

What is a Monetary Policy Shock?

Paul Bousquet*

July 31, 2026

preliminary

[\[Click here for latest version\]](#)

Abstract

This paper explores a new microfoundation for monetary policy shocks: randomness in reaching a consensus. I argue high-frequency movements in interest rate futures around monetary policy announcements plausibly identify these shocks, but show the measurement can be contaminated even absent "information effects" or risk premia. Empirical tests show suggestive evidence of this new form of contamination. Because sizable shocks are rare, there are unavoidable difficulties for these high-frequency changes to be useful for estimating the effects of monetary policy using low-frequency macroeconomic aggregates, even absent measurement issues.

*University of Virginia, pbousquet@virginia.edu. This draft will be updated regularly; the Appendix is to be completed.

1 Introduction

The transmission of monetary policy is typically characterized with an impulse response function, the response of the macroeconomy to a random deviation from a fixed decision rule. Though this disturbance to the systematic behavior of policymakers, a "structural policy shock", can always be coherently defined within a model, it's unclear why we should expect such an object to exist in reality. Policy shocks are usually represented as i.i.d. (a trembling hand), but this is equivalent to a policymaker who "flips coins" and is highly counterfactual relative to how involved policy deliberations are (Leeper et al., 1996; Nakamura and Steinsson, 2018b). Many common rationalizations are explicitly not exogenous changes in policy. For example, shocks are sometimes described as corresponding to discretionary acts, but the timing of discretion is not chosen arbitrarily (Debortoli and Lakdawala, 2016; King and Lu, 2025).

This paper provides a microfoundation for policy shocks and gives explicit conditions under which popular approaches identify them. Policy shocks occur because there is randomness in reaching consensus about the correct interpretation of data.¹ Beyond being an intuitive description of policy by committee, this mechanism is reflected in the data: monetary policy uncertainty is larger when there is more public disagreement about the trajectory of the macroeconomy. Structurally, the distinguishing feature of these shocks is higher-order moments of their distribution depend on the state of the world. These shocks also provide a theory of forward guidance. Away from the zero lower bound, forward guidance is typically modeled as a collection of orthogonal news shocks, but the news is not random and certainly correlated across the term structure (Bianchi et al., 2024). In contrast, the interpretation of data the central bank arrives at today will affect how data are interpreted in the future, generating a multi-horizon shock from a single stochastic event.

"Policy surprises", changes in policy expectations measured at high frequency in a narrow window around policy announcements, are popular instruments for estimating the effects of policy. Current work using high-frequency instruments takes the existence of a policy shock as given and imposes strong assumptions on the measurement of shocks via surprises (Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023a; Casini and McCloskey, 2026; Casini et al., 2026). This paper formally shows general conditions under which this approach is valid. For monetary policy specifically, the policy shocks in this paper are consistent with the interpretation of high-frequency surprises considered by Caballero and Simsek (2022); Sastry (2026) and explain why surprises have a state-dependent distribution. Of course, policy surprises can still occur for other reasons that will not satisfy the assumptions for validity. A commonly voiced concern is measurement can be compromised by superior information that policymakers have about the economy or their systematic policy

¹More generally, the ideas are also found in the literature on stochastic choice (Luce, 1959; Block and Marschak, 1960; Caplin and Martin, 2015; Fudenberg et al., 2015; Flynn and Sastry, 2023, 2024); even with a willingness to incur any reasoning or attention costs, a decisionmaker's uncertainty can lead to state-dependent randomization over options.

response.² In companion work (Bousquet, 2026a,b), I find these factors are not quantitatively important threats to identification, in line with existing comprehensive overviews of U.S. data (Bauer and Swanson, 2023a,b; Sastry, 2026; Acosta, 2026). But contrary to existing conjecture (Nakamura and Steinsson, 2018a, fn. 20; Bauer and Swanson, 2023a, p. 119), I show the closing off of this "learning from the policymaker" channel is necessary but not sufficient for high-frequency identification to work. Even if the central bank's beliefs are viewed as purely subjective, differences in beliefs themselves are not valid instruments and can contaminate the measurement of shocks. The most straightforward test of this channel is whether high-frequency surprises satisfy properties of rational expectations, but ostensible violations can appear ex post even when expectations were formed optimally in real time (Bauer and Swanson, 2023a; Eva and Winkler, 2023). For U.S. surprises, it appears that tail events do not occur at a frequency consistent with market expectations. Though sample size limitations make these exercises only suggestive in nature, there will still be problems for empirical estimation even if expectations are rational because forecast errors about beliefs can act as additional measurement noise. I conclude with a simple quantification of how an econometrician will struggle to estimate the effects of monetary policy when shocks are small. Policy shocks themselves are small, meaning there is a low ceiling on the signal-to-noise ratio (Ramey, 2016; Nakamura and Steinsson, 2018a,b) and any contamination exacerbates this problem.

2 A Short Primer on Identification

Shocks are commonly defined as disturbances to a fixed policy decision rule (a random change in the rule's intercept). At face value, we might think the effects of shocks and systematic policy are disconnected, so it's useful to establish the means and challenges for shocks to be helpful. To fix ideas, let $\mathbf{U}_t$ be latent states driving the economy, $\mathbf{I}_t$ be the anticipated path of interest rates, and $\mathbf{Y}_t$ be future economic conditions. We want to estimate the effects of $\mathbf{I}_t$ on $\mathbf{Y}_t$, but the classic simultaneity problem is $\mathbf{U}_t$ drives both. This problem is solved if we find some instrument z_t that affects $\mathbf{I}_t$ but is not caused by $\mathbf{U}_t$.




Figure 1: Causal diagram.

z_t here is really a "meta-instrument" – it satisfies the no simultaneous determination property for *all* future economic conditions. For all practical purposes, we can think of "meta-instruments" as equivalent

²Notable contributions for each case include Romer and Romer (2000); Nakamura and Steinsson (2018a); Jarociński and Karadi (2020); Miranda-Agrippino and Ricco (2021) and Bauer and Swanson (2023b); Cieslak et al. (2019).

to the previous definition of policy shocks, but this structure is more than is really required.³ In the spirit of Sims (1998) or the Cochrane (2004) discussion of Romer and Romer (2004), instruments can be valid for one outcome variable in $\mathbf{Y}_t$ and not another (e.g., a supply shock to trace out the demand curve). Such an instrument allows researchers to recover summary statistics of the causal effects of policy. Suppose we are interested in the effects of i_t on $Y_t^H = \{Y_{t+1}, \dots, Y_{t+H}\} \in \mathbf{Y}_t$. Define marginal policy effects as $\theta_t^h(i) = \frac{\partial Y_t^H(i)}{\partial i}$. If we can find some z_t^Y that is a determinant of i_t but is unrelated to the future determination of Y_t^H , then we can estimate something that will be informative about the effects of i_t on Y_t^H .

