A NOTE ON VARIANCE DECOMPOSITION WITH LOCAL PROJECTIONS

Yuriy Gorodnichenko
University of California – Berkeley
and NBER

Byoungchan Lee
University of California – Berkeley

October 27, 2017

Abstract: We propose and study properties of several estimators of variance decomposition in the local-projections framework. We find for empirically relevant sample sizes that, after being bias corrected with bootstrap, our estimators perform well in simulations. We also illustrate the workings of our estimators empirically for monetary policy and productivity shocks.

JEL: E37, E47, C53

Key words: local projections, variance decomposition.

We thank Oscar Jorda and Mikkel Plagborg-Moller for comments on an earlier version of the paper.

I. Introduction

Macroeconomists have been long interested in estimating dynamic responses of output, inflation and other aggregates to structural shocks. While many analyses use vector autoregressions (VARs) or dynamic stochastic general equilibrium (DSGE) models to construct estimated responses, an increasing number of researchers focus on a single structural shock and employ single-equation methods to study the dynamic responses. This approach allows concentrating on well-identified shocks and leaving other sources of variation unspecified. In addition, these approaches often impose no restrictions on the shape of the impulse response function. As a result, the local projections method (Jordà 2005, Stock and Watson 2007) has gained prominence in applied macroeconomic research.

The properties of impulse responses estimated with these methods are well studied (see e.g. Coibion 2012) but little is known about how one can estimate quantitative significance of shocks in the single-equation framework. Specifically, the vast majority of studies using single-equation approaches do not report variance decomposition for the variable of interest and hence one does not know if a given shock accounts for a large share of variation for the variable.[1] This practice contrasts sharply with the nearly universal convention to report variance decompositions in VARs and DSGE models. In this paper, we propose several methods to construct variance decomposition in the local projection framework.

We show that local projections lead to a simple and intuitive way to assess the contribution of identified shocks to variation at different horizons. However, there are several options to implement this insight. While the details of implementation do not matter in large samples, we observe heterogeneity in the performance of various options in small, empirically relevant samples. To illustrate the properties of various methods, we use several data generating processes which cover main profiles of variance decompositions documented in previous works. We show that estimated contributions to variation may be biased in small samples and one should use bootstrap to correct for possible biases in the local projections' estimates of variation decompositions. We also demonstrate how our method works in settings with multiple identified shocks. We illustrate the performance of our method using actual data and commonly used identified shocks as well as data simulated according to the Smets and Wouters (2007) DSGE model. Our work is concurrent and

[1] Coibion et al. (2017) is among the very few papers reporting variance decomposition in the local projection method. More precisely, if we use definitions of Plagborg-Møller and Wolf (2017), the object of our analysis is forecast variance ratio.

complementary to Plagborg-Møller and Wolf (2017) who provide set-identified variance decompositions in the local projections framework.

The rest of the paper is structured as follows. Section II lays out a basic setting to derive the estimators. Section III presents simulation results for bivariate and multivariate settings. Section IV provide an application of our estimators to estimate the contribution of monetary policy and productivity shocks to variation of output and inflation in the local projections framework. Section V concludes.

II. Basics of variance decomposition

Consider a generic setup encountered in studies using local projections. Let y_t be an endogenous variable of interest. We assume that variation in y_t has two components: an identified white-noise shocks series x_t with mean zero and variance σ_x^2 and the "rest" captured by series z_t so that

$$y_t = \sum_{i=0}^{\infty} \psi_{x,i} x_{t-i} + z_t = \psi_x(L) x_t + z_t. \quad (1)$$

We are interested in estimating coefficients in the lag polynomial $\psi_x(L)$ which provides us with the impulse response function of variable y_t to shock x_t. We make only a few assumptions about properties of x_t and z_t. Specifically, we assume that z_t admits an integrated $MA(\infty)$ representation,

$$\Delta z_t = g_y + \psi_e(L) e_t \quad (2)$$

where e_t is a zero-mean white noise series with variance σ_e^2. Following the conventions of local projection applications, we assume that x_t and e_t are uncorrelated and that $\sum_{i=0}^{\infty} \psi_{x,i}^2 < \infty$ and $\sum_{i=0}^{\infty} \psi_{e,i}^2 < \infty$. Without loss of generality we set $\psi_{e,0} = 1$. We assume that $\{(x_t, \Delta y_t)\colon t = 1, \dots, T\}$ is observable.

Forecast error for *h*-period ahead value of the endogenous variable is given by

$$f_{t+h|t-1} \equiv y_{t+h} - y_{t+h|t-1} = (y_{t+h} - y_{t-1}) - E[y_{t+h} - y_{t-1} | \Omega_{t-1}]$$

where $y_{t+h|t-1} \equiv E[y_{t+h}|\Omega_{t-1}]$ is the prediction of y_{t+h} given information set $\Omega_{t-1} \equiv \{\Delta y_{t-1}, x_{t-1}, \Delta y_{t-2}, x_{t-2}, \dots\}$. We can decompose forecast error due to innovations in x_t and other sources of variation as follows[2]

$$f_{t+h|t-1} = \psi_{x,0} x_{t+h} + \cdots + \psi_{x,h} x_t + v_{t+h|t-1}. \quad (3)$$

[2] If $\psi_e(L)$ is invertible, $v_{t+h|t-1}$ is equal to $\psi_{e,0} e_{t+h} + \cdots + (\psi_{e,0} + \cdots + \psi_{e,h}) e_t$. This representation in e_t's is obtained, because $e_t \in \Omega_t$. See Appendix A for details. Note that we do *not* need invertibility of $\psi_e(L)$ to construct the contribution of x_t to variability in y_t. Intuitively, we only need an estimate of either $Var(f_{t+h|t-1})$ in equations (4) and (4'), or $Var(v_{t+h|t-1})$ in equation (4'') which does not require us separating $\psi_e(L)$ and e_t.

Following Sims (1980), we can define the population share of variance explained by the future innovations in x_t to the total variations in $f_{t+h|t-1}$:

$$s_h = \frac{Var\left(\psi_{x,0}x_{t+h} + \cdots + \psi_{x,h}x_t\right)}{Var\left(f_{t+h|t-1}\right)} \tag{4}$$

$$= \frac{\left(\sum_{i=0}^{h}\psi_{x,i}^2\right)\sigma_x^2}{Var\left(f_{t+h|t-1}\right)} \tag{4'}$$

$$= \frac{\left(\sum_{i=0}^{h}\psi_{x,i}^2\right)\sigma_x^2}{\left(\sum_{i=0}^{h}\psi_{x,i}^2\right)\sigma_x^2 + Var\left(v_{t+h|t-1}\right)}. \tag{4''}$$

Equation (4) demonstrates that we have several options for estimating s_h and these options vary in their reliance on imposing parametric structure. In what follows, we propose and evaluate several methods to estimate s_h.

A. R^2 method

Let $X_t^h = (x_{t+h}, \dots, x_t)'$. It can be shown with some algebra that equation (4) can be written as

$$s_h = \frac{Cov\left(f_{t+h|t-1}, X_t^h\right)\left[Var\left(X_t^h\right)\right]^{-1}Cov(X_t^h, f_{t+h|t-1})}{Var(f_{t+h|t-1})}. \tag{5}$$

This quantity can be understood as an R^2 of the population projection of $f_{t+h|t-1}$ on X_t^h, or probability limit of sample R^2. This observation suggests a natural estimator for s_h. First, the forecast errors for each horizon h are estimated using local projections. Second, forecast error for horizon h at time t $f_{t+h|t-1}$ is regressed on shocks x that happen between t and $t+h$. The R^2 in this regression is an estimate of s_h.

