CENTRE FOR MACROECONOMICS

Economic and Social Research Council

Controls, Not Shocks: Estimating Dynamic Causal Effects in Macroeconomics

CFM-DP2024-22

Simon Lloyd[1,2] and Ed Manuel[3]

[1]Bank of England, [2]Centre For Macroeconomics, [3]London School of Economics and Political Science

Estimates of the causal effects of policy are essential for effective policymaking, but identifying them is complicated by the presence of confounding factors (i.e., variables that simultaneously drive policy and macro outcomes). Within empirical macro, it is common for researchers to deal with this via a 'two-step' procedure. First, they regress the causal variable (e.g., interest rates) on the set of confounding factors (e.g. central bank forecasts for GDP and inflation) and save the residuals—i.e., estimating shocks as deviations from an estimated policy 'reaction function'. Second, armed with these shocks, they identify causal effects by estimating the association between the constructed shock and the outcome variable of interest, typically through a local projection (LP) or vector auto-regression (VAR). However, intuitively, another way to account for the effect of confounding factors is through a simple 'one-step' regression that includes confounders as control variables, without the need to first estimate a series of shocks.

In this paper, we compare the two and show that the two-step approach is problematic. We do so by considering a general econometric framework, with limited assumptions on the structure of the economy or the specific estimation procedure. Within our framework, we derive an analytical relationship between the one- and two-step estimators. Across a range of common applications (incl. OLS, IV, VARs, LPs, Quantile Regression) we show that, relative to the simple one-step regression, the two-step results in some combination of mis-estimation of standard errors, inefficiency and/or bias. In the best cases, limited to specific OLS/IV settings, the one- and two-step approaches will deliver identical coefficient estimates. But, even in this case, the two-step has practical drawbacks: estimated standard errors will be larger than their true values, so inference will therefore be too conservative. In other cases (e.g., where additional controls are included in the second stage, including VARs) and for other estimation methods (e.g., quantile regressions), differences between the one-step and two-step approach can be expressed in terms of an omitted-variable bias in the latter.

We demonstrate our results with an empirical application, estimating the dynamic effects of US monetary policy on US CPI inflation when controlling for central-bank forecasts, as in Romer and Romer (2004). To date, this literature has relied extensively on the two-step approach (i.e., first estimating a series of monetary policy shocks and then employing them in a secondary regression). For a range of estimation methods, we show that the two-stage approach can give rise to excessively conservative standard errors and can generate bias.

Correcting for this through a one-step approach removes a significant portion of the 'price puzzle'—a long-standing finding in some monetary policy applications that, counterintuitively, estimates that tighter monetary policy 'causes' higher inflation in the near term. As a result, our application delivers results that suggest that estimates of monetary policy's causal effects are more robust than previously realised.