Proposition 1 (summarized): If z_t^Y is orthogonal to the joint distribution of $Y_t^H(\cdot)$ (given the time t filtration), then a linear projection of Y_t^H on z_t^Y estimates a positively weighted average of $\theta_t^h(\cdot)$.

Because similar results are well-known (Imbens and Angrist, 1994; Yitzhaki, 1996; Rambachan and Shephard, 2025; Kolesár and Plagborg-Møller, 2025), an abbreviated result is presented here with more formal assumptions and statement included in the Appendix. It's worth taking the time to state such a result nonetheless because it makes the conditions needed for identification clear. For example, this statement is entirely in terms of "nature's" expectations of Y_t^H (formed with the full information set and true probability measure). Later, we will consider candidate instruments $\hat{z}_t$ defined with different expectations (e.g., with limited information about i_t directly) that do not estimate a summary statistic of policy's systematic effects. In fact, what is estimated may be divorced from the transmission of policy entirely.

Because policy shocks are not necessary to estimate the effects of policy, an alternative approach would be dispensing with the shock paradigm altogether. But with notable exceptions (Cochrane, 2004; Cogleanese et al., 2025; Manuel and Wolf, 2026), virtually all identification approaches in macroeconomics rely on policy shocks. Despite the lack of theory on their origins, current treatments take their existence as given (Rambachan and Shephard, 2025; Kolesár and Plagborg-Møller, 2025; Casini and McCloskey, 2026; ?).⁴ The next section provides a microfoundation for policy shocks and links it explicitly to popular identification approaches.

3 Consensus Monetary Policy Shocks

Why might there be randomness in monetary policy? Monetary policy is set by a committee. If the outcome of committee deliberations was perfectly predictable, there would be no point in convening. Even if the decision for the current interest rate target is essentially a given, policy committees also convene to think

³Shocks are actually a special case of a "meta-instrument" because the policymaker could systematically respond to something unrelated to its objectives, like how many stars it could see in the sky last night. This distinction is certainly unimportant in application, so we will think of policy shocks as being the same as z_t in Figure 1.

⁴Christiano et al. (1999) suggest random shifts in preferences and measurement error.

about future decisions. In practice, these discussions boil down to how committee members expect the economy to evolve, meaning that each member’s optimal decision will generally differ from everyone else’s. So committees can be thought of as a single entity with an economic outlook that evolves with persistence but also with some intrinsic randomness (Blinder, 2007; Riboni and Ruge-Murcia, 2010; Laarits et al., 2026). Section 3.1 provides some supporting evidence for this interpretation. Section 3.2 formalizes a special case of how forming consensus could lead to policy shocks. Section 3.3 establishes why popular identification approaches would plausibly measure these shocks.

3.1 Motivating Empirical Evidence

Monetary policy committees presumably should be representative of the population (Reis, 2013). If disagreement about the economy’s trajectory is a source of randomness, a representative monetary policy committee should be more unpredictable when there is more disagreement within the public itself. To test this theory, we can look at the distribution of changes in interest rate expectations in a narrow window around policy announcements. Under an assumption that policy announcements do not reveal information about the economy, a significant high-frequency shift in expectations about future rates implies that policy decisions were uncertain.

To measure belief dispersion, we construct a disagreement measure using panel data from the Blue Chip Financial Forecast survey for 1993-2025. For inflation and real growth forecasts at horizons 1, 2, and 3 quarters ahead, the cross-sectional standard deviation is calculated and standardized. Then these six measures are aggregated together to create a single metric of macro-disagreement. For interest rate uncertainty, we look at the absolute value of the change in 1-quarter-ahead Eurodollar futures (ED2) in a narrow window around Federal Reserve events (Jarociński, 2024; Acosta et al., 2026). The zero lower bound is excluded since movement in short-term interest rate futures was negligible during this time, and unscheduled meetings are excluded to prevent results from being driven by outliers.⁵ Figure 2 suggests higher-order moments of these high-frequency surprises depend on underlying disagreement. It’s commonly noted that interest rate surprises are typically negligible, but the frequency of non-negligible events and the mean of absolute surprises clearly increase with disagreement. An R^2 of 10% from a single variable is substantial relative to existing numbers in the literature on interest rate surprises (Acosta, 2026; Sastry, 2026).

⁵Another motivation for excluding unscheduled meetings is Fed staff forecasts are not available.




Figure 2

To avoid false inference, we should establish whether disagreement about the state of the world is truly the primitive driving force. The Appendix discusses some exercises to provide additional context, which are summarized here. In general, we would expect disagreement to be driven by times when forecasting itself is difficult. One indicator of a difficult forecasting environment is the size of forecast revisions. A similar composite measure of forecast revision magnitude from the Blue Chip is indeed strongly correlated (.83) with the disagreement measure. Though the difficulty of forecasting is endogenous to the state of the world, I don't find a control variable that attenuates the relationship. The positive relationship remains significant when I control for the VIX, the real-time output gap, the economic policy uncertainty index (Baker et al., 2016), and consensus forecasts themselves.

There are still some caveats on these results. Monetary policy uncertainty is endogenous to the communication about the policymaker's beliefs. The 1990s are a good example of why forecast dispersion does not explain all of monetary policy uncertainty. Figure 3 shows that at the beginning of the 90s, monetary policy uncertainty outweighed macro uncertainty, presumably because the Fed was still in the early stages of becoming more transparent. A few years later, while there was extensive disagreement about whether Alan Greenspan's "productivity miracle" would materialize, the view of the Fed was resolute and known (low policy uncertainty). The remaining sample shows a much stronger association between disagreement and surprises. Next, we derive a simple policy rule in which the central bank puts weight on its own view and on the public average, with these weights evolving randomly.




Figure 3: Blue Chip forecast dispersion vs. ED2 monetary policy surprises. NBER recessions are shaded in gray and periods at the zero lower bound are shaded in blue.

3.2 A Special Case of Consensus Shocks

There are several ways to represent a decision as containing a random component because of a stochastic aggregation of disparate views. A transparent, low-dimensional rule is the central bank puts some weight on its own view and some weight on the prevailing public view, as in [Caballero and Simsek \(2022\)](#). For simplicity, suppose the central bank follows a [Taylor \(1993\)](#) rule over X_t , where its own belief is $\hat{X}_t^{CB}$ and the average public view is $\bar{X}_t$. A rule of the form of [Caballero and Simsek \(2022\)](#) is

$$i_t = \omega_t \phi X_t^{CB} + (1 - \omega_t) \phi \bar{X}_t.$$

It's plausible that the accommodation of public opinion is not a constant over time. So let

$$\omega_t = (1 - \rho_\omega) \bar{\omega} + \rho_\omega \omega_{t-1} + \varepsilon_t^\omega.$$

Using a quasi-local approximation, we can consider this a special case of the rule above, but with a state-dependent monetary policy shock

$$i_t = \bar{\omega} \phi X_t^{CB} + (1 - \bar{\omega}) \phi \bar{X}_t + \tilde{\omega}_t (X_t^{CB} - \bar{X}_t),$$

where $\tilde{\omega}_t = \frac{\omega_t - \bar{\omega}}{\bar{\omega}}$. The primary difference between this shock $\varepsilon_t^m = \tilde{\omega}_t (X_t^{CB} - \bar{X}_t)$ and a traditional monetary policy shock is the endogenous dependence of its higher-order moments. Whereas an i.i.d. white noise shock amounts to consulting a "coin flip" when setting i_t ([Leeper et al., 1996](#); [Nakamura and Steinsson, 2018b](#)), randomness that occurs as a part of a deliberative process will generally not leave the same clean residual.