More precisely, the estimated forecast error $\hat{f}_{t+h|t-1}$ is the residual of the following regression:

$$y_{t+h} - y_{t-1} = c_h + \sum_{i=1}^{L_y}\gamma_i^h\,\Delta y_{t-i} + \sum_{i=1}^{L_x}\beta_i^h\,x_{t-i} + f_{t+h|t-1}, \tag{6}$$

which is an approximation to $y_{t+h} - y_{t-1} = c_h + \sum_{i=1}^{\infty}\gamma_i\,\Delta y_{t-i} + \sum_{i=1}^{\infty}\beta_i^h\,x_{t-i} + f_{t+h|t-1}$ in population. Then we run the following regression and calculate its R^2:

$$\hat{f}_{t+h|t-1} = \alpha_{x,0}x_{t+h} + \cdots + \alpha_{x,h}x_t + \tilde{v}_{t+h|t-1}. \tag{7}$$

Thus, our first estimator is $\hat{s}_h^{R2} = R^2$ which, by construction, is between 0 and 1. Note that $\alpha_{x,i}$ in equation (7) corresponds to $\psi_{x,i}$ in equation (1). Because $\hat{f}_{t+h|t-1}$ in equation (7) is a residual of an OLS regression with the intercept in equation (6) and x_t is assumed to be zero mean, an

intercept term in equation (7) is not required. Moreover, the population mean of both $f_{t+h|t-1}$ and X_t^h are zeros, so both centered and non-centered R^2's are the same in population. We report results for the non-centered R^2, but properties are similar when we use the centered R^2. The following proposition derives the asymptotic distribution of the estimator.

Proposition 1. Suppose $f_h = \left(f_{T|T-h-1}, f_{T-1|T-h-2}, \dots, f_{L_{max}+h+1\,|L_{max}}\right)'$ and $X_h = \left(X_{T-h}^h, X_{T-1}^h, \dots, X_{L_{max}+1}^h\right)'$ for all $h \geq 0$ where $L_{max} = \max\{L_x, L_y\}$. Then the R^2 of the regression of $f_{t+h|t-1}$ on X_t^h, given by $\left(f_h' P_{X_h} f_h\right)/(f_h' f_h)$ where $P_{X_h} = X_h(X_h'X_h)^{-1}X_h'$, has the following asymptotic distribution for some V_{h,R^2}:

$$\sqrt{T}\left(\frac{f_h' P_{X_h} f_h}{f_h' f_h} - s_h\right) \xrightarrow{d} \mathcal{N}\left(0, \quad V_{h,R^2}\right).$$

Proof. See Appendix B1.

In practice, we may plug the estimated forecast errors from equation (6) in the place of f_h. Appendix B1 contains details of implementation. Note that, instead of using shocks $x_t, \dots, x_{t+h}$ in equation (7), one may want to use residuals from projecting $x_t, \dots, x_{t+h}$ on lags of x_t and Δy_t from equation (6) to guarantee that one does not use forecastable movements in $x_t, \dots, x_{t+h}$ to account for variation in $\hat{f}_{t+h|t-1}$. In practice, however, shocks x_t are constructed in ways to ensure that x_t is not predictable by lags of macroeconomic variables. As a result, we find in our simulations and applications that purifying structural shocks make little difference. Relatedly, one may implement this estimator by augmenting equation (6) with shocks $x_t, \dots, x_{t+h}$ and calculating partial R^2. This insight also justifies using $\hat{f}_{t+h|t-1}$ instead of $f_{t+h|t-1}$ in Proposition 1.

B. Local projection based methods

The R^2 approach requires estimation of two regressions for each horizon (first, construct forecast errors; second, compute the contribution of shocks x between t and $t+h$). However, one can estimate variance decomposition from the local projection directly. Following Jordà (2005), we can estimate $\psi_{x,h}$ from the following equation:

$$y_{t+h} - y_{t-1} = c_h^{\text{LP}} + \sum_{i=1}^{L_y} \gamma_i^{h,LP} \Delta y_{t-i} + \sum_{i=0}^{L_x} \beta_i^{h,LP} x_{t-i} + r_{t+h|t-1} \qquad (8)$$

where $\hat{\beta}_0^{h,LP}$ is an estimate of $\psi_{x,h}$. Since we can estimate σ_x^2 directly from x_t, we can calculate $\left(\sum_{i=0}^{h}\psi_{x,i}^2\right)\sigma_x^2$ in the numerator of equation (4'). To compute the denominator in equation (4'), we note that the residual in equation (8) can be related to the forecast error $f_{t+h|t-1}$ in equation (6). For example, $\hat{f}_{t|t-1} = \hat{\beta}_0^{0,LP}x_t + \hat{r}_{t|t-1}$, that is, a part of forecast error $f_{t|t-1}$ is explained by shock x happening at time t which is now included as one of the regressors in equation (8). In a similar spirit, we can use equation (3) to compute $\hat{f}_{t+h|t-1} = \hat{\beta}_0^{h,LP}x_t + \hat{r}_{t+h|t-1}$. With these estimates of $\hat{f}_{t+h|t-1}$, we can compute $\widehat{Var}(\hat{f}_{t+h|t-1})$ where $\widehat{Var}(\cdot)$ denotes a sample variance. Using these insights, we define a local projection estimator of variance decomposition "LPA" as

$$s_h^{LPA} = \frac{\left(\sum_{i=0}^{h}\left\{\hat{\beta}_0^{i,LP}\right\}^2\right)\hat{\sigma}_x^2}{\widehat{Var}\left(\hat{\beta}_0^{h,LP}x_t + \hat{r}_{t+h|t-1}\right)} \tag{9}$$

where $\hat{\sigma}_x^2 \equiv \widehat{Var}(x_t)$.

Although simple, LPA estimator does not guarantee that in small samples the estimated s_h is between 0 and 1. A simple solution to this issue is to split the denominator into variation due to x and due to v so that $\left(\sum_{i=0}^{h}\psi_{x,i}^2\right)\sigma_x^2$ appears in both the numerator and denominator as in equation (4''). Note that

$$\begin{aligned}\hat{v}_{t+h|t-1} &= \hat{f}_{t+h|t-1} - \hat{\beta}_0^{h,LP}x_t - \hat{\beta}_0^{h-1,LP}x_{t+1} - \cdots - \hat{\beta}_0^{0,LP}x_{t+h} \\ &= \hat{r}_{t+h|t-1} - \hat{\beta}_0^{h-1,LP}x_{t+1} - \cdots - \hat{\beta}_0^{0,LP}x_{t+h}\end{aligned}$$

so that

$$\widehat{Var}\left(f_{t+h|t-1}\right) = \widehat{Var}\left(\hat{r}_{t+h|t-1} - \hat{\beta}_0^{h-1,LP}x_{t+1} - \cdots - \hat{\beta}_0^{0,LP}x_{t+h}\right) + \sum_{i=0}^{h}\left(\hat{\beta}_0^{i,LP}\right)^2\hat{\sigma}_x^2$$

which we use to define another local projection estimator of variance decomposition "LPB":

$$s_h^{LPB} = \frac{\left(\sum_{i=0}^{h}\left\{\hat{\beta}_0^{i,LP}\right\}^2\right)\hat{\sigma}_x^2}{\sum_{i=0}^{h}\left(\hat{\beta}_0^{i,LP}\right)^2\hat{\sigma}_x^2 + \widehat{Var}\left(\hat{r}_{t+h|t-1} - \hat{\beta}_0^{h-1,LP}x_{t+1} - \cdots - \hat{\beta}_0^{0,LP}x_{t+h}\right)}. \tag{9'}$$

Using tools from Proposition 1, we can derive the asymptotic distribution of the LPA and LPB estimators.

