Bootstrapping Structural VARs: Avoiding a Potential Bias in Confidence Intervals for Impulse Response Functions

Phillips, Kerk L. and Spencer, David E.

Brigham Young University

February 2010

Bootstrapping Structural VARs:

Avoiding a Potential Bias in Confidence Intervals for Impulse Response Functions[1]

Kerk L. Phillips
Department of Economics
P.O. Box 22363
Brigham Young University
Provo, UT 84602-2363
phone: (801) 422-5928
fax: (801) 422-0194
email: kerk_phillips@byu.edu

David E. Spencer
Department of Economics
P.O. Box 22363
Brigham Young University
Provo, UT 84602-2363
phone: (801) 422-7277
fax: (801) 422-0194
email: david_spencer@byu.edu

February 2010

JEL Codes: E32, E37, C32

Keywords: impulse response function, structural VAR, bias, bootstrap

Abstract

Constructing bootstrap confidence intervals for impulse response functions (IRFs) from structural vector autoregression (SVAR) models has become standard practice in empirical macroeconomic research. The accuracy of such confidence intervals can deteriorate severely, however, if the bootstrap IRFs are biased. In this paper, we document an apparently common source of bias in the estimation of the VAR error covariance matrix. The bias is easily corrected with a straightforward scale adjustment. This bias is often unrecognized because it only affects the bootstrap estimates of the error variance, not the original OLS estimates. Nevertheless, as we illustrate here, analytically, with sampling experiments, and in an example from the literature, the bootstrap error variance bias can have significant distorting effects on bootstrap IRF confidence intervals even if the original IRF estimate relies on unbiased parameter estimates.

[1] We are grateful to Richard W. Evans, Lutz Kilian, and James B. McDonald for their comments on previous versions of this paper. We are also grateful to Mark Hendricks for excellent research assistance.

1. INTRODUCTION

Impulse response functions (IRFs) from structural vector autoregression (SVAR) models are widely employed to investigate the response of macroeconomic variables to identified structural shocks. Leading and influential examples of such studies include Blanchard and Quah (1989) examining the effects of aggregate demand and aggregate supply shocks on output and unemployment, Galí (1999) which investigates the effects of technology shocks, and Christiano, Eichenbaum, and Evans (1999) which assesses the effects of monetary policy shocks.

To assess uncertainty and draw inferences, these and other studies construct confidence intervals (CIs) around the estimated IRF. Increasingly, these intervals are constructed using bootstrap techniques.[2] In this paper we document a commonly occurring, but easily corrected, source of bias in bootstrap estimates of IRFs from SVAR models.[3] Given the pervasiveness of the techniques that lead to this bias, it has important implications. For example, it can lead to sufficiently distorted CIs with such severe spurious asymmetry that the bootstrap CIs do not even include the estimated IRF. Sims and Zha (1999, p. 1125, fn 13) note that some SVAR studies have found it necessary to "use a modification of [the bootstrap confidence interval] that makes *ad hoc* adjustments to prevent the computed bands from failing to include the point estimates."

This bias-caused distortion can be seen in the results reported by Blanchard and Quah (1989); see especially their Figures 3 and 5. Our Figure 1 is a reestimated[4] version of their Figure 3 with asymmetric one standard deviation bands.[5] Notice that the upper one standard deviation band actually lies below the original estimated IRF over the early horizon interval.[6]

---

[2] See, e.g., Runkle (1987) and Berkowitz and Kilian (2000).

[3] This bias arises from the downward bias in the standard bootstrap estimate of the reduced form VAR error covariance matrix. Any object that depends on these estimates will be affected. This includes not only IRFs but bootstrap confidence intervals for error variance decompositions and bootstrap prediction intervals as well.

[4] We make the same data adjustments made by Blanchard and Quah and estimate the model over the same sample period. Our results differ slightly because we use revised data.

[5] We compute our asymmetric one standard deviation bands by obtaining 1000 bootstrap IRFs and then taking, in each direction, the square root of the mean squared deviation from the mean bootstrap IRF.

[6] The fact that the corresponding Blanchard-Quah IRF does not actually cross the bounds is due to the way they compute their one standard deviation bands. They obtain 1000 bootstrap IRFs which, for each horizon, they separate into those above and those below the original IRF. They then compute the standard deviation for each class to obtain the asymmetric one standard deviation bounds. This procedure assures that the IRF will not "cross"

Anticipating our later discussion, Figure 2 shows the same impulse response function with bias-corrected one standard deviation bands. The original asymmetry is greatly attenuated reflecting the fact that it is a spurious consequence of bias in the bootstrap estimates of the IRF.

If, as in the case of Galí (1999), researchers do not allow for asymmetric confidence intervals and simply plot error bands that are the estimated IRFs plus or minus one or two standard deviations, then the CIs are symmetric by construction, any bias is completely invisible, and the reported error bands are incorrect.[7]

Not all researchers attribute this odd behavior of IRFs completely to skewness. Christiano, Eichebaum and Vigfusson (2006), for example, note that, in their case, the mean value of the bootstrapped IRFs is not the same as the IRF from the original estimation. They plot both of these along with confidence intervals and note that the "asymmetric percentile confidence intervals show that when data are generated by these [bootstrap] VARs, … the impulse response functions have a downward bias."[8]

The bias we examine arises from the fact that the *bootstrap* IRF for a SVAR depends on the *bootstrap* OLS estimate of the error covariance matrix in the reduced form vector autoregression (VAR), standard estimates of which are biased downward. This bias is apparently common[9] but easily corrected by a degrees of freedom adjustment. It is not corrected in practice because it is generally unrecognized as it only affects the *bootstrap* estimates of the error variance, not the original OLS estimates. Nevertheless, bootstrap IRF CIs can be substantially distorted even if the original IRF estimate relies on unbiased parameter estimates.

