Non-technical summary for 'The Unit-effect Normalisation in Set-identified Structural Vector Autoregressions'

By Matthew Read

Estimating the response of the economy to macroeconomic shocks (e.g. unexpected changes in monetary policy) is difficult, because it requires disentangling the effects of the specific shock from the effects of other shocks hitting the economy. Macroeconomists often do this using econometric models called structural vector autoregressions (SVARs). Roughly speaking, these models decompose the correlations among a group of macroeconomic variables such that the components have a causal interpretation. This decomposition requires assumptions about the economy's structure, which are known as 'identifying restrictions'.

A popular type of identifying restriction is a 'sign restriction', which is an assumption about the sign of an economic relationship. For example, researchers might assume that inflation doesn't increase following an unexpected increase in the policy rate, which is consistent with predictions from conventional macroeconomic theory. These types of assumptions are arguably more believable than other commonly used assumptions. However, imposing these 'weaker' restrictions comes at the expense of only being able to determine a *set* of possible responses. In this case, the responses to the shocks are 'set identified'.

When using these models, researchers typically assume that macroeconomic shocks have a standard deviation of one (the 'standard deviation normalisation'). This means that the estimated responses are to shocks that are one standard deviation in magnitude. However, answering important economic questions often requires knowing the effects of a 'unit' shock. For example, central bankers want to answer questions like 'what are the effects of a 100 basis point change in the policy rate?' These responses are obtained by computing the responses to a shock that raises a particular variable by one unit (the 'unit-effect normalisation'), which involves dividing the response of the variable of interest (e.g. GDP) by the response of another variable (e.g. the policy rate).

The problem that arises in this setting is that set-identifying restrictions may admit the possibility that a variable (e.g. the policy rate) does not respond to its own shock (e.g. the monetary policy shock). This means that we may be dividing by something that equals (or is arbitrarily close to) zero. As a consequence, the set of responses to a unit shock may be 'unbounded' (i.e. infinite in length), in which case we may learn very little – and sometimes nothing – about the responses to a unit shock.

The potential for set-identifying restrictions to be extremely uninformative about the responses to unit shocks has been underappreciated. A key contribution of this paper is to highlight this issue and explain why it arises. Additionally, I argue that researchers should transparently report the extent to which the set of responses may be unbounded, since this makes it clear what we actually learn about the effects of macroeconomic shocks given the data that we observe and the identifying restrictions we impose. To assist with this, I derive a condition under which the identifying restrictions *always* admit the possibility that a variable does not respond to its own shock.

I illustrate these issues by estimating the response of US output to a monetary policy shock that raises the federal funds rate by 100 basis points on impact. The existing literature contains a very wide range of estimates for this response; for example, an influential survey paper reports estimates for the peak decline in output following a contractionary shock ranging from 0.6 per cent to 5 per cent (Ramey 2016). The existing evidence therefore does not really tell us whether the macroeconomic effects of US monetary policy are mild, potent or somewhere in between. This has implications for how aggressively the federal funds rate needs to respond to any particular shock. I show that some identifying restrictions that have been used previously in the literature always admit the possibility that the federal funds rate does not respond on

impact to a monetary policy shock. Those restrictions are consequently extremely uninformative about the output effects of a 100 basis point shock. I then impose richer sets of identifying restrictions and find evidence to suggest that the output response lies towards the smaller end of the range of existing estimates.

Reference

Ramey VA (2016), 'Macroeconomic Shocks and Their Propagation', in JB Taylor and H Uhlig (eds), *Handbook of Macroeconomics: Volume 2A*, Handbooks in Economics, Elsevier, Amsterdam, pp 71–162.