Controls, Not Shocks: Estimating Dynamic Causal Effects in Macroeconomics*

Simon Lloyd† Ed Manuel‡

April 25, 2024

Abstract

A common approach to estimating causal effects in macroeconomics involves constructing orthogonalised 'shocks' then integrating them into local projections or vector autoregressions. For a general set of estimators, we show that this two-step 'shock-first' approach can be problematic for identification and inference relative to a one-step procedure which simply adds appropriate controls directly in the outcome regression. We show this analytically by comparing one- and two-step estimators without assumptions on underlying data-generating processes. In simple OLS settings, the two approaches yield identical coefficients, but two-step inference is unnecessarily conservative. More generally, one- and two-step estimates can differ due to omitted-variable bias in the latter when additional controls are included in the second stage or when employing non-OLS estimators. In monetary-policy applications controlling for central-bank information, one-step estimates indicate that the (dis)inflationary consequences of US monetary policy are more robust than previously realised, not subject to a 'price puzzle'.

Key Words: Identification; Instrumental Variables; Local Projections; Omitted-Variable Bias; VARs.

JEL Codes: C22, C26, C32, C36, E50, E60.

*This paper was formerly presented under the title: *Controls, Not Shocks: Estimating Dynamic Causal Effects in the Face of Confounding Factors*. We thank Saleem Bahaj, Ambrogio Cesa-Bianchi, Robin Braun, Maarten De Ridder, Álvaro Fernández-Gallardo, Georgios Georgiadis, Maximilian Grimm, Sinem Hacioglu Hoke, Joe Hazell, Daniel Lewis, Silvia Miranda-Agrippino, Daniel Ostry, Steve Pischke, Ricardo Reis, Rana Sajedi, Alan Taylor, Silvana Tenreyro and an anonymous referee, as well as presentation attendees at the Bank of England, Econometric Society Winter Meetings 2022 (Berlin), 30th Symposium of the Society for Nonlinear Dynamics and Econometrics (Orlando), Royal Economic Society Annual Conference 2023 (Glasgow), Workshop in Empirical and Theoretical Macroeconomics (King's College London), UEA 1st Time Series Workshop (Norwich), Siena Workshop in Econometric Theory and Applications, International Association for Applied Econometrics Annual Conference 2023 (Oslo), International Conference on Empirical Economics (PSU-Altoona), Money, Macro and Finance Annual Conference 2023 (U. Portsmouth), EABCN Conference on Advances in Local Projections and Empirical Methods for Central Banking (Pompeu Fabra), CEPR Paris Symposium 2023, London School of Economics, and U. Glasgow for helpful comments. The views expressed here are solely those of the authors and so cannot be taken to represent those of the Bank of England.

†Bank of England and Centre for Macroeconomics. Email Address: simon.lloyd@bankofengland.co.uk.

‡London School of Economics. Email Address: e.manuel@lse.ac.uk.

1 Introduction

Identifying the causal impact of economic events or policy interventions is crucial in dynamic macroeconomics (Frisch, 1933). A key challenge in any study of causal inference is the presence of confounding factors that simultaneously drive the causal variable and outcome variable of interest. One way to 'partial out' the effect of confounding factors is to include them as control variables in a regression with the outcome variable. However, the macroeconomics literature has typically taken an alternate route in two steps: first estimating 'shocks' as the residuals from a regression of the causal variable of interest on the set of confounding factors, and then integrating these shocks in a local projection (LP) or vector auto-regression (VAR).[1]

This two-step 'shock-first' approach to controlling for confounding factors is widespread in macroeconomics. Indeed, the construction of a series of shocks is typically viewed as an essential first step for the study of dynamic causal effects (e.g., Ramey, 2016; Nakamura and Steinsson, 2018b). The approach has been used extensively in the monetary-policy literature—most famously in Romer and Romer (2004) and subsequent studies that control for central-bank forecasts or other policy-rule variables to identify shocks.[2] A similar two-step approach is applied in studies that use high-frequency monetary-policy 'surprises' and first 'clean' the surprises of any predictability before using them as an instrumental variable (IV) in a regression or VAR (e.g., Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2022; Karnaukh and Vokata, 2022). The approach has also been used to estimate the effects of other policies—including fiscal, macroprudential, and trade—as well as non-policy variables—like oil-price, technology, sentiment and climate shocks.[3] While many of these studies have typically relied on OLS (or IV) estimates, the two-step 'shock-first' approach has also proved popular in quantile regression (QR) settings when estimating the causal effects of various macroeconomic policies on conditional quantiles of outcome variables.[4]

In this paper, we argue that this popular two-step 'shock-first' approach is problematic. Our argument is a simple one. It is well-known that in a setting with confounding factors,

[1] In the literature estimating the effects of policy, this first step is sometimes described as capturing the policy 'reaction function'. Alternatively, this step can be interpreted as 'cleaning' or 'purging' the variable of endogeneity. This is also analogous to constructing a shock via a recursive vector auto-regression.

[2] See, for example: Coibion (2012); Cloyne and Hürtgen (2016); Tenreyro and Thwaites (2016); Coibion, Gorodnichenko, Kueng, and Silvia (2017); Chen, Ren, and Zha (2018); Champagne and Sekkel (2018); Cloyne, Ferreira Mayorga, and Surico (2020); Falck, Hoffmann, and Hürtgen (2021); Holm, Paul, and Tischbirek (2021); Cloyne, Hürtgen, and Taylor (2022); Coglianese, Olsson, and Patterson (2023). However, this two-step approach has a much longer tradition in the monetary-policy literature going back to at least Barro (1977).

[3] Examples for fiscal policy include: Corsetti, Meier, and Müller (2012); Auerbach and Gorodnichenko (2013); Miyamoto, Nguyen, and Sergeyev (2018). Macroprudential policy examples include: Forbes, Reinhardt, and Wieladek (2017); Ahnert, Forbes, Friedrich, and Reinhardt (2021); Chari, Dilts-Stedman, and Forbes (2022). Metiu (2021) applies the two-step approach to trade policy. Other examples include studies on the effects of oil-price (e.g., Kilian, 2009), sentiment (e.g., Al-Amine and Willems, 2023), technology (e.g., Miranda-Agrippino, Hacioglu-Hoke, and Bluwstein, 2020) and temperature (e.g., Nath, Ramey, and Klenow, 2023) shocks.

[4] See, for example: Linnemann and Winkler (2016); Brandão-Marques, Gelos, Narita, and Nier (2021); Gelos, Gornicka, Koepke, Sahay, and Sgherri (2022).

identification can be achieved by using a simple one-step multivariate regression that uses confounding factors as control variables. So in this paper we compare the difference in regression coefficients (and standard errors) between the 'one-step' and 'two-step' approaches. Conventional wisdom holds that these two approaches are simply equivalent, by the Frisch-Waugh-Lovell Theorem (Frisch and Waugh, 1933; Lovell, 1963), but we highlight that this equivalence result only holds for estimated OLS coefficients (not standard errors), and only in simple settings rarely used in the applied literature. Across a range of applications, we demonstrate that the two-step approach can result in issues for both identification and inference, with important consequences for the estimation of dynamic causal effects in macroeconomics. We further demonstrate how a one-step approach can be applied in a wide range of common macroeconomic applications, covering estimation by both LPs and VARs.

An 'Omitted-Variable Bias' (OVB) Result. Our results are grounded purely in the properties of various estimators, without any assumptions about the true data-generating process or underlying causal structures. Our key insight is to note that the difference between estimated coefficients in the one- and two-step approaches can always be expressed in terms of an omitted-variable bias (OVB) formula.[5] The issue arises as the two-step approach excludes potentially relevant variables (i.e., the confounding factors) that are included in the first stage, but are then excluded from the outcome regression. This result is general, applying to a range of estimators defined as the unique solution to a minimisation problem of some function of the residuals. Armed with this result, we demonstrate the implications of OVB across a range of settings covering OLS, IV and QR.

In a simple OLS setting, the Frisch-Waugh-Lovell Theorem provides our point of departure. When the outcome variable is directly regressed on the shock (without auxiliary controls), the OVB term is zero such that the one- and two-step approaches yield identical point estimates. However, the two-step approach still has practical drawbacks. Most notably, the estimated standard errors from two-step estimation will be *over*-estimated if the confounding factors have explanatory power for the outcome variable. This result follows directly from comparing well-known standard-error formulas, although we are not aware that this point has been previously noted in the macroeconomics literature.[6]

When auxiliary controls are used in the second stage—a common approach in any LP or VAR setting—the equivalence in coefficients no longer applies. If the shock is orthogonal to the auxiliary controls, then the OVB term is zero and so the one- and two-step approaches identify

[5]Throughout, we refer to OVB as the mechanical difference in regression coefficients between a 'short' and 'long' regression, where the latter includes more covariates than the former (see Angrist and Pischke, 2009).

[6]Specifically, this follows almost directly from the regression-anatomy standard-error formula in Angrist and Pischke (2014). The fact the variance of the residuals in the two-step approach are larger than the one-step approach is also discussed in Lovell (1963).

the same population parameter—though issues with standard-errors in the two-step approach still apply. In contrast, if the shock is correlated with the auxiliary controls, then OVB can be non-zero as the two-step effectively fails to partial out the confounding factors used in the first stage. We also show that these drawbacks of the two-step approach carry over to IV settings—since IV estimates are the ratio of OLS estimates—covering cases where orthogonalised shocks are used as 'external' instruments in LPs or VARs.

In more general settings, including QR, OVB in the two-step approach is more problematic. We provide an explicit formula for OVB in this setting, and demonstrate that it will non-zero even in applications without auxiliary controls.

Applications. Our OVB result has important implications for a range of techniques used in the applied macroeconomics literature to estimate dynamic causal effects. We illustrate this by revisiting popular approaches to estimate the dynamic effects of US monetary policy on US Consumer Price Index (CPI) inflation when controlling for central-bank information. We begin with the approach of Romer and Romer (2004): we construct monetary policy 'shocks' by regressing changes in the Federal Funds target rate on Greenbook forecasts and forecast revisions. We then employ these shocks in a range of LP and VAR specifications to estimate dynamic causal effects, and compare this with a 'one-step' approach that controls directly for central-bank information within the LP or VAR. Figure 1 summarises some of our key findings.

First, and most directly, our results have implications for estimation via LPs (Jordà, 2005). Using an orthogonalised shock directly in a second-stage LP delivers identical IRF estimates to simply estimating a one-step LP with appropriate controls—although the two-step approach leads to overly wide standard errors, as demonstrated in Figure 1a. Projecting h-period-ahead inflation on the Romer and Romer (2004) shock does not produce a significant response at any horizon. In contrast, estimates are highly significant when using a one-step approach that regresses inflation on the change in the Federal Funds target rate controlling for the Greenbook forecasts. A more common approach in the literature is to include additional controls in LPs alongside a measure of a shock. In this case, the two-step approach can suffer from a form of OVB, since the auxiliary controls can serve to 'undo' some of the orthogonalisation from the first-stage. We demonstrate this in Figure 1b, using lagged CPI, industrial production and unemployment as auxiliary controls. Here, the OVB is non-zero and removing it with a one-step approach removes a significant portion of the near-term 'price puzzle' identified in previous studies when using this shock series (e.g., Ramey, 2016) and—unlike the two-step procedure—delivers a significant reduction in CPI at longer horizons.

Second, our results extend to estimation via recursive structural VARs (SVARs). As in the LP with auxiliary controls, an OVB can arise in a VAR setting since lags of the endogenous variables in the VAR can serve to 'undo' some of the orthogonalisation from the first-

Figure 1: Estimated response of US CPI to a US monetary policy shock

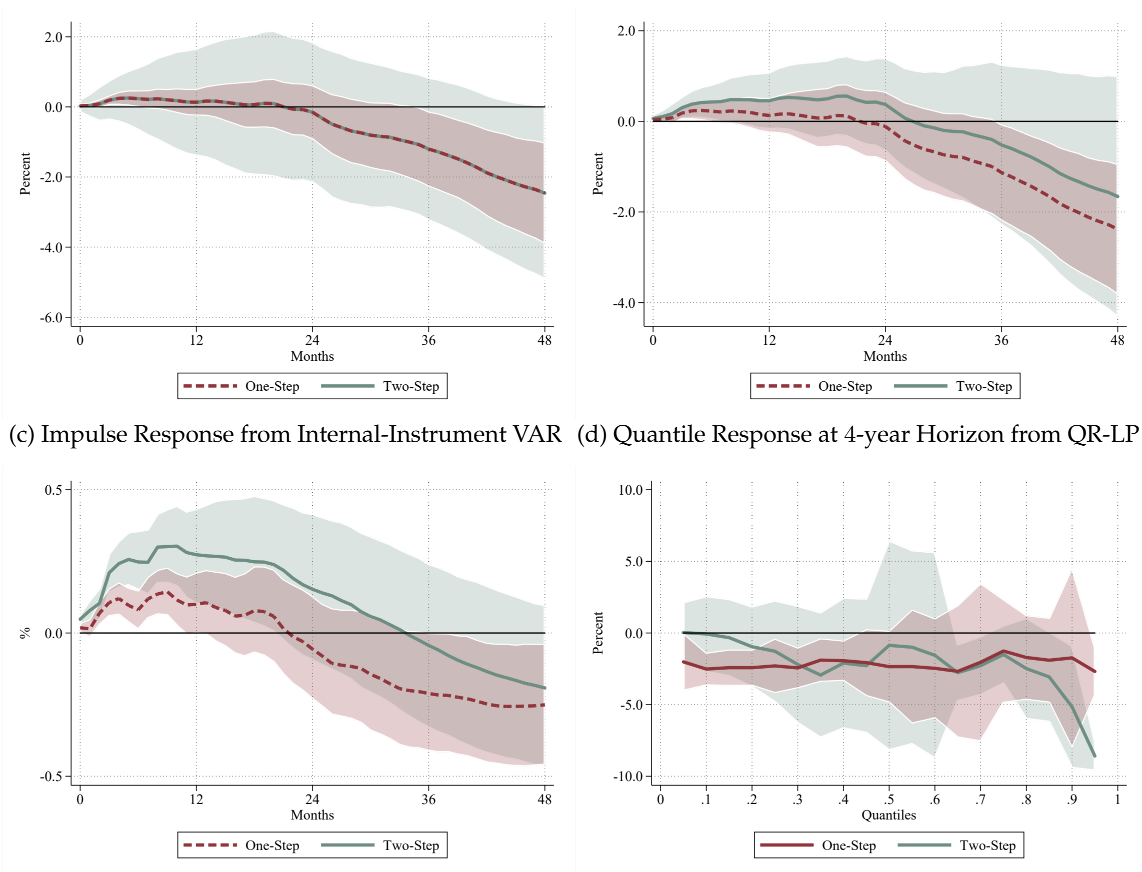

(a) Impulse Response from LP No Controls (b) Impulse Response from LP With Controls

(c) Impulse Response from Internal-Instrument VAR (d) Quantile Response at 4-year Horizon from QR-LP

Notes: Estimated response of US $\ln(CPI)$ to US monetary policy shock using two-step shock-identification strategy, as well as alternative one-step estimator that controls directly for central-bank information. Estimated using monthly data for the period 1972:01-2007:12. Shaded area denotes: 90% confidence bands from Newey and West (1987) standard errors in panels (a) and (b); 68% confidence bands from wild bootstrap in panel (c); 90% confidence bands from block bootstrap in panel (d). For more details, see Section 4 and Appendix D.

stage construction of the shock. We demonstrate this by employing the orthogonalised Romer and Romer (2004) shock as an ‘internal instrument’ (Plagborg-Møller and Wolf, 2021) in a recursively-identified SVAR. We order the shock first in the VAR, which also includes the Federal Funds target rate, consumer prices, industrial production, unemployment and commodity prices. We contrast this with a one-step approach that adds the Greenbook forecasts directly as additional variables into the recursive SVAR, ordered before the Federal Funds target rate, in turn ordered before the macro variables. As Figure 1c shows, controlling directly for the Greenbook forecasts within the VAR mitigates the price puzzle in the near term and, again, delivers a significant reduction in CPI further out, unlike the ‘shock-first’ internal-instrument approach. Strikingly, these results indicate that price puzzles can be resolved using the Romer and Romer identification strategy in a VAR without appealing to a 'recursiveness assumption' where monetary policy is prevented from impacting real variables contemporaneously.

