

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

Simon Lloyd[†] Ed Manuel[‡]

January 29, 2026

Abstract

Many studies in macroeconomics first construct orthogonalised ‘shocks’, then integrate them into local projections or vector autoregressions. For a general set of estimators, we show analytically that this two-step approach can be problematic for identification and inference relative to a one-step procedure that adds controls directly to outcome regressions. Estimated coefficients from a two-step approach suffer from omitted variable bias when including additional controls in the second stage or when employing non-OLS estimators. When coefficient estimates are equivalent, two-step inference is unnecessarily conservative. One-step estimates indicate that the (dis)inflationary consequences of US monetary policy are more robust than previously realised.

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, Christian Wolf and anonymous referees, 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 causal effects is crucial in dynamic macroeconomics. A key challenge 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 controls in a regression with the outcome variable. However, practitioners often take an alternate two-step route: first estimating ‘shocks’ as residuals from a regression of the causal variable on confounding factors, and then using these shocks in local projections (LPs) or vector auto-regressions (VARs).

This two-step ‘shock-first’ approach is widespread in macroeconomics, where the construction of a ‘shock’ series is typically viewed as essential when estimating dynamic causal effects (e.g., [Ramey, 2016](#); [Nakamura and Steinsson, 2018b](#)). The approach is used extensively in the monetary-policy literature—most famously in [Romer and Romer \(2004\)](#) and subsequent studies controlling for central-bank forecasts or other policy-rule variables.¹ A similar two-step method is used to ‘clean’ high-frequency monetary-policy ‘surprises’ of predictability before using them in instrumental variable (IV) regressions or VARs (e.g., [Miranda-Agrippino and Ricco, 2021](#); [Bauer and Swanson, 2023](#); [Karnaugh and Vokata, 2022](#)). The method is also used to estimate the effects of other policies (e.g., fiscal, macroprudential, and trade) and non-policy variables (e.g., oil-price, technology, sentiment and climate shocks), as well as to identify parameters in structural equations.² While many studies have relied on OLS (or IV), two-step approaches have also proved popular in quantile-regression (QR) settings for estimating the causal effects on conditional quantiles.³

In this paper, we argue that the popular two-step approach is problematic and propose an alternative. The crux of our argument is simple. It is well-known that, with confounding factors, identification can be achieved via one-step multivariate regression using confounders as controls. So, we compare coefficients and standard errors from one- and two-step approaches. Conventional wisdom holds that the two are equivalent by the Frisch-Waugh-Lovell (FWL) Theorem ([Frisch and Waugh, 1933](#); [Lovell, 1963](#)), but this equivalence only holds for estimated OLS coefficients (not standard errors), and only in simple settings rarely used in the literature. Across a range of applications and estimators, we demonstrate that two-step estimation

¹E.g.: [Coibion \(2012\)](#); [Cloyne and Hürtgen \(2016\)](#); [Tenreyro and Thwaites \(2016\)](#); [Coibion et al. \(2017\)](#); [Champagne and Sekkel \(2018\)](#); [Chen et al. \(2018\)](#); [Cloyne et al. \(2020\)](#); [Falck et al. \(2021\)](#); [Holm et al. \(2021\)](#); [Cloyne et al. \(2022\)](#); [Coglianese et al. \(2023\)](#). More generally, the two-step estimation has a longer tradition in the monetary-policy literature going back to at least [Barro \(1977\)](#).

²Fiscal policy examples include: [Corsetti et al. \(2012\)](#); [Auerbach and Gorodnichenko \(2013\)](#); [Miyamoto et al. \(2018\)](#); [Barattieri et al. \(2023\)](#). Macropurudential policy examples include: [Forbes et al. \(2017\)](#); [Ahnert et al. \(2021\)](#); [Chari et al. \(2022\)](#). [Barattieri and Cacciatore \(2023\)](#) and [Metiu \(2021\)](#) apply the two-step approach to trade policy. Other examples include studies on the effects of shocks to: oil prices (e.g., [Kilian, 2009](#)), sentiment (e.g., [Al-Amine and Willems, 2023](#)), technology (e.g., [Miranda-Agrippino et al., 2020](#)) and temperature (e.g., [Nath et al., 2023](#); [Bilal and Käning, 2024](#)). The [Romer and Romer \(2004\)](#) shock has also been used to identify the parameters of the Philips Curve (e.g., [Barnichon and Mesters, 2020, 2021](#); [Lewis and Mertens, 2022](#)).

³E.g.: [Linnemann and Winkler \(2016\)](#); [Brandão-Marques et al. \(2021\)](#); [Gelos et al. \(2022\)](#).

can impair identification and inference. Further, we describe how one-step estimation can be implemented in common macroeconomic applications, covering both LPs and VARs.

We set the scene by first outlining identifying assumptions necessary for uncovering dynamic causal effects with control variables. We provide sufficient conditions under which one-step regression coefficients can be interpreted as impulse responses to a structural shock, even though this method does not explicitly construct a ‘shock’ series. Since a one-step OLS regression can be viewed as a special case of identification with external instruments (Stock and Watson, 2018), we highlight that the one-step approach identifies impulse responses even when the researcher is unable to recover exactly the structural shock of interest.

We then formalise the difference between coefficients estimated via the one- and two-step approaches without assumptions about the true data-generating process or underlying causal structures. We note that in general settings the difference between estimated one- and two-step coefficients can be expressed via an omitted-variable bias (OVB) formula.⁴ This arises as the two-step approach excludes potentially relevant variables (i.e., the confounding factors the researcher wishes to partial out) that are included in the first stage, but then excluded from the second. This result is general, applying to a range of estimators defined as the unique minimum 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.

The differences between the one- and two-step approaches in OLS and IV settings follow from essentially elementary applications of FWL and OVB. A direct application of FWL reveals that (in OLS/IV settings) when the outcome variable is 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, two-step standard errors will be *over-estimated* if the confounders have explanatory power for the outcome variable. This result follows directly from simply comparing well-known standard-error formulas, although we are not aware that this point has been previously noted in the macroeconomics literature.

When auxiliary controls are used in the second stage—common in LPs or VARs—the coefficient equivalence no longer applies. If the shock is orthogonal to the auxiliary controls, the OVB term is zero and the one- and two-step approaches will identify the same population parameter—though issues with two-step standard-errors will remain. If the shock is correlated with auxiliary controls, then OVB can be non-zero as the two-step effectively fails to partial out the confounders used in the first stage. We show that these drawbacks of the two-step approach also carry over to IV settings, covering cases where orthogonalised shocks are used as ‘external’ instruments in LPs or VARs.

In more general settings, including QR, the two-step approach is more problematic. We

⁴Throughout, 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).

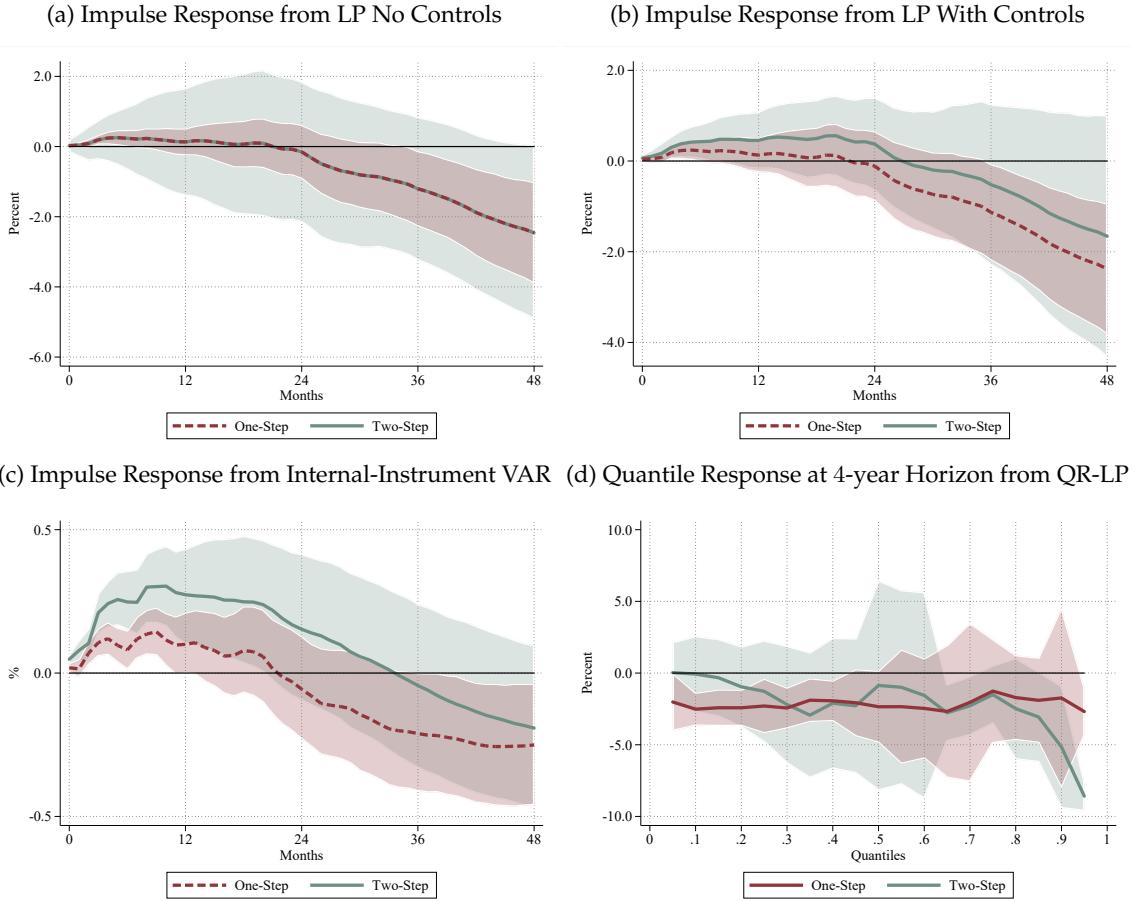
use our OVB result to provide an explicit formula for the bias in the two-step approach in QR settings, and demonstrate that it will non-zero even in applications without auxiliary controls.

Main Applications. Our discussion of OVB in the two-step approach has implications for a range of techniques used in applied macroeconomics 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 [Romer and Romer \(2004\)](#) approach: constructing monetary ‘shocks’ by regressing changes in the Federal Funds target rate on Greenbook forecasts. 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 IRFs to simply estimating a one-step LP with appropriate controls—although the two-step approach leads to overly wide standard errors (Figure 1a). Projecting h -period-ahead inflation on the [Romer and Romer \(2004\)](#) shock does not produce significant responses 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 (lagged) controls in LPs alongside the shock measure. 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 first-stage orthogonalisation. We demonstrate this in Figure 1b, using lagged CPI, industrial production and unemployment as controls. Here, 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, delivers a significant reduction in CPI at longer horizons. This issue is particularly relevant in light of recent literature that highlights the importance of including a rich set of covariates in local projections for both inference and identification (e.g., [Montiel Olea and Plagborg-Møller, 2021](#); [Montiel Olea et al., 2024](#)). We stress that including many controls in a LP as part of a two-step procedure can, in fact, fail to identify dynamic causal effects under standard identifying assumptions.

Second, our results extend to recursive structural-VAR (SVAR) estimation. As in the LP with auxiliary controls, OVB can arise in VARs since lags of the endogenous variables can serve to ‘undo’ some of the first-stage orthogonalisation. We demonstrate this using 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 uses the Green-

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



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 6 and Appendix F.

book forecasts as variables in the recursive SVAR, ordered before the Federal Funds target rate, in turn ordered before the macro variables. 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 (Figure 1c).

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 external-instruments applications use constructed residuals as instruments, which can generate issues for both inference and identification. Within LPs, 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 use 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 then 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 OLS standard-error results, the two-step approach underestimates 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 delivers F -statistics that are around twice as large. This has important implications for the monetary-policy high-frequency literature, where current best practice orthogonalises surprises with respect to macro-financial data before integrating them in LPs or VARs, and where weak instruments and low power are a perennial issue ([Bauer and Swanson, 2023](#)).

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, 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 (Figure 1d). 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.

One issue that can arise in monetary policy applications is that the frequency of FOMC meetings (typically, 6-8 weekly) differs from that of common macroeconomic outcomes of interest (typically, monthly or quarterly). ⁵ We sidestep this issue in our main application by simply converting all data to the desired frequency prior to the shock-estimation stage. For papers that employ a two-step procedure, a common approach is to instead first estimate shocks at meeting frequency, then convert these shocks to the correct frequency to estimate IRFs by variously summing values within months or imposing zero values. ⁶ We show these two approaches yield almost identical estimates for the [Romer and Romer \(2004\)](#) shock. A full

⁵As we detail in footnotes 1 and 2, the two-step procedure is common in monetary policy settings, although it is also commonly used in various other settings even when mixed-frequency data is no concern.

⁶Although it remains common in applied work, we are not aware of any theoretical justification for this approach to handling mixed-frequency data. There exists a broad literature on temporal aggregation and mixed-frequency estimation that suggests alternate methods (see [Chudik and Georgiadis, 2022](#), for a recent example).

treatment of the optimal methods for handling mixed frequency data in macroeconomics are beyond the scope of this paper. However, in Appendix F we discuss the issue in more depth, where we additionally demonstrate a simple alternate approach for handling meeting- and monthly-frequency data which permits estimation via a one-step local projection without the need to first convert data between frequencies. We also discuss potential avenues for future research.

Further Applications and Simulations. How significant and widespread are the issues we highlight? We provide exact analytical formulas for the difference between one- and two-step standard errors and coefficients, highlighting that any differences will depend crucially on the researcher’s choice of how to partition controls across the first- and second-stages of the two-step approach. Nevertheless, we seek to provide some guidance on this question in two ways. First, we conduct a simulation study using a medium-scale VAR model with recursively identified monetary policy shocks to assess the quantitative significance of our results in an empirically realistic setting. We estimate impulse responses to a monetary policy shock using both a one-step and two-step LP estimator, generating hundreds of comparisons across outcome variables and horizons. We find that, without auxiliary controls, estimated standard errors for the two-step estimator are considerably larger than the one-step estimator: over 50% larger for a stationary DGP and over 500% larger for a non-stationary DGP. We then consider a case where a researcher constructs a monetary policy shock by regressing the Fed Funds rate only on the (appropriate) contemporaneous variables in the VAR and then relies on including lagged controls in the second-stage LP to achieve identification. In this case, the finite-sample bias of the two-step LP estimator is around *ten* times larger across DGPs than the bias of the one-step estimator at short horizons, and 40 – 100% larger at longer horizons. Second, we extend our empirical applications, revisiting the identification strategy for a total of five well-known studies that employed a two-step approach, generating over 500 comparisons between the one-step and two-step coefficients and standard errors. We find that, on average, without auxiliary controls, standard errors are around 50 – 70% larger in the two-step approach (reaching 200 – 400% on average over short and long horizons in some cases), broadly in line with our simulation study for stationary DGPs. Somewhat reassuringly, we find that with auxiliary controls, across these studies, estimated coefficients between a one-step and two-step approach differ by less than half a standard error on average. We stress however that even for these studies, the bias can be large for particular variables and horizons – up to around 3 standard errors in worst cases.