Proposition 2. The local projections based estimators when $f_{t+h|t-1}$ is observable have the following asymptotic distributions for some $V_{h,LPA}$ and $V_{h,LPB}$:

$$\sqrt{T}\left(\frac{\sum_{i=0}^{h}\left(\hat{\beta}_0^{i,LP}\right)^2\hat{\sigma}_x^2}{\widehat{Var}(f_{t+h|t-1})}-s_h\right)\xrightarrow{d}\mathcal{N}\left(0,\ V_{h,LPA}\right),\quad \text{and}$$

$$\sqrt{T}\left(\frac{\sum_{i=0}^{h}\left(\hat{\beta}_0^{i,LP}\right)^2\hat{\sigma}_x^2}{\sum_{i=0}^{h}\left(\hat{\beta}_0^{i,LP}\right)^2\hat{\sigma}_x^2+\widehat{Var}\left(r_{t+h|t-1}-\sum_{i=0}^{h-1}\hat{\beta}_0^{i,LP}\,x_{t+h-i}\right)}-s_h\right)\xrightarrow{d}\mathcal{N}\left(0,\quad V_{h,LPB}\right).$$

Proof. See Appendix B2.

C. Small-sample refinements

To correct for potential small-sample biases in the estimates of s_h and to enhance coverage rates for confidence bands, we bootstrap $\hat{s}_h^{R2}$, $\hat{s}_h^{LPA}$, and $\hat{s}_h^{LPB}$ using an estimated VAR model which includes two variables $\{x_t, \Delta y_t\}$. While our implementation of bootstrap is aimed to remove potential biases, alternative implementations may also refine asymptotic inference. Details on how bootstrap is implemented are relegated to Appendix E.

D. Extension

While our analysis has focused on the bivariate case, the framework can be easily generalized to include more controls in equation (6):

$$y_{t+h}-y_{t-1}=\sum_{i=1}^{L_x}\beta_i^h\,x_{t-i}+\sum_{i=1}^{L_C}C_{t-i}'\Gamma_i^h+f_{t+h|t-1} \tag{10}$$

where C_t is the vector of control variables which may include structural shocks other than x_t. In the base case, C_t consists only of Δy_t. Note that for VAR-based bootstrap, one has to include x_t and *all* variables in C_t to simulate data.[3] Similar adjustments are also possible for LPA and LPB methods.

One should bear in mind that, although including or excluding C_t or changing the composition of variables in C_t should make little difference of impulse responses estimated with local projections (provided x is uncorrelated with other shocks), what goes in C_t is potentially important for variance decomposition. Intuitively, by including more controls in C_t, we (weakly) reduce the size of the forecast error (that is, information set Ω_t expands) and hence the amount of variation to be explained shrinks. In other words, the regressand in equation (7) and therefore s_h

[3] As the number of variables in C_t increases, the number of parameters in the VAR increases rapidly. When C_t is a large vector, or when a VAR is not a good representation of the DGP for control variables, VAR-based bootstrap might not be an appealing option. In such a case, one may correct for biases by simulating asymptotic distribution of primitive quantities in (4) such as $\hat{\psi}_{x,i}, \hat{\sigma}_x^2$, and $\widehat{Var}(\hat{v}_{t+h|t-1})$. By considering s_h as a non-linear function of those parameters, such simulations would detect biases due to the non-linearity. See Appendix B for implementation and E and F for the results.

change with a change in the list of variables included in C_t. Thus, one should not be surprised to observe that the share of variation explained by x may be sensitive to changes in C_t.

III. Simulations

This section presents two sets of simulations. The first set shows results for the baseline bivariate case and studies the performances of the three estimators for various profiles of contribution of x to variance of y at different horizons. The second set uses the estimated Smets and Wouters (2007) model to investigate the performance in a setting with many control variables.

For each data generating process (DGP), we simulate data 2,000 times. When we employ bootstrap to correct for biases, the number of bootstrap replications is set to *B*=2,000. As a benchmark, we report results based on a corresponding VAR. This benchmark corresponds to the practice of including shocks into VARs directly (e.g., Basu et al. 2006, Ramey 2011, Barakchian and Crowe 2013, Romer and Romer 2004, 2010). We choose the Hannan-Quinn information criterion (HQIC) as our benchmark criterion to determine the number of lags in VAR. To make VAR and LP models comparable, we use HQIC number of lags in the VAR to set L_x and L_y. Results are similar when we use AIC instead of HQIC.

The sample size for simulated data is $T = 160$. Results for other sample sizes are reported in Appendices E and F. Standard errors are computed as the standard deviation of estimates across bootstrapped samples. The coverage rates are calculated as $\Pr\left(\left|\frac{\hat{s}_h - s_h}{s.e.(\hat{s}_h)}\right|\right) \leq 1.65$.

A. Bivariate Data Generating Processes

We study three DGPs to cover different shapes of s_h. The basic structure is as follows:

$$
\begin{aligned}
&y_t = \psi_x(L)x_t + z_t \\
&z_t = p_t + a_t, \\
&(\Delta p_t - g_y) = \rho_p(\Delta p_{t-1} - g_y) + e_t^p, \qquad e_t^p \sim iid\ N(0, \sigma_p^2), \\
&a_t = \rho_a a_{t-1} + e_t^a, \qquad e_t^a \sim iid\ N(0, \sigma_a^2), \\
&x_t \sim iid\ N(0, \sigma_x^2),
\end{aligned}
$$

where x_t, e_t^p and e_t^a are mutually independent, p_t and a_t are permanent and transitory components of z_t. Appendix C derives the population $MA(\infty)$ representation of Δz_t.

DGP1 is characterized by hump-shaped ψ_x and s_h. We assume that $\psi_x(L)x_t$ follows an $MA(100)$ process with the maximum response set to 3 after 8 periods.[4] DGP2 has a strong response of y to x only in the short-run and thus the shape of s_h is downward-sloping. Finally, DGP3 assumes that $\psi_x(L)$ has a unit root so that x has persistent effects on y and the shape of s_h is upward-sloping. Table 1 reports parameter values for each DGP. Figure 1 plots true impulse responses of y to x (Panel A) and the contribution of x to variation in y (Panel B).

For DGP1, we find (Table 2) that local projections capture the hump-shaped impulse response correctly but s^{R2}, s^{LPA} and s^{LPB} fail to match the hump-share dynamics of s_h. s^{R2}, s^{LPA} and s^{LPB} tend to monotonically increase with the horizon. The VAR misses the hump both in the impulse response and variance decomposition as HQIC selects too few lags (on average the number of lags is 1.27). Confidence bands yield poor coverage rates. This performance reflects the fact that, by construction, shock x contributes zero variation in y for this DGP at short horizons. Since s_h is between zero and one, we effectively have estimates close to the boundary and, therefore, standard methods are likely to fail. While bootstrap appears to provide some improvement (e.g., the bias at long horizons when x accounts for a larger share of variance in y is corrected)[5], it does not perform consistently better because the parameter is at the boundary. When we allow x to explain 5 percent or more of the variation in y at short horizons, bootstrap brings coverage rates close to nominal (results are available upon request). Note that, although VAR is strongly biased, the VAR estimates tend to have smaller variance so that the root mean squared error (RMSE) is similar in magnitude to RMSE of the s^{R2}, s^{LPA} and s^{LPB} estimators. Finally, we observe that the s^{R2}, s^{LPA} and s^{LPB} estimators have similar performance.

