High-Frequency Identification*

Refet S. Gürkaynak Eric T. Swanson Jonathan H. Wright

This version: August 2026

Abstract

In macroeconomics and finance, it is very difficult to estimate causal effects with quarterly, monthly or similar lower frequency data because the relevant variables all tend to respond to each other. By contrast, in half-hour observation windows even monetary policy can be regarded as predetermined, which allows establishing causality. We discuss how this high-frequency identification idea is implemented for financial market event studies and macroeconomic vector autoregressions (VARs) and local projections (LPs), addressing finer points of data and measurement as well as conceptual issues and applications.

JEL Classification: E43, E44, E52, E58, G14

Keywords: High-frequency identification, event study, VAR, local projections, monetary policy

*Refet S. Gürkaynak: Bilkent University and CEPR. E-mail: refet@bilkent.edu.tr. Eric T. Swanson: University of California, Irvine and NBER. E-mail: eric.swanson@uci.edu. Jonathan H. Wright: Johns Hopkins University and NBER. E-mail: wrightj@jhu.edu. We are grateful to Muhammet Ali Canşı and Mahmut Sefa İpek for excellent research assistance.

1 Introduction

A perennial issue in applied economics research is the identification of causal effects. In macroeconomics and finance, the economic variables of interest, such as interest rates, stock prices, economic activity, and others are all interrelated and affect each other. In a macro-finance world where everything affects everything else, how is causality to be established?

In finance, the identification problem is typically one of estimating the effect of some change, such as a monetary policy decision, on a financial market variable, such as the 10-year Treasury yield. The system in question is:

$$\begin{aligned} \Delta y_t &= \alpha\,\Delta a_t + \boldsymbol{\gamma}\mathbf{w}_t + \eta_t \\ \Delta a_t &= \beta\,\Delta y_t + \boldsymbol{\delta}\mathbf{w}_t + \nu_t \end{aligned} \tag{1}$$

where y_t is an asset price, a_t the event, $\mathbf{w}_t$ a vector of additional variables that potentially affect a_t and y_t, and η_t and ν_t are iid shocks. The econometrician wishes to estimate α, but there are clear problems of reverse causality and simultaneity, since y_t may also affect a_t, and $\mathbf{w}_t$ may affect both.

In macroeconomics, the identification problem is typically one of tracing out the effects of a shock, perhaps again to the stance of monetary policy, on macroeconomic variables. The system of equations in this case is typically a monthly or quarterly vector autoregression (VAR) or local projection (LP) for a vector of variables $\mathbf{x}$,

$$\mathbf{x}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{u}_t, \tag{2}$$

here shown as a first-order VAR without loss of generality,[1] with parameter matrix $\mathbf{B}$ of appropriate dimensions and

$$\mathbf{u}_t \overset{iid}{\sim} (\mathbf{0}, \boldsymbol{\Omega})\,. \tag{3}$$

In general, innovations to the variables $\mathbf{x}_t$ each month t may be correlated, so that $\boldsymbol{\Omega}$ is not a diagonal matrix. The econometrician typically assumes that the variables $\mathbf{x}_t$ are driven by a set of exogenous structural *shocks* each month, denoted $\boldsymbol{\varepsilon_t}$, which are conceptually independent of each

[1] Any VAR can be rewritten as a first-order VAR in companion form, as discussed in Hamilton (1994) and Lütkepohl (2006).

other (e.g., Bernanke, 1986; Ramey, 2016) and thus have an identity covariance matrix:

$$\boldsymbol{\varepsilon_t} \overset{iid}{\sim} (\mathbf{0}, \boldsymbol{I})\,. \tag{4}$$

We write

$$\mathbf{u}_t \;=\; \mathbf{S}\varepsilon_t, \tag{5}$$

to describe the relationship between the reduced-form innovations $\mathbf{u}_t$ and the structural shocks $\boldsymbol{\varepsilon_t}$ each month, where $\mathbf{S}$ is a matrix of appropriate dimensions.[2]

Estimation of equation (2) is relatively straightforward using equation-by-equation ordinary least squares (OLS). The identification problem arises because the reduced-form innovations $\mathbf{u}_t$ are not conceptually independent, and although the structural shocks $\boldsymbol{\varepsilon_t}$ are independent, they are generally unobserved and thus their effects cannot be estimated without additional assumptions. In other words, although the reduced-form innovations $\mathbf{u}_t$ and the matrix $\boldsymbol{\Omega}$ can be estimated from the data, the structural shocks $\boldsymbol{\varepsilon_t}$ and matrix $\mathbf{S}$ cannot be.[3]

High-frequency identification begins with the observation that many data releases or policy announcements can be viewed as natural experiments with the proper choice of observation window, an idea that goes back at least to Fama et al. (1969). On the day that inflation, GDP, or unemployment data are released, for example, the information in that release corresponds to a variable that was determined in the past—the previous month or quarter—and is just being reported now. Thus, the release cannot be affected by any financial developments or other news that day, so the causality must flow from the release to the financial variables that change in response. In the model given by equation (1), treating w_t as scalar for ease of notation, Rigobon and Sack (2004) show that if we regress Δy_t on Δa_t by OLS, then the probability limit of the slope coefficient is

$$\alpha \,+\, (1-\alpha\beta)\frac{\beta(\sigma_\eta^2+\sigma_w^2)+\delta\sigma_w^2}{\sigma_\nu^2+\beta^2\sigma_\eta^2+(\beta+\delta)^2\sigma_w^2}\,, \tag{6}$$

[2]When $\mathbf{u}_t$ and $\boldsymbol{\varepsilon}_t$ have the same dimensions and $\mathbf{S}$ is invertible, then the VAR in equation (2) is also said to be *invertible*. However, high-frequency identification generally does not require this assumption, so $\boldsymbol{\varepsilon_t}$ could in principle cover a large set of structural shocks.

[3]For example, if $\boldsymbol{\Omega}$ and $\mathbf{S}$ are $n \times n$ matrices, $\mathbf{S}$ has n^2 unknown elements while $\boldsymbol{\Omega}$ has only $\frac{n(n+1)}{2}$ independent entries due to symmetry, requiring $\frac{n(n-1)}{2}$ additional restrictions to achieve identification.

where σ_η^2, σ_ν^2, and σ_w^2 denote the variances of η_t, ν_t, and w_t, respectively. If β and σ_w^2 (or δ itself) are both negligible (as would be the case with a tight observation window), then the probability limit is α. This is what it means to say that, at high frequency, the release is a natural experiment.

Policy announcements, such as monetary policy announcements, also can be viewed as natural experiments with a sufficiently narrow observation window. On the day of a policy announcement, there is typically some lag between the time when the decision is made and when the policy is announced, making the announcement predetermined as long as the observation window begins after the time of the decision. Ideally, an intradaily observation window would be used if the policy is decided earlier on the same day as the announcement, but even with a daily window, identification is achieved so long as the policy decision does not react to that day's developments (i.e., β and δ are negligible on policy days).

The idea of using these information releases at high frequency to make causal inference is called an *event study*. A similar idea can be used to identify the causal effects of a shock on the macroeconomic variables in (2). In this case, the change in asset prices around the high-frequency event is used as an instrumental variable for the unobserved monthly structural VAR shock. Whereas the classic event-study literature looked just at high-frequency effects of surprises on asset prices, once we use these events as instruments in a VAR, we have a tool to study the causal effects of shocks on macroeconomic variables that are observed at lower frequency. The high-frequency event study answers questions that are of interest to financial economists, such as the response of the yield curve to news about inflation.[4] The high-frequency-identified VAR answers questions of interest to macroeconomists, such as the effect of monetary policy on GDP, employment, and inflation.

Gürkaynak and Wright (2013) provide a review of the event-study literature at that time, but the literature has greatly expanded since then and links to the structural VAR literature have been developed. The goal of the present paper is to provide a review of the state of the art in these literatures. Throughout, our main running example of macroeconomic shocks being identified is U.S. monetary policy surprises, but other shocks are also considered. We will focus on frequentist

[4]There is an obvious macroeconomic aspect of this as well, since the yield curve reaction will depend on market participants' beliefs about future macroeconomic outcomes and risk pricing.

applications of high-frequency identification due to space constraints, although Bayesian treatments also exist, both for proxy-VARs and local projections with high-frequency instruments (Miranda-Agrippino and Ricco, 2021, 2023). While the approach to inference is different, the idea behind identification using high-frequency surprises remains the same.

2 Event Studies

Gürkaynak and Wright (2013) and Gürkaynak, Kısacıkoğlu, and Wright (2020) provide introductions to event studies. We provide a brief summary and discussion of recent applications here. As discussed above, the causality established by event studies rests on the simple point that during a narrow enough observation window, the announcement is not affected by financial or other developments. As a result, reverse causality and simultaneity can be ruled out. It is the high-frequency observation window that achieves identification.

An important issue is the measurement of the news, or surprise, component of the announcement. Financial markets are forward looking and asset prices already reflect expected future developments—if not, there would be a trading opportunity that would be profitable on average as the news is revealed. Studying the effects of announcements on asset prices thus requires having measures of the unexpected news that market participants receive when the announcement is made. One way to do this is to measure the surprise by comparing actual releases to survey expectations and taking the difference between the two, which works well when survey data are available. Another approach is to rely on the change in an asset price that is closely connected to the news being released, such as using short-term interest rates to measure monetary policy surprises.

The choice of event window is clearly important. In a liquid financial market, announcements cause a jump in the conditional mean of asset prices that is generally completed within a matter of minutes (Andersen et al., 2003; Gürkaynak et al., 2005a). For this reason, it is best to use a short, intradaily window when such data are available. Using daily data adds additional noise to the change in asset prices, making identification less precise.[5] Indeed, using daily data can even

[5] For example, Andersen et al. (2003) were able to find clear evidence of effects of news announcements on exchange rates using intradaily data; the evidence is much weaker with daily data.

break the exogeneity of a measured shock, if other news came out earlier in the day that affected the policy decision and thus the announcement.[6] In this case, daily changes will represent a mix of the policy announcement and the news that influenced that policy. A sufficiently narrow intradaily observation window avoids this problem.

Sometimes the researcher may be concerned that an announcement's effects may dissipate (or increase) over time. To check this possibility, the econometrician can regress the change in an asset price from shortly before the announcement to h days afterward on the announcement surprise, with the latter still measured using a narrow intradaily window. Bauer, Bernanke, and Milstein (2023) and Swanson (2021) are examples that use this strategy to measure the dynamic effects of FOMC announcements. As long as the right-hand-side surprise is correctly measured, attenuation bias is avoided; wider measurement windows for the left-hand-side variable may add noise but not bias.