Literature. To the best of our knowledge, our paper is the first to explicitly highlight these specific drawbacks of the two-step approach—both that it *overestimates* standard errors and generally suffers from OVB, which biases coefficient estimates. It is clearly well-established

that the one- and two-step approach yield identical coefficient estimates in simple OLS settings without auxiliary controls, a result which follows directly from FWL.⁷ However, we are not aware of previous discussion of the bias in coefficient estimates that can arise with a two-step approach. We derive a formula for this bias for a two-step OLS estimator with auxiliary controls which drops out of simple OLS-algebra. Further, we derive novel analytical formulas for the bias in two-step approaches in more complex settings—covering, e.g., Proxy-VARs and QR-LPs—which go beyond any previously established results in the literature.

Our findings for standard-error estimation also relate to a broad literature on inference in econometric models with generated regressors. It is well-known that, as a general rule, generated regressors and two-step estimation affects sampling distributions and variance matrices (Hansen, 2022). We show that, in an OLS setting, for a two-step procedure that projects on a constructed shock without auxiliary controls, OLS standard errors will be over-estimated (*vis-à-vis* a one-step regression). This follows directly from simply comparing well-known standard error formulas—and is often taught as a close-corollary of FWL—but, as far as we are aware, has not previously been discussed in the literature. Indeed, previous literature on generated regressors instead stresses that when projecting on a constructed residual (without controls), OLS standard errors may be valid without need for any correction.⁸ Another key result from this literature is that, under appropriate assumptions, standard errors for OLS-generated regressors (i.e. for either OLS fitted values or OLS residuals) do not need to be adjusted when the researcher is only concerned with testing the null hypothesis of a zero coefficient.⁹ These results are frequently appealed to in applied work to justify the use of OLS-standard errors for a two-step approach. Crucially, however, these results are valid under the assumption that the first-stage controls are exogenous (i.e., do not have explanatory power for the outcome variable)—which we argue is generally not applicable in common macroeconomic applications. We show that if in fact the first-stage controls have explanatory power for the outcome variable, then using unadjusted standard-error formulas in the two-step approach delivers inference that is unnecessarily conservative (i.e., the opposite of the typical concern raised with inference for generated regressors). More recently, Breitung and Brüggemann (2023) also consider alternate estimation strategies in applied work that employs generated residuals as shocks. Our paper is complementary in so far as we advocate for a ‘one-step’ approach, while they advocate for what we term a ‘hybrid’ approach – which follows the ‘two-step’ procedure and then adds on as additional controls in the second-stage regression the variables from the shock-estimation stage – which we show is exactly equivalent. However, while Breitung and Brüggemann (2023) are concerned primarily with establishing the asymptotic properties of

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

⁸See discussion of Model 4 in Pagan (1984). This is also discussed in Hansen (2022) pp.368-369.

⁹See discussion of Model 1 and 4 in Pagan (1984), as well as general discussion in Hansen (2022) (pp. 363-368).

their proposed ‘hybrid’ estimator – and in particular proving that inference for their estimator is not affected by the generated regressor – we simply note it is equivalent to a one-step approach for a broad class of estimators (and so immediately avoiding the need for ‘special treatment’ vis-a-vis standard error calculation stemming from generated regressors). Unlike them, we are concerned instead with establishing the differences between the one-step and two-step estimator, which allows us to highlight formally the drawbacks of the two-step procedure common in applied work (our key contribution).¹⁰

Outline. The remainder of the paper is structured as follows. Section 2 presents the assumptions underpinning identification with controls in a time series setting. Section 3 formalises our key insight in general form. Section 4 discusses the implications of the OVB in different settings. Section 6 presents empirical applications. Section 7 concludes.

Notation. We use y_t to denote the outcome variable of interest at time t and r_t as the causal variable (e.g., a policy indicator). Throughout, our central focus is on the dynamic causal effect of r on y —i.e., the impulse response of y to a ‘shock’ to r . 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 r_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., $K_1 > 0$). We use $A^{\perp B}$ to denote the OLS population residual from a regression of A on B , and we omit constant terms in regressions throughout for convenience.

2 Identification of Dynamic Causal Effects with Control Variables

In this section, we highlight the conditions under which one-step regressions with appropriate controls can be interpreted as capturing dynamic causal effects. We do so through the lens of a structural moving-average model (as in [Stock and Watson, 2018](#)).

2.1 Structural Moving-Average Model

We start by expressing macroeconomic aggregates \mathbf{q}_t (an $N \times 1$ vector) as functions of all current and past primitive structural shocks. That DGP has the form:

$$\mathbf{q}_t = \Theta(L)\epsilon_t \quad (1)$$

where $\Theta(L)$ is the lag polynomial and $\epsilon_t = [\epsilon_{1,t}, \dots, \epsilon_{M,t}]'$ comprises $M \times 1$ independent and identically distributed (i.i.d.) structural shocks (and measurement errors). Without loss of

¹⁰Note in particular that none of the drawbacks of the two-step procedure that we detail in this paper – the over-estimation of standard errors, bias with additional control variables, and bias for quantile-regression estimators – are discussed formally in [Breitung and Brüggemann \(2023\)](#).

generality, we assume that r_t and y_t are the first and second variables in \mathbf{q}_t , respectively.

The parameter of interest is the impulse response of y_t to the first shock, $\epsilon_{1,t}$ (i.e., entries in the second row and first column in matrices comprising $\Theta(L)$). For each horizon h , this is:

$$\Theta_{h,21} \equiv \mathbb{E}[y_{t+h} | \epsilon_{1,t} = 1] - \mathbb{E}[y_{t+h} | \epsilon_{1,t} = 0] \quad (2)$$

We also normalise $\epsilon_{1,t}$ to imply a unit on-impact increase in r_t for interpretability, so $\Theta_{0,11} = 1$.

As noted in [Stock and Watson \(2018\)](#), the structural impulse response (2) in this setting is simply the population coefficient in the OLS regression:

$$y_{t+h} = \Theta_{h,21}\epsilon_{1,t} + u_{t+h} \quad (3)$$

But, since $\epsilon_{1,t}$ is not directly observed, this regression is not feasible. An alternate feasible approach discussed in [Stock and Watson \(2018\)](#) is to identify a plausible shock measure—with, e.g., narrative or high-frequency techniques—which may potentially be measured with error. Since these typically do not map one-to-one to the variable that the unit-shock normalisation is applied to, IV methods are needed to identify impulse response (2).

A commonly applied ‘two-step’ approach involves first *estimating* (rather than directly constructing) a plausible shock measure by regressing r_t on observable covariates \mathbf{x}_t and then running regression (3) with $\epsilon_{1,t}$ replaced with the estimated residual, denoted by $r_t^{\perp \mathbf{x}_t}$:

$$y_{t+h} = \beta_{2S}^h(r_t^{\perp \mathbf{x}_t}) + u_{t+h}^{2S} \quad (4)$$

where β_{2S}^h captures the impulse response of interest. By FWL, this estimated impulse response is equivalent to β_{1S}^h in the following one-step regression:

$$y_{t+h} = \beta_{1S}^h r_t + \delta \mathbf{x}_t + u_{t+h}^{1S} \quad (5)$$

2.2 Identification in Structural Moving-Average Model

Under what conditions does the one-step regression (5) identify the object of interest? Note that, since β_{1S}^h and β_{2S}^h are equivalent, one can immediately dismiss the claim that β_{1S}^h cannot be interpreted as capturing an impulse response function to a structural shock since it loads directly on a policy variable, not a shock measure. In fact, we now demonstrate that the one-step regression (5) can uncover the impulse response (2), even when the researcher is unable to directly recover exactly the shock of interest. This includes cases where $r_t^{\perp \mathbf{x}_t}$ is not exactly equivalent to the structural shock of interest $\epsilon_{1,t}$ —e.g., due to some measurement error that is uncorrelated with other shocks. Specifically, the following conditions are sufficient for β_{1S}^h in (5) to capture the impulse response of interest:

Condition 1. (Identification with Controls in Structural MA Model) (a) $\mathbb{E}[\epsilon_{1,t}^{\perp x_t} r_t^{\perp x_t}] \neq 0$ (*relevance*); (b) $\mathbb{E}[\epsilon_{i,t}^{\perp x_t} r_t^{\perp x_t}] = 0, \forall i \neq 1$ (*contemporaneous exogeneity*); (c) $\mathbb{E}[\epsilon_{t-j}^{\perp x_t} r_t^{\perp x_t}] = 0, j > 0$ (*lag exogeneity*); (d) $\mathbb{E}[\epsilon_{t+j}^{\perp x_t} r_t^{\perp x_t}] = 0, j > 0$ (*lead exogeneity*).

We prove this condition, and its sufficiency, in Appendix A.1. It follows from noting that the OLS regression (5) is equivalent to an IV regression where r_t instruments for itself, and so identifying assumptions reduce to the LP-IV conditions with controls in Stock and Watson (2018) treating r_t as the instrument. Intuitively, exogenous variation in r_t is generated purely by control variables x_t , rather than in combination with an IV as in standard LP-IV.

The relevance condition, (a), is somewhat trivial here, with the unit-shock normalisation reducing to the familiar OLS collinearity assumption that $r_t^{\perp x_t}$ has non-zero variance.¹¹ In addition, lead exogeneity, (d), will generally be satisfied since structural shocks are i.i.d., such that future shocks should be uncorrelated with time- t variables. However, the contemporaneous and lag exogeneity assumptions, (b) and (c), are strong, clearly depending on the choice of controls x_t . Intuitively, the set of controls must ‘partial out’ all other shocks from r_t —in a monetary policy setting, e.g., by capturing variables in the policy ‘reaction function’.

Although Condition 1 follows from the observation that OLS can be viewed as a special case of IV, the implications of this do not appear to have been generally appreciated in the applied literature. Conventional wisdom seems to hold that LP-OLS regressions of the form (4) and (5) are only valid when the researcher is able to exactly observe the shock of interest (i.e., when $r_t^{\perp x_t} = \epsilon_{1,t}$), and that only LP-IV is robust to cases where the researcher only observes a ‘proxy’ for the shock. However, Condition 1 highlights that LP-OLS is also robust to cases where the researcher only observes variation that is correlated with the structural shock of interest and cannot observe directly the shock itself due to measurement error.¹² Additionally, it is also widely held that identification via recursive-VARs requires (partial) invertibility to estimate the impulse response of interest (see, e.g., Stock and Watson, 2016). But noting that recursive-VARs and LP-OLS identify exactly the same impulse response (Plagborg-Møller and Wolf, 2021), it follows that identification via recursive-VARs only requires the exogeneity conditions detailed above, without necessarily appealing to invertibility.

The discussion thus far has highlighted the conditions under which one-step regressions with control variables, like (5), can be interpreted as capturing dynamic causal effects—or impulse responses to structural shocks. We now turn to consider the difference between one-step and two-step regressions in a general setting.

¹¹The unit-shock normalisation, (1) and the exogeneity conditions (b) and (c) imply $\mathbb{E}[\epsilon_{1,t}^{\perp x_t} r_t^{\perp x_t}] = \mathbb{E}[(r_t^{\perp x_t})^2]$.

¹²In particular, the above conditions are satisfied when $r_t^{\perp x_t} = \epsilon_{1,t} + \zeta_t$ where ζ_t reflects measurement error uncorrelated with other shocks. Note that measurement error resulting in $r_t^{\perp x_t}$ not scaling one-to-one with $\epsilon_{1,t}$ is explicitly ruled out here by the unit shock normalisation.

3 Omitted-Variable Bias in the Two-Step Approach

In this section, we highlight the drawbacks of the two-step approach for inference and identification. We show that the two-step approach can in general fail to identify dynamic causal effects, even when Condition 1 is satisfied. We do so within a general setup, without assumptions on the underlying DGP, presenting our key insight: that differences between two- and one-step approaches can be expressed in terms of an OVB formula stemming from the exclusion of potentially relevant variables from the outcome regression in the two-step approach.

3.1 General Setup

We first define the two-step approach in its general form. Here, a ‘shock’ to r_t is defined as:

$$\varepsilon_t = r_t - \mathbf{x}'_{1,t} \boldsymbol{\delta} \quad (6)$$

For now, we only impose the restriction that $\boldsymbol{\delta}$ is a vector of real numbers, $\boldsymbol{\delta} \in \mathbb{R}^{K_1}$. However, we predominantly focus on cases where ε_t is defined as a population OLS residual from a regression of r_t on $\mathbf{x}_{1,t}$, where the following holds by construction: $\mathbb{E} [\varepsilon_t \mathbf{x}'_{1,t}] = 0$, where ε_t can be thought of as ‘purging’ r_t of confounding effects from $\mathbf{x}_{1,t}$.

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

$$y_t = \mathbf{x}'_{2,t} \boldsymbol{\alpha} + \varepsilon_t \beta_{2S} + u_t^{2S} \quad (7)$$

where $\boldsymbol{\alpha}$ is a $K_2 \times 1$ vector of population regression coefficients pertaining to the controls $\mathbf{x}_{2,t}$. The coefficients in equation (7) satisfy the following (population) minimisation problem:

$$\{\boldsymbol{\alpha}, \beta_{2S}\} = \arg \min_{\mathbf{a}, b} \{\mathbb{E} (f(y_t - \mathbf{x}'_{2,t} \mathbf{a} - \varepsilon_t b))\} \quad (8)$$

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

It is important to note the generality of our setup. In addition to subsuming multiple estimators (i.e., different functional forms for $f(\cdot)$), it applies to a range of macroeconomic applications. For instance, when $\mathbf{x}_{1,t}$ includes lags of r_t plus lagged (and some contemporaneous) values of other variables, then (6) can be viewed as an equation from a recursively-identified SVAR. Alternatively, $\mathbf{x}_{1,t}$ may contain variables from a ‘reaction function’ as in, e.g., [Romer and Romer \(2004\)](#). Outcomes can also be defined as h -period ahead values of y , such that (7) amounts to a single h -specific regression from a LP with auxiliary controls $\mathbf{x}_{2,t}$ that were not used when constructing the shock. When $\mathbf{x}_{2,t}$ includes p lagged values of y_t alongside p lags

of other macroeconomic variables, then regression (7) can be thought of as a single equation from a VAR with β_{2S} capturing the contemporaneous response of y_t to ε_t .¹³

In line with the discussion in Section 2, the one-step approach is defined as the (population) regression of y_t on r_t and the full set of controls \mathbf{x}_t :

$$y_t = \underbrace{\mathbf{x}'_{1,t} \boldsymbol{\theta}_1 + \mathbf{x}'_{2,t} \boldsymbol{\theta}_2}_{\equiv \mathbf{x}'_t \boldsymbol{\theta}} + r_t \beta_{1S} + u_t^{1S} \quad (9)$$

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:

$$\{\boldsymbol{\theta}, \beta_{1S}\} = \arg \min_{\boldsymbol{\vartheta}, b} \{\mathbb{E}(f(y_t - \mathbf{x}'_t \boldsymbol{\vartheta} - r_t b))\} \quad (10)$$

where the function $f(\cdot)$ matches that used in the second stage of the shock-first regression (8).