Because DGP2 permits an exact, finite-order VAR representation,[6] VAR has good properties in terms of bias, RMSE and coverage rates (Table 3). The local projections recover the share of the impulse response correct, but the estimates of contribution of x to variance of y again overstate the contribution at long horizons. Bootstrap can correct this bias. Given that VAR nests

[4] This value and pattern is motivated by a 3 percent response of real GDP to a 100bp monetary policy shock estimated in Coibion (2012).

[5] The bias can be further reduced by using higher values of L_x and L_y by reducing errors in $\hat{f}_{t+h|t-1}$ due to the truncation.

[6] Given the parameter values in Table 1, $\Delta y_t = g_y + (1-L)(1-0.9L)^{-1}x_t + (1-0.9L)^{-1}e_t^p$. By pre-multiplying $(1-0.9L)$, we have $\Delta y_t = 0.1g_y + 0.9\Delta y_{t-1} - x_{t-1} + x_t + e_t^p$.

the DGP and that VAR is more parsimonious than local projections, VAR has a better performance than the s^{R2}, s^{LPA} and s^{LPB} estimators.

Because x has long-lasting effects on y in DGP3, the VAR underestimates the responses at long horizons in small samples. Impulse responses estimated with local projections perform better but also exhibit a downward bias at long horizons. In a similar spirit, $\hat{s}_h$ shows a strong downward bias for VAR and a smaller, but still considerable bias for the s^{R2}, s^{LPA} and s^{LPB} estimators (this is the case even after we use bootstrap to correct for possible biases). This performance reflects the fact that HQIC chooses a low number of lags (1.34 lags on average across simulations). As a result, VARs used to simulate bootstrap samples fail to capture the degree of persistence in the data. To demonstrate the importance of the lag order, we report results (Table 4) when we use VAR(5) and VAR(10) for bootstrap. As the number of lags increases, we observe some improvement but these enhancements are achieved at the price of higher variance in the estimates. These results suggest that one may want to overfit VAR for persistent processes at the bootstrap stage.

In summary, we find for small samples that the s^{R2}, s^{LPA} and s^{LPB} estimators perform reasonably well across the DGPs and that bootstrap helps to improve the estimators' properties. In addition, there is relatively little difference between the s^{R2}, s^{LPA} and s^{LPB} estimators. In contrast, VARs that include structural shock x tend to perform poorly when a DGP is not nested in a small-order VAR.

B. Smets-Wouters model

While the bivariate DGPs provide important insights about how the R^2, LPA and LPB estimators perform, researchers face potentially more complex DGPs and often have more information in practice. In this section, we use the Smets and Wouters (2007) model to study performance of our estimators in an environment with multiple shocks and many control variables.

As discussed above, different information sets determine different population s_h. In the simulations, we assume that the researcher is interested in explaining variation in output and that the researcher observes output growth rate, inflation, federal funds rate, and monetary policy shocks.[7] This choice of variables is motivated by the popularity of small VARs which include

[7] For this information set, we construct the true variance decomposition using a stationary Kalman filter similar to the method in Appendix C. We also tried various combinations of shocks and endogenous variables and found similar results. Figures for inflation and results with large samples are in Appendix F. Note that monetary policy shocks are *nearly* invertible in the Smets-Wouters model (see Wolf (2017) for more details). While this may be a problem if we

output, inflation and a policy rate to study effects of monetary policy on the economy. In this exercise, the shock is ordered first because the Smets-Wouters model allows contemporaneous responses of macroeconomic variables to policy shocks. When estimating impulse responses using local projections, we augment equation (8) with inflation and federal funds rate as controls.

We find (Figure 2) that local projections correctly recover the response of output to monetary policy shocks, while a low order VAR (lag length is chosen with HQIC) fails to capture the transitory effect of monetary shocks on output. Consistent with our bivariate analysis, $\hat{s}_h^{R2}$, $\hat{s}_h^{LPA}$ and $\hat{s}_h^{LPB}$ increase with the horizons while the true s_h exhibits hump-shaped dynamics. s_h estimated with a VAR also fails to capture the true dynamics as $\hat{s}_h$ flattens out after about $h = 5$. Similar to our results in the previous section, we find that bias correction helps $\hat{s}_h^{R2}$, $\hat{s}_h^{LPA}$ and $\hat{s}_h^{LPB}$ to recover the true hump-shaped profile of s_h. Coverage rates (after bias correction) are 10 percentage points lower their nominal values at short horizons ($h \leq 5$) but the coverage rates are close to nominal at longer horizons. Again, although VAR estimates of s_h are strongly biased, the variance of estimates is low so that RMSE is broadly similar cross methods. We conclude that our proposed methods to estimate variance decomposition work well in more complex settings.

IV. Application

To illustrate the properties our estimators, we use two structural shocks identified in the literature. The first shock is the monetary policy innovation identified as in Romer and Romer (2004) and extended in Coibion et al. (2017). The second shock is the total factor productivity (TFP) shock identified as in Fernald (2014).[8] The correlation between the shocks is -0.059. Our objective is to quantify the contribution of these shocks to variation of output and inflation. The sample covers 1969Q1-2007Q4 which excludes the period of binding zero lower bound. The set of variables for local projections includes inflation (annualized growth rate of GDP deflator, i.e. $400\Delta\ln(P_t)$), annual GDP growth rate ($400\Delta\ln(Y_t)$), federal funds rate, and the two-shock series. We set $L_C = L_x = 4$ in equation (10) and add control variables similarly when estimating impulse responses. In the benchmark VAR, we have all five variables and allow four lags.[9]

use shocks identified and recovered from a DSGE model, the spirit of our exercise is to assume that we have access to other information (as in e.g. Romer and Romer (2004)) so that we can observe monetary policy shocks directly.

[8] When we use empirically identified shocks, measurement errors might be an issue. Given measurement errors, we show that asymptotic biases of our estimators are negative in Appendix D. Therefore, results here can be understood as conservative estimates. In addition, shocks are often estimated and thus are generated regressors, but if the researcher is interested in testing the null of no response then there is no need to adjust inference (Pagan 1984).

[9] The ordering of shocks in the VAR is TFP shock, output growth rate, inflation, monetary policy shock, fed funds rate.

Consistent with previous studies, we find (Figures 3 and 4) that a contractionary monetary policy shocks lowers output and prices, and that a positive TFP shock raises output and lowers prices. Impulse responses estimated with VAR and local projections are similar. VAR estimates for variance decomposition suggest that each of the shocks accounts for approximately 10 percent of variation in output. According to the VAR estimates, monetary policy shocks account for approximately 25 percent of variation in inflation at long horizons and little variation at short horizons while the contribution of TFP shocks is generally small. Bias correction makes no material difference for the variance decomposition estimates for all cases but one: the bias-corrected estimate of the contribution of monetary policy shocks to variation of inflation at long horizons is reduced to about 10 percent.

