

Structural VARs and the Monetary Transmission Mechanism in Low-Income African Countries¹

Bin Grace Li (IMF), Christopher Adam (University of Oxford, UK), Andrew Berg (IMF)

Peter Montiel (Williams College, MA), Stephen O’Connell (Swarthmore College, PA)

July 2, 2018

Revised November 30, 2018

Abstract

Structural Vector Autoregression (SVAR) methods suggest the monetary transmission mechanism may be weak and unreliable in many low-income African countries. But are structural VARs identified via short-run restrictions capable of detecting a transmission mechanism when one exists, under research conditions typical of low-income countries (LICs)? Using a small DSGE as our data-generating processes, we assess the impact on VAR-based inference of short data samples, measurement error, high-frequency supply shocks, and other features of the LIC environment. The impact of these features on finite-sample bias appears to be relatively modest when identification is valid—a strong caveat, especially in LICs. Nonetheless many of these features undermine the precision of estimated impulse responses to monetary policy shocks, and cumulatively they suggest that statistically and economically insignificant results can be expected even when the underlying transmission mechanism is strong. These data features not only undermine the efficacy of the SVAR methodology for research and policy-making, but are also severe enough to motivate a continued search for monetary policy rules that are robust to these limitations.

¹ This paper is part of a research project on macroeconomic policy in low-income countries supported by the U.K. Department for International Development (DFID). The views expressed herein are those of the authors and should not be attributed to the IMF, its Executive Board, its management, or to DFID. We are grateful to Olivier Blanchard, Prakash Loungani, Sam Ouliaris, Chris Papageorgiou, Rafael Portillo, and Romain Houssa for helpful discussions and to seminar participants at the Centre for Study of African Economies, Midwest Macro conference, IMF, Tri-Co summer seminar, Bank of Tanzania, and University of Maryland. Any errors are our own.

1. Introduction

To conduct monetary policy effectively, central banks need a reliable understanding of the monetary transmission mechanism (MTM) describing the links between their policy instruments and the ultimate economic outcomes they are seeking to influence. The task of generating robust estimates of the direction, strength and speed of the MTM is demanding in all economies, but particularly so in low-income countries where financial markets tend to be thin, economies are undergoing rapid structural change, and monetary policy regimes are evolving. This task is made all the harder by the relatively poor quality of macroeconomic data in many such countries. Few low-income countries produce high-frequency data on nominal and real GDP, and where quarterly data are available these rarely date back more than a decade. For example, only thirteen of the forty-five Sub Saharan African countries in the IMF databases have any quarterly GDP data and only five have data on both nominal and real GDP (Botswana, Mauritius, Rwanda, Seychelles and South Africa). Excluding South Africa, which has consistent nominal and real quarterly GDP data back to 1980, the median span of quarterly data is less than nine years. Even apparently long samples of annual data are often subject to serious limitations; as recent high-profile cases have illustrated, the lack of regular updating to underlying survey instruments means that when they do occur, revisions to national accounts data are often very large and extremely difficult to reconcile with earlier-vintage data.²

Despite these difficulties, for most central banks and many researchers in low-income countries (LICs) the method of choice employed to assess the strength and reliability of the MTM remains the structural vector auto-regression (SVAR) model, estimated using aggregate macroeconomic data. The SVAR approach works by imposing just enough restrictions on the reduced-form dynamic relationship between a monetary policy instrument and key macroeconomic outcomes to recover the responses of these outcomes to unanticipated monetary policy shocks. The predominant approach in the literature

² The rebasing of Nigerian GDP data in 2014, for example, suggested the economy was almost 90% larger than previously thought, while similar exercises in Ghana in 2010 and Kenya in 2015 resulted in increases in estimated GDP of 60% and 30% respectively. The scale of African GDP revisions reflect both rapid structural change in these economies as well as the fact that, for many, the basic surveys from which GDP data are constructed – household and labor market surveys, agricultural censuses and even population censuses -- are often incomplete and woefully outdated (see African Development Bank, 2013).

relies on short-run restrictions, defined as restrictions on the contemporaneous interactions between the reduced-form residuals of the VAR and a set of unobserved structural shocks. Christiano and Eichenbaum (2007) convey the attraction of this approach by observing that “if the relevant short-run restrictions are satisfied in the data-generating process, then standard structural VAR methods reliably uncover and identify the dynamic effects of shocks to the economy” (p. 3). The impulse responses of inflation and output to monetary-policy shocks identified via short-run restrictions have constituted the dominant form of econometric evidence on the monetary transmission mechanism in high-income countries.

The empirical harvest from SVAR applications, however, appears to differ sharply by level of development. In a survey of SVAR-based evidence on the monetary transmission mechanism in LICs, emerging-market economies, and advanced countries, Mishra, Montiel, and Spilimbergo (2012) conclude that the MTM is routinely found to be weaker and less reliable in LICs than it is in higher-income economies. By weaker, they mean that monetary policy instruments tend to have small estimated effects on aggregate demand. By less reliable, they mean that these impacts are not precisely estimated, leaving considerable statistical uncertainty about how the economy responds to monetary-policy shocks.

Mishra *et al.* (2013) suggest two broad possible explanations for these findings. One is that the SVAR findings reflect the *facts on the ground*: the MTM is simply weaker in LICs than it is in more advanced economies. Formal financial markets are small and poorly arbitrated in these countries, and many low-income countries (LICs) maintain fixed or heavily managed exchange rates. As a consequence, the link between the short-term interest rates that central banks can control and the variables that matter for aggregate demand (e.g., longer-term interest rates, the exchange rate) may in fact be weak or absent on average, as well as highly uncertain. Even the bank lending channel may tend to be weak and unreliable when the formal financial sector is small, financial frictions are severe, and the banking industry is characterized by imperfect competition.

The second explanation, however, is that the SVAR approach may be poorly suited to the research environment in LICs. As Mishra and Montiel (2013) emphasize, this explanation points to *limitations of the method* and suggests that the SVAR evidence in LICs may be relatively uninformative about the facts on the ground. Evidence for an effective MTM may be missing not because the MTM is actually weak

and unreliable but because a data-intensive approach that was developed for industrial countries is incapable of uncovering the MTM under the research conditions typical of LICs. If this explanation is correct, then it is the SVAR evidence itself that is weak and unreliable in LICs.

Sub-Saharan African countries are prominent among the SVAR-based studies reviewed by Mishra *et al.* (2012). Of the 11 African LICs covered in that study, none displays a strong and statistically significant response of both output and prices to monetary policy shocks, and only Cheng (2006) shows strong evidence for one of these two variables (monetary policy transmits strongly to prices in Kenya). Our own review of the small number of SVAR papers not included in the Mishra/Montiel study corroborates this conclusion. Tsangarides (2010), for example, finds results for Mauritius that are similar to Cheng (2006); but in contrast to Cheng, Misate *et al.* (2012) finds a weak MTM for both prices and output in Kenya. Gottschalk (2008) reports statistically significant but weak transmission to prices in Sierra Leone, in a system that excludes output altogether due to data limitations.

Our primary aim in this paper is to evaluate limitations of the method as explanations for the ‘missing’ MTM in LICs. We do this by asking whether the characteristics of a typical LIC research environment are particularly hostile to SVAR-based approaches to identification. The stakes here seem high, certainly for the weight policymakers should place on SVAR results and possibly also for the conduct of monetary policy. If the facts on the ground explanation is correct, the SVAR evidence suggests that managing monetary policy successfully may be particularly challenging in low-income countries. Along with other features of the LIC environment such as frequent large supply shocks, weak and unreliable transmission may make it more difficult for policymakers to keep inflation within narrow bounds and to stabilize activity in the face of demand shocks (Mishra *et al.* 2012). Passive monetary policies might be viewed as giving up very little in terms of macroeconomic performance.

On the other hand, if the missing MTM mainly reflects methodological limitations, then the results of the SVAR-based literature should be suitably discounted. Researchers evaluating the strength and reliability of the MTM should seek to complement VAR-based analyses with approaches that are likely to be more robust to the peculiar weaknesses of these methods in LIC-like environments.³ Policymakers, meanwhile, should incorporate the implications of a hostile data environment for policy itself, seeking

³ The use of bank-level or even loan-level data to investigate the strength of the bank-lending channel is an obvious candidate (e.g., Mbowe 2012, Abuka *et al* 2015).

approaches that are agnostic about the strength of transmission but that are robust to measurement error and parameter uncertainty.

We investigate the limitations of the method explanation by applying SVAR methods to a world in which a strong MTM indeed exists, but where the research environment is characterized by several features typical of LICs. These include poorly-understood economic structure and non-transparent central banks; short data samples due to missing data or major structural breaks arising from significant changes in the institutional and regulatory environment or as the result of reforms affecting prevailing policy regimes; large real-time measurement errors; and a high volatility of macroeconomic shocks (especially the prevalence of large temporary supply shocks).

To implement this program we set up a Monte Carlo experiment to assess the statistical properties of the kind of SVAR models typically deployed on low-income country data. Underlying the data generating process is a dynamic stochastic general equilibrium (DSGE) model of a small open economy that embodies a well-defined MTM with an interest-rate channel and an exchange-rate channel. We use the solution to this DSGE to generate multiple independent runs of data, and then within each of these runs, mimic the process of an empirical researcher using SVAR-based methods to infer the nature of the MTM. In particular, we examine the properties of the impulse response functions (IRFs) that she would produce. We compare her median estimated IRFs to the true one, study the spread of estimated IRFs across simulations, and examine the power of conventional significance tests against the hypothesis of a zero response.

In Section 2 we introduce our DSGE model in four macroeconomic variables – the GDP gap, the inflation rate, the real exchange rate, and the nominal interest rate – and discuss the relationship between the impulse responses implied by a DSGE model and those generated by a structural VAR identified using restrictions on the contemporaneous interactions between the variables. Section 3 begins by documenting the empirical success of the VAR-based approach when the researcher has chosen a valid identification scheme and is operating in a favourable research environment (as may be found in a mature open economy such as Canada, for example). Section 4 briefly quantifies the effects on inference of weak transmission – the ‘facts on the ground’ interpretation of the weak MTM – while the core of the paper's results are presented in Section 5. Here we focus on the ‘limitations of method’ interpretation by examining Monte Carlo evidence on the performance of SVAR estimates under the

research conditions listed above. In doing so, we address in turn the implications of short data samples, volatility and measurement error, and high-frequency supply shocks, both individually and in combination. We conclude this section by returning to the question of mis-identification. Section 6 concludes the paper with a summary of findings and a discussion of possible extensions and policy implications. In particular, we highlight the implications for monetary policy formulation of the types of data environments that tend to be typical in low-income countries.

2. DSGEs as a data-generating process

The MTM concerns the ability of monetary policy to exert a temporary effect on nominal aggregate demand.⁴ To focus on these effects we begin by ignoring stochastic trends in the data, implicitly assuming that these can be estimated with reasonable statistical confidence so that the stationary part of the data is cleanly isolated. Our DSGE models will therefore generate a stationary vector $x_t = [\tilde{y}_t, \pi_t, \tilde{e}_t, i_t]'$ of quarterly values for the GDP gap (\tilde{y}_t , defined as the gap between actual GDP and unobservable potential GDP), the inflation rate (π_t), the real exchange rate (\tilde{e}_t , with an increase being a real appreciation), and the annualized nominal interest rate (i_t). We introduce an underlying trend in Section 4.3, where we argue that LIC applications confront particular difficulties in inferring the GDP gap from observed measures of output. For the bulk of the paper, however, we treat the model-generated GDP gap as observable.