For monetary policy, the event-study literature dates back to Cook and Hahn (1989), who used the raw change in the federal funds rate target as the exogenous impulse. Kuttner (2001) refined this analysis by using federal funds futures price changes on the days of the Federal Reserve's interest rate announcements (FOMC announcements) to decompose the funds rate changes into expected and unexpected components. The literature grew rapidly after this. In this survey, like Kuttner (2001), we focus on using changes in asset prices, such as federal funds futures, that are very closely tied to the monetary policy decision to measure monetary policy surprises. But there are also papers that use regressions (e.g., Romer and Romer, 2004), textual analysis (e.g., Lucca and Trebbi, 2009; Handlan, 2022), or both (Aruoba and Drechsel, 2026) to estimate the unexpected component of monetary policy announcements.

As an example, we perform a standard event-study regression. The independent variable is the surprise change in the federal funds rate target, measured as in Kuttner (2001) but using an intradaily, 30-minute window around each FOMC announcement for both the dependent and

[6] This is usually not a problem in practice, but it can happen. For example, in the 1990s, there were several cases when a weak employment report in the morning led the FOMC to cut interest rates later that same day—see Gürkaynak et al. (2005a).

independent variables.[7] The high-frequency observation window allows us to estimate

$$\Delta y_t = \alpha + \beta \Delta a_t + \eta_t \tag{7}$$

directly by OLS, as simultaneity and reverse causality are ruled out due to the monetary policy announcement surprise measured at this frequency being clearly predetermined. The sample includes all FOMC announcements from July 1991 to December 2019.[8] Note that, although there is a t subscript in regression (7) as monetary policy events have dates, the data in this regression have no time series dependence. These are theoretically uncorrelated surprises, so that the observations could be reordered without affecting the analysis. In fact, this is one of the best examples of data that are suitable for OLS analysis, albeit with heteroskedastic errors.

The results from this event study are reported in Table 1 and are standard. Federal funds rate surprises strongly affect short-term interest rates, but their effect on longer-term yields is much smaller or even zero. Stock prices fall and the dollar appreciates. The coefficients are in units of basis points (bp) per 100-basis-point surprise change in the federal funds rate, except for the S&P 500 and exchange rate responses, which are in units of percent change per 100-bp surprise change in the funds rate. The estimated constants are numerically small and statistically insignificant.

Short-term Treasury yields respond by about 40% as much as the change in the funds rate. The numbers here are not 100% because in many cases market participants expected the Fed to change rates sooner or later, and so a surprise change in the funds rate today often just brings forward an interest rate change that had been expected to occur in the next few months, anyway. The 10- and 30-year Treasury yields do not respond significantly to surprises in the federal funds rate. However, these long-term yields effectively average the response of expected interest rates over a very long 10- or 30-year horizon; Gürkaynak, Sack, and Swanson (2005b) instead studied the response of

[7] The 30-minute window begins 10 minutes before each announcement and ends 20 minutes after, following Gürkaynak et al. (2005a).

[8] There are eight regularly scheduled FOMC announcements per year. In addition, the FOMC occasionally changed the federal funds rate between regularly scheduled meetings when the economy was weakening rapidly and the committee did not want to wait until the next scheduled meeting to cut interest rates. As is standard in the literature, we drop the unscheduled FOMC announcement on 9/17/2001, as it occurred before most financial markets had opened and after they had been closed for several days following the 9/11 terrorist attacks in New York City. Finally, note that before 1994, the FOMC did not explicitly announce its decision for the federal funds rate, but implicitly announced it through the size and type of open market operation the following morning, as discussed in Kuttner (2001) and Gürkaynak et al. (2005a).

Table 1
Asset-Price Responses to Monetary Policy Surprises, 1991–2019

	Constant	FFR Target Surprise	R^2
3-month Treasury	−0.21	40.9***	0.55
	(0.15)	(7.58)	
6-month Treasury	−0.24	42.5***	0.48
	(0.17)	(8.23)	
2-year Treasury	−0.22	36.7***	0.24
	(0.29)	(9.02)	
5-year Treasury	−0.13	20.2**	0.08
	(0.32)	(8.30)	
10-year Treasury	−0.16	6.9	0.01
	(0.30)	(5.88)	
30-year Treasury	−0.06	−4.2	0.01
	(0.24)	(4.78)	
S&P 500	−0.03	−3.01**	0.16
	(0.03)	(1.22)	
USD/EUR	−0.04	1.24**	0.04
	(0.03)	(0.52)	
USD/JPY	0.01	1.21**	0.04
	(0.03)	(0.47)	

Notes: Each row reports a separate event-study regression $\Delta y_t = \alpha + \beta \Delta a_t + \eta_t$, where Δa_t is the surprise change in the federal fund rate target. Coefficients are in units of basis points per 100bp surprise; S&P 500 and exchange rate responses are in percent per 100bp surprise. Sample: FOMC announcements from July 1991 through December 2019 (250 total). Heteroskedasticity-consistent standard errors in parentheses. *, **, and *** denote significance at the .10, .05, and .01 levels. See text for details.

forward nominal interest rates and showed that surprise monetary policy tightenings significantly *lowered* far-ahead forward nominal interest rates by a moderate amount (which can also be seen in the negative coefficient for the 30-year yield in Table 1). Gürkaynak et al. (2005b) and Gürkaynak, Levin, and Swanson (2010) presented substantial evidence that this is at least partly due to FOMC tightenings better anchoring market participants' long-term inflation expectations. However, other authors (Beechey and Wright, 2009; Hanson and Stein, 2015) have emphasized a more prominent role for real rates and risk premia.

The stock market's response to monetary policy has been closely studied since at least Rigobon and Sack (2004) and Bernanke and Kuttner (2005), who estimate that the stock market falls about 3–4 percent per percentage point increase in the federal funds rate. Researchers have attempted to pinpoint the reason why, decomposing the effect into higher discount rates, higher risk premia, and lower dividend expectations. Bernanke and Kuttner (2005) found risk premia to be the main driver, but using more recently available dividend strips data, Nagel and Xu (2026) find a larger role for discount rates, though confidence intervals for the decomposition are very wide.

Many papers in the event-study literature have documented the effects of monetary policy announcements on risk premia and the market's risk appetite. Indeed, Borio and Zhu (2012) termed this the "risk-taking channel" of monetary policy. When the Fed lowers interest rates, it increases the values of bonds, stocks, and many other assets on financial institutions' balance sheets, which could plausibly increase their ability and willingness to take on risk. Bauer et al. (2023) construct an index of measures of risk appetite and show that monetary policy easings lead to an increase in this index. Hanson and Stein (2015) also find that monetary policy easings lead to declines in bond term premia, which they attribute to "reaching for yield" behavior on the part of investors—essentially an increase in risk appetite. Note that this risk-taking channel of monetary policy generally amplifies monetary policy's effects on the real economy.

The monetary policy event-study literature has also estimated the effects of U.S. monetary policy announcements on other countries' financial markets. Researchers generally find that U.S. monetary policy announcements have much bigger effects on the rest of the world than the other way around, although spillovers can go in both directions (Ehrmann and Fratzscher, 2005; Rogers et al., 2014; Miranda-Agrippino and Nenova, 2022). This asymmetry is a key piece of evidence that U.S. monetary policy is a main driver of the global financial cycle (Miranda-Agrippino and Rey, 2020).

Our event study in Table 1 considers the case of a single independent driving variable. However, one can imagine that monetary policy announcements by the Fed are multi-dimensional, in which case the econometrician might want to include several measures of monetary policy on the right-hand side of regression (7). Gürkaynak et al. (2005a) were the first to seriously consider this question, and found that monetary policy announcements at the time were well represented by two dimensions, corresponding to the surprise change in the current federal funds rate target and an orthogonal surprise driven by policy communication, now called "forward guidance".[9] Swanson

[9] Gürkaynak et al. (2005a) originally referred to forward guidance as the future "path" of monetary policy.

(2021) extended this approach to include a third dimension of monetary policy covering the Fed's large-scale asset purchases, and Altavilla et al. (2019) consider similar dimensions of ECB monetary policy announcements.

Gürkaynak et al. (2020) show that macroeconomic data releases are generally multi-dimensional as well, since those releases typically contain many components, such as expenditure categories in a GDP release. In general, only one or two of these components, the headline numbers, have corresponding survey expectations, so the econometrician does not observe the news component of much of the overall data release. Gürkaynak et al. (2020) show how those unobserved components may be estimated using latent variable methods. Boehm and Kroner (2024) apply this methodology to identify a component of monetary policy announcements that affects risky assets like stocks, credit spreads, and exchange rates without changing the yield curve at all—what Boehm and Kroner call the Fed "non-yield shock".

Recently, central banks have begun making databases of intraday asset price changes around monetary policy announcements available as a service to researchers. These include Altavilla et al. (2019) for the euro area, Braun et al. (2025) for the Bank of England, and Acosta et al. (2025) for the Federal Reserve. There is also a dataset of *daily* monetary policy surprises for a range of advanced and emerging market economies in Choi et al. (2024) and Bolhuis et al. (2024). Many of these databases are updated regularly.

Researchers have also begun to collect data for a wider range of monetary policy announcement types, such as press conferences and speeches by the Fed Chair (Swanson and Jayawickrema, 2024; Swanson, 2023; Bauer and Swanson, 2023a), FOMC minutes releases and speeches by the Federal Reserve Board Vice Chair (Swanson and Jayawickrema, 2024), speeches by ECB policymakers (Altavilla et al., 2026), and speeches by Bank of Canada policymakers (Sekkel et al., 2026). These databases are still in early stages but offer the prospect of a much richer set of events with which to conduct monetary policy event studies going forward.

Finally, although we have focused on monetary policy as a concrete event-study application in this survey, high-frequency event studies are an effective tool for studying the effects of other

The term "forward guidance" has been used to refer both to information about policymakers' *forecast* for the target rate (so-called "Delphic" forward guidance) and to a *commitment* by policymakers to follow a *particular* policy rate path ("Odyssean" forward guidance). See Campbell et al. (2012).

types of announcements as well. In fact, the event-study literature began in finance and there are numerous finance and corporate finance studies analyzing the effects of announcements such as corporate earnings releases, stock splits, takeover offers, merger announcements, and stock market index additions and removals (Fama et al., 1969).

