model help us identify belief updates about future short rates? That is, can we obtain

$$z_\tau^{r,t} := \mathbb{E}[r_{\tau+t} \mid X_\tau] - \mathbb{E}[r_{\tau+t} \mid X_{\tau-1}] = (0,1,0)A^t B\Delta W_\tau \tag{B.13}$$

from asset-price data?

Obtaining surprises about future rates is actually good enough for many purposes, for the following reason. Conditional on $X_{\tau-1}$, the correlation between $z_\tau^{r,0}$ and $z_\tau^{r,t}$ is imperfect:

$$\operatorname{corr}\left[z_\tau^{r,t}, z_\tau^{r,0} \mid X_{\tau-1}\right] = \frac{[A^t BB']_{2,2}}{\sqrt{[A^t BB' A^t]_{2,2}[BB']_{2,2}}}$$

Unless A is a diagonal matrix, which is the uninteresting case where monetary policy has no effects, this correlation will be below one for $t > 0$. Therefore, replacing z_τ^f in Procedure 1 with $z_\tau^{r,t}$ (for some $t > 0$) allows us to recover all desired objects: the coefficients on z_τ^r and $z_\tau^{r,t}$ will be the short-rate and forward-guidance effects, respectively. Intuitively, there is a two-factor structure to the short-rate path, and so z_τ^r and $z_\tau^{r,t}$ pick up different factors. In a richer model with additional latent factors driving short rates, we may want to include an entire collection $(z_\tau^{r,t})_{t=1}^T$ of surprises. The fact that these reduced-form surprises to future rates are "good enough" focuses our attention in the remainder of the paper on trying to extract the $z_\tau^{r,t}$ surprises, rather than forward guidance per se. For emphasis, we collect this discussion in the following procedure:

Procedure 2. *The estimate of* $(b_h^g, b_h^r, b_h^{r,t})$ *from the regression*

$$g_{\tau+h} = a_h + b_h^g z_\tau^g + b_h^r z_\tau^r + b_h^{r,t} z_\tau^{r,t} + c_h' X_{\tau-1} + \epsilon_{\tau+h} \tag{B.14}$$

is a consistent estimator of D_h.

To implement Procedure 2, we need a proxy for $z_\tau^{r,t}$. If investors are *risk-neutral*, we can recover these surprises. 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. In that case, we can recover the surprises in (B.13) via $z_\tau^{r,t} = \mathbb{E}^*[r_{\tau+t} \mid X_\tau] - \mathbb{E}^*[r_{\tau+t} \mid X_{\tau-1}]$, given estimates of A^* and B. Intuitively, risk-neutrality allows recovery because it implies an "Expectations Hypothesis" but under the investor subjective belief.

A relaxed version of risk-neutrality also permits shock recovery. Assume

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

Condition (B.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 it is both necessary and sufficient for identification of $B\Delta W$ in this model.

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

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

If we know A^*, then we know A, which allows us to obtain perceived shocks as $X_{t+1} - AX_t = B\Delta W_{t+1}$. 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. The constant shifter is differenced out when studying belief surprises, rather than belief levels.

Necessity is harder to see. The important fact is that condition (B.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 S in (B.11) as

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

where $v := -[I - (A - B\Pi)']^{-1}(0,1,0)'$. In our general environment of Section 2, 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 $\pi_t + L_t = \pi_0 + [\Pi + B^{-1}(A - A^o)]X_t$ to be independent of X_t, which translates to condition (B.15).

As it turns out, this example sneakily permits looser identification conditions than the more general case covered in Section 2. In particular, we have assumed a very strong structure where *monetary shocks are iid*: they happen every period with the same time-invariant distribution. As our more general model in Section 2 shows, 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*.

