

School of Economics Working Paper
2020-06



**SCHOOL OF
ECONOMICS**

An Alternative Bootstrap for Proxy Vector Autoregressions

Martin Bruns *

Helmut Luetkepohl **

*University of East Anglia

**DIW Berlin and Freie Universitaet Berlin

School of Economics
University of East Anglia
Norwich Research Park
Norwich NR4 7TJ
United Kingdom
www.uea.ac.uk/economics

An Alternative Bootstrap for Proxy Vector Autoregressions

Martin Bruns¹

University of East Anglia, School of Economics,
Norwich Research Park, NR4 7TJ, Norwich, United Kingdom
email: martin.j.bruns@gmail.com

and

Helmut Lütkepohl

DIW Berlin and Freie Universität Berlin, Mohrenstr. 58, 10117 Berlin,
Germany
email: hluetkepohl@diw.de

November 9, 2020

Abstract. We propose a new bootstrap for inference for impulse responses in structural vector autoregressive models identified with an external proxy variable. Simulations show that the new bootstrap provides confidence intervals for impulse responses which often have more precise coverage than and similar length as the competing moving-block bootstrap intervals. An empirical example shows how the new bootstrap can be applied in the context of identifying monetary policy shocks.

Key Words: Bootstrap inference, structural vector autoregression, impulse responses, instrumental variable

JEL classification: C32

¹The authors thank Carsten Jentsch for helpful comments on an earlier version of the paper.

1 Introduction

In structural vector autoregressive (VAR) analysis one strand of the literature uses external instruments, also called proxies, to identify shocks of interest (e.g., Stock and Watson (2012), Mertens and Ravn (2013), Piffer and Podstawski (2018), Kilian and Lütkepohl (2017, Chapter 15)). The related models and methods are often abbreviated as proxy VARs. In this context, frequentist inference for impulse responses is typically based on bootstrap methods. In some of the literature, wild bootstraps (WB) are used (e.g., Mertens and Ravn (2013), Gertler and Karadi (2015), Carriero, Mumtaz, Theodoridis and Theophilopoulou (2015)). However, based on work by Brüggemann, Jentsch and Trenkler (2016), Jentsch and Lunsford (2019) show that wild bootstrap methods are not asymptotically valid in this context and they propose a moving-block bootstrap (MBB) which provides asymptotically correct confidence intervals for impulse responses under very general conditions. It can cope, for example, with conditional heteroskedastic (ARCH) VAR errors which is an advantage in many applied studies where financial data are of interest. On the other hand, Lütkepohl and Schlaak (2019) demonstrate by simulations that the MBB can result in confidence intervals with low coverage rates in small samples.

In this study, we propose an alternative bootstrap method for proxy VARs which is based on resampling not only the VAR residuals but also the residuals of a model for the proxy and is therefore signified as RBB (residual-based bootstrap). We show by simulation that it leads to quite precise confidence intervals for impulse responses in small samples which makes it attractive for macroeconomic analysis where often smaller samples with less than 200 observations are available only. In contrast to the MBB, it cannot cope with ARCH features in the data and it is based on a specific assumption on the generation mechanism of the proxy variable. Its advantage is, however, that it yields more precise inference in small samples.

The remainder of the paper is structured as follows. The proxy VAR model is presented in the next section. Estimation of proxy VAR models is considered in Section 3. The alternative bootstrap methods considered in this study are presented in Section 4 and a small sample Monte Carlo comparison of the bootstraps is discussed in Section 5. An illustrative example is given in Section 6 and conclusions are provided in Section 7.

2 The Proxy VAR Model

A K -dimensional reduced-form VAR process,

$$y_t = \nu + A_1 y_{t-1} + \cdots + A_p y_{t-p} + u_t, \quad (2.1)$$

is considered. Here ν is a $(K \times 1)$ constant term and the A_i , $i = 1, \dots, p$, are $(K \times K)$ slope coefficient matrices. The reduced-form error, u_t , is a zero mean white noise process with covariance matrix Σ_u , i.e., $u_t \sim (0, \Sigma_u)$. The vector of structural errors, $w_t = (w_{1t}, \dots, w_{Kt})'$, is such that $u_t = Bw_t$, where B is the nonsingular $(K \times K)$ matrix of impact effects of the shocks on the observed variables y_t . Thus, $w_t \sim (0, \Sigma_w = B^{-1}\Sigma_u B^{-1'})$, where Σ_w is a diagonal matrix.

If the first column, say b , of B is known, the structural impulse responses of the first shock, $\theta_i = (\theta_{11,i}, \dots, \theta_{K1,i})'$, can be computed as

$$\theta_i = (\theta_{11,i}, \dots, \theta_{K1,i})' = \Phi_i b, \quad i = 1, \dots, H,$$

where the $(K \times K)$ matrices $\Phi_i = \sum_{j=1}^i \Phi_{i-j} A_j$ can be obtained recursively from the VAR slope coefficients using $\Phi_0 = I_K$ (e.g., Lütkepohl (2005, Chapter 2)). In the following, the $(K \times (H + 1))$ matrix of impulse responses,

$$\Theta(H) = [\theta_0, \theta_1, \dots, \theta_H] = [b, \Phi_1 b, \dots, \Phi_H b], \quad (2.2)$$

is of interest. It is assumed that the first shock increases the first variable by one unit on impact. In other words, the first component of $b = \theta_0$ is assumed to be 1.

If b , Σ_u and the reduced-form errors are given, the first structural shock can be obtained as

$$w_{1t} = b' \Sigma_u^{-1} u_t / b' \Sigma_u^{-1} b \quad (2.3)$$

(see Stock and Watson (2018, Footnote 6, p. 933)).

Suppose there is an instrumental variable z_t satisfying

$$\mathbb{E}(w_{1t} z_t) = c \neq 0 \quad (\text{relevance}), \quad (2.4)$$

$$\mathbb{E}(w_{kt} z_t) = 0, \quad k = 2, \dots, K, \quad (\text{exogeneity}). \quad (2.5)$$

These conditions imply that

$$\mathbb{E}(z_t u_t) = B \mathbb{E}(z_t w_t) = cb.$$

In other words, the proxy z_t identifies a multiple of b .

In line with some of the proxy VAR literature, the proxy z_t is assumed to be generated as

$$z_t = D_t(\phi w_{1t} + \eta_t), \quad (2.6)$$

where D_t is a random 0-1 variable which determines the number of nonzero values of the proxy. It is assumed to have a Bernoulli distribution, $B(d)$, with parameter d , $0 < d \leq 1$, and captures the fact that many proxies are measured only at certain announcement days or when special events occur (see, e.g., Jentsch and Lunsford (2019)). It is assumed to be stochastically independent of w_{1t} and the error term η_t which is thought of as representing measurement error. This error term is assumed to have mean zero and variance σ_η^2 , i.e., $\eta_t \sim (0, \sigma_\eta^2)$, and it is distributed independently of w_{1t} . The parameter ϕ , the error η_t and the distribution of the Bernoulli random variable D_t (i.e., the parameter d) determine the strength of the correlation between z_t and w_{1t} and, hence, the strength of the proxy as an instrument.