3.2 Omitted-Variable Bias Result

We now seek to understand the difference between β_{1S} and β_{2S} . The following Proposition clarifies that this can always be expressed in terms of an OVB formula impacting the latter.

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 :*

$$y_t = \underbrace{\mathbf{x}'_{1,t} \boldsymbol{\phi}_1 + \mathbf{x}'_{2,t} \boldsymbol{\phi}_2}_{\equiv \mathbf{x}'_t \boldsymbol{\phi}} + \varepsilon_t \beta_{Hyb} + u_t^{Hyb} \quad (11)$$

where ε_t is defined as in equation (6) 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:

$$\{\boldsymbol{\phi}, \beta_{Hyb}\} = \arg \min_{\boldsymbol{\varphi}, b} \{\mathbb{E}(f(y_t - \mathbf{x}'_t \boldsymbol{\varphi} - \varepsilon_t b))\} \quad (12)$$

where the population objective function $f(\cdot)$ matches that used in two- and one-step regressions (8) and (10). Then the following holds:

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

where Ω is defined as omitted-variable bias term that arises from the exclusion of $\mathbf{x}_{1,t}$ in (7).

¹³We explicitly extend our setting to impulse-response estimation via SVARs in Appendices C and D.

Proof: Substituting ε_t from (6) into the hybrid estimator (12) and rearranging, we have:

$$\{\phi, \beta_{Hyb}\} = \arg \min_{\varphi, b} \left\{ \mathbb{E} (f(y_t - \mathbf{x}'_{1,t}(\varphi_1 - b\delta) - \mathbf{x}'_{2,t}\varphi_2 - r_t b)) \right\}$$

When $f(\cdot)$ has a unique minimand, we know from the one-step minimisation (10) that the solution to the above minimisation is given by: $\phi_2 = \theta_2$, $\phi_1 = \theta_1 + \beta_{Hyb}\delta$, and $\beta_{Hyb} = \beta_{1S}$.

Since (11) is the same as (7), albeit with additional covariates, the difference in coefficients can be expressed in terms of an omitted-variable bias term (i.e., the difference between a ‘long’ and ‘short’ regression that excludes some covariates): $\beta_{2S} = \beta_{Hyb} + \Omega \implies \beta_{2S} = \beta_{1S} + \Omega$. \square

Since these results rely only on the properties of optimisation, they carry over to in-sample estimated coefficients. That is, $\mathbb{E}(f(\cdot))$ can be redefined as the in-sample objective function (e.g., for some sample of length T , as: $\frac{1}{T} \sum_{t=1}^T f(x)$).

3.3 Discussion

The result establishes a mechanical link between the one- and two-step coefficients, in a setting that remains agnostic about the ‘true’ model. Whether the OVB term Ω reflects genuine bias is context-dependent.¹⁴ But the result is useful for two reasons.

First, it complements discussion in Section 2 by reiterating that differences between one- and two-step approaches do not hinge on coefficients being interpretable as capturing the effect of ‘shocks’. Rather, the two-step approach depends on the restriction that first-stage controls $\mathbf{x}_{1,t}$ do not feature in the regression with the outcome variable. It is hard to understand an economic rationale for imposing this restriction in settings where the two-step approach is typically applied—as is made clear in our primary empirical application, where the two-step approach amounts to assuming that central bank forecasts for inflation are exactly orthogonal to actual future inflation outcomes. Indeed, the motivation for the first-stage regression is typically to isolate exogenous variation in r_t by partialling out $\mathbf{x}_{1,t}$. But 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 r_t .

Second, this result provides an analytical formula for differences between one- and two-step coefficients across a range of settings, highlighting exactly when OVB may be genuinely problematic. This is true even in cases where no analytical expression exists for individual coefficients (as in QR).

¹⁴Note that, under the conditions provided in Section 2, the one-step regression can be interpreted as capturing the causal effect of r_t on y_t —and so under these same conditions the Ω term here reflects genuine bias.

4 Omitted-Variable Bias in Specific Settings

In this section, we set out the practical implications of Proposition 1 for numerous estimators where the two-step approach has been widely applied, specifically: OLS, IV, and QR.

4.1 Ordinary Least Squares (OLS)

Let equations (6), (7) and (9) all be OLS population regressions. Using Proposition 1 and the OVB formula for OLS, the following Result expresses the difference between coefficients:

Result 1. (OLS) *Define coefficients across regressions (6), (7), (9) and (11) as OLS population regression coefficients. This formula relates β_{2S} from equation (7) and β_{1S} from equation (9):*

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

where $\mathbf{B} \equiv [\mathbb{E} [\varepsilon_t^2] \mathbb{E} [\mathbf{x}_{2,t} \mathbf{x}'_{2,t}] - \mathbb{E} [\varepsilon_t \mathbf{x}_{2,t}] \mathbb{E} [\varepsilon_t \mathbf{x}'_{2,t}]]^{-1} \mathbb{E} [\mathbf{x}_{2,t} \mathbf{x}'_{1,t}]$ and $\boldsymbol{\phi}_1$ are coefficient loadings on $\mathbf{x}_{1,t}$ in equation (11).

Proof: First equality is from Proposition 1. Second uses the OLS OVB formula and standard matrix-partition algebra. Third uses $\mathbb{E} [\varepsilon_t \mathbf{x}'_{1,t}] = \mathbf{0}$, as ε_t is an OLS population residual. \square

To discuss the implications of this result, 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 (7). This is a common approach in applied work since, if the first stage (6) 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. (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.¹⁵ \square

In this case, point estimates from the one- and two-step approaches are equivalent. But, although sample estimates $\hat{\beta}_{2S}$ and $\hat{\beta}_{1S}$ are mathematically equivalent, the following Corollary shows that the two-step approach delivers wider standard-error estimates:

¹⁵ Angrist and Pischke (2009) refer to this formulation of Frisch-Waugh-Lovell (FWL) Theorem as the regression-anatomy formula (p. 27). 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 r_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 and the partial R-squared of $\mathbf{x}_{1,t}$ in regression (12) is non-zero, 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 $\hat{\beta}_{1S}$:*

$$\widehat{\text{Var}}(\hat{\beta}_{2S}) \geq \widehat{\text{Var}}(\hat{\beta}_{1S})$$

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:

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

Note also that the sample residual from the hybrid regression (11), \hat{u}_t^{Hyb} , and the sample residual from the one-step regression (9), \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: $\text{Var}(\hat{u}_t^{2S}) \geq \text{Var}(\hat{u}_t^{Hyb})$, where this holds with strict inequality when the partial R-squared on $\mathbf{x}_{1,t}$ is non-zero in the hybrid regression (11). And so when T is large relative to K_1 , then: $\widehat{\text{Var}}(\hat{\beta}_{2S}) \geq \widehat{\text{Var}}(\hat{\beta}_{1S})$. \square

The issue with inference arises since excluding $\mathbf{x}_{1,t}$ from the second stage 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 $\mathbf{x}_{1,t}$ has no explanatory power for y_t then the standard-error estimates will be equivalent between the one-step and two-step approaches.¹⁶ 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 r_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 with an application in Section 6, when $\mathbf{x}_{1,t}$ includes forecasts at time t of the outcome variable of interest.

Although Corollary 2 is written for a specific (homoskedastic) case, the 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 and [Newey and West \(1987\)](#) autocorrelation-robust standard errors.

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

¹⁶This relates to the discussion in [Pagan \(1984\)](#) pp. 222-233, showing that in models where only a generated residual appears on the right-hand side, standard-error formulas are valid for a two-step approach.

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 C, this case generalises to 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. In general, there are two key reasons researchers seek to include additional controls in this second stage. First, as we 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 $\mathbb{E} [\varepsilon_t \mathbf{x}'_{2,t}] = \mathbf{0}$, then $\beta_{2S} = \beta_{1S}$.*

Proof: This follows directly from Result 1. □

In practice, as long as the first-stage shock is *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 practical drawbacks of the two-step approach in this case. As before, we have that $\text{Var}(\hat{u}_t^{2S}) \geq \text{Var}(\hat{u}_t^{1S})$ 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.¹⁷

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 r_t of any further confounding factors—a common approach in the macroeconomics literature, both in estimation via LPs and VARs.¹⁸ In this case, by standard Frisch-Waugh-Lovell arguments, the two-step approach can alternatively be understood as estimating coefficients via a *three*-step procedure: (i) regress z_t on $\mathbf{x}_{1,t}$ and recover the residual ε_t ; (ii) regress ε_t on $\mathbf{x}_{2,t}$ and recover the residual $\varepsilon_t^{\perp \mathbf{x}_{2,t}}$; (iii) regress y_t on $\varepsilon_t^{\perp \mathbf{x}_{2,t}}$. This procedure has the unfortunate property that the effective shock measure the researcher is projecting on in the final stage – i.e. $\varepsilon_t^{\perp \mathbf{x}_{2,t}}$ – may be correlated with $\mathbf{x}_{1,t}$, exactly the variables the researcher hoped to partial out.

¹⁷Pagan (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.

¹⁸Recent 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’.

Result 1 then provides a formula for the exact bias that arises by re-introducing correlation with $\mathbf{x}_{1,t}$ in this way.

From an economic perspective, the issue can be understood as a failure to correctly identify a first-stage shock. Viewing r_t as a policy variable, suppose the researcher uses the two-step approach (which includes only $\mathbf{x}_{1,t}$ in the first 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, 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 since it does not rely on the correct partitioning of \mathbf{x}_t into $\mathbf{x}_{1,t}$ and $\mathbf{x}_{2,t}$.

4.2 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, 2023](#); [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.

To discuss the implications for IV, we develop our general setting. We introduce an additional variable m_t , which the researcher wishes to instrument for ε_t and may differ from r_t .¹⁹ The two-step approach is now defined as the following (population) IV regression:

$$y_t = m_t \beta_{2S}^{IV} + \mathbf{x}'_{2,t} \boldsymbol{\alpha} + u_t^{2S} \quad (13)$$

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

$$y_t = m_t \beta_{1S}^{IV} + \mathbf{x}'_t \boldsymbol{\theta} + u_t^{1S} \quad (14)$$

with r_t as an instrument for m_t . Because just-identified IV coefficients can be written as the ratio of OLS coefficients from a ‘first-stage’ and ‘reduced-form’ regression (see, e.g., [Angrist](#)

¹⁹Note that when $m_t \equiv r_t$, 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).

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:

$$m_t = \varepsilon_t \beta^{FS} + \mathbf{x}_{2,t} \boldsymbol{\pi}^{FS} + e_t^{FS} \quad \text{and} \quad y_t = \varepsilon_t \beta^{RF} + \mathbf{x}_{2,t} \boldsymbol{\pi}^{RF} + e_t^{RF}$$

Given this, the results from Section 4.1 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 potentially 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 (6), 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: $\widehat{\text{Var}}(\beta_{2S}^{FS}) \geq \widehat{\text{Var}}(\beta_{1S}^{FS})$ and $\widehat{\text{Var}}(\beta_{2S}^{IV}) \geq \widehat{\text{Var}}(\beta_{1S}^{IV})$.*

Proof: Coefficient equivalence follows by applying Frisch-Waugh-Lovell Theorem to the numerator and denominator of the IV estimator, the ratio of OLS coefficients. The first result on estimated standard errors follows from Corollary 2 as it refers to standard errors on OLS coefficients from a first-stage regression. The proof for IV standard errors is in Appendix B. \square

These implications of these results are 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.

It is also straightforward to show that other results from OLS carry over to IV. In particular, if auxiliary controls $\mathbf{x}_{2,t}$ were included in the IV regression (13), 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 r_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 demonstrate in Appendix D and in our empirical application in Section 6).

4.3 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.

Mirroring our discussion of identification in OLS settings in Section 2, in Appendix A.2 we additionally discuss identification in quantile regression settings. We do so through the lens of the potential outcomes framework, where we highlight that one-step quantile regressions of the form (9) have a causal interpretation under an appropriate ‘selection-on-observables’ assumption: that the policy variable of interest r_t is independent of the outcome variable y_{t+h} conditional on the controls \mathbf{x}_t . We now show that the two-step and one-step approach yield different estimates, implying the two-step approach generally fails to have a causal interpretation under this same assumption:

Result 2. (QR) Define the coefficients in equation (6) as OLS population regression coefficients, and the coefficients in equations (7), (9) and (11) 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 (8), (10) and (12): $f(x) = \rho_\tau(x)$, where $\rho_\tau(x) = x(\tau - \mathbb{1}(x \leq 0))$ is the check function.