Third, our results have implications for estimation of dynamic causal effects via 'external instruments' in either LPs or VARs (see Stock and Watson, 2018, for a review). Common macroeconomic applications of external-instrument methods use constructed residuals as instruments, which we show can generate issues for both inference and identification. Within LP settings, we demonstrate that our empirical results with and without auxiliary controls using LP-OLS (Figures 1a and 1b), carry over to an LP-IV setup in which the Romer and Romer (2004) shock is employed as an instrument. We also demonstrate the implications of our findings by using the Romer and Romer (2004) shock within a proxy-SVAR (popularised, e.g., by Mertens and Ravn, 2013; Gertler and Karadi, 2015) and, again, present an alternate 'one-step' approach that alleviates the issues induced by the two-step approach. We also show how these issues play out in an application where we use monetary-policy surprises from financial markets as instruments, but account for central-bank information, akin to Miranda-Agrippino and Ricco (2021). Similar to our results for OLS standard errors, we find the two-step approach over-estimates first-stage F-statistics relative to a one-step IV regression. This is significant in our application: using the Miranda-Agrippino and Ricco (2021) shock as an instrument delivers a first-stage F-statistic of close to, or even a little below, 10, while using monetary-policy surprises as an instrument and controlling for Greenbook forecasts within the IV regression delivers F-statistics that are around twice as large (and so sit comfortably above the common threshold of 10). This has important implications for the monetary-policy high-frequency literature, where current best practice involves orthogonalising surprises with respect to macroeconomic and financial data before integrating them in LPs or VARs, and where weak instruments and low power are a perennial issue (Bauer and Swanson, 2022).

Fourth, our results have implications for recent literature that goes beyond standard linear-regression techniques to study the probability of financial crises and drivers of macroeconomic tail risk (Schularick and Taylor, 2012; Adrian et al., 2019). In particular, our findings are relevant for recent attempts to identify the causal effect of policies on tail risk (see, e.g., Linnemann and Winkler, 2016; Brandão-Marques et al., 2021; Gelos et al., 2022). These studies have relied on a two-step approach, which we highlight suffers from QR-OVB relative to a one-step regression. In our empirical application, we consider the response of quantiles of future CPI to the monetary shock. Focusing on the 4-year-ahead horizon, where the average effects of monetary policy peak in the other panels, Figure 1d demonstrates how estimates from the two-step approach imply that changes in monetary policy significantly affect the right-tail of the inflation distribution much more so than the median. However, removing the OVB through the one-step approach reveals a different conclusion: that monetary policy instead acts as a 'location shifter' of the entire inflation distribution. More generally, these findings have important implications for effective policymaking when policymakers are seeking to contain future risks.

Literature. To the best of our knowledge, our paper is the first to explicitly highlight these drawbacks of the two-step approach popular in the macroeconomics literature. The starting point for our paper—that the two-step shock-first approach is equivalent to a regression-control strategy—follows from the well-known Frisch-Waugh-Lovell theorem.[7] Our key contribution is to derive an explicit formula linking the one- and two-step approaches that can be applied across a range of settings, allowing us to demonstrate where this equivalence breaks down and clearly highlight the drawbacks of the two-step approach with respect to both inference and identification.

Our findings for standard-error estimation may, at first glance, appear related to results from a broad literature on inference in econometric models with generated regressors (see, e.g. Hansen, 2022, for a review). A key result from this literature is that standard-error formulas provide a consistent estimator of true standard errors specifically when testing the null hypothesis that the coefficient on a generated regressor is zero (Pagan, 1984). More generally, a number of papers highlight that standard-error formulas typically under-state the variance associated with estimated coefficients on generated regressors (see, e.g., Murphy and Topel, 2002, who provide alternate standard-error formulas for this case). We depart from this literature by arguing that controlling for confounding factors via a two-step approach amounts to estimating a misspecified model since doing so omits a potentially relevant variable from the regression with the outcome variable. We show that if in fact the confounding factors have explanatory power for the outcome variable (as would be expected), then using unadjusted standard-error formulas in the two-step approach delivers inference that is *unnecessarily conservative* (i.e., the exact opposite of the concern raised in previous literature), even when the researcher is only concerned with testing the null hypothesis of a zero coefficient.

Outline. The remainder of the paper is structured as follows. Section 2 sets out our key insight in general form. Section 3 discusses the implications of the OVB in different settings. Section 4 presents empirical applications. Section 5 concludes.

2 Omitted-Variable Bias in the Two-Step Approach

In this section, after defining notation, we describe our general setup. Using this, we present our key insight: the difference between the two- and one-step approach can always be ex-

[7] This equivalence has been noted in a range of the applied macro literature e.g. variously in: Angrist and Kuersteiner (2011), Jordà and Taylor (2016), Angrist, Jordà, and Kuersteiner (2018), Barnichon and Brownlees (2019), Jordà, Schularick, and Taylor (2020) and Plagborg-Møller and Wolf (2021).

pressed in terms of an OVB formula stemming from the exclusion of potentially relevant variables from the outcome regression in the two-step approach.

2.1 Notation and General Setup

We label y_t as the outcome variable of interest at time t and z_t is the causal variable (e.g., a policy indicator). Throughout, our central focus is on the causal effect of z_t on y_t.

We define $\mathbf{x}_t$ as a $(K_1 + K_2) \times 1$ vector of (non-perfectly-collinear) observable control variables that potentially drive y_t and z_t. We partition these controls into $\mathbf{x}_{1,t}$ and $\mathbf{x}_{2,t}$, which are $K_1 \times 1$ and $K_2 \times 1$ vectors, respectively. We allow $\mathbf{x}_{2,t}$ to be potentially empty (i.e., $K_2 \geq 0$), but restrict $\mathbf{x}_{1,t}$ to be non-empty (i.e., $K_1 > 0$). Our key results do not rely on any assumptions around the causal structure. But to aid with intuition, one can think of $\mathbf{x}_t$ as potentially 'confounding factors' (i.e., variables that simultaneously drive y_t and z_t).[8]

We now introduce the widely used two-stage 'shock-first' approach to estimate the effect of z_t on y_t in its general form.

First, a 'shock' to z_t is defined as:

$$\varepsilon_t = z_t - \mathbf{x}_{1,t}'\boldsymbol{\delta} \tag{1}$$

where for now we only impose the restriction that $\boldsymbol{\delta}$ is a vector of real numbers, $\boldsymbol{\delta} \in \mathbb{R}^{K_1}$. However, throughout the paper we will be concerned with the specific case where ε_t is defined as a population OLS residual from a regression of z_t on $\mathbf{x}_{1,t}$, in which case the following holds by construction:

$$\mathbb{E}\left[\varepsilon_t \mathbf{x}_{1,t}'\right] = 0$$

where, given this orthogonality condition, ε_t can be thought of as changes in z_t 'purged' of the confounding effect of $\mathbf{x}_{1,t}$.

The coefficient of interest is then defined by a second-step (population) regression of the outcome variable y_t on the shock ε_t and (potentially) additional controls $\mathbf{x}_{2,t}$:

$$y_t = \mathbf{x}_{2,t}'\boldsymbol{\alpha} + \varepsilon_t \beta_{2S} + u_t^{2S} \tag{2}$$

where $\boldsymbol{\alpha}$ is a $K_2 \times 1$ vector of population regression coefficients pertaining to the controls $\mathbf{x}_{2,t}$, and β_{2S} is the population regression coefficient of interest. The regression coefficients underpinning equation (2) satisfy the following (population) minimisation problem:

$$\{\boldsymbol{\alpha}, \beta_{2S}\} = \arg\min_{\mathbf{a},b} \left\{\mathbb{E}\left(f\left(y_t - \mathbf{x}_{2,t}'\mathbf{a} - \varepsilon_t b\right)\right)\right\} \tag{3}$$

[8] Alternatively, z_t can be thought of as a 'policy variable' and $\mathbf{x}_t$ as variables that potentially feature in the policy 'reaction function', driving endogeneity with respect to y_t.

where $f(x)$ summarises the objective function of the regression. For example, for OLS $f(x) = x^2$ and for QR $f(x) = \rho(x)$, where $\rho(x)$ is the check function. For our purposes, the function $f(\cdot)$ must be well-defined and have a unique minimisation point. Beyond that, we place no further restrictions on $f(\cdot)$ in this section.

It is important to note the generality of our setup. In addition to subsuming a range of estimation methods (i.e., different functional forms for $f(\cdot)$), this framework applies to a wide range of applications in macroeconomics. For instance, when the set of covariates $\mathbf{x}_{1,t}$ includes lags of z_t plus lagged (and some contemporaneous) values of other variables, then ε_t from equation (1) is equivalent to a shock from a recursively-identified SVAR model. Alternatively, when $\mathbf{x}_{1,t}$ contains forecasts made at time t then ε_t can, e.g., be thought of as a 'narrative' shock as in Romer and Romer (2004). In addition, the outcome variable can be defined as h-period ahead values of y, in which case the second-stage regression (2) amounts to a single h-specific regression from a LP model. Alternatively, when $\mathbf{x}_{2,t}$ includes p lagged values of y_t alongside p lags of other macroeconomic variables, then regression (2) can be thought of as a single equation from a VAR model and β_{2S} captures the contemporaneous response of y_t to ε_t.[9]

The key question in this paper is whether this two-step shock-first approach to controlling for confounding factors is appropriate. Intuitively, in instances where all confounding factors are assumed to be observable, identification of the causal effect of interest can instead be achieved through a simpler one-step regression.

The alternative one-step approach to controlling for confounding factors involves the (population) regression of y_t on z_t and the full set of controls $\mathbf{x}_t$:

$$y_t = \underbrace{\mathbf{x}'_{1,t}\boldsymbol{\theta_1} + \mathbf{x}'_{2,t}\boldsymbol{\theta_2}}_{\equiv \mathbf{x}'_t\boldsymbol{\theta}} + z_t\beta_{1S} + u_t^{1S} \tag{4}$$

where $\boldsymbol{\theta}$ is a $K \times 1$ vector of population regression coefficients (where $K = K_1 + K_2$) and β_{1S} is a scalar population regression coefficient. Combined, these coefficients are defined by:

$$\{\boldsymbol{\theta}, \beta_{1S}\} = \arg\min_{\boldsymbol{\vartheta},b} \left\{\mathbb{E}\left(f\left(y_t - \mathbf{x}'_t\boldsymbol{\vartheta} - z_t b\right)\right)\right\} \tag{5}$$

where the function $f(\cdot)$ matches that used in the second stage of the shock-first regression (3).

2.2 Omitted-Variable Bias Result

Intuitively, like β_{2S}, we seek to interpret β_{1S} as the effect of z_t after partialling out any confounding effects of $\mathbf{x_t}$. The following Proposition clarifies that the difference between these

[9] We explicitly extend our setting to impulse-response estimation via SVARs in Appendices B and C.

one- and two-step coefficients can always be expressed in terms of an OVB formula that impacts the two-step approach.

Proposition 1. (OVB in the General Two-Step Approach) *Consider the following 'hybrid' population regression of y_t on ε_t and the full set of K controls $\mathbf{x}_t$:*

$$y_t = \underbrace{\mathbf{x}_{1,t}'\boldsymbol{\phi}_\mathbf{1} + \mathbf{x}_{2,t}'\boldsymbol{\phi}_\mathbf{2}}_{\equiv \mathbf{x}_t'\boldsymbol{\phi}} + \varepsilon_t \beta_{Hyb} + u_t^{Hyb} \tag{6}$$

where ε_t is defined as in equation (1) for any real vector of coefficients $\boldsymbol{\delta} \in \mathbb{R}^{K_1}$, where $K_1 \leq K$, and $\boldsymbol{\phi}$ is a $K \times 1$ vector of population regression coefficients that solves:

$$\{\boldsymbol{\phi}, \beta_{Hyb}\} = \arg\min_{\boldsymbol{\varphi}, b} \left\{\mathbb{E}\left(f\left(y_t - \mathbf{x}_t'\boldsymbol{\varphi} - \varepsilon_t b\right)\right)\right\} \tag{7}$$

where the population objective function $f(\cdot)$ matches that used in two- and one-step regressions (3) and (5). Then the following holds:

$$\beta_{Hyb} = \beta_{1S} \quad \textit{and} \quad \beta_{2S} = \beta_{1S} + \Omega$$

where Ω is defined as omitted-variable bias term that arises from the exclusion of $\mathbf{x}_{1,t}$ in (2).

Proof: Substituting the definition of ε_t from equation (1) into the hybrid estimator (7), we have:

$$\begin{aligned}
\{\boldsymbol{\phi}, \beta_{Hyb}\} &= \arg\min_{\boldsymbol{\varphi}, b} \left\{\mathbb{E}\left(f\left(y_t - \mathbf{x}_t'\boldsymbol{\varphi} - (z_t - \mathbf{x_{1,t}}\boldsymbol{\delta})b\right)\right)\right\} \\
&= \arg\min_{\boldsymbol{\varphi}, b} \left[\mathbb{E}\left(f\left(y_t - \mathbf{x}_{1,t}'\boldsymbol{\varphi}_\mathbf{1} - \mathbf{x}_{2,t}'\boldsymbol{\varphi}_\mathbf{2} - (z_t - \mathbf{x_{1,t}}\boldsymbol{\delta})b\right)\right)\right\} \\
&= \arg\min_{\boldsymbol{\varphi}, b} \left\{\mathbb{E}\left(f\left(y_t - \mathbf{x}_{1,t}'(\boldsymbol{\varphi}_\mathbf{1} - b\boldsymbol{\delta}) - \mathbf{x}_{2,t}'\boldsymbol{\varphi}_\mathbf{2} - z_t b\right)\right)\right\}
\end{aligned}$$

When $f(\cdot)$ has a unique minimand, we know from the one-step minimisation problem (5) that the solution to the above minimisation problem is given by:

$$\begin{aligned}
\boldsymbol{\phi}_\mathbf{2} &= \boldsymbol{\theta}_\mathbf{2} \\
\boldsymbol{\phi}_\mathbf{1} - \beta_{Hyb}\boldsymbol{\delta} &= \boldsymbol{\theta}_\mathbf{1} \quad \Longrightarrow \quad \boldsymbol{\phi}_\mathbf{1} = \boldsymbol{\theta}_\mathbf{1} + \beta_{Hyb}\boldsymbol{\delta} \\
\beta_{Hyb} &= \beta_{1S}
\end{aligned}$$

Since regression (6) is the same as regression (2), albeit with additional covariates, the difference between coefficients can be expressed in terms of an omitted-variable bias term (that is, the difference in coefficients between a 'long' and 'short' regression that excludes some covariates):

$$\beta_{2S} = \beta_{Hyb} + \Omega \quad \Longrightarrow \quad \beta_{2S} = \beta_{1S} + \Omega$$

which completes the proof. □

In line with the setup described in Section 2.1, this Proposition is general. In addition to the requirements described there, since our results rely only on the properties of optimisation, they carry over to in-sample estimated coefficients.[10]

The result establishes a mechanical link between the one- and two-step coefficients, in a setting where we remain agnostic about the 'true' model. Whether the OVB term Ω reflects genuine bias is, of course, ultimately context-dependent. But the result is useful for two reasons. First, it highlights that the two-step approach hinges on the restriction that the control variables from the first stage $\mathbf{x}_{1,t}$ do not feature in the regression with the outcome variable. It is hard to understand an economic rationale for imposing such a restriction in settings where the two-step approach is typically applied.[11] Indeed, the motivation for the first-stage regression is typically to isolate exogenous variation in z_t by partialling out $\mathbf{x}_{1,t}$. But note that if $\mathbf{x}_{1,t}$ does not have any explanatory power for the outcome variable, this first step will generally not partial out *endogenous* variation in z_t. Second, this result allows us to derive an exact analytical formula for the difference between the one-step and two-step coefficients across a range of settings, and highlight exactly when such OVB may be genuinely problematic. This is even true in cases where no analytical expression exists for either coefficient individually (as in QR). We turn to this in the next section.

3 Omitted-Variable Bias in Specific Settings

In this section, we set out the practical implications of the OVB formula from Proposition 1 for a range of commonly used estimators where the two-step shock-first approach has been widely applied, specifically: OLS, IV, and QR.