The variance of z_t is $\text{Var}(z_t) = d(\phi^2 \text{Var}(w_{1t}) + \sigma_\eta^2)$ for $0 < d \leq 1$. Moreover, the covariance between w_{1t} and z_t is

$$\mathbb{E}(w_{1t}z_t) = \mathbb{E}(w_{1t}D_t(\phi w_{1t} + \eta_t)) = \phi d \text{Var}(w_{1t})$$

so that the correlation between w_{1t} and z_t is

$$\text{Corr}(w_{1t}, z_t) = \phi \sqrt{d} \sqrt{\text{Var}(w_{1t})} / \sqrt{\phi^2 \text{Var}(w_{1t}) + \sigma_\eta^2}.$$

Thus, the correlation between the proxy and the first shock declines with declining d and increasing σ_η^2 .

3 Estimation

Suppose an effective sample y_1, \dots, y_T of size T is available for the model variables, plus all required presample values, y_{-p+1}, \dots, y_0 . Moreover, a corresponding sample z_1, \dots, z_T is available for the proxy.

Then the VAR(p) is estimated by bias-adjusted least squares (LS) giving estimates $\hat{\nu}$, $\hat{A}_1, \dots, \hat{A}_p$, residuals $\hat{u}_1, \dots, \hat{u}_T$ and an error covariance matrix estimator

$$\hat{\Sigma}_u = \frac{1}{T} \sum_{t=1}^T \hat{u}_t \hat{u}_t'$$

based on mean-adjusted residuals. Kilian (1998) shows that employing bias-adjusted LS estimators improves inference for impulse responses. Therefore

we use the bias-adjustment based on Pope (1990), as proposed by Kilian (1998), throughout the paper.

The first column b of B is estimated using the proxy z_t ,

$$\hat{b} = \sum_{t=1}^T \hat{u}_t z_t \bigg/ \sum_{t=1}^T \hat{u}_{1t} z_t, \quad (3.1)$$

where \hat{u}_t are the residuals corresponding to bias-adjusted LS estimation. The impulse response matrix $\Theta(H)$ is estimated as

$$\hat{\Theta}(H) = [\hat{b}, \hat{\Phi}_1 \hat{b}, \dots, \hat{\Phi}_H \hat{b}],$$

where

$$\hat{\Phi}_i = \sum_{j=1}^i \hat{\Phi}_{i-j} \hat{A}_j, \quad i = 1, \dots, H.$$

Moreover, the first shock is estimated as

$$\hat{w}_{1t} = \hat{b}' \hat{\Sigma}_u^{-1} \hat{u}_t / \hat{b}' \hat{\Sigma}_u^{-1} \hat{b}, \quad t = 1, \dots, T, \quad (3.2)$$

and ϕ is estimated by LS from

$$z_t = \mu + \phi \hat{w}_{1t} + \eta_t$$

for all $t \in \mathcal{T}_D$, where $\mathcal{T}_D = \{t | D_t = 1\}$. The estimate of ϕ is denoted by $\hat{\phi}$ and the residuals are $\hat{\eta}_t$ for $t \in \mathcal{T}_D$ and $\hat{\eta}_t = 0$ for $t \notin \mathcal{T}_D$.

4 Bootstraps

As mentioned in the introduction, the WB and the MBB are the bootstrap methods most frequently used in the proxy VAR literature for frequentist inference for impulse responses. The WB generates asymptotically invalid confidence intervals while the MBB yields confidence intervals with the correct coverage level asymptotically under quite general conditions (Jentsch and Lunsford (2019)). It may be imprecise in small samples, however, and therefore we propose the RBB which turns out to have better properties in small samples. The three bootstrap versions differ in the way they generate bootstrap samples of y_t and z_t . Based on N bootstrap samples $y_{-p+1}^{(n)}, \dots, y_0^{(n)}, y_1^{(n)}, \dots, y_T^{(n)}$ and $z_1^{(n)}, \dots, z_T^{(n)}$, $n = 1, \dots, N$, they all use the following steps to determine bootstrap impulse responses and confidence intervals:

1. A VAR(p) model is fitted to the sample by LS and bias-adjusted, giving bootstrap estimates $\hat{A}^{(n)}$,

$$\hat{\Phi}_i^{(n)} = \sum_{j=1}^i \hat{\Phi}_{i-j}^{(n)} \hat{A}_j^{(n)}, \quad i = 1, \dots, H, \quad \text{with } \hat{\Phi}_0^{(n)} = I_K,$$

and residuals $\hat{u}_t^{(n)}$.

2. Then bootstrap estimates

$$\hat{b}^{(n)} = \sum_{t=1}^T \hat{u}_t^{(n)} z_t^{(n)} \bigg/ \sum_{t=1}^T \hat{u}_{1t}^{(n)} z_t^{(n)}.$$

of the structural parameters are determined.

3. Finally bootstrap estimates of the impulse responses of interest are computed as

$$\hat{\Theta}(H)^{(n)} = [\hat{b}^{(n)}, \hat{\Phi}_1^{(n)} \hat{b}^{(n)}, \dots, \hat{\Phi}_H^{(n)} \hat{b}^{(n)}]$$

and stored.

The N bootstrap estimates $\hat{\Theta}(H)^{(1)}, \dots, \hat{\Theta}(H)^{(N)}$ are used to construct pointwise confidence intervals based on the relevant quantiles of the bootstrap distributions. Alternatively, percentile- t or Hall intervals could be used (see Kilian and Lütkepohl (2017, Section 12.2)). However, the intervals based on quantiles are quite common in practice and the relative performance of the alternative bootstrap versions is not expected to depend on the type of interval used.

The samples are generated by one of the three alternative bootstrap methods, WB, MBB and RBB, as follows:

WB: For $t = 1, \dots, T$, independent standard normal variates ψ_t , $\psi_t \sim \mathcal{N}(0, 1)$, are drawn and bootstrap residuals and proxy variables are generated as

$$\begin{pmatrix} u_t^{WB} \\ z_t^{WB} \end{pmatrix} = \psi_t \begin{pmatrix} \hat{u}_t \\ z_t \end{pmatrix}.$$

The u_t^{WB} are de-meant and multiplied by $\sqrt{T/(T - Kp - 1)}$, as in Davidson and MacKinnon (2004, p. 597), and they are used to generate $y_t^{WB} = \hat{\nu} + \hat{A}_1 y_{t-1}^{WB} + \dots + \hat{A}_p y_{t-p}^{WB} + u_t^{WB}$, $t = 1, \dots, T$, starting from $y_{-p+1}^{WB}, \dots, y_0^{WB}$, which are obtained as a random draw of p consecutive values from the original sample.