Assuming that the conditional quantile function of y_t given r_t, \mathbf{x}_t is linear,²⁰ then the following formula relates the two-step population coefficient at the τ -th quantile $\beta_{2S}(\tau)$ in equation (7) and the one-step population coefficient $\beta_{1S}(\tau)$ in equation (9):

$$\begin{aligned}\beta_{2S}(\tau) &= \beta_{1S}(\tau) + \Omega_{QR}(\tau) \\ &= \beta_{1S}(\tau) + [\mathbf{A} \mathbb{E}[w_\tau \varepsilon_t \mathbf{x}'_{1,t}] + \mathbb{E}[w_\tau \varepsilon_t \mathbf{x}'_{2,t}] \mathbf{B}] \phi_1(\tau)\end{aligned}$$

where: $\mathbf{A} \equiv [\mathbb{E}[w_\tau \varepsilon_t^2]^{-1} + \mathbb{E}[w_\tau \varepsilon_t^2]^{-2} \mathbb{E}[w_\tau \varepsilon_t \mathbf{x}_{2,t}] [\mathbb{E}[w_\tau \mathbf{x}_{2,t} \mathbf{x}'_{2,t}]]^{-1} \mathbb{E}[w_\tau \varepsilon_t \mathbf{x}_{2,t}] - \mathbb{E}[w_\tau \varepsilon_t^2]^{-1} \mathbb{E}[w_\tau \varepsilon_t \mathbf{x}_{2,t}] \mathbb{E}[w_\tau \varepsilon_t \mathbf{x}'_{2,t}]]^{-1} \mathbb{E}[w_\tau \varepsilon_t \mathbf{x}_{2,t}]$

$$\mathbf{B} \equiv [\mathbb{E}[w_\tau \mathbf{x}_{2,t} \mathbf{x}'_{2,t}] - \mathbb{E}[w_\tau \varepsilon_t^2]^{-1} \mathbb{E}[w_\tau \varepsilon_t \mathbf{x}_{2,t}] \mathbb{E}[w_\tau \varepsilon_t \mathbf{x}'_{2,t}]]^{-1} \mathbb{E}[w_\tau \mathbf{x}_{2,t} \mathbf{x}'_{1,t}] \mathbb{E}[w_\tau \varepsilon_t^2]$$

and $\phi_1(\tau)$ are coefficient loadings on $\mathbf{x}_{1,t}$ in equation (11). In addition the ‘importance weights’ are defined as:

²⁰Linearity here merely helps to simplify the expression for OVB included below. A more general expression under non-linearity can be found in [Angrist et al. \(2006\)](#).

$$w_\tau = \int_0^1 f_{u_\tau^{Hyb}}(u [\mathbf{x}'_{2,t} \boldsymbol{\pi}(\tau) + \varepsilon_t \beta(\tau) - \mathbf{x}'_t \boldsymbol{\phi}(\tau) - \varepsilon_t \beta_{hyb}(\tau)]) | \mathbf{x}_t, \varepsilon_t) du / 2.$$

Proof: Line 1 follows from Proposition 1. Line 2 follows from the OVB formula for QR (Angrist, Chernozhukov, and Fernández-Val, 2006) and standard matrix-partition algebra. \square

To discuss the practical implications of this result, we again first consider a simple setting without auxiliary controls in the second-stage regression. In Case 1, unlike in OLS, there can still be OVB in the two-step estimator. The below Corollary formalises this:

Corollary 5. (OVB in QR without Controls) *Under Case 1, the following holds:*

$$\begin{aligned} \beta_{2S}(\tau) &= \beta_{1S}(\tau) + \Omega_{QR}(\tau) \\ &= \beta_{1S}(\tau) + \boldsymbol{\phi}_1(\tau) \frac{\mathbb{E} [w_\tau \varepsilon_t \mathbf{x}'_{1,t}]}{\mathbb{E} [w_\tau \varepsilon_t^2]} \end{aligned}$$

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

Proof: This follows from Result 2. \square

Unlike with OLS, even though $\mathbb{E} [\varepsilon_t \mathbf{x}'_{1,t}] = \mathbf{0}$ by construction, this does not imply $\Omega_{QR}(\tau) = 0$. We can still have $\mathbb{E} [w_\tau \varepsilon_t \mathbf{x}'_{1,t}] \neq \mathbf{0}$. In order for OVB to be zero, we need additional assumptions, which are unlikely to hold in practical applications. For example, if the importance weights w_τ are constant across $[\varepsilon_t, \mathbf{x}'_{1,t}]$ then it is straightforward to show that $\Omega_{QR}(\tau) = 0$. As Angrist et al. (2006) explain, the importance weights will be constant across right-hand side variables 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 (7) 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 settings where auxiliary controls are included in the second-stage regression. In this Case 2, unlike in OLS, OVB can be non-zero in QR when the shock is constructed to be uncorrelated with other drivers of y_t , i.e.: $\mathbb{E} [\varepsilon_t \mathbf{x}'_{2,t}] = \mathbf{0} \nRightarrow \beta_{2S}(\tau) = \beta_{1S}(\tau)$.

5 Simulation Study

We now summarise the results from a simulation study to assess the quantitative significance of the differences between the one- and two-step approaches in an empirically realistic setting. For brevity, we relegate full details of the study to Appendix E and provide only a brief description here. We employ the nine-variable recursive VAR(4) of Christiano et al. (2005) as

a medium-scale, empirically realistic data-generating process (DGP).²¹ We focus on estimating the effects of a monetary policy shock on the nine outcome variables in the model: real GDP, prices, consumption, investment, wages, labour productivity, the federal funds rate, real profits and money growth. Data is quarterly and we assume the econometrician has access to a sample of 121 quarters, as in [Christiano et al. \(2005\)](#). We use our DGP to compute biases and standard errors for different estimation procedures across 2,000 Monte Carlo simulations. We consider two DGPs - a 'stationary' DGP where we first stationarise all data before including it in the VAR, and a 'non-stationary' DGP where some variables enter the VAR directly in log-levels (as in [Christiano et al. \(2005\)](#)).

We consider estimation of impulse responses via either a one- or two-step local projection using OLS. The one-step estimator regresses the outcome variable on the federal funds rate, controlling for four lags of all variables in the VAR plus contemporaneous values of variables ordered before the federal funds rate (i.e., all variables except real profits and money growth). As in our theoretical exposition, we then distinguish two cases for the two-step estimator. In Case 1, the shock is estimated 'correctly', by regressing the federal funds rate on four lags of all variables in the VAR plus contemporaneous values of variables ordered before the federal funds rate, and then this shock is employed in a local projection without any auxiliary controls. In Case 2, the shock is identified by regressing the federal funds rate on *only* the contemporaneous values of variables ordered before the federal funds rate (i.e., without lags). The IRF is then estimated by employing the shock in a local projection, with four lags of all the outcome variables as controls. We view this as a reasonable approximation to common applications in macroeconomics, where researchers identify 'shocks' by cleaning the causal variable of the key variables believed to drive reverse causality (typically contemporaneous variables or those at very short lags), and then estimate local projections with significant numbers of lags of the outcome variable plus 'standard' macroeconomic variables as auxiliary controls.²²

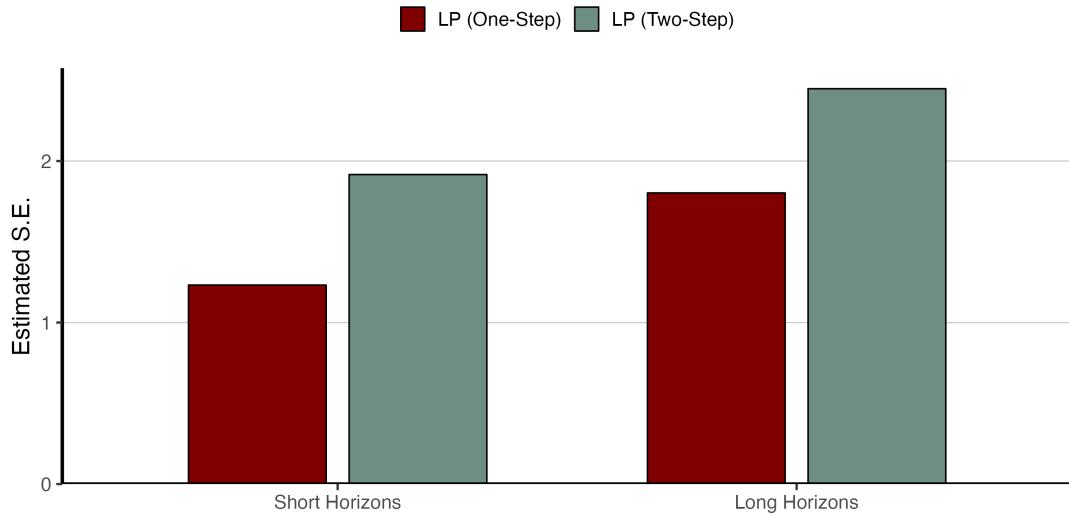
In Case 1, the one-step and two-step estimators are identical from the perspective of identification, but differ in terms of estimated standard errors (see Corollary 1 and 2). Figures 2a and 3a below quantifies the extent of these differences in standard errors. We compute the expected value of estimated standard errors for each estimator across simulations, and then take the median across outcome variables and horizons. For comparison (and to ensure appropriate scaling across outcomes), before taking the median, we divide the width of the expected standard error to those estimated from a correctly-specified VAR(4). We find the one-step LP estimator produces larger standard errors than a VAR - around 20% larger on average at short horizons (< 1 year) and around 50-80% larger at long horizons (1-4 years) across DGPs — similar to the discussion of standard error estimates for LPs vs VARs in [Montiel Olea et al.](#)

²¹We follow the choice of DGP from the simulation study in [Herbst and Johannsen \(2024\)](#).

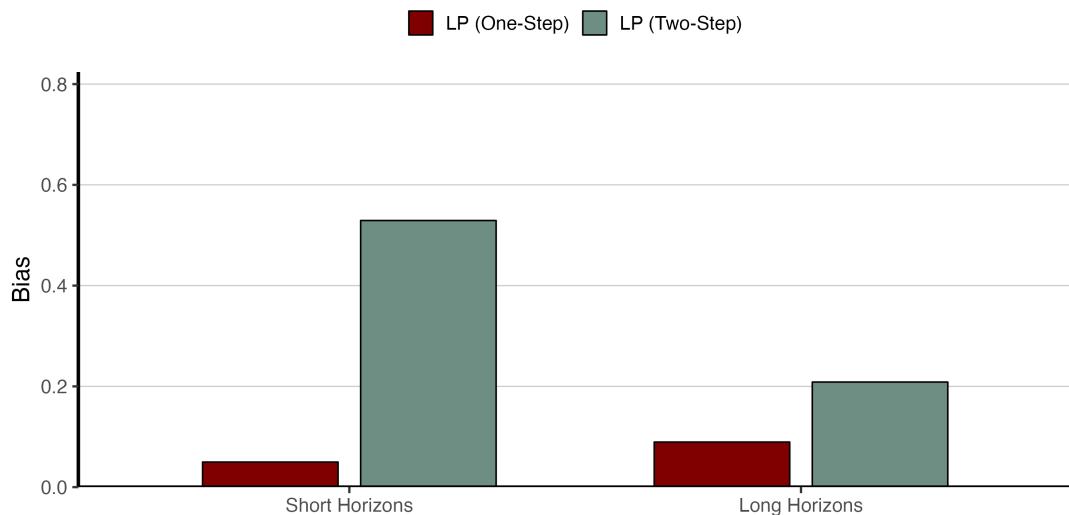
²²This is similar in spirit to the approach in monetary policy settings (see, e.g., [Romer and Romer, 2004](#); [Miranda-Agrippino and Ricco, 2021](#); [Bauer et al., 2023](#)), as well as fiscal policy settings (see e.g. [Corsetti et al., 2012](#); [Auerbach and Gorodnichenko, 2013](#)).

Figure 2: Results from local-projections simulation study (Stationary DGP)

(a) Estimated Standard Errors (no auxiliary controls)

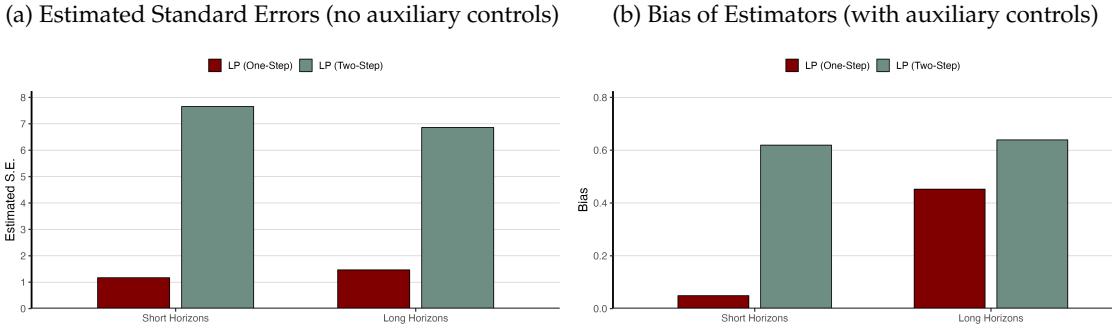


(b) Bias of Estimators (with auxiliary controls)



Notes: Results from simulation study using medium-scale VAR(4) with stationary data as DGP. Figure (a) displays the median (across outcomes and horizons) of estimated standard errors from a one-step and two-step LP estimator (without auxiliary controls) relative to a VAR. Figure (b) displays median (across outcomes and horizons) of absolute bias of one-step and two-step LP estimator (with auxiliary controls) relative to the root mean squared value of the true impulse response function out to the maximum horizon (4-years).

Figure 3: Results from local-projections simulation study (Non-Stationary DGP)



Notes: Results from simulation study using medium-scale VAR(4) with non-stationary data as DGP. Figure (a) displays the median (across outcomes and horizons) of estimated standard errors from a one-step and two-step LP estimator (without auxiliary controls) relative to a VAR. Figure (b) displays median (across outcomes and horizons) of absolute bias of one-step and two-step LP estimator (with auxiliary controls) relative to the root mean squared value of the true impulse response function out to the maximum horizon (4-years).

(2025). While meaningful, these differences generally appear small compared to the two-step procedure which produces standard errors that are about 100% and 150% larger than a VAR at short and long horizons for the stationary DGP, and around 600% larger than a VAR at all horizons for the non-stationary DGP.

In Case 2, with auxiliary controls, the one and two-step estimator do not identify the same IRF, where Result 1 provides the formula for the (population) bias in the two-step estimator. Figures 2b and 3b compare quantitatively the small-sample bias of the one- and two-step estimators in this case. As in Li et al. (2021), we display the absolute median bias of each estimator across all outcome variables and horizons where, before taking the median, we cancel out the units of the outcome variables by dividing by the root mean squared value of the true IRF over the impulse horizon. The results for the one-step LP estimator for our non-stationary DGP are similar to the broader simulation study in Li et al. (2021): as a fraction of the root mean squared value of the true IRF, median bias is under 0.1 at short horizons (<1 year), rising to around 0.5 at longer horizons (1-4 years). The two-step estimator performs much worse: at short horizons, the median bias is over *ten* times larger than for the one-step estimator, and around 40% larger at longer horizons. These differences are quantitatively large, e.g. exceeding any difference in bias between the range of LP and VAR estimators considered in Li et al. (2021). We find similar results for the stationary DGP – the median bias is over ten times larger than for the one-step estimator at short horizons, and over 100% larger at longer horizons. These results underscore the practical significance of the discussion from the previous section.

6 Empirical Applications

We now illustrate our theoretical results with a series of empirical applications, analysing the dynamic effects of US monetary shocks controlling for central-bank information. Focusing on the responses of US CPI, we illustrate our theoretical results for each of the three estimators discussed in Section 4 in turn: OLS, IV, and QR, covering estimation by both LPs and VARs. We explain how these applications have important implications for the vast empirical literature studying the causal effects of monetary policy.

For most applications, we use common data. To estimate the first-stage US monetary shocks $\hat{\varepsilon}_t^{mp}$, orthogonalised with respect to central-bank information, we use the [Romer and Romer \(2004\)](#) specification. The shock is constructed by regressing the change in the Federal Funds target rate Δi_t on the previous target rate i_{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:

$$\begin{aligned} \Delta i_t = & \delta_0 + \delta_1 i_{t-1} + \sum_{i=-1}^2 [\delta_{2,i} \Delta y_{t,i}^e + \delta_{3,i} (\Delta y_{t,i}^e - \Delta y_{t-1,i}^e) + \delta_{4,i} \pi_{t,i}^e + \delta_{5,i} (\pi_{t,i}^e - \pi_{t-1,i}^e)] \\ & + \delta_6 u_{t,0}^e + \hat{\varepsilon}_t^{mp} \end{aligned} \quad (15)$$

This is analogous to equation (6) from our theoretical results defining r_t as the change in the Federal Funds target rate and $\mathbf{x}_{1,t}$ as a vector of controls including the previous target rate alongside Greenbook forecasts and forecast revisions. We discuss the construction of these shocks in further detail in [Appendix F](#).