In other macroeconomic event studies, Bianchi et al. (2023) look at how asset prices moved in short windows around tweets by Donald Trump about the Fed. Snowberg, Wolfers, and Zitzewitz (2013) examine stock market responses to perceptions of who is going to win an election, using political prediction markets. Hazell and Hobler (2024) use the 2021 Georgia senate runoff election for an event study, as this was a close election which gave Democrats control of the Senate and enabled a large fiscal stimulus. Wiegand (2025) cleverly uses the timing of the budget resolution process to understand how asset prices respond to fiscal policy news.

3 Panel Event Studies

There has been a trend over the past twenty years toward using more disaggregated, microeconomic data to analyze macroeconomic questions, and this trend has been present in the high-frequency event-study literature as well. The idea is that more detailed, microeconomic data may shed light on the transmission mechanisms of macroeconomic shocks that we would not be able to observe by looking at macroeconomic data alone. Note that the event-study identification is typically very strong in this case because macroeconomic developments are typically exogenous with respect to the microeconomic data being analyzed.

The general methodology in these studies involves using panel data where the time dimension is different events (such as monetary policy announcements) and the cross-section covers households, firms, or industries. One can also consider panels of countries as well—for example, to study monetary policy spillovers from the U.S. to the rest of the world. Let y_{it} be some attribute of the ith cross-sectional unit (such as firm investment), Δa_t be the announcement surprise (common to all units), and $\mathbf{X}_{it}$ a vector of characteristics of the unit (such as age or riskiness of a firm). The methodology involves some variant of the panel event-study regression

$$\Delta y_{it} \;=\; \alpha_i \,+\, \beta_t \,+\, \boldsymbol{\gamma}'\mathbf{X}_{it}\Delta a_t \,+\, \eta_{it}, \tag{8}$$

where α_i and β_t denote cross-sectional unit and time fixed effects, the announcement surprise Δa_t itself is not included in (8) because it is absorbed by the time fixed effects, and the parameter of interest is $\boldsymbol{\gamma}$, which describes how the announcement affects units differently depending on the characteristics $\mathbf{X}_{it}$.

Examples of panel studies of firm responses to monetary policy include Ippolito, Ozdagli, and Perez-Orive (2018), who show that the stock prices of firms with more bank loans are more sensitive to monetary policy, and Ottonello and Winberry (2020), who analyze the investment response of firms to monetary policy and find bigger responses in firms that are lower risk (have a higher distance to default). Gürkaynak, Karasoy-Can, and Lee (2022) show that the Ippolito et al. (2018) result is driven by firms' cash flow sensitivity to monetary policy, which depends on the amount and maturity of their floating-rate loans and also affects their investment response. Cloyne et al. (2023) find that younger firms adjust their investment and borrowing more in response to monetary policy than older firms, relating this to their sensitivity to collateral valuation changes.

Panel studies of financial intermediaries include English, Van den Heuvel, and Zakrajšek (2018), who study the effects of monetary-policy-induced changes to the slope of the yield curve on bank equity and bank profits, and relate this to the duration mismatch between banks' assets and liabilities. Altavilla, Gürkaynak, and Quaedvlieg (2024) find that monetary policy in the euro area transmits differentially to banks and firms of different types, with quantitative easing (QE) helping narrow spreads on loans extended by more fragile banks and to more vulnerable firms. Altavilla et al. (2024) show that monetary policy easings, which normally increase bank equity values, lowered them during the period of negative interest rates in the euro area, as rates were cut deeper into negative values.

Panel event studies can also be used to study the effects of QE on individual securities. For example, in the QE literature there is a debate as to whether QE broadly lowers the yields on all long-term bonds, or instead more narrowly lowers the yields on targeted individual securities and their closest substitutes. D'Amico and King (2013) and D'Amico et al. (2012) study the effects of quantitative easing announcements on panels of individual bonds and find evidence for both a broad duration channel and a narrow local supply channel, with local supply effects being predominant.

Finally, given a large enough cross-sectional dimension, it is even possible to perform event-study analysis with a single significant event. Foley-Fisher, Ramcharan, and Yu (2016) study the effects of the Federal Reserve's September 2011 Maturity Extension Program announcement on individual firms' stock prices. The announcement resulted in a flattening of the yield curve and led firms with long-term debt to issue more debt, expand investment, and increase employment relative to firms with only short-term debt. In the same spirit, Oehler, Horn, and Wendt (2017) study the characteristics of firms whose stock prices were most affected by the surprise Brexit vote in 2016.

4 High-Frequency Identification of VARs

The literature on structural VARs, going back to the seminal paper of Sims (1980), evolved quite separately from the event-study literature. However, high-frequency data can be used in some cases to identify the effects of structural shocks in a VAR, and this approach offers several advantages over older, lower-frequency approaches. For example, recursive ordering assumptions—that some macro variables are not affected by a shock within the month or quarter of impact—are often not plausible. Sign restrictions—that some macro variables must respond in a certain direction to a shock on impact—are often more plausible but are not very restrictive, leaving very wide ranges of estimated impulse responses. High-frequency identification is more plausible than recursive orderings and provides much tighter identification than sign restrictions.

The idea of high-frequency identification is that high-frequency observations of policy announcements are predetermined with respect to the macroeconomic variables in the VAR in equation (2), for the same reasons as discussed in the previous sections. If one of the monthly structural shocks $\boldsymbol{\varepsilon}_t$ in the VAR corresponds to this policy (such as a monetary policy shock), then the high-frequency observations can plausibly be used as an instrumental variable that is both correlated with the unobserved monthly structural shock and at the same time exogenous or orthogonal with respect to other potential structural shocks hitting the economy.

One of the earliest examples of high-frequency identification of a VAR is Faust, Swanson, and Wright (2004b). The instrumental variables framework for this identification was developed in more detail by Stock and Watson (2012, 2018). Mertens and Ravn (2013) is an early application of the Stock-Watson framework to a practical VAR, although their instrument was narratively identified rather than based on high-frequency data. The first application of the Stock-Watson framework to a monetary policy VAR using high-frequency data is Gertler and Karadi (2015), which has become a standard benchmark in the literature.

Although we will continue focusing on interest rate surprises in the discussion below, high-frequency identification of VARs is broader. For example, Känzig (2021) measures oil supply shocks from jumps in oil futures prices around OPEC meetings and uses these as impulses. This is an example where oil prices are used to measure shocks to oil supply, very much like using interest rate derivatives to measure shocks to monetary policy. High-frequency instruments may also come from further afield. Gürkaynak et al. (2026) note that China is the dominant buyer of copper and use high-frequency changes in copper futures prices around monetary policy announcements in China (which involve changes to a long list of policy instruments) to measure Chinese monetary policy surprises, as copper futures, traded in London, are liquid with a deep market that is open when policy is announced, unlike Chinese markets themselves.

4.1 External Instruments

The first step in high-frequency identification is to convert the high-frequency announcement observations into a lower-frequency instrumental variable z_t that aligns with the frequency of the VAR in equation (2). The standard approach is simply to sum up the high-frequency observations every month (for a monthly VAR) to arrive at a cumulative effect of high-frequency policy announcements each month (and analogously for a quarterly- or annual-frequency VAR).[10] In months with

[10]Some authors have tried other ways of aggregating the high-frequency observations to lower frequency, as it does seem reasonable that a shock occurring at the very end of the month might have a smaller impact effect than one occurring at the start of the month. For example, Gertler and Karadi (2015) take a moving average of high-frequency observations that takes account of when within the current or previous month each announcement occurred; however, this approach has the drawback of inducing serial correlation in the monthly instrument z_t, correlation between z_t and ε^P_{t-1}, and correlation between z_t and $\mathbf{x}_{t-1}$, and is thus

no high-frequency policy observations, the instrument z_t is set to zero. Stock and Watson (2012) refer to z_t as an *external instrument* because the data underlying z_t are from outside the variables $\mathbf{x}_t$ in the VAR, while Mertens and Ravn (2013) call z_t a *proxy* for the structural shock and the external instrument identification a *proxy VAR*.

Let ε_t^P denote the element of the structural shock vector $\boldsymbol{\varepsilon_t}$ that corresponds to the policy shock of interest, and let $\boldsymbol{\varepsilon}_t^{-P}$ denote the vector of all other structural shocks in $\boldsymbol{\varepsilon_t}$. To be a valid instrument, z_t must be both *relevant* for ε_t^P and *exogenous*, that is:

$$E\big[\varepsilon_t^P z_t\big] \;\neq\; 0 \tag{9}$$

and

$$E\big[\boldsymbol{\varepsilon}_t^{-P} z_t\big] \;=\; \mathbf{0}. \tag{10}$$

Together, these two conditions are sufficient to identify the effects of ε_t^P in the VAR. Stock and Watson (2012) note that

$$E[\mathbf{u}_t z_t] \;=\; \mathbf{S}^P E\big[\varepsilon_t^P z_t\big], \tag{11}$$

where $\mathbf{S}^P$ is the column of $\mathbf{S}$ corresponding to the policy shock. Let $\boldsymbol{\alpha} \equiv E[\mathbf{u}_t z_t]$, where $\mathbf{u}_t$ denotes the reduced-form innovations from equation (2). Then the VAR's impulse response function for the ith element of $\mathbf{x}_t$ at horizon h to the shock ε_t^P is given by

$$\mathrm{IRF}(i,h) \;=\; \frac{\mathbf{e}_i' \mathbf{B}^h \boldsymbol{\alpha}}{\mathbf{e}_P' \boldsymbol{\alpha}}\,, \tag{12}$$

where $\mathbf{B}$ is the coefficient matrix from (2) and $\mathbf{e}_i$ is a vector of zeros with a one in the ith position. The denominator of (12) is a scalar that normalizes the size of the IRF to correspond to an impact effect of one unit on the Pth variable of $\mathbf{x}_t$, which we assume is the policy variable (e.g., the short-term interest rate).

The econometrician can estimate $\hat{\mathbf{B}}$ and $\hat{\mathbf{u}}_t$ from equation (2), estimate $\hat{\boldsymbol{\alpha}} = \frac{1}{T}\sum \hat{\mathbf{u}}_t z_t$, and substitute these into (12) to estimate the impulse response function, $\widehat{\mathrm{IRF}}(i,h)$. However, rather than estimating $\hat{\mathbf{u}}_t$ and then $\hat{\boldsymbol{\alpha}}$ in two steps, Stock and Watson (2018) recommend estimating $\boldsymbol{\alpha}$

not recommended.

directly from the regression

$$\mathbf{x}_t = \mathbf{B}\mathbf{x}_{t-1} + \boldsymbol{\alpha} i_t + \tilde{\mathbf{u}}_t \tag{13}$$

using two-stage least squares (2SLS), where i_t denotes the policy variable, and with z_t as the instrument for i_t. In regression (13), the point estimate of $\boldsymbol{\alpha}$ is the same as in the two-step procedure, but the standard errors are wider and correctly account for the sampling uncertainty about the residuals $\hat{\mathbf{u}}_t$ that the two-step procedure ignores.