The four endogenous variables in the model will in turn be functions of a vector $\varepsilon_t = [\varepsilon_t^y, \varepsilon_t^\pi, \varepsilon_t^e, \varepsilon_t^i]'$ of structural shocks that are not directly observable by the researcher. The objects of interest to the researcher are the responses of x_{t+j} to a one-time unit-value shock to monetary policy ($\Delta \varepsilon_t^i = 1$). To estimate these, the researcher starts by estimating a reduced-form VAR of the form

$$x_t = A(L)x_{t-1} + u_t, \quad (1)$$

where $A(L)$ contains enough lags to render the reduced-form innovations u_t approximately white noise. In the absence of measurement error or inappropriate truncation, this produces consistent estimates of the lag parameters in $A(L)$ and the covariance matrix Ω of the reduced-form innovations. The

⁴ Weak and unreliable transmission in the short run is, of course, perfectly consistent with monetary policy providing an effective long-run anchor for inflation, as Friedman recognized long ago.

researcher then imposes enough restrictions on the reduced form to identify the structural shocks to monetary policy. In our case, these take the common form of zero restrictions on elements of the square and invertible matrix B in

$$u_t = B\varepsilon_t. \quad (2)$$

Conditional on identification, the impulse responses (IRs) can be calculated as nonlinear functions of the estimated lag parameters and reduced-form shock covariances. The researcher computes these estimated IRs and, in a final step, bootstraps their standard errors and calculates t ratios for each impulse-response step. When a ‘true’ MTM is present in the data-generating process, the researcher should see impulse responses that are appropriately signed and shaped, of economically significant magnitudes, and estimated with reasonable statistical precision. We loop over multiple simulated datasets in order to study the population distribution of estimated impulse responses, the associated t -ratios, and the power of the t -ratio test in a wide variety of specific environments.

2.1 The DSGE model

The model we employ for our experiments is a canonical New Keynesian open-economy model that combines an IS curve, a New Keynesian Phillips curve, an interest-parity condition, and a Taylor Rule for monetary policy (e.g., Berg, Karam and Laxton 2006). There is no empirical consensus on the appropriate parameterization of such a model for LICs, but in choosing parameters we can draw on recent research by Berg, Portillo and Unsal (2010), who develop DSGEs with similar 4-equation structure for Kenya, Tanzania and Uganda. Our basic model, complete with parameter values, is:

IS equation:

$$\tilde{y}_t = 0.5 \cdot E_t[\tilde{y}_{t+1}] + 0.5 \cdot \tilde{y}_{t-1} - 0.2 \cdot [0.5 \cdot (i_t - E_t[\pi_{t+1}] - \bar{r}) + 0.5 \cdot \tilde{e}_t] + \varepsilon_t^y, \quad (3)$$

New Keynesian Phillips curve:

$$\pi_t = 0.5 \cdot E_t[\pi_{t+1}] + 0.5 \cdot \pi_{t-1} + 0.15 \cdot \tilde{y}_t - 0.15 \cdot \tilde{e}_t + \varepsilon_t^\pi, \quad (4)$$

Uncovered interest parity equation:

$$\tilde{e}_t = 0.5 \cdot E_t[\tilde{e}_{t+1}] + 0.5 \cdot \tilde{e}_{t-1} + 0.25 \cdot [i_t - E_t[\pi_{t+1}] - \bar{r}^*] + \varepsilon_t^e, \quad (5)$$

Taylor-type rule for monetary policy:

$$i_t = 0.5 \cdot (\bar{r} + 1.4 \cdot E_t[\pi_{t+1}] + 0.5 \cdot E_t[\tilde{y}_{t+1}]) + 0.5 \cdot i_{t-1} + \varepsilon_t^i, \quad (6)$$

Structural shocks:

$$\varepsilon_t \sim i.i.d. N(0, I_4) \quad (7)$$

Here E_t denotes an expectation conditional on information available at time t , while \bar{r} and \bar{r}^* denote the natural, or equilibrium, rate of interest in the domestic and world economy respectively. As explained below, we allow information sets to vary across equations, reflecting differences in the information available to agents. Many DSGEs are too simultaneous to allow monetary policy shocks to be identified via short-run restrictions (Canova 2007). To avoid well-known issues of inappropriate specification, we impose information restrictions directly on the DSGE in order to validate identification through short-run restrictions.⁵

In our model, monetary policy follows a Taylor-style rule, even though many LICs use the monetary base rather than a policy interest rate as the main operational instrument. We also omit a banking sector from the model, even though the nature of the credit channel may differ in LICs as compared to more advanced countries. These simplifications reflect our focus on aspects of the research environment, rather than on model structure. Our Monte Carlo approach can, of course, be applied to *any* structural model, a topic to which we return in the concluding section.

Our data-generating process is the solution to the DSGE model in terms of the endogenous state variables and the shock vector ε_t . This introduces a set of technical issues that are well understood in the DSGE and VAR literatures. First, for VAR-based methods to have a chance of uncovering the features of the MTM, the solution to the DSGE must be representable, at least approximately, as a finite-order VAR in observable variables. As discussed in detail in Li *et al.*, (2016) our model solutions have *exact* representations as finite-order VARs. Second, the monetary policy shocks must be identifiable through the imposition of conventional structural-VAR restrictions on this representation. We focus on short-run restrictions, because these remain the dominant approach to identification in the

⁵ Note also that equation (7) departs from the bulk of the DSGE literature by assuming *i.i.d.* shocks, in preference to the standard AR(1) structure: our version allows for distributed lag responses similar to those in the literature, but these are governed completely by the lags within the behavioural equations.

applied literature reviewed by Mishra *et al.* (2012). As discussed below, such restrictions work by limiting the contemporaneous interactions among the variables in the VAR.

2.2 Identifying Restrictions in the SVAR

The restrictions we impose at the estimation stage are typically motivated in the structural VAR literature by appealing to a structural simultaneous equations model of the form

$$B_0 x_t = B(L)x_{t-1} + \varepsilon_t. \quad (8)$$

The shocks ε_t are *i.i.d.* and mutually uncorrelated variables that can be normalized without loss of generality to have unit variances ($E[\varepsilon_t \varepsilon_t'] = I$).⁶ As long as B_0 is invertible, equation (8) implies the reduced-form VAR representation in equation (1), with $A(L) = B_0^{-1}B(L)$. The relationship between the structural and reduced-form innovations is then given by equation (2), with $B = B_0^{-1}$.

Within the class of short-run restrictions, the most common are those that impose a recursive structure on B_0 . Cholesky decompositions assume that the model is fully recursive, so that B_0 is lower triangular. As Christiano, Eichenbaum and Evans (1999) have shown, however, if the focus is on the impulse responses just to monetary policy shocks, these can be recovered from the reduced-form VAR under the considerably weaker condition that the system be *contemporaneously block-lower-triangular, with the interest rate occupying its own diagonal block*. We refer to any system that can be ordered into two or more block-recursive segments, with the interest rate occupying its own diagonal block, as *CEE-recursive*. When a structural VAR model is CEE-recursive, the impulse responses to monetary policy shocks can be recovered from the reduced-form VAR even if the remaining impulse responses cannot.⁷

A glance at equations (3) – (7), however, confirms that the solution to our DSGE will not, in fact, exhibit the block-recursive property under full information. Instead, it will tend to be highly simultaneous. This is in part because monetary policy is assumed to affect all endogenous variables contemporaneously, so the interest rate does not occupy its own diagonal block. However, it also reflects the role of expectation variables, since any endogenous variable that is in the information set of

⁶ A one-unit shock to ε_t^i is then equivalent to one standard deviation of the structural shock to monetary policy.

⁷ For the same reason, the ordering of the variables within each of the recursively prior and posterior blocks is irrelevant to obtaining the responses to interest-rate shocks (Christiano, Eichenbaum and Evans 1999).

a particular class of agents will contemporaneously affect all of the endogenous variables that are influenced by the forecasts formulated by those agents. Since all of the equations in our model contain such expectation variables, under full information *all* of the model's endogenous variables would tend to appear in every equation.

We therefore have to impose additional restrictions on our DSGE in order to produce a data-generating process that is identifiable via short-run restrictions. We retain the simultaneity of the structural model and obtain exclusion restrictions through assumptions about the information sets available to the private sector and the central bank. Specifically, we place the interest rate first in the CEE block-recursive ordering so that it affects all other variables contemporaneously. What this means is that the private sector has full information, but the central bank can only observe the endogenous variables with a lag. This structure can be rationalized as follows: the central bank (strictly its monetary policy committee) sets its policy interest rate at the beginning of each period. Non-policy shocks hit the system during the period, and are observed both by the private sector and the monetary policy committee. But while the private sector can react within the period, the central bank cannot do so until the beginning of the next period (i.e. at the next MPC meeting). The central bank is, therefore, setting its policy on the basis of $t-1$ information, while the private sector is behaving on the basis of time- t information. We suggest this structure has plausibility in a LIC context, where the central bank has less access to timely information on the state of the economy than in higher-income countries.

As explained in the Appendix to Li *et al.*, (2016), we solve the model using a version of the method of undetermined coefficients. The solutions have exact VAR(1) representations with CEE-recursive B matrices.

3. Facts on the ground

3.1 Strong VAR performance under baseline conditions

Figure 1 sets the stage for the subsequent discussion. Here we report the performance of a validly identified CEE-recursive VAR using 40 years of quarterly data measured without error. The notional

researcher estimates the VAR with four lags.⁸ To focus on parameters of interest we report only the impulse response functions for the monetary policy shock, defined as a one unit (i.e. one hundred basis point) positive shock to the policy interest rate. Since the components of Figures like 1a and 1b will appear throughout the paper, we begin by describing their content.

The researcher is trying to uncover the true, model-based impulse responses (IRFs), which appear as the bold lines marked with circles in Figure 1a. As shown in the figure, these IRFs display the conventional hump-shaped responses of the real exchange rate, inflation and output to a monetary contraction. On impact, a 100 basis point increase in the interest rate leads to a one percent contraction in inflation and a reduction in the output gap by around 0.7% of GDP, values that are broadly in line with Christiano *et al* (2005). To examine whether VAR methods can uncover these responses, we generate 1,000 data samples from our model, based on independent simulations of the DSGE solution, each generated by 160 quarters of independent draws on the shock vector ε_t . For each data sample, our researcher estimates a VAR and constructs IRs by imposing the CEE-recursive identifying restrictions. The empirical performance of these IRs is summarized by the three lighter lines in Figure 1a. These lines show the 5th, 50th, and 95th percentiles of the population distribution of simulated point estimates for the impulse responses (with percentiles computed separately for each impulse-response step).⁹

For each of the 1,000 simulations, the researcher computes the VAR coefficients and standard errors using conventional Bayesian estimation methods. Figure 1b shows the probability of rejecting the (incorrect) null hypothesis of a zero impulse-response coefficient at each step. We assume that the researcher treats the t ratios as asymptotically normal and applies the relatively undemanding hurdle of 10 percent significance.

Figures 1a and 1b establish that with appropriate identification and in the presence of ample and high-quality data, the VAR methodology does very well at uncovering strong monetary transmission when it is present. The estimated impulse responses for output and inflation show only a trivial degree of small-

⁸ Given the true model's VAR(1) structure, assuming the researcher estimates a VAR(4) model leads to some loss of efficiency. In principle it would be straightforward to embed a data-driven choice of lag length, but we leave this for future work. We can show, however, that the loss of efficiency is not large (results available on request).

⁹ The line corresponding to the median response is virtually invisible in the figure because it overlaps almost exactly with the true response in each case.

sample attenuation at the median, and for the first few quarters fully 90 percent or more of the point estimates lie on the correct side of zero.

The researcher's own inference will of course frequently be less confident than suggested by Figure 1b, because the researcher has only one data sample. Figure 1c reflects this by showing the full distribution of t ratios across the 1,000 runs. The structure of the exercise suggests that the width of bootstrapped confidence intervals for the IR coefficients will not be far from that implied by the population distribution of impulse responses, and the comparison of Figures 1a and 1c bears this out. When one end of the population distribution of IRFs is close to zero in Figure 1a, roughly half of the t statistics reported in Figure 1c fail to reject the null.¹⁰

The overall impression from Figure 1, then, gives weight to a pure facts on the ground interpretation of the missing MTM puzzle. Given a valid block-recursive identification scheme and sufficient data, structural VARs identified through short-run restrictions do well at uncovering the true MTM in the strict sense that the median estimated IRs track those generated by the true MTM very closely.

3.2 Detecting weak transmission

Before we investigate the 'limitations of methods' diagnosis, we briefly consider the facts on the ground and show that the power of the SVAR method to reject the null of no transmission deteriorates significantly when true transmission is present, but weak. To do so we consider the impact of small true effects driven by model parameters.¹¹ There are many parameters that may differ substantially between LIC and higher-income applications. The private-sector block incorporates both an interest-rate channel that operates through the IS curve and an exchange-rate channel that branches off from the interest parity condition to the IS and Phillips curves, while the monetary policy rule incorporates feedback from both inflation and the GDP gap along with a parameter that governs the degree of interest-rate smoothing. Based on Mishra *et al.* (2012, 2013), we focus here on two simple experiments.