Local projections estimate that the contribution of the two shocks to variation of output is approximately twice as large as the contribution in VAR estimates. Consistent with simulations, bias correction tends to generate lower contributions but generally the magnitudes are similar. Specifically, when we use the s^{R2} estimator, monetary policy shocks account for approximately 20 percent of variation in output according to local projection estimates (25 percent without bias correction) and approximately 10 percent according to VAR estimates. While the LPB estimator yields similar results, the LPA estimator assigns a much larger role to the monetary policy shocks. This pattern reflects the fact that $\hat{s}_h^{LPA}$ may be greater than 1 in finite samples. Also, note that, in contrast to the profile of s_h estimated with VAR for output (which is generally flat after $h = 5$), $\hat{s}_h^{R2}$, $\hat{s}_h^{LPA}$ and $\hat{s}_h^{LPB}$ have richer dynamics.

In a similar spirit, the contribution of TFP and monetary policy shocks to variation in inflation is much greater according to our local-projections estimates. The difference is particularly large for monetary shocks: $\hat{s}_h^{R2}$ and $\hat{s}_h^{LPB}$ are close to 40 percent (after bias correction) and $\hat{s}_h^{VAR}$ is about 10 percent at long horizons. Again, $\hat{s}_h^{LPA}$ estimates an even greater contribution of monetary shocks and confidence intervals are much wider for $\hat{s}^{LPA}$ than for $\hat{s}^{LPB}$ or $\hat{s}^{R2}$. Again, this stems from the fact that $\hat{s}_h^{LPA}$ may be greater than 1 in finite samples.

V. Concluding remarks

Single-equation methods can offer flexibility and parsimony that many economists seek. The increasing popularity of these methods, specifically the local projections, calls for further development of these tools. An important limitation for practitioners using this framework has been a lack of simple tools to assess quantitative significance of a given set of shocks, that is, the

contribution of the shocks to variance of the variable of interest. We propose several methods to provide such a metric. In a series of simulation exercises, we document that these methods have good small-sample properties. We also show that conventional approaches to assess the quantitative significance of two popular structural shocks (monetary policy shocks and total factor productivity shocks) could have understated the importance of these two shocks.

References

Barakchian, S. Mahdi, and Christopher Crowe, 2013. "Monetary policy matters: Evidence from new shocks data," *Journal of Monetary Economics*, 60(8): 950-966.

Basu, Susanto, John G. Fernald, and Miles S. Kimball, 2006. "Are Technology Improvements Contractionary?" *American Economic Review*, 96(5): 1418–1448.

Coibion, Olivier, 2012. "Are the Effects of Monetary Policy Shocks Big or Small?" *American Economic Journal: Macroeconomics*, 4(2): 1–32.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia, 2017. "Innocent Bystanders? Monetary policy and inequality," *Journal of Monetary Economics*, 88(C): 70–89.

Fernald, John, 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Working Paper 2012-19.

Jordà, Oscar, 2005. "Estimation and Inference of Impulse Responses by Local Projections," A*merican Economic Review*, 95(1): 161–182.

Pagan, Adrian, 1984. "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review*, 25(1): 221-247.

Plagborg-Møller, Mikkel, and Christian K. Wolf, 2017. "Instrumental Variable Identification of Dynamic Variance Decompositions," manuscript.

Ramey, Valerie A., 2011. "Identifying Government Spending Shocks: It's all in the Timing," *Quarterly Journal of Economics*, 126(1): 1–50.

Romer, Christina D., and David H. Romer, 2010. "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, 100(3): 763-801.

Romer, Christina, D., and David H. Romer, 2004. "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94(4): 1055-1084.

Sims, Christopher A., 1980. "Macroeconomics and reality." *Econometrica*, 48(1): 1-48.

Smets, Frank, and Rafael Wouters, 2007. "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American Economic Review*, 97(3): 586-606.

Stock, James, and Mark Watson, 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Banking and Credit*, 39(1): 3–33.

Wolf, Christian K., 2017. "Masquerading Shocks in Sign-Restricted VARs," manuscript.

Table 1. Parameter values for data generating processes (DGPs) used in simulations.

	$\psi_x(L)$	σ_x	g_y	ρ_p	σ_p	ρ_a	σ_a
DGP1	Hump-shaped	1	0.5	0.9	0.5	0.9	3
DGP2	$(1-0.9\mathrm{L})^{-1}$	3	0.5	0.9	1.5	-	-
DGP3	$(1-\mathrm{L})^{-1}(1-0.9\mathrm{L})^{-1}$	1	0.5	0.5	2	0.9	3

Table 2. Simulation results for DGP 1.

	Horizon h					
	0	4	8	12	16	20
Impulse response						
True	0.00	1.39	3.00	2.06	0.88	0.29
Local projections	0.00	1.36	2.99	2.03	0.85	0.29
VAR(HQIC)	0.00	0.18	0.26	0.27	0.27	0.27
Variance decomposition						
True	0.00	0.04	0.19	0.21	0.18	0.14
Average estimate						
R2	0.01	0.06	0.20	0.25	0.26	0.27
LP A	0.01	0.04	0.18	0.23	0.23	0.23
LP B	0.01	0.04	0.17	0.22	0.21	0.21
VAR(HQIC)	0.01	0.02	0.02	0.03	0.03	0.03
Root mean squared error						
R2	0.01	0.05	0.11	0.16	0.19	0.22
LP A	0.01	0.04	0.11	0.15	0.18	0.20
LP B	0.01	0.04	0.10	0.14	0.15	0.15
VAR(HQIC)	0.01	0.03	0.17	0.20	0.16	0.14
Coverage (90 % level) (asymptotic)						
R2	1.00	0.94	0.74	0.71	0.75	0.73
LP A	1.00	0.94	0.53	0.57	0.74	0.80
LP B	1.00	0.93	0.53	0.55	0.70	0.78
VAR(HQIC)	1.00	0.57	0.06	0.05	0.07	0.09
Variance decomposition (bias corrected, VAR(HQIC))						
True	0.00	0.04	0.19	0.21	0.18	0.14
Average estimate						
R2	0.00	0.02	0.13	0.16	0.13	0.11
LP A	0.00	0.02	0.14	0.17	0.16	0.14
LP B	0.00	0.02	0.14	0.17	0.15	0.13
VAR(HQIC)	0.00	0.00	0.01	0.02	0.02	0.02
Root mean squared error						
R2	0.01	0.05	0.13	0.16	0.17	0.18
LP A	0.01	0.04	0.12	0.16	0.17	0.18
LP B	0.01	0.04	0.12	0.14	0.15	0.15
VAR(HQIC)	0.01	0.04	0.19	0.21	0.18	0.15
Coverage (90 % level)						
R2	1.00	0.93	0.59	0.61	0.67	0.79
LP A	1.00	0.82	0.45	0.49	0.59	0.79
LP B	1.00	0.81	0.46	0.47	0.57	0.73
VAR(HQIC)	1.00	0.41	0.05	0.05	0.06	0.08

