Online Supplemental Appendix to Asymptotically Valid Bootstrap Inference for Proxy SVARs*

Carsten Jentsch†
TU Dortmund University

Kurt G. Lunsford‡
Federal Reserve Bank of Cleveland

May 2019

C Additional Monte Carlo Results

This appendix presents additional Monte Carlo results for "Asymptotically Valid Bootstrap Inference for Proxy SVARs" by Carsten Jentsch and Kurt G. Lunsford. The baseline Monte Carlo results in the paper show the coverage rates of 95% confidence intervals from the residual-based moving block bootstrap (MBB) and the residual-based Rademacher wild bootstrap. In this online appendix, we present four additional sets of results. In Section C.1, we show the coverage rates of 68% confidence intervals from the residual-based MBB and the residual-based Rademacher wild bootstrap. In Section C.2, we show the coverage rates of 68% and 95% confidence intervals from the residual-based wild bootstrap with standard normal multipliers. In Section C.3, we show the coverage rates of 68% and 95% confidence intervals from the MBB and the Rademacher wild bootstrap when the proxy variable is often censored to zero. In Section C.4, we show the coverage rates of 68% and 95% confidence intervals from the MBB and the Rademacher wild bootstrap when Hall's percentile intervals are used (Hall (1992) and Lütkepohl (2005, Appendix D)).

*The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

†TU Dortmund University, Faculty of Statistics, D-44221 Dortmund, Germany; jentsch@statistik.tu-dortmund.de

‡Federal Reserve Bank of Cleveland, P.O. Box 6387, Cleveland, OH 44101; kurt.lunsford@clev.frb.org

C.1 Results for 68% Confidence Intervals

This section presents the coverage rates of 68% confidence intervals from the residual-based MBB and the residual-based Rademacher wild bootstrap for DGP1, DGP2 and DGP3. The figures in the section correspond to Figures 2 and 3, Figures 5 through 7, and Figures 9 through 12 in the paper, but with 68% confidence intervals instead of 95% confidence intervals.

C.1.1 Results for DGP1

We begin with results for DGP1. Figure C.1 shows the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs under DGP1. Similar to Figure 2 in the paper, it shows that the coverage rates of the MBB are close to the intended levels on impact, displayed as horizon 0. In contrast, the coverage rates for the wild bootstrap are very undersized and shrink as the sample size increases.

For IRF horizons 1 through 20, C.1 shows results quite similar to the results in Figure 2. The MBB's confidence intervals gradually become undersized; however, this is much less of a problem with large sample sizes. Conversely, the coverage rates of the wild bootstrap's confidence intervals gradually rise and become quite similar to those of the MBB.

Figure C.2 shows the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs under DGP1. Note that $y_{1,t}$ is normalized to fall by 1 on impact and that this normalization is made within every bootstrap loop. Hence, the coverage rates for $y_{1,t}$ are 1 at horizon 0 for both the MBB and the wild bootstrap. For $y_{2,t}$, the MBB is appropriately sized on impact. However, the wild bootstrap is very undersized on impact for $y_{2,t}$, and the coverage rates of its confidence intervals shrink when the sample size increases.

As in Figure 3, the coverage rates for the MBB fall over horizons 1 through 20 with $T = 400$. However, this fall is less noticeable at earlier horizons than for one standard deviation IRFs. For $T = 2000$, the MBB performs well. In contrast, the wild bootstrap's confidence intervals are undersized at both sample sizes, and its coverage rates do not fully converge with those of the MBB by horizon 20.

Figure C.3 shows the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the FEVDs under DGP1. As in Figure 5 of the paper, all panels of Figure C.3 show that the coverage rates of the MBB are generally close to the intended levels, but may be slightly undersized at smaller sample sizes. In contrast, all panels of Figure C.3

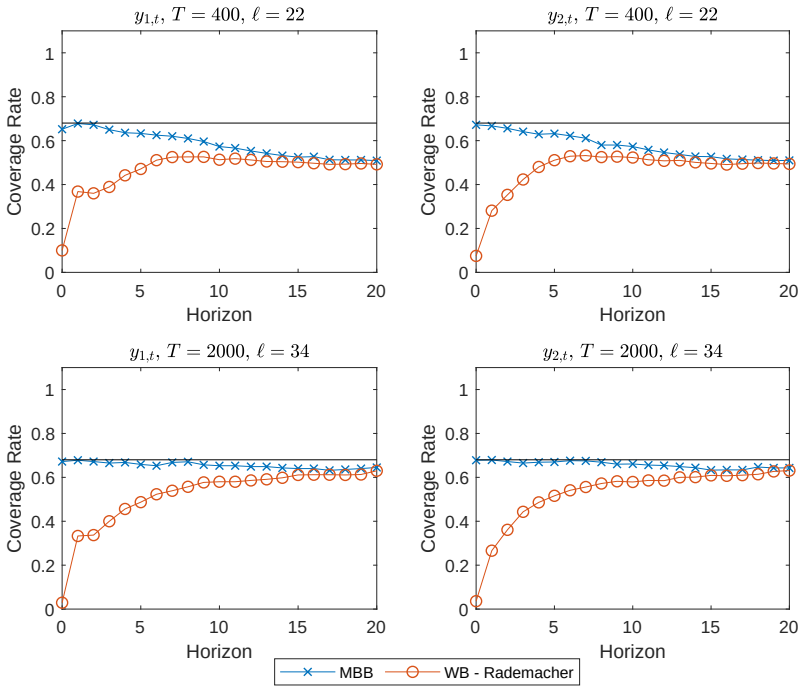

Figure C.1: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP1. The solid horizontal line shows the 0.68 target level.

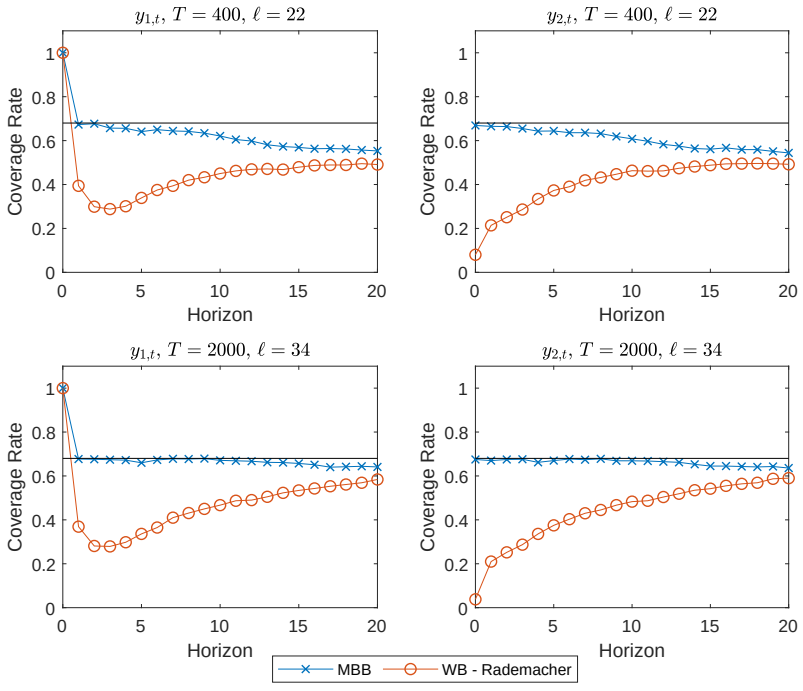

Figure C.2: Coverage rates of 68% confidence intervals for normalized IRFs under DGP1. The solid horizontal line shows the 0.68 target level.

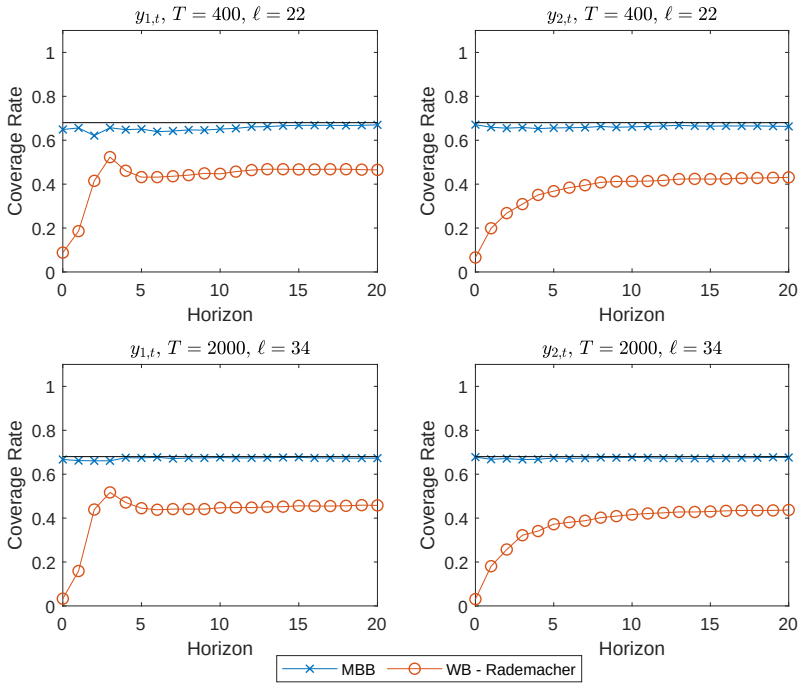

Figure C.3: Coverage rates of 68% confidence intervals for FEVDs under DGP1. The solid horizontal line shows the 0.68 target level.

show that the coverage rates of the wild bootstrap are very undersized. The coverage rates for the wild bootstrap are smallest at horizon 0 and shrink when the sample size increases. While the coverage rates of the wild bootstrap generally rise with the horizon, they do not converge to the MBB at a horizon of 20.

C.1.2 Results for DGP2

Figure C.4 shows the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs under DGP2. It shows that the MBB provides very similar coverage rates when compared to DGP1 in Figure C.1. On impact, the coverage rates are close to the intended levels. With $T = 400$, these rates gradually decline to about 0.51 at a horizon of 20 under DGP1 and DGP2. Under both DGP1 and DGP2, the MBB stays close to the intended levels with $T = 2000$.

In contrast to DGP1, the wild bootstrap coverage rates do not fully converge with the MBB converge rates under DGP2. With $T = 400$, the coverage rates gradually rise to about 0.45 at a horizon of 20 under DGP2. While the coverage rates for $T = 2000$ rise to slightly higher levels, they remain lower than the MBB.

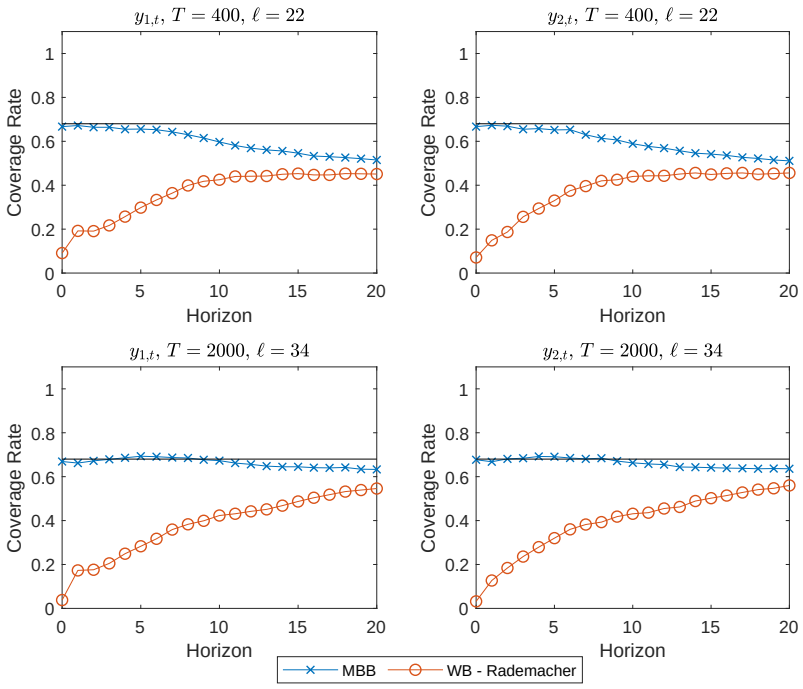

Figure C.4: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP2. The solid horizontal line shows the 0.68 target level.

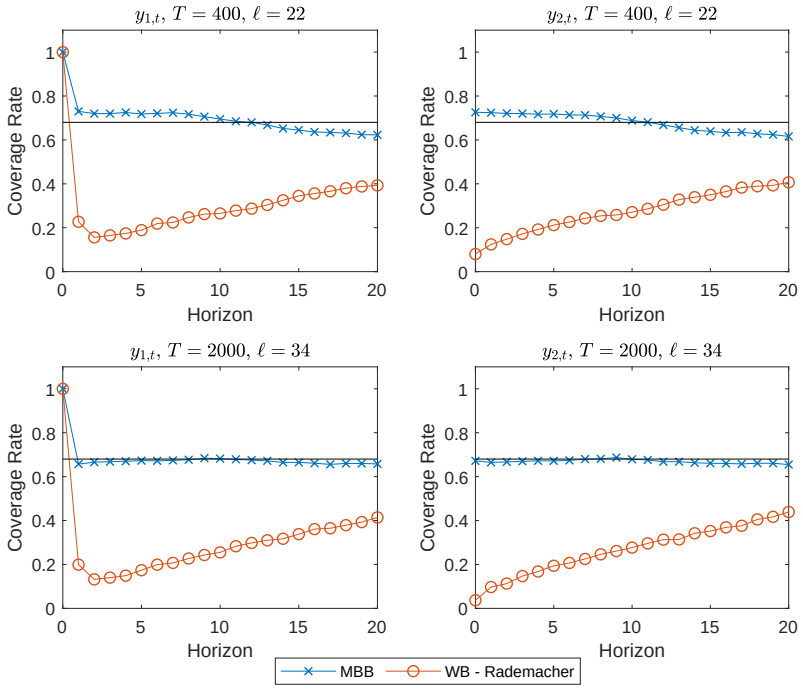

Figure C.5: Coverage rates of 68% confidence intervals for normalized IRFs under DGP2. The solid horizontal line shows the 0.68 target level.

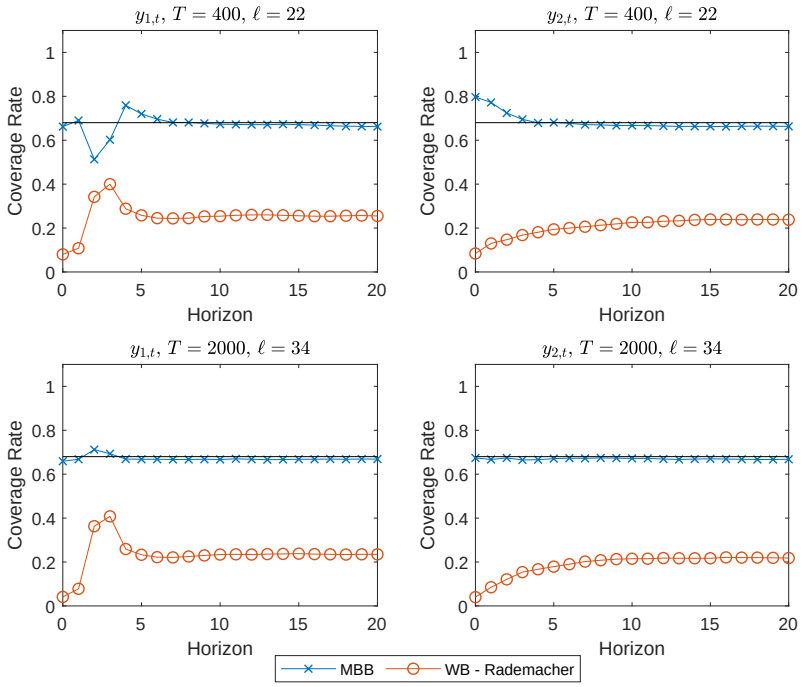

Figure C.6: Coverage rates of 68% confidence intervals for FEVDs under DGP2. The solid horizontal line shows the 0.68 target level.

Figure C.5 shows the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs under DGP2. As with the normalized IRFs under DGP1, the coverage rates for $y_{1,t}$ are 1 at horizon 0 for both the MBB and the wild bootstrap. For $T = 400$, the MBB's coverage rates are slightly too high at low horizons, with rates between 0.72 and 0.73. However, these coverage rates fall as the horizon increases. The MBB also performs well when $T = 2000$.

In Figure C.5, the wild bootstrap's coverage rates stay persistently too low out to horizon 20 and do not rise above 0.45 for either sample size.

Figure C.6 shows the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the FEVDs under DGP2. While the MBB's coverage rates for IRFs were quite similar for DGP1 and DGP2, we see that the coverage rates for FEVDs deteriorate slightly under DGP2 relative to DGP1 when $T = 400$. As in Figure 9 in the paper, this is particularly true for $y_{1,t}$ at horizons 2 through 4, which corresponds to a hump in the true FEVD. However, the MBB does much better when $T = 2000$.