In our specifications, we also include additional auxiliary controls alongside $\hat{\varepsilon}_t^{mp}$ —i.e., $\mathbf{x}_{2,t}$ from our general setup. Throughout, we include lags of the following in the set of auxiliary controls: industrial production, consumer prices and the unemployment rate.

6.1 Ordinary Least Squares

We first estimate US monetary policy's effects on US CPI by employing our shocks in LPs, estimated by OLS, and as internal instruments in recursive SVARs.

Local Projections (LP-OLS). We estimate the effects of the shocks using a second-stage LP:²³

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h + \mathbf{x}'_{2,t} \boldsymbol{\alpha}^h + \hat{\varepsilon}_t^{mp} \beta_{2S}^h + u_{t+h}^{2S} \quad (16)$$

²³Unlike [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.

where $h = 0, 1, \dots, 48$.²⁴ This is an analog to equation (7). We compare this to the one-step LP:

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h + \mathbf{x}'_{1,t} \boldsymbol{\theta}_1^h + \mathbf{x}'_{2,t} \boldsymbol{\theta}_2^h + \Delta i_t \beta_{1S}^h + u_{t+h}^{1S} \quad (17)$$

an analog to equation (9).

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, where $\mathbf{x}_{2,t}$ is non-empty (i.e. $K_2 > 0$). In the latter, $\mathbf{x}_{2,t}$ contains one-month lags of month-on-month changes in (log) industrial production, (log) CPI and the unemployment rate.

Table 1 presents the results, which support our main findings from Section 4.1. Focusing on Case 1, where $\mathbf{x}_{2,t}$ is empty, columns (1)-(3) present β_i estimates for $i = \{2S, 1S\}$ and $h = 0, 12, 24, 36, 48$.²⁵ Here, point estimates from the one- and two-step estimators are identical, per Corollary 1. They, unsurprisingly, indicate that a US monetary-policy shock is associated with negative lagged effects on US CPI.

However, as Corollary 2 states, estimated standard errors from the one- and two-step estimators differ. In all cases, as columns (1) and (2) show, OLS standard errors, calculated assuming homoskedasticity, are smaller for the one-step estimates relative to the two-step. This carries over to other standard errors too, including [White \(1980\)](#) heteroskedasticity-robust standard errors and [Newey and West \(1987\)](#) standard 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. 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 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 significantly negative using the one-step (and hybrid) approach.

Internal Instruments in a Recursive SVAR. Next, we estimate the effects of [Romer and Romer \(2004\)](#) shocks when used as ‘internal instruments’ in a VAR ([Plagborg-Møller and Wolf, 2021](#)). Here, the two-step approach uses the shocks directly in a recursive SVAR, with the shock ordered first such that shocks to subsequent variables do not impact it contemporaneously.²⁶

²⁴Strictly, 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.

²⁵These impulse response functions are also presented in Figures 1a and 1b

²⁶Note 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.

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	
	Two-Step	One-Step	Two-Step	One-Step
$h = 0$	0.03 (0.05)	0.03 (0.04)	0.06 (0.04)	0.04 (0.03)
OLS s.e.				
N-W s.e.				
Rob. s.e.				
$h = 12$	0.13 (0.50)	0.13 (0.22)	0.46 (0.38)	0.13 (0.22)
OLS s.e.				
N-W s.e.				
Rob. s.e.				
$h = 24$	-0.15 (0.90)	-0.15 (0.41)	0.37 (0.74)	-0.11 (0.41)
OLS s.e.				
N-W s.e.				
Rob. s.e.				
$h = 36$	-1.20 (1.26)	-1.20** (0.53)	-0.52 (1.07)	-1.13** (0.53)
OLS s.e.				
N-W s.e.				
Rob. s.e.				
$h = 48$	-2.46 (1.57)	-2.46*** (0.62)	-1.66 (1.36)	-2.38*** (0.63)
OLS s.e.				
N-W s.e.				
Rob. s.e.				

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.

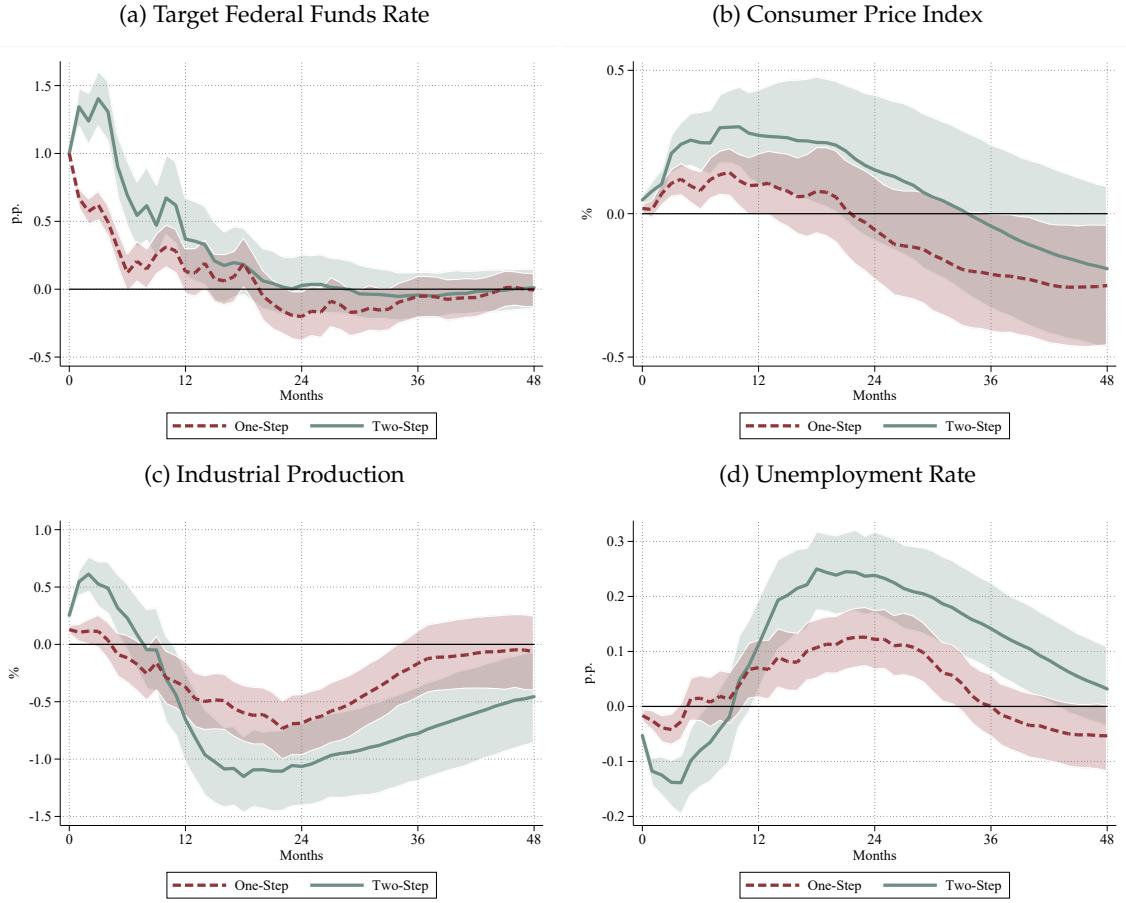
Our application uses the VAR specification of [Coibion \(2012\)](#). For the two-step approach, we estimate a VAR(12) with (in order): estimated Romer-Romer shock $\hat{\varepsilon}_t^{mp}$, Federal Funds target rate (i_t), and macroeconomic variables (specifically, consumer prices, industrial production, unemployment and commodity prices). We compare this to a one-step approach where the VAR(12) includes the Greenbook variables in $x_{1,t}$ as additional endogenous variables.²⁷ We use the new target rate (i_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.²⁸

In Appendix C, we analytically demonstrate how our OVB result carries over to this case.

²⁷To 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).

²⁸As 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 4: 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

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 4 presents the estimated impulse responses from the one- and two-step recursive VARs. As in the LP-OLS application, one-step estimates again suggest that US monetary policy has significant (dis)inflationary consequences. The two-step VAR points to a pronounced price 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.

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). They also differ from [Cochrane \(2004\)](#), who employs the [Romer and Romer \(2004\)](#) shocks in a LP and concludes that they provide only very weak evidence for the effects of monetary policy on the basis the response of prices are barely significant. Instead, our results highlight that appropriately controlling for the Greenbook forecasts in a one-step procedure need not result in puzzling findings and can deliver highly significant estimated coefficients at medium-term horizons.

6.2 Instrumental Variables

As we outlined in Section 4.2, 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 the implications of our IV results, 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 F, in particular Figure E2.

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 stage one, 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 $r_t := mp_t^{surp}$ 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 an IV regression of the form:

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h + m_t \beta_{2S}^{IV,h} + u_{t+h,IV}^{2S} \quad (18)$$

to estimate the dynamic effects of US monetary policy on US CPI. This is an analog to equation (13). We compare this to a one-step LP-IV regression:

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h + \mathbf{x}'_{1,t} \boldsymbol{\theta}_1^h + m_t \beta_{1S}^{IV,h} + u_{t+h,IV}^{1S} \quad (19)$$

in which the surprise mp_t^{surp} instruments for m_t to make this the analog to equation (14). Throughout, we use the 1-year Treasury yield as our policy indicator m_t and, due to the availability of high-frequency surprises, start our sample in 1990:01—but still end it in 2007:12.

We report the impulse responses of consumer prices from this specification in Figure E3 of Appendix F. 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 (see e.g., [Nakamura and Steinsson, 2018a](#)), this finding has particular importance. ‘Best practice’ for identification in that 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, 2023](#)). However, doing so in a two-step approach implies that tests of instrument strength and coefficient significance will be unnecessarily conservative.

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

	Two-Step	One-Step
<i>A: LP-IV</i>		
Case 1: No Auxiliary Controls	11.756	19.571
Case 2: With Auxiliary Controls	11.266	19.578
<i>B: External Instruments SVAR</i>		
	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-Agrifino and Ricco \(2021\)](#) using one- and two-step approaches. Panel A reports results from LP-IV applications, equations (18) and (19), 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.

External Instruments SVAR. We also demonstrate that our results in Section 4.2 have implications for estimation via Proxy SVARs. As we show formally in Appendix D, 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 instruments 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.

Appendix F, Figure E4 presents impulse responses for the 1-year yield and CPI from the one- and two-step Proxy SVARs. The two-step approach shows a large and significant rise in prices 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. [Stock and Watson \(2018\)](#) propose comparing Proxy-SVAR and LP-IV estimators as a direct test of VAR-invertibility, implying that the sharp contrast in impulse responses from our LP-IV and Proxy-SVAR applications strongly suggests that invertibility fails in this setting.

6.3 Quantile Regression

Finally, we study the dynamic response of conditional quantiles of US CPI to US monetary shocks. We estimate QR analogs to equations (16) and (17), again employing $\hat{\varepsilon}_t^{mp}$ as our shock measure. The two-step approach involves estimating the following second-stage LP-QR :

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h(\tau) + \hat{\varepsilon}_t^{mp} \beta_{2S}^h(\tau) + u_{t+h}^{2S}(\tau) \quad (20)$$

and the one step is:

$$\ln(CPI_{t+h}) - \ln(CPI_{t-1}) = \alpha_0^h(\tau) + \mathbf{x}'_{1,t} \boldsymbol{\theta}_1^h(\tau) + \Delta i_t \beta_{1S}^h(\tau) + u_{t+h}^{1S}(\tau) \quad (21)$$

As discussed in Section 4.3, OVB 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—especially the differences in point estimates from the two estimation approaches.

While the one-step estimates indicate that a US monetary 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 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—a point made forcibly by Greenspan (2004).²⁹ 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 one-step approach we advocate for here can be applied 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.

²⁹Greenspan (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

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

6.4 Further Applications

We now consider the broader relevance of our results by using data from other well-known studies. We revisit a total of five studies for which replication data were available, summarised in Table 3. The studies we consider have been published in top economic journals and have a high citation count (averaging almost 1600 citations). They cover applications estimating the effects of: monetary policy, oil and temperature shocks. In total, across applications, we generate over 500 comparisons between one-step and two-step coefficients and standard errors.

Table 4: Summary of Further Applications

Paper	Publication		Sample		
	Journal	Cit.	Period	Region	Freq.
Romer and Romer (2004)	AER	2399	1972-2007	USA	Monthly
Miranda-Agrippino and Ricco (2021)	AEJ:Macro	813	1990-2007	USA	Monthly
Barnichon and Mesters (2021)	JME	80	1972-2007	USA	Monthly
Kilian (2009)	AER	5305	1973-2007	USA	Monthly
Bilal and Kängig (2024)	NBER WP	120	1960-2019	Global	Annual

Notes: Citation count measured by Google Scholar as of Aug 2025.

We follow each paper in the construction of the shock, and then estimate local projections of the form (17) and (16) to examine the differences between the one-step and two-step approach. As before, 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, where $\mathbf{x}_{2,t}$ is non-empty (i.e. $K_2 > 0$). In the latter, $\mathbf{x}_{2,t}$

contains lags of standard macroeconomic variables (e.g. inflation, industrial production and unemployment), as in our main application. Across studies, we estimate the response of key macroeconomic outcomes (e.g. consumer prices, industrial production, unemployment, GDP). We provide full details of this exercise in Appendix G.

Table 5 presents our first set of results examining the difference in estimated Newey and West (1987) standard errors across studies for the case when no second-stage controls are included. The differences are generally large and meaningful: the two-step approach generates standard errors that are on-average about 72% (mean), or 47% (median), larger than the one-step approach. But we find the difference between the one-step and two-step standard error estimates varies substantially across studies, horizons and outcome variables. As in our main application, the differences can be much larger: in the worst cases, around 200 – 400% larger.

Table 5: Relative Standard Errors (No Auxiliary Controls)

Study	Horizon	Outcome Variable	Mean	Median
All	All	All	1.727	1.473
Barnichon and Mesters (2021)	Short	Philips Multiplier	4.220	4.365
Romer and Romer (2004)	Short	Consumer Prices	3.656	3.810
Romer and Romer (2004)	Long	Consumer Prices	3.378	2.501
Miranda-Agrippino and Ricco (2021)	Long	Unemployment	2.089	2.090
Barnichon and Mesters (2021)	Long	Philips Multiplier	1.747	1.539

Notes: Table presents mean- and median- average relative Newey and West (1987) standard errors across one-step and two-step estimators. The first row computes averages across all studies, variables and horizons. The remaining rows display the five worst cases with largest (mean) relative standard errors across studies, variables and short and long horizons, where horizons less than 1-year are defined as 'short', and longer than 1-year are defined as 'long'.