Table 3. Simulation results for DGP 2

	Horizon h					
	0	4	8	12	16	20
Impulse response						
True	3.00	1.97	1.29	0.85	0.56	0.36
Local projections	2.99	1.88	1.18	0.71	0.43	0.22
VAR(HQIC)	2.96	1.98	1.40	1.04	0.82	0.67
Variance decomposition						
True	0.80	0.25	0.10	0.05	0.03	0.02
Average estimate						
R2	0.79	0.27	0.15	0.14	0.15	0.18
LP A	0.80	0.27	0.13	0.10	0.09	0.09
LP B	0.79	0.26	0.13	0.09	0.09	0.09
VAR(HQIC)	0.80	0.27	0.13	0.08	0.06	0.05
Root mean squared error						
R2	0.03	0.11	0.12	0.14	0.17	0.21
LP A	0.03	0.09	0.08	0.08	0.09	0.11
LP B	0.03	0.08	0.07	0.08	0.09	0.10
VAR(HQIC)	0.03	0.08	0.07	0.06	0.05	0.05
Coverage (90 % level) (asymptotic)						
R2	0.92	0.87	0.92	0.89	0.85	0.80
LP A	0.93	0.90	0.94	0.95	0.93	0.93
LP B	0.90	0.88	0.93	0.94	0.92	0.90
VAR(HQIC)	0.90	0.88	0.91	0.97	0.99	0.99
Variance decomposition (bias corrected, VAR(HQIC))						
True	0.80	0.25	0.10	0.05	0.03	0.02
Average estimate						
R2	0.81	0.25	0.09	0.04	0.01	0.00
LP A	0.79	0.25	0.10	0.05	0.03	0.02
LP B	0.81	0.25	0.10	0.05	0.03	0.02
VAR(HQIC)	0.80	0.25	0.10	0.05	0.03	0.02
Root mean squared error						
R2	0.03	0.10	0.09	0.09	0.11	0.13
LP A	0.03	0.08	0.07	0.07	0.07	0.08
LP B	0.03	0.08	0.07	0.07	0.07	0.08
VAR(HQIC)	0.03	0.08	0.06	0.05	0.04	0.04
Coverage (90 % level)						
R2	0.91	0.90	0.97	0.98	0.98	0.97
LP A	0.93	0.88	0.91	0.98	0.97	0.97
LP B	0.89	0.83	0.89	0.97	0.97	0.96
VAR(HQIC)	0.89	0.88	0.89	0.92	0.99	1.00

Table 4. Simulation results for DGP 3 with alternative lag orders in VARs.

	Horizon h					
	0	4	8	12	16	20
Impulse response						
True	1.00	4.10	6.13	7.46	8.33	8.91
Local projections	0.99	3.96	5.78	6.86	7.44	7.66
VAR(5)	0.93	3.74	4.76	5.01	5.10	5.14
VAR(10)	0.92	3.65	5.34	6.05	6.17	6.23
Variance decomposition (bias corrected, VAR(5))						
True	0.06	0.29	0.47	0.58	0.65	0.70
Average estimate						
R2	0.06	0.26	0.41	0.50	0.55	0.58
LP A	0.05	0.24	0.38	0.48	0.54	0.58
LP B	0.06	0.25	0.40	0.49	0.54	0.57
VAR(5)	0.06	0.24	0.33	0.36	0.38	0.39
Root mean squared error						
R2	0.04	0.11	0.16	0.19	0.21	0.23
LP A	0.04	0.11	0.17	0.21	0.25	0.28
LP B	0.04	0.11	0.16	0.19	0.20	0.22
VAR(5)	0.04	0.11	0.19	0.26	0.31	0.34
Coverage (90 % level) (asymptotic)						
R2	0.77	0.80	0.80	0.81	0.82	0.83
LP A	0.85	0.83	0.82	0.83	0.84	0.85
LP B	0.84	0.79	0.78	0.78	0.79	0.80
VAR(5)	0.83	0.78	0.66	0.51	0.40	0.33
Variance decomposition (bias corrected, VAR(10))						
True	0.06	0.29	0.47	0.58	0.65	0.70
Average estimate						
R2	0.07	0.30	0.47	0.57	0.63	0.66
LP A	0.05	0.23	0.37	0.47	0.52	0.55
LP B	0.05	0.27	0.44	0.54	0.60	0.63
VAR(10)	0.06	0.27	0.42	0.50	0.54	0.56
Root mean squared error						
R2	0.05	0.12	0.16	0.19	0.21	0.22
LP A	0.04	0.12	0.17	0.21	0.25	0.29
LP B	0.04	0.12	0.16	0.18	0.20	0.21
VAR(10)	0.04	0.11	0.15	0.19	0.21	0.23
Coverage (90 % level)						
R2	0.72	0.76	0.78	0.80	0.82	0.83
LP A	0.87	0.87	0.89	0.90	0.91	0.92
LP B	0.85	0.77	0.75	0.76	0.78	0.79
VAR(10)	0.83	0.83	0.81	0.79	0.77	0.74