3.3 Are Surprises Shocks?: Linking Consensus Shocks to Common Practice

A large literature beginning with [Kuttner \(2001\)](#); [Cochrane and Piazzesi \(2002\)](#) estimates the effects of monetary policy using policy surprises like the aforementioned ED2, broadly defined as high-frequency changes in the nominal yield curve in a narrow window around central bank announcements and press events. Suppose the central bank really does follow the interest rate rule of the preceding section. Will the high-frequency surprises measure these shocks? In the context of Proposition 1, this may seem to be an overly strong requirement, as surprises could in principle be a valid instrument without measuring a shock (a "meta-instrument"). But Proposition 1 picked out a specific instrument z_t^Y for a specific outcome variable Y_t^H . Policy surprises are simply measuring the change in expectations about future interest rates. Because they are not targeted towards a specific variable, they must satisfy valid instrument conditions for *any* variable we are interested in studying. So high-frequency identification must in general measure policy shocks.

Though a special case of the consensus shock environment can lead to a correct measurement of shocks, certain information structures can also pose problems. The consensus shock rule considered here can be used to illustrate why information asymmetries have led to concerns about high-frequency identification. First, X_t^{CB} is not necessarily public knowledge. In this case, forecast errors of i_t , a particular type of monetary policy surprise, will contain traces of X_t^{CB} . Many go a step further, arguing that surprises are not valid instruments if X_t^{CB} is formed with information not spanned by common knowledge ([Nakamura and Steinsson, 2018a](#); [Jarociński and Karadi, 2020](#); [Miranda-Agrippino and Ricco, 2021](#)). Though it appears central banks are good but not above median forecasters ([Bauer and Swanson, 2023b](#)), i_t may still be taken as a public signal about X_t if information is dispersed throughout the economy ([Morris and Shin, 2002](#); [Melosi, 2017](#); [Acosta, 2026](#)). Also, the central bank may have private knowledge about ω_t . In the simple case here, this would translate to private knowledge about the form of its policy rule ([Bauer and Swanson, 2023b](#)). Broadly defined, these two possibilities constitute a "learning" channel for monetary policy. Another possibility is that because surprises are measured using financial assets, risk premia may distort measurement. It has been conjectured that the absence of the "learning" and premia channels is a sufficient condition for surprises to be valid instruments ([Nakamura and Steinsson, 2018a](#), fn. 20; [Bauer and Swanson, 2023a](#), p. 119).

In companion work ([Bousquet, 2026a,b](#)), I look at the typically discussed channels for measurement contamination and find no role for a learning channel. For the purposes of this paper, we will take this conclusion as given. Contrary to the earlier conjecture, [Section 4](#) shows that the absence of a "learning effect" in surprises is necessary but not sufficient for estimating the effects of policy. Returning to the earlier illustration, when X_t^{CB} is private information, it can still contaminate measurement even if it is

regarded as a purely subjective belief that we don't learn from. [Section 5](#) estimates a model of dispersed information in order to capture this belief disagreement channel. An econometrician in this environment will struggle to estimate the effects of policy using surprises because policy shocks have low variance and small contamination can have an outsized effect. Before proceeding, it's worth exploring evidence for this new form of contamination.

3.4 Are Surprises Shocks?: Evidence from Data

Arguably the most attractive feature of identification through high-frequency surprises is not having to specify a policy rule, which inevitably will lead to extensive misspecification ([Chari et al., 2009](#)). The true decision rule from the universe of data to a target rate decision is high-dimensional, and while low-dimensional representations offer tremendous insight ([Taylor, 1993](#)), the fit is generally poor even at low frequencies ([Nakamura et al., 2025](#); [Bousquet, 2026b](#)). Because changes in systematic policy seem to occur infrequently ([Romer and Romer, 2023](#); [Bianchi et al., 2026](#)), market expectations should roughly track the high-dimensional mapping between the universe of data and policy. At the same time, since central banks do not follow a literal mathematical formula and there are frictions in communicating high-dimensional beliefs, there is still ample room for market misspecification.

A corollary of rational expectations is any information available at the time of the forecast cannot have been used to improve the forecast. An ostensible violation of this has been documented for monetary policy surprises, which are correlated with information publicly available prior to the announcement ([Bauer and Swanson, 2023b](#); [Miranda-Agrippino and Williams, 2026](#)).⁶ But these estimated correlations are unstable, sensitive to outliers, and weak out of sample ([Acosta, 2026](#); [Sastry, 2026](#)), strongly suggesting they only appear ex post ([Eva and Winkler, 2023](#)). When learning is slow, such finite sample artifacts will be abundant ([Bauer and Swanson, 2023a](#)) and are not reflective of any nuisance to be corrected. Though this particular piece of evidence may not be informative of whether market expectations are misspecified, we can consider a weaker test. Higher-order moments of the distribution of forecast errors can be state-dependent, but the mean should not be. [Figure 4](#) shows monetary policy surprises (ED2) are conditionally skewed – they are likely to share the sign of the policy adjustment that actually occurred.

⁶[Sastry \(2026\)](#) points out that the central bank's private information cannot generate this fact on its own.




Figure 4: High-frequency changes (ED2) in futures from 1990-2026, cumulated by month. Note: the y-axis is truncated to have uniform scaling across panels; there are over 100 "no surprise" events when the overnight rate target is unchanged.

The skew is severe enough that ED2 is not conditionally mean zero. One explanation is risk premia. But for this to be the case, there must be not only time-varying risk premia but endogenous changes in risk premia as a result of policy announcements. Though this is certainly plausible (Bianchi et al., 2026; Cieslak and Khorrami, 2026), it may not be enough to explain these patterns. Because risk premia on short-term interest rate futures contracts are small, if present at all Piazzesi and Swanson (2008), they relate to surprises clustered around zero. In order to have a conditionally skewed distribution that is mean-zero, we would need more frequent realizations of lower probability events (e.g., tightening surprises during easing cycles). The Appendix discusses further tests for the absence of these infrequent events using second moments of the surprise distribution, reaching similar conclusions. It's also possible the sample is simply too small to contain enough realizations of these low probability events. But taken at face value, it appears that market expectations do not perfectly measure a policy shock. Section 4 discusses a possible contributing factor.