Figure C.6 shows that the wild bootstrap generally performs worse under DGP2 than under DGP1. This is consistent with the wild bootstrap's worse performance for the IRFs

under DGP2 relative to DGP1. In particular, the coverage rates for $y_{2,t}$ never rise to be even half of the intended level.

C.1.3 Results for DGP3

Figure C.7 shows the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs under DGP3. Consistent with Figure 10 in the paper, the coverage rates for the MBB are generally undersized and may even be undersized on impact. However, Figure C.7 shows a noticeable improvement in coverage rates rates from $T = 400$ to $T = 2000$. Further, the MBB usually performs much better and never worse than the wild bootstrap in terms of statistical size under DGP3. As with the previous DGPs, the wild bootstrap can be badly undersized, especially at low horizons.

Figure C.8 shows the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs under DGP3. The MBB's coverage rates are very good at low horizons, in contrast to its coverage rates for the one standard deviation IRFs. While the MBB does become undersized at long horizons for $T = 400$, it performs very well for $T = 2000$. In contrast, the wild bootstrap is very undersized at low horizons and persistently undersized at long horizons, consistent with its coverage of the normalized IRFs under DGP1 and DGP2.

Figure C.9 shows the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the FEVDs under DGP3. The MBB's coverage rates are consistently too low at every horizon with rates generally between 0.59 and 0.62. The exception is for horizons 2 through 4, where the coverage rates for $y_{1,t}$ become even smaller. These rates change little from $T = 400$ to $T = 2000$. Despite these low coverage rates, the MBB consistently dominates the wild bootstrap, which is persistently undersized, in terms of statistical size.

C.2 Results for a Residual-Based Wild Bootstrap with Standard Normal Multiplier

This section presents the coverage rates of 68% and 95% confidence intervals from the residual-based wild bootstrap with standard normal multipliers. For comparison, we also show the coverage rates from the MBB and the Rademacher wild bootstrap. We present results for DGP1, DGP2 and DGP3, and the figures containing the 95% confidence intervals correspond to Figures 2 and 3, Figures 5 through 7, and Figures 9 through 12 in the paper.

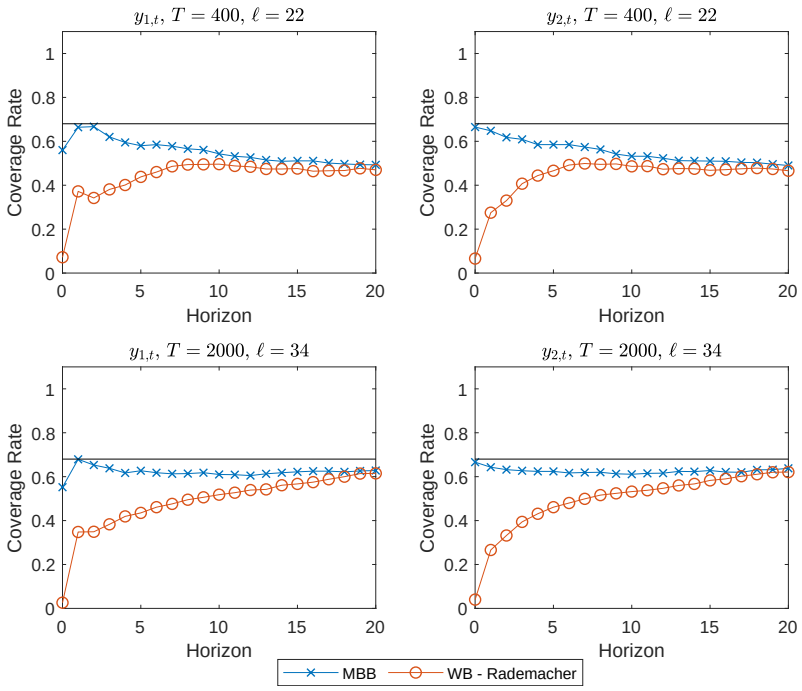

Figure C.7: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP3. The solid horizontal line shows the 0.68 target level.

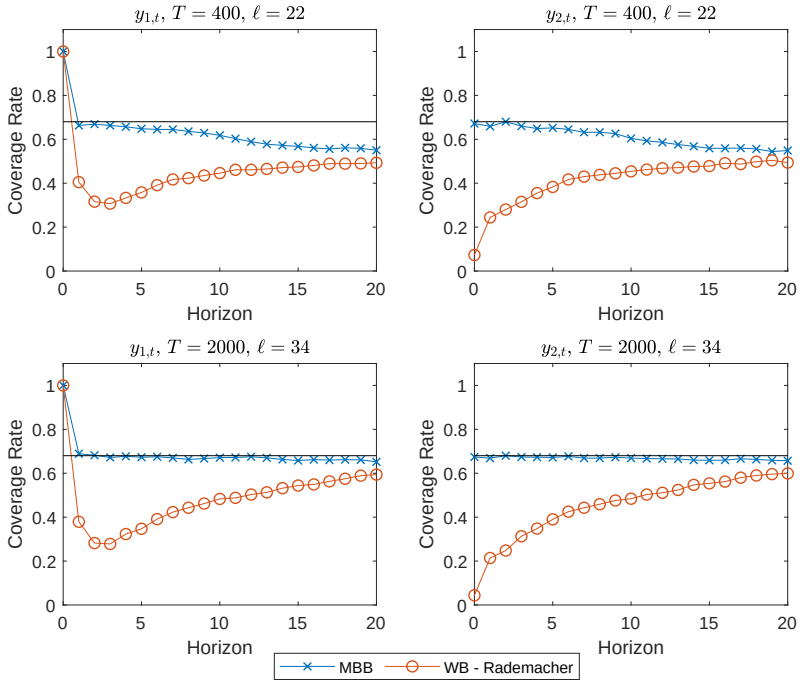

Figure C.8: Coverage rates of 68% confidence intervals for normalized IRFs under DGP3. The solid horizontal line shows the 0.68 target level.

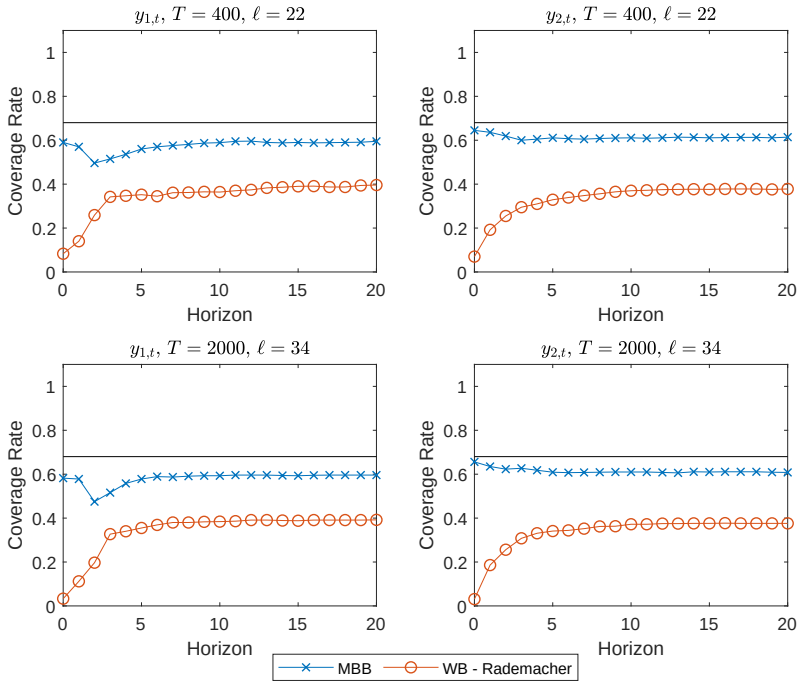

Figure C.9: Coverage rates of 68% confidence intervals for FEVDs under DGP3. The solid horizontal line shows the 0.68 target level.

In general, the normal wild bootstrap's confidence intervals are poorly sized. This follows from Theorem 3.1 in the paper, which shows that wild bootstraps are generally invalid for inference on smooth functions of all of the proxy SVAR parameters. This section highlights an important feature of Theorem 3.1, which is that the invalidity of the wild bootstrap does not depend on the choice of bootstrap multiplier, η_t. Multipliers from the Rademacher distribution, the standard normal distribution, or other distributions that satisfy the conditions in Section 3.1.1 of the paper will generally be invalid and produce confidence intervals that are poorly sized.

We also note that the normal wild bootstrap's confidence intervals are often oversized. In the context of Cholesky-identified SVARs, Brüggemann, Jentsch, and Trenkler (2016) show that a normal wild bootstrap can produce oversized confidence intervals when VAR innovations are iid. The results presented here are consistent with their result.

C.2.1 Results for DGP1

Figures C.10 and C.11 show the confidence interval coverage rates from the normal wild bootstrap along with coverage rates from the MBB and the Rademacher wild bootstrap

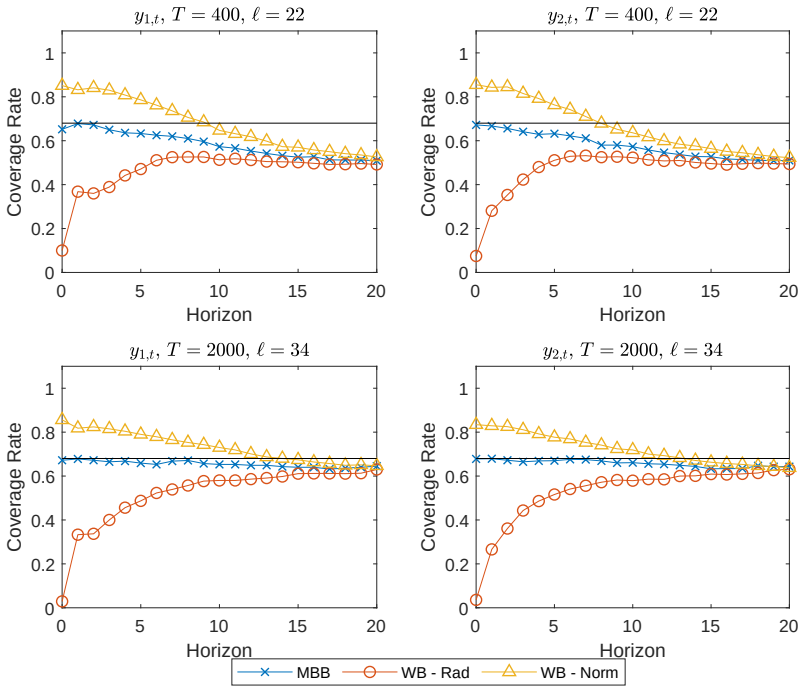

Figure C.10: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP1. The solid horizontal line shows the 0.68 target level.

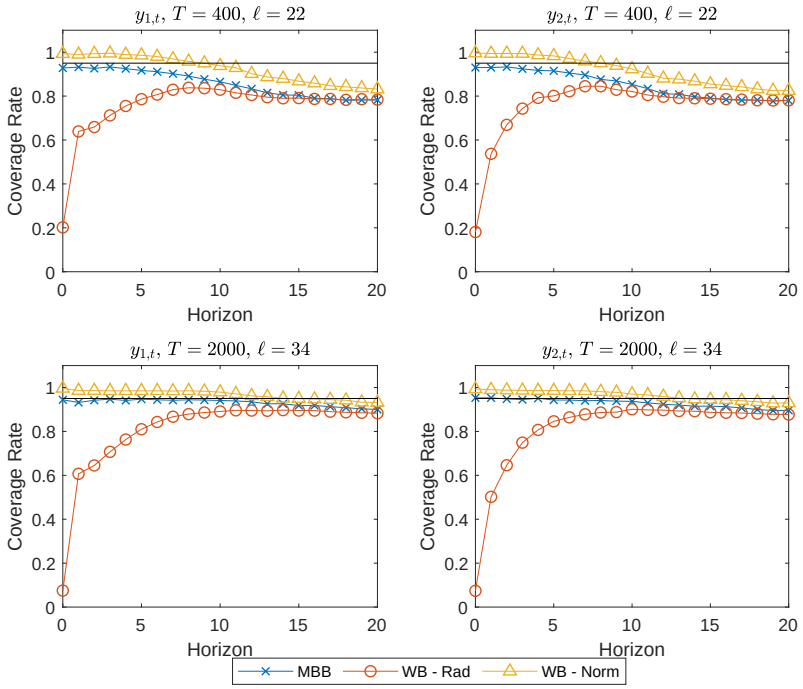

Figure C.11: Coverage rates of 95% confidence intervals for one standard deviation IRFs under DGP1. The solid horizontal line shows the 0.95 target level.

for the one standard deviation IRFs at the 68% and 95% levels, respectively, under DGP1. They show that the normal wild bootstrap has coverage rates that deviate from intended levels and are often too big, even with large sample sizes. As with the Rademacher wild bootstrap, the normal wild bootstrap's coverage rates can converge to those of the MBB at longer horizons.

Figures C.12 and C.13 show the confidence interval coverage rates from the normal wild bootstrap along with coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs at the 68% and 95% levels, respectively, under DGP1. The confidence intervals of the normal wild bootstrap are more persistently oversized with the normalized IRFs than with the one standard deviation IRFs. With large sample sizes, they may be oversized out to horizon 20. This shows that persistently mis-sized confidence intervals can be a general feature of wild bootstraps. That is, they are not unique to the Rademacher wild bootstrap and can occur as well with other distributions of the bootstrap multiplier.

Figures C.14 and C.15 show the confidence interval coverage rates from the normal wild bootstrap along with coverage rates from the MBB and the Rademacher wild bootstrap for FEVDs at the 68% and 95% levels, respectively, under DGP1. The confidence intervals of the normal wild bootstrap are very persistently oversized, again showing that persistently mis-sized confidence intervals are not unique to the Rademacher distribution.

C.2.2 Results for DGP2

Figures C.16 and C.17 show the confidence interval coverage rates from the normal wild bootstrap along with coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs at the 68% and 95% levels, respectively, under DGP2. As with DGP1, the normal wild bootstrap has coverage rates that deviate from intended levels and are often too big under DGP2. Also, the normal wild bootstrap's coverage rates deviate from those of the MBB more persistently under DGP2 than under DGP1.

Figures C.18 and C.19 show the confidence interval coverage rates from the normal wild bootstrap along with coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs at the 68% and 95% levels, respectively, under DGP2. The confidence intervals of the normal wild bootstrap are more persistently oversized with the normalized IRFs than with the one standard deviation IRFs, and this persistence appears to increase with large sample sizes.

Figures C.20 and C.21 show the confidence interval coverage rates from the normal wild

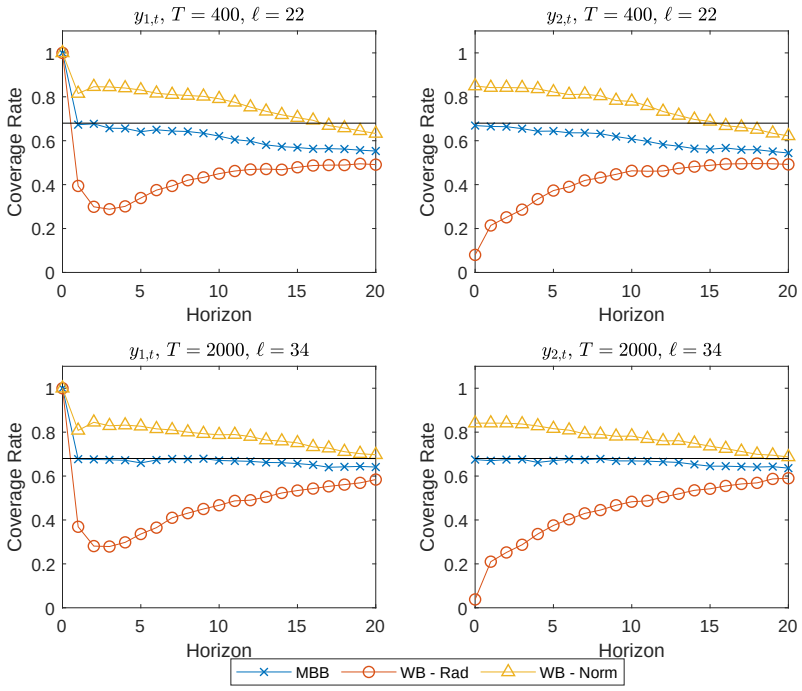

Figure C.12: Coverage rates of 68% confidence intervals for normalized IRFs under DGP1. The solid horizontal line shows the 0.68 target level.