In the next section, we illustrate the specific source of this bias in the bootstrap estimate of error variances in the context of simple models and confirm its impact. In Section 3 we show how this bias in bootstrap error variance estimates effects the bootstrap IRFs and thus

---

the bounds. A bound that is coincident with the original IRF indicates that, at that horizon, *none* of the bootstrap IRFs were above (or below) the original IRF.

[7] This is practice followed in some econometric software packages like EViews.

[8] Christiano, Eichenbaum, and Vigfusson (2006), p. 26.

[9] Of course, we have not documented this for all or even most SVAR studies. We have, however, examined programs that authors have posted on web sites. In *none* of the cases was the appropriate bias-adjusted bootstrap error covariance estimator used. Some programs (including those which use the standard VCV instruction in RATS) calculate the MLE of the bootstrap covariance matrix and thus make no degrees of freedom adjustment at all. We therefore conclude that this bias is quite common in practice.

the bootstrap confidence intervals for the original IRF. In Section 4 we illustrate how correcting for this bias affects the IRF confidence intervals obtained in a widely-cited previous study. The final section offers a brief conclusion.

2. THE SOURCE OF BIAS

2.1. Standard Regression Models

The simplest way to illustrate the bias under investigation is to examine a standard linear regression model with nonstochastic regressors. We first consider a univariate regression model represented by

(1) $y = X\beta + u$

where *y* is a $T \times 1$ vector of observations on a dependent variable, *X* is an $T \times R$ matrix of observations on *R* nonstochastic regressors (perhaps including a constant), β is an $R \times 1$ vector of regression coefficients, and *u* is an $T \times 1$ vector of errors. We assume that $E(u) = 0$ and $E(uu') = \sigma^2 I_T$. Applying ordinary least squares (OLS), we obtain coefficient and error variance estimates: $\hat{\beta} = (X'X)^{-1} X'y$, $\hat{\sigma}^2 = \frac{1}{T-R}(\hat{u}'\hat{u})$, where $\hat{u} = y - X\hat{\beta}$. The indicated degrees of freedom correction makes $\hat{\sigma}^2$ an unbiased estimator for σ^2.

To help us understand the key argument to follow, it is useful to interpret the degrees of freedom adjustment from the perspective that it is necessary to compensate for the fact that the OLS residuals tend to be smaller than the error terms. Note that the expected value of the average squared error is σ^2; i.e., $E\left(\frac{u'u}{T}\right) = \sigma^2$. On the other hand, $E\left(\frac{\hat{u}'\hat{u}}{T}\right) = \left(\frac{T-R}{T}\right) E\left(\frac{u'u}{T}\right)$, which reflects that, on average, the squared residuals are $\left((T-R)/T\right)$ times as large as the squared errors[10]. Thus, to obtain an unbiased estimate, we must rescale each residual by $\left(\frac{T}{T-R}\right)^{1/2}$ and then compute the average squared *rescaled*

[10] See Davidson and MacKinnon (1993), pp. 69-70.

residual giving the usual unbiased estimate for σ^2, $\hat{\sigma}^2 = \frac{1}{T-R}(\hat{u}'\hat{u})$.

As noted by Freedman and Peters (1984, p. 99) and Peters and Freedman (1984, p. 408), a similar issue arises when obtaining an unbiased *bootstrap* estimate of $\hat{\sigma}^2$. This follows from the analogy on which the bootstrap is based and the insight above. Suppose we obtain *bootstrap* estimates of the error variance as follows. For bootstrap replications b=1,…, B, generate

(2) $y_b^* = X\hat{\beta} + u_b^*$

where the elements of u_b^* are drawn with replacement from the OLS residuals, $\hat{u}$. Then, apply OLS to equation (2) to get bootstrap estimates of $\hat{\beta}$, which we denote $\tilde{\beta}_b$, and bootstrap residuals, $\tilde{u}_b$. In the bootstrap, the variance estimate, $\tilde{\sigma}_b^2$, is an estimate of $\hat{\sigma}^2$, the "population" error variance in the pseudo-population given by the original OLS residuals, $\hat{u}$. The usual bootstrap variance estimate, $\tilde{\sigma}_{b,1}^2 = \frac{1}{T-R}(\tilde{u}_b{}'\tilde{u}_b)$, is biased for $\hat{\sigma}^2$.

Proceeding as above, we note that though $E\left(\frac{\hat{u}'\hat{u}}{T-R}\right) = \sigma^2$,

$E\left(\frac{\tilde{u}_b{}'\tilde{u}_b}{T-R}\right) = \left(\frac{T-R}{T}\right) E\left(\frac{u_b^{*\prime}u_b^*}{T-R}\right) = \left(\frac{T-R}{T}\right)\sigma^2$ since the elements of u_b^* are drawn randomly from $\hat{u}$. This reflects that, on average, the squared bootstrap residuals are $\left((T-R)/T\right)$ times as large as the squared OLS residuals which are the pseudo-population errors. Consequently, $\tilde{\sigma}_{b,1}^2$ yields a biased estimate of $\hat{\sigma}^2$ which, in turn, is an unbiased estimate of σ^2. It follows that an unbiased bootstrap estimate is $\tilde{\sigma}_{b,2}^2 = \frac{T}{(T-R)^2}(\tilde{u}_b{}'\tilde{u}_b) = \frac{T}{T-R}\tilde{\sigma}_{b,1}^2$. An analogous rescaling has been suggested by Stine (1987, p. 1074) and Berkowitz and Kilian (2000, p. 5) in the case of a univariate autoregressive model of order *p*, AR(p), in order to obtain the "desired variance."