There are several approaches to computing standard errors around the $\widehat{\mathrm{IRF}}(i, h)$ estimates in (12). The most common is the wild bootstrap approach used by Mertens and Ravn (2013) and Gertler and Karadi (2015), as these are benchmark papers in the literature. The main advantage of the wild bootstrap is that it is robust to heteroskedasticity of arbitrary form. However, Jentsch and Lunsford (2019) point out that this approach is not asymptotically valid because it fails to replicate the covariance between VAR residuals and the external instrument and argue instead for a moving blocks bootstrap procedure. Montiel Olea, Stock, and Watson (2021) propose a delta method approach, simply leveraging the fact that the estimates of $\mathbf{B}$ and $\boldsymbol{\alpha}$ in (12) are jointly asymptotically normal. An additional important issue to consider when computing standard errors of these estimates is whether the instrument z_t is weak, which we discuss further in Section 5, below.

In practice, it is often the case that the low-frequency data for the variables $\mathbf{x}_t$ in the VAR are available over a longer sample than the high-frequency data for the instrument z_t. An interesting feature of the external instruments approach is that the econometrician can estimate $\mathbf{B}$ in (2) over the longer sample for which $\mathbf{x}_t$ is available, and estimate $\boldsymbol{\alpha}$ in (13) over the shorter sample for which z is available. In fact, this is generally desirable as it reduces the estimation uncertainty surrounding $\hat{\mathbf{B}}$ and thus $\widehat{\mathrm{IRF}}(i, h)$, making this a common approach in the literature (e.g., Gertler and Karadi, 2015; Eberly, Stock, and Wright, 2019; Bauer and Swanson, 2023b).

4.2 Internal Instruments

An alternative way to perform high-frequency identification of a VAR is an *internal instruments* approach, in which the set of variables $\mathbf{x}_t$ is expanded to include the instrument z_t. The expanded

VAR is then estimated in the standard way, and identified by imposing a recursive (Cholesky) ordering of the variables with z_t ordered first. Because the instrument z_t satisfies the exogeneity condition (10), the other structural shocks $\boldsymbol{\varepsilon}_t^{-P}$ have no contemporaneous effect on z_t, which justifies ordering it first. Meanwhile, the ordering of the other variables in the VAR does not affect their responses to a shock to z_t, as discussed in Christiano, Eichenbaum, and Evans (1999). Thus, the recursive ordering assumption is clearly valid for estimating the IRFs to a shock to z_t.

The internal instruments approach has the advantage that all of the econometrician's standard toolkit for estimating, identifying, and analyzing recursive VARs carries over immediately. The approach is thus very simple. Paul (2020) and Plagborg-Møller and Wolf (2021) show that the methods of external and internal instruments are equivalent to each other (and to local projections) asymptotically. However, in finite samples, Li, Plagborg-Møller, and Wolf (2024) find that internal instruments produce estimates with greater dispersion.

The primary disadvantage of the internal instruments approach is that the extended VAR can only be estimated over the sample for which *both* the original variables $\mathbf{x}_t$ and the instrument z_t are available. As discussed in the previous section, it is often the case in practice that z_t is available over a shorter sample than the lower-frequency variables in $\mathbf{x}_t$. In this case, the internal instruments approach essentially throws away the additional data, increasing the estimation uncertainty surrounding the VAR dynamics $\mathbf{B}$.

In principle, the high-frequency instrument z_t is not only exogenous in the sense of (10), but is also unpredictable with information from previous months, including the lagged values $\mathbf{x}_{t-1}$. This implies that the first row of $\mathbf{B}$ in (2) is identically zero. Some researchers, such as Paul (2020) and Jarociński and Karadi (2020), impose this restriction in their internal instruments implementations, but most do not. If the restriction is correct, imposing it improves econometric efficiency, so it is recommended, although the econometrician must be careful because instrument predictability is sometimes present in practice, an issue that we discuss further in Section 5, below.

4.3 Local Projections

Since Jordà (2005), a popular way of estimating impulse response functions is to estimate them directly in a local projection (LP) framework instead of a VAR. In LP, the econometrician estimates

$$\mathbf{x}_{t+h} \;=\; \boldsymbol{a}^{(h)} + \mathbf{B}^{(h)}\,\mathbf{x}_{t-1} + \boldsymbol{\gamma}^{(h)}\varepsilon_t^P \;+\; \boldsymbol{\eta}_t^{(h)} \tag{14}$$

by OLS separately for each horizon h, assuming the policy shock ε_t^P is observed.[11] The coefficients $\boldsymbol{\gamma}^{(h)}$ then trace out the IRFs for the effects of the shock ε^P on $\mathbf{x}$. In practice, the policy shock ε_t^P is typically not observed, but high-frequency identification can be used to estimate $\boldsymbol{\gamma}^{(h)}$ via instrumental variables. Instead of (14), the econometrician estimates the LP-IV specification

$$\mathbf{x}_{t+h} \;=\; \boldsymbol{a}^{(h)} + \mathbf{B}^{(h)}\,\mathbf{x}_{t-1} + \boldsymbol{\gamma}^{(h)}i_t \;+\; \boldsymbol{\eta}_t^{(h)} \tag{15}$$

by 2SLS, where i_t is the policy variable (e.g., the short-term interest rate), using z_t as an instrument.

Note that many authors directly include z_t instead of i_t on the right-hand side of (15), effectively assuming that z_t and ε_t^P are the same. We discourage this practice for several reasons. First, it is typically not true that z_t and ε_t^P are the same—for example, in the case of monetary policy shocks, there is often substantial news about monetary policy each month that is not released around FOMC announcements—so z_t is usually best thought of as an instrument that is positively correlated with ε_t^P but not exactly equal to it. Second, thinking of z_t as an instrument forces the econometrician to confront the issue of whether the instrument is potentially weak, which is often a very important issue in practice, as we discuss in Section 5, below. Third, 2SLS standard errors for $\boldsymbol{\gamma}^{(h)}$ correctly account for the additional uncertainty that arises from z_t not being perfectly correlated with i_t. For all of these reasons, the LP-IV specification (15) is almost always better than including z_t in (15) directly.[12]

Both VARs and LPs are valid ways of estimating IRFs. Plagborg-Møller and Wolf (2021) show that VAR and LP IRF estimates are asymptotically equivalent as long as enough lags are included

[11] For a Bayesian approach to LP, see Ferreira, Miranda-Agrippino, and Ricco (2025).

[12] The main exception is the case when the instrument z_t may affect the variables $\mathbf{x}_t$ through multiple channels rather than through the single policy variable i_t. In that case, the LP-IV specification (15) would only pick up one out of the potentially many channels through which z_t affects $\mathbf{x}_t$.

in each case. Jordà (2005) argues that LPs are more robust than VARs to model misspecification, and Montiel Olea et al. (2026a) prove this claim formally under certain assumptions. However, LPs are generally more volatile and have wider standard errors than IRFs from a VAR, implying a bias-variance tradeoff in practice (Li et al., 2024; Montiel Olea et al., 2026b; Baumeister, 2026).[13]

One of the primary advantages of LPs is that they are extremely flexible and are thus very easy to extend to more general specifications than (14) and (15). For example, the econometrician can easily divide shocks into positive and negative realizations and allow for asymmetries in their effects on $\mathbf{x}_t$, or allow for nonlinear interaction effects between the policy shock and the state variables $\mathbf{x}_{t-1}$. However, with this flexibility comes the potential for abuse in terms of optimizing the specification to achieve the lowest p-values, a practice that we obviously strongly discourage.

The primary disadvantage of LPs is the same as for internal instruments: they can only be estimated over the sample for which *both* the original variables $\mathbf{x}$ and the instrument z are available. In practice, $\mathbf{x}$ is typically available over a longer sample than z, but LPs essentially discard that additional data, increasing estimation uncertainty.

4.4 Identification by Heteroskedasticity

A different but related approach to identification in a structural VAR is identification by heteroskedasticity (Rigobon, 2003; Rigobon and Sack, 2004; Wright, 2012). This requires the variances of structural shocks to change over time. If there are days (like FOMC days) when the variance of the monetary policy shock is elevated, while the variances of other shocks are unchanged, and if our VAR (2) is at daily frequency, we can let $\mathbf{\Omega}^E$ and $\mathbf{\Omega}^{NE}$ denote the variance-covariance matrix of the reduced-form residuals $\mathbf{u}_t$ on event and non-event days, respectively. Then

$$\mathbf{\Omega}^E - \mathbf{\Omega}^{NE} = \mathbf{S}^P \mathbf{S}^{P\prime} \tag{16}$$

[13]IRF point estimates from a VAR often have lower mean-squared errors for the true IRF (Li et al., 2024; Montiel Olea et al., 2026b), but the standard errors around those IRFs are often distorted and can lead to incorrect statistical inference relative to LPs (Montiel Olea et al., 2026a). De Graeve and Westermark (2025) and Baumeister (2026) argue for using long-lag VARs to address some of these issues.

provides an alternative way of identifying $\mathbf{S}^P$ using high-frequency data.[14] This has the advantage that it does not require the policy shock to be the only one in the event window, but the disadvantage of being less efficient.

A subtle but important issue, discussed by Gürkaynak et al. (2020), is that identification by heteroskedasticity identifies the response to an event, such as a monetary policy announcement or data release, but not a particular component of that event, such as the surprise in the federal funds rate target or headline inflation. This mechanically produces a larger estimated effect but can leave some ambiguity as to its exact cause.

4.5 Empirical Illustration

We conclude this section by illustrating the high-frequency identification method in a simple 13-lag trivariate VAR with year-on-year growth rates of the U.S. industrial production (IP) and the consumer price index (CPI), and the two-year zero-coupon Treasury yield of Gürkaynak et al. (2007). Following Gertler and Karadi (2015), Ramey (2016), and Bauer and Swanson (2023b), we use the two-year Treasury yield rather than the federal funds rate as the measure of monetary policy because it was essentially unconstrained by the zero lower bound from 2009–15 (Swanson and Williams, 2014) and because it allows for a more prominent role for unconventional monetary policies like forward guidance, which affect the two-year yield but not the current federal funds rate.