¹⁰ The inference plots are very consistent with this effect. This suggests that the bootstrapped standard errors calculated by the researcher on each run of data tend to closely approximate the spread of the population distribution.

¹¹ Li *et al.* (2016) also examine the case where the source of weak transmission emerges from a setting in which the central bank has an informational advantage over the private sector which, for a constant model parameterization will attenuate the MTM.

In the first, reported in Figure 2, we scale down the transmission elasticities in the IS and Phillips curves by a uniform 75 percent relative to the baseline model. In other words, the parameter multiplying the policy interest rate in the IS curve, equation (3), is reduced from 0.5 to 0.125, while the parameter on exchange rate in the Phillips curve, equation (5), is reduced from 0.15 to 0.0375. In the second experiment we allow for a weaker pass-through from the central bank's policy rate to long rates and spending. Mishra, Montiel, and Spilimbergo (2012) find that the correlation between short-term interest rates and lending rates tends to become progressively weaker at lower levels of development. While equations (3) – (6) do not directly incorporate a lending channel, the IS curve and interest-parity condition can be solved forward to express the levels of the current GDP gap and real exchange rate gap as functions of current and expected future short-term interest rates. As emphasized by Woodford (2001), monetary policy shocks affect the 'tilt' of the spending and real exchange rate gaps via the short-term interest rate, but they alter the equilibrium *level* of these variables only to the degree that they change current long-term rates. The pass-through of short rates to long rates is in turn governed both by the parameters of the private sector block and by the degree of interest-rate smoothing implemented by the central bank. To investigate the role of the latter, in Figure 3, we leave the transmission elasticities unchanged and reduce the smoothing parameter in the monetary policy rule by 75 percent.

The results of these modifications are shown in Figures 2 and 3. As before, the lighter lines show the 5th, 50th, and 95th percentiles of the population distribution of simulated point estimates. For reference, the true IRs with the original model parameters are retained (the heavy line with diamonds). In both cases the change in the parameters weakens the 'true' transmission to output and to inflation to a rather similar degree. Notably, however, the estimated impulse responses continue to show very strong fidelity at the median, with the median IRs corresponding very closely to the true ones. Moreover, there is also no discernible impact on the spread of estimated impulse responses. To a first approximation, therefore, the impact of weak transmission elasticities operates exclusively through the impact of small true effect sizes on the power of *t*-ratio tests against the null hypothesis of no effect. That impact is substantial, however, with the scope for confident inference cut roughly in half (bottom row of Figure 2).¹²

¹² The exchange-rate channel is quantitatively important in our model. In simulations not reported here, we show that if it is only the interest-rate elasticity that differs between LIC and non-LIC applications, and not the exchange rate elasticity, the deterioration in the inference about the MTM is mild.

The overall impression from Figures 1 to 3, then, gives weight to a pure ‘facts on the ground’ interpretation of the ‘missing MTM’ puzzle. Given a valid block-recursive identification scheme and sufficient data, structural VARs identified through short-run restrictions do well at uncovering the true MTM—whether it is strong or weak—in the strict sense that the median estimated IRs track those generated by the true MTM very closely. Therefore, with appropriate identification and sufficient data, the SVAR approach is able to correctly discriminate between strong and weak monetary transmission. At the same time, however, we have found that the population dispersion in estimated IRs is not substantially affected by the strength of the underlying “true” MTM. Therefore, the weaker the true MTM, the harder it is for the data to reject the null of no response—i.e., the weaker the power of tests of the null hypothesis that the MTM is entirely missing. Even where a (plausibly weak) MTM exists, the researcher armed with only one data set, even a pristine one that is 40 years long, may well conclude that it is missing.

4. Limitations of the method

4.1 Small sample sizes generate low precision

In the remainder of the paper we turn to the challenges posed for the banks and researchers using SVARs by the specific research environment that characterizes LICs. We begin with short data samples, which are surely among the most daunting if perhaps mundane constraints in the LIC research environment. Except as noted, we assume the underlying strong transmission mechanism of the baseline, presented in section 3.1.

Structural economic reforms (e.g., financial liberalization) and changes in the monetary policy regime (e.g., a move to a flexible exchange rate, a very different smoothing parameter for interest rates, or a much different role for monetary aggregates) are often of sufficiently recent vintage that the researcher cannot expect that the current data-generating process has been in place for very long. Data-collection limitations also undermine the availability of long data samples in LICs; as noted in the Introduction, quarterly data on the real economy, for example, may be unavailable well through the 1990s or even more recently.

Figure 4 shows an experiment that is identical to the baseline with the exception that the researcher has 10 years rather than 40 years of quarterly data. As in the case of more abundant data, the basic shape of the true impulse responses is strongly reproduced, although the median estimated impulse response is discernibly attenuated over the first half-dozen periods relative to Figure 1. More importantly, the reduction in sample size produces a substantial widening in the population distribution of IRF point estimates. Coupled with the attenuation, this produces a sharp deterioration in the scope for confident inference about the MTM. As shown in lower half of Figure 2, the power of statistical tests to reject the null of no monetary policy effect based on bootstrapped standard errors is 30 to 50 percent smaller with 10 years of data (solid line) than with 40 years of data (bold line with dots, which reproduces the lower half of Figure 1).

The upshot is that estimation of the strength of the MTM in LICs with VAR methods is likely to produce wide confidence bands around estimated IRs when the available span of data is limited. In combination with the somewhat attenuated point estimates of the strength of policy transmission, judgments about unreliability based on the width of confidence bands around estimated IRs are unwarranted when data samples are as short as they often tend to be in LIC applications. When data samples are short, wide confidence bands around estimated IRs should be interpreted as uncertainty about the estimates rather than as confirmation of on-the-ground instability or absence of policy effects.

4.2 Volatility does not undermine VAR-based inference, but measurement error does

Macroeconomic variables display greater volatility at business-cycle frequencies in LICs than in higher-income countries. Variation is of course crucial to uncovering the MTM in an SVAR context: even a perfectly-specified SVAR will fail if the variance of monetary policy shocks is small enough relative to that of other shocks. It is not necessary to illustrate this case. Here we focus instead on two drivers of data variability that clearly differ systematically between developing and advanced countries. We show that *true economic volatility* and *measurement error* have sharply different impacts on VAR-based inference about the MTM. To a first approximation, economic volatility leaves VAR-based inference unchanged, while measurement error rapidly undermines it.

The intuition for the neutral effect of volatility is straightforward. It is well understood in the context of the stochastic-regressor model $y_t = \beta x_t + \varepsilon_t$, that the precision of the OLS estimator in any finite sample depends approximately on the relative variances of x_t and ε_t . These variances have the familiar

effects: volatility in the disturbance term undermines inference about β , while volatility in the independent variable enhances inference.

A similar expression characterizes the OLS estimator in the AR(1) model $x_t = \alpha x_{t-1} + \varepsilon_t$, viewed here as the simplest possible stationary VAR (ε_t is white noise with variance σ_ε^2 and finite higher-order moments, and $|\alpha| < 1$):

$$\sqrt{T}(a_T - \alpha) \xrightarrow{L} N(0, \sigma_\varepsilon^2 / \text{Var}[x_{t-1}]) = N(0, 1 - \alpha^2). \quad (9)$$

In contrast to the conventional stochastic-regressor case, therefore, the variance of x_t in a VAR model is a function only of the volatility of the shocks. Thus any factor that uniformly scales up the variances of the shocks in a VAR has little effect on inference about the VAR coefficients, because the sampling variances of the estimated coefficients are simultaneously pushed upwards by the variances of the shocks, σ_ε^2 , and downwards by the variances of the lagged variables in the VAR, $\text{Var}[x_{t-1}]$. These effects cancel, leaving the finite-sample variances of VAR coefficients approximately invariant to the variances of the structural shocks. This property carries over to the variances of the IRF coefficients, because the latter are continuous functions of the VAR coefficients.

Measurement error, in contrast, unambiguously undermines the accuracy of VAR estimates. Figure 5 quantifies this effect for the simple case of classical measurement error that affects the researcher but not the agents in the model.¹³ For this and subsequent experiments, we return to the section 3.1 baseline with a full sample of 40 years of quarterly data. We induce classical measurement error in the GDP gap and inflation, by adding white noise measurement error to the model-generated true values of these variables, with a variance that is equal to 20 percent of the variance of the structural shocks.¹⁴ As shown in the top row, the point estimates show attenuation towards zero in the first two quarters (the

¹³ To contemplate measurement error is to open a can of worms when considering not just the VAR but the underlying economy, as captured by the DSGE model. To assess the impact of measurement error more fully, a natural approach might be to confront the agents in the model with the same measurement error faced by the econometrician. This seems likely both to weaken the true MTM as well as to undermine the ability of VAR methods to discover it, including by inducing infinite lags in the DSGE solution and therefore exposing the VAR results to truncation bias.

¹⁴ Reliable estimates on the scale of measurement error in low-income countries are rare but there is reason to believe that our value of 20 percent is on the conservative end of the spectrum (see, for example, Jerven, 2013, and Ley and Misch, 2014).

estimated effects are slightly more than half of the true effects) for the GDP gap, but little impact in later quarters. Attenuation effects are larger and more persistent on inflation, with such effects persisting after the first two quarters. In both cases, there is little effect on the dispersion of estimated IRs but the power of t -ratio tests deteriorates badly during the quarters when attenuation bias is most pronounced.

4.3 High-frequency supply shocks obscure transmission to output, but not to other variables

We have been assuming that the GDP gap is observable. In reality, the gap must be inferred by the researcher—even when actual output is measured without error—by developing an empirical proxy for potential GDP. In this section we show that the high-frequency real-side shocks that are characteristic of LICs exacerbate measurement error in the GDP gap. When we follow typical practice and use a one-sided filter to measure the gap, this effect substantially weakens inference about the real side of the monetary transmission mechanism. The transmission to inflation remains surprisingly robust, however.

The GDP gap is typically measured in empirical applications by assuming that potential GDP follows a slow-moving trend. This trend is then either extracted from the actual series by differencing or filtering (leaving a stationary gap variable that can be used in the VAR) or controlled for within the VAR by expanding the set of variables to include proxies for aggregate supply (e.g., a deterministic trend or the economy-wide capital stock and labour force). As emphasized in the literature on the New Keynesian Phillips curve, however, the concept of potential GDP that matters for inflation dynamics is the natural or ‘flex-price equilibrium’ level of GDP. Natural GDP is a function of slow-moving processes like factor accumulation and technological change, but it also depends on transitory supply-side shocks that can affect output in the absence of sticky prices. Such shocks may play a greater role in determining natural GDP in LICs than in high-income countries. Droughts, for example, are likely to have larger effects on natural GDP in countries with larger rain-fed agricultural sectors. Supply-side shocks may also be more important in countries with less diversified non-agricultural sectors. If this is indeed the case, then the practice of proxying potential GDP with a slow-moving trend will induce more serious measurement errors in LICs than elsewhere.

To formalize this idea, we rely on the fact that the DSGE specifies the stationary interactions of macroeconomic variables once stochastic trends have been removed. This leaves us free to construct

the path of actual output as the sum of our model-generated GDP gap (now denoted \tilde{y}_t^M) and a new, exogenous stochastic process for natural GDP, y_t^n :

$$y_t = y_t^n + \tilde{y}_t^M. \quad (10)$$

In each of our 1,000 simulation runs, therefore, we construct output from the two components on the right-hand side of equation (10). The researcher then receives the vector $[i_t, \tilde{e}_t, y_t, \pi_t]'$, which includes observable output y_t rather than the unobservable gap \tilde{y}_t^M , and then approximates the gap by applying a one-sided Hodrick-Prescott (HP) filter to actual GDP, with the standard quarterly smoothing parameter of 1600. The VAR is then estimated on $[i_t, \tilde{e}_t, \tilde{y}_t^R, \pi_t]'$, where the researcher's GDP gap, \tilde{y}_t^R , is the HP-cycle in actual GDP.