Remark 3 (Risk-neutral shocks)**.** *One limitation of our simple three-factor VAR is that, taking the setup literally, one could get away with using* risk-neutral surprises *rather than investor's perceived surprises. In particular, suppose we obtain $z_\tau^{r,t,*} := \mathbb{E}^*[r_{\tau+t} \mid X_\tau] - \mathbb{E}^*[r_{\tau+t} \mid X_{\tau-1}]$ from bond markets and run a modified version of Procedure 2 with $z_\tau^{r,t,*}$ in place of $z_\tau^{r,t}$. If the true model has (g, r, f) as the factors, then the modified procedure works—i.e., we recover both the short-rate and forward-guidance effects—with two possible caveats.*

First, while seemingly mundane, it is absolutely critical to control for $X_{\tau-1}$ when running the variant of Procedure 2 with risk-neutral surprises. Controlling for $X_{\tau-1}$ soaks up the wedge between risk-neutral, subjective, and objective beliefs, so one can recover the correct structural effects. The literature often overlooks controls for drivers of belief distortions and risk adjustments; for example, Gürkaynak et al. (2005b) and Swanson (2021), by decomposing the effects of monetary policy on various asset markets, effectively employ risk-neutral surprises, but they do not control for $X_{\tau-1}$. Second, the term structure of investor forecast revisions $(z_\tau^{r,t})_{t>0}$ are closer to structural objects of interest than their risk-neutral counterparts $(z_\tau^{r,t,*})_{t>0}$. If so, measuring $(z_\tau^{r,t})_{t>0}$ is preferable. Within this second caveat, we see two specific issues: risk premia effects and multi-dimensional policy.

Suppose forward guidance has risk premium effects in addition to its traditional impact on future short rates (Bauer et al., 2023). For instance, signals of higher future rates may lower growth but also reduce risks (like inflation risks in a richer model), with an ambiguous combined impact on long-term bond yields. The risk-neutral surprises $z_\tau^{r,t,*}$ always contain this joint effect; a regression of economic outcomes onto $z_\tau^{r,t,*}$ cannot isolate the pure impact of foreward guidance. Depending on the size of risk premia effects, $z_\tau^{r,t,*}$ may even take an opposite sign to $z_\tau^{r,t}$.[18]

[18] To take an extreme but enlightening example where risk premia complicate inference, suppose A^* is a diagonal matrix. Bond yields $y_t^{(n)}$ will not reflect forward guidance in this world: the loading $\mathcal{B}^{(n)}$ will have zeros everywhere except for the short-rate entry. In fact, one can verify that $\text{corr}[z_\tau^{r,t,*}, z_\tau^{r,0,*} \mid X_{\tau-1}] = 1$ when A^* is diagonal, so that $z_\tau^{r,t,*}$ provides no additional information relative to the short-rate shock. This offsetting issue is related to the discussion surrounding "unspanned factors" in the term structure literature (e.g., yields may not contain all information about the SDF as in Collin-Dufresne and Goldstein, 2002; Cochrane and Piazzesi, 2005; Andersen and Benzoni, 2010; Duffee, 2011; Joslin et al., 2014).

In a world of multi-dimensional policy, risk-neutral surprises struggle to correctly attribute the relative impacts of various interventions. To see this, consider that in a general K-factor VAR, the risk-neutral and investor forecast revisions differ by

$$z_\tau^{r,t,*} - z_\tau^{r,t} = (0,1,0,\ldots,0)\Big\{\big[(A^*)^t - A^t\big]B\Delta\tilde{W}_\tau + (A^*)^t B\big[\pi_0 + \big[\Pi + B^{-1}(A - A^o)\big]X_{\tau-1}\big]\Big\}.$$

The second term involving $X_{\tau-1}$ can be managed, assuming one controls for the lagged state as in Procedure 2. But the first term involves the perceived shocks ΔW_τ and cannot be controlled. To the extent that A^* differs from A^o, the risk-neutral and investor surprises can load very differently on the underlying shocks. In this world, the risk-neutral surprises cannot help answer questions like: at which horizon is forward guidance most effective? how does uncertainty management compare to traditional forward guidance? how does quantitative easing compare to forward guidance?

B.4 Additional details

B.4.1 Calibration of 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 misspecification wrongly assuming rational expectations impacts the estimated IRF of growth to a forward-guidance shock. The result is Figure B.1 in the main text.

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

$$r_{t+j} = \beta_0^{(j)} + \beta_g^{(j)} g_t + \beta_r^{(j)} r_t + \text{residual} \tag{B.16}$$

$$\mathbb{E}_t[r_{t+j}] = \beta_0^{(j)} + \beta_g^{(j)} g_t + \beta_r^{(j)} r_t + \text{residual} \tag{B.17}$$

where $\mathbb{E}_t[r_{t+j}]$ denotes subjective expectations of interest rates, obtained from the Blue Chip Financial Forecasts. The results of these regressions, in the data, are presented in the bottom row of Table B.1.