4 A New Wedge Between Surprises and Shocks

In companion work (Bousquet, 2026a,b), I conclude the learning channel is not prevalent during monetary policy announcements. Informally, it has been suggested this is a sufficient condition for monetary policy surprises to identify policy shocks. This conjecture is untrue because, as discussed at the end of Section 3, surprises must in general measure shocks to be valid instruments.⁷ This section will present several examples

⁷There are exceptions, as discussed in Section 4.3. More generally, identification need not be exact: an excluded coefficient need not be exactly zero but "small relative to the other coefficients" (Leeper et al., 1996), and the excluded coefficients do not interact with a persistent part of the data generating process (Miranda-Agrippino and Ricco, 2023).

of why surprises need not measure shocks even absent the learning channel.

It may be useful to first consider a more informal line of argument. When the central bank committee is deliberating, they can observe market expectations. Therefore when their target decision is released, they know exactly how much markets will be "surprised" at the shortest end of the yield curve. So this "surprise" is in the central bank's information set when they make a decision. Monetary policy in practice largely boils down to a forecasting problem. Insofar as the forecast errors of the central bank will affect the macroeconomy (Friedman, 1968), these real effects should be orthogonal to the surprise under a rational expectations benchmark because the central bank knew what the surprise would be when they set the interest rate. This is developed further in the remainder of the section.

4.1 Using Signals to Generate Policy Surprises

There are several ways to introduce unpredictable movements in policy other than the central bank introducing arbitrary disturbances to a policy rule. Fundamentally, all such approaches amount to the public lacking some form of information. For example, in a Markov switching setup (Davig and Leeper, 2007), there would be no pure policy surprises if the latent states were known in advance. But one reason surprises need not measure shocks is because unpredictable policy can be generated by a dispersed information environment with no stochastic component in the central bank's decision. In such an environment, it may still be the case that the information the central bank uses to make their decision is not perceived as useful by the private sector. It's been suggested that policy shocks might be recoverable because, even though the public understands the central bank is trying to act optimally according to its objectives, the exact inclination behind any particular decision remains opaque. The following framework seeks to replicate as closely as possible the line of justification for treating surprises as shocks.

To generate surprises, we consider a class of models with a particular constraint on the central bank and a general information structure. The central bank must announce a policy function at the start of period t that will determine the realization of policy instruments at $t + 1$. The policy function is measurable with respect to all common knowledge at t and a signal the central bank receives between periods. This structure is purposefully broad, nesting cases where there are several policy instruments, signals, or different assumptions about what information is common to all agents. Without loss of generality, we will represent a scalar case with policy function

$$i_{t+1}(s) : \bar{\mathcal{I}}_t \times \mathbb{R} \rightarrow \mathbb{R} \tag{1}$$

where $\bar{\mathcal{I}}_t$ is the common time t information set implicit in the time $t + 1$ subscript of the function. For expositional purposes, consider this policy function the mapping that determines tomorrow's interest rate.

If agents knew ahead of time the central bank would receive signal s between periods, they would know tomorrow's interest rate exactly. The central bank can decide to make tomorrow's interest rate known regardless of the signal's realization by making the function constant in s . More traditional uses of the signal would, for instance, include committing to use the signal as prescribed by a standard Kalman filter. Under these restrictions, surprises from optimal rules will be caused by intentional behavior that nonetheless is driven by some level of randomness — under standard assumptions, the central bank's signal could fall *anywhere* on the real line. This is precisely the black box element the earlier informal conjecture about surprises and shocks relates to.

4.2 Irrelevance Results for Surprises

Let $\mathbb{E}[\cdot]$ denote the objective expectations operator with respect to a given data generating process. To distinguish it from the notation in [Section 2](#), we will use $\hat{z}_t$ as the monetary policy surprise (since it is a candidate instrument, but not necessarily a valid instrument).

Proposition 2: From [\(1\)](#), let $\hat{z}_t = i_t - \mathbb{E}[i_t \mid \bar{\mathcal{I}}_{t-1}]$ and let $s \notin \bar{\mathcal{I}}_{t-1}$ denote the additional random variable i_t can be conditioned on. If $\mathbb{E}[y_{t+h} \mid \bar{\mathcal{I}}_{t-1}, s] = \mathbb{E}[y_{t+h} \mid \bar{\mathcal{I}}_{t-1}]$, then $\text{Cov}(y_{t+h}, z_t) = 0$.

This is a simple corollary of rational expectations, but applied to policy surprises it is arguably harder to reconcile with standard intuition. First, we can consider a purely statistical explanation by the law of iterated expectations

$$\text{Cov}(y_{t+h}, \hat{z}_t) = \mathbb{E} \left[\hat{z}_t \left(\mathbb{E}[y_{t+h} \mid \bar{\mathcal{I}}_{t-1}, s] - \mathbb{E}[y_{t+h} \mid \bar{\mathcal{I}}_{t-1}] \right) \right] = 0,$$

where the last step follows from the sufficient condition.⁸ Note that $\hat{z}_t$ is a known object given the central bank's between-period information set. Optimal policy for the central bank, because of its timing constraint, involves a heavy forecasting component. Because $\hat{z}_t$ cannot be correlated with the central bank's forecast error, this property will extend to many of the aggregates y_t in a broad class of models. For example, consider the simplest "divine coincidence" equilibrium in a New Keynesian model ([Benigno and Woodford, 2004](#)) where inflation π_t and the output gap $\tilde{y}_t$ are directly proportional to $i_t - r_t^*$, the gap between the interest rate and the natural rate of interest. If the central bank follows a rule where it seeks to track the natural rate, it will always announce a policy function of $i_{t+1} = \mathbb{E}[r_{t+1}^* \mid \bar{\mathcal{I}}_t, s]$ under the constraints in this setup. This means future inflation and output gap are always expected to be zero, both before and after the central bank receives its signal. So policy surprises $\hat{z}_t$ must be uncorrelated with the realized output gap and inflation at all horizons. This result generalizes broadly in linear settings. Notice the logic is not altered even

⁸This shows the sufficient condition of Proposition 1 can be weakened slightly.

if there exists a "central bank information effect" because the commitment to the policy rule already tells private agents all they need to know about the future. The result can also apply under more complicated policy functions and models (e.g., if we include cost-push shocks and break the "divine coincidence"). The sole requirement is the information structure is such that expectations about aggregates do not change once the central bank receives its signal.