MBB: A block length $\ell < T$ has to be chosen for the MBB. The blocks of length ℓ of the estimated residuals and proxies are arranged in the form of the matrix

$$\begin{bmatrix} \begin{pmatrix} \hat{u}_1 \\ z_1 \end{pmatrix} & \begin{pmatrix} \hat{u}_2 \\ z_2 \end{pmatrix} & \cdots & \begin{pmatrix} \hat{u}_\ell \\ z_\ell \end{pmatrix} \\ \begin{pmatrix} \hat{u}_2 \\ z_2 \end{pmatrix} & \begin{pmatrix} \hat{u}_3 \\ z_3 \end{pmatrix} & \cdots & \begin{pmatrix} \hat{u}_{1+\ell} \\ z_{1+\ell} \end{pmatrix} \\ \vdots & \vdots & & \vdots \\ \begin{pmatrix} \hat{u}_{T-\ell+1} \\ z_{T-\ell+1} \end{pmatrix} & \begin{pmatrix} \hat{u}_{T-\ell+2} \\ z_{T-\ell+2} \end{pmatrix} & \cdots & \begin{pmatrix} \hat{u}_T \\ z_T \end{pmatrix} \end{bmatrix}.$$

The bootstrap residuals and proxy are recentered columnwise by constructing

$$\tilde{u}_{j\ell+i} = \hat{u}_{j\ell+i} - \frac{1}{T-\ell+1} \sum_{r=0}^{T-\ell} \hat{u}_{i+r}$$

and

$$\tilde{z}_{j\ell+i} = z_{j\ell+i} - \frac{1}{T-\ell+1} \sum_{r=0}^{T-\ell} z_{i+r}$$

for $i = 1, 2, \dots, \ell$ and $j = 0, 1, \dots, s-1$. Then $s = \lceil T/\ell \rceil$ of the recentered rows of the matrix are drawn with replacement, where $\lceil \cdot \rceil$ denotes the smallest number greater than or equal to the argument such that $\ell s \geq T$. These randomly drawn blocks are joined end-to-end and the first T bootstrap residuals and proxies are retained,

$$\begin{pmatrix} u_t^{MBB} \\ z_t^{MBB} \end{pmatrix}, \quad t = 1, \dots, T.$$

Finally, the u_t^{MBB} are de-means, multiplied by $\sqrt{T/(T-Kp-1)}$, and used to generate $y_t^{MBB} = \hat{\nu} + \hat{A}_1 y_{t-1}^{MBB} + \cdots + \hat{A}_p y_{t-p}^{MBB} + u_t^{MBB}$, $t = 1, \dots, T$, starting from $y_{-p+1}^{MBB}, \dots, y_0^{MBB}$, which are obtained as a random draw of p consecutive values from the original sample.

RBB: Samples

$$\begin{pmatrix} u_t^{RBB} \\ \eta_t^{RBB} \\ w_{1t}^{RBB} \end{pmatrix}, \quad t = 1, \dots, T, \text{ are drawn from } \begin{pmatrix} \hat{u}_1 \\ \hat{\eta}_1 \\ \hat{w}_{11} \end{pmatrix}, \dots, \begin{pmatrix} \hat{u}_T \\ \hat{\eta}_T \\ \hat{w}_{1T} \end{pmatrix},$$

with replacement. Bootstrap samples $y_t^{RBB} = \hat{\nu} + \hat{A}_1 y_{t-1}^{RBB} + \dots + \hat{A}_p y_{t-p}^{RBB} + u_t^{RBB}$, $t = 1, \dots, T$, are generated, starting from $y_{-p+1}^{RBB}, \dots, y_0^{RBB}$, which are obtained as a random draw of p consecutive values from the original sample. Samples of the proxy are generated as

$$z_t^{RBB} = \tilde{D}_t(\hat{\phi} w_{1t}^{RBB} + \eta_t^{RBB}),$$

where \tilde{D}_t is a random 0-1 variable following a Bernoulli distribution, $B(\hat{d})$, with \hat{d} being the share of non-zero observations in the original sample.

We emphasize again that the WB does not result in asymptotically valid confidence intervals but is presented here and included in the simulation comparison in Section 5 because it has been used in the proxy VAR literature. The WB and the MBB draw the proxies directly from observed values and, hence, do not make assumptions on the exact DGP of z_t . In contrast, the RBB samples from the residuals of the assumed DGP for z_t in (2.6) and constructs new proxy values in each bootstrap replication. In addition, while the WB design sets the share of non-zero observations of the bootstrap proxy equal to the share in the original sample, this is not generally the case in the MBB and the RBB design. Apart from that, all three bootstraps are recursive-design residual based bootstraps for generating the y_t samples.

In all three bootstrap algorithms, the initial values $y_{-p+1}^{(n)}, \dots, y_0^{(n)}$ are a random draw of p consecutive values from the original y_t sample. Alternatively, the original initial values y_{-p+1}, \dots, y_0 could have been used as initial values for each bootstrap sample. If the y_t are mean-adjusted, one could even simply use zero initial values if stationary models are under consideration. For example, Jentsch and Lunsford (2019) used zero initial values, generate more than T sample values and then drop some burn-in values.

For the MBB a decision on the block length ℓ is needed. To make the asymptotic theory work, it has to be chosen such that $\ell \rightarrow \infty$ and $\ell^3/T \rightarrow 0$ as $T \rightarrow \infty$ (see Jentsch and Lunsford (2019)). The choice is less clear in small samples. Choosing ℓ too small, the blocks may not capture the data features well and may result in poor confidence intervals. On the other hand, choosing ℓ large undermines inference precision because there are too few blocks to choose from. Note that the number of available blocks is $T - \ell + 1$ and, hence, depends on the block length. Jentsch and Lunsford (2019) mention a block length of $\ell = 5.03T^{1/4}$ as a rule of thumb and we use this rule of thumb in our simulations in Section 5 and in the empirical example in Section 6.

5 Small Sample Comparison of Bootstraps

In this section a small sample simulation comparison of the three bootstrap methods is presented. The simulation design is considered first and then the simulation results are discussed.

5.1 Monte Carlo Design

5.1.1 DGP 1

The first data generating process (DGP1) is similar to a DGP that has been used frequently in related work on comparing inference methods for impulse responses (e.g., Kilian (1998), Kilian and Kim (2011), Lütkepohl, Staszewska-Bystrova and Winker (2015a, 2015b)). It is a two-dimensional VAR(1) of the form:

$$y_t = \begin{bmatrix} a_{11} & 0 \\ 0.5 & 0.5 \end{bmatrix} y_{t-1} + u_t, \quad (5.1)$$

where $0 < a_{11} < 1$. The process is stable with more persistence for a_{11} closer to one.