	A. Dependent variable r_{t+j}				B. Dependent variable $\mathbb{E}_t[r_{t+j}]$			
	$j = 1$ quarter		$j = 4$ quarters		$j = 1$ quarter		$j = 4$ quarters	
	$\beta_g^{(1)}$	$\beta_r^{(1)}$	$\beta_g^{(4)}$	$\beta_r^{(4)}$	$\beta_g^{(1)}$	$\beta_r^{(1)}$	$\beta_g^{(4)}$	$\beta_r^{(4)}$
Model:	0.195	0.902	0.547	0.493	0.145	0.952	0.215	0.746
Data:	0.180	0.880	0.650	0.500	0.100	0.930	0.120	0.840

Table B.1. Short-rate forecasts in the model and data. 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). The calibration of B is a diagonal matrix with $\sigma_g = 0.0025/4$, $\sigma_r = 0.010/4$, and $\sigma_f = 0.050/4$. 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.

We set A^o and A in the model to roughly match these empirical results. We calibrate

$$A^o = \begin{bmatrix} 0.90 & -0.05 & 0 \\ 0.50 & 0.75 & 0.05 \\ 0 & 0 & 0.80 \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} 0.90 & -0.05 & 0 \\ 0.45 & 0.80 & 0.05 \\ 0 & 0 & 0.99 \end{bmatrix}. \tag{B.18}$$

The comparison between our model and the data is presented in Table B.1. The fit, while not perfect, is good enough to illustrate our main points here. Relative to the econometric transition matrix, agents perceive a smaller feedback of growth to future short rates, as well as higher persistences of r and especially of f. The true and perceived dynamics of short rates in the model, plotted in Figure B.2, illustrate the discrepancy between A^o and A, especially the greater perceived persistence of the forward-guidance shock.

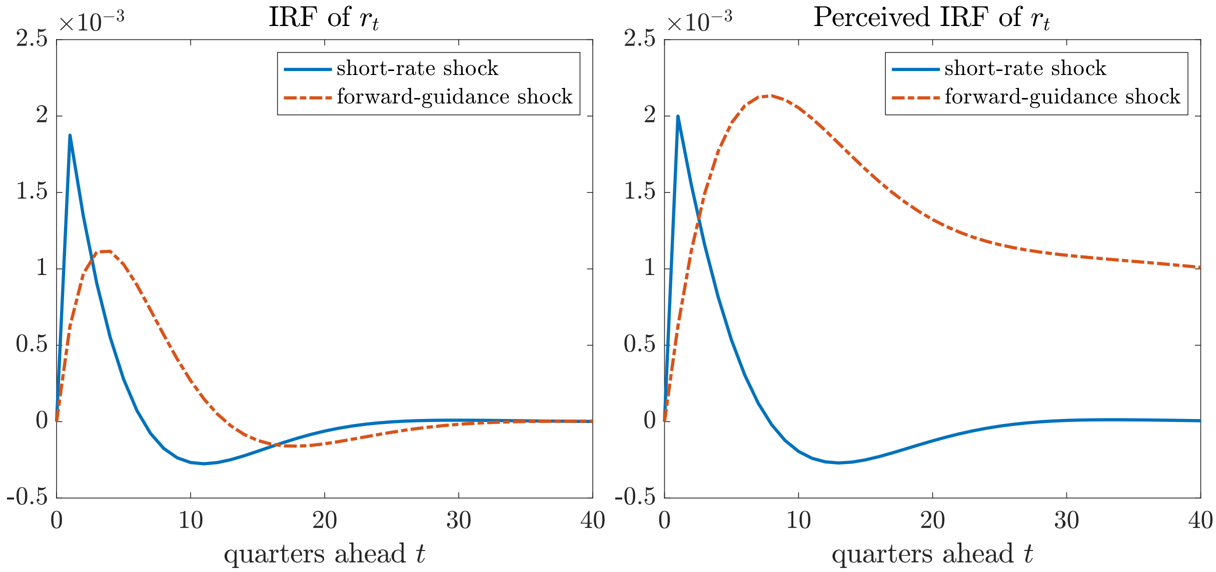

Figure B.2. IRFs of short rates to the two monetary shocks. The left panel plots the true IRFs. The right panel plots the perceived IRFs. The calibrations of A^o and A are given in (B.18). The calibration of B is a diagonal matrix with $\sigma_g = 0.0025/4$, $\sigma_r = 0.010/4$, and $\sigma_f = 0.050/4$.