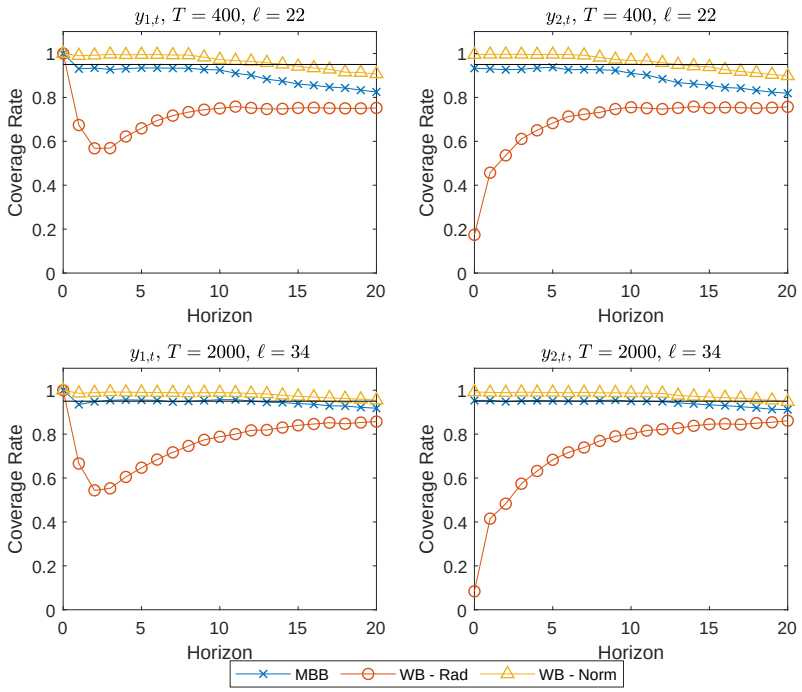

Figure C.13: Coverage rates of 95% confidence intervals for normalized IRFs under DGP1. The solid horizontal line shows the 0.95 target level.

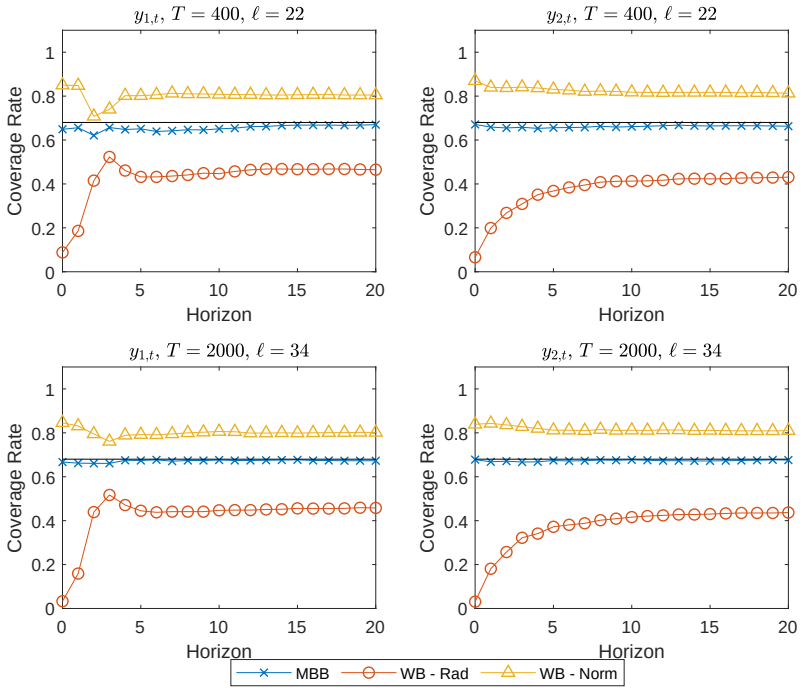

Figure C.14: Coverage rates of 68% confidence intervals for FEVDs under DGP1. The solid horizontal line shows the 0.68 target level.

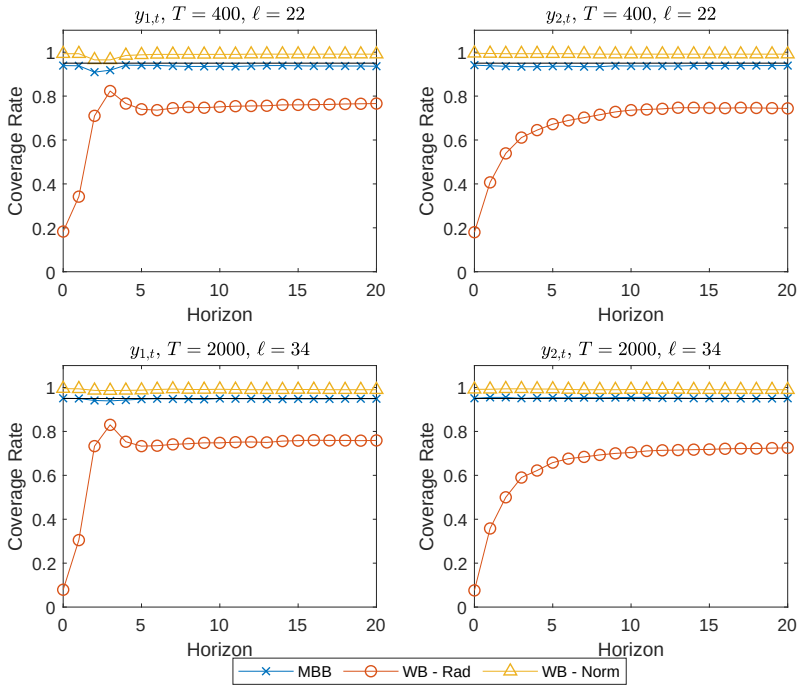

Figure C.15: Coverage rates of 95% confidence intervals for FEVDs under DGP1. The solid horizontal line shows the 0.95 target level.

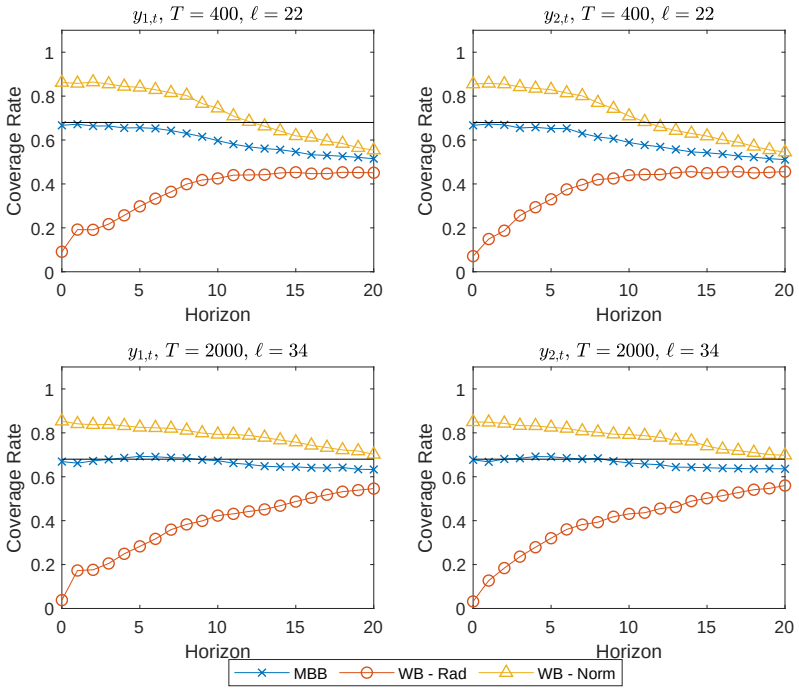

Figure C.16: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP2. The solid horizontal line shows the 0.68 target level.

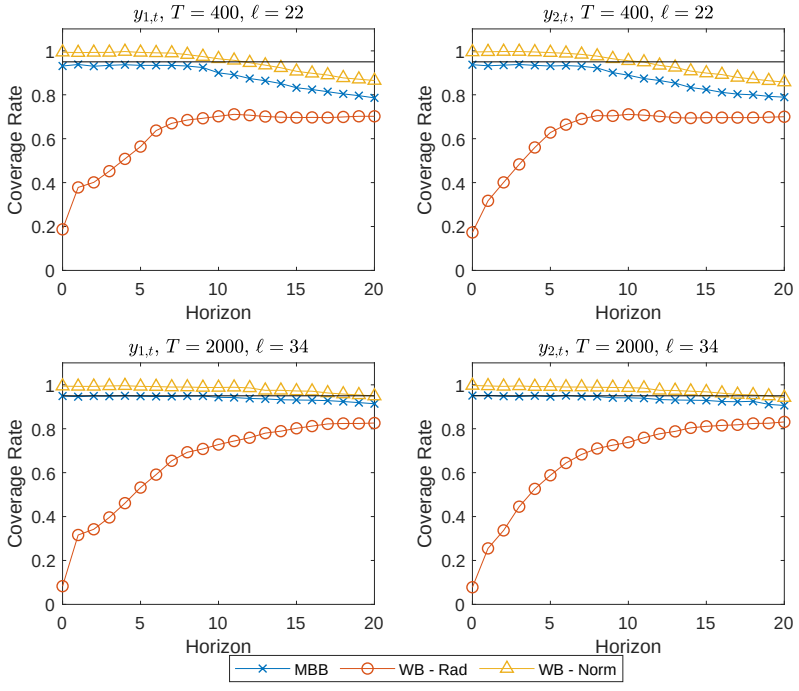

Figure C.17: Coverage rates of 95% confidence intervals for one standard deviation IRFs under DGP2. The solid horizontal line shows the 0.95 target level.

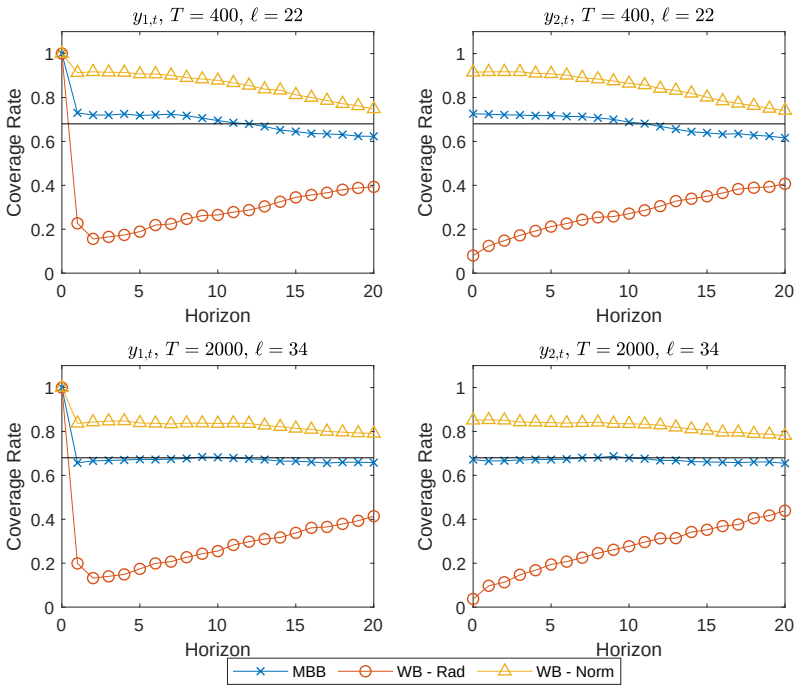

Figure C.18: Coverage rates of 68% confidence intervals for normalized IRFs under DGP2. The solid horizontal line shows the 0.68 target level.

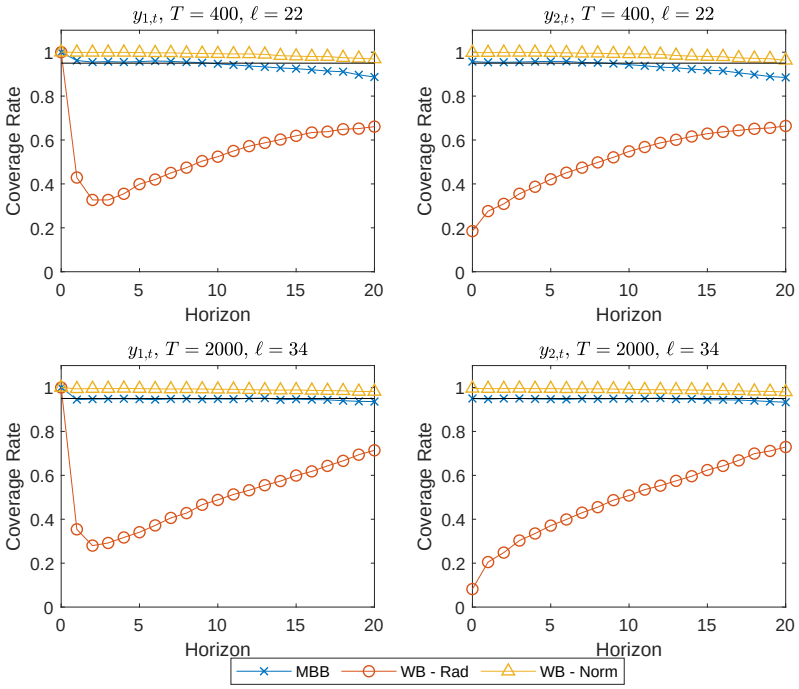

Figure C.19: Coverage rates of 95% confidence intervals for normalized IRFs under DGP2. The solid horizontal line shows the 0.95 target level.

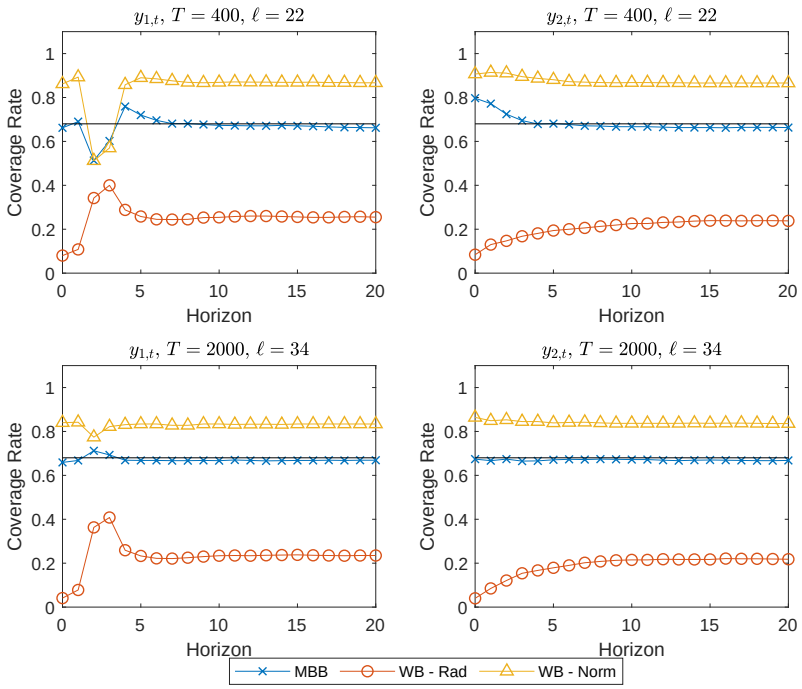

Figure C.20: Coverage rates of 68% confidence intervals for FEVDs under DGP2. The solid horizontal line shows the 0.68 target level.

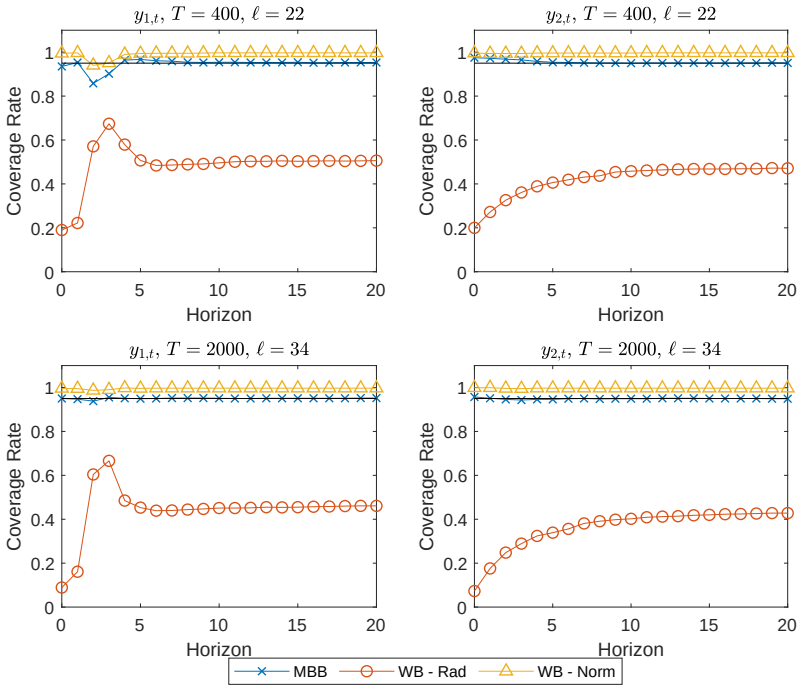

Figure C.21: Coverage rates of 95% confidence intervals for FEVDs under DGP2. The solid horizontal line shows the 0.95 target level.

bootstrap along with coverage rates from the MBB and the Rademacher wild bootstrap for FEVDs at the 68% and 95% levels, respectively, under DGP2. The confidence intervals of the normal wild bootstrap are almost always oversized. Exceptions may occur around humps in the FEVD.

C.2.3 Results for DGP3

Figures C.22 and C.23 show the confidence interval coverage rates from the normal wild bootstrap along with coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs at the 68% and 95% levels, respectively, under DGP3. With the conditional heteroskedasticity in DGP3, the normal wild bootstrap gives better coverage rates than under DGP1 or DGP2. However, it remains oversized at some horizons.