Table 6 presents results examining the difference in estimated coefficients across studies for the case when second-stage controls are included. We compute the (absolute) difference between coefficient estimates from the one-step and two-step estimator, and then scale this difference by the estimated standard error of the one-step estimator. On average, one- and two-step coefficients differ by around 0.2 – 0.4 standard errors. In classic studies on the effects of monetary policy, these differences can be much larger: around 0.7 – 1.5 standard errors. While Table 6 presents results on average bias across horizons, we find individual coefficient estimates for specific horizons can differ by over 3 standard deviations in some cases.

7 Conclusions

A common approach to estimating dynamic causal effects in macroeconomics involves estimating the 'shocks first': orthogonalising causal variables of interest with respect to confounders; then, using the orthogonalised variables in a second-stage LP or VAR. As we have explained, this approach subsumes multiple identification approaches and has been applied

Table 6: Difference in Coefficient Estimates (With Auxiliary Controls)

Study	Horizon	Outcome Variable	Mean	Median
All	All	All	0.444	0.255
Miranda-Agrippino and Ricco (2021)	Long	Unemployment	1.529	1.858
Miranda-Agrippino and Ricco (2021)	Long	Consumer Prices	0.836	0.813
Romer and Romer (2004)	Long	Unemployment	0.713	0.725
Miranda-Agrippino and Ricco (2021)	Long	Industrial Production	0.403	0.137
Barnichon and Mesters (2021)	Long	Philips Multiplier	0.294	0.280

Notes: Table presents mean- and median- average differences in coefficient estimates (scaled by the estimated one-step standard error) across one-step and two-step estimators. The first row computes averages across all studies, variables and horizons. The remaining rows display the five worst cases with largest (mean) difference in coefficient estimates across variables and short and long horizons, where horizons less than 1-year are defined as 'short', and longer than 1-year are defined as 'long'.

in a wide range of settings. An alternate one-step approach involves simply including confounders as control variables in a regression for the outcome variable.

We have shown, for a wide set of estimators, that the two-step 'shock-first' approach can be problematic for both identification and inference relative to the one-step method. 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 OVB 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. One-step LPs and VARs remove a significant portion of the near-term price puzzle identified in previous studies analysing the response of prices to monetary shocks. Our one-step results do not rely on 'recursiveness assumptions' to resolve the price puzzle, and yield estimated impulse responses that are highly significant at medium-term horizons. Moreover, we provide new evidence that monetary policy acts as a "location shifter" of the entire inflation distribution, a result that was missed when implementing a two-step procedure. Together, our 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.

BARATTIERI, A. AND M. CACCIATORE (2023): "Self-Harming Trade Policy? Protectionism and Production Networks," *American Economic Journal: Macroeconomics*, 15, 97–128.

BARATTIERI, A., M. CACCIATORE, AND N. TRAUM (2023): "Estimating the Effects of Government Spending Through the Production Network," NBER Working Papers 31680, National Bureau of Economic Research, Inc.

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.

——— (2021): "The Phillips multiplier," *Journal of Monetary Economics*, 117, 689–705.

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

BAUER, M. D., B. S. BERNANKE, AND E. MILSTEIN (2023): "Risk Appetite and the Risk-Taking Channel of Monetary Policy," *Journal of Economic Perspectives*, 37, 77–100.

BAUER, M. D. AND E. T. SWANSON (2023): "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.

BILAL, A. AND D. R. KÄNZIG (2024): "The Macroeconomic Impact of Climate Change: Global vs. Local Temperature," Working Paper 32450, National Bureau of Economic Research.

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.

BREITUNG, J. AND R. BRÜGGEMANN (2023): "Projection Estimators for Structural Impulse Responses," *Oxford Bulletin of Economics and Statistics*, 85, 1320–1340.

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.

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

CHUDIK, A. AND G. GEORGIADIS (2022): "Estimation of Impulse Response Functions When Shocks Are Observed at a Higher Frequency Than Outcome Variables," *Journal of Business & Economic Statistics*, 40, 965–979.

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.

COCHRANE, J. H. (2004): "Comments on "A New Measure of Monetary Shocks: Derivation and implications"," *Comments at NBER EFG Meeting*, July 2004.

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, Á., 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.

FORONI, C. AND M. MARCELLINO (2016): "Mixed frequency structural vector auto-regressive models," *Journal of the Royal Statistical Society Series A*, 179, 403–425.

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*, Princeton University Press.

HERBST, E. P. AND B. K. JOHANSEN (2024): "Bias in local projections," *Journal of Econometrics*, 240, None.

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.

LI, D., M. PLAGBORG-MØLLER, AND C. K. WOLF (2021): "Local Projections vs. VARs: Lessons From Thousands of DGPs," Working Papers 2021-55, Princeton University. Economics Department.

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.

MONTIEL OLEA, J. L., M. PLAGBORG-MØLLER, E. QIAN, AND C. K. WOLF (2024): "Double Robustness of Local Projections and Some Unpleasant VARithmetic," NBER Working Papers 32495, National Bureau of Economic Research, Inc.

——— (2025): “Local Projections or VARs? A Primer for Macroeconomists,” in *NBER Macroeconomics Annual 2025, Volume 40*, National Bureau of Economic Research, Inc, NBER Chapters.

MONTIEL OLEA, J. L. AND M. PLAGBORG-MØLLER (2021): “Local Projection Inference Is Simpler and More Robust Than You Think,” *Econometrica*, 89, 1789–1823.

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.

NEWHEY, 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. A. (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. M. 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 (2016): “Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics,” Elsevier, vol. 2, chap. 8, 415–525.

——— (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.

Supplemental Appendix

A Identification With Controls

A.1 Structural Moving-Average Model

First, we prove that Condition 1 in the structural moving-average model from Section 2.1 is sufficient for estimation of impulse responses using a one-step regression.

The OLS estimand from the 'one-step' regression (5) is:

$$\beta_{1S}^h = \frac{\mathbb{E} [y_{t+h}^{\perp \mathbf{x}_t} r_t^{\perp \mathbf{x}_t}]}{\mathbb{E} [r_t^{\perp \mathbf{x}_t} r_t^{\perp \mathbf{x}_t}]}$$

The numerator can be written as:

$$\mathbb{E} [y_{t+h}^{\perp \mathbf{x}_t} r_t^{\perp \mathbf{x}_t}] = \mathbb{E} [(\Theta_{h,21} \epsilon_{1,t}^{\perp \mathbf{x}_t} + u_{t+h}^{\perp \mathbf{x}_t}) r_t^{\perp \mathbf{x}_t}] = \Theta_{h,21} \mathbb{E} [\epsilon_{1,t}^{\perp \mathbf{x}_t} r_t^{\perp \mathbf{x}_t}]$$

where the first line substitutes in expression (3) and the second line uses exogeneity conditions (and uses that u_{t+h} contains all other past, contemporaneous and future shocks that affect y_{t+h}). The denominator can then be written as:

$$\begin{aligned} \mathbb{E} [r_t^{\perp \mathbf{x}_t} r_t^{\perp \mathbf{x}_t}] &= \mathbb{E} [(\Theta_{01} \epsilon_t^{\perp \mathbf{x}_t} + \Theta_{11} \epsilon_{t-1}^{\perp \mathbf{x}_t} + \dots) r_t^{\perp \mathbf{x}_t}] \\ &= \Theta_{0,11} \mathbb{E} [\epsilon_{1,t}^{\perp \mathbf{x}_t} r_t^{\perp \mathbf{x}_t}] = \mathbb{E} [\epsilon_{1,t}^{\perp \mathbf{x}_t} r_t^{\perp \mathbf{x}_t}] \end{aligned}$$

where the second line uses the exogeneity conditions and final line follows from the unit shock normalisation. From here we have: $\beta_{1S}^h = \Theta_{h,21}$. \square

A.2 Potential-Outcomes Framework

We now summarise how identification with controls can alternatively be understood through the lens of the potential outcomes framework (Angrist and Kuersteiner, 2011; Angrist et al., 2018), without appealing to the specific structural moving-average model of Section 2.1.

We define the potential outcome $y_{t,h}(r)$ as the value that the observed outcome variable y_{t+h} would have taken if $r_t = r$ for all $r \in \mathcal{R}$ where \mathcal{R} denotes the space of possible values for the policy. The average causal effect on y_{t+h} of setting $r = 1$ in time t relative to $r = 0$ is then:

$$\mathbb{E} [y_{t,h}(1) - y_{t,h}(0)] \tag{A1}$$

This is analogous to equation (2), reframed in terms of counterfactual outcomes for y under different levels of r , rather than structural shocks to r . Just as we cannot directly observe structural shocks, we cannot observe counterfactual outcomes and so the expectation in equation (A1) cannot be estimated directly. However, it is straightforward to show that the one-step regression (5) recovers exactly this causal effect of interest under the following assumption:

Condition 2. (Conditional Independence) $y_{t,h}(r) \perp\!\!\!\perp r_t | \mathbf{x}_t \forall r \in \mathcal{R}$.

To see this, first note that observed outcomes y_{t+h} relate to potential outcomes $y_{t,h}(r)$ via:

$$y_{t+h} = \sum_{r \in \mathcal{R}} y_{t,h}(r) \mathbb{1}[r_t = r]$$

Using this relation plus Condition 2, we can then write the causal effect of interest conditional on r_t and \mathbf{x}_t as the difference in observable conditional means:

$$\mathbb{E} [y_{t,h}(1) - y_{t,h}(0) | r_t, \mathbf{x}_t] = \mathbb{E} [y_{t+h} | r_t = 1, \mathbf{x}_t] - \mathbb{E} [y_{t+h} | r_t = 0, \mathbf{x}_t]$$

Then when the conditional expectation function is linear, recognise that the right hand-side is exactly the regression coefficient on r_t from the one-step OLS regression (5). Finally the unconditional average causal effect is simply recovered by the law of iterated expectations:

$$\begin{aligned}\mathbb{E}[y_{t,h}(1) - y_{t,h}(0)] &= \mathbb{E}[\mathbb{E}[y_{t,h}(1) - y_{t,h}(0)|r_t, \mathbf{x}_t]] \\ &= \mathbb{E}[\beta_{1S}^h] = \beta_{1S}^h\end{aligned}$$

This ‘selection-on-observables’ condition captures the idea that, conditional on \mathbf{x}_t , r_t is as good as random—i.e., that \mathbf{x}_t captures all relevant ‘confounding factors’ that simultaneously drive r_t and the outcome variable of interest. This is the counterpart to the exogeneity assumptions discussed in Condition 1, and similarly reflects a strong assumption, requiring specific knowledge of the process determining r_t in order to be plausible.

This setup can also be extended to cases where average causal effects are not the focus. If we define a different object of interest, the causal effect of r_t on conditional quantiles of y_{t+h} :

$$\mathbb{Q}_\tau(y_{t,h}(1)|\mathbf{x}_t, r_t) - \mathbb{Q}_\tau(y_{t,h}(0)|\mathbf{x}_t, r_t) \quad (\text{A2})$$

we can show that, under linearity of the conditional quantile function and Condition 2, the causal effect of interest (A2) is equal to $\beta_{1S}^h(\tau)$ in the following quantile regression:

$$y_{t+h} = \beta_{1S}^h(\tau)r_t + \delta(\tau)\mathbf{x}_t + u_{t+h}^{1S}(\tau) \quad (\text{A3})$$

Proceeding similarly to the discussion above, note that we can write the causal effect of r_t on conditional quantiles of y_{t+h} as the difference in observable conditional quantiles:

$$\mathbb{Q}_\tau(y_{t,h}(1)|\mathbf{x}_t, r_t) - \mathbb{Q}_\tau(y_{t,h}(0)|\mathbf{x}_t, r_t) = \mathbb{Q}_\tau[y_{t+h}|r_t = 1, \mathbf{x}_t] - \mathbb{Q}_\tau[y_{t+h}|r_t = 0, \mathbf{x}_t]$$

where we have used Condition 2 and the relation between potential and observed outcomes. Then again when the conditional quantile function is linear the right-hand side is simply the coefficient on r_t from the one-step quantile regression (A3), giving the desired result:

$$\mathbb{Q}_\tau(y_{t,h}(1)|\mathbf{x}_t, r_t) - \mathbb{Q}_\tau(y_{t,h}(0)|\mathbf{x}_t, r_t) = \beta_{1S}^h(\tau)$$

Note that, unlike for OLS, the causal effect on unconditional quantiles is not recoverable through the law of iterated expectations. \square

B Additional Results on Standard Errors

B.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 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 one- and two-step estimates $\hat{\beta}_{1S}$ and $\hat{\beta}_{2S}$ using [White \(1980\)](#) robust standard-error formulas are:

$$\widehat{\text{Var}}(\hat{\beta}_{1S}) = (\hat{\varepsilon}'_t \hat{\varepsilon}_t)^{-1} \hat{\varepsilon}'_t \hat{\Sigma}_{1S} \hat{\varepsilon}_t (\hat{\varepsilon}'_t \hat{\varepsilon}_t)^{-1} \quad (\text{B1})$$

$$\widehat{\text{Var}}(\hat{\beta}_{2S}) = (\hat{\varepsilon}'_t \hat{\varepsilon}_t)^{-1} \hat{\varepsilon}'_t \hat{\Sigma}_{2S} \hat{\varepsilon}_t (\hat{\varepsilon}'_t \hat{\varepsilon}_t)^{-1} \quad (\text{B2})$$

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

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

where $\mathbf{a} = (\hat{\varepsilon}'_t \hat{\varepsilon}_t)^{-1} \hat{\varepsilon}'_t$ and $\hat{\Sigma} = \text{diag}((\hat{\mathbf{u}}_t^{2S})^2 - (\hat{\mathbf{u}}_t^{1S})^2)$. This difference is weakly positive if and only if the diagonal matrix $\hat{\Sigma}$ 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 $\hat{\Sigma}$ is positive. However, since $(\hat{u}_t^{2S})^2$ 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 $\widehat{\text{Var}}(\hat{\beta}_{2S})$ will indeed be larger than $\widehat{\text{Var}}(\hat{\beta}_{1S})$.

Heteroskedastic-and-Autocorrelation-Robust Standard Errors. Estimated standard errors for the one- 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 (B1) and (B2), but with the following form for $\hat{\Sigma}_{1S}$ and $\hat{\Sigma}_{2S}$:

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

Like 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 $(\hat{u}_t^{1S})^2$ is on average larger than $(\hat{u}_t^{2S})^2$ by construction will generally tend to inflate standard errors for the two-step vs. the one-step approach.