B.4.2 Identifying latent factors from term structure models

We consider the three-factor VAR (B.1), paired with the SDF (B.10). In this model, the forward guidance variable f is latent. One would hope that asset prices, paired with the pricing model, allow us to "invert bond yields" for the latent factor, but this is not the case as we explain here.

For example, consider using the transformed state variable $\hat{X}_t = UX_t$, where

$$U = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ u_g & u_r & 1 \end{bmatrix}.$$

While the subsequent argument works for any U, this particular example acknowledges that (g, r) are observable and lets the latent factor be a linear combination of forward guidance with the observables. At the same time, suppose the risk price vector is $\hat{\pi}_t = \hat{\pi}_0 + \hat{\Pi}\hat{X}_t$, where $\hat{\pi}_0 = \pi_0$ and $\hat{\Pi} = \Pi U^{-1}$. Then, the SDF is identical to the original specification (i.e., $\hat{\pi}_t = \pi_t$) and the risk-neutral dynamics of $X_t = U^{-1}\hat{X}_t$ are identical to the original specification.[19]

[19]Indeed, the physical dynamics of $\hat{X} := UX$ are $\hat{X}_{t+1} = (UAU^{-1})\hat{X}_t + (UB)\Delta W^o_{t+1}$, so the risk-neutral dynamics for $\hat{X}$ are

$$\hat{X}_{t+1} = -UB\hat{\pi}_0 + (UAU^{-1} - UB\hat{\Pi})\hat{X}_t + (UB)\Delta\hat{W}^*_{t+1}$$

To address the generic non-identification of latent factors, we must effectively make some assumption about U. A canonical approach starting with Hamilton and Wu (2012) assumes the state vector $\hat{X}$ is chosen such that its risk-neutral persistence is a diagonal matrix. This can be done by implicitly picking U^{-1} as the matrix of eigenvectors of $A^* = A^o - B\Pi$, since UA^*U^{-1} is the risk-neutral persistence of $\hat{X}$. But after making this choice, only $\hat{X}_t = UX_t$ is recovered from bond yields, not X_t itself. Without additional assumptions, bond yields do not reveal *particular latent factors*.

What can be done in light of the challenges to observing f? On the one hand, we are still able to extract $\hat{X}$ from bond yields, a state vector which spans the same space as X. This is good enough for the purpose of controlling for the past state: controlling for $\hat{X}_{\tau-1} = UX_{\tau-1}$ in Procedure 1 instead of $X_{\tau-1}$ will lead to identical inference for b_h. Thus, we may still recover the desired IRF in principle. On the other hand, failure to specifically recover f prevents us from constructing its shock z^f, which prevents us from actually implementing Procedure 1 as stated. We must necessarily exclude the forward-guidance shock z^f_τ and so cannot measure any forward-guidance effect, at least not in this way.[20]

Although it is not our focus, excluding z^f_τ may also bias estimates of the causal impact of short-rate shocks. Indeed, note that $\text{Cov}[z^r_\tau, z^f_\tau \mid X_{\tau-1}] = (0,1,0)BB'(0,0,1)'$. So if we run Procedure 1 without z^f_τ, any correlation between short-rate and forward-guidance shocks will be impounded into the coefficient on z^r_τ.

What solves some issues in the simple baseline model is an assumption that forward-guidance shocks are orthogonal to short-rate and growth shocks. In other words, if we assume that the third row of B is proportional to $(0,0,1)$, then we can recover f from bond yields. This is because the unique matrix U that preserves this orthogonality is $U = I$, and so $\hat{X} = X$ is uniquely pinned down by yields. However, even armed with this orthogonality assumption, we re-encounter difficulties as soon as there are additional latent factors present. Examples of such problematic environments would be if forward guidance was described by a multi-factor structure, or if monetary policy managed either uncertainty or risk premia in addition to future short rates. In such case, f is once again unidentified without additional assumptions.