To implement this approach, we require a measure of natural GDP. We assume that natural GDP consists of two components: a stochastic trend y_t^{nT} that follows an integrated random walk with deterministic drift, and a stationary component y_t^{nS} that is present only in a LIC environment. Thus

$$y_t^n = y_t^{nT} + y_t^{nS}, \quad \Delta y_t^{nT} = \bar{g} + \varepsilon_t^{nT}, \quad y_t^{nS} = \varepsilon_t^{nS}, \quad (11)$$

where $\bar{g} = 0.02$, ε_t^{nT} is a white noise stochastic shock and ε_t^{nS} a mutually uncorrelated stationary AR(1) process, and where the variance of ε_t^{nS} is zero outside of a LIC environment. The researcher in these experiments estimates \tilde{y}_t^M as the HP cycle in actual GDP, y_t . This induces some degree of measurement error in the GDP gap even when there is no high-frequency component of natural GDP.¹⁵ To benchmark this 'unavoidable' deterioration, we first show results for a hypothetical case where the GDP gap is unobservable but estimated—under optimal conditions for a one-sided filter—by the researcher but where we assume that the true variance of ε_t^{nS} is zero. This is shown in Figure 6 which suggests that the impact on inference is confined to the GDP responses, which show modest initial attenuation towards zero. The t ratios are not very strongly affected. These results suggest that despite unavoidable measurement error, considerable scope remains for confident and qualitatively accurate inference

¹⁵ The HP filter yields $y_t = y_t^{HPT} + \tilde{y}_t^R$ where HPT denotes the Hodrick-Prescott trend. Combining this with the equation used to construct y_t , we can see that $\tilde{y}_t^R = \tilde{y}_t^M + [y_t^n - y_t^{HPT}]$. The measurement error in \tilde{y}_t^R is therefore the difference between two highly persistent series. Estimating the gap in this way tends to produce a measurement error that is persistent, even in the 'best' of circumstances, in which actual GDP is measured without error.

about the transmission of monetary policy to output, at least for the impact effect and the early steps of the process. Inference about inflation is even less affected.

This conclusion, however, depends strongly on the absence of transitory shocks to potential GDP. Figure 7 illustrates the case where there are transitory supply shocks with variance equal to that of the true model-based GDP gap.¹⁶ The researcher proceeds as before, but since natural GDP now includes a high-frequency component, the HP filter over-smooths the GDP series and thereby exacerbates the measurement error in the GDP gap. Inference deteriorates correspondingly, so that Figure 7 portrays a transmission mechanism that is counterfactually both weak and unreliable with respect to the (true) GDP gap, slightly more so if there is persistence in the measurement error.

4.4 Assembling the pieces

We have so far considered a set of LIC-specific features of the data and environment that may weaken the power of SVARs to detect transmission that exists, taking those features one at a time. Of course, these problems do not come in isolation. In this section, we get a feel of how these various factors work together. First, we combine several of the most plausible data/environmental factors to see what happens when the transmission mechanism is strong (as in Figure 1) but data series are only 10 years long, there is 20 percent measurement error, and there are i.i.d shocks to potential output that complicate estimation of the output gap. As is shown in Figure 8, in this case the SVAR estimates are substantially biased downwards, with the median estimated impulse response of output to an interest rate shock at about half of its true value. Indeed, for inflation the median estimated impact effect of a monetary policy contraction is zero. And uncertainty is high; the SVAR has very little power (only about one third of what we had in the baseline) to extract the true MTM.

4.5 Identification may be especially tricky in LICs

Up to this point, we have maintained the assumption of correct identification. The perils of incorrect SVAR identification are of course neither surprising nor LIC-specific. As emphasized by Canova and Pina (2005) and many others, impulse responses can be highly misleading when monetary policy shocks are (mis-)identified via inappropriate short-run restrictions. We have avoided this obvious pitfall by

¹⁶ In Figure 7 we assume white noise shocks. Results are similar where they follow an AR(1) process with for example $\rho = 0.9$.

incorporating information restrictions directly into the data-generating process. In contrast to industrialized countries, however, where there is an extensive literature on the transmission mechanism, there tends to be no consensus on the nature of behavioral lags and information sets in the LIC environment. The identification of monetary policy shocks is therefore likely to prove far more difficult in this environment than in the better-understood environment of high-income countries.

Erroneous identification can occur in many ways. Here we focus on a case in which the researcher incorrectly assumes the central bank has an informational advantage over the private sector and therefore places the interest rate ‘too late’ in the block recursive structure of the model. This error might be a natural one for a researcher trained in the advanced-country literature. The outcome is shown in Figure 9. The solid line with dots in bold represents the true impulse responses (reproduced directly from Figure 1) while the light lines once again represent the 5th, 50th, and 95th percentiles of the estimated IRFs when the researcher mistakenly imposes the wrong ‘interest-rate-last’ identification scheme.

The result of this error is sufficient to produce impulse responses that are ‘weak and unreliable’ in the extreme: they are essentially zero, both economically and statistically. Not surprisingly, statistical tests based on bootstrapped standard errors for the estimated IRs will have essentially no power to reject the null of zero monetary policy effects in this case, even though the true effects are in fact extremely powerful. This suggests that incorrect identification of monetary policy shocks is itself a prime suspect in the case of the missing MTM.

5 Conclusions

In an effort to come to grips with the “missing MTM” in the empirical literature on LICs, we have reversed the standard dialogue in which SVAR-based impulse responses provide an empirical standard against which DSGEs or other theory-based models can be evaluated. We have instead used DSGE models as a data-generating process, in order to ask a question about the validity of SVAR-based impulse responses. If a strong MTM is present in the data, can standard SVAR methods uncover it?

No parametric method will do very well if it mis-specifies the data-generating process. This is the basis of Sims’ (1992) critique of structural econometric modelling, and as long as the data are generated by a stable but unknown data-generating process, this critique favours the use of as few structural

restrictions as possible to identify the MTM. SVARs identified via short-run restrictions are very widely used in the literature on LICs, and perhaps even more intensively there than elsewhere given the relative dearth of structural modelling in these countries. Within this class, we have focused on block-recursive VARs, which impose just enough recursive structure to identify the monetary policy impulse responses, while leaving the other responses potentially unidentified. An off-the-shelf DSGE model will not generate a solution with this property, but the otherwise-canonical DSGE we have employed in this paper is capable of doing so under mixed-information assumptions of the type often seen in the structural VAR literature.

When the VAR researcher imposes a valid identification scheme and has access to ample and high-quality data, the virtues of the VAR approach come through strongly. The LIC environment nonetheless poses a set of well-defined challenges to a strategy that ‘lets the data speak’: we investigate short samples arising from missing data or structural breaks, volatile data, measurement error, and high-frequency shocks to natural GDP.

We find a surfeit of plausible explanations for the ‘missing MTM’. Specifically, likely LIC-specific features of the estimation environment such as short time series, measurement error, and the need to estimate the output gap greatly weaken the power of the VAR methodology to reliably uncover the MTM. The small samples moderately attenuate the median point estimates. In addition, the resulting uncertainty is high. Each of these features reduces sharply the power of standard *t*-ratio tests against a null hypothesis of zero response, even in properly-identified structural VARs with a strong MTM. Together, they are devastating. A combination of plausibly short time series, a modicum of measurement error, and an estimated output gap with some high-frequency supply shocks almost entirely eliminates the power of an SVAR, even when the underlying MTM is strong.

This paper has thus demonstrated that even when identification is proper and the MTM strong, features of the environment such as short data series and measurement error will likely make it hard to reliably hear the data speaking with a SVAR approach. A potential extension would address one of the simplest questions posed by this analysis: why not use higher-frequency – e.g., monthly – data instead? There are two powerful reasons to do so in a LIC environment, and one potentially daunting constraint. The first reason is that ten years of monthly data yield 120 observations rather than 40. The sample information will of course increase much less than three-fold, because the data cover the same time-

period and the monthly model will display greater short-run persistence. But other things equal, we would expect sharper inference from the increase in sample size. In a structural VAR context, however, the more important reason for going to monthly data is that contemporaneous timing restrictions become more plausible. Block-recursive methods are the least demanding and therefore most *a priori* plausible approach for a researcher seeking to identify the MTM via recursive structure, but it strains credulity when applied to quarterly data. By underscoring the leverage of valid identification, our results suggest that the gains in small-sample bias from better identification in monthly data may be very substantial. A major constraint, however, is that measurement error is likely to be greater in monthly data, especially for measures of real activity. Our Monte Carlo approach is well suited to running this horse race, as a natural extension of the current work.

Altogether, all these VAR based methods seem worth pursuing, but most effectively so in tandem with other methods that help triangulate insights. Further country-specific VAR analyses that carefully tailor the approach to the particular country-specific institutional framework for monetary policy may permit progress in this respect. Even in the U.S., with its unusually long, stable data series and policy regimes, economists experimented for many years before reliably generating acceptable results, only eventually solving the “liquidity puzzle” that interest rates tended to rise in response to an increase in the money supply and the “price puzzle” that inflation seemed to rise after a shock that tightened monetary policy.

But this paper does suggest that the challenges to achieving similar success in LICs will be severe. Other empirical methods may thus be helpful. The case study approach adopted by Berg *et al* (2018) examines the implications of a large monetary policy shock identified through a narrative approach. Abuka *et al*. (2015) use loan-level data to assess the bank-lending channel in Uganda. Staying within the SVAR space, sign restrictions on the impulse responses to monetary and non-monetary shocks may provide a promising alternative to recursive identification (Canova 2007), although their robustness to LIC-type data challenges remains to be explored. Perhaps more radically, the imposition of more economic structure and use of Bayesian techniques is a natural way to make use of available data while accepting its scarcity.¹⁷ It seems clear, however, that despite the possible ways forward uncertainty about the MTM is likely to continue to face LIC central banks.

¹⁷ See for example Berg *et al*. (2010) and Peiris and Saxegaard (2007).

What implications follow from these conclusions? One view is that, in the face of this uncertainty, monetary policy should be passive. In particular, frameworks that do not require confident knowledge of the MTM may be more appropriate, notably fixed exchange rate regimes or perhaps textbook money targeting regimes that aim to keep money growth at a constant predetermined rate. This conclusion may be too extreme, however. Recent research has investigated the implications of real-time data limitations for the conduct of monetary policy.¹⁸ While this work suggests that policy activism should indeed be restrained in the presence of measurement error (to avoid instability arising from policy responses to what amounts to simple noise in the data), it does not imply that active policies should be avoided altogether. In other words, noisy data does not mean abandoning feedback rules, but rather tempering the degree of feedback embedded in such rules, and possibly incorporating greater inertia into the policy interest rate.

A broader perspective is appropriate when uncertainty applies not just to the quality of the data available to policymakers in real time, but also to the parameters of the underlying model of the economy. Such uncertainty is hardly unique to LICs (Friedman 1953) but is likely to be acute as these countries implement new policy frameworks, often in the face of rapid structural change or financial crises. The endogeneity of the MTM to the policy regime and financial environment suggests that learning-by-doing is a necessary counterpart to the conduct of monetary policy in LICs.¹⁹ This broader uncertainty may well justify caution, but it is not clear that it justifies passivity.²⁰ Effective monetary policy frameworks incorporate a critical element of ‘tatonnement’ in all countries, including LICs: assess the state of the economy and the outlook; adjust policy cautiously if it seems too tight or too loose; then wait for new information and repeat. For this, only a view about the sign of the effects of monetary policy is required. However, clearly further empirical work is a critical component of this learning process. We hope that the results of this paper will help guide the agenda.

¹⁸ There is a broad recent literature on the implications of data imperfections for monetary policy design. Orphanides (2003), Aoki (2003), Taylor and Williams (2010), and Boehm and House (2014), for example, study the consequences of measurement error in observable real-time macroeconomic data such as inflation and real output for optimal Taylor-type monetary policy rules. Gust *et al.* (2015) focus on the implications for policy of imperfect knowledge about the (unobservable) natural real interest rate in the vicinity of the zero lower bound.

¹⁹ See Batini and Laxton (2007) for the case of inflation targeting in emerging markets. IMF (2015) discusses some of these issues for the case of low and lower-middle-income countries.

²⁰ The recent experience with quantitative easing in advanced countries is a (distant) case in point. Furthermore, uncertainty about the transmission mechanism is one of the main reasons that central banks in advanced countries use simple policy rules (like the Taylor rule) as benchmarks in determining the course of monetary policy (see Levin, 2014).