The structural errors, w_t , are normally distributed with mean zero and identity covariance matrix, $w_t \sim \mathcal{N}(0, I_2)$, and $u_t = Bw_t$ with

$$B = \begin{bmatrix} 2 & 0 \\ 1 & 3 \end{bmatrix},$$

such that $b = (1, \frac{1}{2})'$. These u_t errors are used to generate the y_t as in equation (5.1), starting from a standard normal y_0 , i.e., $y_0 \sim \mathcal{N}(0, I_2)$. In the simulations, we fit VAR models of order $p = 1$ and $p = 12$, without constant term, to de-meaned data.

In line with the related literature (e.g., Jentsch and Lunsford (2019)), the proxy z_t is generated as in equation (2.6), i.e., $z_t = D_t(\phi w_{1t} + \eta_t)$, where D_t , ϕ and the error η_t determine the strength of the correlation between z_t and w_{1t} and, hence, the strength of the proxy which is important for how well the impact effects of the shock can be estimated and these estimates are of central importance for estimating the impulse responses. The error term η_t is generated independently of w_{1t} as $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$, with different values of σ_η^2 . The random variable D_t has a Bernoulli distribution with parameter d , $B(d)$, which specifies the average proportion of nonzero z_t variables. D_t is stochastically independent of η_t and w_{1t} . For $d = 1$, the proxy variable is nonzero with probability one for all sample periods $t = 1, \dots, T$.

Table 1: Design Parameters for DGP1

T	p	a_{11}	B	Σ_w	d	ϕ	σ_η^2	$\text{Corr}(w_{1t}, z_t)$	H
100	12	0.5	$\begin{bmatrix} 2 & 0 \\ 1 & 3 \end{bmatrix}$	I_2	1	1	0.2346	0.9	20
250		0.9						0.5	
500		0.95						0.5692	
								0.3162	

The parameter values used in our simulations are summarized in Table 1. We use propagation horizons of $H = 20$ to capture not only the short-term effects of a shock but also the longer-term effects which may still be a bit away from zero for the more persistent processes. Sample sizes $T = 100, 250$ and 500 are considered. The number of replications for each Monte Carlo design is $R = 1000$ and we use $N = 2000$ bootstrap repetitions within each replication.

Our criteria for evaluating the bootstrap methods are the coverage precision and the widths of the confidence intervals obtained from the bootstraps. These criteria capture main features of interest in related empirical studies and they have also been used in related small sample comparisons of bootstrap inference (e.g., Kilian and Kim (2011), Lütkepohl and Schlaak (2019)).

5.1.2 DGP 2

Our second DGP (DGP2) mimics a VAR model from a study of Gertler and Karadi (2015). It is based on parameters estimated from their dataset. One of the models used by Gertler and Karadi is a four-dimensional US monthly model. We use their data from 1990M1 to 2016M6 and fit a VAR(1) model with constant term to the data. Using bias-adjusted estimates, the reduced-form parameters of DGP2 are $\nu = 0$,

$$A_1 = \begin{bmatrix} 0.97 & 0.00 & 0.00 & -0.13 \\ 0.01 & 1.00 & 0.00 & -0.09 \\ -0.03 & 0.00 & 1.00 & -0.53 \\ 0.02 & 0.00 & 0.00 & 0.91 \end{bmatrix} \text{ and } \Sigma_u = \begin{bmatrix} 0.04 & 0.00 & 0.01 & -0.01 \\ 0.00 & 0.11 & 0.03 & 0.00 \\ 0.01 & 0.03 & 0.37 & -0.01 \\ -0.01 & 0.00 & -0.01 & 0.07 \end{bmatrix}.$$

The maximum eigenvalue of A_1 has modulus 0.9997. Thus, DGP2 is stable but very persistent. These parameters are used to generate the y_t based on $u_t \sim \mathcal{N}(0, \Sigma_u)$ and starting from $y_0 = 0$, the unconditional mean of the y_t .

We also use a proxy with similar properties as the proxy for monetary policy shocks constructed by Gertler and Karadi (2015). More precisely, we estimate the b vector of impact effects of the first shock, giving a vector

$b = (1, -0.14, 0.70, 0.24)'$, and estimate the parameters ϕ and σ_η^2 of the model (2.6) as described in Section 3 from the Gertler/Karadi data with nonzero z_t values and the first shock obtained from equation (2.3). This yields values $\phi = 0.1019$ and $\sigma_\eta^2 = 0.0020$ that are used for generating z_t as in equation (2.6) with D_t having a $B(0.82)$ distribution. The parameter, $d = 0.82$, of the Bernoulli distribution is chosen because the Gertler-Karadi proxy has nonzero values for 82% of the sample periods. The implied correlation between proxy and shock is 0.36 and, hence, it is rather low.

Note that the generation mechanism for DGP2 differs from that of DGP1, where the structural shocks are generated directly and the reduced-form data as well as the proxy are computed from the generated structural shocks and the generated η_t series. In contrast, we generate the reduced-form errors for DGP2, construct the first structural shock from the structural parameters b and the error covariance matrix Σ_u as in equation (3.2) and then generate z_t as in equation (2.6) with an additionally generated $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$.

The rationale behind using this DGP2 is that we will also use the Gertler/Karadi data for an illustrative example in Section 6 and, hence, the simulation results for DGP2 may be indicative of what to expect in the example. Moreover, it is, of course, of interest to see whether the results for our small bivariate process underlying DGP1 carry over to a higher-dimensional DGP.

For DGP2, we fit VAR models of orders $p = 1$ and $p = 12$, including a constant term, to samples of size $T = 200$ and 500 . The smallest sample size considered is a bit larger than for DGP1 to account for the larger model dimension. It is not far from the sample size used in the example in Section 6. The number of bootstrap replications is again $N = 2000$ and the number of Monte Carlo repetitions is $R = 1000$, as for DGP1.

5.2 Small Sample Results

5.2.1 Results for DGP1

Some main findings from simulating DGP1 are presented in Figures 1 and 2. Specifically, in Figure 1 the implications of changing the VAR order p , the persistence of the VAR process (a_{11}) and the strength of the proxy reflected in the Bernoulli parameter d on the coverage and average length of the bootstrap confidence intervals can be seen for relatively short samples of size $T = 100$. In Figure 2, the impact of increasing the sample size is presented.