The size of the (proportional) bias for the natural estimator is $-R/T$[11]. While this vanishes asymptotically, it can be important in small samples when *R* is large relative to *T*. To illustrate, we conduct a Monte Carlo experiment in which we simulate obtaining bootstrap estimates of the error variance in a univariate regression model like (1). We estimate models with nine regressors including a constant term, $R=9$, for three sample sizes: $T=30,\ 50,\ 100$. Consequently, the true bias for $\tilde{\sigma}_{b,1}^2$ is -30%, -18% and -9% respectively. For each sample size, we draw 1000 samples of size *T* from a normal distribution with mean zero and variance 0.81. For each of these Monte Carlo draws we generate observations for *y*, estimate (1) by OLS, and compute the usual unbiased estimate of the error variance, $\hat{\sigma}^2$. The average estimate is given in Table 1. To examine the bias of the two bootstrap error variance estimates, $\tilde{\sigma}_{b,1}^2$ and $\tilde{\sigma}_{b,2}^2$, we take each of the 1000 Monte Carlo samples and obtain 1000 bootstrap estimates in each case. The average values are reported in Table 1 for our three sample sizes.

The results in Table 1 confirm the theory very nicely. The "natural" bootstrap estimator, $\tilde{\sigma}_{b,1}^2$, has bias approximately equal to $-R/T$ while the other estimators are unbiased.

This bias in the standard bootstrap "error" variance carries over exactly to the case of a multivariate seemingly unrelated regression model with nonstochastic regressors. To confirm the theory, we have conducted simple Monte Carlo experiments similar to those undertaken for the univariate regression model discussed above. To save space, we do not report the results here but simply indicate that the conclusions are the same.[12]

2.2. Autoregressive Models

Consider a univariate AR(p) with a constant term, ν, so that $R=p+1$:

---

[11] It should be noted that bias arising from maximum likelihood estimation (MLE) of the error variance will be even larger. As is well known, the MLE of σ^2, $\breve{\sigma}^2=\frac{1}{T}\left(\hat{u}'\hat{u}\right)$, is biased; i.e., $E\left(\breve{\sigma}^2\right)=\left(\frac{T-R}{T}\right)\sigma^2$. Thus, the proportional bias is $-R/T$. Now, when we bootstrap and obtain the MLE of $\breve{\sigma}^2$, $\tilde{\breve{\sigma}}_b^2=\frac{1}{T}\left(\tilde{u}_b{}'\tilde{u}_b\right)$, the bias is magnified since we have a biased estimate of a biased estimate. $\tilde{\breve{\sigma}}_b^2=\frac{\left(T-R\right)^2}{T^2}\tilde{\sigma}_{b(2)}^2$, so $E\left(\tilde{\breve{\sigma}}_b^2\right)=\frac{\left(T-R\right)^2}{T^2}\sigma^2$ and the proportional bias is $\left(\frac{\left(T-R\right)^2}{T^2}\right)-1=\frac{R^2-2TR}{T^2}$ which is negative and larger (in absolute value) than $-R/T$.

[12] The results are available on request.

(3) $y_t = \nu + \phi_1 y_{t-1} + ... + \phi_p y_{t-p} + u_t; \quad t = -p+1, ..., 0, 1, ..., T$

where u_t is white noise with variance σ^2 and *T* is the number of *usable* observations. Because the regressors are stochastic, the finite sample theory of the previous section does not apply. However, following Stine (1987), we might speculate (correctly) that similar bias problems exist for bootstrap estimators of the error variance in this case.

Since exact analytical results are not available, we examine the small-sample bias issue for the AR(*p*) model using a Monte Carlo exercise similar to the one described above. We generate data for, and estimate, a model like (3) in which $p = 8$ so $R = 9$. For each of three sample sizes, $T = 30,\ 50,\ 100$, we draw 1000 samples for u_t of size *T+p* from a normal distribution with mean zero and variance 0.81. For each of these Monte Carlo draws we generate observations for *y*, estimate (3) by OLS, and compute the usual estimate of the error variance, $\hat{\sigma}^2$. The average estimate is given in Table 2. To examine the bias of the two bootstrap error variance estimates, $\tilde{\sigma}^2_{b,1}$ and $\tilde{\sigma}^2_{b,2}$, we obtain 1000 bootstrap estimates for each of the Monte Carlo samples[13]. The average values are reported in Table 2 for each of our three sample sizes.

The results are quite informative. The exact theoretical bias for the corresponding standard linear regression is a rather good guide for the bias in the AR(*p*) model. We confirm that the bootstrap estimator of the error variance given by $\tilde{\sigma}^2_{b,1}$ is biased and thus likely to result in significant distortion when the number of slope coefficients is large relative to the sample size.

We expect these bias results to carry over to the case of a VAR(*p*) with *K* variables. In that case, our interest is the $K \times K$ error (innovation) covariance matrix Σ. Assuming a constant term, the usual degrees-of-freedom-corrected OLS estimator for Σ is $\hat{\Sigma} = \left(\frac{1}{T - Kp - 1} \right) \hat{U}'\hat{U}$ where $\hat{U}$ is the $T \times K$ matrix of OLS residuals. The "natural" but biased

[13] For each bootstrap iteration, we obtain the initial *p* observations $\{ y_{-p+1}, ..., y_o \}$ by drawing (with replacement) from the original generated sample $\{ y_t \}_{-p+1}^{T}$.

bootstrap estimator of Σ is $\tilde{\Sigma}_{b,1} = \left(\frac{1}{T-Kp-1}\right)\tilde{U}_b{}'\tilde{U}_b$ where $\tilde{U}_b$ is the $T \times K$ matrix of bootstrap residuals from the b^{th} bootstrap iteration. The unbiased bootstrap estimator of Σ is $\tilde{\Sigma}_{b,2} = \left(\frac{T}{(T-Kp-1)^2}\right)\tilde{U}_b{}'\tilde{U}_b$ and $R = Kp+1$.

We have investigated the bootstrap error variance bias for a two-equation VAR(8) model with a constant term using Monte Carlo methods similar to those described above and find the bias to be quite close to the theoretical bias from the corresponding multivariate regression model. To conserve space, we do not report the results here since they are quite similar to those reported for the AR(8) model above[14]. In particular, the bias for $\tilde{\Sigma}_{b,1}$ is approximately $-\left(\frac{Kp+1}{T}\right)$ where *K* is the number of equations (variables) in the VAR(*p*). For a two-equation VAR(8) model, this implies an approximate bias of -17% for each element of Σ when $T = 100$.[15]

3. BOOTSTRAPPING IRFS FOR SVARS

The downward bias of the standard bootstrap estimator of the VAR error covariance matrix is of particular concern when we are interested in drawing inferences about IRFs from a SVAR model since the IRFs are nonlinear functions of both VAR slope parameters *and* the elements of the error covariance matrix.[16] In this section we show how bias in the bootstrap estimate of the VAR error covariance matrix creates a specific bias in bootstrap IRFs and thus bootstrap CIs.