References

Abuka, Charles, Ronnie K. Alinda, Camelia Minoiu, Jose-Luis Peydro and Andrea F. Presbitero (2015), "Monetary Policy in a Developing Country: Loan Applications and Real Effects", *IMF Working Paper* 15/270.

African Development Bank (2013) "Situational Analysis of the Reliability of Economic Statistics in Africa: Special Focus on GDP measurement." Tunis: African Development Bank.

Aoki, K. (2003) "On the Optimal Monetary Policy Response to Noisy Indicators," *Journal of Monetary Economics*, 50(3), 501-523.

Batini, Nicoletta and Douglas Laxton, 2007, "Under What Conditions Can Inflation Targeting be Adopted? The Experience of Emerging Markets," Central Banking, Analysis, and Economic Policies Book Series 11: pp. 467-506.

Berg, Andrew, Rafael Portillo and D. Filiz Unsal (2010) "On the Optimal Adherence to Money Targets in a New-Keynesian Framework: An Application to Low-Income Countries" *IMF Working Paper* 10/134.

Berg, Andrew, Luisa Charry, Rafael A. Portillo, and Jan Vlcek (2018) "The Monetary Transmission Mechanism in the Tropics: Lessons from a Dramatic Effect" in Berg, A. and R.Portillo (eds) *Monetary Policy in Sub-Saharan Africa* Oxford: Oxford University Press..

Berg, Andrew, Philippe D. Karam and Douglas Laxton (2006) "A Practical Model-Based Approach to Monetary Policy Analysis" *IMF Working Paper* 06/80.

Boehm, C. E., & House, C. L. (2014), "Optimal Taylor Rules in New Keynesian Models," National Bureau of Economic Research Working Paper No. w20237.

Canova, Fabio (2007) *Methods for Applied Macroeconomic Research* (Princeton: Princeton University Press)

Canova, Fabio and Joaquim Pires Pina (2005) "What VAR Tell Us About DSGE Model?" in C. Diebolt and C. Kyrtsov, eds, *New Trends in Macroeconomics* (Berlin: Springer): 89-123

Charalambos Tsangarides (2010) "Monetary Policy Transmission in Mauritius using a VAR analysis", *IMF Working Paper* 10/36

Cheng, Kevin C. (2006) "A VAR Analysis of Kenya's Monetary Transmission Mechanism: How Does the Central Bank's REPO Rate Affect the Economy?" *IMF Working Paper* 06/300.

Christiano, Lawrence J. and Martin Eichenbaum (2007) "Assessing Structural VARs" in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds, *NBER Macroeconomics Annual 2006*, Volume 21 (Cambridge, MA: MIT Press)

Christiano, Lawrence J. Martin Eichenbaum, and Charles L. Evans (1999) "Monetary Policy Shocks: What Have We Learned?" In John B. Taylor and Michael Woodford, eds, *Handbook of Macroeconomics, Vol 1A* (Amsterdam: North Holland)

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005) "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy" *Journal of Political Economy* 113(1): February: 1-45

Friedman, M. (1953) *Essays in Positive Economics* Chicago: University of Chicago Press.

Gust, C. J., Johannsen, B. K., & Lopez-Salido, D. (2015), "Monetary Policy, Incomplete Information, and the Zero Lower Bound," Finance and Economics Discussion Series, Divisions of Research & Statistics and Monetary Affairs Federal Reserve Board, Working Paper No. 2015-099.

IMF, 2015, *Evolving Monetary Policy Frameworks in Low-Income and Other Developing Countries*, Washington: International Monetary Fund.

Jerven, Morten (2013) *Poor Numbers: How We are Misled by African Development Statistics and What to do about it*. Ithaca, N.Y. Cornell University Press.

Levin, Andrew T., 2014, "The Design and Communication of Systematic Monetary Policy Strategies," *Journal of Economic Dynamics and Control*, Vol 49, pp. 52-69.

Ley, Eduardo. and Florian Misch (2014) "Output data revisions in Low-Income Countries", conference presentation, IMF-DFID conference on *Macroeconomic Challenges Facing Low-Income Countries: New Perspective* Washington DC (Jan 2014).

Li, Bin Grace, Stephen O'Connell, Christopher Adam, Andrew Berg and Peter Montiel (2016) "VAR meets DSGE: Uncovering the Monetary Transmission Mechanism in Low-Income Countries" *IMF Working Paper* 16/90.

Mbowe, Wilfred (2012) "The Bank Lending Channel of Monetary Policy Transmission: Dynamic Bank-Level Panel Data Analysis on Tanzania" Dar es Salaam: Bank of Tanzania.

Mishra, Prachi and Peter Montiel (2013) "How Effective is Monetary Transmission in Low-Income Countries? A Survey of the Empirical Evidence" *Economic Systems*.

Misati, Roseline Nyakerario, Esman Morekwa Nyamongo, Lucas Kamau Njoroge and Sheila Kaminchia (2012) "Feasibility of inflation targeting in an emerging market: evidence from Kenya" *Journal of Financial Economic Policy* 4(2): 146-159

Mishra, Prachi, Peter Montiel, Peter Pedroni, and Antonio Spilimbergo (2014), "Monetary Policy and Bank Lending Rates in Low-Income Countries: Heterogeneous Panel Estimates," *Journal of Development Economics*.

Mishra, Prachi, Peter J. Montiel, and Antonio Spilimbergo (2012) "Monetary Transmission in Low-Income Countries: Effectiveness and Policy Implications," *IMF Economic Review* 60 (2): 270-302.

Orphanides, A. (2003), "Monetary Policy Evaluation with Noisy Information," *Journal of Monetary economics*, 50(3), 605-631.

Peiris, Shanaka J. and Magnus Saxegaard (2007) "An Estimated DSGE Model for Monetary Policy Analysis in Low-Income Countries," *IMF Working Paper* 07/282.

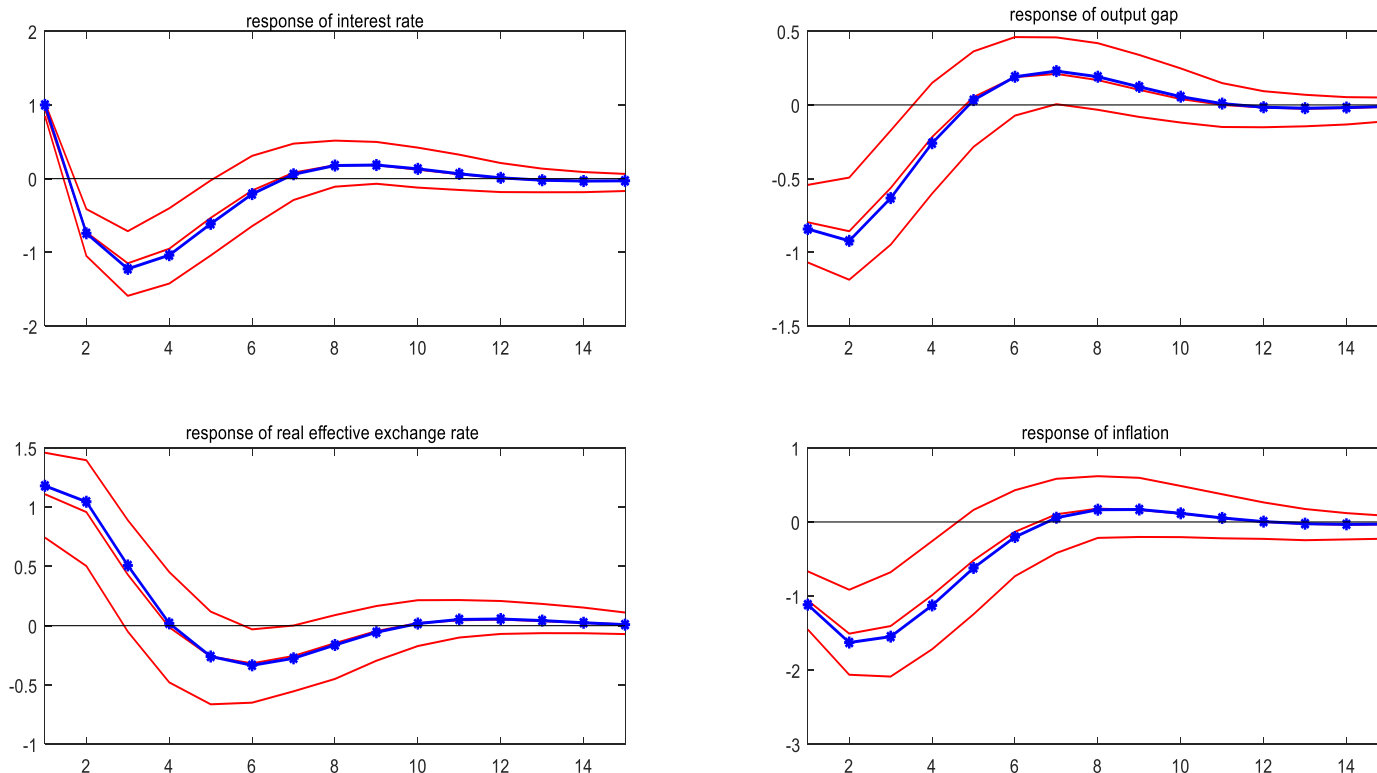
Sims, Christopher, 1992, "Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy," *European Economic Review* 36, No. 5: pp. 975-1000.

Taylor, J. B., & Williams, J. C. (2010). Simple and robust rules for monetary policy," In *Handbook of Monetary Economics*, Vol. 3 {Elsevier), 829-859.

Woodford, Michael (2001) "Monetary Policy in the Information Economy" in *Economic Policy for the Information Economy: Proceedings of the Federal Reserve Bank of Kansas City Annual Symposium Conference* Jackson Hole, Wyoming, August 30 – September 1: 297-370.