We note that the improved coverage rates of the normal wild bootstrap are accidental here. Our GARCH DGP produces quite a bit of uncertainty, and even with $T = 2000$ and $\ell = 34$ the MBB may not capture all of this uncertainty and thus have low coverage rates. As seen for DGP1 and DGP2, the normal wild bootstrap captures more uncertainty than the MBB and so accidentally gives decent coverage for DGP3. However, the normal wild bootstrap remains asymptotically invalid for functions of all of the proxy SVAR parameters. This will be apparent next when we discuss normalized IRFs.

Figures C.24 and C.25 show the confidence interval coverage rates from the normal wild bootstrap along with coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs at the 68% and 95% levels, respectively, under DGP3. As with DGP1 and DGP2, the normal wild bootstrap's confidence intervals are persistently oversized. This highlights that although the normal wild bootstrap may have decent coverage rates for one standard deviation IRFs under heteroskedasticity, these decent coverage rates may not hold across all statistics of interest.

Figures C.26 and C.27 show the confidence interval coverage rates from the normal wild bootstrap along with coverage rates from the MBB and the Rademacher wild bootstrap for FEVDs at the 68% and 95% levels, respectively, under DGP3. As with the one standard deviation IRFs, the normal wild bootstrap's coverage rates are better under DGP3 than under DGP1 or DGP2. The coverage rates for $y_{1,t}$ are close to the intended levels. As with the one standard deviation IRFs, this decent coverage is accidental and we note that the coverage rates for $y_{2,t}$ remain persistently too large.

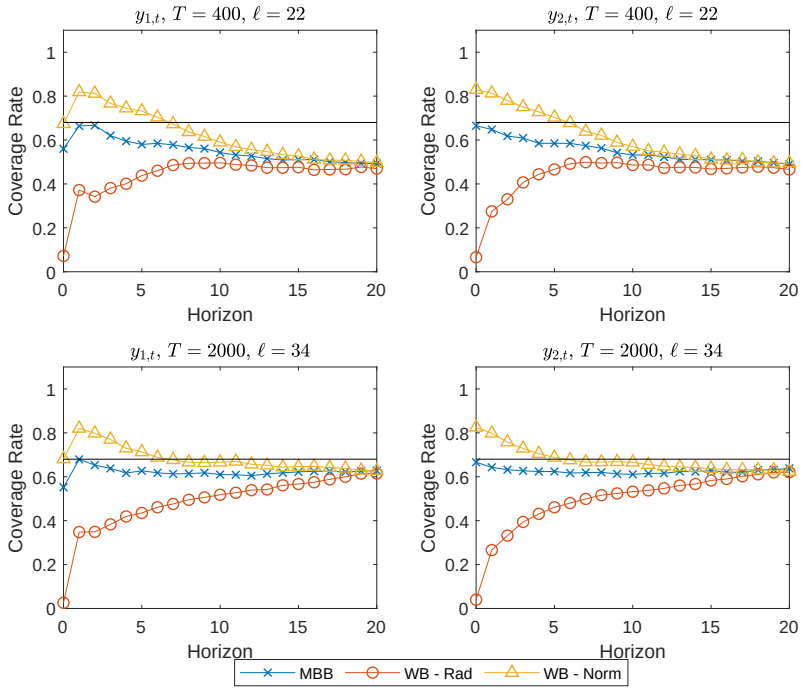

Figure C.22: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP3. The solid horizontal line shows the 0.68 target level.

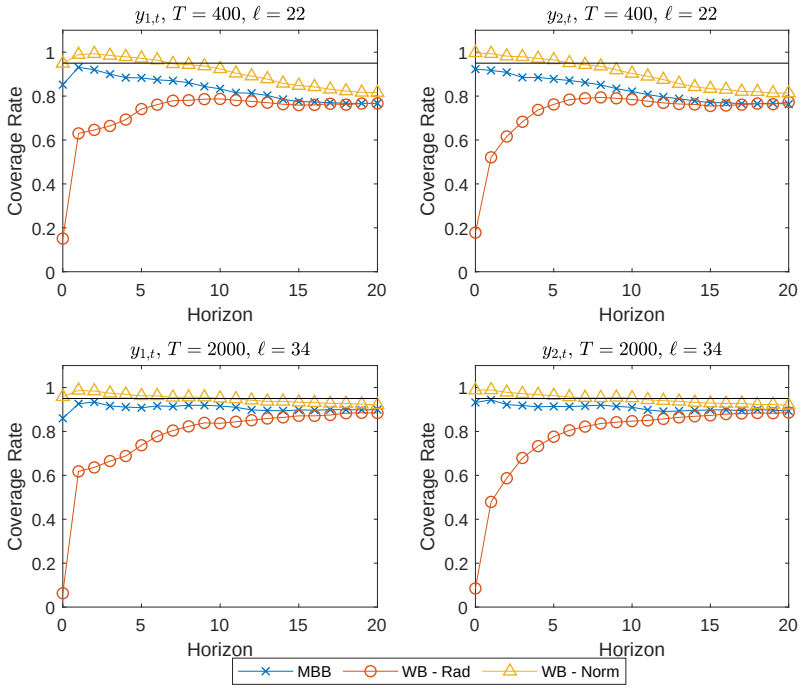

Figure C.23: Coverage rates of 95% confidence intervals for one standard deviation IRFs under DGP3. The solid horizontal line shows the 0.95 target level.

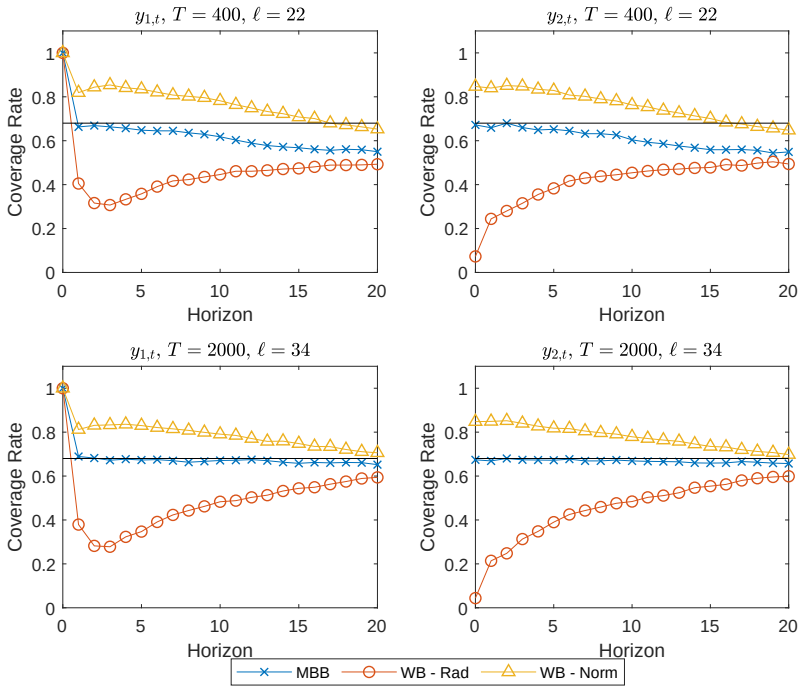

Figure C.24: Coverage rates of 68% confidence intervals for normalized IRFs under DGP3. The solid horizontal line shows the 0.68 target level.

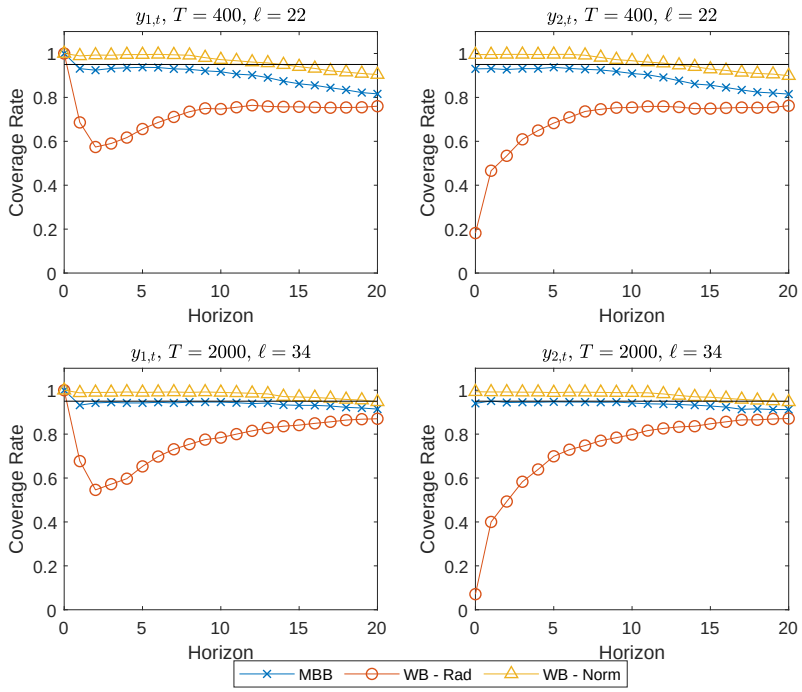

Figure C.25: Coverage rates of 95% confidence intervals for normalized IRFs under DGP3. The solid horizontal line shows the 0.95 target level.

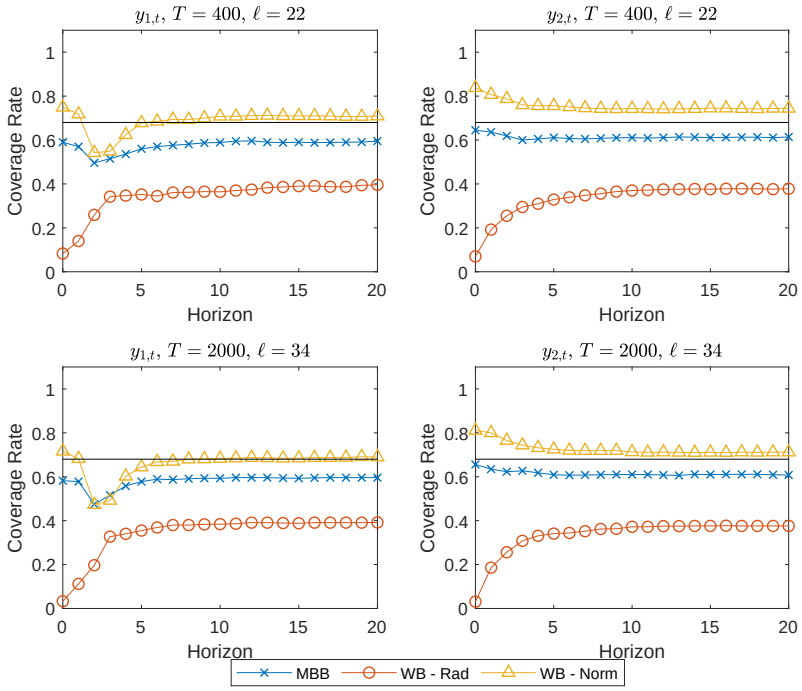

Figure C.26: Coverage rates of 68% confidence intervals for FEVDs under DGP3. The solid horizontal line shows the 0.68 target level.

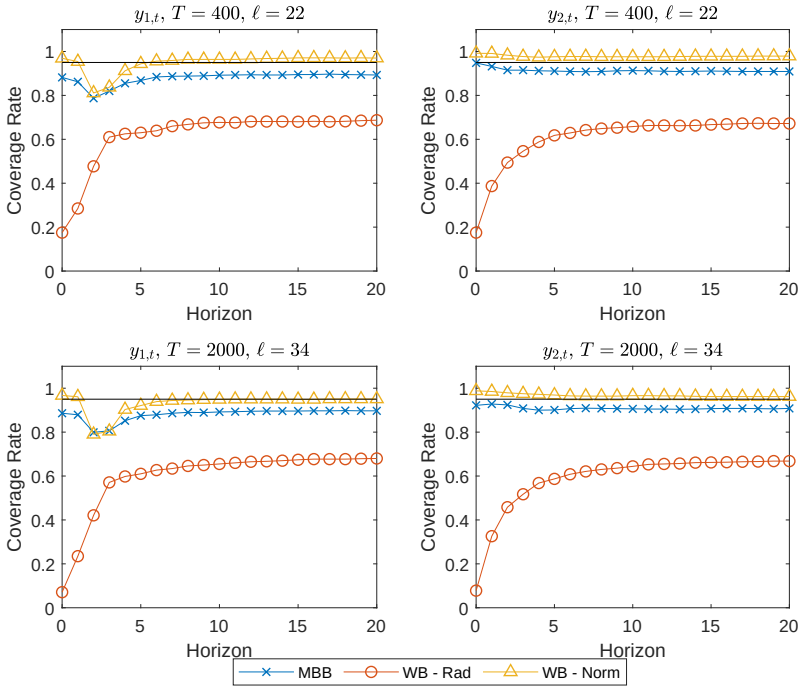

Figure C.27: Coverage rates of 95% confidence intervals for FEVDs under DGP3. The solid horizontal line shows the 0.95 target level.

C.3 Results for a Data Generating Process with Censored Proxy Variables

In applications, it is often the case that proxy variables are censored to zero. For example, see Mertens and Ravn (2013) and Gertler and Karadi (2015). In this section, we modify our Monte Carlo simulations to account for the possibility that the proxy variables may be censored to zero. Specifically, instead of using $m_t = \Psi\epsilon_t^{(1)} + v_t$ with $v_t \sim \mathcal{N}(0,1)$ as in the paper, we now use $m_t = d_t(\Pi\epsilon_t^{(1)} + v_t)$, with $v_t \sim \mathcal{N}(0,1)$. This follows Equation (8) in Mertens and Ravn (2013). We assume that $d_t \in \{0,1\}$ is iid. We choose $\Pr(d_t = 0) = 0.8$ and $\Pr(d_t = 1) = 0.2$. This indicates that 80% of the proxy variables will be censored to zero in population. This falls in between the rate of censoring observed in Mertens and Ravn (2013) and Gertler and Karadi (2015). Following Jentsch and Lunsford (2019), we adjust Π so that $\mathbb{E}(m_t\epsilon_t^{(1)}) = \Psi$, where $\Psi = 0.5$ in DGP1 and DGP3 and $\Psi = 0.2$ in DGP2. This gives $\Pi = 2.5$ in DGP1 and DGP3 and $\Pi = 1$ in DGP2. We do this to isolate the effect of censoring on the bootstrap algorithms independent of the effect of lowering $\mathbb{E}(m_t\epsilon_t^{(1)})$.

For all figures in this section, we present four coverage rates. We present the coverage rates of the confidence intervals from the MBB and the Rademacher wild bootstrap produced from our baseline Monte Carlo simulations presented in the paper. We also present the coverage rates produced from our Monte Carlo simulations with censored proxy variables as discussed in the previous paragraph. We present the coverage rates of 68% and 95% confidence intervals, and the figures containing the 95% confidence intervals correspond to Figures 2 and 3, Figures 5 through 7, and Figures 9 through 12 in the paper. In general, the coverage rates with censored proxy variables closely align with the coverage rates when proxy variables are not censored. This is consistent with the findings of Jentsch and Lunsford (2019), who run Monte Carlo simulations with censoring with a different DGP. However, we note that the effects from censoring appear larger under DGP2 when Ψ is smaller.

C.3.1 Results for DGP1

Figures C.28 and C.29 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs at the 68% and 95% levels, respectively, under DGP1 without and with censoring. They show that censoring has very little effect on the coverage rates as the coverage rates with censored proxy variables lie essentially on top of the coverage rates in the baseline simulations.

Figures C.30 and C.31 show the confidence interval coverage rates from the MBB and

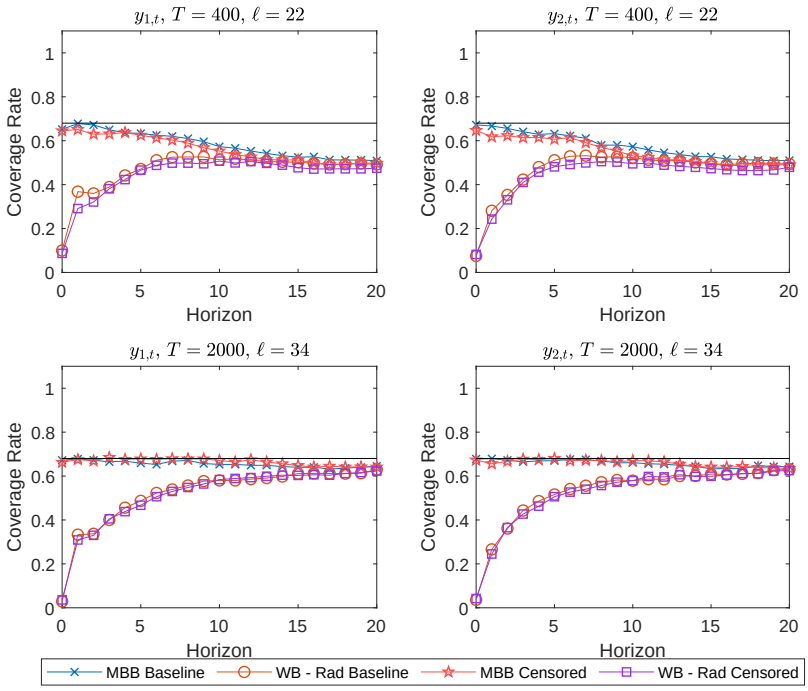

Figure C.28: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP1. The solid horizontal line shows the 0.68 target level.