Proposition 2 is an "irrelevance result" because it shows a broad class of environments where policy surprises have no propagation. But it may be useful to consider a different kind of irrelevance result: one that illustrates the propagation of belief surprises can go directly against the underlying monetary transmission mechanism. Again, we consider the aforementioned New Keynesian model but explicitly add cost-push shocks⁹

$$\tilde{y}_t = \mathbb{E}_t[\tilde{y}_{t+1}] - i_t + r_t^* + \mathbb{E}_t[\pi_{t+1}] \quad (2)$$

$$\pi_t = \beta \mathbb{E}_t[\pi_{t+1}] + \kappa \tilde{y}_t + u_t. \quad (3)$$

We will consider two ways to get a "puzzling" comovement between policy surprises and inflation. The first will involve a hypothetical where the central bank follows a known, non-optimal forecast rule. As discussed further in Online Appendix B, this first result has a precedent in [Caballero and Simsek \(2022\)](#) where the central bank follows a non-optimal forecast rule because of its dogmatic beliefs, making the correlation between policy surprises and aggregates essentially unrestricted. Suppose $\text{Var}(u_t) = 0$, r_{t+1}^* follows a (known) linear law of motion with Gaussian innovations, and the central bank receives a signal about r_{t+1}^* between t and $t + 1$. The aforementioned policy rule in this model is for the central bank to commit to setting $i_{t+1} = b_t(s)$, where b_t is the objective forecast for r_{t+1}^* prescribed by the Kalman filter after observing signal s .

Example 1: Suppose the policy rule is $i_{t+1} = \alpha b_t(s)$ with $\alpha \in (0, 1)$. In the equilibrium in (2), (3) generated by $\text{Var}(u_t) = 0$, inflation and the output gap are directly proportional to $i_t - r_t^*$ and $\text{Cov}(\pi_t, \hat{z}_t), \text{Cov}(\tilde{y}_t, \hat{z}_t) > 0$.

Example 1 inverts the logic of Proposition 1: when the private sector forecast of interest rates is too high, interest rates were on average realized below the optimal level. This is because the central bank is following an "attenuated" forecast rule, so the corollaries of rational expectations no longer hold. Similar results hold if the central bank instead follows an inertial forecast rule. There is also a connection to optimal policy in an unconstrained environment. If a central bank follows the real rate rule in [Nakamura and Steinsson \(2018a\)](#) $i_t - \mathbb{E}_t[\pi_{t+1}] = \psi r_t^* + \phi \pi_t$ with $\psi < 1$, then rates adjust less than one-for-one with the natural rate. This

⁹Throughout, we make the implicit assumption that the unique equilibrium minimizing $\sum_t \beta^t (\tilde{y}_t^2 + \pi_t^2)$ is always selected.

means an impulse to r_t^* will generate positive responses of inflation and the output gap in textbook New Keynesian models.

Example 1 considers an environment with suboptimal policy, but a similar irrelevance result can still hold under optimal policy once we change what the central bank receives a signal about. For simplicity, assume $\{r_t^*\}_{t=0}^\infty$ is known at $t = 0$. Instead, the central bank faces the same timing constraint but receives a between-period signal s about u_{t+1} . Appendix A.2 shows optimal policy leads to the following equilibrium expressions, largely determined by the central bank's forecast $m_t = \mathbb{E}[u_t \mid \bar{\mathcal{I}}_{t-1}, s]$

$$i_t = r_t^* + cm_t \quad \tilde{y}_t = -cm_t \quad \pi_t = u_t - \kappa cm_t, \quad (4)$$

for a specific $c \in \mathbb{R}$ pinning down the optimal response parameter. If u_t is i.i.d. normal with zero mean, the policy surprise is simply $\hat{z}_t = i_t - r_t^* = cm_t$. This $\hat{z}_t$ has a very different relationship with inflation than a hypothetical structural policy shock in this model would.

Example 2: In the equilibrium described by (4), $\text{Cov}(\pi_t, \hat{z}_t) > 0$.

4.3 The Generality of Irrelevance

Section 3.2 shows belief surprises are fundamentally different objects than structural shocks, and not for the reasons often given. Nakamura and Steinsson (2018a); Miranda-Agrippino and Ricco (2021) note that when the public uses monetary policy decisions to extract information about economic fundamentals, there is scope for "puzzling" correlations. By contrast, these results apply even when private agents perceive the central bank's signal as containing only information spanned by the public information set. This can be easily accomplished through modifications in the timing and signal structure, as discussed in Section 5. This alternative structure is meant to mimic the availability of signal aggregators like financial markets and surveys, which presumably should depress the weight put on the subjective assessment of one particular agent.

Though these results are motivated by fairly general settings, they are confined to an implicit or explicit assumption of linearity. The "no response" result of Proposition 2 is vulnerable to nonlinearities because Jensen's inequality will tend to break the sufficient condition. At face value, the wealth of nonzero responses estimated in the literature (Ramey, 2016; Montiel Olea et al., 2025) might appear congruent with this consequence. Yet Proposition 2 is simply a relationship that holds *in population*. Data generating processes with an abundance of persistence can likewise create an excess of finite sample correlations that vanish asymptotically. Bousquet (2026b) considers an environment with slow learning rates and shows "hump-shaped" responses can exist in small samples even when there will be no response with sufficient data.

To zero in on the crux of the issue, we can again consider the simple divine coincidence structure but relax the assumptions on the information structure by letting policy surprises be defined by potentially non-nested information sets.

Example 3: Define the policy rule $i_t(s) = \mathbb{E}[r_t^* | \mathcal{F}_{t-1}, s]$, expected policy $i_t^e = \mathbb{E}[i_t(s) | \mathcal{G}_{t-1}]$, and policy surprise $\hat{z}_t = i_t(s) - i_t^e$. In the $\text{Var}(u_t) = 0$ equilibrium for (2), (3), an IV regression of $\tilde{y}_t$ on $\hat{z}_t$ will be biased for the correct loading (-1) of the policy rate in (2): $(-1 - \beta^{IV}) \propto [\text{Var}(i_t) - \text{Cov}(i_t^e, r_t^*)]$. When $\mathcal{G}_t \subset \mathcal{F}_t$, $\beta^{IV} = 0$. In the limit where $|r_t^* - \mathbb{E}[r_t^* | \mathcal{G}_{t-1}]| \rightarrow 0$ but $|r_t^* - i_t(s)| > 0$ almost surely, $\beta^{IV} \rightarrow -1$.

The model of Sims (1987) features an "as good as random" policy instrument because the central bank acts based on an advance estimate of a data release, whereas private agents form expectations and make decisions using the actual data release. Example 3 recovers this result as a special case, showing $\hat{z}_t$ becomes an arbitrarily good instrument as the private information set $\mathcal{G}_{t-1}$ expands while holding the central bank's information set $\mathcal{F}_t$ fixed. In reality, there is a range of possibilities, as in Caballero and Simsek (2022). Taken together, Examples 1-3 show the extent to which belief surprises can function as an instrument depends on a delicate interconnectedness of the equilibrium conditions, the policy rule, and the exact information and signal structure. The exact wedge is that forecast errors of policy are not sufficient on their own to solve an omitted variable problem. Sims (1998) hones in on a useful analogy between using demand shifters to trace out the supply curve and vice versa. In this framework, $\hat{z}_t$ does not act as a pure shifter because it is generically correlated with r_t^* .