---

where $\Delta\hat{W}^*_{t+1} := \Delta W^o_{t+1} + \hat{\pi}_t$ is the risk-neutral shock. If $\hat{\pi}_0 = \pi_0$ and $\hat{\Pi} = \Pi U^{-1}$, then risk prices are

$$\hat{\pi}_t = \hat{\pi}_0 + \hat{\Pi}\hat{X}_t = \pi_0 + \Pi U^{-1} U X_t = \pi_t.$$

As a result, the risk-neutral shocks coincide: $\Delta\hat{W}^*_{t+1} = \Delta W^*_{t+1}$. Thus, the risk-neutral dynamics of $\hat{X}$ are

$$\hat{X}_{t+1} = UA^*_0 + UA^*U^{-1}\hat{X}_t + (UB)\Delta W^*_{t+1}$$

These dynamics imply the same risk-neutral dynamics for X_t displayed in (B.12). One can also easily show that, as a consequence, the solution for equilibrium bond yields is invariant to the choice of u_g and u_r.

[20]A similar critique applies to the VAR approach in place of the Jorda projection. If we fully observed X, then we would not need to run Jorda projections; we could estimate the VAR directly via time series regressions, obtain A^o, and then construct the desired IRF $(A^o)^h$. But if we obtain only $\hat{X} = UX$, we would not recover A and would therefore obtain an incorrect IRF.

C Long-run risks model

We write and solve a standard long-run risks model as a laboratory to evaluate various procedures for extracting monetary policy's impacts on future interest rates. This model is a generalized version of the examples in Section 4.

Aggregate consumption follows

$$\log C_{t+1} = \log C_t + g_t + \sqrt{v_t}\Delta W_{t+1}^{(1)},$$

where the expected growth rate g_t and growth variance v_t follow

$$\begin{aligned} g_{t+1} &= (1-a_g)\mu_g + a_g g_t + \sigma_g\sqrt{v_t}\Delta W_{t+1}^{(2)} \\ v_{t+1} &= (1-a_v)\mu_v + a_v v_t + \sigma_v\sqrt{v_t}\Delta W_{t+1}^{(3)}. \end{aligned}$$

The three shocks $\Delta W := (\Delta W^{(1)}, \Delta W^{(2)}, \Delta W^{(3)})' \sim N(0, I)$. In this discrete-time model, it is technically possible that v becomes negative in simulations. However, because the model solution is identical between discrete and continuous time, we will address this problem by simulated using a continuous-time version of the above dynamics.[21] We may stack the stationary state variables as $X := (g, v)$ and write their dynamics in vector form as

$$\begin{aligned} X_{t+1} &= (I-A)\mu + AX_t + \sqrt{X_t^{(2)}}\Sigma\Delta W_{t+1} \\ \mu &:= \begin{bmatrix} \mu_g \\ \mu_v \end{bmatrix} \\ A &:= \begin{bmatrix} a_g & 0 \\ 0 & a_v \end{bmatrix} \\ \Sigma &:= \begin{bmatrix} 0 & \sigma_g & 0 \\ 0 & 0 & \sigma_v \end{bmatrix} \end{aligned}$$

The fact that all conditional volatilities are proportional to $\sqrt{v_t}$ will make the model analytically tractable.

The value function recursion (Bellman equation) for the representative agent with unitary EIS is given by

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

Guess that the value function is given by $U_t = \log C_t + u_0 + u \cdot X_t$ for some scalar u_0 and vector $u := (u_g, u_v)$. Substituting this guess into the Bellman equation gives the following three

[21] The continuous-time versions of these processes are

$$\begin{aligned} d\log C_t &= g_t dt + \sqrt{v_t}dW_t^{(1)} \\ dg_t &= -\lambda_g(g_t - \mu_g)dt + \sigma_g\sqrt{v_t}dW_t^{(2)} \\ dv_t &= -\lambda_v(v_t - \mu_v)dt + \sigma_v\sqrt{v_t}dW_t^{(3)}, \end{aligned}$$

where $W := (W^{(1)}, W^{(2)}, W^{(3)})$ is a Brownian motion. We will have to calibrate λ_g and λ_v here, in order to match the same rate of mean-reversion from the discrete-time model. This is done by noting that $\exp(-\lambda_g) = a_g$ and $\exp(-\lambda_v) = a_v$ are the annual persistences.

equations

$$\begin{aligned} u_0 &= \frac{\beta u'(I-A)\mu}{1-\beta} \\ u_g &= \frac{\beta}{1-\beta a_g} \\ 0 &= (\gamma-1)\sigma_v^2 u_v^2 + 2[\beta^{-1} - a_v]u_v + (\gamma-1)[1+\sigma_g^2 u_g^2] \end{aligned}$$

As explained in Hansen and Scheinkman (2009), the appropriate solution for u_v is the larger root of the above quadratic equation. This solution will be negative for $\gamma > 1$.