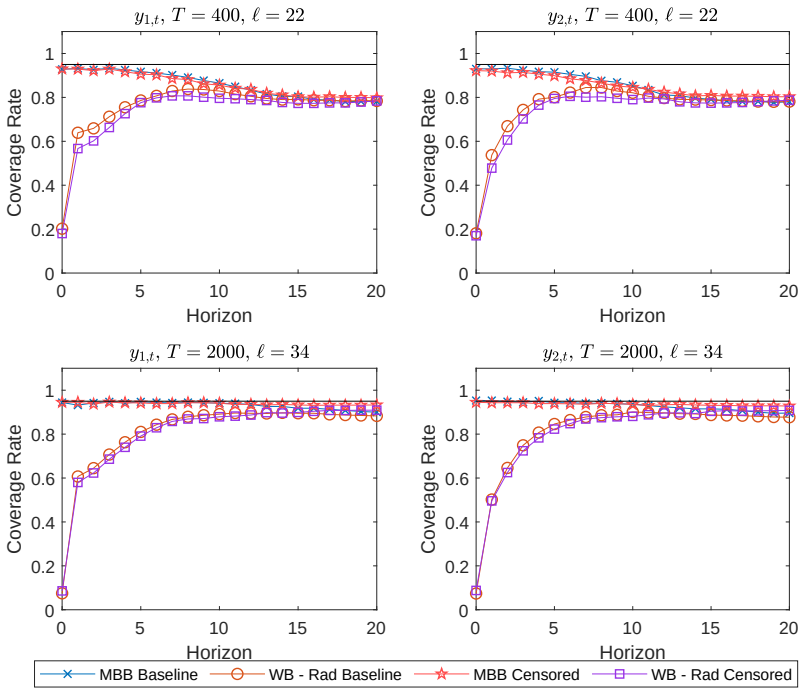

Figure C.29: Coverage rates of 95% confidence intervals for one standard deviation IRFs under DGP1. The solid horizontal line shows the 0.95 target level.

the Rademacher wild bootstrap for normalized IRFs at the 68% and 95% levels, respectively, under DGP1 without and with censoring. As with the one standard deviation IRFs, censoring appears to have very little impact on the coverage rates for normalized IRFs.

Figures C.32 and C.33 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for FEVDs at the 68% and 95% levels, respectively, under DGP1 without and with censoring. Again, censoring has little impact on these coverage rates.

C.3.2 Results for DGP2

Figures C.34 and C.35 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs at the 68% and 95% levels, respectively, under DGP2 without and with censoring. With this DGP, censoring has a modest impact on the results. The MBB's confidence intervals become slightly undersized, but only with small sample sizes. Also, the coverage rates for the wild bootstrap increase slightly. However, they remain persistently too low.

Figures C.36 and C.37 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs at the 68% and 95% levels, respectively, under DGP2 without and with censoring. As with the one standard deviation IRFs, censoring causes the coverage rates to fall for the MBB and rise for the wild bootstrap. However, the wild bootstrap remains persistently undersized and the MBB continues to dominate in terms of statistical size.

Figures C.38 and C.39 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for FEVDs at the 68% and 95% levels, respectively, under DGP2 without and with censoring. Again, censoring generally reduces coverage rates for the MBB and increases coverage rates for the wild bootstrap, but the wild bootstrap remains persistently very undersized.

C.3.3 Results for DGP3

Figures C.40 and C.41 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs at the 68% and 95% levels, respectively, under DGP3 without and with censoring. As with DGP1, censoring only has a small impact on coverage rates.

Figures C.42 and C.43 show the confidence interval coverage rates from the MBB and

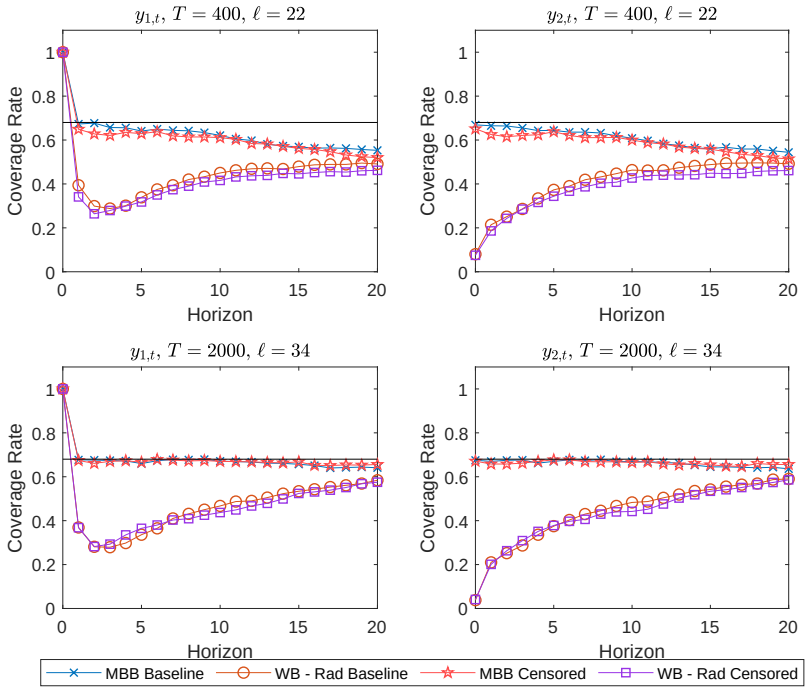

Figure C.30: Coverage rates of 68% confidence intervals for normalized IRFs under DGP1. The solid horizontal line shows the 0.68 target level.

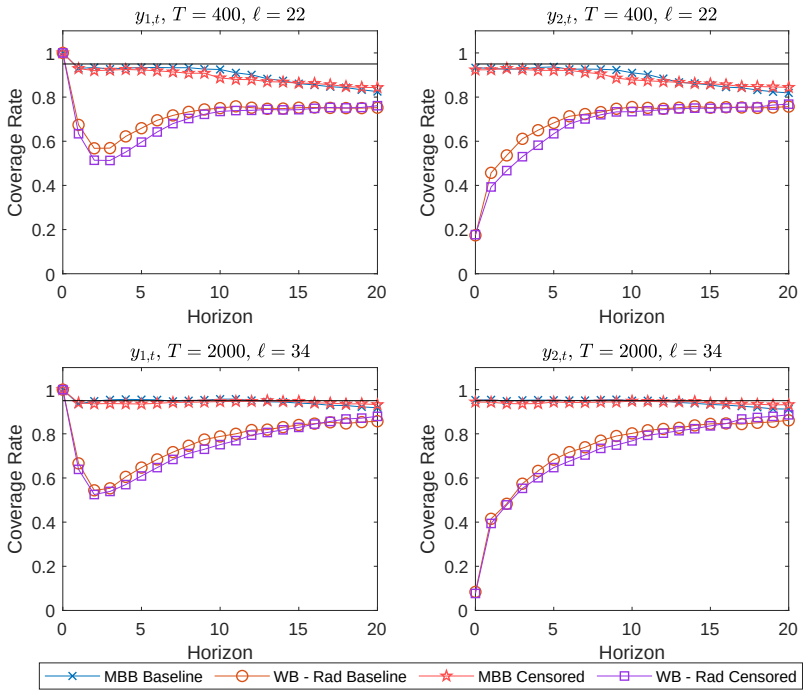

Figure C.31: Coverage rates of 95% confidence intervals for normalized IRFs under DGP1. The solid horizontal line shows the 0.95 target level.

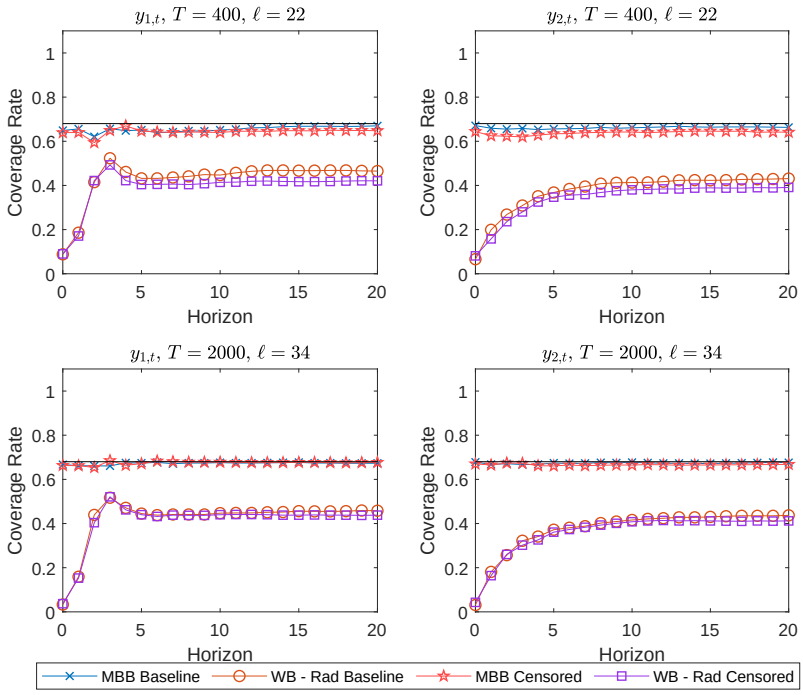

Figure C.32: Coverage rates of 68% confidence intervals for FEVDs under DGP1. The solid horizontal line shows the 0.68 target level.

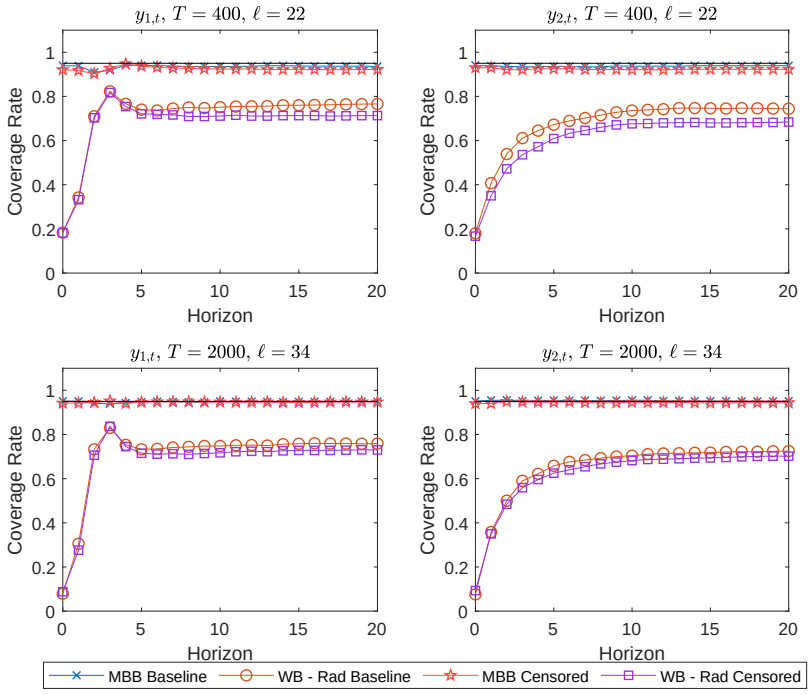

Figure C.33: Coverage rates of 95% confidence intervals for FEVDs under DGP1. The solid horizontal line shows the 0.95 target level.

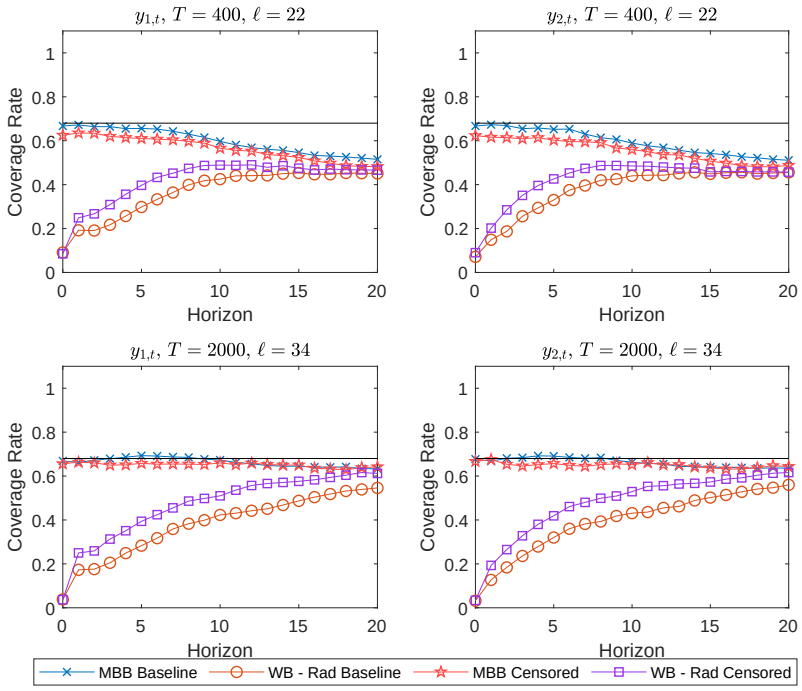

Figure C.34: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP2. The solid horizontal line shows the 0.68 target level.

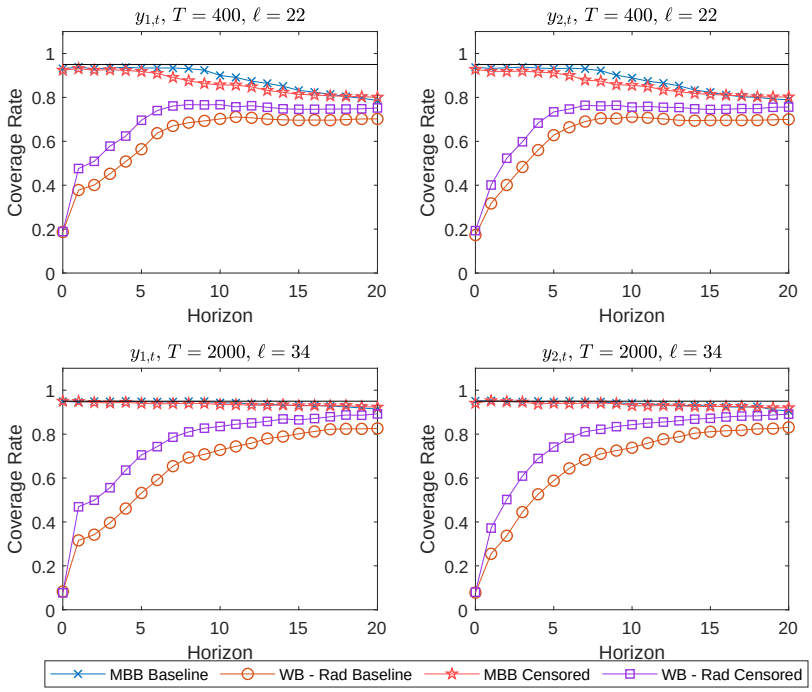

Figure C.35: Coverage rates of 95% confidence intervals for one standard deviation IRFs under DGP2. The solid horizontal line shows the 0.95 target level.

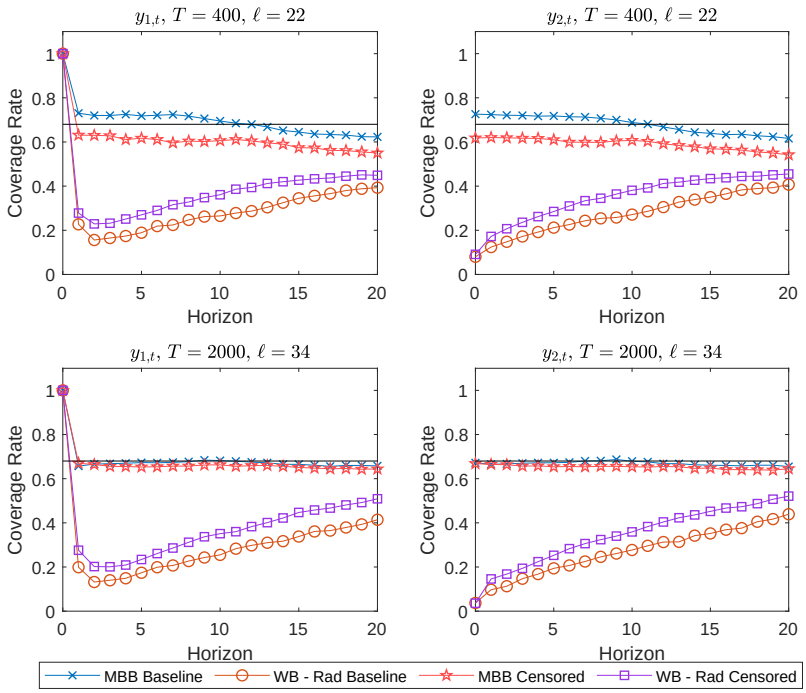

Figure C.36: Coverage rates of 68% confidence intervals for normalized IRFs under DGP2. The solid horizontal line shows the 0.68 target level.