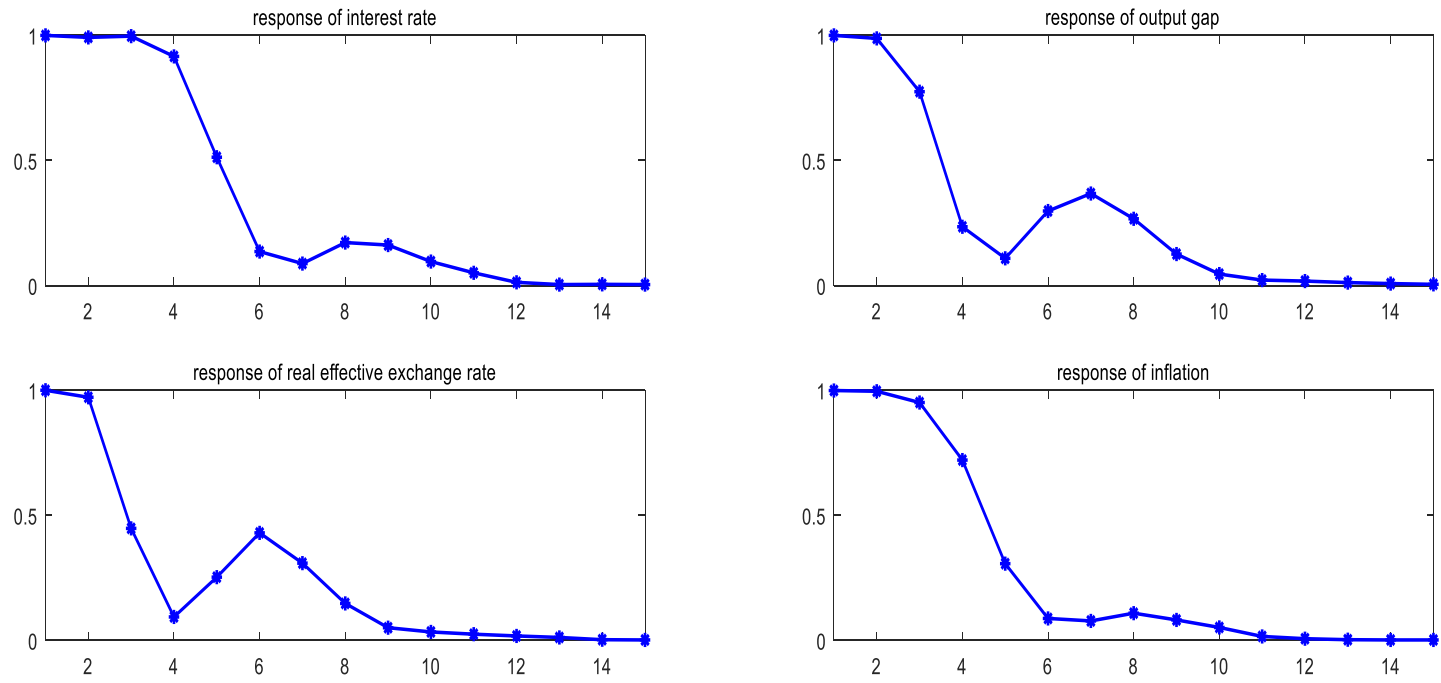
Figures

Figure 1(a) Impulse Responses to Monetary Policy Shock: Baseline



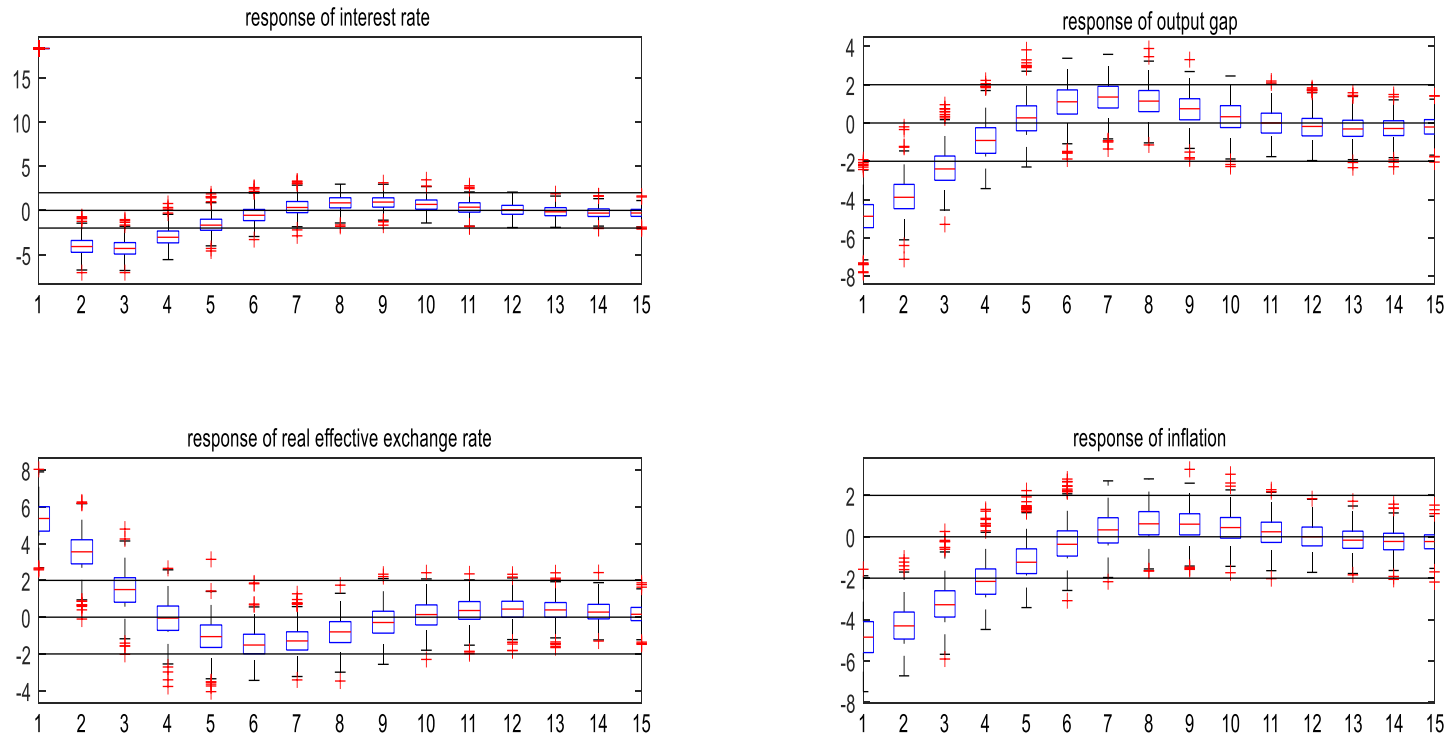
Notes: The blue-coloured lines with circles indicate the 'true' impulse response functions (IRFs) generated by the DSGE model. The three red lines correspond to the 5th, 50th and 95th percentiles of the distribution of estimated IRFs (see text). The percentiles are computed separately for each impulse-response step, from SVAR models estimated on 1,000 independent simulations of the DSGE model, each generated by 160 quarters of independent draws on the shock vector ε_t .

Figure 1(b) Power Functions for Monetary Policy Shock: Baseline



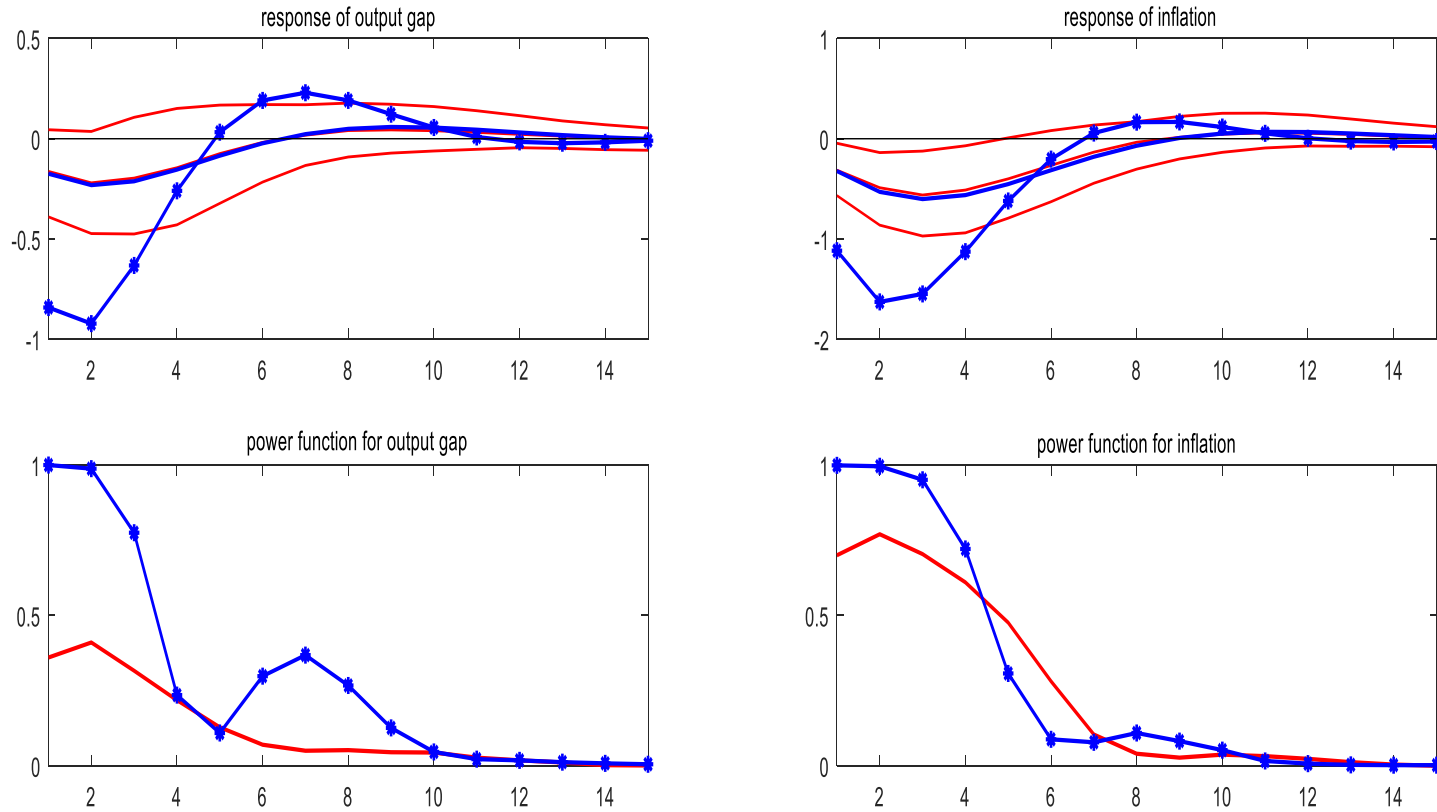
Notes: Lines indicate the probability of rejecting the (incorrect) null hypothesis of a zero impulse-response coefficient at each step based on SVAR coefficients and standard errors computed from each of the 1,000 simulations using standard Bayesian VAR estimation methods. Rejection of the null is defined against an asymptotically normal t ratio with 10 percent size.

Figure 1(c) T-stats for Monetary Policy Shocks: Baseline



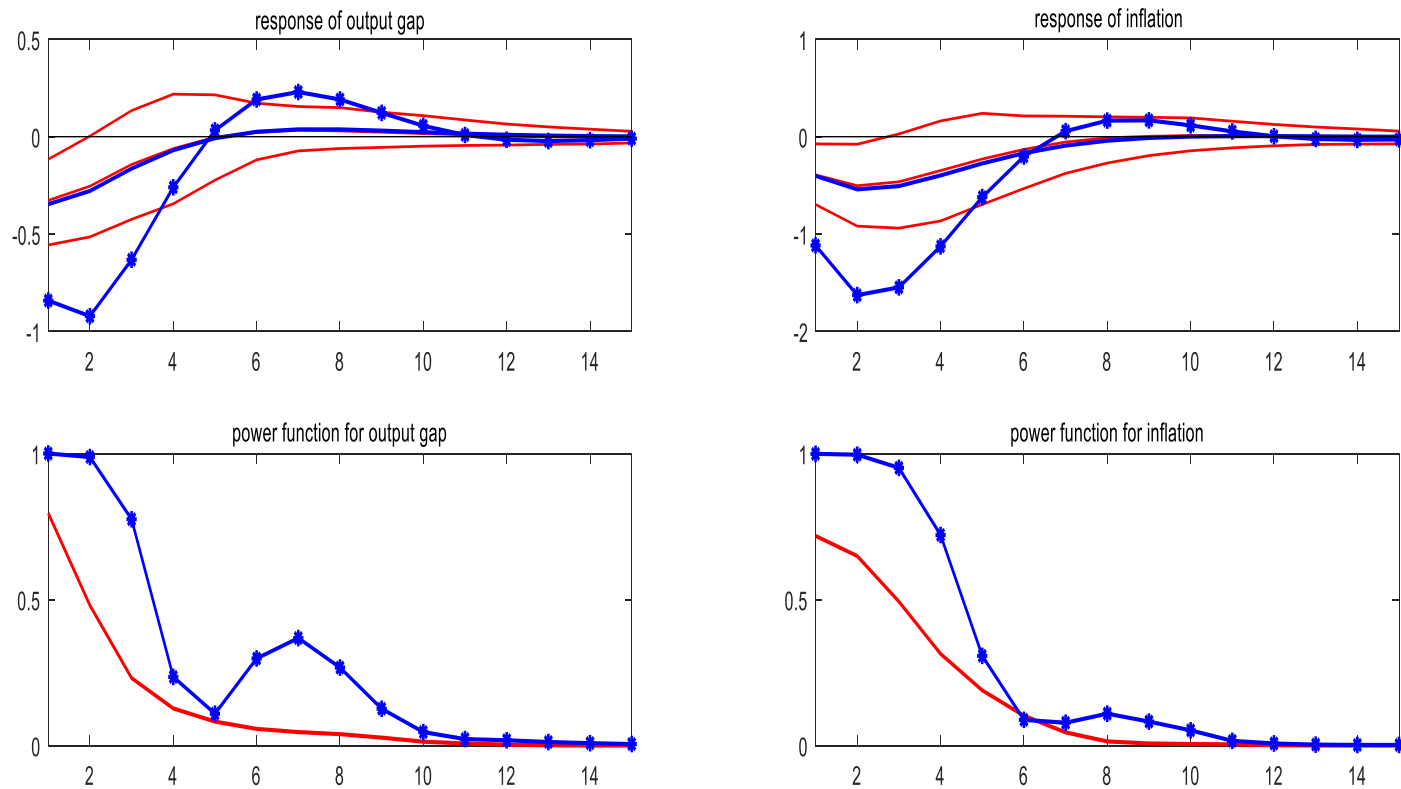
Notes: Each panel reports box-and-whisker plots to summarise the full distribution of t ratios across the 1,000 simulation runs.