We use two identification schemes for monetary policy. The first is the usual Cholesky ordering with monetary policy last, imposing the restriction that monetary policy does not affect the CPI and IP within the month. The second is high-frequency identification with the orthogonalized (discussed below) shock of Bauer and Swanson (2023b) around FOMC announcements as the instrument, which utilizes changes in short rates up to a year in maturity and therefore captures some forward guidance effects as well. The impulse response estimates, along with 90 percent standard error bands from the wild bootstrap are shown in the left column of Figure 1, while the panels in the right column provide the same IRFs with the weak-instrument-robust confidence intervals of Montiel Olea et al. (2021) discussed further in Section 5, below.

[14] See Wright (2012) for details.

Fig. 1. VAR Impulse Responses to Monetary Policy Shock

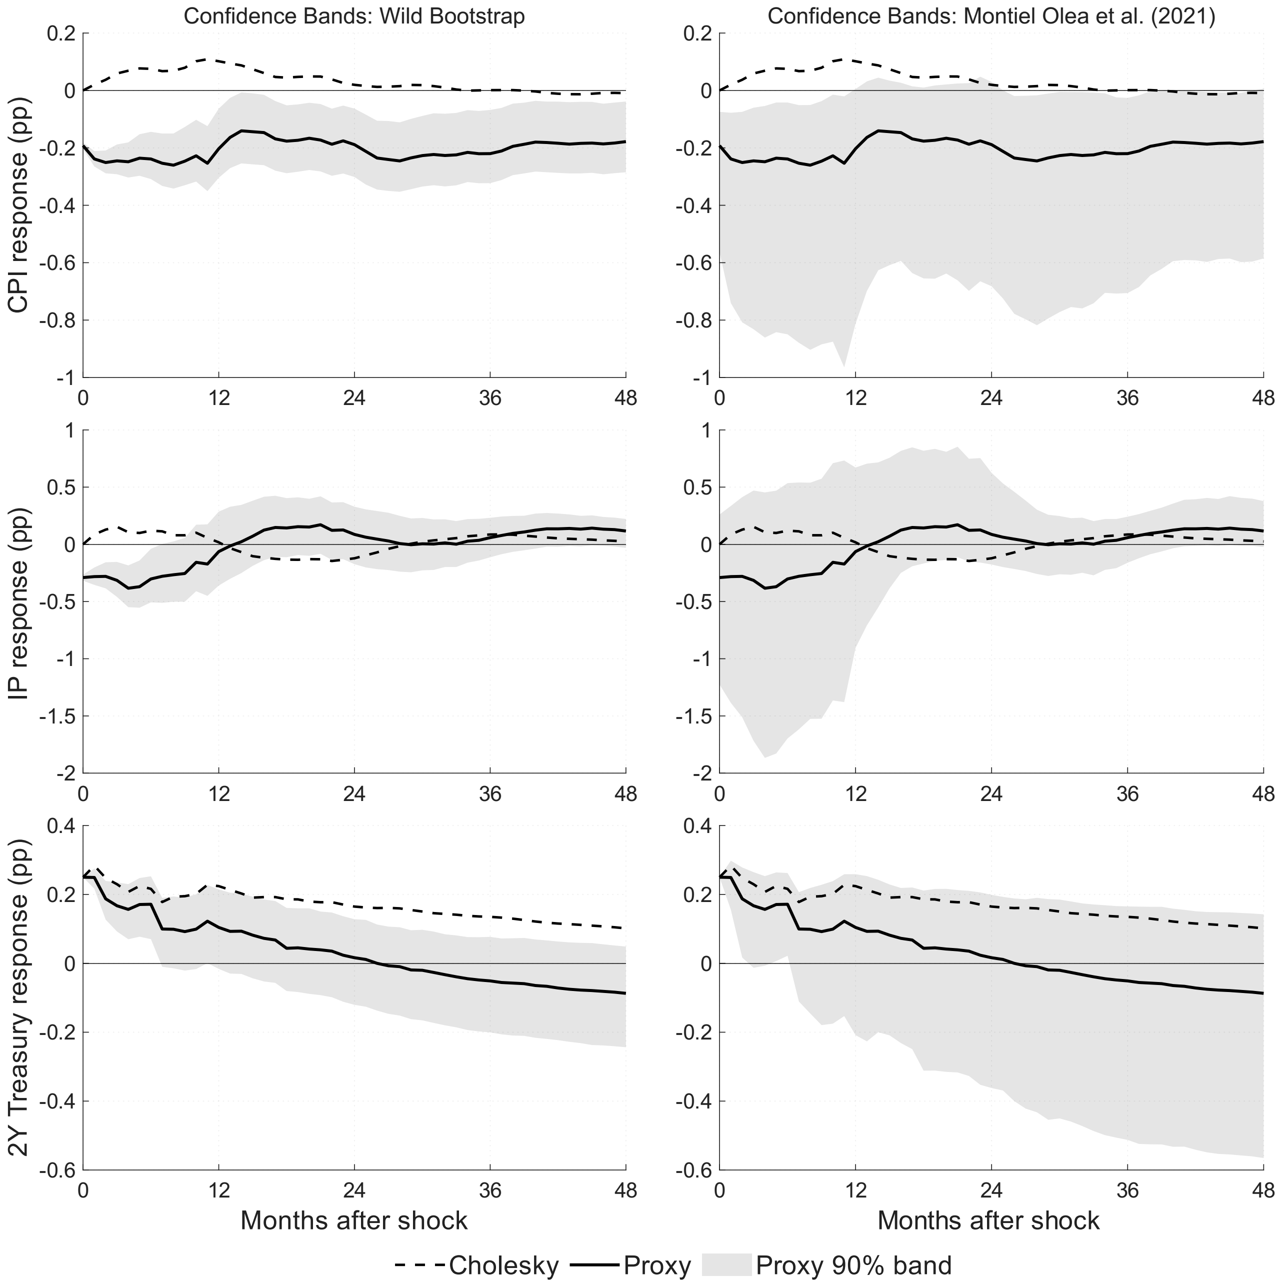

Note: This figure reports the impulse responses from a VAR with 13 lags identified alternatively by Cholesky decomposition and the method of external instruments, estimated from 1973:01 to 2019:12. The instrument is available for a shorter sample. The variables are the U.S. IP and CPI in annual growth rates and the two-year yield, in that order. Bauer and Swanson (2023b) orthogonalized monetary shocks are the external instrument. For the proxy-VAR, in the panels in the left column confidence intervals are formed by the wild bootstrap as in Bauer and Swanson and in the right column by the weak-identification-robust procedure of Montiel Olea et al. (2021), both with 90 percent coverage.

For the monetary policy shocks with the Cholesky identification, there is a price puzzle (top row), with inflation initially increasing in response to a monetary policy tightening. This is avoided by the use of the external instrument. Similarly, a surprise tightening of 25 basis points in the two-year yield leads to an immediate drop in industrial production (middle row) when using the instrument and the effect remains significantly negative for about 6 months. Comparing the left and right columns, the weak-instrument-robust confidence intervals of Montiel Olea et al. (2021) are much wider, but it remains the case that the monetary policy shock significantly lowers inflation for some time.

The identification assumption behind the external-instrument monetary policy shock is that the high-frequency policy surprise cannot be a function of industrial production and the CPI. This is a more credible assumption than the Cholesky one, that monetary policy cannot affect industrial production and the CPI within the month. Even if the Cholesky assumption were true (and we see that the data reject this), there would be no way of verifying that without using a more credible identification scheme like the high-frequency instrument and comparing the results.

Note that it is not only the high-frequency measurement, but doing that measurement around the policy decision, that identifies the monetary policy shock. Using the changes in short rates in non-policy event windows would not have identified monetary policy effects. For example, if one takes the changes in federal funds futures rates around CPI announcements and uses these as instruments, the impulse is news about the CPI, not monetary policy.

An active strand of literature asks whether monetary policy surprises may also be affected by problems such as containing information that violates the exclusion restriction, and whether these behave sufficiently like surprises in being unforecastable. There are indeed various ways in which high-frequency identification may not be straightforward. We now turn to complications and possible problems of this identification approach.

5 Complications and Pathologies

The methods above have been described for the case of about ideal econometric conditions. However, there are many ways in which these conditions could be more complicated in practice, such as asymmetric or nonlinear interaction effects of announcements, weak instruments, multiple instruments, instruments that are predictable based on past information, and information effects. Here we consider these issues in turn.

5.1 Nonlinearities

One possible issue in linear specifications like (1) and (2) is that the announcement Δa_t may in fact have asymmetric effects for increases vs. decreases, or may interact with the control variables $\mathbf{w}_t$ or VAR variables $\mathbf{x}_{t-1}$. In high-frequency event studies like (1), it is typically straightforward for the econometrician to add decompositions of Δa_t, nonlinear interaction terms, or threshold effects into the regression and estimate the corresponding coefficients—see, e.g., Altavilla et al. (2019), who investigate whether the high-frequency effects of euro-area monetary policy announcement surprises on financial markets are asymmetric and do not reject symmetry.

For VARs like (2), it is typically much more difficult to add nonlinearities like these, but LP specifications as in Section 4, above, are generally very amenable to these complications—see, e.g., Tenreyro and Thwaites (2016), Ramey and Zubairy (2018), and Barnichon and Matthes (2018), who consider asymmetric effects of monetary policy shocks and interactions between monetary or fiscal policy shocks and the state of the economy (although these authors use narrative instruments rather than high-frequency ones in their analyses).

Of course, a significant concern with extending the specifications (1) and (15) to include additional nonlinear terms is the loss of econometric power. High-frequency announcement surprises are often small already, which can make estimation challenging, and dividing them into smaller subsets only increases those challenges. This is particularly true for macroeconomic specifications like (15), for which weak instruments are often a significant problem.

5.2 Weak Instruments

It is often the case in practice that an instrument is exogenous, satisfying equation (10), but may not be very relevant, so that the expectation in equation (9) is small. For example, high-frequency changes in short-term interest rates around FOMC announcements are typically only a few basis points per announcement, so the econometrician may reasonably be concerned that the explanatory power of these high-frequency interest rate changes for monthly innovations in i_t is small. A large econometrics literature on weak instruments shows that, even if the instrument z_t is exogenous, if its relevance is weak then coefficients may be biased and standard errors seriously understated (Staiger and Stock, 1997; Stock, Wright, and Yogo, 2002; Stock and Yogo, 2005).