A main observation from Figure 1 is that, for some designs and propagation horizons, there are clear differences in the coverage of the confidence intervals of the three bootstrap variants. The RBB yields overall the coverage results closest to the desired 90% while the MBB tends to yield coverage

rates a bit smaller and, in particular for short propagation horizons, the WB yields intervals with partly substantially larger coverage rates. These differences in coverage rates, in some cases, translate into shorter intervals for the MBB and longer intervals for WB in comparison to the RBB. Interestingly, the average lengths of the intervals of all three bootstrap methods are often very similar despite substantial differences in coverage. For example, the RBB intervals for short propagation horizons are in most cases very similar to the MBB intervals although the latter have a coverage which is below the RBB coverage and it is also lower than the nominal 90%. Thus, Figure 1 clearly shows that the RBB tends to be more precise in terms of coverage and often it does so without sacrificing much interval length. Thus, under these two criteria it is preferable to the MBB and the WB. The latter bootstrap is often conservative and also yields larger confidence intervals than the other two bootstrap methods.

There are also some more specific results related to the VAR order and the proxy strength that can be seen in Figure 1.

- The VAR order p has an important impact on both the coverage and average lengths of the confidence intervals. In particular, considering the order $p = 12$ of a short-memory process with $a_{11} = 0.5$ results in substantial over-coverage, especially for longer propagation horizons, for all three bootstrap variants. The interval lengths tend to increase for all propagation horizons and substantially so for the longer propagation horizons for all three bootstraps if the VAR order increases from $p = 1$ to $p = 12$ (compare the second and fourth columns of Figure 1).
- Comparing panels (c) and (e) as well as (d) and (f) in Figure 1, it is apparent that the proxy strength does not have much of an effect on the coverage but partly leads to larger intervals (see in particular the average lengths of θ_{21} intervals for short horizons). In panels (e) and (f) in Figure 1, the proxy has a lower correlation with the structural shock of interest due to the reduced number of event dates, d , for which the proxy is constructed. A similar result is obtained, however, if $d = 1$ is maintained but the correlation between proxy and shock is reduced due to a larger variance σ_η^2 of the error term in equation (2.6), as can be seen in Figure A.1 in the Appendix.
- The impact of higher persistence (a larger a_{11} parameter) of the process can be seen by comparing panels (a) and (c) as well as (b) and (d) in Figure 1. Generally the coverage is reduced and the intervals become larger, especially for longer propagation horizons, if a_{11} increases from 0.5 to 0.95. The reduction in coverage is most severe for the MBB, while

the RBB continues to have acceptable coverage for persistent processes. In Figure A.2 of the Appendix, additional results for $a_{11} = 0.9$ are presented and it can be seen that the results for $a_{11} = 0.9$ are similar to those of $a_{11} = 0.95$.

In Figure 2, the impact of the sample size on the confidence intervals is exhibited for the case of a persistent process with $a_{11} = 0.95$ and a relatively strong proxy with correlation 0.9 with the shock and $d = 1$. As we saw in Figure 1 already, in this situation the MBB has a coverage clearly smaller than the nominal 90% for $p = 1$ and all three bootstrap methods tend to yield under-coverage for $T = 100$. In Figure 2 it can be seen that the coverage clearly improves for $T = 250$ already and the coverage deficiencies largely disappear for $T = 500$. Also, the interval lengths for all three methods become very similar and are reduced for larger sample sizes, as one would expect. Only the WB intervals for some short propagation horizons remain wider and less precise for larger samples. This result may be a reflection of the asymptotic inadmissibility of the WB.

5.2.2 Results for DGP2

Coverage and average interval lengths for DGP2 are depicted in Figure 3 for sample size $T = 200$ and in Figure 4 for $T = 500$. Even for the smaller sample size $T = 200$, all coverage rates of the nominal 90% confidence intervals of WB and RBB are between 80% and 100%, except for the long-run response of the third variable. In other words, the two bootstrap methods yield rather precise confidence intervals for three out of four variables across our Monte Carlo designs. Given the asymptotic inadmissibility of the WB, this result may, of course, not be generalizable to other simulation designs.

Even the MBB has coverage rates above 80% for variables 1, 2 and 4 and propagation horizons up to 30 periods when $T = 200$. Thus, even the MBB is relatively precise in terms of coverage for three out of four variables. There are, however, differences in interval lengths among the three bootstraps. Typically, the WB intervals are a bit longer than the MBB and RBB intervals which are often close together on average. Overall the performance of the WB is inferior to MBB and RBB. Thus, although there is often not much to choose between MBB and RBB in terms of coverage and interval length, it is remarkable that the RBB tends to have typically coverage rates closer to 90% than the MBB. Thus, even for the higher-dimensional DGP2, RBB performs well relative to its competitors, at least for three of the four variables.

For the third variable, VAR order $p = 1$ and $T = 200$ all three bootstrap methods yield coverage rates below 80% for a propagation horizon of 48

periods. For $T = 500$ and $p = 1$, only MBB still has a coverage below 80% for long horizons (see Figure 4). The coverage rates for long horizons are actually a bit closer to 90% for $p = 12$, although one might expect lower precision for the larger VAR order as the larger VAR order implies a model with substantially more parameters. Even for $p = 12$, the RBB outperforms the MBB in terms of coverage and is almost as good in terms of interval length.

In summary, our simulations show that the WB is often conservative and yields more than the nominal coverage. In turn, its confidence intervals are often considerably larger than those of the MBB and RBB. Thus, the WB is overall inferior to the MBB and the RBB. Between the latter two, the RBB is preferable because it yields typically coverage rates closer to the nominal rate than the MBB. Moreover, the RBB confidence intervals are often about as long on average as those of the MBB. Hence, our simulations show that the RBB has merit. In the next section, it will be applied to an illustrative example model from the literature.

6 Empirical Example

We consider an example based on the study of Gertler and Karadi (2015) mentioned earlier to illustrate the differences between the three bootstrap methods. One of the models used by Gertler and Karadi is a four-dimensional US monthly model for the variables (1) one-year government bond rate, (2) log consumer price index (CPI), (3) log industrial production (IP) and (4) excess bond premium. They employ the three months ahead federal funds rate future surprises as the baseline proxy to identify a monetary policy shock and they find their proxy to be a strong instrument. We re-estimate their model, shortening the sample to include only periods for which all four variables and the proxy are available. This leaves us with a sample running from 1990M1 through 2016M6, i.e., the sample size is $T = 270$. The proxy is available in $d = 82\%$ of these periods. Following their baseline model specification we include a constant and 12 lags in the VAR.

Figure 5 shows the pointwise 90% confidence bands of the impulse responses to a monetary policy shock that increases the one-year-rate by 25 basis points on impact. Such a shock corresponds roughly to a one standard deviation shock in Gertler and Karadi (2015). The point estimates of the impulse responses are qualitatively in line with the findings by Gertler and Karadi (2015).² A monetary tightening induces declining point estimates of

²Differences compared to Gertler and Karadi (2015) are the result of a shorter sample and differences in the bootstrap procedures.

the response of industrial production and consumer prices and an increase in the excess bond premium by slightly more than 10 basis points on impact. However, the bootstrap confidence intervals indicate that the responses of industrial production and the CPI may not be significant, although the choice of the bootstrap procedure affects the widths of the confidence intervals, in line with the simulation results reported in Section 4.