B.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 IV settings. Specifically, we show the following for estimated standard errors for $\hat{\beta}_{1S}^{IV}$ and $\hat{\beta}_{2S}^{IV}$ defined in equations (13) and (14):

$$\widehat{\text{Var}}(\hat{\beta}_{2S}^{IV}) \geq \widehat{\text{Var}}(\hat{\beta}_{1S}^{IV}) \quad (\text{B3})$$

Estimated variances for $\hat{\beta}_{1S}^{IV}$ and $\hat{\beta}_{2S}^{IV}$ have the following form under homoskedasticity (see,

e.g., [Angrist and Pischke, 2014](#), pp. 140):

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

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

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

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

$$\hat{m}_t^{2S} = \hat{\varepsilon}_t \hat{\beta}_{2S} \quad \text{and} \quad \hat{m}_t^{1S} = r_t \hat{\beta}_{1S} + \mathbf{x}'_{1,t} \hat{\delta}_1$$

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

Finally, $\text{Var}((\hat{m}_t^{1S})^{\perp \mathbf{x}_{1,t}}) = \text{Var}(\hat{m}_t^{2S})$ which ensures the inequality (B3) holds:

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

where line 1 substitutes in definitions, line 2 follows from standard OLS-algebra, line 3 substitutes in definitions and line 4 follows from Frisch-Waugh-Lovell Theorem. \square

C Internal Instruments in Structural VARs

Here, we analytically demonstrate how our results for estimation in an OLS setting carry over to SVAR settings in which identification is achieved by using orthogonalised shocks as ‘internal instruments’ ([Plagborg-Møller and Wolf, 2021](#)) in a recursive SVAR. We derive exact expressions linking estimates of *contemporaneous* responses from a two-step internal-instrument approach to a one-step approach that includes confounders directly in a recursive SVAR.

Recursive-SVAR Setting. Let ε_t be the OLS-population residual from a regression of r_t on $\mathbf{x}_{1,t}$. We also define a vector of outcomes $\mathbf{y}_t = [y_{1,t}, \dots, y_{n,t}]'$ 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 $i = 1, \dots, n$:

$$y_{i,t} = \varepsilon_t \beta_{2S} + \underbrace{\sum_{j=1}^p \Gamma_j^{2S} \mathbf{y}_{t-j} + \sum_{j=1}^p \pi_j^{2S} \varepsilon_{t-j} + u_t^{2S}}_{= \mathbf{x}'_{2,t} \boldsymbol{\alpha}} \quad (\text{C1})$$

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

$$y_{i,t} = \varepsilon_t \beta_{Hyb} + \sum_{j=0}^p \Phi_j \mathbf{x}'_{1,t-j} + \underbrace{\sum_{j=1}^p \Gamma_j^{Hyb} \mathbf{y}_{t-j} + \sum_{j=1}^p \pi_j^{Hyb} \varepsilon_{t-j} + u_t^{Hyb}}_{= \mathbf{x}'_{2,t} \boldsymbol{\phi}_2} \quad (\text{C2})$$

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 :

$$y_{i,t} = r_t \beta_{1S} + \sum_{j=0}^p \Theta_j \mathbf{x}'_{1,t-j} + \sum_{j=1}^p \Gamma_j^{1S} \mathbf{y}_{t-j} + \sum_{j=1}^p \pi_j^{1S} r_{t-j} + u_t^{1s} \quad (\text{C3})$$

since controlling for r_{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 r_t in a standard recursive SVAR with r_t ordered after $\mathbf{x}_{1,t}$ and before \mathbf{y}_t .

Recursive-SVAR Results. Defining $\mathbf{x}_{1,t}^p$ as the vector $[\mathbf{x}'_{1,t}, \mathbf{x}'_{1,t-1}, \dots, \mathbf{x}'_{1,t-p}]$, we have the following relationship between β_{2S} from regression (C1) and β_{1S} from regression (C3):

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

where the \mathbf{B} -matrix is defined analogously to Result 1, and Φ^p collects the vector of coefficients $[\Phi_1, \Phi_2, \dots, \Phi_p]$ from regression (C2). The OVB term now captures the omission of contemporaneous *and lagged* $\mathbf{x}'_{1,t}$ from the two-step approach.

D External Instruments in Structural VARs

Here, we demonstrate how our results for IV estimation carry over to SVAR settings where 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 exact expressions 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 4.2, as the OLS-population residual from a regression of r_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 de-

fined via the following IV-regression for each variable $w_{i,t} \in \mathbf{w}_t$, where $i = 1, \dots, n, n+1$:

$$w_{i,t} = m_t \beta_{0,i1}^{2S} + \underbrace{\sum_{j=1}^p \mathbf{w}'_{t-j} \boldsymbol{\alpha}_j}_{\equiv \mathbf{x}'_{2,t} \boldsymbol{\alpha}} + u_t^{2S} \quad (\text{D1})$$

with ε_t as an instrument for m_t . It is immediately obvious that equation (D1) is just a special case of equation (13), 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:

$$\Phi_{h,1}^{2S} = \mathbf{C}_h \beta_{0,1}^{2S} \quad (\text{D2})$$

where $\beta_{0,1}^{2S}$ 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:

$$A(L)\mathbf{w}_t = \boldsymbol{\eta}_t \quad (\text{D3})$$

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

Following the logic of Section 4.2, the contemporaneous responses of each variable could instead be recovered via a one-step IV (population) regression with $\mathbf{x}_{1,t}$ as controls:

$$w_{i,t} = m_t \beta_{0,i1}^{1S} + \mathbf{x}'_{1,t} \boldsymbol{\theta}_1 + \underbrace{\sum_{j=1}^p \mathbf{w}'_{t-j} \boldsymbol{\theta}_j}_{\equiv \mathbf{x}'_{2,t} \boldsymbol{\theta}_2} + u_t^{1S} \quad (\text{D4})$$

with r_t as an instrument for m_t . Again, this is just a special case of equation (14) setting $\mathbf{x}_{2,t}$ as the p vectors of lagged controls $[\mathbf{w}_{t-1}, \dots, \mathbf{w}_{t-p}]$.

In this case, the entire impulse response can then be constructed as before using the same reduced-form VAR coefficients as equation (D2) to project-out across horizons:

$$\Phi_{h,1}^{1S} = \mathbf{C}_h \beta_{0,1}^{1S} \quad (\text{D5})$$

SVAR-IV Results. Comparing equations (D5) and (D2), impulse responses from a one- and two-step approach differ only in their construction of the contemporaneous coefficients $\beta_{0,1}$. As in Section 4.2, 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:

$$\Phi_{h,1}^{2S} = \Phi_{h,1}^{1S} + \Omega_h^{SVAR-IV} = \Phi_{h,1}^{1S} + \mathbf{C}_h \frac{\Omega_w^{OLS}}{\Omega_m^{OLS}}$$

where Ω_w^{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:

$$\Omega_{w_i}^{OLS} = \mathbb{E} [\varepsilon_t \mathbf{x}_{2,t}] \mathbf{B}_{w_i}^{-1} \mathbb{E} [\mathbf{x}_{2,t} \mathbf{x}'_{1,t}] \phi_{w_i} \quad \text{and} \quad \Omega_m^{OLS} = \mathbb{E} [\varepsilon_t \mathbf{x}_{2,t}] \mathbf{B}_m^{-1} \mathbb{E} [\mathbf{x}_{2,t} \mathbf{x}'_{1,t}] \phi_m$$

where $\Omega_{w_i}^{OLS}$ is the i -th element of Ω_w^{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.

ity 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 w_t) that affect the outcome variable. Although these variables are included as controls in equation (D1), this serves to reintroduce correlation with $x_{1,t}$, 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 $x_{1,t}$ as exogenous variables in an SVAR-IV since, in our case, $x_{1,t}$ 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.

E Simulation Study

Our simulation study employs the 9-variable VAR(4) from [Christiano et al. \(2005\)](#) as a medium-scale, empirically-realistic DGP. Specifically, we use quarterly data over 1965Q3-1995Q3, and estimate a recursive VAR using data on: real GDP, real consumption, real investment, GDP deflator prices, real wages, labor productivity, federal funds rate, real profits, and the growth rate of M2 - where we list the variables listed in their recursive order. All data is sourced from FRED.³⁰ In the original paper, all variables except the federal funds rate and M2 growth rate enter the VAR in log levels. In our baseline results, we instead stationarise all data before including it in the VAR – so enter real GDP, real consumption, real investment, GDP deflator prices, real wages, labor productivity, and real profits in log differences, alongside the federal funds rate (in levels) and M2 growth rate. We also display results for the original specification with relevant variables in log levels. We then estimate ‘true’ IRFs for each variable in our VAR to the identified monetary policy shock (i.e. the shock associated with the 7th equation in the VAR, for the federal funds rate) out to 4-years. We simulate data from the VAR to construct 2000 Monte Carlo simulations, each of sample-length $T=121$ to match [Christiano et al. \(2005\)](#). Across simulations, we estimate IRFs for all outcome variables to a monetary policy shock using various estimators:

1. A ‘one-step’ local projection estimator, that estimates IRFs by regressing each outcome variable on the federal funds rate, controlling for four lags of the nine variables in the VAR plus contemporaneous values of the variables that appear before the federal funds rate in the recursive order.

³⁰We follow the simulation study from [Herbst and Johannsen \(2024\)](#). Full details on variable construction can be found in the Supplemental Material from that paper.

2. A 'two-step' local projection estimator, that (i) first constructs an estimate for the monetary policy shock by regressing the federal funds rate on four lags of the nine variables in the VAR plus contemporaneous values of the variables that appear before the federal funds rate in the recursive order; then (ii) estimates IRFs by regressing each outcome variable on the constructed shock (without any auxiliary controls).
3. A 'two-step' local projection estimator, that (i) first constructs an estimate for the monetary policy shock by regressing the federal funds rate on *only* the contemporaneous values of the variables that appear before the federal funds rate in the recursive order; then (ii) estimates IRFs by regressing each outcome variable on the constructed shock controlling for four lags of the nine variables in the VAR.
4. A correctly specified VAR (4), that estimates IRFs in exactly the same way as the original construction of the DGP.

Note that the second and third estimator correspond to a version of the 'two-step' estimator considered in the main text under Case 1 (with empty $\mathbf{x}_{2,t}$) and Case 2 (with non-empty $\mathbf{x}_{2,t}$) respectively. We seek to compare these to the one-step local projection estimator, and additionally estimate a VAR(4) purely as a comparison against which to benchmark differences between the one-step and two-step local projections. For the local projection estimators, we compute Newey-West standard errors, and we compute standard errors for the VAR(4) using a bootstrap procedure.³¹ We compute Monte Carlo estimates of the absolute bias of estimated IRFs – $|E[\hat{\beta}^h - \beta^h]|$ – as well as estimated standard errors – $E[\widehat{Var}(\hat{\beta}^h)]$ – for each outcome variable and horizon. In order to display a summary of our results, we then compute appropriately scaled median estimates for both bias and standard errors across outcome variables and horizons. Specifically, we scale the estimated standard errors – $E[\widehat{Var}(\hat{\beta}^h)]$ – for each estimator relative to the estimates from the VAR(4), and then take the median of this ratio across outcome variables and across horizons, separately for short horizons (< 1 year) and longer horizons (1 – 4 years). And, following Li et al. (2021), we scale the estimated absolute bias – $|E[\hat{\beta}^h - \beta^h]|$ – by the root mean squared value of the true impulse response function out to the four-year horizon – i.e. dividing by $\sum_{h=0}^{16} \sqrt{\beta_h}$ – before taking the median across outcomes and horizons.

Results We report relative standard errors and biases of estimators in the main text.

³¹For the local projections, standard errors are computed using a Newey-West HAC estimator with a horizon-dependent truncation lag, without prewhitening or finite-sample adjustments. For the VAR, the bootstrap is constructed through (i.i.d.) sampling with replacement of the residuals (without block- or wild- bootstrapping).

F Empirical Application: Additional Results

F.1 Data Sources

We use monthly data. Our dependent variable, the US Consumer Price Index (CPI), is sourced from *FRED*, and we use additional macroeconomic controls in $x_{2,t}$ —seasonally-adjusted US industrial production and the US unemployment rate—from the same source. To estimate [Romer and Romer \(2004\)](#) 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.

F.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—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. To do so, for months in which no FOMC announcements occurred, we set the Federal Funds target rate change to zero, Greenbook forecasts equal to their last value, and forecast revisions to zero. 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. As we go on to discuss, this constitutes a minimal difference relative to previous work. Following [Romer and Romer \(2004\)](#), we construct the shock by estimating (15). Column (1) of Table [E1](#) presents the estimated coefficients from this regression.

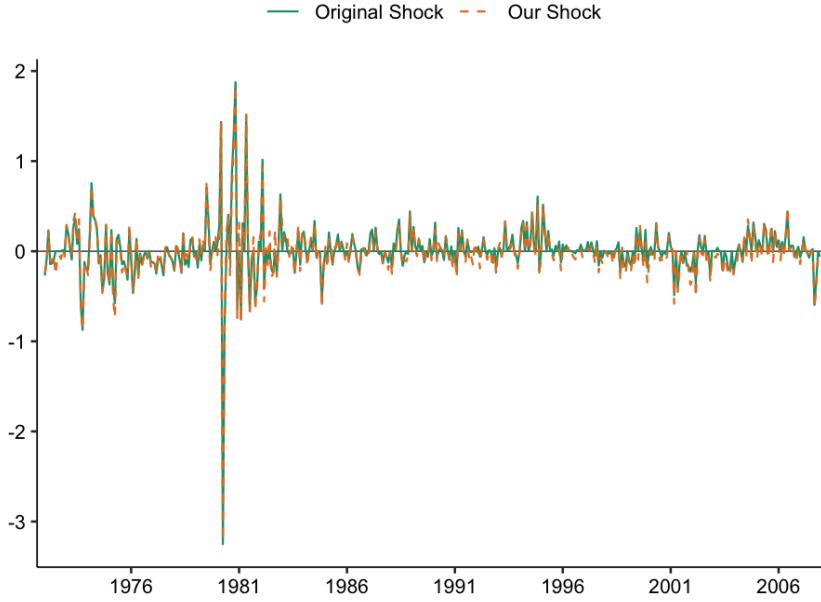
Monthly vs. Meeting Frequency. As discussed, all regressions in the main body are estimated at monthly frequency to foster comparability between LPs and VARs. The approach taken in previous work is instead to construct the shock using regressions at meeting frequency, and then convert this shock-series to monthly by imposing zeros for non-meeting months (and summing any shocks within a month). We show that in practice this constitutes a negligible difference. To support this, we first show in Column (2) of Table [E1](#) that the first-stage regression coefficients with the policy rule estimated at meeting frequency are similar to monthly-frequency estimates in Column (1). We also compare our monthly-constructed shock to the updated shock from [Wieland and Yang \(2020\)](#), which uses the full sample and estimates the regression at meeting frequency then converts to monthly. We plot the two series in Figure [E1](#) which almost exactly overlap. We find a correlation of 0.98 between the two shock series.