Consider a SVAR model explaining the behavior of a $K \times 1$ vector of variables, y_t. The IRFs are obtained from the moving average representation of the model:

(4) $y_t = A(L)\varepsilon_t$

[14] Results are available on request.

[15] Note that if, for this VAR model, we had computed MLE rather than OLS estimates of Σ in both the initial and bootstrap stages, the approximate bias for the elements of the bootstrap estimate of Σ would have been magnified to -31%. See footnote 4.

[16] Other objects of frequent interest that are also nonlinear functions of VAR slope parameters and elements of the error covariance matrix are forecast error variance decompositions and measures of predictability. Thus, related bootstrap confidence or prediction intervals would also suffer from the bias we discuss here. See Inoue and Kilian (2002).

where ε_t is a vector of *K structural* shocks and we make the standard assumption that $E(\varepsilon_t \varepsilon_t') = I_K$. This assumption provides a normalization as well as a set of identifying restrictions. The elements of the matrix polynomial $A(L)$ give the impulse response functions: $a_{ij,l}$ $(i, j = 1,...K;\ l = 0,1,...)$ indicates the response of variable *i* in *l* periods to a one unit (standard deviation) movement in the *j*th structural shock today. Though the IRFs are frequently the objects of interest in macroeconomic analysis, they cannot generally be estimated directly from time series data since the SVAR model (4) is not identified without further restrictions.

To estimate the SVAR and thus the IRFs, we begin by specifying a finite-order *reduced form* VAR model which can always be estimated:

(5) $B(L) y_t = u_t$

where $B(L)$ is a matrix of polynomials of order *p* and $E(u_t u_t') = \Sigma$. In general, OLS estimates of $B(L)$ and Σ can be obtained, $\hat{B}(L)$ and $\hat{\Sigma}$, where $\hat{\Sigma} = \left(\frac{1}{T - Kp - 1} \right) \hat{U}'\hat{U}$ and $\hat{U}$ is the $T \times K$ matrix of OLS residuals.

The reduced form moving average representation is obtained by inverting (5):

(6) $y_t = B(L)^{-1} u_t = C(L) u_t$

Equating terms in (4) and (6) allows us to conclude the following:

(7) $u_t = A_0 \varepsilon_t$

(8) $A_l = C_l A_0 \quad l = 1,...$

Thus, it is clear that knowledge of the K^2 elements of A_0 is sufficient to obtain the IRF.

From (7) we infer the key relationship between the covariance matrices of the structural and reduced form errors:

(9) $\Sigma = A_0 A_0'$

Symmetry of Σ provides $\left(\frac{K(K+1)}{2} \right)$ restrictions on A_0. With $\left(\frac{K(K-1)}{2} \right)$ additional restrictions, A_0 can be identified and IRFs computed. Equations (8) and (9) assure us that the estimated IRFs depend on the estimates of both $B(L)$ and Σ. I.e.,

$\hat{a}_{ij,l} = g(\hat{\beta}, \hat{\sigma})$ $(i, j = 1,...K;\ l = 1,...)$ where $\hat{\beta} = vec(\hat{B})$, a $K(Kp+1)\times 1$ vector, and $\hat{\sigma} = vech(\hat{\Sigma})$, a $\left(\frac{K(K+1)}{2}\right)\times 1$ vector. Consequently, the properties of the IRFs depend on the properties of $\hat{\beta}$ and $\hat{\sigma}$. Similarly, the properties of the bootstrap IRFs depend in the same way on the properties of the bootstrap estimates of β and σ:
$\tilde{a}_{ij,l} = g(\tilde{\beta}, \tilde{\sigma})$ $(i, j = 1,...K;\ l = 0,1,...)$.

We can see from this that there are several potential sources of bias for the bootstrapped IRFs and, thus, bootstrap confidence intervals for the original IRFs. The source we focus on here arises when the *bootstrap* estimate of σ, $\tilde{\sigma}$, is biased for $\hat{\sigma}$. How much difference does the appropriate bootstrap estimation of the error covariance matrix make for bootstrap estimates of the IRF? We can obtain an exact analytical answer to this question.

From equation (8) we infer that the bootstrap estimates of the IRF are given by

(10) $\tilde{A}_l = \tilde{C}_l \tilde{A}_0 \quad l = 1,...$

where $\tilde{A}_l$, $\tilde{C}_l$, and $\tilde{A}_0$ are bootstrap estimates. If $\tilde{A}_0$ is biased, all $\tilde{A}_l$ will be affected. We see from equation (9) that $\tilde{A}_0$ depends only on the bootstrap estimate of Σ, $\tilde{\Sigma}$. We first derive the "bias"[17] for $\tilde{A}_0$ that arises from a biased estimate of Σ and then derive the resulting bias for the IRF.

Let $\hat{\Sigma} = \hat{A}_0 \hat{A}_0{}'$ be an unbiased estimate of Σ so that $\hat{A}_0$ is the corresponding "unbiased" estimate of A_0. Then $\tilde{\Sigma} = \tilde{A}_o \tilde{A}_o{}' = (1+b)\hat{\Sigma}$ where the scalar *b* reflects the proportional bias in $\tilde{\Sigma}$, a potentially biased bootstrap estimate. Thus,

(11) $\tilde{A}_0 = (1+b)^{1/2} \hat{A}_0 = (1+a)\hat{A}_0$

where *a* is the "bias" in $\tilde{A}_0$. Equating $(1+b)^{1/2}$ and $(1+a)$ in (11) implies that the bias of $\tilde{\Sigma}$, *b*, and the "bias" of $\tilde{A}_0$, *a*, are related by

(12) $a = (1+b)^{1/2} - 1$

[17] We put the term "bias" in quotes here because $\tilde{A}_0$ and $\tilde{\Sigma}$ are related by a quadratic equation. Consequently, we cannot infer the true bias of $\tilde{A}_0$ from the bias of $\tilde{\Sigma}$. The bias we derive is therefore only approximate. We will use this notational device for the next few paragraphs.