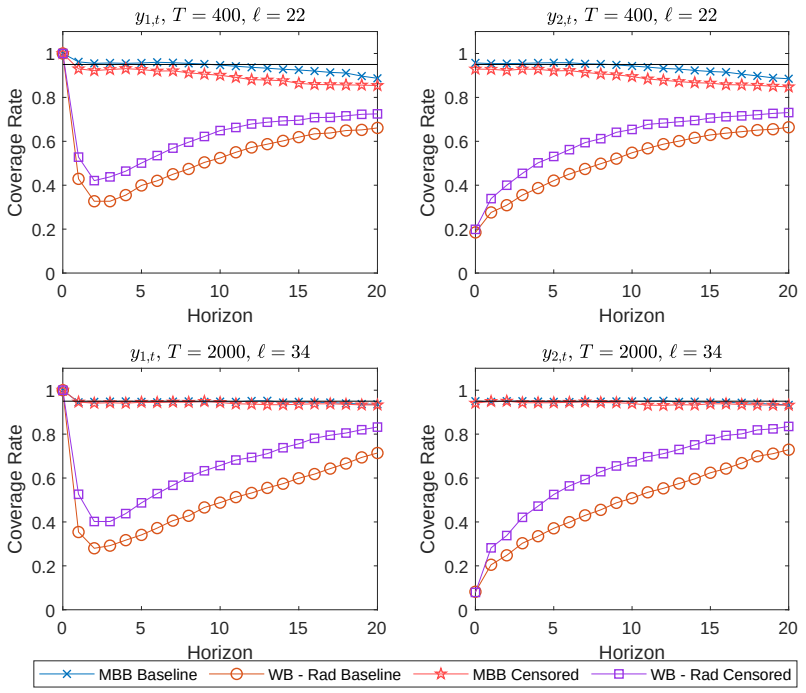

Figure C.37: Coverage rates of 95% confidence intervals for normalized IRFs under DGP2. The solid horizontal line shows the 0.95 target level.

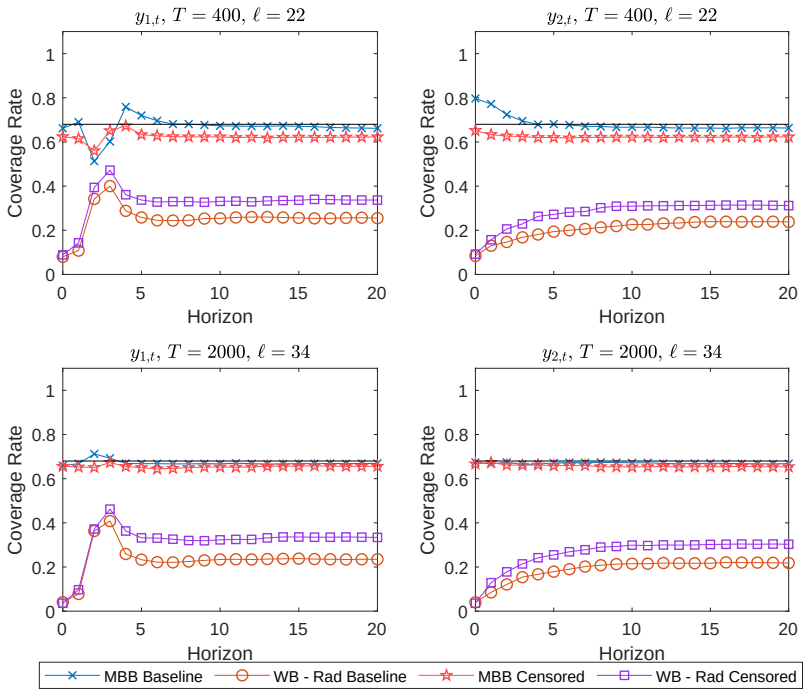

Figure C.38: Coverage rates of 68% confidence intervals for FEVDs under DGP2. The solid horizontal line shows the 0.68 target level.

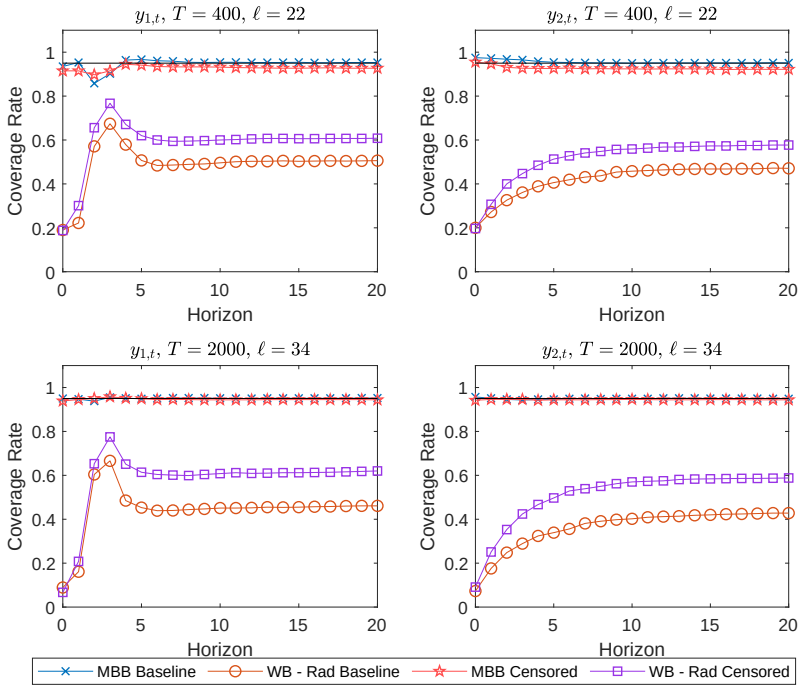

Figure C.39: Coverage rates of 95% confidence intervals for FEVDs under DGP2. The solid horizontal line shows the 0.95 target level.

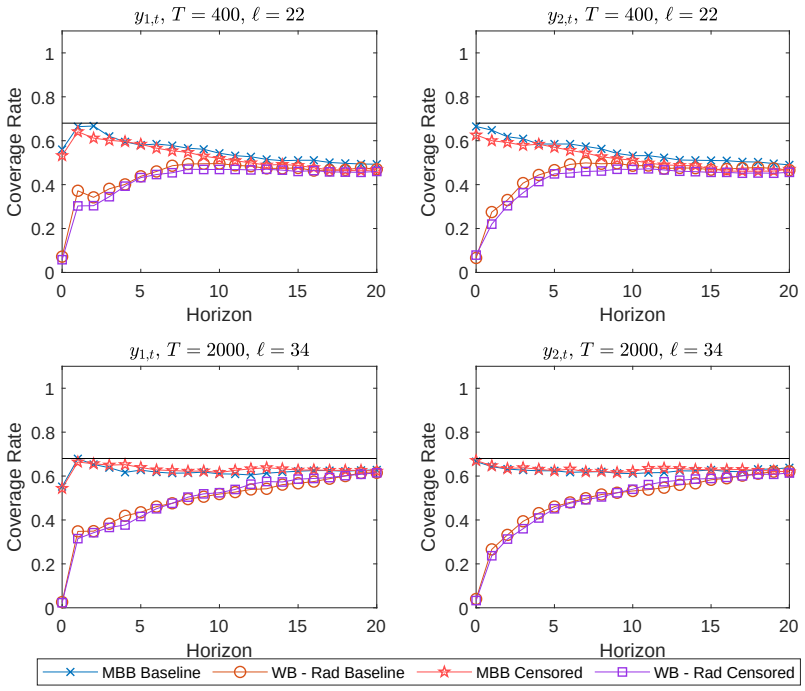

Figure C.40: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP3. The solid horizontal line shows the 0.68 target level.

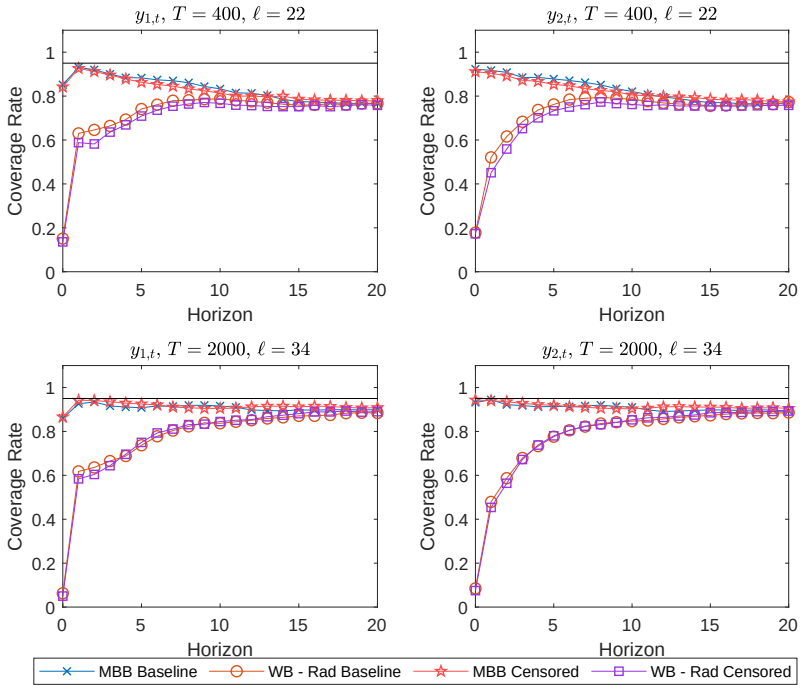

Figure C.41: Coverage rates of 95% confidence intervals for one standard deviation IRFs under DGP3. The solid horizontal line shows the 0.95 target level.

the Rademacher wild bootstrap for normalized IRFs at the 68% and 95% levels, respectively, under DGP3 without and with censoring. Censoring appears to cause a very small decrease in coverage rates for the MBB. However, this is only with small sample sizes, and the overall effects of censoring appear minimal.

Figures C.44 and C.45 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for FEVDs at the 68% and 95% levels, respectively, under DGP3 without and with censoring. Censoring has little impact on these coverage rates.

C.4 Results for Hall's Percentile Intervals

We close by studying the coverage rates of confidence intervals produced by Hall's percentile intervals (Hall (1992) and Lütkepohl (2005, Appendix D)), which were used in Brüggemann, Jentsch, and Trenkler (2016). We continue to produce data from DGP1, DGP2 and DGP3 as in the paper. However, in addition to using standard percentile intervals to form confidence intervals, we also use Hall's percentile intervals.

In the following figures, we show the coverage rates of standard percentile intervals as the baseline results in addition to coverage rates from Hall's percentile intervals. The figures containing the 95% confidence intervals correspond to Figures 2 and 3, Figures 5 through 7, and Figures 9 through 12 in the paper. Consistent with Kilian (1999), we find that Hall's percentile intervals are not systematically better than the standard percentile intervals. In some cases, they appear to be worse. The coverage rates for Hall's percentile intervals are similar to the standard percentile intervals for IRFs when using the MBB. This is especially true with large sample sizes. However, standard percentile intervals appear to be generally better than Hall's percentile intervals in terms of statistical size for FEVDs when using the MBB. When using the Rademacher wild bootstrap, there is very little difference in coverage rates between standard percentile intervals and Hall's percentile intervals.

C.4.1 Results for DGP1

Figures C.46 and C.47 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs at the 68% and 95% levels, respectively, under DGP1 using both the standard percentile interval and Hall's percentile interval. When using the MBB, the two percentile intervals are virtually indistinguishable with large sample sizes. For small sample sizes, Hall's percentile intervals appear to give slightly better coverage at longer horizons. The same is true when using the Rademacher

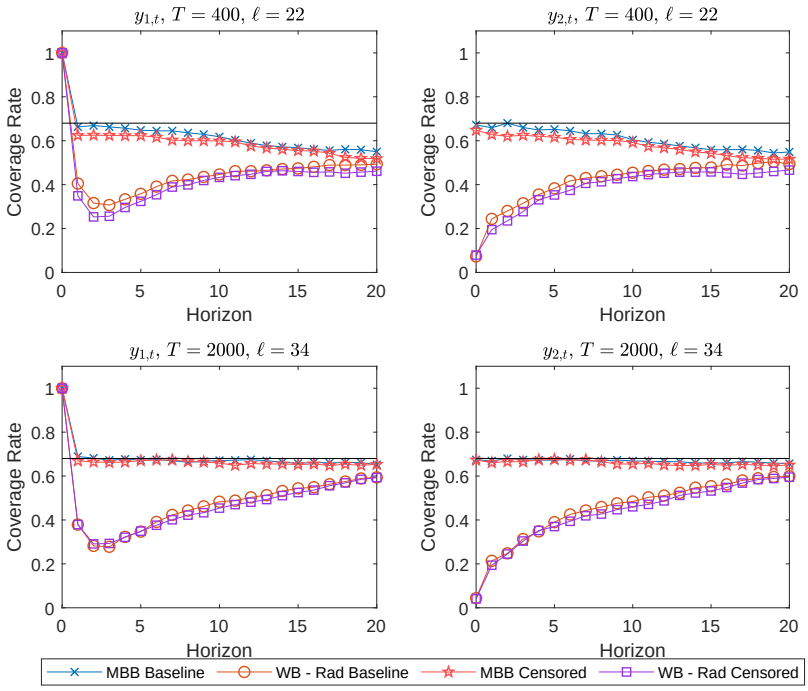

Figure C.42: Coverage rates of 68% confidence intervals for normalized IRFs under DGP3. The solid horizontal line shows the 0.68 target level.

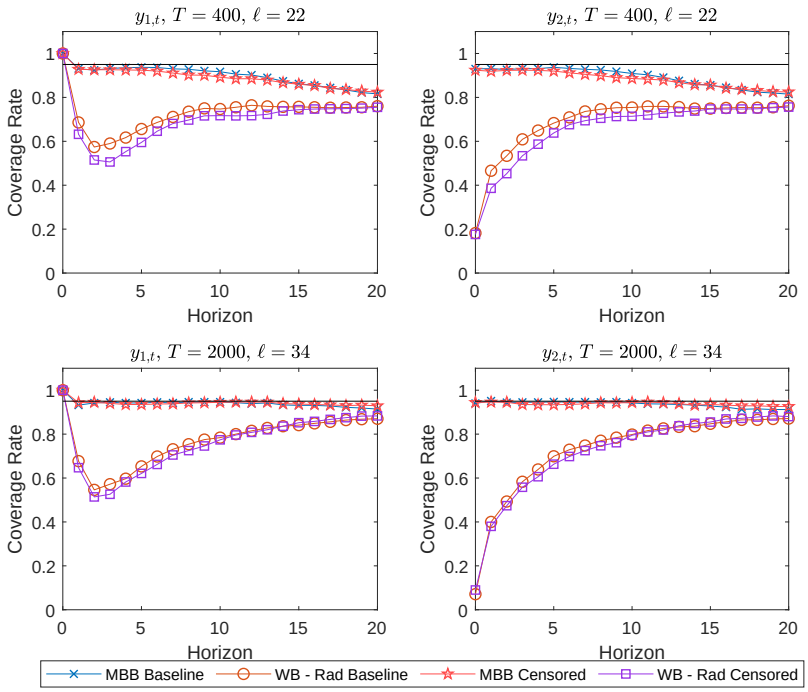

Figure C.43: Coverage rates of 95% confidence intervals for normalized IRFs under DGP3. The solid horizontal line shows the 0.95 target level.

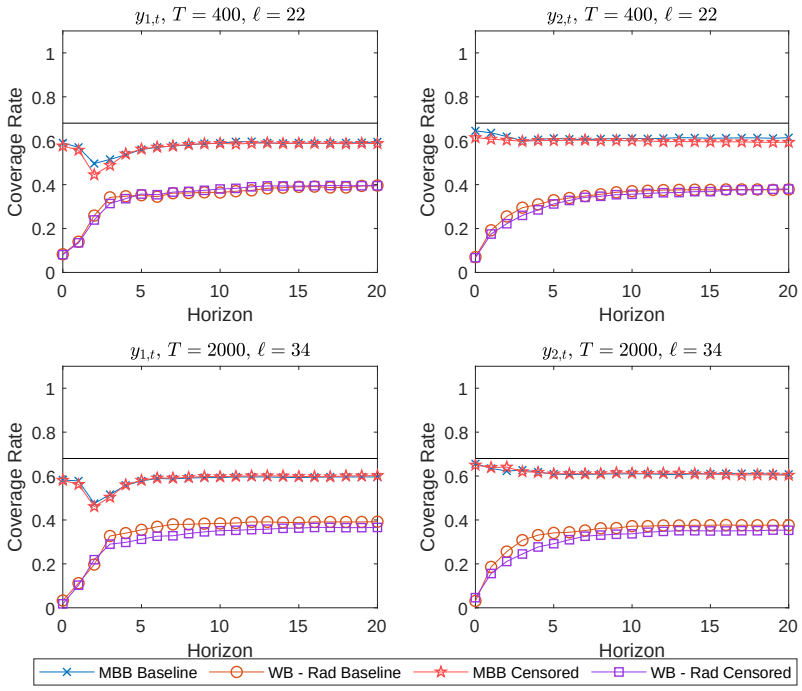

Figure C.44: Coverage rates of 68% confidence intervals for FEVDs under DGP3. The solid horizontal line shows the 0.68 target level.