3.1 Ordinary Least Squares (OLS)

Consider a case where all of equations (1), (2) and (4) are defined as OLS population regressions. From Proposition 1 and the formula for OVB in OLS, we can express the difference between regression coefficients in the following Result:

[10] Specifically, Proposition 1 holds with when $\mathbb{E}\left(f(\cdot)\right)$ is re-defined as the in-sample objective function—e.g., for some sample of length T, re-defining it as: $\frac{1}{T}\sum_{t=1}^{T} f(x)$.

[11] This point is particularly clear in our empirical application, where the two-step approach amounts to assuming that central bank forecasts for inflation are exactly orthogonal to actual future inflation outcomes—an assumption that the data strongly rejects.

Result 1. (OLS) *Define coefficients across regressions (1), (2), (4) and (6) as OLS population regression coefficients. Then the following formula relates the two-step coefficient β_{2S} in equation (2) and the one-step coefficient β_{1S} in equation (4):*

$$\begin{aligned}\beta_{2S} &= \beta_{1S} + \Omega_{OLS} \\ &= \beta_{1S} + \left[\mathbf{A}\,\mathbb{E}\left[\varepsilon_t \mathbf{x}_{1,t}'\right] + \mathbb{E}\left[\varepsilon_t \mathbf{x}_{2,t}'\right]\mathbf{B}\right]\boldsymbol{\phi}_1 \\ &= \beta_{1S} + \mathbb{E}\left[\varepsilon_t \mathbf{x}_{2,t}'\right]\mathbf{B}\,\boldsymbol{\phi}_1\end{aligned}$$

where $\mathbf{B} \equiv \left[\mathbb{E}\left[\varepsilon_t^2\right]\mathbb{E}\left[\mathbf{x}_{2,t}\mathbf{x}_{2,t}'\right] - \mathbb{E}\left[\varepsilon_t \mathbf{x}_{2,t}\right]\mathbb{E}\left[\varepsilon_t \mathbf{x}_{2,t}'\right]\right]^{-1}\mathbb{E}\left[\mathbf{x}_{2,t}\mathbf{x}_{1,t}'\right]$ *and* $\boldsymbol{\phi}_1$ *are coefficient loadings on* $\mathbf{x_{1,t}}$ *in equation (6).*

Proof: The first line of this Result follows directly from Proposition 1. The second line then follows from the formula for OVB in OLS and standard matrix-partition algebra. The final line uses the fact that since ε_t is an OLS population residual, by construction, $\mathbb{E}\left[\varepsilon_t \mathbf{x}_{1,t}'\right] = \mathbf{0}$. □

We now discuss the implications of this result. To do so, we consider two cases.

Case 1. (No Auxiliary Controls) $\mathbf{x}_{2,t}$ *is an empty vector such that* $K_2 = 0$.

We first consider a case where no auxiliary controls are included in the second-stage regression (2). This is a common approach in applied work since, if the first stage (1) is thought to adequately identify a 'shock', then no auxiliary controls are needed in the second stage to identify the causal effect of interest. In this case, there will be no OVB in the two-step estimator. The below Corollary formalises this:

Corollary 1. (General OLS Equivalence without Controls) *Under Case 1,* $\beta_{2S} = \beta_{1S}$ *and* $\hat{\beta}_{2S} = \hat{\beta}_{1S}$.

Proof: This follows directly from the Frisch-Waugh-Lovell Theorem.[12] □

In this case, point estimates from the one- and two-step approaches will be exactly equivalent. However, although sample estimates $\hat{\beta}_{2S}$ and $\hat{\beta}_{1S}$ are mathematically equivalent, standard-error formulas deliver wider standard-error estimates for the two-step approach. The following Corollary shows this:

[12] Angrist and Pischke (2009) refer to this formulation of Frisch-Waugh-Lovell (FWL) Theorem as the regression-anatomy formula (p. 27). In our setting, the 'standard' FWL formulation instead states equivalence between β_{1S} and the coefficient from a regression of y_t orthogonalised with respect to $\mathbf{x}_{1,t}$ on ϵ_t. This approach—orthogonalising z_t *and* y_t with respect to $\mathbf{x}_{1,t}$ and then regressing orthogonalised variables on each other—would deliver identical point estimates *and identical standard error estimates* as the one-step approach (see Ding, 2021).

Corollary 2. (OLS Standard Errors without Controls) *Under Case 1, when the number of observations T is large relative to the number of regressors K_1, for homoskedastic-only standard error formulas, the estimated variance of the two-step coefficient $\hat{\beta}_{2S}$ is weakly greater than the estimated variance of the one-step coefficient $\hat{\beta}_{1S}$:*

$$\widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{2S}) \geq \widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{1S})$$

Proof: This follows from that fact that the regression anatomy formula carries over to estimated standard errors (see, e.g., Angrist and Pischke, 2014), and so:

$$\widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{2S}) \geq \widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{1S}) \quad \iff \quad \frac{1}{T-1}\mathbb{V}\text{ar}(\hat{u}_t^{2S}) \geq \frac{1}{T-(K_1+1)}\mathbb{V}\text{ar}(\hat{u}_t^{1S})$$

Note also that the sample residual from the hybrid regression (6), $\hat{u}_t^{Hyb}$, and the sample residual from the one-step regression (4), $\hat{u}_t^{1S}$, are the same—i.e., $\hat{u}_t^{Hyb} = \hat{u}_t^{1S}$. In addition, since adding covariates reduces sample residual variance in OLS, then the following must hold: $\mathbb{V}\text{ar}(\hat{u}_t^{2S}) \geq \mathbb{V}\text{ar}(\hat{u}_t^{Hyb})$. And so when T is large relative to K_1, then: $\widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{2S}) \geq \widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{1S})$. □

Note the issue arises since excluding $\mathbf{x}_{1,t}$ from the second stage in the two-step approach excludes a variable that is, by construction, completely uncorrelated with the variable of interest ε_t, while (potentially) having explanatory power for y_t. If $\mathbf{x}_{1,t}$ in fact has no explanatory power for y_t then the standard-error estimates will be equivalent between the one-step and two-step approaches.[13] However, since the variables in $\mathbf{x}_{1,t}$ are typically selected by the researcher to capture confounding factors (i.e., variables that drive not just z_t but also y_t) this is unlikely to hold. Indeed, in practical applications this over-estimation of standard errors can be large since $\mathbf{x}_{1,t}$ may have significant explanatory power for y_t. We demonstrate this point with an application in Section 4, when $\mathbf{x_{1,t}}$ includes forecasts at time t of the outcome variable of interest (as in, e.g., Romer and Romer, 2004).

Although Corollary 2 is written for a specific (homoskedastic) case, the fundamental insight carries over to other standard-error formulas. Here, again, the estimated variance of the two-step coefficient will typically be over-estimated relative to the one-step coefficient since in effect the two-step approach excludes an explanatory variable that is uncorrelated with ε_t, while having explanatory power for y_t. We again demonstrate this in our empirical application where standard errors remain significantly wider for the two-step approach for both White (1980) heteroskedastic-robust standard errors and Newey and West (1987) autocorrelation-robust standard errors.

[13]This is related to discussion in Pagan (1984) pp. 222-233, who shows that in models where only a generated residual appears on the right-hand side, standard-error formulas are valid for a two-step approach.

Case 2. (Auxiliary Controls) $\mathbf{x}_{2,t}$ *is a non-empty vector such that* $K_2 > 0$.

We now consider an alternative case in which controls are included in the second stage, such that $\mathbf{x}_{2,t}$ is non-empty. As we explain in Appendix B, this case generalises to incorporate VARs identified using recursive restriction schemes—including identification in VARs with 'internal instruments' (Plagborg-Møller and Wolf, 2021). In these cases, lags of endogenous variables in the VAR, in effect, act as auxiliary controls that can create non-zero OVB. Notwithstanding this, there are two further reasons researchers may typically include additional controls in this stage. First, as we explicitly demonstrated above, adding controls that are uncorrelated with ε_t but have explanatory power for y_t can lower estimated standard errors by 'mopping up' variance in the error term. Second, researchers may be interested in providing 'additional robustness' by controlling for other variables that may correlate with ε_t in their second-stage regression which were not included in the first stage.

We discuss each of these possible reasons in turn. In the first case, the OVB can again be zero, as the below Corollary states:

Corollary 3. (Specific OLS Equivalence with Controls) *Under Case 2, if* $\mathbb{E}\left[\varepsilon_t \mathbf{x}'_{2,t}\right] = \mathbf{0}$, *then* $\beta_{2S} = \beta_{1S}$.

Proof: This follows directly from Result 1. □

This Corollary implies that, in practice, as long as the shock is constructed correctly in the first stage to be genuinely exogenous (i.e., uncorrelated with any other drivers of y), then the OVB at the population level will be zero. But there are still important practical drawbacks of the two-step approach in this case. As before, we have that $\mathbb{V}\mathrm{ar}(\hat{u}_t^{2S}) \geq \mathbb{V}\mathrm{ar}(\hat{u}_t^{1S})$ which mechanically inflates estimated standard errors in the two-step approach. In addition, unlike for Case 1, in-sample estimated coefficients will likely differ across the two- and one-step approaches, even when the shock identified in the first stage is genuinely exogenous in population. Hence there may be efficiency costs to using the two-step approach.[14]

Suppose instead that the researcher includes controls in the second-stage regression that are potentially correlated with the shock constructed in the first stage in an attempt to additionally purge z_t of any further confounding factors—a common approach in the macroeconomics literature, both in estimation via LPs and VARs.[15] In this case, by a standard Frisch-Waugh-Lovell argument, the coefficient from the one-step approach will be identical to a two-step approach

[14] Pagan (1984) and Murphy and Topel (2002) discuss potential losses in efficiency from two-step procedures in settings where additional controls are used in the second-stage that are not included in the first-stage.

[15] Recent examples include Cloyne, Hürtgen, and Taylor (2022) who include a rich set of controls in their LP precisely to purge their constructed shock of any remaining predictability, and also in McKay and Wolf (2023) who place the Romer and Romer (2004) monetary-policy shock second-to-last in their VAR to additionally purge it of predictability with respect to contemporaneous variables as a means of securing 'exogeneity insurance'.

where the shock was constructed to be orthogonal to $\mathbf{x}_{1,t}$ *and* $\mathbf{x}_{2,t}$ in the first-stage. Crucially though, the two-step approach (which excludes $\mathbf{x}_{1,t}$ from the second-stage) will not be equivalent to the one-step approach, and Result 1 provides an exact formula for the difference in coefficients. Intuitively, the two-step approach suffers from a form of bias since the inclusion of $\mathbf{x}_{2,t}$ in the outcome regression serves to 'undo' some of the orthogonalisation (with respect to $\mathbf{x}_{1,t}$) from the first stage.

From an economics perspective, the issue can be understood as a failure to correctly identify a shock in the first stage. Viewing z_t as a policy variable, suppose the researcher uses the two-step approach (which includes only $\mathbf{x}_{1,t}$ in the shock-identification stage), but the 'true' policy reaction function includes both $\mathbf{x}_{1,t}$ and $\mathbf{x}_{2,t}$. If the researcher thinks this is a genuine risk, then Result 1 highlights that it is *not* sufficient to simply include $\mathbf{x}_{2,t}$ as controls in the second-stage regression, and instead the shock must be re-estimated with $\mathbf{x}_{1,t}$ and $\mathbf{x}_{2,t}$ as covariates. The one-step approach avoids this issue since it does not rely on the researcher correctly partitioning $\mathbf{x}_t$ into $\mathbf{x}_{1,t}$ and $\mathbf{x}_{2,t}$ such that only the former appear in the reaction function.

3.2 Instrumental Variables (IV)

We now consider the implications of Proposition 1 and Result 1 for estimation via IV. In particular, we consider a setting where OLS residuals from a first-stage regression are then used as instruments in an IV regression for the outcome variable of interest y_t. This approach has been used in a variety of studies that employ generated shocks as external instruments in either LPs or VARs (e.g., Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2022; Miranda-Agrippino and Ricco, 2023). It also appears in studies that use constructed shocks as instruments to identify structural macroeconomic equations (e.g., Barnichon and Mesters, 2020; Lewis and Mertens, 2022). As before, we compare this to a 'one-step' IV regression with control variables. Since IV regression coefficients can be expressed as the ratio of OLS regression coefficients, the problems associated with using a generated residual in a second-stage OLS regression discussed in the previous sub-section carry over directly to IV.

In order to discuss the implications for IV we develop our general setting. In particular we introduce an additional variable m_t—which the researcher wishes to instrument for ε_t—which may differ from z_t.[16] The two-step approach is now defined as the following (population) IV regression:

$$y_t = m_t \beta_{2S}^{IV} + \mathbf{x}_{2,t}' \boldsymbol{\alpha} + u_t^{2S} \tag{8}$$

where ε_t (defined as in the previous section as an OLS population residual) instruments for m_t and β_{2S}^{IV} and $\boldsymbol{\alpha}$ are population regression IV coefficients. We are interested in comparing β_{2S}^{IV}

[16]Note that when $m_t \equiv z_t$, the IV regression coefficients in this section will be (mechanically) equivalent to the OLS coefficients in the previous section (and so the results from the previous section hold).

to β_{1S}^{IV} from the following population IV-regression:

$$y_t = m_t \beta_{1S}^{IV} + \mathbf{x}_t' \boldsymbol{\theta} + u_t^{1S} \tag{9}$$

with z_t as an instrument for m_t.

Because IV coefficients can be written as the ratio of OLS coefficients from a 'first-stage' and 'reduced-form' regression (see, e.g., Angrist and Pischke, 2009, p. 122), we can write:

$$\beta_{2S}^{IV} \equiv \frac{\beta^{RF}}{\beta^{FS}}$$

where β^{RF} and β^{FS} are defined from the following OLS population regressions:

$$\begin{aligned} m_t =& \varepsilon_t \beta^{FS} + \mathbf{x}_{2,t} \boldsymbol{\pi}^{FS} + e_t^{FS} \\ y_t =& \varepsilon_t \beta^{RF} + \mathbf{x}_{2,t} \boldsymbol{\pi}^{RF} + e_t^{RF} \end{aligned}$$

Given this, the results from Section 3.1 carry over almost directly to this setting. Starting with Case 1, where auxiliary controls $\mathbf{x}_{2,t}$ are not included in the second-stage, we show that the two-step approach delivers identical coefficient estimates as the one-step approach, while over-estimating the degree of uncertainty around these estimates:

Corollary 4. (IV without Controls) *Under Case 1, then, if ε_t is defined as in equation (1), then: $\beta_{2S}^{IV} = \beta_{1S}^{IV}$ and $\hat{\beta}_{2S}^{IV} = \hat{\beta}_{1S}^{IV}$. In addition, when the number of observations T is large relative to the number of regressors $K1$, for homoskedastic-only standard error formulas, then:*

$$\begin{aligned} \widehat{\mathbb{V}\text{ar}}(\beta_{2S}^{\hat{F}S}) \geq \widehat{\mathbb{V}\text{ar}}(\beta_{1S}^{\hat{F}S}) \\ \widehat{\mathbb{V}\text{ar}}(\beta_{2S}^{\hat{I}V}) \geq \widehat{\mathbb{V}\text{ar}}(\beta_{1S}^{\hat{I}V}) \end{aligned}$$

Proof: Equivalence of coeffients follows by applying Frisch-Waugh-Lovell Theorem to the numerator and denominator of the IV estimator, defined as the ratio of OLS coefficients. The first result on estimated standard errors follows directly from Corollary 2 as it refers to standard errors on OLS coefficients from a first-stage regression. We relegate the final proof on IV standard errors to Appendix A. □

These implications of these results are very similar to the OLS setting: the two- and one-step deliver identical coefficient estimates, but with larger estimated standard errors in the former. Unlike in OLS, there are now *two* implications of note related to standard-error calculations. The first is that the two-step approach leads to an under-estimation of F-statistics from the first stage, implying a tendency to mistakenly reject 'strong' instruments as 'weak', while the second relates to an over-estimation of standard errors on the IV-coefficient of interest.