The question naturally arises: how does the econometrician know if an instrument z_t is weak or strong? Stock and Yogo (2005) provide formal tests and critical values for instrument weakness that depend on the specifics of the econometric model, but a standard rule of thumb in the literature is that if the marginal F-statistic for the instrument(s) in the first-stage regression is less than 10, then weak instruments are more likely to be a concern (Staiger and Stock, 1997; Stock and Watson, 2018).[15] Note that, for a single instrument, this corresponds to a t-statistic in the first-stage regression that is less than about 3.16.

If the first-stage F-statistic is substantially above 10, then in practice, most studies ignore issues related to weak instruments and report the usual standard error bands around their estimates, as described in the previous section. However, if the first-stage F-statistic suggests the instrument may be weak, the econometrician has essentially two options. The first is simply to try to improve the instrument or find a better one. In the case of monetary policy announcements, Swanson and Jayawickrema (2024), Bauer and Swanson (2023b), and Swanson (2024) improve the instrument by expanding the set of monetary policy announcements to include post-FOMC press conferences, speeches, and Congressional testimony by the Fed Chair as well as official FOMC announcements.

Swanson and Jayawickrema (2024) show that speeches and testimony by the Fed Chair are even more important than FOMC announcements for interest rates at all but the very shortest maturities, so that extending the set of monetary policy announcements in this way leads to a much stronger high-frequency instrument z_t. Doing the same exercise for the euro area, Altavilla et al. (2026) show that the high-frequency instrument is similarly improved, although the ECB's official Governing Council announcement remains more important than the ECB President's speeches.

[15] Note that it is the *marginal* F-statistic for the instruments (the squared t-statistic for the case of a single instrument) that is the relevant metric; many authors incorrectly report that the *omnibus* F-statistic for their first-stage regression is larger than 10, but even for an instrument uncorrelated with the variable of interest, the omnibus F-statistic can be large if other control variables included in the first stage have substantial explanatory power.

(The reason for the difference between the two central banks in this regard remains an interesting open question.)

The second approach available to the econometrician is to report standard errors that are robust to weak instruments. For the case of a single instrument identifying a single shock, Montiel Olea et al. (2021) provide a relatively simple formula that is asymptotically correct whether the instrument is weak or strong. Of course, these weak-instrument-robust standard errors are typically and sometimes dramatically wider than conventional standard errors in finite samples, but that is the appropriate statement of the econometrician's uncertainty (Staiger and Stock, 1997; Montiel Olea et al., 2021).

5.3 Multiple Instruments

Our instrumental variables discussion in Section 4 focused on the case of a single instrument z_t, both for simplicity and because that is the most common case in practice. However, the econometrician sometimes has more than one instrument available, either for a single policy variable or for multiple policy variables.

The case where the econometrician has multiple instruments available for a single policy variable is straightforward. 2SLS in the external instruments regression (13) or LP-IV (15) can handle any number of instruments for i_t. In fact, Gertler and Karadi (2015) considered multiple instruments for the monetary policy shock in their benchmark analysis, although they ultimately settled on a single instrument for their baseline specification. However, the econometrician should take care to ensure that any additional instruments used in their analysis are sufficiently strong, as the biases caused by weak instruments depend on the average strength of the instruments (Stock et al., 2002). It is better to have just one strong instrument than two instruments, one of which is strong and one of which is weak.

In other cases, the econometrician may be interested in the effects of multiple policy variables—e.g., the federal funds rate, forward guidance, and large-scale asset purchases—and have a single instrument for each. If the instruments are orthogonal to each other, as in Swanson's (2024) study of multiple Fed monetary policy tools, then this is again relatively straightforward, since each 2SLS regression (13) or (15) can be run separately for each policy shock. However, when the instruments

z_t are correlated with each other, as in the fiscal policy studies of Mertens and Ravn (2013) and Mertens and Montiel Olea (2018), then the regressions cannot be run independently and there are potential complications, discussed in Arias, Rubio-Ramirez, and Waggoner (2021).

A final caveat is that Stock and Watson (2012) catalog many variables that are often used as instruments for fundamentally different shocks hitting the economy (fiscal policy, monetary policy, oil, productivity, etc.) and find that they are surprisingly correlated with each other across different shock types. This is a concern because we typically think of different structural shocks as being primitive underlying causes of economic fluctuations, and hence mutually independent (Bernanke, 1986; Ramey, 2016).

5.4 Predictable Instruments

One of the most appealing features of high-frequency identification is that high-frequency changes in interest rates (or other asset prices) around announcements are very plausibly unpredictable based on any information predating the announcement. If not, then there would be a trading strategy that market participants could exploit to earn profits on average, which should drive any such predictability away.

Despite this powerful argument, several authors have found that interest rate changes around central bank announcements are in fact correlated with information predating those announcements (Cieslak, 2018; Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023a,b; Sastry, 2026). This finding raises the obvious question as to what could be driving that correlation, to which there are four basic answers: (i) the finding is a statistical fluke and the interest rate changes actually were unexpected; (ii) the correlation represents a real predictability that traders somehow failed to exploit, forgoing a profitable trading opportunity; (iii) the correlation represents a real predictability but is not exploited by traders because it reflects a premium for bearing risk; and (iv) the interest rate changes really were unexpected, but because of slow learning about regime changes, there is correlation ex post (Timmermann, 1993; Bauer and Swanson, 2023a,b).

It is difficult to rule out any of these explanations definitively, but (i) is statistically unlikely and (ii) seems unlikely because financial firms and traders devote tremendous resources to trading strategies in interest rate markets and are unlikely to ignore profitable opportunities. Explanation

(iii) is somewhat more plausible, but the risk premia involve remarkable Sharpe ratios. Explanation (iv) seems the most plausible. The idea is that when there is an unobserved regime change, such as in the parameters of the Fed's monetary policy reaction function, it takes investors time to learn that the change has occurred. During that time, investors' forecast errors can be serially correlated and correlated with the data even though those forecasts were perfectly rational (and forecast errors unpredictable) *ex ante* at each date. Bauer and Swanson (2023a,b) provide a simple model of this process and present evidence that in fact the Fed's responsiveness to economic news has increased over time, consistent with this learning story that they label the "Fed Response to News" channel.

The next question is, what should the econometrician do about the correlation between high-frequency interest rate changes and the data? The answer depends on the application. For high-frequency identification of a VAR, as in Section 4, the correlation between z_t and macroeconomic and financial market data suggests that the exogeneity condition (10) is likely to be violated. Regardless of the reason for this violation, 2SLS coefficient estimates will be biased because z_t is correlated with the residual $\tilde{\mathbf{u}}_t$ in (13). To address this problem and reduce or eliminate the bias, Bauer and Swanson (2023b) recommend orthogonalizing the instrument z_t with respect to the correlated variables, resulting in a cleaner instrument $z_t^{\perp}$ that is more likely to satisfy the exogeneity condition (10).[16]

For high-frequency event-study regressions like those in Sections 2 and 3, the high-frequency interest rate changes often do not need any correction. For example, in the case of explanation (iv), the high-frequency changes in interest rates and other asset prices around each announcement are true surprises and can be studied in regressions like (7) without any adjustment. Bauer and Swanson (2023b) demonstrate the validity of these high-frequency regressions in their simple model and show that orthogonalizing the high-frequency changes would in fact reduce the econometric power of regressions like (1) because it removes variation in the independent variable that is systematically related to the dependent variable.

Similarly, in the case of explanation (i), the high-frequency changes in interest rates and other asset prices are again true surprises, and regressions like (1) are valid without any adjustment.

[16] They also note that a downside of this correction is that it typically weakens the relevance of the instrument, making the potential weak-instrument problem substantially worse.

However, if explanations (ii) or (iii) are driving the correlation in the data, then it is not clear how to correct them for event studies like (1). For example, if short-term interest rates have large risk premia around FOMC announcements, does the econometrician think changes in those risk premia are related or unrelated to changes in the dependent variable in (1)? The answer determines whether the interest rate changes should be corrected or not. A similar concern arises for the case of explanation (ii). Unfortunately, there is no simple answer for event-study regressions like (1) that applies in every situation for every variable, but if the econometrician views explanation (iv) as the most likely, then it is safe to run the event studies as usual without any adjustment to the included variables, as discussed above.

5.5 Information Effects

An issue closely related to high-frequency instrument predictability is the "Fed Information Effect". When the Federal Reserve raises interest rates, standard macroeconomic models and VARs predict that GDP, employment, and inflation should all fall over the next several months or quarters in response. However, regressions of changes in the monthly Blue Chip consensus forecast (an average of professional forecasts) for these variables on high-frequency interest rate changes around FOMC announcements often produce coefficients that are positive, and sometimes statistically significantly so.

In order to explain this puzzling result, several researchers have argued that: a) the Fed is a better forecaster than the private sector; b) when the Fed raises interest rates, the private sector infers from this that the economy must be stronger than they thought; and c) the private sector then *increases* rather than decreases its forecasts for those variables as a result (Romer and Romer, 2000; Campbell et al., 2012; Nakamura and Steinsson, 2018). This story is known as the "Fed Information Effect".

This issue is related to instrument predictability because it implies that the exogeneity condition (10) is likely to be violated. High-frequency interest rate changes in the instrument z_t would contain information about current output and/or inflation, biasing the estimate of $\boldsymbol{\alpha}$ in (13).

If the Fed Information Effect is true, the instrument z_t can be corrected in much the same way as for predictable interest rate changes, above. Miranda-Agrippino and Ricco (2021) propose orthogonalizing z_t with respect to the Fed's internal "Greenbook" or "Tealbook" forecasts just before each FOMC meeting, with the idea of projecting out any information about these forecasts that might be contained in the interest rate. The orthogonalized instrument $z_t^{\perp}$ is then more likely to satisfy the exogeneity condition (10). Thus, the problems and solutions for high-frequency identification are largely the same as for the predictable instruments case, above, except that instead of orthogonalizing with respect to recent publicly available macroeconomic and financial data, the econometrician must orthogonalize with respect to internal Federal Reserve forecast data, which are typically available with a five-year lag.