Figure 1. Population impulse responses and variance decomposition for each DGP

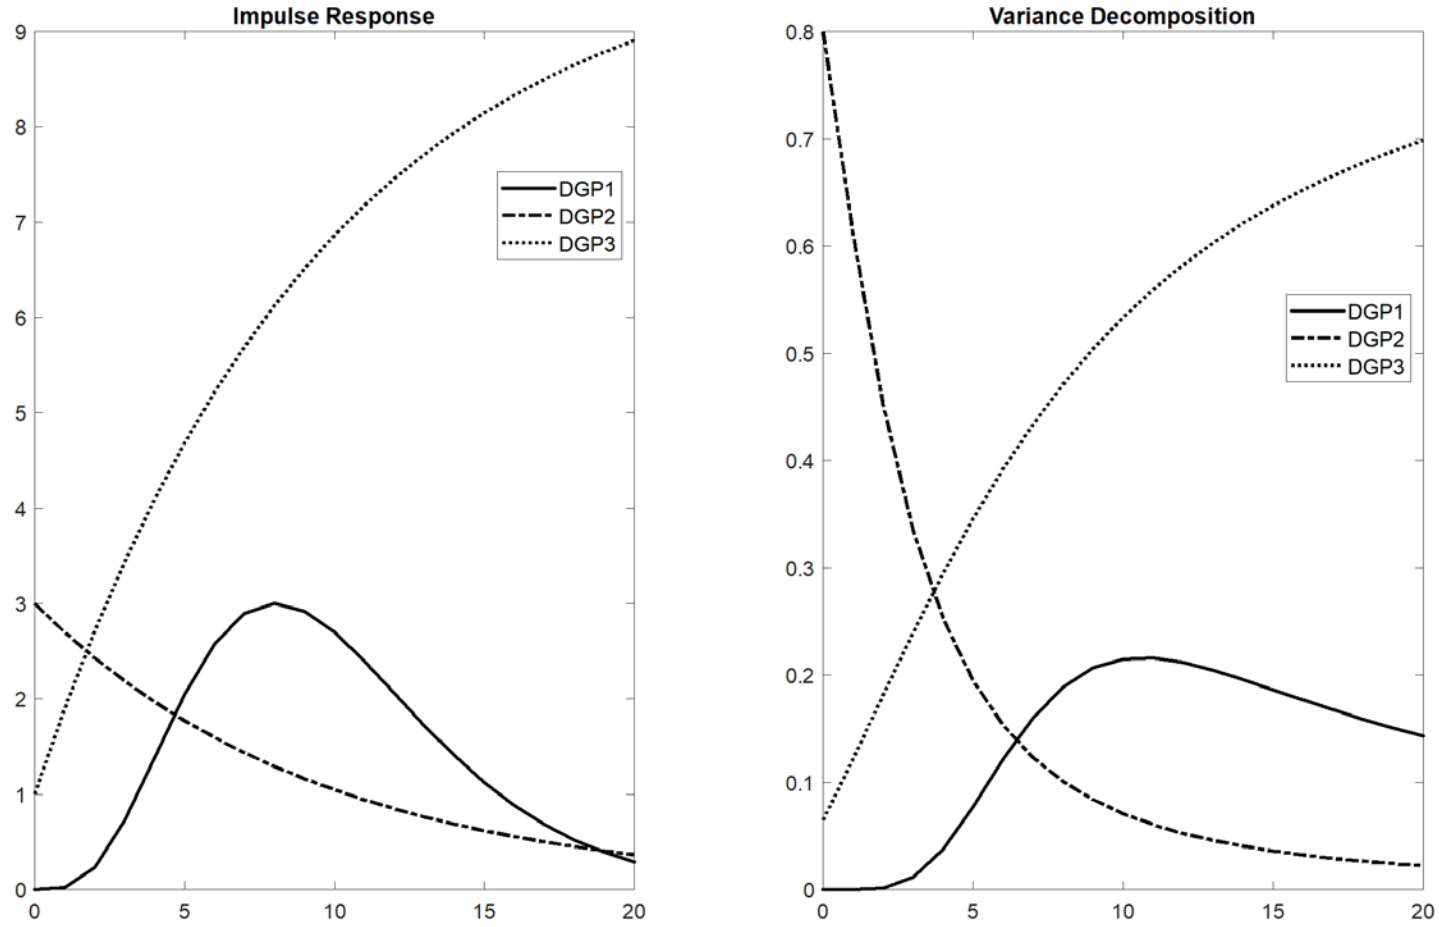

Notes: the left panel shows the impulse response functions for three bivariate data generating processes (DGPs). The right panel shows the contribution of the identified shock to variation of an outcome variable for the DGPs.

Figure 2: Smets and Wouters (2007) model, real GDP and monetary policy shock, T = 160.

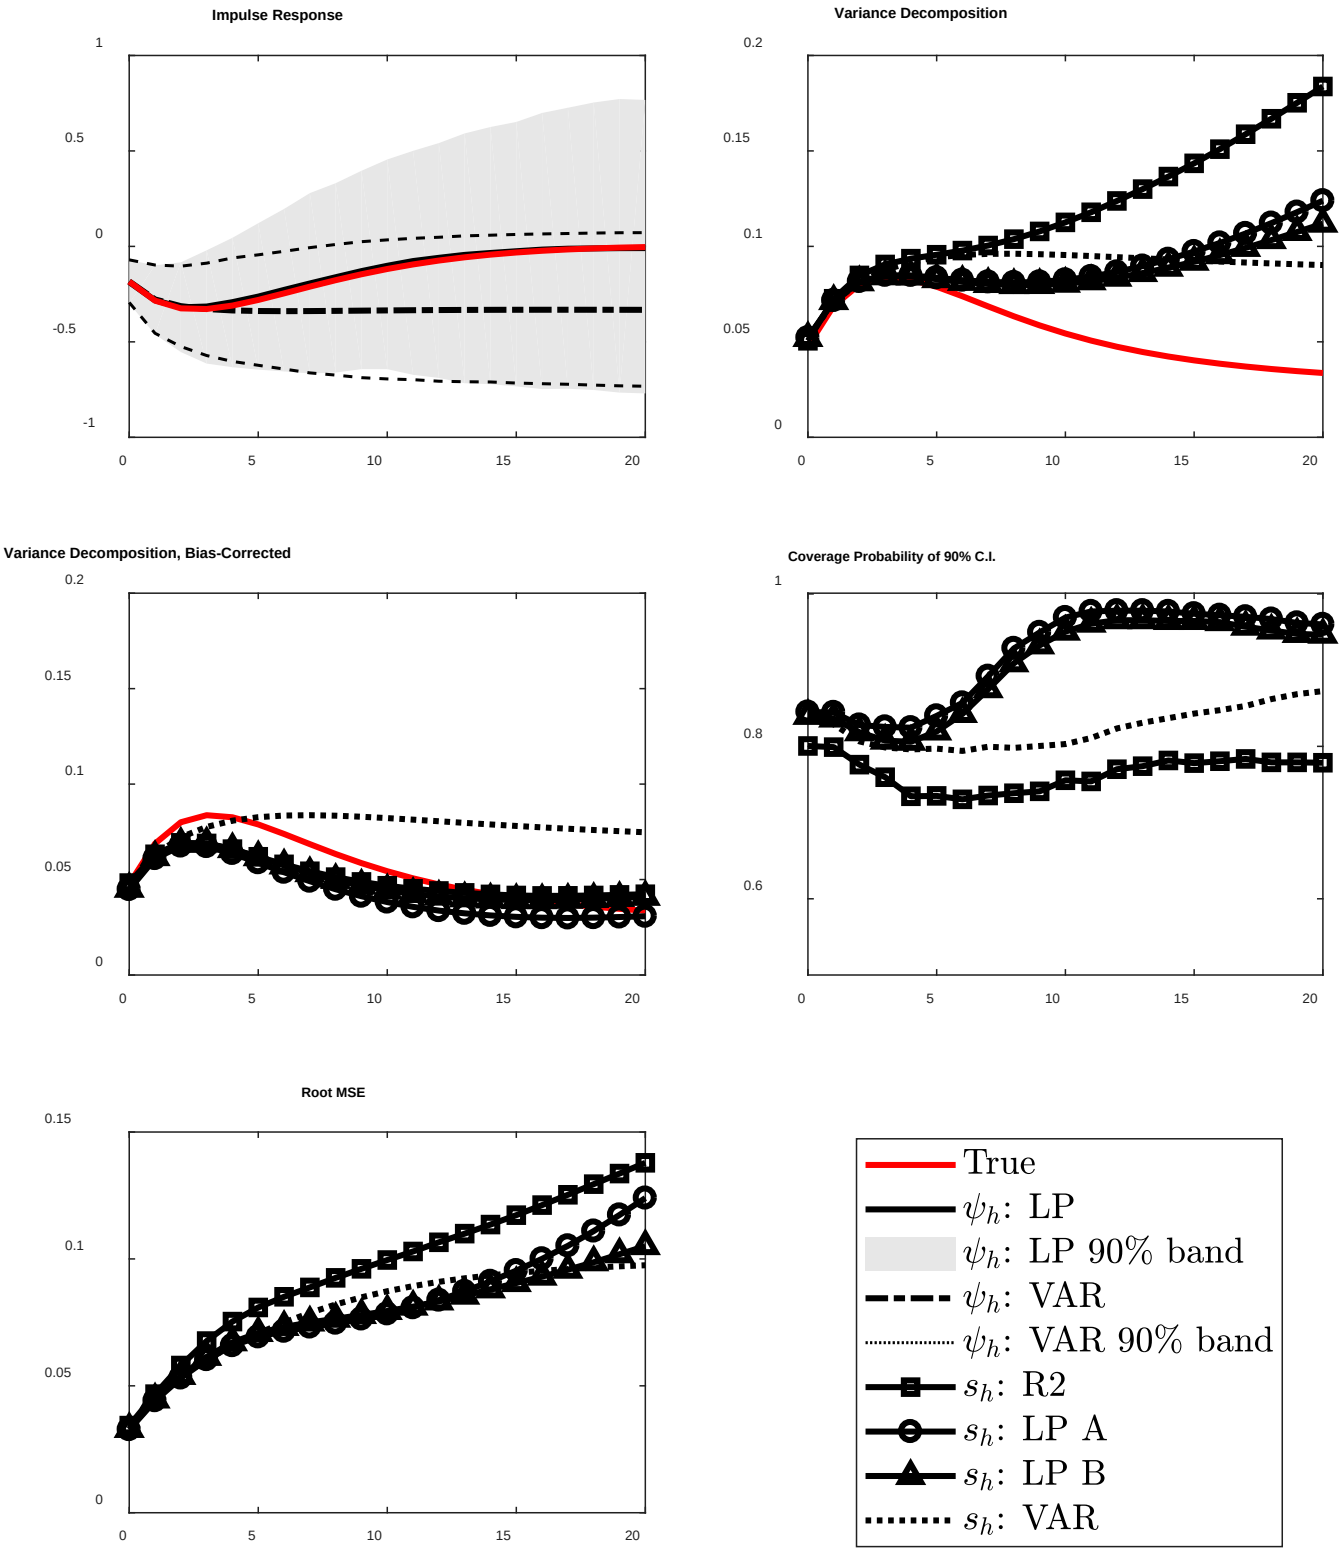

Figure 3 Real GDP. 1969:Q1-2007:Q4.