Likewise, it is straightforward to show that other results for OLS carry over to IV regression. In particular, if auxiliary controls $\mathbf{x}_{2,t}$ were included in the IV regression (8), as in Case 2, then β_{2S}^{IV} will suffer from OVB if $\mathbb{E}[\varepsilon_t \mathbf{x}_{2,t}'] \neq \mathbf{0}$. An explicit expression for OVB can be found by simply applying the formula from Result 1 to the numerator and denominator of the IV estimand. Intuitively, the bias term here relates to a potential failure of the IV-exogeneity condition. The two-step approach is equivalent to regressing ε_t on $\mathbf{x}_{2,t}$ and using the residual as an instrument, but since this residual may be correlated with $\mathbf{x}_{1,t}$ it may not be a valid instrument. In contrast, the one-step approach effectively partials out the full vector $\mathbf{x}_t$ from z_t and uses this as the instrument.

This result applies most directly to using generated residuals in an LP-IV, although carries-over to estimation via SVARs identified with external instruments (as we explicitly demonstrate in Appendix C and in our empirical application in Section 4).

3.3 Quantile Regression (QR)

To consider settings where the conditional-expectation function may not be the object of interest, in this sub-section we specifically focus on a case in which, in the first stage, the shock is constructed via OLS and is then used in a second-stage *quantile* regression. This approach has been adopted to study the effects of various policy 'shocks' on conditional quantiles of outcome variables of interest. For example, Linnemann and Winkler (2016) use it to assess the effects of fiscal policy on the GDP distribution, Gelos et al. (2022) apply it to assess the effects of capital-flow measures on 'capital flows at risk', and Brandão-Marques et al. (2021) study the influence of macroprudential policy on growth-at-risk.

In this QR setting, we have the following expression for the difference between population coefficients from a one-step and two-step quantile regression:

Result 2. (QR) *Define the coefficients in equation* (1) *as OLS population regression coefficients, and the coefficients in equations* (2), (4) *and* (6) *as QR population coefficients for some specific quantile* $\tau \in (0, 1)$, *implying the following functional form for the objective function* $f(\cdot)$ *in equations* (3), (5) *and* (7):

$$f(x) = \rho_\tau(x)$$

where $\rho_\tau(x) = (\tau - \mathbb{1}(x \leq 0))$ is the check function. Assuming that the conditional quantile function of y_t conditional on z_t, x_t is linear,[17] *then the following formula relates the two-step population coefficient at the τ-th quantile $\beta_{2S}(\tau)$ in equation (2) and the one-step population coefficient $\beta_{1S}(\tau)$ in equation (4):*

$$\begin{aligned}\beta_{2S}(\tau) &= \beta_{1S}(\tau) + \Omega_{QR}(\tau) \\ &= \beta_{1S}(\tau) + \left[\mathbf{A}\,\mathbb{E}\left[w_\tau \epsilon_t \mathbf{x}'_{1,t}\right] + \mathbb{E}\left[w_\tau \varepsilon_t \mathbf{x}'_{2,t}\right]\mathbf{B}\right]\boldsymbol{\phi}_\mathbf{1}(\tau)\end{aligned}$$

where:

$$\begin{aligned}\mathbf{A} \equiv &[\mathbb{E}\left[w_\tau \varepsilon_t^2\right]^{-1} + \mathbb{E}\left[w_\tau \varepsilon_t^2\right]^{-2}\mathbb{E}\left[w_\tau \varepsilon_t \mathbf{x}_{2,t}\right]\left[\mathbb{E}\left[w_\tau \mathbf{x}_{2,t}\mathbf{x}'_{2,t}\right]\right. \\ &\left. - \mathbb{E}\left[w_\tau \varepsilon_t^2\right]^{-1}\mathbb{E}\left[w_\tau \varepsilon_t \mathbf{x}_{2,t}\right]\mathbb{E}\left[w_\tau \varepsilon_t \mathbf{x}'_{2,t}\right]\right]^{-1}\mathbb{E}\left[w_\tau \varepsilon_t \mathbf{x}_{2,t}\right] \\ \mathbf{B} \equiv &\left[\mathbb{E}\left[w_\tau \mathbf{x}_{2,t}\mathbf{x}'_{2,t}\right] - \mathbb{E}\left[w_\tau \varepsilon_t^2\right]^{-1}\mathbb{E}\left[w_\tau \varepsilon_t \mathbf{x}_{2,t}\right]\mathbb{E}\left[w_\tau \varepsilon_t \mathbf{x}'_{2,t}\right]\right]^{-1}\mathbb{E}\left[w_\tau \mathbf{x}_{2,t}\mathbf{x}'_{1,t}\right]\mathbb{E}\left[w_\tau \varepsilon_t^2\right]\end{aligned}$$

and $\boldsymbol{\phi}_1(\tau)$ are coefficient loadings on $\mathbf{x_{1,t}}$ in equation (6). In addition the 'weights' are defined as: $w_\tau = \int_0^1 f_{u_\tau^{Hyb}}(u\,[\mathbf{x}'_{2,t}\boldsymbol{\pi}(\tau) + \varepsilon_t\beta(\tau) - \mathbf{x}'_t\boldsymbol{\phi}(\tau) - \varepsilon_t\beta_{hyb}(\tau)]|\mathbf{x}_t, \varepsilon_t)\,\mathrm{d}u/2$.

Proof: The first line follows directly from Proposition 1. The second line follows from the formula for OVB for QR (Angrist, Chernozhukov, and Fernández-Val, 2006) and standard matrix parition algebra. □

We now consider the practical implications for this result in different settings. As before, we first consider a simple setting where no auxiliary controls are included in the second-stage regression (i.e., Case 1).

In Case 1, unlike in OLS, there can still be omitted-variable bias in the two-step estimator. The below Corollary formalises this:

Corollary 5. (OVB in QR without Controls) *Under Case 1, the following relates $\beta_{2S}(\tau)$ and $\beta_{1S}(\tau)$:*

$$\begin{aligned}\beta_{2S}(\tau) &= \beta_{1S}(\tau) + \Omega_{QR}(\tau) \\ &= \beta_{1S}(\tau) + \boldsymbol{\phi}_\mathbf{1}(\tau)\frac{\mathbb{E}\left[w_\tau \varepsilon_t \mathbf{x}'_{1,t}\right]}{\mathbb{E}\left[w_\tau \varepsilon_t^2\right]}\end{aligned}$$

where $w_\tau = \int_0^1 f_{u_\tau^{Hyb}}\left[u\,\left(\varepsilon_t\beta(\tau) - \mathbf{x}'_{1,t}\boldsymbol{\phi}_1(\tau) - \varepsilon_t\beta_{hyb}(\tau)\right)|\varepsilon_t, \mathbf{x}'_{1,t}\right]\,\mathrm{d}u/2$.

Proof: This follows from Result 2. □

[17] Linearity here merely helps to simplify the expression for OVB included below. A more general expression under non-lineairty can be found in Angrist et al. (2006).

Note that unlike in an OLS-setting, even though $\mathbb{E}\left[\varepsilon_t \mathbf{x}_{1,t}'\right] = \mathbf{0}$ by construction, this does not imply $\Omega_{QR}(\tau) = 0$. We can still have $\mathbb{E}\left[\omega_\tau \varepsilon_t \mathbf{x}_{1,t}'\right] \neq \mathbf{0}$. In order for OVB to be zero, we need additional assumptions, which are unlikely to hold in practical applications. For example, if the weights w_τ are constant across $\left[\varepsilon_t, \mathbf{x}_{1,t}'\right]$ then it is straightforward to show that $\Omega_{QR}(\tau) = 0$. As Angrist et al. (2006) explain, this will be approximately true when the model for y_t is a pure location model. But this assumption is unlikely to hold in practice given the motivation for using quantile regression to estimate equation (2) rests on the idea that the covariates have differing effects across quantiles, which would be missed by estimation via OLS.

The situation is similar when moving to a setting where auxiliary controls are included in the second-stage regression (i.e., Case 2). In this case, unlike in OLS, we can still get OVB in a quantile regression setting when the shock is constructed correctly to be uncorrelated with other drivers of y_t:

$$\mathbb{E}\left[\varepsilon_t \mathbf{x}_{2,t}'\right] = \mathbf{0} \quad \not\Rightarrow \quad \beta_{2S}(\tau) = \beta_{1S}(\tau)$$

This follows directly from the discussion above: uncorrelatedness is not sufficient to remove OVB in QR. Intuitively, the bias in the two-step procedure in relation to the identification of causal effects arises from two sources. First, it is well-known that identification of quantile treatment effects requires that treatment is *fully* independent (i.e., not just mean-independent) of potential outcomes. This assumption can fail for the two-step approach even when z_t is fully independent of potential outcomes conditional on $\mathbf{x}_{1,t}$, since ε_t may not be fully independent of $\mathbf{x}_{1,t}$ (as would be the case under heteroskedasticity in the first-stage shock-identification regression). Second, even when ε_t is fully independent of potential outcomes, the two-step procedure does not capture the effect on *conditional* quantiles for the same conditioning set as the one-step procedure since $\mathbf{x}_{1,t}$ are excluded from the regression.[18]

4 Empirical Applications

We illustrate our theoretical results with a series of empirical applications, which analyse the dynamic effects of US monetary-policy shocks, controlling for central-bank information in the spirit of Romer and Romer (2004). Focusing on the dynamic responses of US CPI, we illustrate our theoretical results for each of the three estimators discussed in Section 3 in turn: OLS, IV, and QR, covering estimation by both LPs and VARs. We explain how these various applications of one- and two-step approaches have important implications for the vast empirical literature studying the causal effects of monetary policy.

[18]These issues are discussed more formally in Fernández-Gallardo, Lloyd, and Manuel (2023) who demonstrate that, e.g., a QR-LP with controls identifies dynamic causal effects under a similar selection-on-observables assumption as OLS-LP.

For most applications, we use a common set of data. To estimate the first-stage US monetary-policy shocks $\hat{\varepsilon}_t^{mp}$, orthogonalised with respect to central-bank information, we estimate the regression specification of Romer and Romer (2004). Specifically, we construct the shock by regressing the change in the Federal Funds target rate Δr_t on the previous target rate r_{t-1}, as well as past and future Greenbook forecasts of GDP growth Δy_t^e, inflation π_t^e and unemployment u_t^e, as well as their revisions. We use the same functional form as Romer and Romer (2004), specifically:

$$\begin{aligned}\Delta r_t =& \delta_0 + \delta_1 r_{t-1} + \sum_{i=-1}^{2} \left[\delta_{2,i}\Delta y_{t,i}^e + \delta_{3,i}(\Delta y_{t,i}^e - \Delta y_{t-1,i}^e) + \delta_{4,i}\pi_{t,i}^e + \delta_{5,i}(\pi_{t,i}^e - \pi_{t-1,i}^e)\right] \\ &+ \delta_6 u_{t,0}^e + \varepsilon_t^{mp} \end{aligned} \tag{10}$$

This regression is analogous to equation (1) from our theoretical results, where z_t can be interpreted as the change in the Federal Funds target rate and $\mathbf{x}_{1,t}$ includes the previous target rate alongside Greenbook forecasts and forecast revisions (plus a constant).

Relative to the original work of Romer and Romer (2004), we make two minor changes to estimate equation (10). First, we estimate the model for a different sample period—specifically 1972:01-2007:12, rather than 1969:01-1994:12. We start the sample a little later to avoid calendar months in which there was more than one FOMC meeting. We end the sample later given data availability, stopping just before the effective lower bound was reached (using updated data from Wieland and Yang, 2020). Second, rather than estimating the model at meeting frequency, we estimate the shocks at monthly frequency.[19] We do this to ensure direct comparability of the conditioning data in the one- and two-step approaches across all our LP and VAR applications.

In the various specifications we consider, we also include additional auxiliary controls alongside $\hat{\varepsilon}_t^{mp}$, which correspond to $\mathbf{x}_{2,t}$ from our theoretical section. Throughout, we include lags of the following in the set of auxiliary controls: industrial production, consumer prices and the unemployment rate.

4.1 Ordinary Least Squares

We begin by estimating US monetary policy's effects on US CPI by employing the Romer and Romer (2004) shocks in LPs, estimated by OLS, and as internal instruments in recursive SVARs.

[19]To do so, for months in which no FOMC decision occurred, we set the change in the Federal Funds target rate to zero, the Greenbook forecasts equal to their last value, and forecast revisions to zero. In Appendix D, we show how this frequency change constitutes a minimal difference for estimated responses.

Local Projections (LP-OLS). First, we estimate the effects of the shocks using a second-stage LP regression:[20]

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h + \mathbf{x}_{2,t}'\boldsymbol{\alpha}^h + \hat{\varepsilon}_t^{mp}\beta_{2S}^h + u_{t+h}^{2S} \quad (11)$$

where $h = 0, 1, ..., 48$.[21] This is an analog to equation (2). We compare this to a one-step LP regression:

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h + \mathbf{x}_{1,t}'\boldsymbol{\theta}_1^h + \mathbf{x}_{2,t}'\boldsymbol{\theta}_2^h + \Delta r_t\beta_{1S}^h + u_{t+h}^{1S} \quad (12)$$

an analog to equation (4), as well as a hybrid regression, in which the estimated shocks are used alongside all controls:

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h + \mathbf{x}_{1,t}'\boldsymbol{\phi}_1^h + \mathbf{x}_{2,t}'\boldsymbol{\phi}_2^h + \hat{\varepsilon}_t^{mp}\beta_{Hyb}^h + u_{t+h}^{Hyb} \quad (13)$$

where this is the analog to equation (6).

As in our theoretical exposition, we consider two cases: Case 1, in which the second-stage controls $\mathbf{x}_{2,t}$ are empty such that $K_2 = 0$; and Case 2, in which $\mathbf{x}_{2,t}$ is non-empty (i.e. $K_2 > 0$). In this latter case, $\mathbf{x}_{2,t}$ contains one-month lags of month-on-month changes in (log) industrial production, (log) CPI and the unemployment rate.

Table 1 presents the empirical results, which support our main findings from Section 3.1. Focusing on Case 1, in which the set of second-stage auxiliary controls $\mathbf{x}_{2,t}$ is empty, columns (1)-(3) present β_i estimates for $i = \{2S, 1S, Hyb\}$, respectively, and $h = 0, 12, 24, 36, 48$.[22] The coefficient point estimates shown here demonstrate that, in this case, point estimates from the one- and two-step estimators are identical, as stated in Corollary 1. The point estimates, unsurprisingly, indicate that a US monetary policy shock is associated with significantly negative lagged effects on US CPI. Comparing columns (2) and (3) further illustrates that both the point estimates and standard-error estimates from the one-step and hybrid estimators are identical in sample, a result stated in Proposition 1.

However, as Corollary 2 states, the estimated standard errors from the one- and two-step estimators are not the same. In all cases, as columns (1) and (2) show, the OLS standard errors, calculated assuming homoskedasticity are smaller for the one-step estimates relative to the two-step. This finding carries over to other standard errors too, including White (1980) robust standard errors—which admit heteroskedasticity—and Newey and West (1987) standard

[20]Unlike Romer and Romer (2004), who estimate a distributed-lag model in their second stage, we utilise the LP methodology of Jordà (2005) to estimate direct forecasts of US CPI across different horizons.

[21]Strictly, to ensure that the one-step regression is estimated using the same control data, we use outcome data from 1972:01-2011:12 to estimate the forward lags of this regression.