From Figure 5(a), it is apparent that the bands estimated via RBB and MBB tend to be either very similar or the RBB intervals are slightly larger than the MBB intervals. This outcome is consistent with the simulation evidence, see for example Figure 3, panel (b). Recall, however, that the shorter MBB intervals in the simulations come at the price of a lower coverage rate which may be below the nominal 90% rate. Although the interpretation of the impulse responses does not depend on the choice of bootstrap in this case, it is, of course, desirable to employ the most reliable inference procedure.

Figure 5(b) compares the RBB intervals to the WB intervals and shows that the intervals estimated via WB tend to be larger than the RBB intervals. In line with our simulations, the RBB intervals are actually substantially shorter in some cases (see, e.g., the IP impulse responses). For the excess bond premium, the WB interval for the impact effect of the monetary policy shock actually includes zero while this is not the case for the RBB interval. Hence, based on the WB one may conclude that there is no significant impact effect of the monetary policy shock and that conclusion may just be due to an excessively large confidence interval.

7 Conclusions

In proxy VAR models, an external proxy variable that is correlated with a structural shock of interest and uncorrelated with all other shocks, is used for inference for the impulse responses. In this study, we propose a new bootstrap for such inference. So far frequentist inference in this context is typically based on the WB or the MBB. The former is not valid asymptotically and often yields rather wide confidence intervals, whereas the latter has poor coverage properties in small samples as they are often encountered in macroeconomic studies. We propose an alternative bootstrap method which assumes a specific model for the DGP of the proxy variable and samples from the estimated reduced-form errors and the residuals of the proxy model to generate bootstrap samples.

We show by simulation that our new RBB method works well in relatively small samples. Specifically, it yields bootstrap confidence intervals for impulse responses with more accurate coverage than and similar length

for comparable coverage as the WB and the MBB. Thus, it has merit for empirical studies for which only relatively small samples are available.

One advantage of the MBB is that it also works asymptotically for conditionally heteroskedastic model errors while the RBB is not designed for such data features. The price to pay for the additional generality of the MBB is its reduced accuracy in small samples.

References

- Brüggemann, R., Jentsch, C. and Trenkler, C. (2016). Inference in VARs with conditional heteroskedasticity of unknown form, *Journal of Econometrics* **191**: 69–85.
- Carriero, A., Mumtaz, H., Theodoridis, K. and Theophilopoulou, A. (2015). The impact of uncertainty shocks under measurement error: A proxy SVAR approach, *Journal of Money, Credit and Banking* **47**: 1223–1238.
- Davidson, R. and MacKinnon, J. G. (2004). *Econometric Theory and Methods*, Oxford University Press, New York.
- Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity, *American Economic Journal: Macroeconomics* **7**: 44–76.
- Jentsch, C. and Lunsford, K. G. (2019). The dynamic effects of personal and corporate income tax changes in the United States: Comment, *American Economic Review* **109**: 2655–2678.
- Kilian, L. (1998). Small-sample confidence intervals for impulse response functions, *Review of Economics and Statistics* **80**: 218–230.
- Kilian, L. and Kim, Y. (2011). How reliable are local projection estimators of impulse responses?, *Review of Economics and Statistics* **93**: 1460–1466.
- Kilian, L. and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*, Cambridge University Press, Cambridge.
- Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*, Springer-Verlag, Berlin.
- Lütkepohl, H. and Schlaak, T. (2019). Bootstrapping impulse responses of structural vector autoregressive models identified through GARCH, *Journal of Economic Dynamics and Control* **101**: 41–61.

- Lütkepohl, H., Staszewska-Bystrova, A. and Winker, P. (2015a). Comparison of methods for constructing joint confidence bands for impulse response functions, *International Journal for Forecasting* **31**: 782–798.
- Lütkepohl, H., Staszewska-Bystrova, A. and Winker, P. (2015b). Confidence bands for impulse responses: Bonferroni versus Wald, *Oxford Bulletin of Economics and Statistics* **77**: 800–821.
- Mertens, K. and Ravn, M. O. (2013). The dynamic effects of personal and corporate income tax changes in the United States, *American Economic Review* **103**: 1212–1247.
- Piffer, M. and Podstawski, M. (2018). Identifying uncertainty shocks using the price of gold, *The Economic Journal* **128**(1549): 3266–3284.
- Pope, A. L. (1990). Biases for estimators in multivariate non-Gaussian autoregressions, *Journal of Time Series Analysis* **11**: 249–258.
- Stock, J. H. and Watson, M. W. (2012). Disentangling the channels of the 2007-09 recession, *Brookings Papers on Economic Activity* pp. 81–135.
- Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments, *Economic Journal* **128**: 917–948.




Figure 1: Coverage and average lengths of alternative pointwise bootstrap 90% confidence intervals for $T = 100$.




Figure 2: Coverage and average lengths of alternative pointwise bootstrap 90% confidence intervals for $d = 1$, $a_{11} = 0.95$, $\text{corr} = 0.9$.