In this model, the SDF is

$$\frac{S_{t+1}}{S_t} = \beta \frac{C_t}{C_{t+1}} \frac{\exp[(1-\gamma)U_{t+1}]}{\mathbb{E}_t[\exp[(1-\gamma)U_{t+1}]]}.$$

Substituting the consumption dynamics and value function, we ultimately obtain

$$\frac{S_{t+1}}{S_t} = \beta \exp\Big\{ -(g_t + \frac{1}{2}v_t) + \gamma v_t - \pi_t \cdot \Delta W_{t+1} - \frac{1}{2}\|\pi_t\|^2 \Big\},$$

where the price of risk π_t is given by

$$\pi_t = \sqrt{v_t}\bar{\pi}, \quad \text{where} \quad \bar{\pi} := \begin{pmatrix} \gamma \\ (\gamma-1)\sigma_g u_g \\ (\gamma-1)\sigma_v u_v \end{pmatrix}.$$

Note that the log riskless rate $r_t := -\log(\mathbb{E}_t[S_{t+1}/S_t])$ is given by

$$r_t = -\log(\beta) + g_t + \frac{1}{2}v_t - \gamma v_t.$$

As we are interested in shocks that impact current and future interest rates, this models offers up $\Delta W^{(2)}$ and $\Delta W^{(3)}$ as candidate shocks.

Just as a brief detour, let us observe that this model contains separate shocks to current and future interest rates. We can write

$$r_{t+1} = r_t - (1-a_g)(g_t - \mu_g) + (\gamma - \frac{1}{2})(1-a_v)(v_t - \mu_v) + \sqrt{v_t}\sigma_r \Delta W^r_{t+1},$$

where

$$\begin{aligned} \Delta W^r_{t+1} &:= \frac{\sigma_g}{\sigma_r}\Delta W^{(2)}_{t+1} - (\gamma - \frac{1}{2})\frac{\sigma_v}{\sigma_r}\Delta W^{(3)}_{t+1} \\ \sigma_r &:= \sqrt{\sigma_g^2 + (\gamma - \frac{1}{2})^2 \sigma_v^2} \end{aligned}$$

Thus, ΔW^r captures short-rate shocks, whereas the shock

$$\begin{aligned} \Delta W^f_{t+1} &:= \alpha \Delta W^{(2)}_{t+1} + \sqrt{1-\alpha^2}\Delta W^{(3)}_{t+1} \\ \alpha &:= (\gamma - \frac{1}{2})\frac{\sigma_v}{\sigma_r} \end{aligned}$$

is an orthogonal shock (i.e., $\Delta W^r \perp \Delta W^f$) that captures variation in *future* interest rates, holding fixed current rates. This latter shock is the one that allows us to conceptualize shocks to beliefs about future rates, as in the case of forward guidance.

Consider the term structure of interest rates. The price of an n-year discount bond is given by

$$\begin{aligned} p_t^{(n)} &= \mathbb{E}_t\Big[\frac{S_{t+1}}{S_t} p_{t+1}^{(n-1)}\Big] \\ &= \beta \mathbb{E}_t\Big[\exp\Big\{ -(g_t + \frac{1}{2}v_t) + \gamma v_t - \pi_t \cdot \Delta W_{t+1} - \frac{1}{2}\|\pi_t\|^2 \Big\} p_{t+1}^{(n-1)}\Big]. \end{aligned}$$

Guess that $\log(p_t^{(n)}) = -\mathcal{B}_0^{(n)} - \mathcal{B}^{(n)} X_t$ for all t and n. Substituting this guess, we obtain the recursions

$$\begin{aligned} \mathcal{B}_0^{(n)} &= -\log(\beta) + \mathcal{B}_0^{(n-1)} + \mathcal{B}^{(n-1)} \cdot (I - A)\mu \\ \mathcal{B}^{(n)} &= \begin{bmatrix} 1 \\ \frac{1}{2} - \gamma \end{bmatrix} + \mathcal{B}^{(n-1)} - \big(\bar{\pi}'\Sigma'\mathcal{B}^{(n-1)}\big) \begin{bmatrix} 0 \\ 1 \end{bmatrix} - \frac{1}{2}(\mathcal{B}^{(n-1)})'\Sigma\Sigma'\mathcal{B}^{(n-1)} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \end{aligned}$$