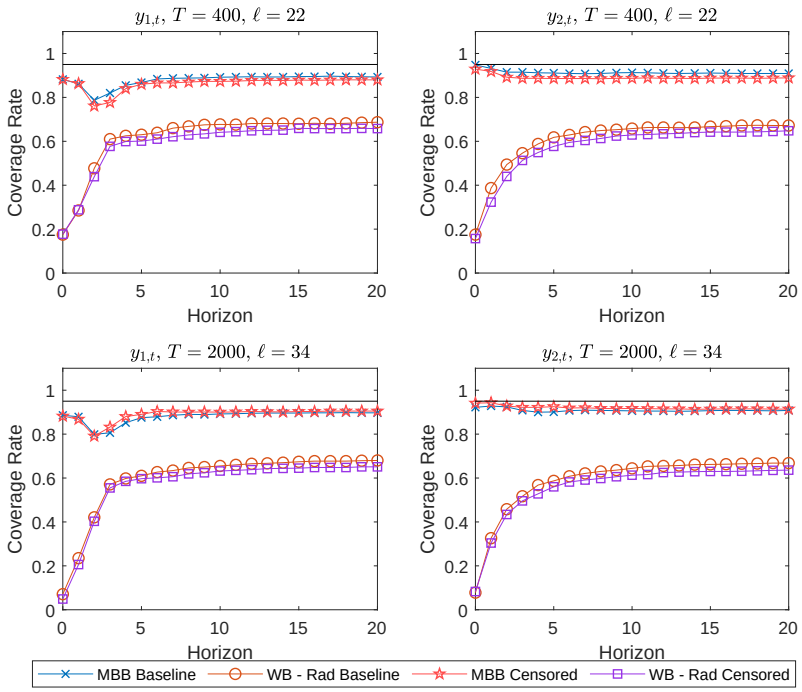

Figure C.45: Coverage rates of 95% confidence intervals for FEVDs under DGP3. The solid horizontal line shows the 0.95 target level.

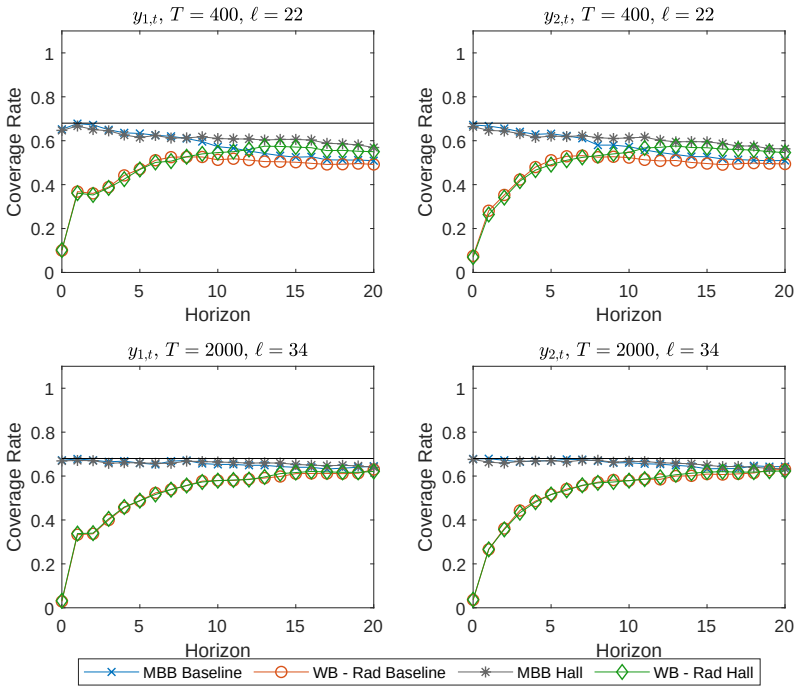

Figure C.46: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP1. The solid horizontal line shows the 0.68 target level.

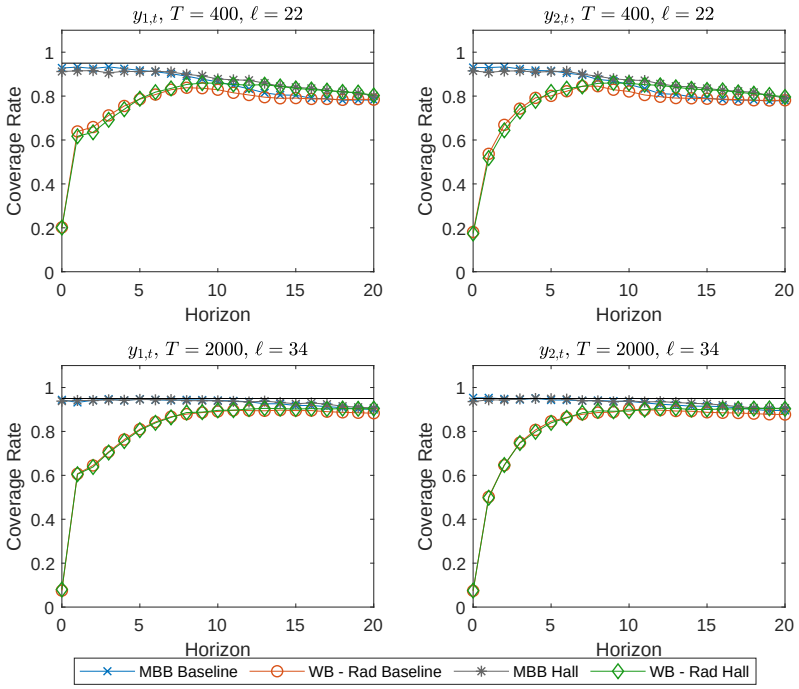

Figure C.47: Coverage rates of 95% confidence intervals for one standard deviation IRFs under DGP1. The solid horizontal line shows the 0.95 target level.

wild bootstrap.

Figures C.48 and C.49 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs at the 68% and 95% levels, respectively, under DGP1 using both the standard percentile interval and Hall's percentile interval. When using the MBB, the two percentile intervals are again virtually indistinguishable with large sample sizes, and Hall's percentile intervals appear to give slightly better coverage at longer horizons for small sample sizes. However, Hall's percentile intervals may also be oversized at small horizons. For the Rademacher wild bootstrap, the two percentile intervals are very similar.

Figures C.50 and C.51 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for FEVDs at the 68% and 95% levels, respectively, under DGP1 using both the standard percentile interval and Hall's percentile interval. When using the MBB with small sample sizes, Hall's percentile interval generally appears to be undersized, and the standard percentile intervals generally give coverage rates closer to the target level.

C.4.2 Results for DGP2

Figures C.52 and C.53 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs at the 68% and 95% levels, respectively, under DGP2 using both the standard percentile interval and Hall's percentile interval. For the MBB with small sample sizes, the results are mixed. The standard percentile interval appears to give better coverage at short horizons, but Hall's percentile interval may give better coverage at long horizons. However, the coverage rates for both intervals are generally similar. For the wild bootstrap, there is very little difference between the two intervals.

Figures C.54 and C.55 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs at the 68% and 95% levels, respectively, under DGP2 using both the standard percentile interval and Hall's percentile interval. For the MBB, the two percentile intervals give very similar coverage in almost all cases. The one exception is for 95% confidence intervals when $T = 400$. In this case, Hall's percentile intervals are generally undersized and the standard percentile interval is closer to the intended level.

Figures C.50 and C.51 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for FEVDs at the 68% and 95% levels, respectively, under

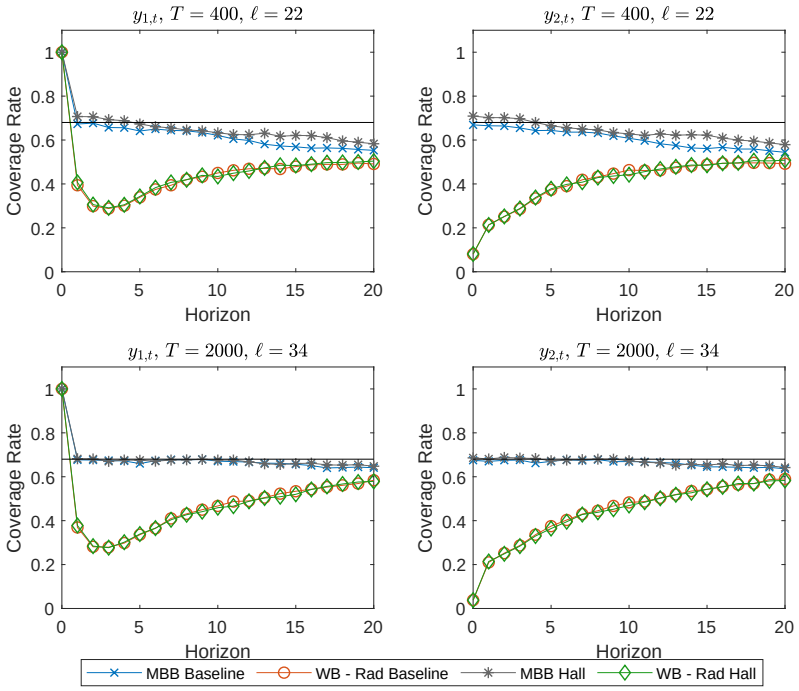

Figure C.48: Coverage rates of 68% confidence intervals for normalized IRFs under DGP1. The solid horizontal line shows the 0.68 target level.

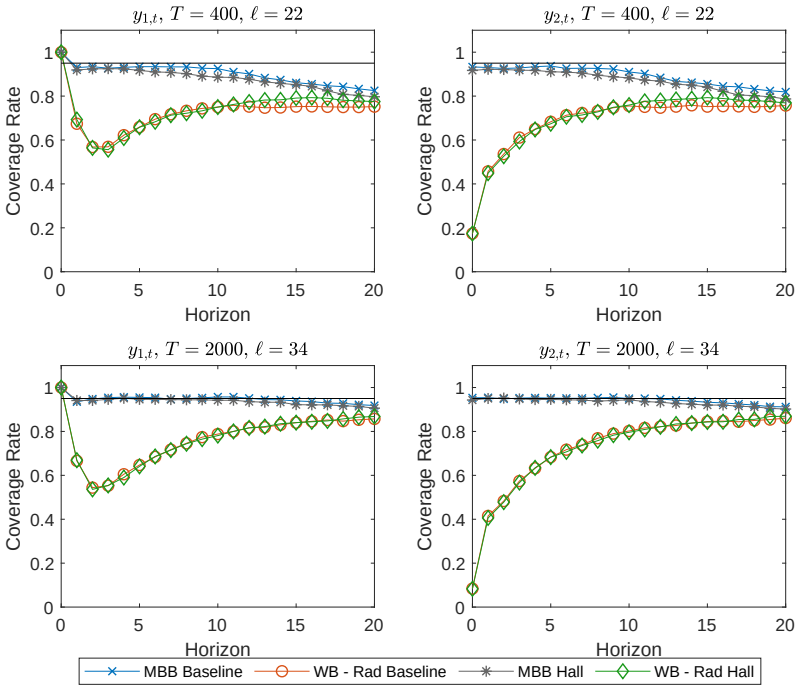

Figure C.49: Coverage rates of 95% confidence intervals for normalized IRFs under DGP1. The solid horizontal line shows the 0.95 target level.

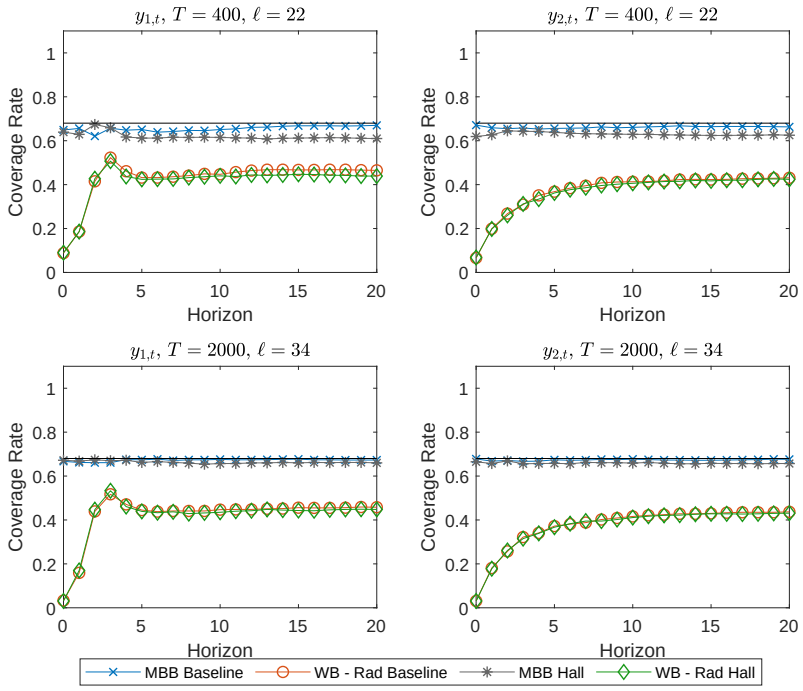

Figure C.50: Coverage rates of 68% confidence intervals for FEVDs under DGP1. The solid horizontal line shows the 0.68 target level.

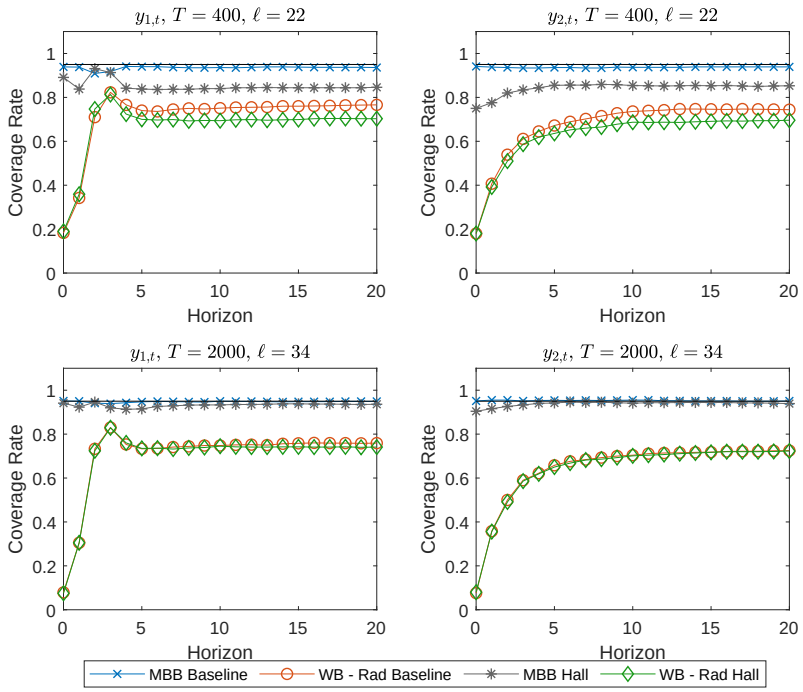

Figure C.51: Coverage rates of 95% confidence intervals for FEVDs under DGP1. The solid horizontal line shows the 0.95 target level.

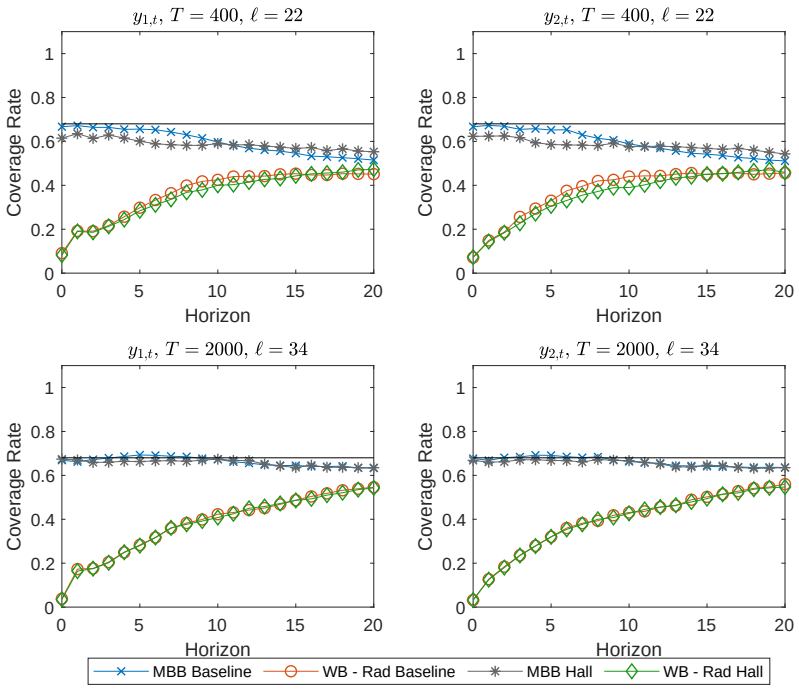

Figure C.52: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP2. The solid horizontal line shows the 0.68 target level.