Efforts in Jarociński and Karadi (2020, 2025) seek to partition surprises into "information shocks" and policy shocks. But unless the central bank has some sort of trembling hand, all policy surprises are information revelations of a certain form (Campbell et al., 2012). This means such partitioning can be useful to separate information revelations we think are more proximate to a textbook policy shock (e.g., an unexpectedly bullish reading of public data) from information the central bank might be revealing about its institutional traits, like acts of discretion (Debortoli and Lakdawala, 2016; King and Lu, 2025) or insufficient commitment to its goals (Cieslak et al., 2026; Ai and Hu, 2025). Still, even what we might call "relevant" information shocks may not be isomorphic to policy rule innovations in a standard model.¹⁰ To satisfy the conditions of Proposition 1 in general, we need to go a step further and impose structure that generates true randomness.

¹⁰In Appendix A.3, I simulate data from Smets and Wouters (2007) with the monetary policy shock variance set to zero and show a "small" VAR system consisting just of output, inflation, and interest rates will nonetheless recover "masquerading" policy shocks (Wolf, 2020). The responses can be similar to those estimated from US data and the VARs themselves pass diagnostic stability checks for all common methods.

5 Estimating Policy Effects with Surprises

Section 4 showed an assumption of randomness is not only sufficient but necessary for high-frequency identification to be valid. The model in Section 3 offers a rationale for such an assumption. However, once we take a stand that high-frequency surprises measure a monetary shock, we are also placing a ceiling on the exogenous variation in policy. This ceiling will be prohibitive, even without the risks that high-frequency identification is contaminated by the policymaker’s internalization of how surprise changes will propagate to the economy. I illustrate this dilemma using a generic structural model.

Suppose we are interested in the effects of i_t on Y_{t+h} . Let the data be generated by

$$Y_{t+h} = \theta_h i_t + u_{t+h} \quad (5)$$

$$i_t = g(\mathbf{X}_t) + \varepsilon_t \quad \varepsilon_t \perp u_{t+h}. \quad (6)$$

Assume we measure $z_t = \varepsilon_t$ exactly. How successful will an econometrician be at recovering a sufficiently precise estimate of θ_h ? This depends on several factors, summarized by a sufficient statistic for statistical power δ_h

$$\delta_h = \sqrt{T} \theta_h \frac{\sigma_\varepsilon}{\sigma_u}. \quad (7)$$

δ_h is increasing in the size of the sample, the strength of the effects, the amount of exogenous variation, and the share of the *total* variation of Y_{t+h} explained by monetary policy at t . The upper bound for δ_h is low. Generous parameterizations for high-frequency analysis would place $T \leq 200$, $\theta_h \leq 1$, $\sigma_u \geq 1$, and $\sigma_\varepsilon \leq .1$. Figure 5 quantifies the argument outlined in Ramey (2016): the less randomness in policy, the more likely any hypothesis test rejection is simply a false positive. At the maximum policy randomness $\sigma_\varepsilon = .1$, there is modest detectability for sufficiently large effect sizes, but most estimates typically place θ_h significantly lower. Of course, many impulse response figures in the empirical macroeconomics literature show confidence bands excluding zero. The power calculations suggest this is caused by underestimated standard errors. Adjustments for multiple-hypothesis testing are seldom implemented, despite important contributions in time series inference Montiel Olea and Plagborg-Møller (2019); Xu (2025); Inoue et al. (2026); Nikolaishvili and Gade (2026). The selection of control variables is also a consequence of both collective and individual search, and standard errors often do not reflect external use of controls.




Figure 5: Rejection probability as a function of effect size in (5). $T = 200$ and $\sigma_u = 1$ are fixed.

6 Conclusion

The paper presents a model of where monetary policy shocks come from — randomness in reaching a consensus — and applies this model to a study of high-frequency identification, which uses measured changes in expectations about future policy action attributable to policymaker communication as an instrument. This paper shows that for the standard implementation of this identification strategy, these measured changes must generally be capturing random behavior even from the policymaker’s perspective. Otherwise, identification is only guaranteed to be valid under the knife-edge assumption that policymakers have strictly less information than the public, as considered in Sims (1987). The "consensus shocks" outlined in this paper give a model under which this assumption is defensible, but they present two potential problems. First, it appears that the public systematically underpredicts how likely the central bank is to choose the most likely choice. Second, even if the identification is valid, these shocks are not large enough to reliably estimate policy effects, expanding on a concern voiced in Ramey (2016); Nakamura and Steinsson (2018a). Future work with the high-frequency approach could include relaxing the current assumptions through considering alternative regression specifications or pursuing a hybrid approach where high-frequency estimates map into a structural model (Fraisse et al., 2026).

References

- Acosta, Miguel.** 2026. “The Perceived Causes of Monetary Policy Surprises.” Working Paper.
- Acosta, Miguel, Andrea Ajello, Michael Bauer, Francesca Loria, and Silvia Miranda-Agrippino.** 2026. “Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database.” working paper.
- Ai, Hengjie, and Xinxin Hu.** 2025. “A Model of Fed Information Effect.” Working paper.
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis.** 2016. “Measuring Economic Policy Uncertainty*.” *The Quarterly Journal of Economics*, 131(4): 1593–1636.
- Bauer, Michael D., and Eric T. Swanson.** 2023a. “A Reassessment of Monetary Policy Surprises and High-Frequency Identification.” *NBER Macroeconomics Annual*, 37.
- 2023b. “An Alternative Explanation for the "Fed Information Effect".” *American Economic Review*, 113(3): 664–700.
- Benigno, Pierpaolo, and Michael Woodford.** 2004. “Optimal Monetary and Fiscal Policy: A Linear-Quadratic Approach.” In *NBER Macroeconomics Annual*. 18.
- Bianchi, Francesco, Sydney C. Ludvigson, and Sai Ma.** 2022. “Belief Distortions and Macroeconomic Fluctuations.” *American Economic Review*, 112(7): 2269–2315.
- 2026. “A Structural Approach to High-Frequency Event Studies: The Fed and Markets as Case History.” Working Paper.
- Bianchi, Francesco, Leonardo Melosi, and Giovanni Nicolò.** 2024. “Is There a Forward Guidance Puzzle?” 7, Working Paper.
- Blinder, Alan S.** 2007. “Monetary policy by committee: Why and how?” *European Journal of Political Economy*, 23(1): 106–123.
- Block, H. D., and J. Marschak.** 1960. “Random Orderings and Stochastic Theories of Response.” In *Contributions to Probability and Statistics: Essays in Honor of Harold Hotelling*. eds. by Ingram Olkin, Sudhish G. Ghurye, Wassily Hoeffding, William G. Madow, and Henry B. Mann.
- Bousquet, Paul.** 2026a. “How Central Bank Announcements Affect Expectations.” Working Paper.
- 2026b. “What Policy Surprises Reveal About the Policy Rule.” Working Paper.