[22]These impulse response functions are also presented in Figures 1a and 1b

Table 1: Response of $\ln(CPI)$ to US monetary policy shock across horizons h from LP-OLS

		Case 1: $x_{2,t}$ empty			**Case 2: $x_{2,t}$ non-empty**		
		(1)	(2)	(3)	(4)	(5)	(6)
		Two-Step	One-Step	Hybrid	Two-Step	One-Step	Hybrid
$h = 0$		0.03	0.03	0.03	0.06	0.04	0.04
	OLS s.e.	(0.05)	(0.04)	(0.04)	(0.04)	(0.03)	(0.03)
	N-W s.e.	(0.06)	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)
	Rob. s.e.	(0.09)	(0.03)	(0.03)	(0.04)	(0.03)	(0.03)
$h = 12$		0.13	0.13	0.13	0.46	0.13	0.13
	OLS s.e.	(0.50)	(0.22)	(0.22)	(0.38)	(0.22)	(0.22)
	N-W s.e.	(0.66)	(0.21)	(0.21)	(0.32)	(0.21)	(0.21)
	Rob. s.e.	(0.78)	(0.20)	(0.20)	(0.32)	(0.19)	(0.19)
$h = 24$		-0.15	-0.15	-0.15	0.37	-0.11	-0.11
	OLS s.e.	(0.90)	(0.41)	(0.41)	(0.74)	(0.41)	(0.41)
	N-W s.e.	(0.93)	(0.44)	(0.44)	(0.57)	(0.44)	(0.44)
	Rob. s.e.	(1.02)	(0.39)	(0.39)	(0.52)	(0.39)	(0.39)
$h = 36$		-1.20	-1.20**	-1.20**	-0.52	-1.13**	-1.13**
	OLS s.e.	(1.26)	(0.53)	(0.53)	(1.07)	(0.53)	(0.53)
	N-W s.e.	(1.20)	(0.67)	(0.67)	(0.97)	(0.65)	(0.65)
	Rob. s.e.	(1.04)	(0.55)	(0.55)	(0.90)	(0.54)	(0.54)
$h = 48$		-2.46	-2.46***	-2.46***	-1.66	-2.38***	-2.38***
	OLS s.e.	(1.57)	(0.62)	(0.62)	(1.36)	(0.63)	(0.63)
	N-W s.e.	(1.67)	(0.93)	(0.93)	(1.49)	(0.91)	(0.91)
	Rob. s.e.	(1.25)	(0.74)	(0.74)	(1.36)	(0.74)	(0.74)

Notes: Estimated response of US $\ln(CPI)$ to US monetary policy shock using Romer and Romer (2004) identification assumptions. Estimated using monthly data for the period 1972:01-2007:12. OLS, Newey and West (1987) and robust standard errors in parentheses. ***, ** and * denote significance at 1, 5 and 10% levels using Newey and West (1987) standard errors, respectively.

errors—which are robust to serial correlation. Because the naïve two-step standard errors are over-estimated, they imply that the dynamic effects of US monetary policy on US CPI are insignificant, even after four years. In contrast, the one-step estimates are significant at the 5%, at least, at the four-year horizon.

For Case 2, where $\mathbf{x}_{2,t}$ is non-empty and the conditions for Corollary 3 are not met, columns (4) and (5) demonstrate that coefficient estimates from the one- and two-step approaches do, in general, differ, confirming Result 1.[23] In this application, the one-step coefficient estimates suggest a limited near-term price puzzle, relative to the two-step estimates—which indicate that a US monetary policy tightening is, counterintuitively, associated with a marginally significant increase in prices in the near term. Further out, the lagged effects of monetary policy are only found to be significantly negative using the one-step (and hybrid) approach.

While comparing columns (1) and (4) demonstrates that adding second-stage controls does generally reduce naïve standard-error estimates from the two-step procedure, the standard errors in column (4) remain greater than those from the one-step approach in column (5). This demonstrates how the implications of Corollary 2 carry over to the case with non-empty $\mathbf{x}_{2,t}$.

[23] Again, consistent with Proposition 1, the point estimates and standard-error estimates from one-step and hybrid approaches are identical.

Internal Instruments in a Recursive SVAR. Next, we estimate the effects of the Romer and Romer (2004) shocks by employing them as 'internal instruments' in a VAR (Plagborg-Møller and Wolf, 2021). The two-step approach we consider amounts to using the shocks directly in a recursive SVAR, with the shock ordered first such that shocks to subsequent variables do not impact it contemporaneously.[24]

Our application uses the VAR specification of Coibion (2012), which Ramey (2016) also uses in their *Handbook of Monetary Economics* chapter. For the two-step approach, we estimate a VAR(12) with (in order): the estimated Romer-Romer shock $\hat{\varepsilon}_t^{mp}$, the Federal Funds target rate (r_t), and the macroeconomic variables (specifically, consumer prices, industrial production, unemployment and commodity prices). We compare this to a one-step approach where the VAR(12) includes the Greenbook variables in $\mathbf{x}_{1,t}$ as additional endogenous variables.[25] We use the new target rate (r_t) as our interest-rate variable, and use the same macroeconomic variables as the two-step. For identification, we place the Federal Funds target rate after the Greenbook forecasts, but before the macroeconomic variables—aligning with the temporal ordering in the two-step, as well as that implied by our selection of controls in the LP-OLS application.[26]

In Appendix B, we analytically demonstrate how our OVB result carries over to this case. Intuitively, bias can arise in the two-step 'internal instruments' approach since the monetary-policy shock in the VAR is constructed by further orthogonalising ε_t^{mp} with respect to lags of the macroeconomic variables, which may in turn be correlated with the Greenbook forecasts used in the first stage. The one-step avoids this issue since the shock estimated in the VAR is by construction orthogonal to the Greenbook forecasts.

Figure 2 presents the estimated IRFs from the one- and two-step recursive VARs. As in the LP-OLS application, the one-step estimates again suggest that US monetary policy has significant (dis)inflationary consequences. The two-step internal-instrument VAR points to a pronounced prize puzzle in the near term, with no significant fall in prices at any subsequent quarter. In contrast, the one-step VAR exhibits a much milder near-term price puzzle, with prices falling significantly in the medium term. Furthermore, the impulse responses for industrial production and unemployment indicate that the two-step estimates also present a more pronounced 'activity puzzle' in the near term (i.e., a significant fall in unemployment and rise in industrial production in response to a monetary tightening), unlike the one-step estimates.

[24]Note that Romer and Romer (2004) order their shock *last* in their VAR when estimating the effects on macro variables. This implies the restriction that monetary policy only affects macro variables with lag, which does not have a solid theoretical basis. We avoid making this assumption throughout.

[25]To avoid including the same variable in both levels and first differences in the VAR, we include only the Greenbook forecasts for that month (not the forecast revisions).

[26]As in the LP-OLS application, we are interested in estimating the effect of monetary policy after partialling-out the contemporaneous Greenbook forecasts and lagged macroeconomic controls. The ordering of variables in the VAR here aligns with the actual temporal ordering of monetary policymaking whereby the Fed's policy decision is made after seeing the Greenbook forecasts (which are produced prior to the FOMC meeting), but before seeing the (end-of-month) readings for each of the macroeconomic variables.

Figure 2: Estimated impulse responses to US monetary policy shock from recursive SVAR

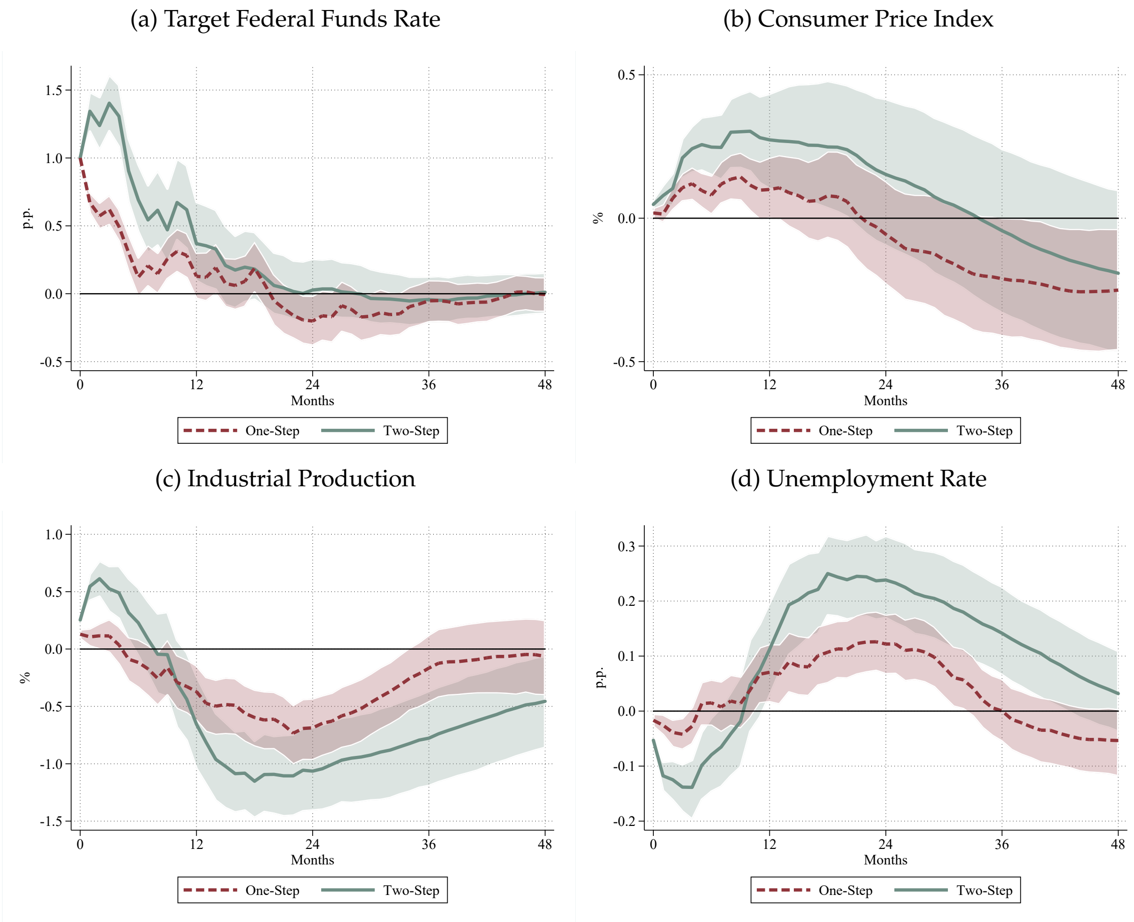

Notes: Estimated response of key US variables to US monetary policy shock, normalised as 1p.p. increase in Federal Funds target rate from one- and two-step recursive SVARs described in main text. Shaded area denotes 68% confidence bands constructed from wild bootstrap. Sample: 1972:01-2007:12

These OLS-based applications highlight how appropriately controlling for Greenbook forecasts can yield estimates of monetary policy's causal effects that are more robust than previously realised. Unlike recent studies that find puzzling results when using the Romer and Romer (2004) shock in both LPs and VARs (Ramey, 2016; Nakamura and Steinsson, 2018a), we find that the response of consumer prices is consistently negative and significant in the medium term—with only a mild near-term 'price puzzle'—when controlling directly for Greenbook forecasts. Crucially, we find these results without resorting to a 'recursiveness assumption' which rules out contemporaneous effects of the monetary shock on macroeconomic variables. In this sense, the implications of our results are very different to Ramey (2016) who finds that "*relaxing the recursiveness assumption imposed by Romer and Romer's hybrid VAR leads to several puzzles*" (p. 111) such that "*even with the Romer and Romer shock, one is forced to make the recursiveness assumption, which does not have a solid economic basis*" (p. 107). Instead, our results highlight that appropriately controlling for the Greenbook forecasts in a one-step procedure need not result in puzzling findings.

4.2 Instrumental Variables

As we outline in Section 3.2, our theoretical results for OLS also carry over to settings in which orthogonalised shocks are used as instruments. We demonstrate this here, considering estimation by both LPs, as well as VARs with external instruments (or 'Proxy SVARs').

Local Projections (LP-IV). To demonstrate how the implications of our IV results from our OLS ones, we employ the Romer and Romer (2004) shock within a LP-IV setting. The two-step approach involves using the shock as an instrument for a monetary-policy indicator, and we compare this to a 'one-step' approach where the Greenbook forecasts are used as controls with the change in the Federal Funds target rate as the instrument. Since this yields similar results to those we described using LP-OLS in the previous sub-section, we defer a discussion of these empirical results to Appendix D, in particular Figure D1.

The more substantive practical implications of our IV results relate to instrument strength. To demonstrate these in the context of US monetary policy, we consider an application that builds on the specification of Miranda-Agrippino and Ricco (2021)—which itself develops that in Romer and Romer (2004). Here, our two-step estimation proceeds as follows. In the first stage, we take high-frequency monetary policy surprises—specifically the move in the third-month-ahead Federal funds futures rate in a 30-minute window around monetary policy announcements, as constructed by Gürkaynak et al. (2005)—at monthly frequency, using the series constructed by Gertler and Karadi (2015). We regress these surprises $z_t := mp_t^{surp}$ on Greenbook forecasts and forecast revisions $\mathbf{x}_{1,t}$ and label the residual from this regression $\hat{\varepsilon}_t^{FF4}$. We then use this orthogonalised residual as an instrument for our monetary-policy indicator m_t in the second stage regression, with the form:

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h + m_t\beta_{2S}^{IV,h} + u_{t+h,IV}^{2S} \quad (14)$$

to estimate the dynamic effects of US monetary policy on US CPI. This is an analog to equation (8). We compare this to a one-step LP-IV regression:

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h + \mathbf{x}_{1,t}'\boldsymbol{\theta}_1^h m_t\beta_{1S}^{IV,h} + u_{t+h,IV}^{1S} \quad (15)$$

in which the surprise mp_t^{surp} is used as an instrument for m_t to make this the analog to equation (9). Throughout, we use the 1-year Treasury yield as our monetary-policy indicator m_t and, due to the availability of high-frequency surprises, we start our sample in 1990:01—but

Table 2: First-stage F-statistics from one- and two-step IV applications

	Two-Step	One-Step
A: LP-IV		
Case 1: No Auxiliary Controls	11.756	19.571
Case 2: With Auxiliary Controls	11.266	19.578
B: External Instruments SVAR		
	7.784	11.358

Notes: First-stage F-statistics from applications where high-frequency monetary-policy surprises are orthogonalised with respect to Greenbook forecasts, à la Miranda-Agrippino and Ricco (2021) using one- and two-step approaches. Panel A reports results from LP-IV applications, equations (14) and (15), with and without auxiliary controls (lagged month-on-month change in (log) industrial production, (log) CPI and unemployment rate. Panel B reports results from structural VAR identified with external instrument. Sample: 1990:01-2007:12.

continue to end it in 2007:12.

We report the impulse responses of consumer prices from this specification in Figure D2 of Appendix D. We find coefficient estimates from the one- and two-step approaches to be very similar, even when including auxiliary controls. The similarity in coefficients between the one- and two-step follows from the fact the Miranda-Agrippino and Ricco (2021) shock has a very low correlation with the lagged macroeconomic variables we employ as auxiliary controls. The standard errors are wider for the two-step approach than the one-step, although responses are generally insignificant for both, consistent with Miranda-Agrippino and Ricco (2021).

However, the differences between first-stage F-statistics from the one- and two-step approaches, shown in Panel A of Table 2, can be substantive. In the two cases, with and without auxiliary controls, we find first-stage F-statistics in this application to be around twice as large using the one-step approach relative to the two-step (around 20 vs. around 10 for the two-step). Panel B of Table 2 also presents the first-stage F-statistics that arise from one- and two-step applications of the Proxy SVAR to the Miranda-Agrippino and Ricco (2021) setting (which we discuss in more detail in the following sub-section). Here, the difference between one- and two-step results is striking: with the two-step, the first-stage F-statistic lies below 10, while the one-step approach yields an F-statistic above that common threshold.

Given the challenge of weak instruments and lack of power in the literature using high-frequency monetary surprises to identify the causal effects of monetary policy (see discussion in, e.g., Nakamura and Steinsson, 2018a), this finding has particular importance. 'Best practice' for identification in the monetary literature typically involves orthogonalising high-frequency surprises with respect to various macroeconomic and financial data and then employing them in LP or VAR specifications (see e.g., Bauer and Swanson, 2022). However, doing so in a two-step approach implies that tests of instrument strength and coefficient significance will be unnecessarily conservative.

External Instruments SVAR. We also demonstrate that our results in Section 3.2 have implications for estimation via Proxy SVARs. As we show formally in Appendix C, using an orthogonalised shock as an external instrument in a SVAR can be viewed as a special case of what we describe as a two-step IV regression with auxiliary controls, and so can generate a form of OVB. We propose an alternate 'one-step' Proxy-SVAR procedure and derive an exact analytical expression for the OVB in the impulse responses from the two-step approach.

For our main application, we employ the Romer and Romer (2004) shocks as external instruments in a VAR(12) with a macro-variable set that includes the 1-year Treasury yield, (log) industrial production, (log) CPI, the unemployment rate and (log) commodity prices for the period 1972:01-2007:12. The two-step approach involves instrumenting the 1-year Treasury yield for the Romer and Romer (2004) shock when estimating contemporaneous responses to changes in monetary policy, then using the reduced-form VAR coefficients to back-out the entire impulse response. In contrast, the one-step approach uses the change in the Federal Funds target rate as an instrument and includes Greenbook forecasts and revisions directly as controls when estimating contemporaneous responses, and then (as in the two-step approach) constructs impulse responses by combining these estimates with the reduced-form VAR coefficients.