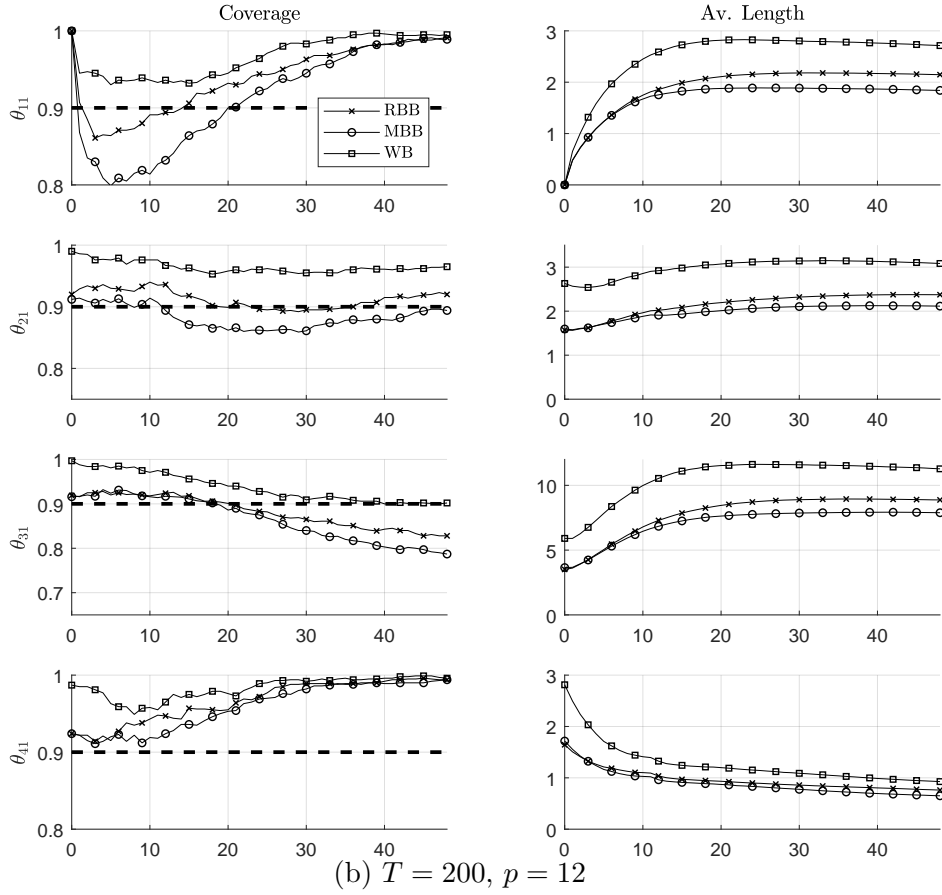
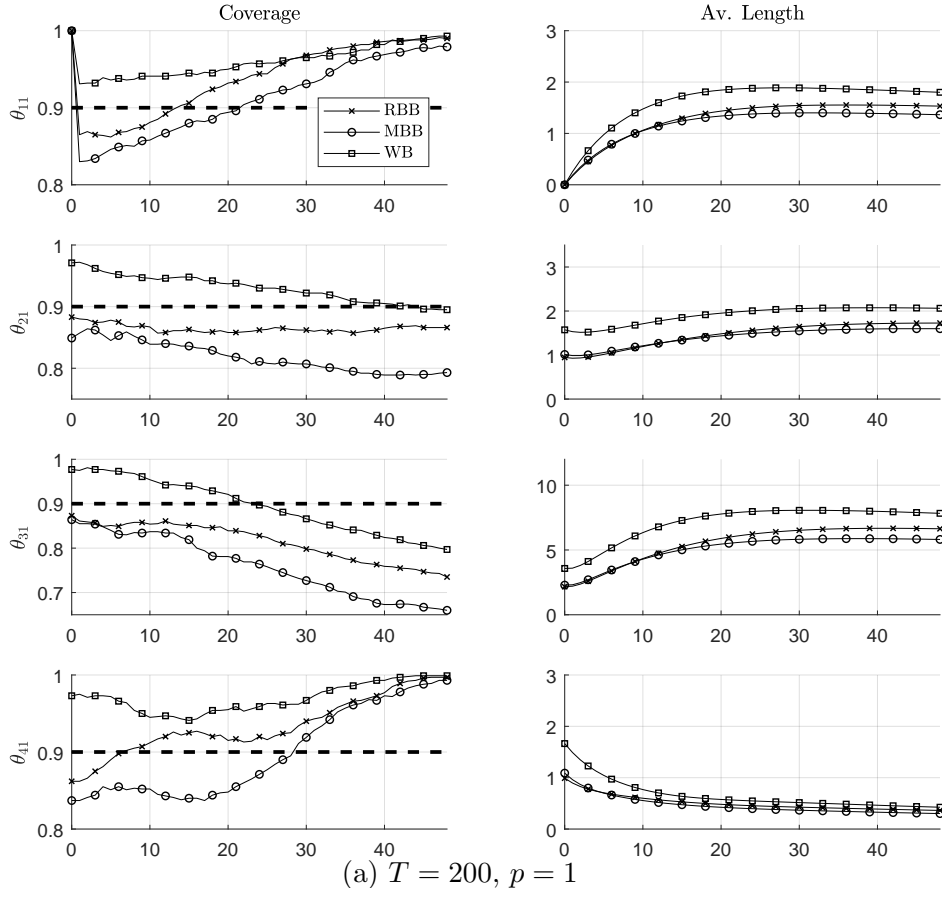


Figure 3: Coverage and average lengths of alternative pointwise bootstrap 90% confidence intervals for DGP2.

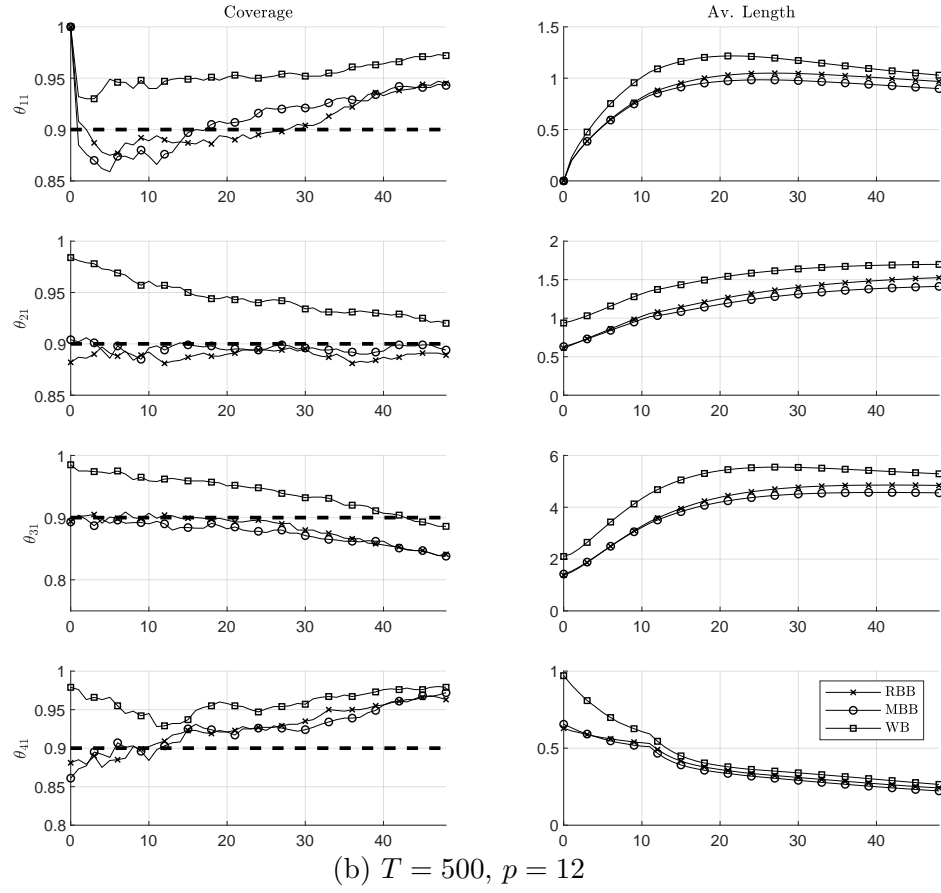
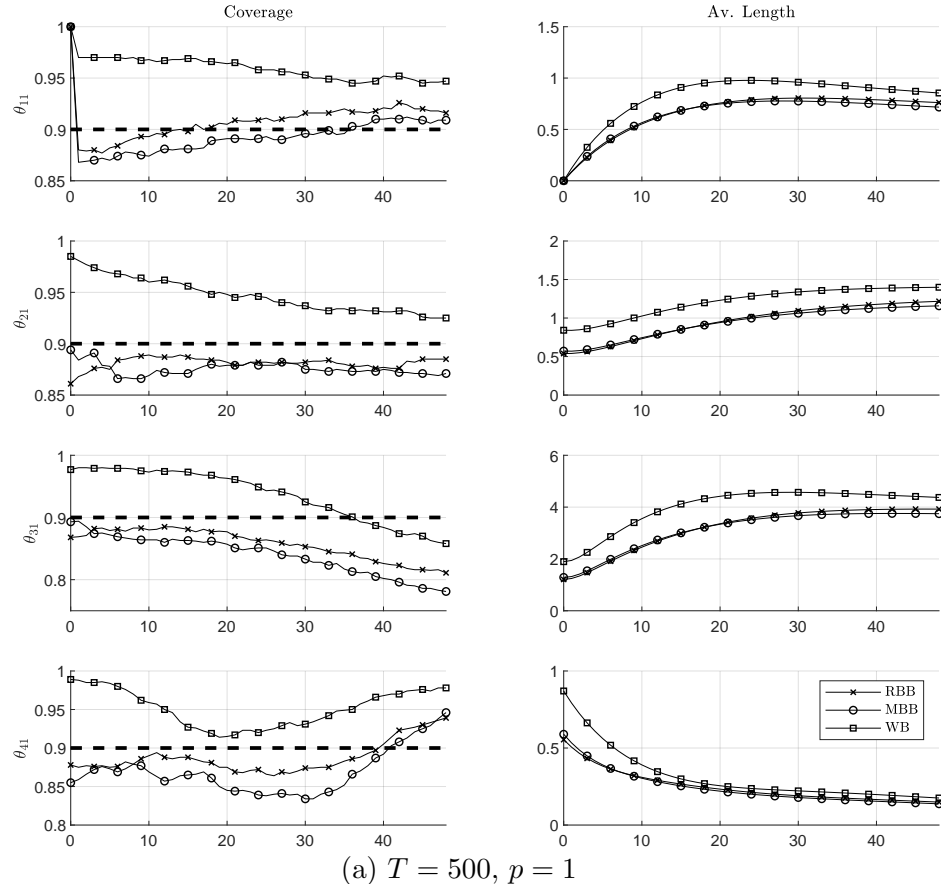