When $-1 < b \leq 0$, as it is in our case, we see than $a \geq b$ and thus the "bias" for $\tilde{A}_0$ is negative but closer to zero than the bias for $\tilde{\Sigma}$.

Now, consider how this "bias" in the bootstrap estimate $\tilde{A}_0$ affects the bootstrap IRF given by equation (10). To isolate the effect of a bias in $\tilde{\Sigma}$, we assume that C_l is known (or at least $\tilde{C}_l$ is unbiased). Suppose $\hat{A}_0$ is an unbiased estimate of A_0. We can rewrite equation (11) as

(13) $$\hat{A}_0 = \left(\frac{1}{1+a}\right)\tilde{A}_0$$

where *a* is the scalar proportional bias for $\tilde{A}_0$. The "unbiased" IRF is then given by

(14) $$\hat{A}_l = C_l\hat{A}_0 = \left(\frac{1}{1+a}\right)C_l\tilde{A}_0,\ \ l = 1,...$$

from which we can infer the proportional "bias" for the terms in the IRF

(15) $$\frac{\tilde{A}_l - \hat{A}_l}{\hat{A}_l} = a,\ \ l = 1,....$$

Thus, the bootstrap IRF proportional "bias" is constant and equal to the "bias" for $\tilde{A}_0$ for the entire IRF horizon. So, for example, if we have a SVAR model with *K=2*, *p=8*, *T=100* and a constant term, the "bias" for $\tilde{\Sigma}_{b(1)}$ is -17% and the "bias" for the IRF is -9%.[18]

4. An Example

The bias discussed here is pervasive in the empirical SVAR literature. To illustrate its effect of this bias in practice, we replicated the biased results obtained in a single influential paper by Chistiano, Eichenbaum, and Evans, CEE, (1999). We then compute the corresponding bias-corrected IRF and associated bootstrap confidence intervals to draw our comparison.

In their paper, CEE examine the effects of monetary policy shocks on several economic variables of interest using models imposing a recursive structure to identify the relevant shocks. Their first benchmark model includes a constant term and four lags (*p=4*) of seven variables

[18] We considered confirming the bias with a SVAR Monte Carlo experiment but quickly realized that such an exercise would be trivial and reveal nothing new. As shown above, the only difference between the biased and unbiased estimates of the IRFs will be a constant scaling factor. Indeed, the computer code used in estimation would be identical in both cases except for this scaling.

(*K=7*) with the federal funds rate as the chosen monetary policy instrument. They estimate their models using quarterly data over the period 1965:3-1995:2. Given the loss of observations due to the four lags in the VAR, *T*=116 in our notation. We replicate their results by estimating their model over the same sample period.[19] For illustrative purposes, we report only the IRF indicating the effects of a negative monetary policy shock on output. While this is an IRF of particular interest, the same bias will be present in all the other 48 IRFs as well. As seen in Figure 3 here and Figure 2 of CEE (1999, p. 86), given a positive federal funds rate shock, "after a delay of 2 quarters, there is a sustained decline in real GDP " (p. 87). We note that CEE use MLE to estimate the VAR error covariance estimate so the estimated IRF will be biased. Furthermore, we see that the bootstrap confidence intervals reflect considerable asymmetry, which we shall see momentarily, is partially due to bias in the confidence intervals arising from biased bootstrap IRF estimates.

To illustrate the effect of bias due to MLE and the further bias due to the CEE bootstrap IRFs, we estimate the CEE model once again but this time including the appropriate degrees of freedom corrections. These results for the first-stage IRF and the bootstrap confidence intervals are also reported in Figure 3. The first thing we notice is that the fundamental conclusion regarding the IRF is unchanged: a contractionary federal funds rate shock will, after a lag, have a sustained negative effect on real GDP.[20] We also notice that, due to the degrees of freedom correction in the original error covariance matrix estimate, the bias-corrected IRF lies entirely below the CEE IRF.

In addition, we see that the confidence intervals also shift significantly when we correct the bias in the bootstrap estimates of the error covariance matrix. We note three consequences. First, we see that for much of the time horizon, the bias-corrected IRF actually lies below the biased CEE 95% confidence intervals. Second, we see that correcting for our bias

[19] Indeed, we have estimated the CEE model using their data and their RATS program which Larry Christiano has generously made available on his website.

[20] Indeed, we will always draw the same conclusion about statistical significance when our interest is in whether or not the IRF is significantly different from zero. This is a consequence of the fact, illustrated in the previous section, that the bias we are reporting is (negative and) *proportional* to the bootstrap IRFs. Accordingly, the biased and bias-corrected confidence interval bounds will cross the horizontal axis (zero line) at exactly the same horizons. This implies that the range over which the IRF is significantly greater or less than zero will be the same whether or not a bias correction is applied. Correcting the bias can lead to a reversal if the null hypothesis takes on a value other than zero.

has greatly reduced the asymmetry in the confidence intervals.[21] Third, we notice that between 2 and 11 quarters, the upper 95% confidence bounds are farther away from zero after bias correction. Thus, correcting the bias allows us to strengthen the conclusion that a contractionary monetary policy has a significant negative effect on output over that horizon.