We present the impulse responses for the 1-year yield and CPI from the one- and two-step Proxy SVARs in Appendix D, Figure D3. The two-step approach shows a large and significant rise in prices across horizons in response to a monetary-policy shock. Unlike with previous estimates, the one-step approach does very little to offset this price puzzle. Note that the one-step procedure only differs from the two-step procedure in the estimation of the *contemporaneous* response. Hence both estimates rely on an invertibility assumption necessary for identification in Proxy SVARs in order to construct impulse responses. And so one explanation for the sharp price puzzle here is a failure of invertibility. Note that Stock and Watson (2018) propose comparing Proxy-SVAR and LP-IV estimators as a direct test of VAR-invertibility. So, the sharp contrast in impulse responses from our LP-IV and Proxy-SVAR applications strongly suggests that invertibility fails in this setting. The difference in results between the Proxy-SVAR and recursive-SVAR also highlights the potential importance of including the Greenbook forecasts directly in the VAR in order to ensure invertibility. Indeed, across all our specifications, a key conclusion that stands out is that the Greenbook forecasts serve to explain significant variation in future economic outcomes—precisely what makes them desirable as controls in both LP and VAR specifications.

4.3 Quantile Regression

Finally, we study the dynamic response of conditional quantiles of US CPI to US monetary policy shocks. To do so, we estimate QR analogs to equations (11) and (12), where we again employ $\hat{\varepsilon}_t^{mp}$ as our measure of monetary-policy shocks. The two-step approach involves estimating the following second-stage LP-QR :

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h(\tau) + \hat{\varepsilon}_t^{mp}\beta_{2S}^h(\tau) + u_{t+h}^{2S}(\tau) \tag{16}$$

and the one-step is:

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h(\tau) + \mathbf{x}_{1,t}'\boldsymbol{\theta}_1^h(\tau) + \Delta r_t\beta_{1S}^h(\tau) + u_{t+h}^{1S}(\tau) \tag{17}$$

As discussed in Section 3.3, omitted-variable bias can arise in the two-step, even under C1. So we restrict our attention to this case here.

Table 3 presents the estimated response of conditional quantiles of US CPI to US monetary policy across horizons from the one- and two-step approaches, with Figure 1d in the Introduction visualising the results for the 4-year horizon. The table illustrates the key insights from Result 2 and Corollary 5—in particular, the differences in point estimates from the two estimation approaches.

While the one-step estimates indicate that a US monetary policy tightening is associated with a reduction in 3- and 4-year-ahead US CPI across all quantiles, the two-step estimates differ in their implications. The two-step estimates indicate that a US monetary policy tightening is associated with a more marked (and significant) reduction in the right-tail of future inflation. On the other hand, one-step estimates - which remove the OVB term - are more similar across other quantiles. In essence, the one-step point estimates imply that tighter US monetary policy shifts the distribution of CPI outturns to the left in a parallel fashion, while the two-step estimates mistakenly imply uneven effects of monetary policy across the inflation distribution.

These findings provide novel evidence on the effects of monetary policy across quantiles of the inflation distribution. Understanding the effects of monetary policy across the entire distribution of macroeconomic variables (i.e, not just at the mean) is important for effective policymaking, especially so when policymakers are seeking to contain future risks—a point made forcibly by Greenspan (2004).[27] A two-step approach mistakenly implies monetary policy is particularly potent at addressing upside tail risks to inflation in the medium-term, while a one-step approach shows it acts more as a location-shifter of the entire distribution. The

[27] Greenspan (2004) writes: "*the conduct of monetary policy in the United States has come to involve, at its core, crucial elements of risk management [... such that] a central bank needs to consider not only the most likely future path for the economy but also the distribution of possible outcomes about that path. The decision makers then need to reach a judgment about the probabilities, costs, and benefits of the various possible outcomes under alternative choices for policy.*"

Table 3: Response of $\ln(CPI)$ quantiles τ to US monetary-policy shock across horizons h from LP-QR

	Two-Step			One-Step		
	(1)	(2)	(3)	(4)	(5)	(6)
	$\tau = 0.05$	$\tau = 0.5$	$\tau = 0.95$	$\tau = 0.05$	$\tau = 0.5$	$\tau = 0.95$
$h = 0$	0.00	-0.00	0.02	-0.03	0.03	0.10
	(0.14)	(0.12)	(0.04)	(0.06)	(0.03)	(0.10)
$h = 12$	0.49	-0.36	0.05	1.11***	-0.02	-0.57
	(0.99)	(0.87)	(0.38)	(0.36)	(0.17)	(0.37)
$h = 24$	0.69	-0.26	1.68***	0.59	-0.11	-1.19***
	(0.92)	(1.70)	(0.31)	(0.55)	(0.88)	(0.44)
$h = 36$	0.64	0.03	-6.67***	-1.37	-0.59	-0.86
	(1.02)	(2.65)	(0.66)	(1.08)	(0.69)	(0.66)
$h = 48$	0.03	-0.86	-8.59***	-2.02*	-2.35*	-2.69***
	(1.08)	(3.74)	(0.52)	(1.03)	(1.28)	(0.86)

Notes: Estimated response of conditional quantiles τ across horizons h US $\ln(CPI)$ to US monetary policy shock using two-step shock-identification strategy, as well as alternative one-step QR estimator. Estimated using monthly data for the period 1972:01-2007:12. ***, ** and * denote significance at 1, 5 and 10% levels using bootstrapped standard errors, respectively.

one-step approach we advocate for here can also be extended to assess the effects of other policies on the distribution of various outcomes—as in, e.g., Fernández-Gallardo et al. (2023) who estimate the effects of macroprudential policies on the GDP-growth distribution.

5 Conclusions

A common approach to estimating dynamic causal effects in macroeconomics involves estimating the 'shocks first': orthogonalising causal variables of interest with respect to confounding factors; then, using the orthogonalised variables in a second-stage LP or VAR. As we have explained in this paper, this approach subsumes a range of identification approaches and has been applied in a wide range of settings. An alternate one-step approach involves simply including confounding factors as control variables in a regression for the outcome variable.

We have shown, for a general set of estimators, that the two-step 'shock-first' approach can be problematic for both identification and inference relative to the simple one-step approach. In simple OLS settings, the two approaches yield identical coefficients, but two-step inference is unnecessarily conservative. More generally, one- and two-step estimates can differ due to omitted-variable bias in the latter when additional controls are included in the second stage (e.g., VARs with internal or external instruments) or when employing non-OLS estimators (e.g., QR).

In practice, this bias can be substantive. In LP and VAR settings, we show that the one-step approach removes a significant portion of the near-term price puzzle identified in previous studies analysing the response of prices to US monetary-policy shocks. Moreover, we show new evidence on how monetary policy weighs on the tails of the inflation distribution. We achieve this without resorting to a 'recursiveness assumption', which prevents monetary policy from impacting macroeconomic outcomes contemporaneous and have weak economic foundations. Instead, our results highlight the importance of appropriately controlling for forecasts to isolate exogenous changes in monetary policy, but also to explain future economic outcomes. Taken together, these applications indicate that the (dis)inflationary consequences of monetary policy are more robust than previously realised.

References

ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–1289.

AHNERT, T., K. FORBES, C. FRIEDRICH, AND D. REINHARDT (2021): "Macroprudential FX regulations: Shifting the snowbanks of FX vulnerability?" *Journal of Financial Economics*, 140, 145–174.

AL-AMINE, R. AND T. WILLEMS (2023): "Investor Sentiment, Sovereign Debt Mispricing, and Economic Outcomes," *The Economic Journal*, 133, 613–636.

ANGRIST, J. D., V. CHERNOZHUKOV, AND I. FERNÁNDEZ-VAL (2006): "Quantile Regression under Misspecification, with an Application to the U.S. Wage Structure," *Econometrica*, 74, 539–563.

ANGRIST, J. D., O. JORDÀ, AND G. M. KUERSTEINER (2018): "Semiparametric Estimates of Monetary Policy Effects: String Theory Revisited," *Journal of Business & Economic Statistics*, 36, 371–387.

ANGRIST, J. D. AND G. M. KUERSTEINER (2011): "Causal Effects of Monetary Shocks: Semiparametric Conditional Independence Tests with a Multinomial Propensity Score," *The Review of Economics and Statistics*, 93, 725–747.

ANGRIST, J. D. AND J.-S. PISCHKE (2009): *Mostly Harmless Econometrics: An Empiricist's Companion*, no. 8769 in Economics Books, Princeton University Press.

——— (2014): *Mastering 'Metrics: The Path from Cause to Effect*, no. 10363 in Economics Books, Princeton University Press.

AUERBACH, A. AND Y. GORODNICHENKO (2013): "Output Spillovers from Fiscal Policy," *American Economic Review*, 103, 141–46.

BARNICHON, R. AND C. BROWNLEES (2019): "Impulse Response Estimation by Smooth Local Projections," *The Review of Economics and Statistics*, 101, 522–530.

BARNICHON, R. AND G. MESTERS (2020): "Identifying Modern Macro Equations with Old Shocks," *The Quarterly Journal of Economics*, 135, 2255–2298.

BARRO, R. J. (1977): "Unanticipated Money Growth and Unemployment in the United States," *American Economic Review*, 67, 101–115.

BAUER, M. D. AND E. T. SWANSON (2022): "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," in *NBER Macroeconomics Annual 2022, volume 37*, National Bureau of Economic Research, Inc, NBER Chapters.

BRANDÃO-MARQUES, L., R. G. GELOS, M. NARITA, AND E. NIER (2021): "Leaning Against the Wind: An Empirical Cost-Benefit Analysis," CEPR Discussion Papers 15693, C.E.P.R. Discussion Papers.

CHAMPAGNE, J. AND R. SEKKEL (2018): "Changes in monetary regimes and the identification of monetary policy shocks: Narrative evidence from Canada," *Journal of Monetary Economics*, 99, 72–87.

CHARI, A., K. DILTS-STEDMAN, AND K. FORBES (2022): "Spillovers at the extremes: The macroprudential stance and vulnerability to the global financial cycle," *Journal of International Economics*, 136.

CHEN, K., J. REN, AND T. ZHA (2018): "The Nexus of Monetary Policy and Shadow Banking in China," *American Economic Review*, 108, 3891–3936.

CLOYNE, J., C. FERREIRA MAYORGA, AND P. SURICO (2020): "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism," *Review of Economic Studies*, 87, 102–129.

CLOYNE, J. AND P. HÜRTGEN (2016): "The Macroeconomic Effects of Monetary Policy: A New Measure for the United Kingdom," *American Economic Journal: Macroeconomics*, 8, 75–102.

CLOYNE, J., P. HÜRTGEN, AND A. M. TAYLOR (2022): "Global Monetary and Financial Spillovers: Evidence from a New Measure of Bundesbank Policy Shocks," NBER Working Papers 30485, National Bureau of Economic Research, Inc.

COGLIANESE, J., M. OLSSON, AND C. PATTERSON (2023): "Monetary Policy and the Labor Market: A Quasi-Experiment in Sweden," Working Paper 2023-123, University of Chicago, Becker Friedman Institute for Economics.

COIBION, O. (2012): "Are the Effects of Monetary Policy Shocks Big or Small?" *American Economic Journal: Macroeconomics*, 4, 1–32.

COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): "Innocent Bystanders? Monetary policy and inequality," *Journal of Monetary Economics*, 88, 70–89.

CORSETTI, G., A. MEIER, AND G. MÜLLER (2012): "What determines government spending multipliers?" *Economic Policy*, 27, 521–565.

DING, P. (2021): "The Frisch–Waugh–Lovell theorem for standard errors," *Statistics & Probability Letters*, 168.

FALCK, E., M. HOFFMANN, AND P. HÜRTGEN (2021): "Disagreement about inflation expectations and monetary policy transmission," *Journal of Monetary Economics*, 118, 15–31.

FERNÁNDEZ-GALLARDO, , S. LLOYD, AND E. MANUEL (2023): "The transmission of macroprudential policy in the tails: evidence from a narrative approach," ESRB Working Paper Series 145, European Systemic Risk Board.

FORBES, K., D. REINHARDT, AND T. WIELADEK (2017): "The spillovers, interactions, and (un)intended consequences of monetary and regulatory policies," *Journal of Monetary Economics*, 85, 1–22.

FRISCH, R. (1933): "Propagation Problems and Impulse Problems in Dynamic Economics," Allen and Unwin, Economic Essays in Honour of Gustav Cassel, 171–203.

FRISCH, R. AND F. V. WAUGH (1933): "Partial Time Regressions as Compared with Individual Trends," *Econometrica*, 1, 387–401.

GELOS, G., L. GORNICKA, R. KOEPKE, R. SAHAY, AND S. SGHERRI (2022): "Capital flows at risk: Taming the ebbs and flows," *Journal of International Economics*, 134.

GERTLER, M. AND P. KARADI (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

GREENSPAN, A. (2004): "Risk and Uncertainty in Monetary Policy," *American Economic Review*, 94, 33–40.

GÜRKAYNAK, R. S., B. SACK, AND E. SWANSON (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 1.

HANSEN, B. E. (2022): *Econometrics*, no. 8769 in Economics Books, Princeton University Press.

HOLM, M. B., P. PAUL, AND A. TISCHBIREK (2021): "The Transmission of Monetary Policy under the Microscope," *Journal of Political Economy*, 129, 2861–2904.

JENTSCH, C. AND K. LUNSFORD (2019): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States: Comment," *American Economic Review*, 109, 2655–78.

JORDÀ, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

JORDÀ, O., M. SCHULARICK, AND A. M. TAYLOR (2020): "The effects of quasi-random monetary experiments," *Journal of Monetary Economics*, 112, 22–40.

JORDÀ, O. AND A. M. TAYLOR (2016): "The Time for Austerity: Estimating the Average Treatment Effect of Fiscal Policy," *Economic Journal*, 126, 219–255.

KARNAUKH, N. AND P. VOKATA (2022): "Growth forecasts and news about monetary policy," *Journal of Financial Economics*, 146, 55–70.

KILIAN, L. (2009): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, 99, 1053–1069.

LEWIS, D. J. AND K. MERTENS (2022): "Dynamic Identification Using System Projections on Instrumental Variables," Working Papers 2204, Federal Reserve Bank of Dallas.

LINNEMANN, L. AND R. WINKLER (2016): "Estimating nonlinear effects of fiscal policy using quantile regression methods," *Oxford Economic Papers*, 68, 1120–1145.

LOVELL, M. C. (1963): "Seasonal Adjustment of Economic Time Series and Multiple Regression Analysis," *Journal of the American Statistical Association*, 58, 993–1010.

MCKAY, A. AND C. K. WOLF (2023): "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" *Econometrica*, 91, 1695–1725.

MERTENS, K. AND M. O. RAVN (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103, 1212–1247.

METIU, N. (2021): "Anticipation effects of protectionist U.S. trade policies," *Journal of International Economics*, 133.

MIRANDA-AGRIPPINO, S., S. HACIOGLU-HOKE, AND K. BLUWSTEIN (2020): "Patents, News, and Business Cycles," CEPR Discussion Papers 15062, C.E.P.R. Discussion Papers.

MIRANDA-AGRIPPINO, S. AND G. RICCO (2021): "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 13, 74–107.

——— (2023): "Identification with External Instruments in Structural VARs," *Journal of Monetary Economics*, 135, 1–19.

MIYAMOTO, W., T. L. NGUYEN, AND D. SERGEYEV (2018): "Government Spending Multipliers under the Zero Lower Bound: Evidence from Japan," *American Economic Journal: Macroeconomics*, 10, 247–277.

MURPHY, K. AND R. TOPEL (2002): "Estimation and Inference in Two-Step Econometric Models," *Journal of Business and Economic Statistics*, 20, 88–97.

NAKAMURA, E. AND J. STEINSSON (2018a): "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *The Quarterly Journal of Economics*, 133, 1283–1330.

——— (2018b): "Identification in Macroeconomics," *Journal of Economic Perspectives*, 32, 59–86.

NATH, I. B., V. A. RAMEY, AND P. J. KLENOW (2023): "How Much Will Global Warming Cool Global Growth? ," Tech. rep.

NEWEY, W. K. AND K. D. WEST (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

PAGAN, A. (1984): "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review*, 25, 221–47.