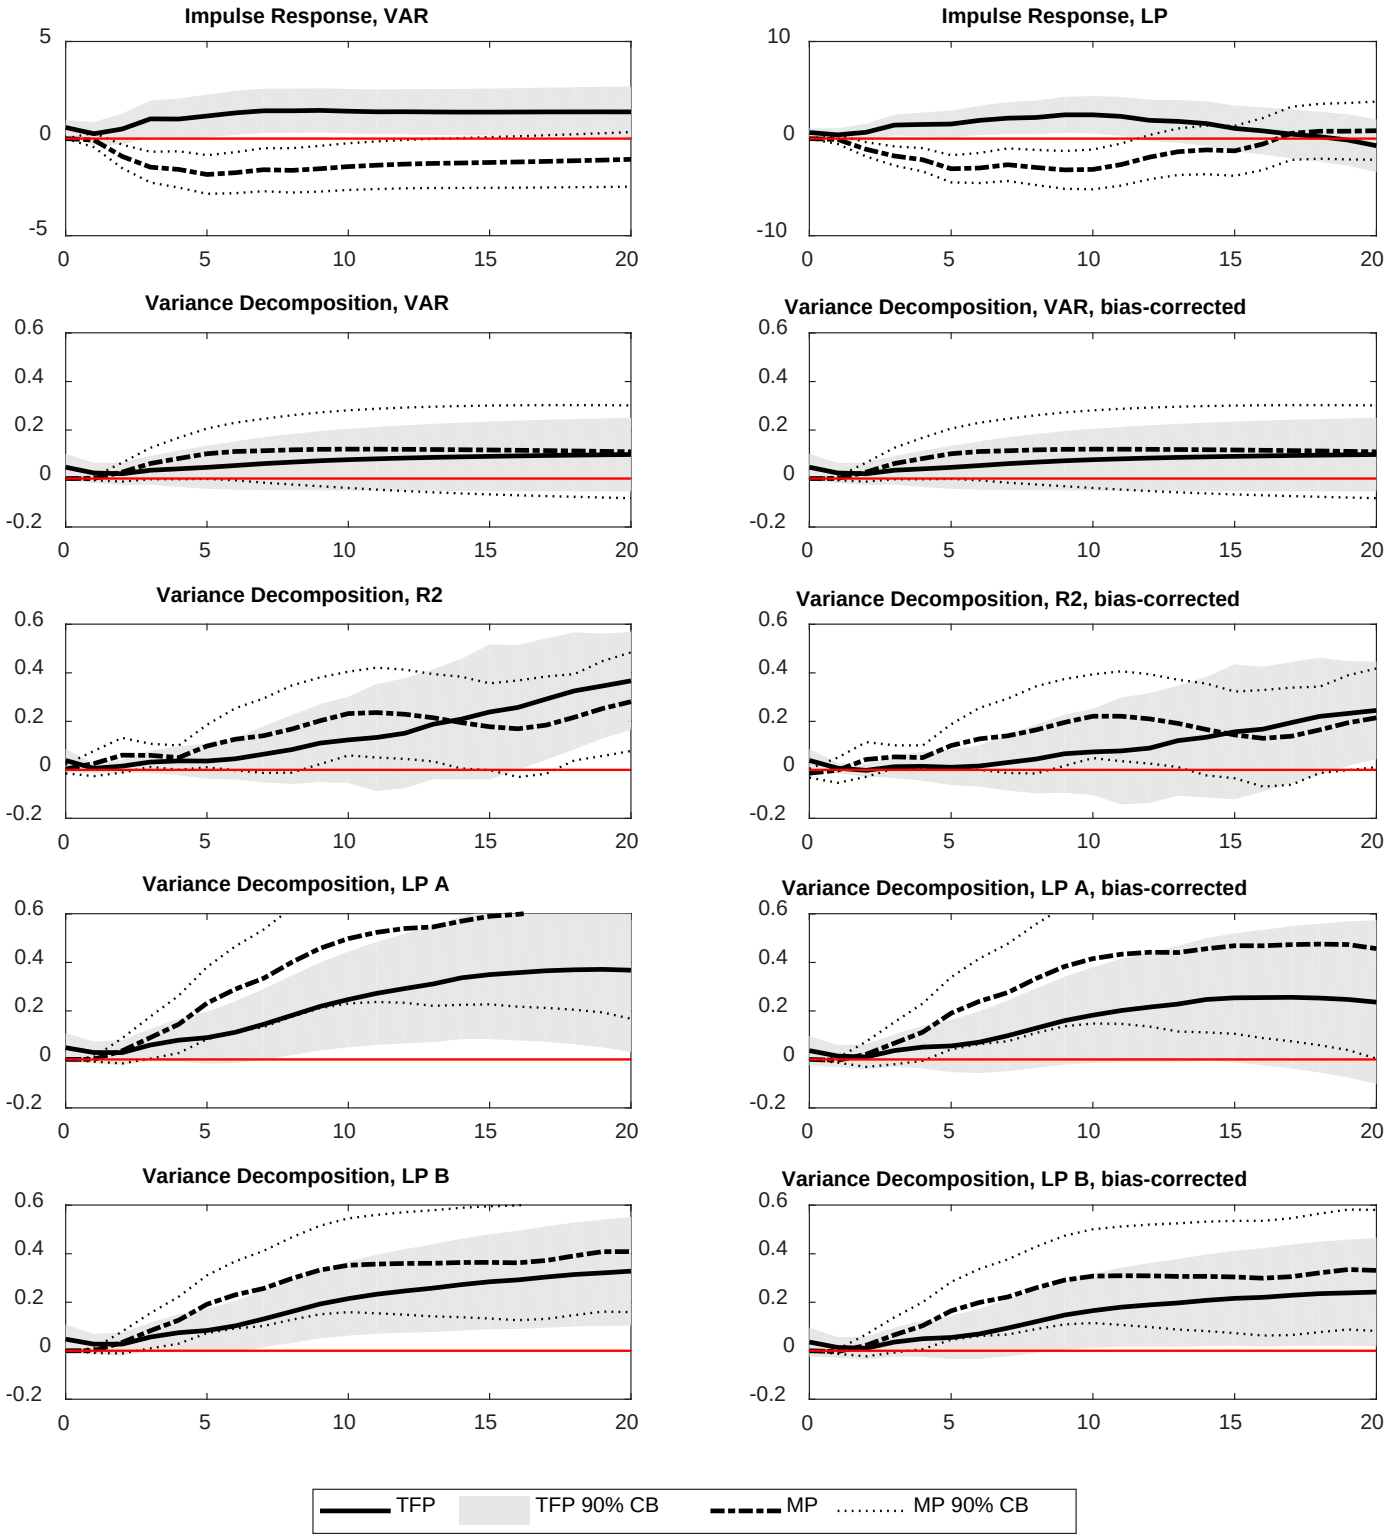

Figure 4. Inflation. 1969:Q1-2007:Q4.