Figure 2: Weak Transmission (elasticities in IS and PC scaled down by 75%)



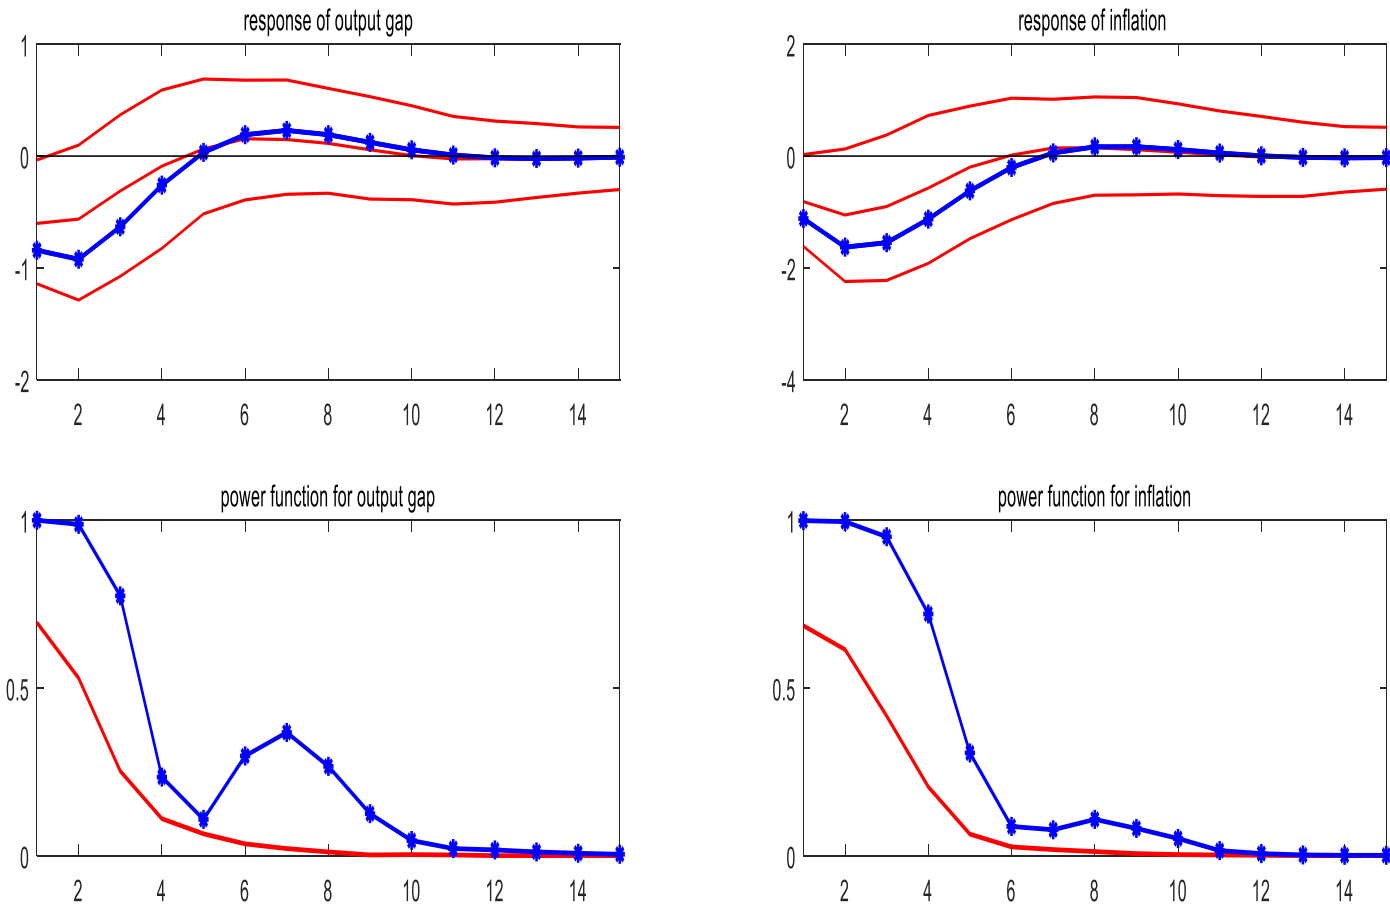
Notes: The blue-coloured lines with circles in the top two panels indicate the original 'true' impulse response functions (IRFs) generated by the DSGE model; the solid blue lines indicate the true model-based IRF with scaled down transmission elasticities. In the lower panels, the solid red lines indicate the corresponding empirical power functions, with the blue lines with dots reproducing the corresponding plots from Figure 1a.

Figure 3: Weak Transmission (smoothing parameter scaled down by 75%)



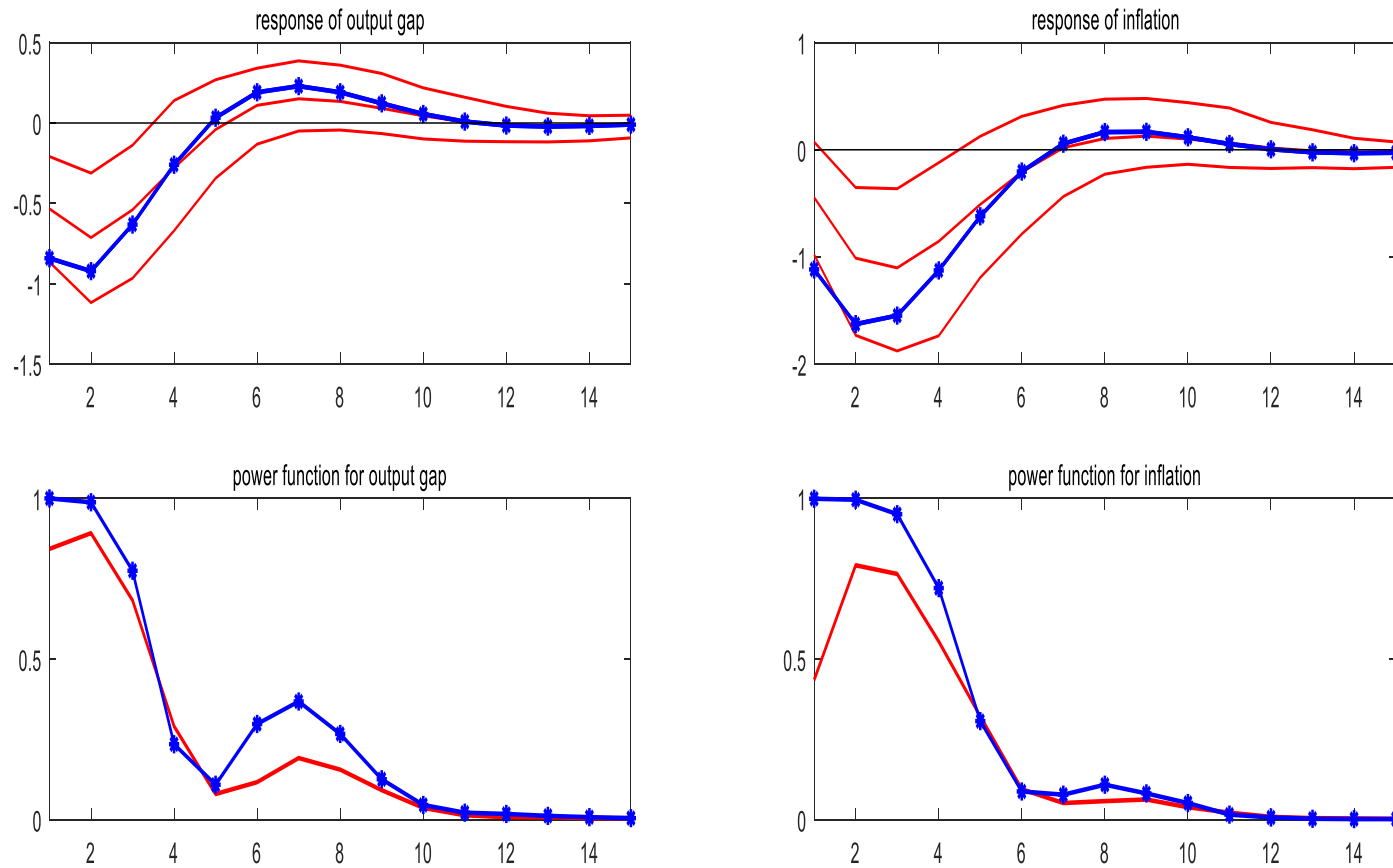
Notes: See text and notes to Figure 2.

Figure 4: Small Sample (sample reduced from 40 to 10 years of quarterly data)



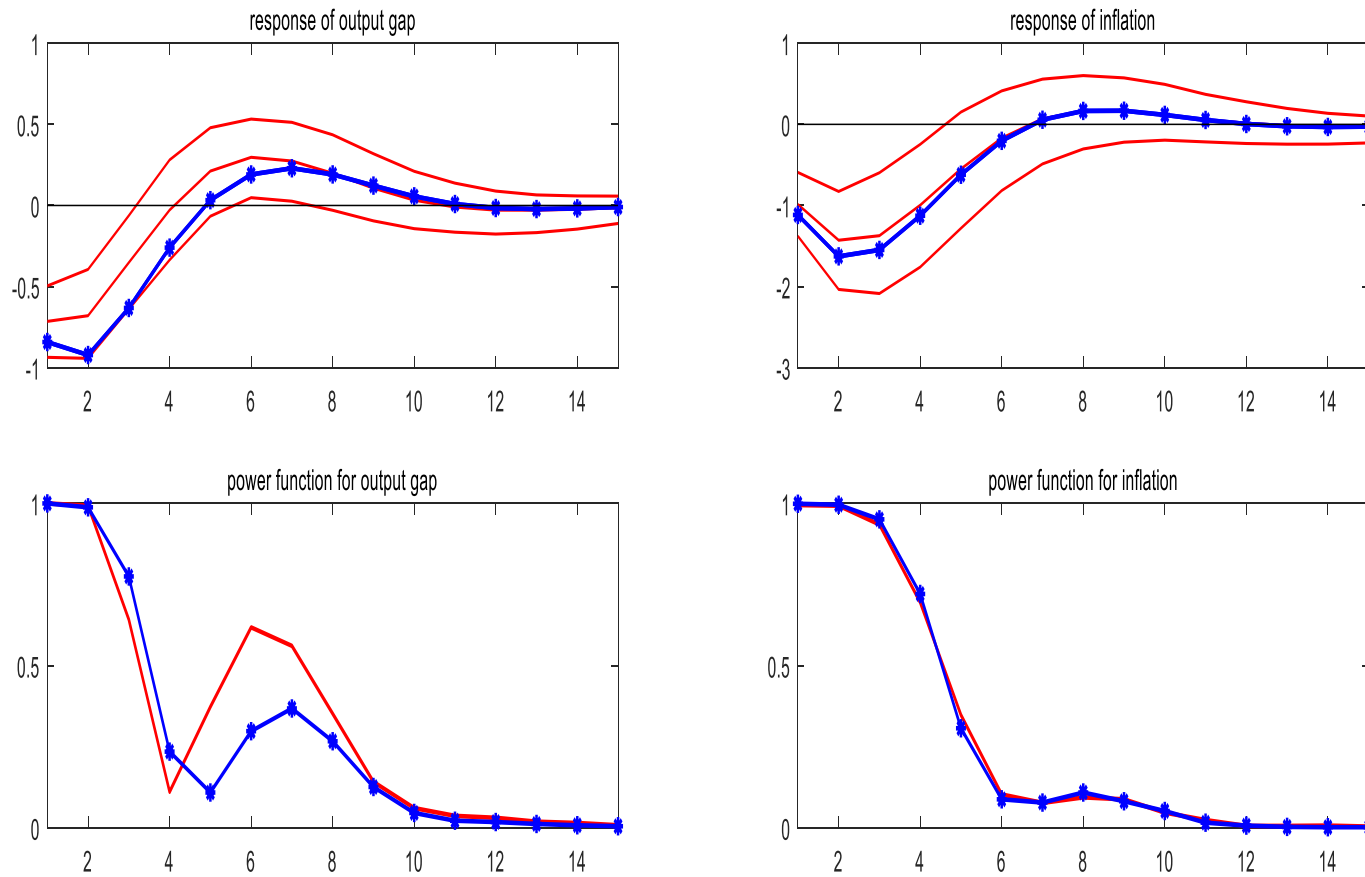
Notes: See text and notes to Figures 1a, 1b and 2.