Since part of the distortion in the CEE results is a consequence of their choice to use MLE estimates of the error covariance matrix, we also illustrate how much distortion remains when we use OLS estimates. The results are reported in Figure 4. In the typical approach incorporating the natural OLS degrees of freedom correction, the original IRF is already bias-corrected so we only have a single IRF estimate. However, the typical procedure does result in biased bootstrap confidence intervals. As in Figure 3, we again see that the typical biased procedure results in quite asymmetric confidence intervals which are, in part, a consequence of the bias; the bias-corrected confidence intervals exhibit much less asymmetry. Also, as noted in the discussion of Figure 3, over a range of intermediate horizons, the upper bound of the bias-corrected confidence intervals lie below their biased counterparts[22] giving us greater confidence in our conclusion that a monetary contraction has a significant negative effect on output.

These examples illustrate that correcting for bias in both the original IRF and especially in the bootstrap confidence intervals can remove distortions that change the quantitative (if not qualitative) conclusions when SVAR models are used.

5. Conclusion

This paper discussed a commonly occurring bias in bootstrap estimates of confidence intervals for IRFs in SVARs. The source of that bias is the downward bias in the traditional bootstrap estimate of the VAR covariance matrix. Since the bootstrap IRFs depend on these biased estimates, they are systematically biased as well. Consequently, the implied bootstrap IRF percentile confidence intervals inherit the same bias. This bias in potentially large but, fortunately, is easily corrected by accounting for the fact that the natural bootstrap estimate of

[21] This leads us to conjecture that the often puzzling asymmetry in IRF CIs found in the literature is largely due to the bias documented in this paper.

[22] As emphasized in the previous footnote, since the bias is negative and proportional, the bias-corrected confidence interval upper bound will lie below the biased confidence interval upper bound whenever the latter is negative. Furthermore, they will be zero at the same horizon.

the VAR covariance matrix must include an additional degrees of freedom adjustment.

References

Berkowitz, J., and Kilian, L. (2000), "Recent Developments in Bootstrapping Time Series," *Econometric Reviews*, 19, 1-48.

Blanchard, O., and Quah, D. (1989), "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review,* 79, 655-673.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (1999), "Monetary Policy Shocks: What Have We Learned and to What End?" in *Handbook of Macroeconomics,* eds. J. B. Taylor and M. Woodford, Amsterdam: Elsevier, 65-148.

Christiano, L. J., Eichebaum, M., and Vigfusson, R. (2006), "Assessing Structural VARs," in *NBER Macroeconomics Annual*., eds. D. Acemoglu, K. Rogoff, and M. Woodford, Cambridge, MA: MIT Press, 1-72.

Davidson, R., and MacKinnon, J. G. (1993), *Estimation and Inference in Econometrics*, New York: Oxford University Press.

Freedman, D. A. and Peters, S. C. (1984), "Bootstrapping a Regression Equation: Some Empirical Results" *Journal of the American Statistical Association,* 79, 97-206.

Galí, Jordi (1999), "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89, 249-271.

Inoue, A., Kilian, L. (2002), "Bootstrapping Smooth Functions of Slope Parameters and Innovation Variances in VAR(∞) Models," *International Economic Review*, 43, 309-331.

Peters, S. C., and Freedman, D. A. (1984), "Some Notes on the Bootstrap in Regression Problems," *Journal of Business and Economics Statistics*, 2, 406-409.

Runkle, D. E. (1987), "Vector Autoregression and Reality," *Journal of Business and Economics Statistics*, 5, 437-442.

Sims, C. A., and Zha, T. (1999), "Error Bands for Impulse Responses, *Econometrica*, 67, 1113-1155.

Stine, R. A. (1987), "Estimating Properties of Autoregressive Forecasts," *Journal of the American Statistical Association*, 82, 1072-1078.

Table 1: Bias of error variance estimate in standard univariate linear regression model with $R = 9$; number of Monte Carlo replications = 1,000, number of bootstrap draws = 1,000. True value of variance =0.81.

Variance estimator	**Sample Size**	**Theoretical bias (%)**	**Mean estimate**	**Bias (%)**
$\hat{\sigma}^2$	30	0	0.8207	1.33%
$\tilde{\sigma}^2_{b(1)}$	30	-30.0%	0.5745	-29.07%
$\tilde{\sigma}^2_{b(2)}$	30	0	0.8208	1.33%
$\hat{\sigma}^2$	50	0	0.8196	1.19%
$\tilde{\sigma}^2_{b(1)}$	50	-18.0 %	0.6720	-17.03%
$\tilde{\sigma}^2_{b(2)}$	50	0	0.8195	1.18%
$\hat{\sigma}^2$	100	0	0.8096	-0.05%
$\tilde{\sigma}^2_{b(1)}$	100	-9.0%	0.7686	-9.0%
$\tilde{\sigma}^2_{b(2)}$	100	0	0.8097	-0.04%

Table 2: Bias of error variance estimate in an AR(8) model with a constant term ($R = 9$); number of Monte Carlo replications = 1,000, number of bootstrap draws = 1,000. True value of variance =0.81.

Variance estimator	Sample Size	"Theoretical" bias (%)[a]	Mean estimate	Bias (%)
$\hat{\sigma}^2$	30	0	0.8323	2.753%
$\tilde{\sigma}^2_{b(1)}$	30	-30.0%	0.5989	-26.06%
$\tilde{\sigma}^2_{b(2)}$	30	0	0.8556	5.63%
$\hat{\sigma}^2$	50	0	0.8316	2.67%
$\tilde{\sigma}^2_{b(1)}$	50	-18.0 %	0.6891	-14.93%
$\tilde{\sigma}^2_{b(2)}$	50	0	0.8404	3.75%
$\hat{\sigma}^2$	100	0	0.8145	0.55%
$\tilde{\sigma}^2_{b(1)}$	100	-9.0%	0.7431	-8.26%
$\tilde{\sigma}^2_{b(2)}$	100	0	0.8166	0.81%

[a] This is the theoretical bias for the corresponding (R=9) standard regression model.

<u>Figure 1</u>: A reestimated version of Figure 3 in Blanchard and Quah (1989). It shows the response of output to aggregate demand shocks with asymmetric one standard deviation bands based on *biased* bootstrap estimates.