PLAGBORG-MØLLER, M. AND C. K. WOLF (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89, 955–980.

RAMEY, V. (2016): "Macroeconomic Shocks and Their Propagation," Elsevier, vol. 2 of *Handbook of Macroeconomics*, 71–162.

ROMER, C. D. AND D. H. ROMER (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94, 1055–1084.

SCHULARICK, M. AND A. TAYLOR (2012): "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008," *American Economic Review*, 102, 1029–61.

STOCK, J. H. AND M. W. WATSON (2018): "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *Economic Journal*, 128, 917–948.

TENREYRO, S. AND G. THWAITES (2016): "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 8, 43–74.

WHITE, H. (1980): "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica*, 48, 817–38.

WIELAND, J. F. AND M.-J. YANG (2020): "Financial Dampening," *Journal of Money, Credit and Banking*, 52, 79–113.

Appendix

A Additional Results on Standard Errors

In this Appendix, we discuss additional results on standard errors.

A.1 Alternative Standard Errors for OLS

First, we discuss extensions of Corollary 2 to other standard-error formulas. We focus on Case 1 (where $\mathbf{x}_{2,t}$ is empty), with all regression coefficients estimated via OLS. Throughout, we use expressions for standard-error formulas for a single variable in a multivariate regression from Ding (2021).

Heteroskedastic-Robust Standard Errors. The estimated standard errors for the estimated one-step and two-step coefficients $\hat{\beta}_{1S}$ and $\hat{\beta}_{2S}$ using White (1980) robust standard-error formulas can be written as:

$$\widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{1S}) = (\hat{\varepsilon}_t'\hat{\varepsilon}_t)^{-1}\hat{\varepsilon}_t'\hat{\mathbf{\Sigma}}_{1S}\hat{\varepsilon}_t(\hat{\varepsilon}_t'\hat{\varepsilon}_t)^{-1} \tag{A1}$$

$$\widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{2S}) = (\hat{\varepsilon}_t'\hat{\varepsilon}_t)^{-1}\hat{\varepsilon}_t'\hat{\mathbf{\Sigma}}_{2S}\hat{\varepsilon}_t(\hat{\varepsilon}_t'\hat{\varepsilon}_t)^{-1} \tag{A2}$$

where $\hat{\varepsilon}_t = [\hat{\varepsilon}_1, ..., \hat{\varepsilon}_T]$, $\hat{\mathbf{\Sigma}}_{1S} = \text{diag}((\hat{\mathbf{u}}_t^{1S})^2)$, $\hat{\mathbf{\Sigma}}_{2S} = \text{diag}((\hat{\mathbf{u}}_t^{2S})^2)$, and $\hat{\mathbf{u}}_t^i = [\hat{u}_1^i, ..., \hat{u}_T^i]'$ for $i = 1S, 2S$. The difference between the estimated variances can then be expressed in the following quadratic form:

$$\widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{2S}) - \widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{1S}) = \mathbf{a}\hat{\mathbf{\Sigma}}\mathbf{a}'$$

where $\mathbf{a} = (\hat{\varepsilon}_t'\hat{\varepsilon}_t)^{-1}\hat{\varepsilon}_t'$ and $\hat{\mathbf{\Sigma}} = \text{diag}((\hat{\mathbf{u}}_t^{2S})^2 - (\hat{\mathbf{u}}_t^{1S})^2)$. This difference is weakly positive if and only if the diagonal matrix $\hat{\mathbf{\Sigma}}$ is positive semi-definite, which requires all elements on the diagonal to be weakly positive, i.e., $(\hat{u}_t^{2S})^2 \gg (\hat{u}_t^{1S})^2$.

While the residual variance is higher for the two-step than the one-step regression by construction (as stated in Corollary 2), this does not imply that each diagonal element in $\hat{\mathbf{\Sigma}}$ is positive. However, since $(\hat{u}_t^{1S})^2$ is on average larger than $(\hat{u}_t^{1S})^2$ by construction—and much larger when $\mathbf{x}_{1,t}$ explains significant variance in y_t—it seems reasonable that in most applications $\widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{2S})$ will indeed be larger than $\widehat{\mathbb{V}\text{ar}}(\hat{\beta}_{1S})$.

Heteroskedastic-and-Autocorrelation-Robust Standard Errors. The estimated standard errors for the estimated one-step and two-step coefficients $\hat{\beta}_{1S}$ and $\hat{\beta}_{2S}$ using Newey and West (1987) autocorrelation-robust standard-error formulas are the same as equations (A1) and (A2), but with the following form for $\hat{\mathbf{\Sigma}}_{1S}$ and $\hat{\mathbf{\Sigma}}_{2S}$:

$$\hat{\mathbf{\Sigma}}_{1S} = (\hat{\mathbf{w}}^{1S}_{|i-j|}(\hat{\mathbf{u}}^{1S}_i)(\hat{\mathbf{u}}^{1S}_j))_{1\leq i,j\leq n}$$
$$\hat{\mathbf{\Sigma}}_{2S} = (\hat{\mathbf{w}}^{2S}_{|i-j|}(\hat{\mathbf{u}}^{2S}_i)(\hat{\mathbf{u}}^{2S}_j))_{1\leq i,j\leq n}$$

Similar to the case above, there is nothing inherent in OLS mechanics to guarantee that the difference between these matrices is positive semi-definite. But, the fact that $(\hat{u}^{1S}_t)^2$ is on average larger than $(\hat{u}^{1S}_t)^2$ by construction will generally tend to inflate standard errors for the two-step vs the one-step approach.

A.2 IV Standard Errors

Here, we prove the final result of Corollary 4 around standard-error formulas for the one- and two-step approach in an IV setting. Specifically, we show the following for estimated standard errors for β^{IV}_{1S} and β^{IV}_{2S} defined in equations (8) and (9):

$$\widehat{\mathbb{V}\text{ar}}(\hat{\beta}^{IV}_{2S}) \geq \widehat{\mathbb{V}\text{ar}}(\hat{\beta}^{IV}_{1S}) \tag{A3}$$

To ease exposition, we introduce the following notation: $A^{\perp B}$ denotes the OLS residual from a regression of A on B. Note, estimated variances for $\hat{\beta}^{IV}_{1S}$ and $\hat{\beta}^{IV}_{2S}$ have the following form under homoskedasticity (see, e.g., Angrist and Pischke, 2014, pp. 140):

$$\widehat{\mathbb{V}\text{ar}}(\hat{\beta}^{IV}_{2S}) = \hat{\sigma}^2_{2S}/\mathbb{V}\text{ar}(\hat{m}^{2S}_t)$$
$$\widehat{\mathbb{V}\text{ar}}(\hat{\beta}^{IV}_{1S}) = \hat{\sigma}^2_{1S}/\mathbb{V}\text{ar}((\hat{m}^{1S}_t)^{\perp \mathbf{x}_{1,t}})$$

where $\hat{\sigma}^2_{2S}$ and $\hat{\sigma}^2_{1S}$ are defined as:

$$\hat{\sigma}^2_{2S} = \mathbb{V}\text{ar}(y_t^{\perp m_t})$$
$$\hat{\sigma}^2_{1S} = \mathbb{V}\text{ar}(y_t^{\perp [m_t, \mathbf{x}_{1,t}]})$$

and $\hat{m}^{2S}_t$ and $\hat{m}^{1S}_t$ are defined as the fitted values from the following OLS regressions:

$$\hat{m}^{2S}_t = \hat{\varepsilon}_t \hat{\beta}_{2S}$$
$$\hat{m}^{1S}_t = z_t \hat{\beta}_{1S} + \mathbf{x}'_{\mathbf{1,t}} \hat{\boldsymbol{\delta}}_1$$

As in Corollary 2, adding covariates to OLS regressions reduces the variance of the error term and so: $\hat{\sigma}^2_{2S} \geq \hat{\sigma}^2_{1S}$.

Finally, we show $\mathbb{V}\text{ar}((\hat{m}^{1S}_t)^{\perp \mathbf{x}_{1,t}}) = \mathbb{V}\text{ar}(\hat{m}^{2S}_t)$ which ensures the inequality in equation

(A3) holds:

$$\begin{aligned}(\hat{m}_t^{1S})^{\perp \mathbf{x}_{1,t}} &= (\hat{\beta}_{1S}z_t + \hat{\boldsymbol{\delta}}_1\mathbf{x_{1,t}})^{\perp \mathbf{x}_{1,t}} \\ &= (\hat{\beta}_{1S}z_t)^{\perp \mathbf{x}_{1,t}} \\ &= \hat{\beta}_{1S}\hat{\varepsilon}_t \\ &= \hat{m}_t^{2S}\end{aligned}$$

where the first line substitutes in definitions, the second line follows from standard OLS-algebra, the third line substitutes in definitions and the final line follows from Frisch-Waugh-Lovell theorem. □

B Internal Instruments in Structural VARs

In this Appendix, we analytically demonstrate how our results for estimation in an OLS setting carry over to SVAR settings in which identification is achieved by employing an orthogonalised shock as an 'internal instrument' (Plagborg-Møller and Wolf, 2021) in a recursive SVAR. In particular, we derive an exact expression linking the estimates of *contemporaneous* responses from a two-step internal-instrument approach to a one-step approach that includes confounding factors directly in a recursive SVAR.

Recursive-SVAR Setting. We define ε_t as in Section 3.1, as the OLS-population residual from a regression of z_t on $\mathbf{x}_{1,t}$. We also define a vector of outcome variables $\mathbf{y}_t = [y_{1,t}, ..., y_{n,t}]'$ which feature in the VAR.

Consider the estimation of impulse responses using ε_t as an internal instrument. The estimated contemporaneous response of the variable $y_{i,t} \in \mathbf{y}_t$ to ε_t is given by the following (population) OLS regression for $i = 1, ..., n$:

$$y_{i,t} = \varepsilon_t \beta_{2S} + \underbrace{\sum_{j=1}^{p} \Gamma_j^{2S}\mathbf{y}_{t-j} + \sum_{j=1}^{p} \pi_j^{2S}\varepsilon_{t-j}}_{=\mathbf{x}'_{2,t}\boldsymbol{\alpha}} + u_t^{2S} \tag{B1}$$

Next, consider a hybrid (OLS) regression for contemporaneous responses defined as:

$$y_{i,t} = \varepsilon_t \beta_{Hyb} + \sum_{j=0}^{p} \boldsymbol{\Phi_j}\mathbf{x}'_{1,t-j} + \underbrace{\sum_{j=1}^{p} \Gamma_j^{Hyb}\mathbf{y}_{t-j} + \sum_{j=1}^{p} \pi_j^{Hyb}\varepsilon_{t-j}}_{=\mathbf{x}'_{2,t}\boldsymbol{\phi_2}} + u_t^{Hyb} \tag{B2}$$

where, as before, the hybrid regression includes additional covariates relative to the two-step and so β_{Hyb} and β_{2S} are related via an OLS-OVB formula.

It is then straightforward to show that β_{Hyb} is equivalent to β_{1S} from a 'one-step' regression that avoids the first-stage construction of ε_t:

$$y_{i,t} = z_t\beta_{1S} + \sum_{j=0}^{p} \Theta_j \mathbf{x}'_{1,t-j} + \sum_{j=1}^{p} \Gamma_j^{1S} \mathbf{y}_{t-j} + \sum_{j=1}^{p} \pi_j^{1S} z_{t-j} + u_t^{1s} \tag{B3}$$

since controlling for z_{t-j} and $\mathbf{x}'_{1,t-j}$ is equivalent to controlling for ε_{t-j} and $\mathbf{x}'_{1,t-j}$, by the same logic that underpins Proposition 1. Note that this 'one-step' regression is equivalent to estimating the contemporaneous responses of $y_{i,t}$ to z_t in a standard recursive SVAR with z_t ordered after $\mathbf{x}_{1,t}$ and before $\mathbf{y}_t$.

Recursive-SVAR Results. Defining $\mathbf{x}_{1,t}^p$ as the vector $[\mathbf{x}'_{1,t}, \mathbf{x}'_{1,t-1}, ..., \mathbf{x}'_{1,t-p}]$, we have the following relationship between β_{2S} from regression (B1) and β_{1S} from regression (B3):

$$\begin{aligned} \beta_{2S} &= \beta_{1S} + \Omega^{SVAR} \\ &= \beta_{1S} + \mathbb{E}\left[\varepsilon_t \mathbf{x}_{2,t}\right] \mathbf{B}^{-1} \mathbb{E}\left[\mathbf{x}_{2,t} \mathbf{x}_{1,t}^p\right] \boldsymbol{\Phi}^p \end{aligned}$$

where the $\mathbf{B}$-matrix is defined analogously to Result 1, and $\boldsymbol{\Phi}^p$ collects the vector of coefficients $[\boldsymbol{\Phi}_1, \boldsymbol{\Phi}_2, ..., \boldsymbol{\Phi}_p]$ from regression (B2). The OVB term now captures the omission of contemporaneous *and lagged* $\mathbf{x}'_{1,t}$ from the two-step approach.

C External Instruments in Structural VARs

In this Appendix, we demonstrate how our results for IV estimation carry over to SVAR settings in which identification is achieved through external instruments (i.e., SVAR-IV / Proxy-SVAR). We demonstrate analytically how a 'one-step' procedure can be implemented as an alternative to the common 'two-step' procedure that first constructs orthogonalised shocks and then using these shocks as instruments in an SVAR-IV. We also derive an exact expression for the OVB for impulse responses estimated from the two-step approach relative to our proposed one-step approach.

SVAR-IV Setting. We are interested in estimating the effect of m_t on a $n \times 1$ vector of variables $\mathbf{y}_t$ and propose doing so using SVAR-IV. We define $\mathbf{w}_t = [m_t, \mathbf{y}'_t]'$. The two-step approach involves using ε_t—defined as in Section 3.2, as the OLS-population residual from a regression of z_t on $\mathbf{x}_{1,t}$—as an external instrument for m_t.

Following Stock and Watson (2018) (p. 932), SVAR-IV coefficients for this two-step-approach can be defined as follows. First the (population) contemporaneous coefficients are defined via the following IV-regression for each variable $w_{i,t} \in \mathbf{w}_t$, where $i = 1, ..., n, n+1$:

$$w_{i,t} = m_t \beta_{0,i1}^{2S} + \underbrace{\sum_{j=1}^{p} \mathbf{w}'_{t-j} \boldsymbol{\alpha}_j}_{\equiv \mathbf{x}'_{2,t}\boldsymbol{\alpha}} + u_t^{2S} \tag{C1}$$

with ε_t as an instrument for m_t. It is immediately obvious that equation (C1) is just a special case of equation (8), setting $\mathbf{x}_{2,t}$ as the p vectors of lagged controls $\mathbf{w}_{t-j}$.

The (population) impulse response of the vector $\mathbf{w}_t$ to a shock to m_t at horizon h is then:

$$\boldsymbol{\Phi}_{h,1}^{2S} = \mathbf{C}_h \boldsymbol{\beta}_{0,1}^{2S} \tag{C2}$$

where $\boldsymbol{\beta}_{0,1}^{2S}$ is a $[(n+1) \times 1]$ vector collecting each $\beta_{0,i1}$ and $\mathbf{C}_h$ is a horizon-specific coefficient matrix formed by inverting the (population) reduced form-VAR:

$$A(L)\mathbf{w}_t = \boldsymbol{\eta_t} \tag{C3}$$

where $A(L) = I - A_1 L – A_2 L^2 - ...$ and L is the lag operator.

Following the logic of Section 3.2, the contemporaneous responses of each variable could instead be recovered via a one-step IV (population) regression with $\mathbf{x}_{1,t}$ as controls:

$$w_{i,t} = m_t \beta_{0,i1}^{1S} + \mathbf{x}'_{1,t}\boldsymbol{\theta}_1 + \underbrace{\sum_{j=1}^{p} \mathbf{w}'_{t-j} \boldsymbol{\theta}_j}_{\equiv \mathbf{x}'_{2,t}\boldsymbol{\theta_2}} + u_t^{1S} \tag{C4}$$

with z_t as an instrument for m_t. Again, this is just a special case of equation (9) setting $\mathbf{x}_{2,t}$ as the p vectors of lagged controls $\mathbf{w}_{t-j}$.