Figure 4: Coverage and average lengths of alternative pointwise bootstrap 90% confidence intervals for DGP2.

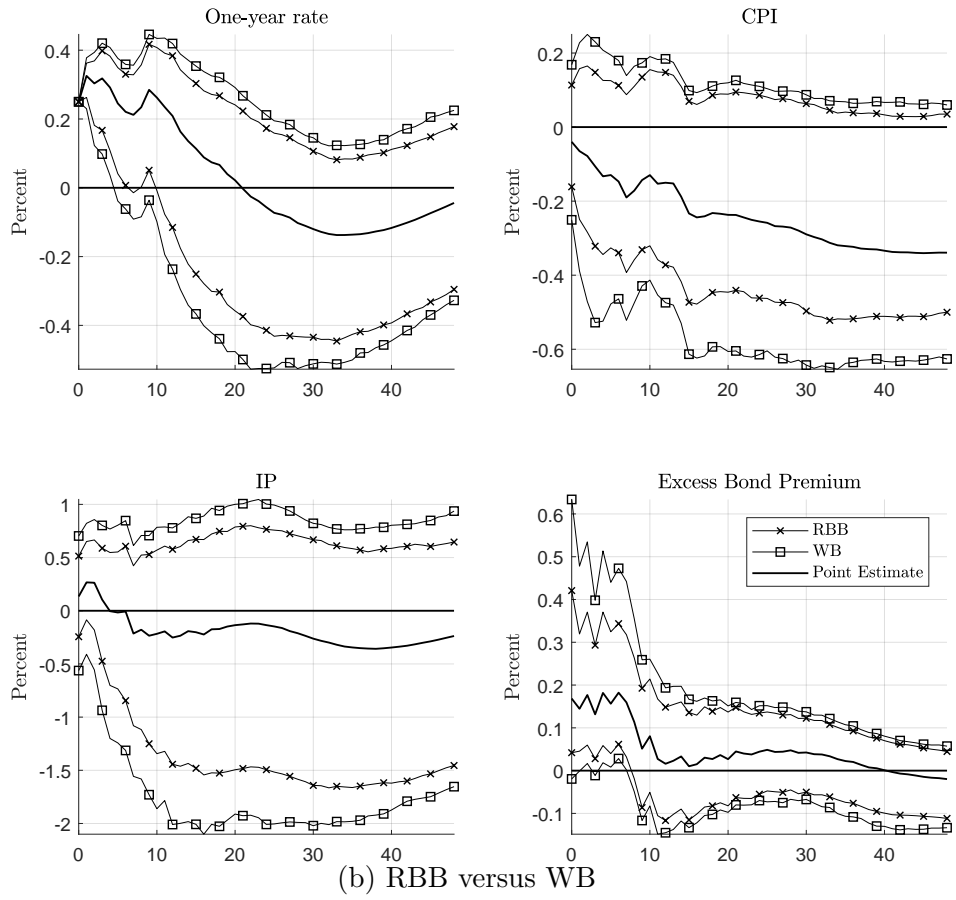







Figure 5: Pointwise bootstrap 90% confidence intervals for the empirical example.



Appendix





(a) $T = 100, p = 1, \text{corr} = 0.9$





(b) $T = 100, p = 12, \text{corr} = 0.9$





(c) $T = 100, p = 1, \text{corr} = 0.5$



(d) $T = 100, p = 12, \text{corr} = 0.5$




(e) $T = 500, p = 1, \text{corr} = 0.5$




(f) $T = 500, p = 12, \text{corr} = 0.5$


Figure A.1: Coverage and average lengths of alternative pointwise bootstrap 90% confidence intervals for $d = 1, a_{11} = 0.95$.




(a) $a_{11} = 0.5$, $p = 1$, $d = 1$, $\text{corr} = 0.9$




(b) $a_{11} = 0.5$, $p = 12$, $d = 1$, $\text{corr} = 0.9$




(c) $a_{11} = 0.9$, $p = 1$, $d = 1$, $\text{corr} = 0.9$



(d) $a_{11} = 0.9$, $p = 12$, $d = 1$, $\text{corr} = 0.9$



(e) $a_{11} = 0.95$, $p = 1$, $d = 1$, $\text{corr} = 0.9$



(f) $a_{11} = 0.95$, $p = 12$, $d = 1$, $\text{corr} = 0.9$

Figure A.2: Coverage and average lengths of alternative pointwise bootstrap 90% confidence intervals for $T = 100$.