Since $p_t^{(0)} = 1$, these recursions have the initial conditions $\mathcal{B}_0^{(0)} = 0$ and $\mathcal{B}^{(0)} = (0,0)'$ from which we can solve the entire term structure of bond prices. Bond yields are then given by $y_t^{(n)} = -\frac{1}{n}\log(p_t^{(n)})$.

Next, consider the permanent-transitory factorization of the SDF. Guess that the transitory piece is given by $e(x) = \exp(\kappa \cdot x)$ for some κ. We have

$$\begin{aligned} \frac{H_{t+1}}{H_t} &= \exp(-\eta)\frac{e(X_{t+1})}{e(X_t)}\frac{S_{t+1}}{S_t} \\ &= \exp\Big\{ -\eta - \log(\beta) - (g_t + \frac{1}{2}v_t) + \gamma v_t - \sqrt{v_t}(\bar{\pi} - \Sigma'\kappa)\cdot \Delta W_{t+1} - \frac{1}{2}v_t\|\bar{\pi}\|^2 + \kappa'(I - A)(\mu - X_t) \Big\} \end{aligned}$$

For this object to be a martingale, we require $\kappa := (\kappa_g, \kappa_v)'$ and η to satisfy

$$\begin{aligned} \eta &= -\log(\beta) + \kappa'(I - A)\mu \\ -1 &= (1 - a_g)\kappa_g \\ 0 &= \frac{1}{2}\sigma_v^2\kappa_v^2 - [1 - a_v + \sigma_v\bar{\pi}_v - \sigma_g\sigma_v\kappa_g]\kappa_v + \frac{1}{2}\sigma_g^2\kappa_g^2 - \sigma_g\bar{\pi}_g\kappa_g + \gamma - \frac{1}{2} \end{aligned}$$

where we have defined $\bar{\pi}_g := (\gamma - 1)\sigma_g u_g$ and $\bar{\pi}_v := (\gamma - 1)\sigma_v u_v$. For the latter quadratic equation in κ_v, we pick the root that creates stochastic stability of the state vector X_t under the probability measure $\hat{\mathbb{P}}$ induced by the resulting H. The resulting H is given by

$$\frac{H_{t+1}}{H_t} = \exp\Big\{ -\frac{1}{2}v_t\|\bar{\pi} - \Sigma'\kappa\|^2 - \sqrt{v_t}(\bar{\pi} - \Sigma'\kappa)\cdot\Delta W_{t+1}\Big\}.$$

To determine which of the two roots to select for κ_v, note that, in this conditionally log-normal environment, the dynamics of X under $\hat{\mathbb{P}}$ are given by

$$X_{t+1} = (I - \hat{A})\hat{\mu} + \hat{A}X_t + \sqrt{X_t^{(2)}}\Sigma\Delta\hat{W}_{t+1}$$

where $\Delta\hat{W}_{t+1} \sim N(0, I)$ under $\hat{\mathbb{P}}$, and the new persistent matrix $\hat{A}$ and long-run mean $\hat{\mu}$ are given by

$$\hat{A} := \begin{bmatrix} a_g & \sigma_g^2\kappa_g - \sigma_g\bar{\pi}_g \\ 0 & a_v + \sigma_v^2\kappa_v - \sigma_v\bar{\pi}_v \end{bmatrix}$$
$$\hat{\mu} := (I - \hat{A})^{-1}(I - A)\mu$$

To check whether X is stable under $\hat{\mathbb{P}}$ amounts to checking the stability of the matrix $\hat{A}$ under each of the two choices for κ_v.

Now we are ready to define the two different approaches to measuring belief surprises. First, we may assume risk neutrality and extract beliefs via the Expectations Hypothesis. Second, we may assume Long-Run Neutrality and extract beliefs via the "recovery" procedure described in Lemma 1. We will compare these two approaches to the correct belief surprises from the calibrated model, assuming monetary policy impacts both growth and uncertainty to various degrees.