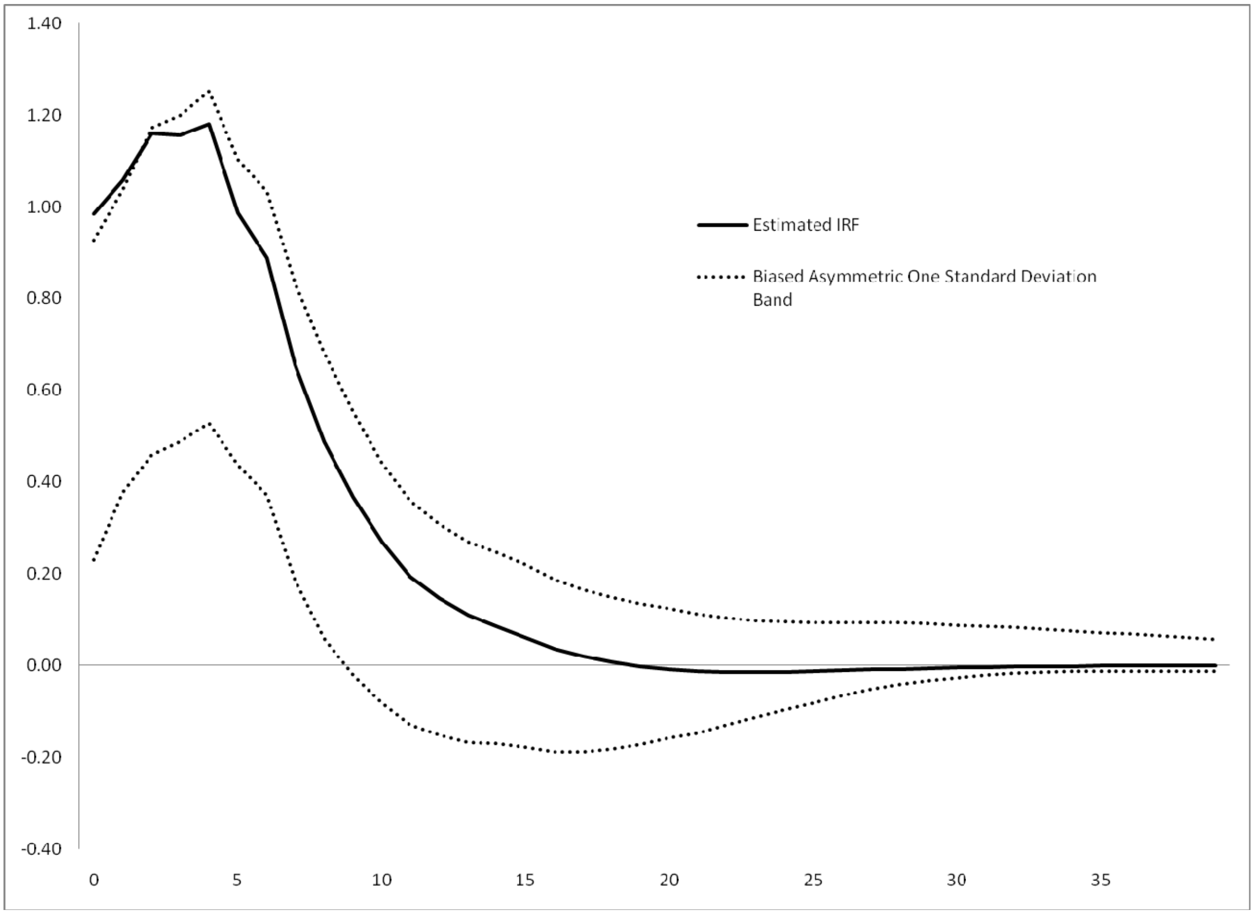

Figure 2: A reestimated version of Figure 3 in Blanchard and Quah (1989) with asymmetric one standard deviation bands based on *biased-corrected* bootstrap estimates.