- Caballero, Ricardo J., and Alp Simsek.** 2022. “Monetary Policy with Opinionated Markets.” *American Economic Review*, 112(7): 2353–2392.
- Campbell, Jeffrey R., Charles L. Evans, Jonas D. M. Fisher, and Alejandro Justiniano.** 2012. “Macroeconomic Effects of Federal Reserve Forward Guidance.” *Brookings Papers on Economic Activity* 1–80, See Chris Sims’ discussion comments.
- Caplin, A., and D. Martin.** 2015. “A testable theory of imperfect perception.” *The Economic Journal*, 125(582): 184–202.
- Casini, Alessandro, and Adam McCloskey.** 2026. “Identification and Estimation of Causal Effects in High-Frequency Event Studies.” June, Working Paper.
- Casini, Alessandro, Adam McCloskey, Luca Rolla, and Raimondo Pala.** 2026. “Dynamic Local Average Treatment Effects in Time Series.” Working paper.
- Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan.** 2009. “New Keynesian Models: Not Yet Useful for Policy Analysis.” *American Economic Journal: Macroeconomics*, 1(1): 242–266.
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James M. Robins.** 2018. “Double/Debiased Machine Learning for Treatment and Structural Parameters.” *The Econometrics Journal*, 21(1): C1–C68.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans.** 1999. “Monetary policy shocks: What have we learned and to what end?” In *Handbook of Macroeconomics*. eds. by J. B. Taylor, and M. Woodford, 1, Chap. 2 65–148.
- Cieslak, Anna, and Paymon Khorrami.** 2026. “Monetary Shocks and Long-Run Neutrality in Asset Pricing.” Working paper.
- Cieslak, Anna, Michael McMahon, and Hao Pang.** 2026. “Did I Make Myself Clear? The Fed and the Market under the 2020 Monetary Policy Framework.” Working paper.
- Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen.** 2019. “Stock Returns over the FOMC Cycle.” *The Journal of Finance*, 74(5): 2201–2248.
- Cochrane, John H.** 2004. “Comments on “A new measure of monetary shocks: Derivation and implications.”” July, NBER EFG meeting.
- Cochrane, John H., and Monika Piazzesi.** 2002. “The Fed and Interest Rates - A High-Frequency Identification.” *American Economic Review*, 92(2): 90–95.

- Coglianese, John, Maria Olsson, and Christina Patterson.** 2025. “Monetary Policy and the Labor Market: A Quasi-Experiment in Sweden.” *American Economic Review*, 115(10): 3451–3486.
- Davig, Troy, and Eric M. Leeper.** 2007. “Generalizing the Taylor Principle.” *American Economic Review*, 97(3): 607–635.
- Debortoli, Davide, and Aeimit Lakdawala.** 2016. “How Credible Is the Federal Reserve? A Structural Estimation of Policy Re-optimizations.” *American Economic Journal: Macroeconomics*, 8(3): 42–76.
- Eva, Kenneth, and Fabian Winkler.** 2023. “A Comprehensive Empirical Evaluation of Biases in Expectation Formation.” Working Paper.
- Flynn, Joel P., and Karthik A. Sastry.** 2023. “Strategic mistakes.” *Journal of Economic Theory*, 212.
- . 2024. “Attention Cycles.” Working paper.
- de Fraisse, Antoine Hubert, Peter Maxted, Kunal Sangani, and David Sraer.** 2026. “Pre-Trends Ahead of Monetary Policy Surprises: Recovering the Response of Corporate Investment.” Working paper.
- Frank, Murray Z., Jing Gao, and Keer Yang.** 2025. “Behavioral Machine Learning? Regularization and Forecast Bias.” Working Paper.
- Friedman, Milton.** 1968. “The Role of Monetary Policy.” *The American Economic Review*, 58(1): 1–17.
- Fudenberg, Drew, Ryota Iijima, and Tomasz Strzalecki.** 2015. “Stochastic Choice and Revealed Perturbed Utility.” *Econometrica*, 83 2371–2409.
- Imbens, Guido W., and Joshua D. Angrist.** 1994. “Identification and Estimation of Local Average Treatment Effects.” *Econometrica*, 62(2): 467–475.
- Inoue, Atsushi, Òscar Jordà, and Guido M. Kuersteiner.** 2026. “Inference for Local Projections.” *The Econometrics Journal*, 29(1): 2–26.
- Jarociński, Marek.** 2024. “Estimating the Fed’s unconventional policy shocks.” *Journal of Monetary Economics*, 144, p. 103548.
- Jarociński, Marek, and Peter Karadi.** 2020. “Deconstructing Monetary Policy Surprises—The Role of Information Shocks.” *American Economic Journal: Macroeconomics*, 12(2): 1–43.
- . 2025. “Disentangling Monetary Policy, Central Bank Information, and Fed Response to News Shocks.” April, Working Paper.

- King, Robert G., and Yang K. Lu.** 2025. “Evolving Reputation for Commitment: Understanding Inflation and Inflation Expectations.” Working Paper.
- Kolesár, Michal, and Mikkel Plagborg-Møller.** 2025. “Dynamic Causal Effects in a Nonlinear World: The Good, the Bad, and the Ugly.” *Journal of Business & Economic Statistics*, 43(4): 737–754.
- Kuttner, Kenneth N.** 2001. “Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market.” *Journal of Monetary Economics*, 47(3): 523–544.
- Laarits, Toomas, Ben Matthies, Kaushik Vasudevan, and Will Yang.** 2026. “Policy by Committee.” Working paper.
- Leeper, E. M., C. A. Sims, and T. Zha.** 1996. “What does monetary policy do?” *Brookings papers on economic activity* 1–78.
- Luce, R. Duncan.** 1959. *Individual Choice Behavior: A Theoretical Analysis*.
- Manuel, Ed, and Christian K. Wolf.** 2026. “Identifying Policy Causal Effects from Rule Changes.” Working Paper.
- Melosi, Leonardo.** 2017. “Signalling Effects of Monetary Policy.” *The Review of Economic Studies*, 84(2): 853–884.
- Miranda-Agrippino, Silvia, and Giovanni Ricco.** 2021. “The Transmission of Monetary Policy Shocks.” *American Economic Journal: Macroeconomics*, 13(3): .
- . 2023. “Identification with External Instruments in Structural VARs.” *Journal of Monetary Economics*, 135 1–19.
- Miranda-Agrippino, Silvia, and John C. Williams.** 2026. “Interest Rate Surprises When the Fed Doesn’t Speak.” Working Paper.
- Montiel Olea, José Luis, and Mikkel Plagborg-Møller.** 2019. “Simultaneous Confidence Bands: Theory, Implementation, and an Application to SVARs.” *Journal of Applied Econometrics*, 34(1): 1–17.
- Montiel Olea, José Luis, Mikkel Plagborg-Møller, Eric Qian, and Christian K. Wolf.** 2025. “Local Projections or VARs? A Primer for Macroeconomists.” In *NBER Macroeconomics Annual 2025*.
- Morris, Stephen, and Hyun Song Shin.** 2002. “The Social Value of Public Information.” *American Economic Review*, 92 1521–1534.