However, there is a substantial debate in the literature about whether a Fed Information Effect actually exists in the data. For example, Bauer and Swanson (2023a) show that the Fed's internal forecasts have almost exactly the same root-mean-squared forecast errors as the Blue Chip consensus forecast, and Gamber and Smith (2009), Edge and Gürkaynak (2010), Paul (2020), and Hoesch, Rossi, and Sekhposyan (2023) likewise find that, to the extent that the Fed ever had a forecasting advantage, it has disappeared over time and was essentially zero by the 1990s. These results call into question the first part of the Fed Information Effect story.

Second, Faust, Swanson, and Wright (2004a) explicitly test whether monetary policy surprises could improve private sector forecasts of upcoming data releases and whether they in fact led the private sector to revise its forecasts for those upcoming data releases. The answer to both questions was negative. Bauer and Swanson (2023a) emailed the chief economists of all the professional forecasting firms in the Blue Chip Consensus survey and asked them directly how they revise their forecasts in response to Fed monetary policy announcements. The answer to that survey was overwhelmingly in the direction predicted by standard macroeconomic models and VARs, and not in the direction predicted by the Fed Information Effect. These results call into question the second and third parts of the Fed Information Effect story.

Finally, Bauer and Swanson (2023a) show that the Fed Information Effect regressions that produce the puzzling results suffer from omitted-variable bias, because the high-frequency monetary policy surprises are correlated with recent macroeconomic and financial data, as discussed in the previous section. Re-running these regressions with the omitted macroeconomic and financial data included on the right-hand side shows that the regression R^2 is dramatically higher, the coefficient standard errors are substantially smaller, and the signs of the coefficients on the high-frequency monetary policy surprises are *reversed*, making them consistent with standard macroeconomic models and VARs, rather than the Fed Information Effect.[17]

For all of these reasons, the evidence for a large, consistent Fed Information Effect is suspect. (Although note that, Miranda-Agrippino and Ricco (2023) show that the proxy-VAR may have major problems even if the first-stage fit—the genuine information about the state the FOMC reveals—is minor.) However, an alternative version of the Fed Information Effect story allows the effect to vary over time, so that it might be small on average but could be important once in a while. Jarociński and Karadi (2020) and Jarociński (2024) provide estimates of this time-varying information effect based on the comovement between stock prices and short-term interest rates and decompose high-frequency monetary policy surprises into "pure monetary policy" and "information" shock components. In this case, the estimated pure monetary policy shocks are cleaned of the information component and can be used in high-frequency identification, since they satisfy the exogeneity condition (10). Gürkaynak et al. (2021) show that many of the "information" events have path or forward guidance components that have the opposite sign of the surprise in the immediate policy setting so would explain the perverse comovement, but it may also be that the measured path component itself is driven by the information effect.

In practice, whether the econometrician orthogonalizes the high-frequency instrument z_t with respect to recent macroeconomic and financial data or with respect to the Fed's internal forecasts makes little difference. To a first approximation, both orthogonalizations achieve the same effect and produce a cleaner instrument $z_t^{\perp}$ that is more likely to satisfy the exogeneity condition (10).[18] Thus, the econometrician generally doesn't need to take a stand on whether the Fed Information Effect is true or not, so long as they acknowledge the potential predictability problem and clean the high-frequency instrument as appropriate.

[17] Then again, Ricco and Savini (2025) also present an elegant, model based breakdown of the problem and make an argument in favor of the information effect.

[18] However, both orthogonalizations also generally reduce the relevance of the instrument, making weak instruments potentially a bigger concern.

6 Conclusions and Open Questions

Many data releases and policy announcements can be viewed as natural experiments within a sufficiently narrow window of time around those announcements. High-frequency identification has long exploited this fact to get causal estimates of the effects of economic data releases or policy announcements on financial market variables using event-study regressions. But over the more recent past, the literature has gone much further than this. High-frequency identification has been used in panel event studies with microeconomic data to shed light on the transmission mechanism of shocks. It has also been incorporated into structural VARs and local projections, where it is arguably the dominant form of identification of the effects of monetary policy and many other shocks. Thus, the technique is being extended beyond event studies to also estimate the effects of policy changes on lower-frequency macroeconomic variables. Although we have taken U.S. monetary policy surprises as our running example, the same event-study approach can be and has been used with great success for other countries and for other types of shocks, including oil price shocks, other commodity price shocks, fiscal policy shocks, tweets by President Trump, and Brexit.

We have surveyed here the main ideas and applications of these techniques, and the main complicating factors that can make their application difficult. Open questions remain. The fact that high-frequency monetary policy surprises are correlated with data that precede those surprises is important and the reasons remain contentious. Similarly, there is a debate over whether monetary policy surprises contain an informational component that changes the private sector's views about the current state of the economy. Recent research has focused on expanding the set of events around which high-frequency surprises are computed, increasing instrument relevance and power. There are also continuing debates about the relative benefits of local projections vs. VARs, different methods of computing impulse response function standard errors, and how best to test for and handle weak instruments. High-frequency identification remains an active and exciting research area.

References

Acosta, Miguel, Andrea Ajello, Michael Bauer, Francesca Loria, and Silvia Miranda-Agrippino, 2025, Financial market effects of FOMC communication: Evidence from a new event-study database, Federal Reserve Bank of San Francisco Working Paper 2025-30.

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

Altavilla, Carlo, Refet S. Gürkaynak, Luc Laeven, and Thilo Kind, 2026, Monetary transmission with frequent policy events, *Journal of Monetary Economics* forthcoming.

Altavilla, Carlo, Refet S. Gürkaynak, and Rogier Quaedvlieg, 2024, Macro and micro of external finance premium and monetary policy transmission, *Journal of Monetary Economics* 147, 103634.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega, 2003, Micro effects of macro announcements: Real-time price discovery in foreign exchange, *American Economic Review* 93, 38–62.

Arias, Jonas, Juan Rubio-Ramirez, and Daniel Waggoner, 2021, Inference in Bayesian proxy-SVARs, *Journal of Econometrics* 225, 88–106.

Aruoba, S. Borağan, and Thomas Drechsel, 2026, Identifying monetary policy shocks: A natural language approach, *American Economic Journal: Macroeconomics* forthcoming.

Barnichon, Regis, and Christian Matthes, 2018, Functional approximation of impulse responses, *Journal of Monetary Economics* 99, 41–55.

Bauer, Michael D., Ben S. Bernanke, and Eric Milstein, 2023, Risk appetite and the risk-taking channel of monetary policy, *Journal of Economic Perspectives* 37, 77–100.

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

Bauer, Michael D., and Eric T. Swanson, 2023b, A reassessment of monetary policy surprises and high-frequency identification, *NBER Macroeconomics Annual* 37, 87–155.

Baumeister, Christiane, 2026, Comment on 'Local projections or vector autoregressions? A primer for macroeconomists', *NBER Macroeconomics Annual* 40, 153–171.

Beechey, Meredith J., and Jonathan H. Wright, 2009, The high-frequency impact of news on long-term yields and forward rates: Is it real?, *Journal of Monetary Economics* 56, 535–544.

Bernanke, Ben S., 1986, Alternative explanations of the money-income correlation, *Carnegie-Rochester Conference Series on Public Policy* 25, 49–99.

Bernanke, Ben S., and Kenneth N. Kuttner, 2005, What explains the stock market's reaction to Federal Reserve policy?, *Journal of Finance* 60, 1221–1257.

Bianchi, Francesco, Roberto Gómez-Cram, Thilo Kind, and Howard Kung, 2023, Threats to central bank independence: High-frequency identification with twitter, *Journal of Monetary Economics* 135, 37–54.

Boehm, Christoph, and T. Niklas Kroner, 2024, Monetary policy without moving interest rates: The Fed non-yield shock, NBER Working Paper 32636.

Bolhuis, Marijn A, Sonali Das, and Bella Yao, 2024, A new dataset of high-frequency monetary policy shocks, IMF Working Paper 2024/224.

Borio, Claudio, and Haibin Zhu, 2012, Capital regulation, risk-taking and monetary policy: A missing link in the transmission mechanism?, *Journal of Financial Stability* 8, 236–251.

Braun, Robin, Silvia Miranda-Agrippino, and Tuli Saha, 2025, Measuring monetary policy in the UK: The UK monetary policy event-study database, *Journal of Monetary Economics* 149, 103645.

Campbell, Jeffrey, Charles Evans, Jonas Fisher, and Alejandro Justiniano, 2012, Macroeconomic effects of Federal Reserve forward guidance, *Brookings Papers on Economic Activity* 1–80.

Choi, Sangyup, Tim Willems, and Seung Yong Yoo, 2024, Revisiting the monetary transmission mechanism through an industry-level differential approach, *Journal of Monetary Economics* 145, 103556.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, 1999, Monetary policy shocks: What have we learned and to what end?, in John B. Taylor, and Harald Uhlig, eds., *Handbook of Macroeconomics*, volume 1, 65–148 (Elsevier).

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

Cloyne, James, Clodomiro Ferreira, Maren Froemel, and Paolo Surico, 2023, Monetary policy, corporate finance, and investment, *Journal of the European Economic Association* 21, 2586–2634.

Cook, Timothy, and Thomas Hahn, 1989, The effect of changes in the Federal Funds rate target on market interest rates in the 1970s, *Journal of Monetary Economics* 24, 331–351.

De Graeve, Ferre, and Andreas Westermark, 2025, Long-lag VARs, *Journal of Monetary Economics* 156, 103831.

D'Amico, Stefania, William English, David López-Salido, and Edward Nelson, 2012, The Federal Reserve's large-scale asset purchase programmes: Rationale and effects, *Economic Journal* 122, F415–F446.

D'Amico, Stefania, and Thomas B. King, 2013, Flow and stock effects of large-scale Treasury purchases: Evidence on the importance of local supply, *Journal of Financial Economics* 108, 425–448.

Eberly, Janice C., James H. Stock, and Jonathan H. Wright, 2019, The Federal Reserve's current framework for monetary policy: A review and assessment, *International Journal of Central Banking* 16, 5–71.

Edge, Rochelle M., and Refet S. Gürkaynak, 2010, How useful are estimated DSGE model forecasts for central bankers?, *Brookings Papers on Economic Activity* 2010, 209–244.

Ehrmann, Michael, and Marcel Fratzscher, 2005, Equal size, equal role? Interest rate interdependence between the euro area and the United States, *Economic Journal* 115, 928–948.

English, William B., Skander J. Van den Heuvel, and Egon Zakrajšek, 2018, Interest rate risk and bank equity valuations, *Journal of Monetary Economics* 98, 80–97.