Table E1: First-Stage Regressions: The Romer-Romer Reaction Function

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

Figure E1: Comparison of Shock Series



Notes: Chart compares our baseline monetary policy shocks series (dashed orange) with the updated [Romer and Romer \(2004\)](#) shock series from [Wieland and Yang \(2020\)](#). The latter is taken from the replication package of [Ramey \(2016\)](#).

F.3 Alternate Approach to Mixed-Frequency Data

We now discuss an alternate way of comparing between one-step and two-step local projections that still allows us to project directly on the shock estimated at meeting-frequency. This extends our discussion of one-step and two-step approaches to a case where the policy-variable is observed at a lower frequency (e.g. 6-weekly) than the outcome variable (monthly). Note that the common approach in applied work is to estimate the shock at meeting-frequency then convert to monthly by imposing zeros for months in which no meeting occurred. An alternate approach does not impose zeros for these months but instead sets data to missing, and then estimates local projections on the constructed shock, with observations removed for months with missing data. When conditional expectations are linear and data are missing at random, this approach estimates exactly the object of interest:

$$E[y_{t+h}|\epsilon_t^m = 1] - E[y_{t+h}|\epsilon_t^m = 0] \quad (E1)$$

without requiring the researcher to impose a value for monetary policy shocks in non-meeting months. Note that this approach is simply equivalent to a one-step local projection where the right-hand side variables are at meeting frequency, and the left-hand side variable denotes the outcome in the corresponding month, shifted forward by h -months. We additionally consider this approach in our empirical application. As well as our baseline 'monthly fre-

Table E2: Response of $\ln(CPI)$ to monthly- and meeting-frequency US monetary policy shocks across horizons h estimated by LP-OLS

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

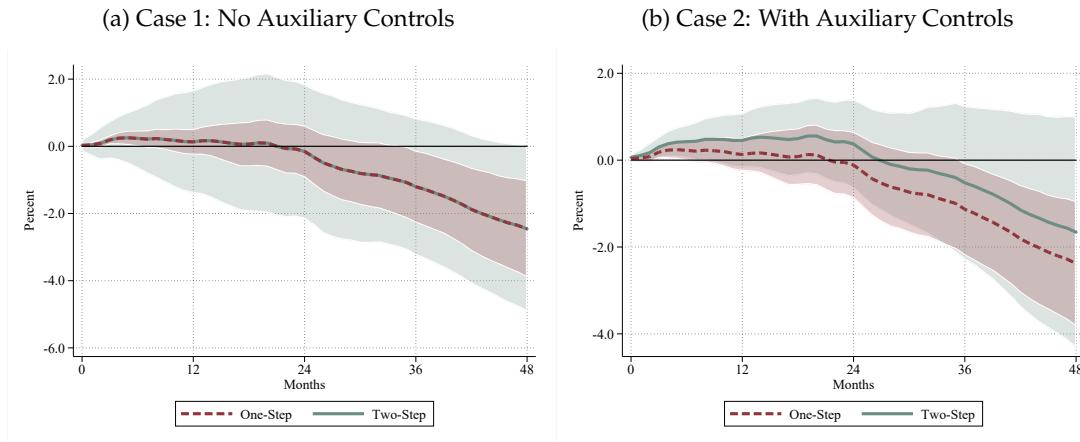
quency approach' we additionally estimate one-step and two-step IRFs from this 'meeting frequency' approach. Table E2 demonstrates how, for the various cases presented in Section 6.1, estimated impulse responses are similar when estimated at monthly and meeting frequency. In particular, as with our baseline estimates, for the meeting-frequency estimates, the near-term 'price puzzle' is mitigated with the one-step approach, and the decline in CPI at longer horizons is highly significant (unlike the two-step approach). A full discussion of the appropriate method for handling mixed-frequency data is beyond the scope of this paper. Various work discuss identification with mixed-frequency data for VAR models (see e.g. [Foroni and Marcellino, 2016](#)) and distributed-lag regressions (see e.g. [Chudik and Georgiadis, 2022](#)). Our results suggest that an interesting avenue for future research is the development of tools for handling mixed-frequency that permits one-step estimation via local projections.

F4 Additional Results for IV Application

LP-IV Application with Romer-Romer Shocks. 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 yield (which we denote m_t to align with Section 4.2). In the two step, we use $\hat{\varepsilon}_t^{mp}$ as an instrument for m_t in regression (18) to estimate the dynamic effects of US monetary policy on US CPI. We compare this to the one-step LP-IV regression (19), in which Δi_t is used as an instrument for m_t .

Figure E2 presents the estimated impulse responses from the one- and two-step. Unsurprisingly, given the discussion in Section 4.2, the estimates for both approaches align closely with those attained from LP-OLS—as a comparison of Figure E2 with Figures 1a and 1b reveals. As before, unlike two-step estimate, the price puzzles in one-step estimates are limited.

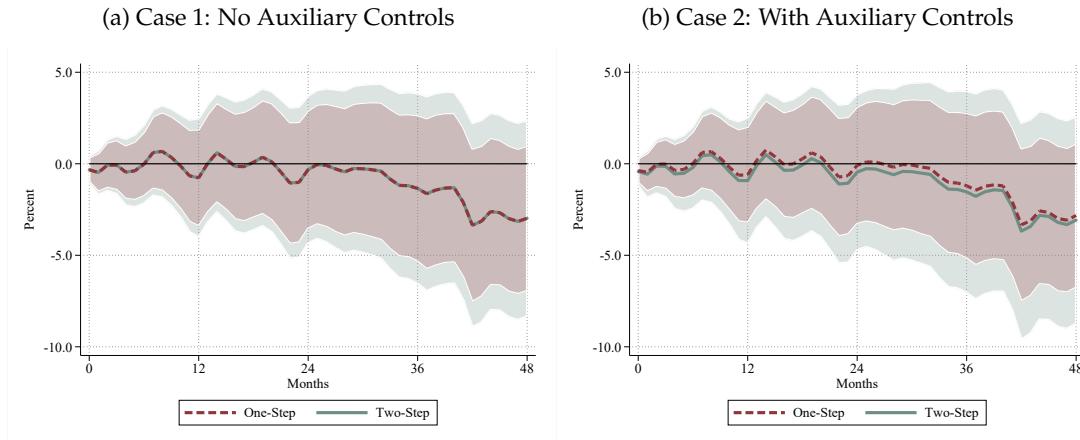
Figure E2: Estimated impulse responses of $\ln(CPI)$ to US monetary policy shock from LP-IV



Notes: Estimated IRF 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. Shaded area denotes 90% confidence bands constructed from heteroskedasticity and autocorrelation robust standard errors. Sample: 1972:01-2007:12

LP-IV Application with Monetary Surprises. Figure E3 presents the estimated IRFs from the [Miranda-Agrippino and Ricco \(2021\)](#) LP-IV application described in Section 6.2.

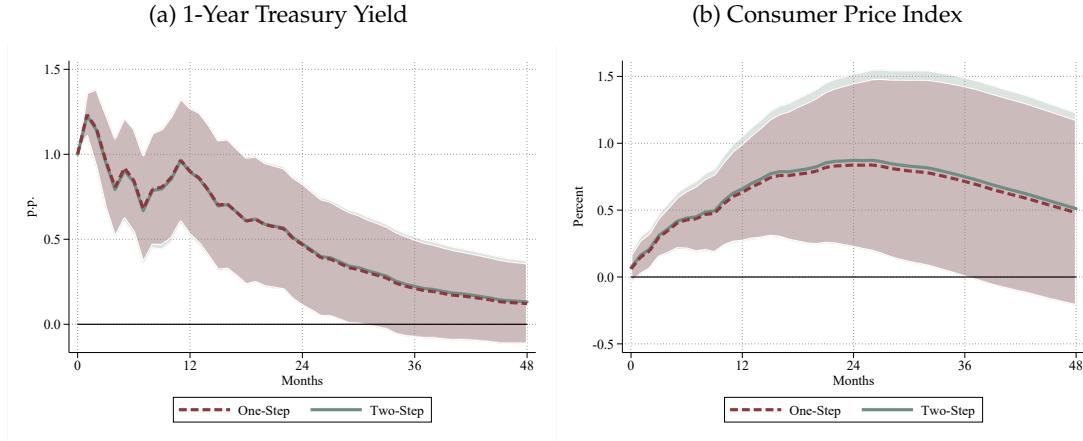
Figure E3: Estimated impulse responses of $\ln(CPI)$ to US monetary policy shock from LP-IV



Notes: Estimated IRF 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. Shaded area denotes 90% confidence bands constructed from heteroskedasticity and autocorrelation robust standard errors. Sample: 1990:01-2007:12.

Proxy-SVAR Application with Romer-Romer Shocks. Figure E4 presents the estimated IRFs from the [Romer and Romer \(2004\)](#) Proxy-SVAR application described in Section 6.2.

Figure E4: Estimated impulse responses to US monetary policy shock from Proxy SVAR



Notes: Estimated IRF 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. Shaded area denotes 90% confidence bands constructed from [Jentsch and Lunsford \(2019\)](#) residual-based moving block bootstrap.

G Further Applications

In this Appendix, we provide further information underpinning the additional applications we summarise in Section 6.4. Rather than replicate exactly the original specifications from these papers, our intention is to demonstrate the differences between one- and two-step approaches in empirical settings that match as closely as possible those in which the two-step procedure has been employed. Following our main application, we estimate two-step and one-step local projections of the form:

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

$$y_{t+h} - y_{t-1} = \alpha^h + \mathbf{x}'_{1,t} \boldsymbol{\zeta}^h + \mathbf{x}'_{2,t} \boldsymbol{\alpha}^h + z_t \beta_{2S}^h + u_{t+h}^{2S} \quad (3)$$

Table 3 summarises the outcome variables and horizons across which we estimate impulse responses for each study. Table 4 summarises the auxiliary control variables $\mathbf{x}_{2,t}$. In each subsection we then describe the construction of the shock – i.e. the choice of z_t and $\mathbf{x}_{1,t}$ alongside additional details on estimation where relevant.

G.1 Oil-Specific Demand Shocks à la [Kilian \(2009\)](#)

[Kilian \(2009\)](#) builds a structural VAR model of the global crude oil market to jointly identify three structural shocks: crude oil supply shocks; shocks to the global demand for all industrial commodities; and demand shocks that are specific to the global crude oil market. The focus of our application, and a key contribution of the original paper, is on the last shock, which is designed to capture shifts in the price of oil driven by higher precautionary demand, for instance those associated with market concerns about the availability of future oil supplies.

Table 3: Outcome Variables and Horizons for Additional Applications

Paper	Outcomes	Horizons
Romer and Romer (2004)	Consumer Prices, IP, Unemployment	0-48 (months)
Miranda-Agrippino and Ricco (2021)	Consumer Prices, IP, Unemployment	0-48 (months)
Barnichon and Mesters (2021)	Philips Multiplier	0-48 (months)
Kilian (2009)	Consumer Prices, IP, Unemployment	0-48 (months)
Bilal and Käenzig (2024)	World GDP	0-10 (years)

Notes: Table summarises details of impulse responses used in each application. IP refers to Industrial Production. IP, Consumer Prices and GDP enter the local projections in logs. All variables are for the US unless otherwise stated.

Table 4: Auxiliary Controls for Additional Applications

Paper	Controls	Lags
Romer and Romer (2004)	Consumer Prices, IP, Unemployment	12
Miranda-Agrippino and Ricco (2021)	Consumer Prices, IP, Unemployment	12
Barnichon and Mesters (2021)	Consumer Prices, Unemployment	12
Kilian (2009)	Consumer Prices, IP, Unemployment	12
Bilal and Käenzig (2024)	World GDP	4

Notes: Table summarises details of auxiliary controls used to construct impulse responses in each application. IP refers to Industrial Production. IP, Consumer Prices and World GDP enter as controls in log-differences; Unemployment enters as controls in first-differences. All variables are for the US unless otherwise stated.

The VAR model is based on monthly data for three variables with 24 lags: $\Delta prod_t$ is the percent change in global crude oil production, $reat_t$ denotes an index of real economic activity constructed and described in Kilian (2009); and $rpot_t$ refers to the real price of oil. The latter two series are expressed in logs. We source this data from the replication package from the original paper.

Structural shocks are identified recursively. The oil-specific demand shock is assumed orthogonal to contemporaneous moves in oil prices or global demand. In turn, innovations to global real activity that cannot be explained based on crude oil supply shocks reflect aggregate demand shocks. The oil supply shock results in an instantaneous change in all variables.

G.2 Philips Multiplier à la Barnichon and Mesters (2021)

Barnichon and Mesters (2021) use the Romer and Romer (2004) shock to identify the "Philips Multiplier". Specifically they first run regressions of the form:

$$\sum_{j=0}^h u_{t+j} = \alpha_u^h + \mathbf{x}'_{2,t} \zeta_u^h + \hat{\varepsilon}_t^{mp} \beta_u^h + u_{t+h} \quad (4)$$

$$\sum_{j=0}^h \pi_{t+j} = \alpha_\pi^h + \mathbf{x}'_{2,t} \zeta_\pi^h + \hat{\varepsilon}_t^{mp} \beta_\pi^h + e_{t+h} \quad (5)$$

where $\hat{\varepsilon}_t^{mp}$ denotes the [Romer and Romer \(2004\)](#) monetary policy shock, u_t and π_t denote unemployment and CPI inflation respectively.³² They then compute the "Philips Multiplier" as:

$$\beta_{pm}^h = \beta_\pi^h / \beta_u^h \quad (6)$$

We construct [Romer and Romer \(2004\)](#) monetary policy shocks as in our main application. We compare β_{pm}^h to a one-step estimate that replaces $\hat{\varepsilon}_t^{mp}$ with the change in the nominal interest rate, and control directly for the Greenbook forecasts in the local projections.

G.3 Temperature Shocks à la [Bilal and Käenzig \(2024\)](#)

[Bilal and Käenzig \(2024\)](#) identify shocks to global temperature $Temp_t$ by regressing $Temp_t$ on its own lags. Their baseline shock series is constructed by regressing $Temp_t$ on $Temp_{t-3}$ and $Temp_{t-4}$ – we instead construct temperature shocks by regressing $Temp_t$ on all its first four lags (which the original study also considers as a robustness exercise). We then run local projections of World GDP on this shock (with lags of World GDP), and compare this to a one-step estimate that regresses world GDP on $Temp_t$ controls directly for four lags of $Temp_t$.³³ We source the data from the online replication materials provided by the authors.

³²The original study works with quarterly data but we use monthly. The original study uses quarterly changes in the PCE price level while we employ monthly changes in CPI

³³In an updated draft, the authors themselves consider this approach and like us find little difference in estimated impulse responses.