In this case, the entire impulse response can then be constructed as before using the same reduced-form VAR coefficients as equation (C2) to project-out across horizons:

$$\boldsymbol{\Phi}_{h,1}^{1S} = \mathbf{C}_h \boldsymbol{\beta}_{0,1}^{1S} \tag{C5}$$

SVAR-IV Results. Comparing equations (C5) and (C2), impulse responses from a one- and two-step approach differ only in their construction of the contemporaneous coefficients $\boldsymbol{\beta}_{0,1}$. As in Section 3.2, our OVB result applies directly, and it is then straightforward to derive an exact expression for the OVB of the entire two-step impulse response:

$$\begin{aligned} \mathbf{\Phi}_{h,1}^{2S} &= \mathbf{\Phi}_{h,1}^{1S} + \Omega_h^{SVAR-IV} \\ &= \mathbf{\Phi}_{h,1}^{1S} + \mathbf{C}_h \frac{\mathbf{\Omega}_{\boldsymbol{w}}^{\boldsymbol{OLS}}}{\Omega_m^{OLS}} \end{aligned}$$

where $\mathbf{\Omega}_{\boldsymbol{w}}^{\boldsymbol{OLS}}$ and Ω_m^{OLS} are $[(n+1)\times 1]$ and scalar OLS-OVB formulas respectively, relating to regressions with the stacked vector $\mathbf{w_t}$ and the scalar m_t as the dependent variables. Specifically, the OVB formulas in this case take the following form:

$$\begin{aligned} \Omega_{w_i}^{OLS} &= \mathbb{E}\left[\varepsilon_t \mathbf{x}_{2,t}\right] \mathbf{B}_{w_i}^{-1} \mathbb{E}\left[\mathbf{x}_{2,t}\mathbf{x}_{1,t}'\right] \phi_{w_i} \\ \Omega_m^{OLS} & \mathbb{E}\left[\varepsilon_t \mathbf{x}_{2,t}\right] \mathbf{B}_m^{-1} \mathbb{E}\left[\mathbf{x}_{2,t}\mathbf{x}_{1,t}'\right] \phi_m \end{aligned}$$

where $\Omega_{w_i}^{OLS}$ is the i-th element of $\mathbf{\Omega}_{\boldsymbol{w}}^{\boldsymbol{OLS}}$, and $\mathbf{B}$-matrices and ϕ-coefficients are defined analogously to Result 1. Intuitively, this bias can be thought of as a potential failure of exogeneity conditions necessary for identification in an IV setting, which arises when the instrument ε_t is in fact correlated with other variables (i.e., lags of $\mathbf{w}_t$) that affect the outcome variable. Although these variables are included as controls in equation (C1), this serves to reintroduce correlation with $\mathbf{x}_{1,t}$, thereby leading to a potential failure of exogeneity. As before, the one-step approach automatically avoids this bias and so can be thought of as more robust.

Note, the one-step SVAR-IV approach we discuss here is distinct from simply including $\mathbf{x}_{1,t}$ as exogenous variables in an SVAR-IV since, in our case, $\mathbf{x}_{1,t}$ are included only to estimate contemporaneous coefficients and do not feature in the estimation of subsequent impulse responses. While the two-step approach will continue to over-state first-stage F-statistics (typically computed using standard OLS formulas), standard errors for SVAR-IV are typically computed using a bootstrap procedure—where it is less clear whether such a procedure would produce wider confidence bands for the one-step or two-step approach.

D Empirical Application: Additional Results

In this Appendix, we provide more information underpinning our empirical application in Section 4.

D.1 Data Sources

As described in Section 4, we use monthly data for our empirical estimation. Our dependent variable, the US Consumer Price Index (CPI), is sourced from *FRED*, and we also use additional macroeconomic controls—specifically seasonally-adjusted US industrial production and the US unemployment rate—from the same source. These controls comprise $\mathbf{x}_{2,t}$ in our study.

To estimate Romer and Romer (2004) US monetary policy shocks, we use Federal Reserve Greenbook forecasts and forecast revisions. We draw on Wieland and Yang (2020) for this, who provide updated Greenbook forecast data up to, and beyond, the end of our sample period, 2007:12.

D.2 Monetary-Policy Shock Construction

To construct the Romer and Romer (2004) shocks, we make two changes relative to the original work. First, and most notably, we estimate the model for a different sample period—specifically 1972:01-2007:12, rather than 1969:01-1994:12. We start the sample a little later to avoid calendar months in which there was more than one FOMC meeting. And we end the sample later given data availability, stopping just before the effective lower bound was reached. Second, rather than estimating the model at meeting frequency, we estimate the shocks at monthly frequency.

Following Romer and Romer (2004), we construct the shock by estimating regression (10). Column (1) of Table D1 presents the estimated coefficients from this regression, estimated for the period 1972:01-2007:12.

Monthly vs. Meeting Frequency. As discussed in Section 4, we estimate all the regressions in the main body of our chapter at monthly frequency to foster comparability between LPs and VARs. To do so, for months in which no FOMC announcements occurred, we set the change in the Federal Funds target rate to zero, the Greenbook forecasts equal to their last value, and forecast revisions to zero.

This change constitutes a minimal difference. To support this, Column (2) of Table D1 presents first-stage regression coefficients estimated at meeting frequency. They are similar to monthly-frequency estimates in Column (1). In addition, the implied impulse responses are very similar too. To show this, we re-estimate equations (11) and (12) using meeting-frequency observations (but continuing to project the LP forward in monthly horizons). Table D2 demonstrates how, for the various cases presented in Section 4.1, estimated impulse responses are similar when estimated at monthly and meeting frequency.

D.3 Additional Results for IV Application

LP-IV Application with Romer-Romer Shocks. For this application, we estimate the Romer-Romer shock $\hat{\varepsilon}_t^{mp}$ as before, but we now use this shock as an instrument for the 1-Year Treasury

Table D1: First-Stage Regressions: The Romer-Romer Reaction Function

	DEP. VAR.: Change FFR Target	
	(1)	(2)
	Monthly Frequency	Meeting Frequency
Old FFR Target	-0.014	-0.018
	(0.009)	(0.013)
Output forecasts		
$k = -1$	0.002	0.001
	(0.008)	(0.011)
$k = 0$	0.008	0.013
	(0.013)	(0.021)
$k = 1$	0.016	0.023
	(0.020)	(0.030)
$k = 2$	0.019	0.016
	(0.023)	(0.032)
Inflation forecasts		
$k = -1$	0.016	0.032
	(0.015)	(0.023)
$k = 0$	-0.029	-0.043
	(0.021)	(0.031)
$k = 1$	0.019	0.028
	(0.041)	(0.066)
$k = 2$	0.030	0.026
	(0.048)	(0.078)
Unemployment forecasts		
$k = 0$	-0.037***	-0.050***
	(0.011)	(0.014)
Output forecast revisions		
$k = -1$	0.039	0.043
	(0.026)	(0.028)
$k = 0$	0.129***	0.128***
	(0.032)	(0.034)
$k = 1$	0.032	0.017
	(0.044)	(0.044)
$k = 2$	0.011	0.014
	(0.046)	(0.049)
Inflation forecast revisions		
$k = -1$	0.069	0.050
	(0.045)	(0.044)
$k = 0$	-0.007	-0.005
	(0.055)	(0.056)
$k = 1$	0.030	0.022
	(0.090)	(0.107)
$k = 2$	-0.054	-0.056
	(0.087)	(0.105)
R^2	0.274	0.294
Observations	432	318

Notes: Estimated policy reaction functions. Column (1) estimated using monthly data for the period 1972:01-2007:12. Column (2) estimated using meeting frequency data over the same period. Meeting-frequency observations converted to monthly frequency by setting change in FFR target and forecast revisions to 0 and forecasts equal to their previous-meeting value in months without FOMC meeting. Robust standard errors in parentheses. ***, ** and * denote significance at 1, 5 and 10% levels, respectively.

Table D2: Response of $\ln(CPI)$ to monthly- and meeting-frequency US monetary policy shocks across horizons h estimated by LP-OLS

	Case 1: $x_{2,t}$ empty				**Case 2: $x_{2,t}$ non-empty**			
	Two-Step		One-Step		Two-Step		One-Step	
Frequency	Monthly	Meeting	Monthly	Meeting	Monthly	Meeting	Monthly	Meeting
$h = 0$	0.03	0.03	0.03	0.03	0.06	0.06	0.04	0.04
	(0.05)	(0.05)	(0.04)	(0.04)	(0.04)	(0.04)	(0.03)	(0.04)
$h = 12$	0.13	0.13	0.13	0.13	0.46	0.41	0.13	0.16
	(0.50)	(0.53)	(0.22)	(0.23)	(0.38)	(0.40)	(0.22)	(0.24)
$h = 24$	-0.15	-0.24	-0.15	-0.24	0.37	0.18	-0.11	-0.15
	(0.90)	(0.97)	(0.41)	(0.43)	(0.74)	(0.79)	(0.41)	(0.44)
$h = 36$	-1.20	-1.38	-1.20**	-1.38**	-0.52	-0.85	-1.13**	-1.25**
	(1.26)	(1.37)	(0.53)	(0.55)	(1.07)	(1.16)	(0.53)	(0.56)
$h = 48$	-2.46	-2.72	-2.46***	-2.72***	-1.66	-2.10	-2.38***	-2.60***
	(1.57)	(1.72)	(0.62)	(0.63)	(1.36)	(1.49)	(0.63)	(0.64)

Notes: Estimated response of US $\ln(CPI)$ to US monetary policy shock using Romer and Romer (2004) identification assumptions. Estimated using monthly- and meeting-frequency data for the period 1972:01-2007:12. OLS standard errors presented here in parentheses. , and denote significance at 1, 5 and 10% levels, respectively.

yield (which we denote m_t to align with Section 3.2). In the two-step, we use $\hat{\varepsilon}_t^{mp}$ as an instrument for m_t in regression (14) to estimate the dynamic effects of US monetary policy on US CPI. We compare this to the one-step LP-IV regression (15), in which Δr_t is used as an instrument for m_t.

Figure D1 presents the estimated impulse responses from the one- and two-step. Unsurprisingly, given the discussion in Section 3.2, the estimates for both approaches align closely with those attained from LP-OLS—as a comparison of Figure D1 with Figures 1a and 1b reveals. As before, unlike two-step estimate, the price puzzles in one-step estimates are limited.

LP-IV Application with Monetary Surprises. Figure D2 presents the estimates impulse responses from the Miranda-Agrippino and Ricco (2021) LP-IV application described in Section 4.2.

Proxy-SVAR Application with Romer-Romer Shocks. Figure D3 presents the estimates impulse responses from the Romer and Romer (2004) Proxy-SVAR application described in Section 4.2.

Figure D1: Estimated impulse responses of $\ln(CPI)$ to US monetary policy shock from in LP-IV

(a) Case 1: No Auxiliary Controls

(b) Case 2: With Auxiliary Controls

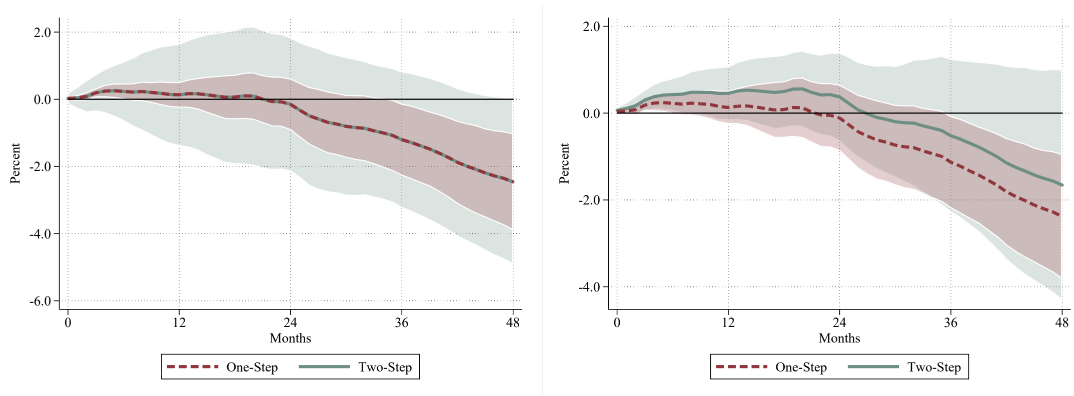

Notes: Estimated response of US $\ln(CPI)$ to US monetary policy shock that leads to 1p.p. increase in 1-year Treasury yield, instrumented with Romer and Romer (2004) shock, as described in main body. Shaded area denotes 90% confidence bands constructed from heteroskedasticity and autocorrelation robust standard errors. Sample: 1972:01-2007:12

Figure D2: Estimated impulse responses of $\ln(CPI)$ to US monetary policy shock from in LP-IV

(a) Case 1: No Auxiliary Controls

(b) Case 2: With Auxiliary Controls

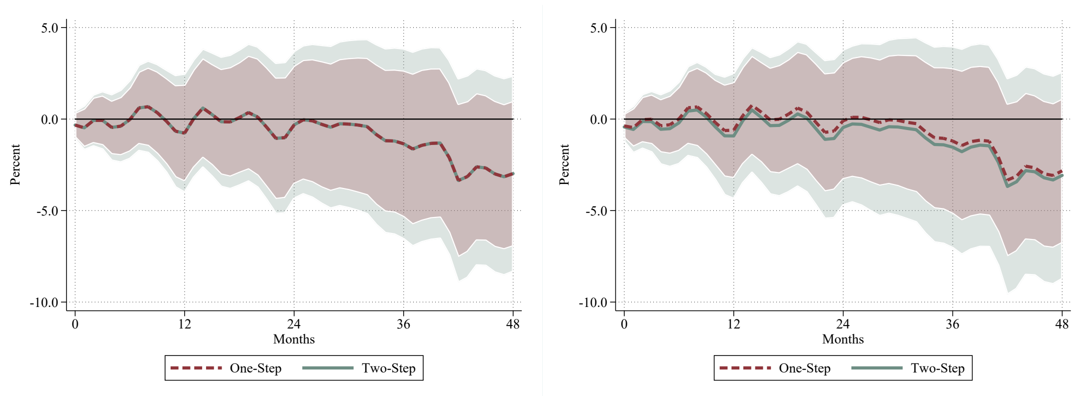

Notes: Estimated response of US $\ln(CPI)$ to US monetary policy shock that leads to 1p.p. increase in 1-year Treasury yield, instrumented with Miranda-Agrippino and Ricco (2021) shock, as described in main body. Shaded area denotes 90% confidence bands constructed from heteroskedasticity and autocorrelation robust standard errors. Sample: 1990:01-2007:12.

Figure D3: Estimated impulse responses to US monetary policy shock from Proxy SVAR

(a) 1-Year Treasury Yield (b) Consumer Price Index

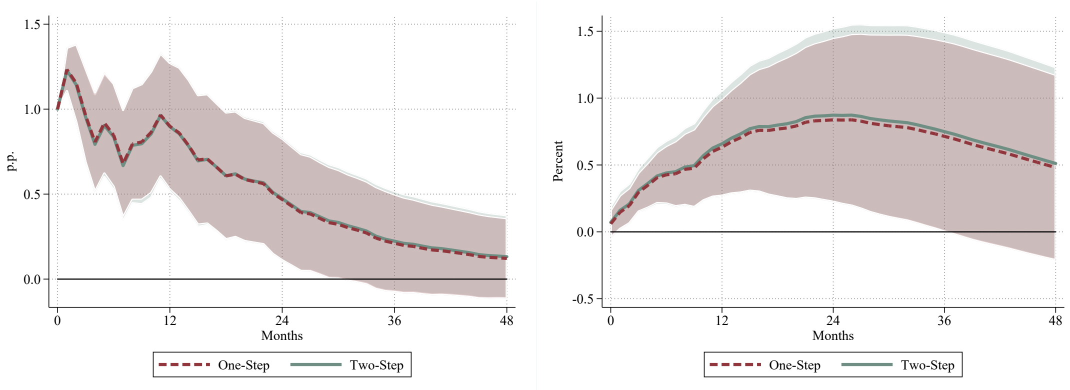

Notes: Estimated response of key US variables to US monetary policy shock, normalised as 1p.p. increase in effective federal funds rate from one- and two-step Proxy-SVAR described in main text. Shaded area denotes 90% confidence bands constructed from Jentsch and Lunsford (2019) residual-based moving block bootstrap.