Fama, Eugene F, Lawrence Fisher, Michael C. Jensen, and Richard Roll, 1969, The adjustment of stock prices to new information, *International Economic Review* 10, 1–21.

Faust, Jon, Eric T. Swanson, and Jonathan H. Wright, 2004a, Do Federal Reserve policy surprises reveal superior information about the economy?, *Contributions to Macroeconomics* 4, 1–29.

Faust, Jon, Eric T. Swanson, and Jonathan H. Wright, 2004b, Identifying VARs based on high frequency futures data, *Journal of Monetary Economics* 51, 1107–1131.

Ferreira, Leonardo N., Silvia Miranda-Agrippino, and Giovanni Ricco, 2025, Bayesian local projections, *The Review of Economics and Statistics* 107, 1424–1438.

Foley-Fisher, Nathan, Rodney Ramcharan, and Edison Yu, 2016, The impact of unconventional monetary policy on firm financing constraints: Evidence from the maturity extension program, *Journal of Financial Economics* 122, 409–429.

Gamber, Edward, and Julie Smith, 2009, Are the Fed's inflation forecasts still superior to the private sector's?, *Journal of Macroeconomics* 31, 240–251.

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

Gürkaynak, Refet S., A. Hakan Kara, Burçin Kısacıkoğlu, and Sang Seok Lee, 2021, Monetary policy surprises and exchange rate behavior, *Journal of International Economics* 130, 103443.

Gürkaynak, Refet S., Burçin Kısacıkoğlu, and Jonathan H. Wright, 2020, Missing events in event studies: Identifying the effects of partially measured news surprises, *American Economic Review* 110, 3871–3912.

Gürkaynak, Refet S., Andrew Levin, and Eric T. Swanson, 2010, Does inflation targeting anchor long-run inflation expectations? Evidence from the US, UK, and Sweden, *Journal of the European Economic Association* 8, 1208–1242.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson, 2005a, Do actions speak louder than words? The response of asset prices to monetary policy actions and statements, *International Journal of Central Banking* 1, 55–93.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson, 2005b, The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models, *American economic review* 95, 425–436.

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

Gürkaynak, Refet S., and Jonathan H. Wright, 2013, Identification and inference using event studies, *The Manchester School* 81, 48–65.

Gürkaynak, Refet S., Gökçe Karasoy-Can, and Sang Seok Lee, 2022, Stock market's assessment of monetary policy transmission: The cash flow effect, *The Journal of Finance* 77, 2375–2421.

Gürkaynak, Refet S., Mahmut S. İpek, and Giovanni Ricco, 2026, A conductive measure of Chinese monetary policy, Working Paper, Bilkent University.

Hamilton, James, 1994, *Time Series Analysis* (Princeton University Press).

Handlan, Amy, 2022, Text shocks and monetary surprises: Text analysis of FOMC statements with machine learning, Working Paper, Brown University.

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

Hazell, Jonathon, and Stephan Hobler, 2024, Do deficits cause inflation?: A high frequency narrative approach, Working Paper, London School of Economics.

Hoesch, Lukas, Barbara Rossi, and Tatevik Sekhposyan, 2023, Has the information channel of monetary policy disappeared? Revisiting the empirical evidence, *American Economic Journal: Macroeconomics* 15, 355–387.

Ippolito, Filippo, Ali K. Ozdagli, and Ander Perez-Orive, 2018, The transmission of monetary policy through bank lending: The floating rate channel, *Journal of Monetary Economics* 95, 49–71.

Jarociński, Marek, 2024, Estimating the Fed's unconventional monetary policy shocks, *Journal of Monetary Economics* 144, 103548.

Jarociński, Marek, and Peter Karadi, 2020, Deconstructing monetary policy surprises—The role of information shocks, *American Economic Journal: Macroeconomics* 12, 1–43.

Jentsch, Carsten, and Kurt G. Lunsford, 2019, The dynamic effects of personal and corporate income tax changes in the United States: Comment, *American Economic Review* 109, 2655–2678.

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

Känzig, Diego R., 2021, The macroeconomic effects of oil supply news: Evidence from OPEC announcements, *American Economic Review* 111, 1092–1125.

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

Li, Dake, Mikkel Plagborg-Møller, and Christian K. Wolf, 2024, Local projections vs. VARs: Lessons from thousands of DGPs, *Journal of Econometrics* 244, 105722.

Lucca, David O., and Francesco Trebbi, 2009, Measuring central bank communication: An automated approach with application to FOMC statements, Technical report, NBER Working Paper 15367.

Lütkepohl, Helmut, 2006, *New Introduction to Multiple Time Series Analysis* (Springer).

Mertens, Karel, and José Luis Montiel Olea, 2018, Marginal tax rates and income: New time series evidence, *Quarterly Journal of Economics* 133, 1803–1884.

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

Miranda-Agrippino, Silvia, and Tsvetelina Nenova, 2022, A tale of two global monetary policies, *Journal of International Economics* 136, 103606.

Miranda-Agrippino, Silvia, and Hélène Rey, 2020, U.S. monetary policy and the global financial cycle, *Review of Economic Studies* 87, 2754–2776.

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

Miranda-Agrippino, Silvia, and Giovanni Ricco, 2023, Identification with external instruments in structural VARs, *Journal of Monetary Economics* 135, 1–19.

Montiel Olea, José Luis, Mikkel Plagborg-Møller, Eric Qian, and Christian K. Wolf, 2026a, Double robustness of local projections and some unpleasant VARithmetic, *Econometrica* 94, 1313–1343.

Montiel Olea, José Luis, Mikkel Plagborg-Møller, Eric Qian, and Christian K. Wolf, 2026b, Local projections or vector autoregressions? A primer for macroeconomists, *NBER Macroeconomics Annual* 40, 111–152.

Montiel Olea, José Luis, James H. Stock, and Mark W. Watson, 2021, Inference in structural vector autoregressions identified with an external instrument, *Journal of Econometrics* 225, 74–87.

Nagel, Stefan, and Zhengyang Xu, 2026, Movements in yields, not the equity premium: Bernanke-Kuttner redux, *Journal of Finance* forthcoming.

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

Oehler, Andreas, Matthias Horn, and Stefan Wendt, 2017, Brexit: Short-term stock price effects and the impact of firm-level internationalization, *Finance Research Letters* 22, 175–181.

Ottonello, Pablo, and Thomas Winberry, 2020, Financial heterogeneity and the investment channel of monetary policy, *Econometrica* 88, 2473–2502.

Paul, Pascal, 2020, The time-varying effect of monetary policy on asset prices, *The Review of Economics and Statistics* 102, 690–704.

Plagborg-Møller, Mikkel, and Christian K. Wolf, 2021, Local projections and VARs estimate the same impulse responses, *Econometrica* 89, 955–980.

Ramey, Valerie, 2016, Macroeconomic shocks and their propagation, in John B. Taylor, and Harald Uhlig, eds., *Handbook of Macroeconomics*, volume 2, 71–162 (Elsevier).

Ramey, Valerie, and Sarah Zubairy, 2018, Government spending multipliers in good times and bad: Evidence from US historical data, *Journal of Political Economy* 126, 850–901.

Ricco, Giovanni, and Giacomo Savini, 2025, Decomposing monetary policy surprises: Shock, information, and policy rule revision, CEPR Discussion Paper DP20166, Centre for Economic Policy Research.

Rigobon, Roberto, 2003, Identification through heteroskedasticity, *Review of Economics and Statistics* 85, 777–792.

Rigobon, Roberto, and Brian Sack, 2004, The impact of monetary policy on asset prices, *Journal of Monetary Economics* 51, 1553–1575.

Rogers, John H., Chiara Scotti, and Jonathan H. Wright, 2014, Evaluating asset-market effects of unconventional monetary policy: A multi-country review, *Economic Policy* 29, 749–799.

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

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

Sastry, Karthik, 2026, Disagreement about monetary policy, *American Economic Journal: Macroeconomics* 18, 486–522.

Sekkel, Rodrigo, Henry Stern, and Xu Zhang, 2026, Money talks: How domestic and foreign monetary policy communications move financial markets, Working Paper.

Sims, Christopher A., 1980, Macroeconomics and Reality, *Econometrica* 48, 1–48.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz, 2013, Prediction markets for economic forecasting, in Graham Elliott, and Allan G. Timmerman, eds., *Handbook of Economic Forecasting*, volume 2, 657–687 (Elsevier).

Staiger, Douglas, and James H. Stock, 1997, Instrumental variables regression with weak instruments, *Econometrica* 65, 557–586.

Stock, James H., and Mark W. Watson, 2012, Disentangling the channels of the 2007–09 recession, *Brookings Papers on Economic Activity* 1, 81–135.

Stock, James H., and Mark W. Watson, 2018, Identification and estimation of dynamic causal effects in macroeconomics using external instruments, *Economic Journal* 128, 917–948.

Stock, James H., Jonathan H. Wright, and Motohiro Yogo, 2002, A survey of weak instruments and weak identification in generalized method of moments, *Journal of business & economic statistics* 20, 518–529.

Stock, James H., and Motohiro Yogo, 2005, Testing for weak instruments in linear IV regression, in Donald Andrews, and James Stock, eds., *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, 80–108 (Cambridge University Press).

Swanson, Eric, 2023, The importance of Fed chair speeches as a monetary policy tool, *AEA Papers and Proceedings* 113, 394–400.

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

Swanson, Eric T., 2024, The macroeconomic effects of the Federal Reserve's conventional and unconventional monetary policies, *IMF Economic Review* 72, 1152–1184.

Swanson, Eric T, and Vishuddh Jayawickrema, 2024, Speeches by the Fed chair are more important than FOMC announcements: An improved high-frequency measure of U.S. monetary policy shocks, Working Paper, UC Irvine.

Swanson, Eric T., and John C. Williams, 2014, Measuring the effect of the zero lower bound on medium- and longer-term interest rates, *American Economic Review* 104, 3154–3185.

Tenreyro, Silvana, and Gregory Thwaites, 2016, Pushing on a string: US monetary policy is less powerful in recessions, *American Economic Journal: Macroeconomics* 8, 43–74.

Timmermann, Allan G., 1993, How learning in financial markets generates excess volatility and predictability in stock prices, *The Quarterly Journal of Economics* 108, 1135–1145.

Wiegand, Courtney, 2025, The effect of fiscal policy shocks on asset prices, Working Paper, University of Pennsylvania.

Wright, Jonathan H., 2012, What does monetary policy do to long-term interest rates at the zero lower bound?, *Economic Journal* 122, F447–F466.