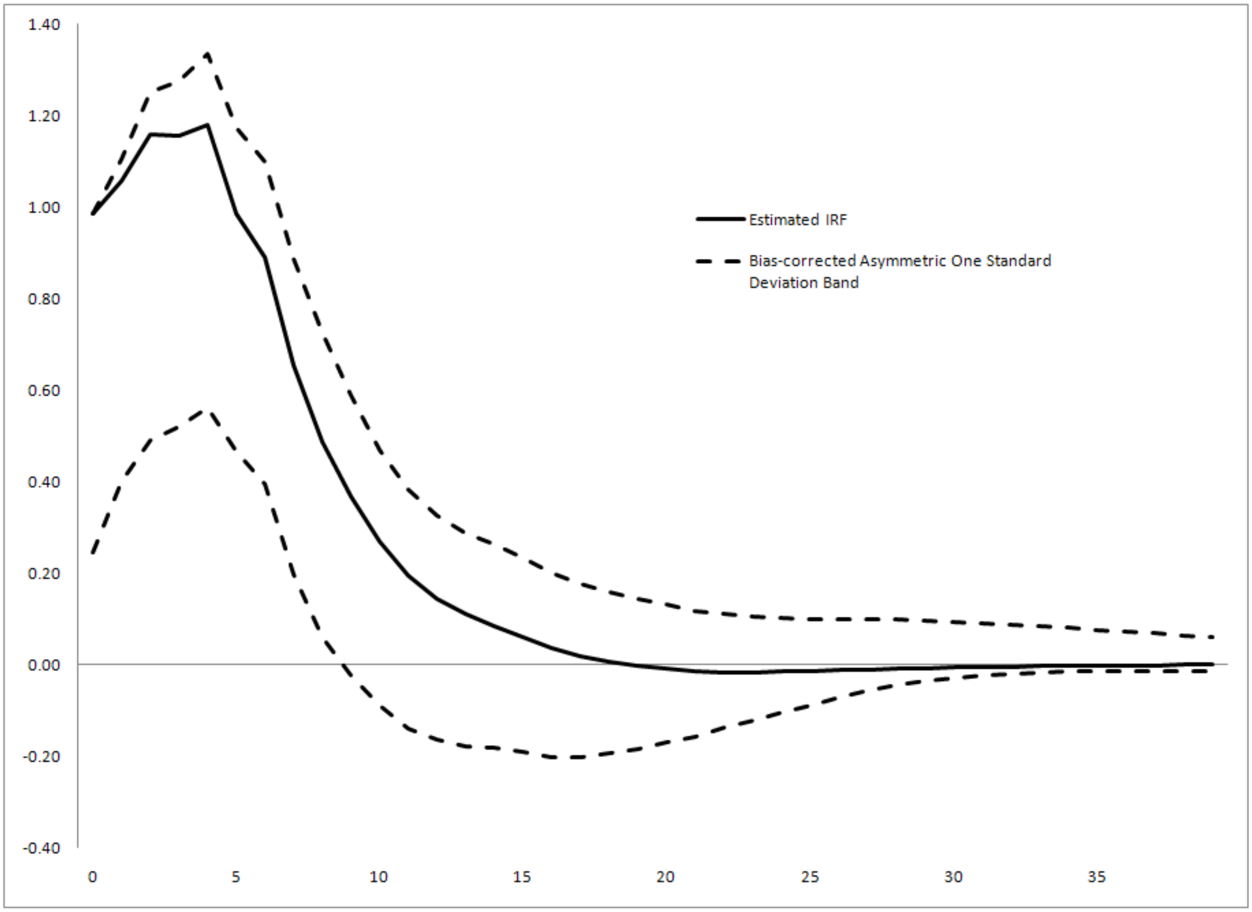

Figure 3: Impulse response functions showing the effect of a contractionary monetary policy on real GDP with 95% confidence intervals. The solid line gives the original MLE IRF and the long-dashed bold line gives the bias-corrected OLS IRF; CEE use MLE. The dotted lines give the MLE bootstrap 95% confidence intervals and the dashed lines give the bias-corrected 95% confidence intervals.

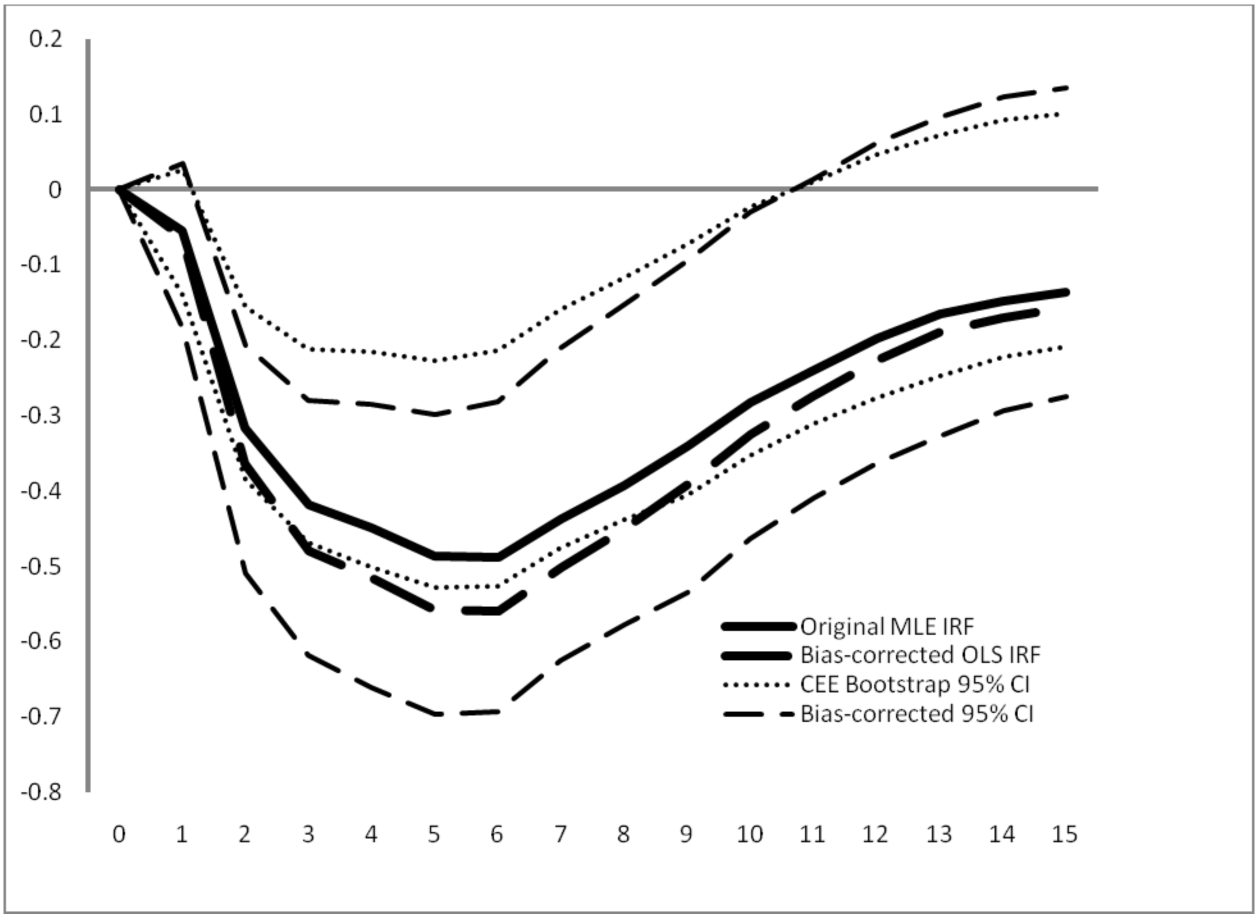

Figure 4: Impulse response function showing the effect of a contractionary monetary policy on real GDP with 95% confidence intervals. The solid line gives the original OLS IRF. The dotted lines give the typical but biased bootstrap 95% confidence intervals and the dashed lines give the bias corrected 95% confidence intervals.