Figure 5: Classical Measurement Error (additional 20% variance on output gap and inflation)



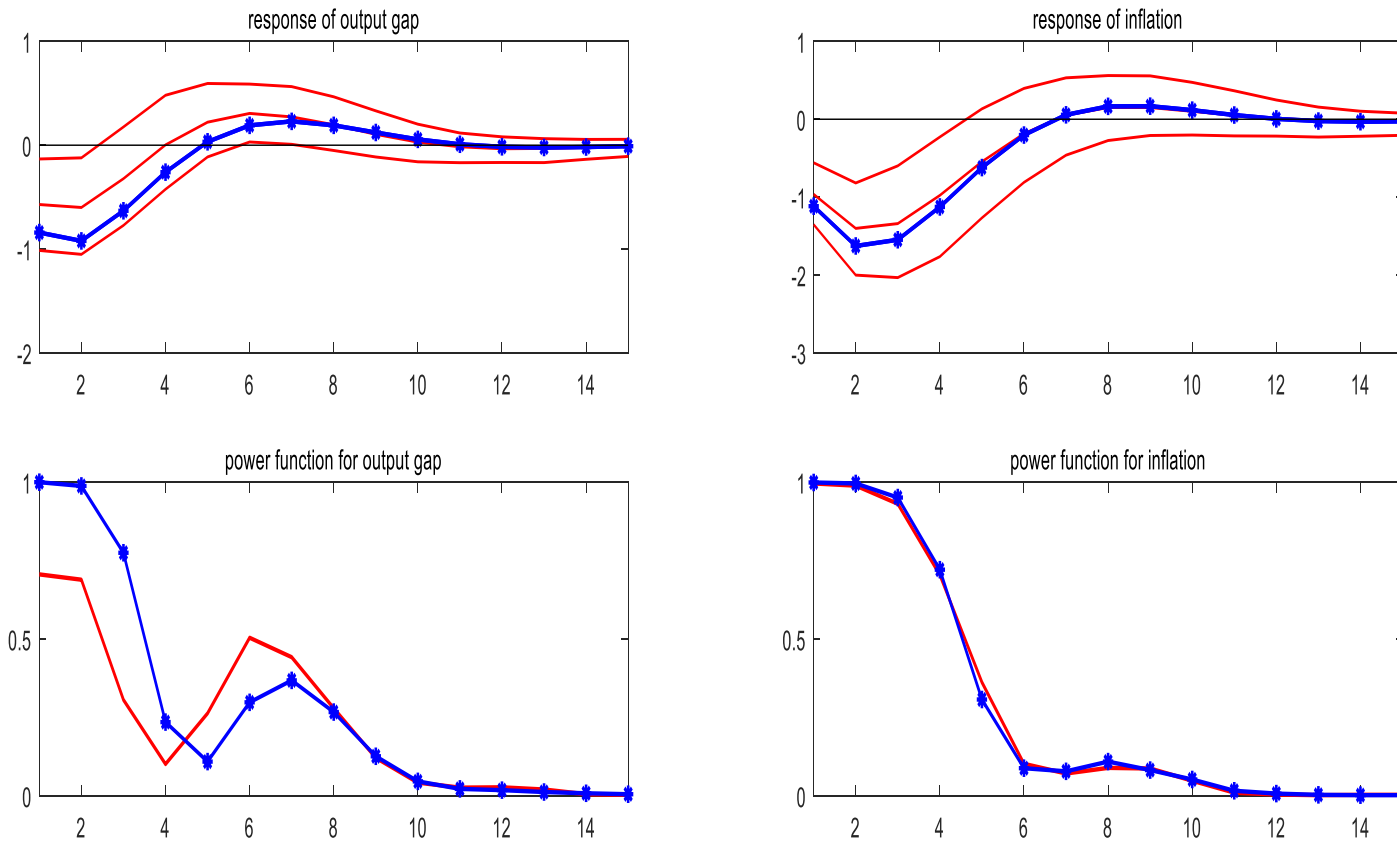
Notes: See text and notes to Figures 1a, 1b and 2..

Figure 6: Output Gap Estimated with One-sided Filter (no supply side shocks)



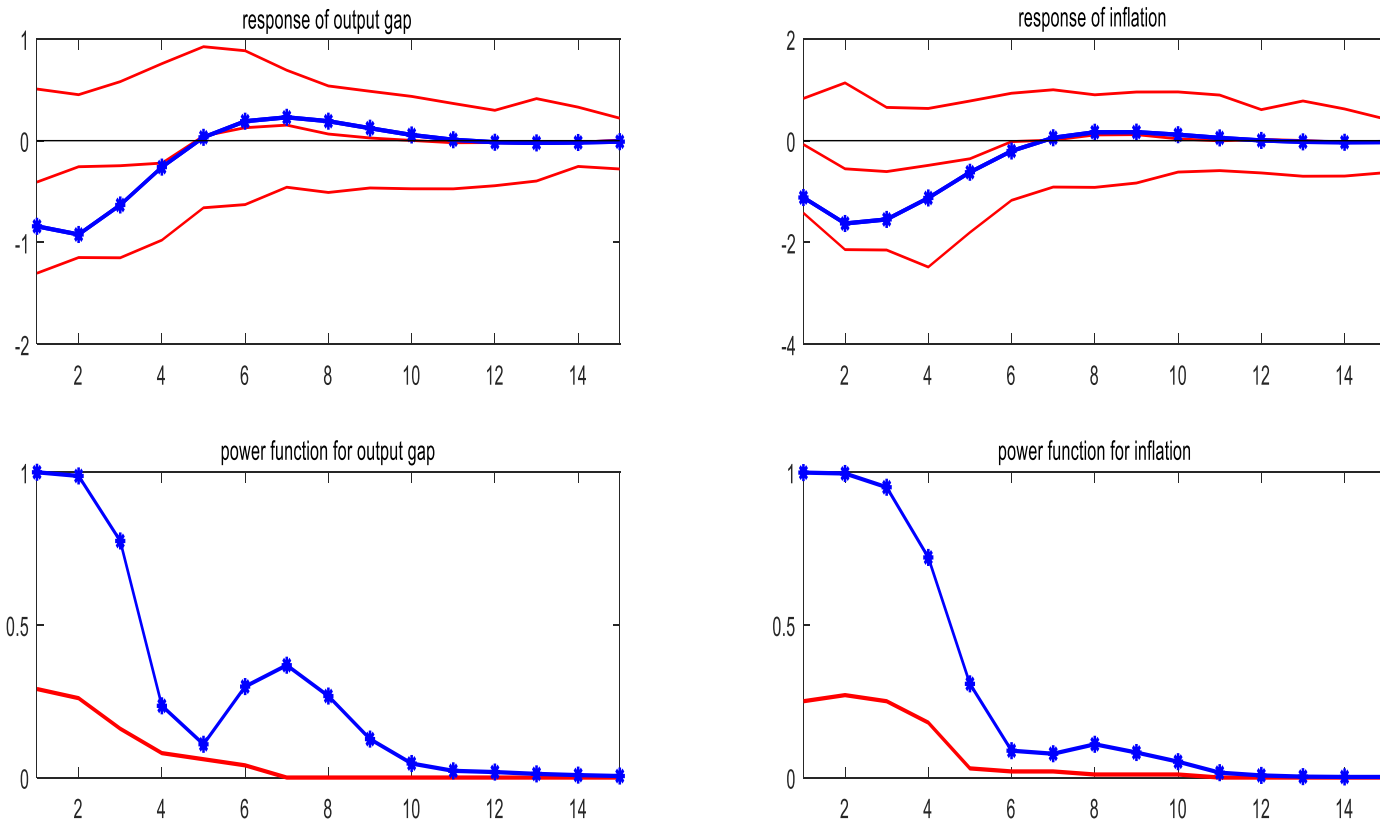
Notes: See text and notes to Figures 1a, 1b and 2.

Figure 7: Output Gap Estimated with One-sided Filter (iid supply side shocks)



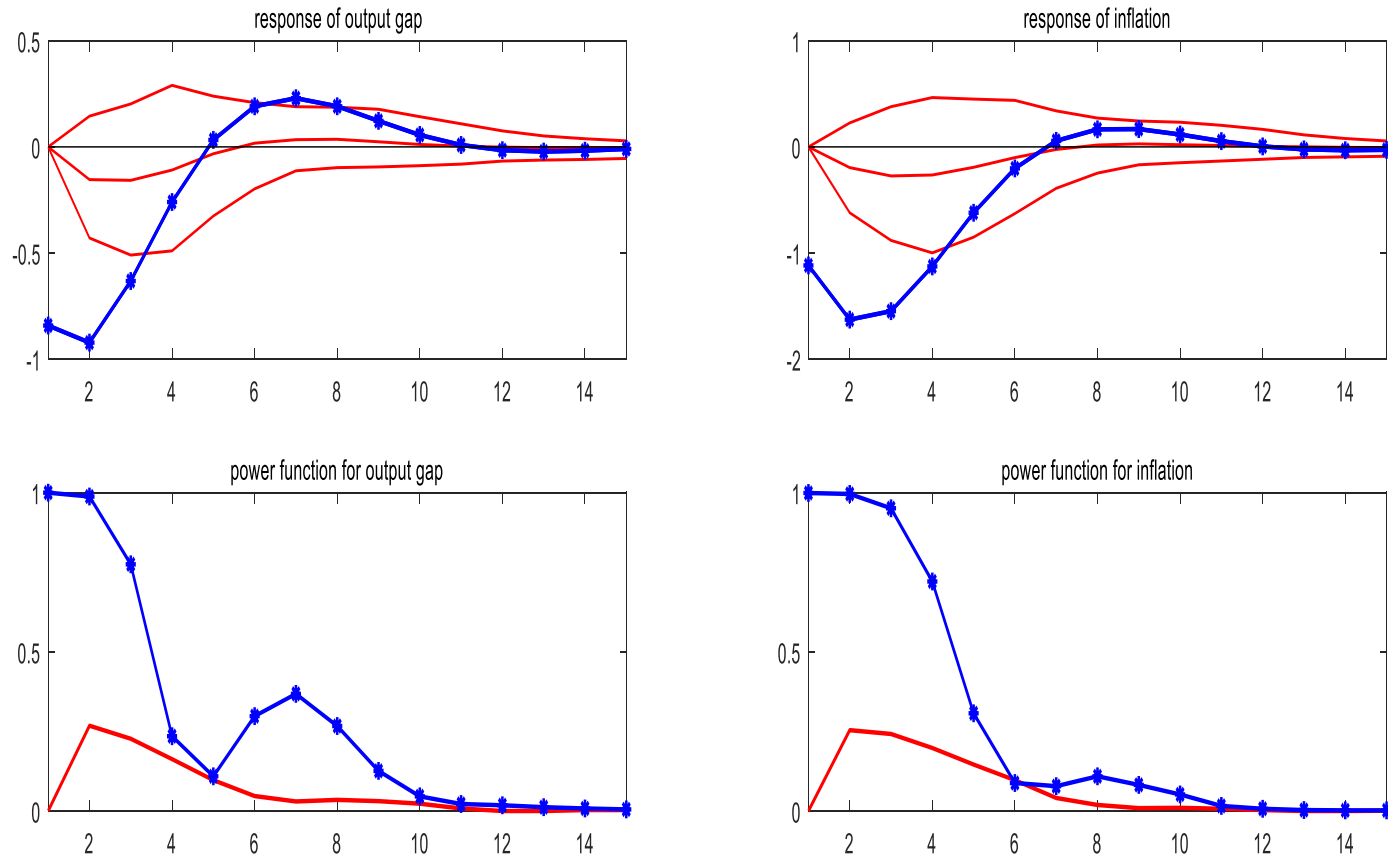
Notes: See text and notes to Figures 1a, 1b and 2.

Figure 8: Combined Scenario (small sample, measurement errors, and iid supply shocks)



Notes: See text and notes to Figures 1a, 1b and 2.

Figure 9: Wrong Identification (assuming central bank has informational advantage)



Notes: See text and notes to Figures 1a, 1b and 2.