- Nakamura, Emi, Venance Riblier, and Jón Steinsson.** 2025. “Beyond the Taylor Rule.” Working Paper.
- Nakamura, Emi, and Jón Steinsson.** 2018a. “High-Frequency Identification of Monetary Non-Neutrality: The Information Effect.” *The Quarterly Journal of Economics*, 133(3): 1283–1330.
- Nakamura, Emi, and Jón Steinsson.** 2018b. “Identification in Macroeconomics.” *Journal of Economic Perspectives*, 32(3): 59–86.
- Nikolaishvili, Giorgi, and Noah D. Gade.** 2026. “Scanning for Significance: False Discovery Control for Impulse Responses.” Working paper.
- Piazzesi, Monika, and Eric T. Swanson.** 2008. “Futures prices as risk-adjusted forecasts of monetary policy.” *Journal of Monetary Economics*, 55(4): 677–691.
- Rambachan, Ashesh, and Neil Shephard.** 2025. “When do common time series estimands have non-parametric causal meaning?”, Working Paper.
- Ramey, Valerie.** 2016. “Macroeconomic Shocks and Their Propagation.” In *Handbook of Macroeconomics*. Chap. 2.
- Reis, Ricardo.** 2013. “Central Bank Design.” *Journal of Economic Perspectives*, 27(4): 17–44.
- Riboni, Alessandro, and Francisco J. Ruge-Murcia.** 2010. “Monetary Policy by Committee: Consensus, Chairman Dominance, or Simple Majority?” *Quarterly Journal of Economics*, 125(1): 363–416.
- Romer, Christina D., and David H. Romer.** 2000. “Federal Reserve Information and the Behavior of Interest Rates.” *American Economic Review*, 90 429–457.
- 2004. “A New Measure of Monetary Shocks: Derivation and Implications.” *American Economic Review*, 94(4): 1055–1084.
- 2023. “Presidential Address: Does Monetary Policy Matter? The Narrative Approach after 35 Years.” *American Economic Review*, 113(6): 1395–1423.
- Sastry, Karthik A.** 2026. “Disagreement about Monetary Policy.” *American Economic Journal: Macroeconomics*, 18(2): 486–522.
- Sims, Christopher A.** 1987. “A Rational Expectations Framework for Short-Run Policy Analysis.” In *New Approaches to Monetary Economics*. eds. by William A. Barnett, and Kenneth J. Singleton, Cambridge, UK: Cambridge University Press, 293–308.

- . 1998. “Comment on Glenn Rudebusch’s “Do Measures of Monetary Policy in a VAR Make Sense?”” *International Economic Review*, 39(4): 933–941.
- Smets, Frank, and Rafael Wouters.** 2007. “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach.” *American Economic Review*, 97(3): 586–606.
- Taylor, John B.** 1993. “Discretion versus policy rules in practice.” *Carnegie-Rochester Conference Series on Public Policy*, 39 195–214.
- Wolf, Christian K.** 2020. “SVAR (Mis-)Identification and the Real Effects of Monetary Policy Shocks.” *American Economic Journal: Macroeconomics*, 12(4): 1–32.
- Xu, Ke-Li.** 2025. “Local Projection Based Inference under General Conditions.” Forthcoming.
- Yitzhaki, Shlomo.** 1996. “On Using Linear Regressions in Welfare Economics.” *Journal of Business & Economic Statistics*, 14(4): 478–486.

Appendix

A.1 Additional Empirical Exercises

A.2 Additional Derivations

A.3 Misidentification of Structural Shocks and Forecast Error Predictability¹¹

A.3.1 Inherency of Ex Post Predictability

When forecast errors are correlated with information in the forecaster's information set, this "predictability" is often described as a rejection of a full information rational expectations benchmark. [Eva and Winkler \(2023\)](#); [Bauer and Swanson \(2023a\)](#); [Sastry \(2026\)](#) note that many documented predictability results are based on correlations that, while sometimes statistically significant ex post, are unstable and with little to no predictive power out of sample. In other words, the estimated relationships do not offer much insight into how to improve forecasts. This suggests the correlations are *ex post* artifacts of an environment where it is difficult to learn. To demonstrate this further, we can extend an analogy made in [Bianchi et al. \(2022\)](#); [Frank et al. \(2025\)](#) – do the forecast errors of machine learning algorithms have similar empirical traits? The answer depends on the application and candidate data generating process (DGP). In general, "predictability" seems to be an inherent trait of high-dimensional estimation in small samples ([Frank et al., 2025](#)).

As an example, consider a classic environment in the machine learning literature ([Chernozhukov et al., 2018](#)) where the underlying DGP is sparse but unknown and the data set is high dimensional $K \gg T$. If the underlying DGP is indeed sparse, a lasso procedure will produce consistent estimation. Suppose we are trying to find the causal effect of a treatment D on an outcome Y , with time series data generated from the following system

$$Y_t = \theta D_t + g(\mathbf{X}_t) + v_t$$

$$D_t = q(\mathbf{X}_t) + u_t.$$

I assume $\dim(\mathbf{X}_t) = K$ is 2 orders of magnitude larger than the length of the time series T , but the underlying DGP is linear and sparse (i.e., $g(\cdot), q(\cdot)$ are linear functions of the same few components of $\mathbf{X}_t$). To estimate θ , lasso estimation is performed on Y_t and D_t , then the residuals of these first stages $(\tilde{Y}_t, \tilde{D}_t)$ are regressed on one another. While the convergence of $\hat{\theta} \rightarrow \theta$ is remarkably fast in this environment, $\{(\tilde{Y}_t, \tilde{D}_t)\}_{t=1}^T$ are ex post predictable. Specifically, if we regress these residuals on the "relevant" components of $\mathbf{X}_t$ as well as a few other randomly selected variables, 100% of these sets of residuals show several statistically significant

¹¹Code for these exercises is available at this paper's GitHub repository, <https://github.com/paulbousquet/WhySurprise>

correlations. Just as in standard rational expectations, these correlations would indeed vanish as $T \rightarrow \infty$. These results extend naturally to the job of professional forecasters and financial market participants; they must sort through reams of data to produce some conditional mean function. This exercise shows even "good" estimation procedures will generate some seeming idiosyncrasies in finite sample.

A.3.2 Masquerading Shocks

Suppose the economy follows a data generating process consistent with a VAR but an econometrician only includes a subset of the variables in estimation. [Wolf \(2020\)](#) shows that the recovered structural shocks will necessarily be functions of several *true* structural shocks, i.e., the recovered structuring shocks are "masquerading" as true primitive objects. This thinking can be extended to the case where there are actually no structural shocks present. [Figure 6](#) shows estimated impulse responses on a small VAR system using data generated by the [Smets and Wouters \(2007\)](#) model at their estimated posterior modes, with the exception that the monetary policy shock variance is set to zero. The results show that the recovered figures will mechanically produce reasonable-looking responses, with the dotted black line represented the actual impulse response [Smets and Wouters \(2007\)](#) estimate. Standard diagnostic tests to check for identification issues do not indicate any issues.




Figure 6