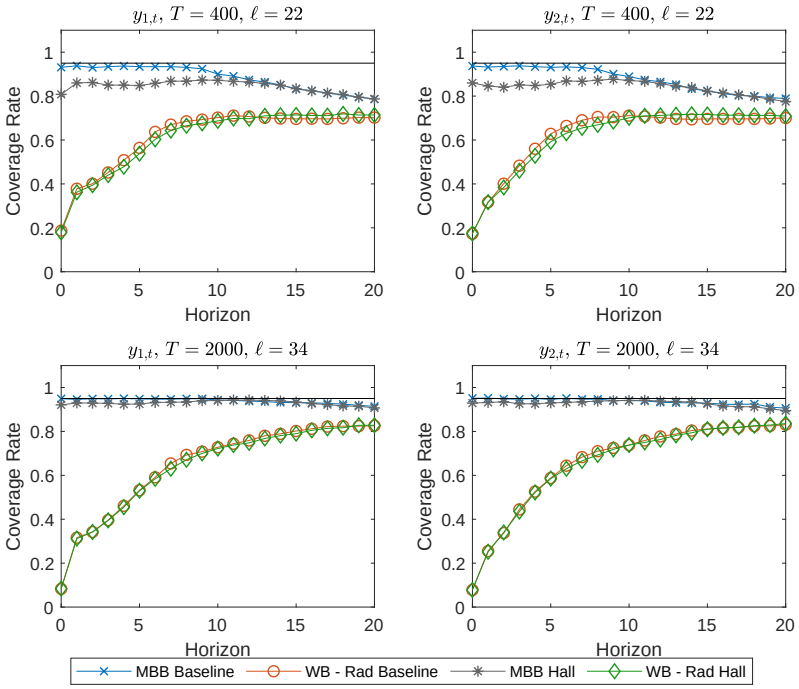

Figure C.53: Coverage rates of 95% confidence intervals for one standard deviation IRFs under DGP2. The solid horizontal line shows the 0.95 target level.

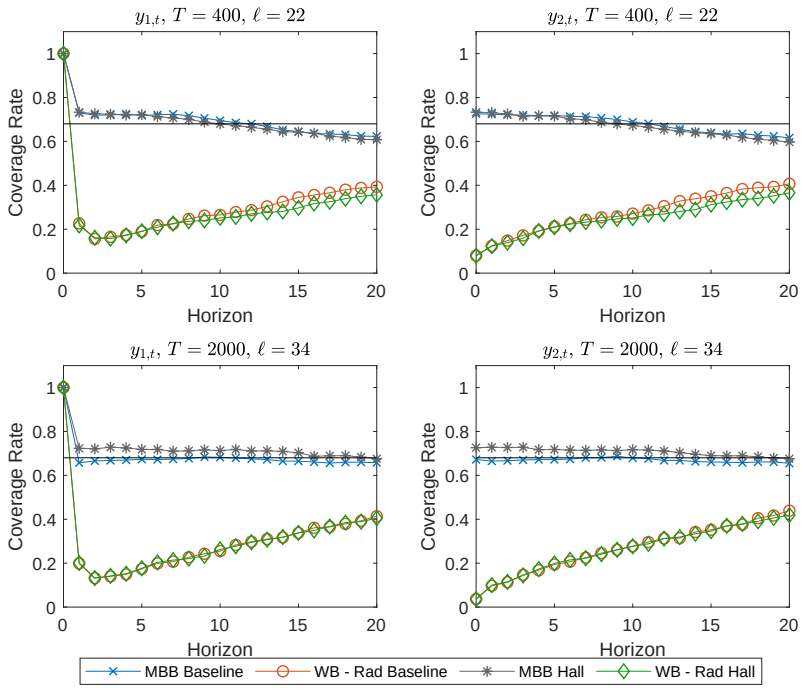

Figure C.54: Coverage rates of 68% confidence intervals for normalized IRFs under DGP2. The solid horizontal line shows the 0.68 target level.

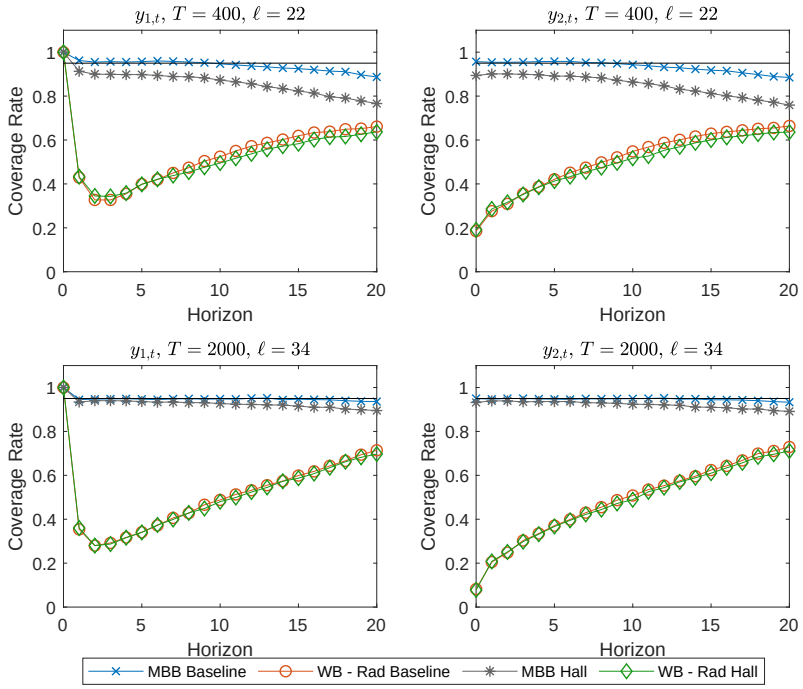

Figure C.55: Coverage rates of 95% confidence intervals for normalized IRFs under DGP2. The solid horizontal line shows the 0.95 target level.

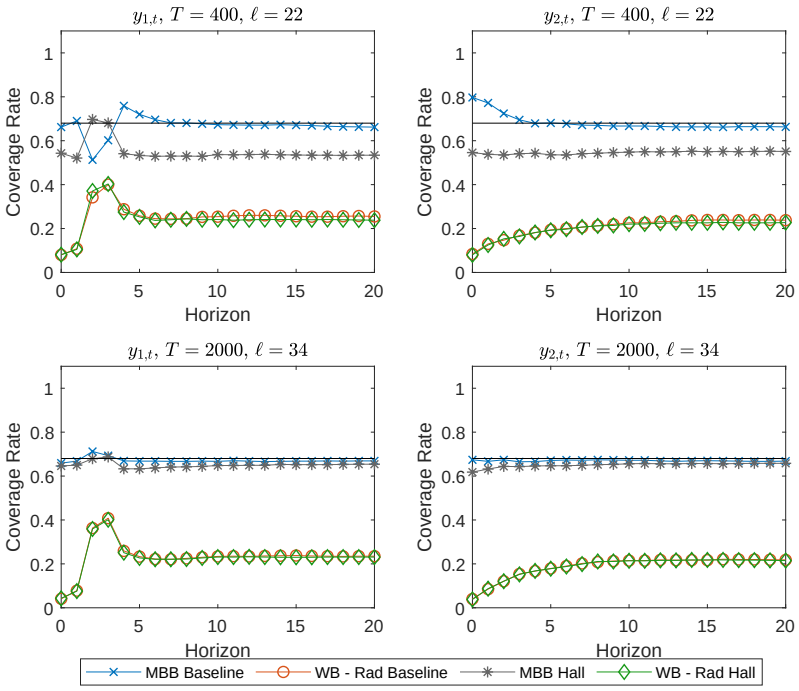

Figure C.56: Coverage rates of 68% confidence intervals for FEVDs under DGP2. The solid horizontal line shows the 0.68 target level.

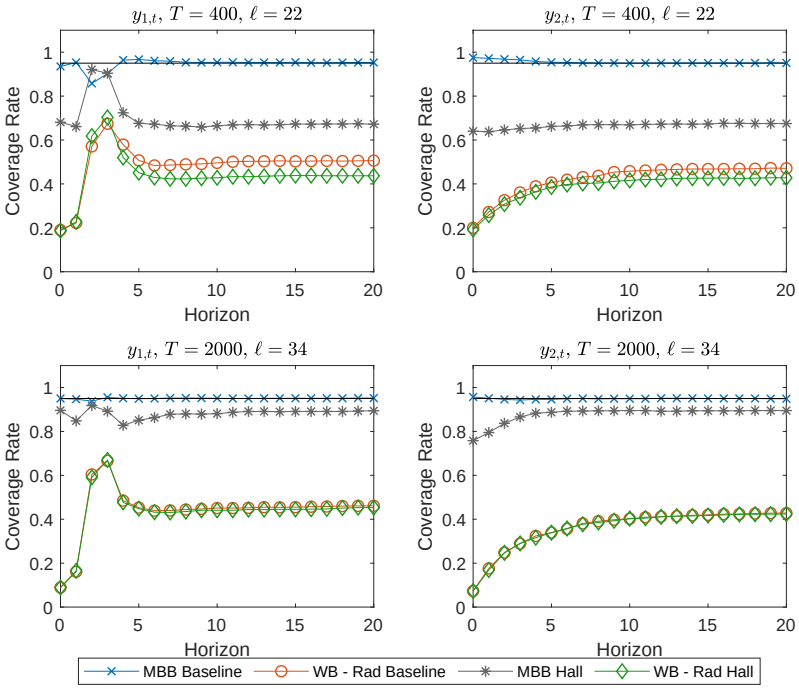

Figure C.57: Coverage rates of 95% confidence intervals for FEVDs under DGP2. The solid horizontal line shows the 0.95 target level.

DGP2 using both the standard percentile interval and Hall's percentile interval. As with DGP1, Hall's percentile interval generally appears to be undersized when using the MBB, and the standard percentile interval generally gives coverage rates closer to the target level. The only exception is for $y_{1,t}$ at horizons 2 and 3.

C.4.3 Results for DGP3

Figures C.58 and C.59 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for the one standard deviation IRFs at the 68% and 95% levels, respectively, under DGP3 using both the standard percentile interval and Hall's percentile interval. They show that the results for DGP1 and DGP3 are quite similar. When using the MBB, the two percentile intervals are virtually indistinguishable with large sample sizes. For small sample sizes, Hall's percentile intervals appear to give slightly better coverage at longer horizons. The same is true when using the Rademacher wild bootstrap.

Figures C.60 and C.61 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for normalized IRFs at the 68% and 95% levels, respectively, under DGP3 using both the standard percentile interval and Hall's percentile interval. When using the MBB, the two percentile intervals are again virtually indistinguishable with large sample sizes. For small sample sizes, Hall's percentile interval appears to give slightly better coverage at longer horizons for 68% confidence intervals but slightly worse coverage for 95% confidence intervals. However, the two intervals generally similar when using the MBB and very similar when using the Rademacher wild bootstrap.

Figures C.62 and C.63 show the confidence interval coverage rates from the MBB and the Rademacher wild bootstrap for FEVDs at the 68% and 95% levels, respectively, under DGP3 using both the standard percentile interval and Hall's percentile interval. As with DGP1 and DGP2, Hall's percentile interval is generally undersized when using the MBB, and the standard percentile interval generally gives coverage rates closer to the target level.

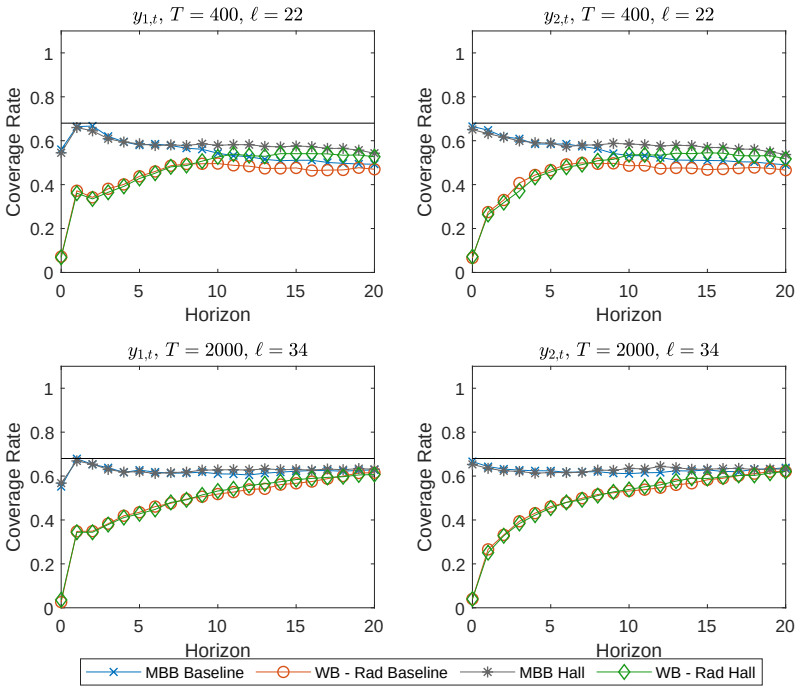

Figure C.58: Coverage rates of 68% confidence intervals for one standard deviation IRFs under DGP3. The solid horizontal line shows the 0.68 target level.

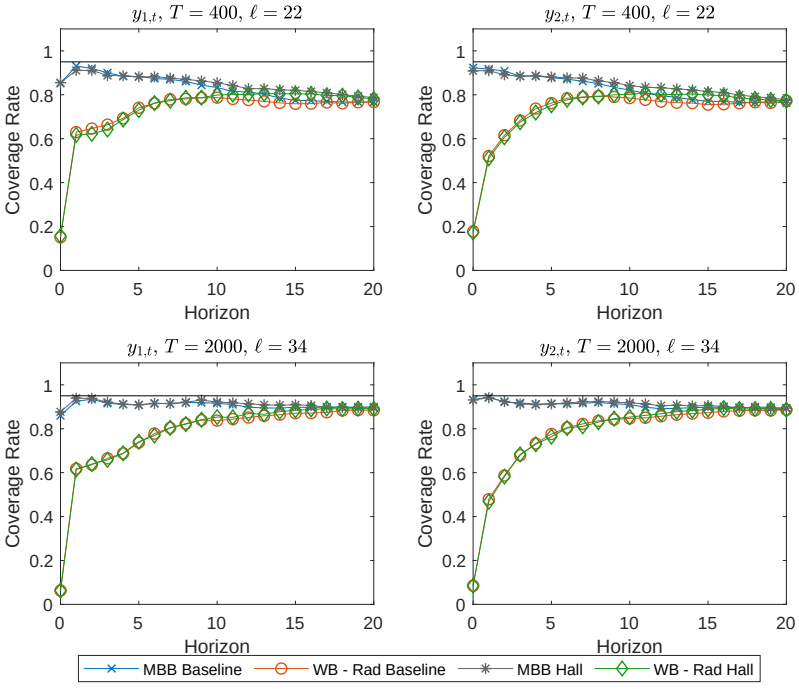

Figure C.59: Coverage rates of 95% confidence intervals for one standard deviation IRFs under DGP3. The solid horizontal line shows the 0.95 target level.

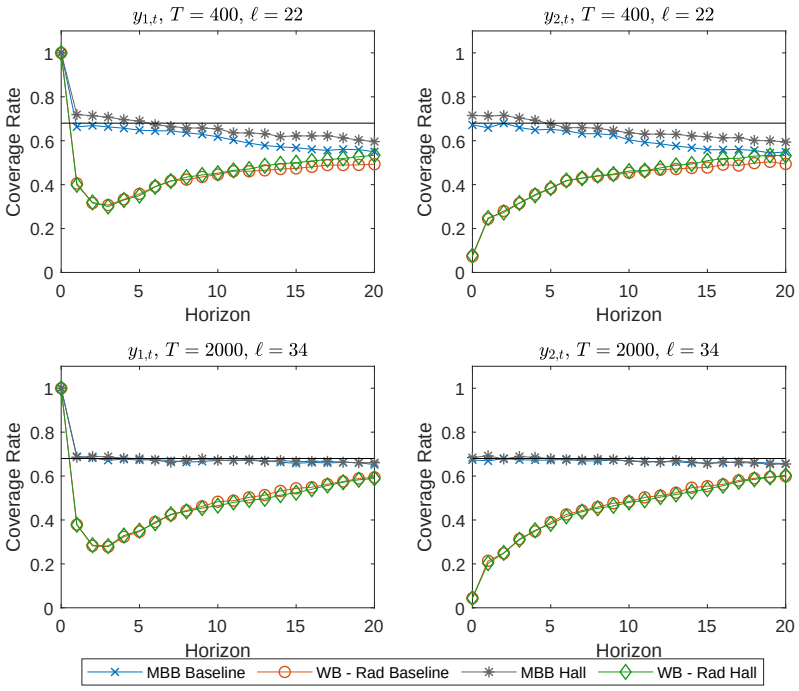

Figure C.60: Coverage rates of 68% confidence intervals for normalized IRFs under DGP3. The solid horizontal line shows the 0.68 target level.

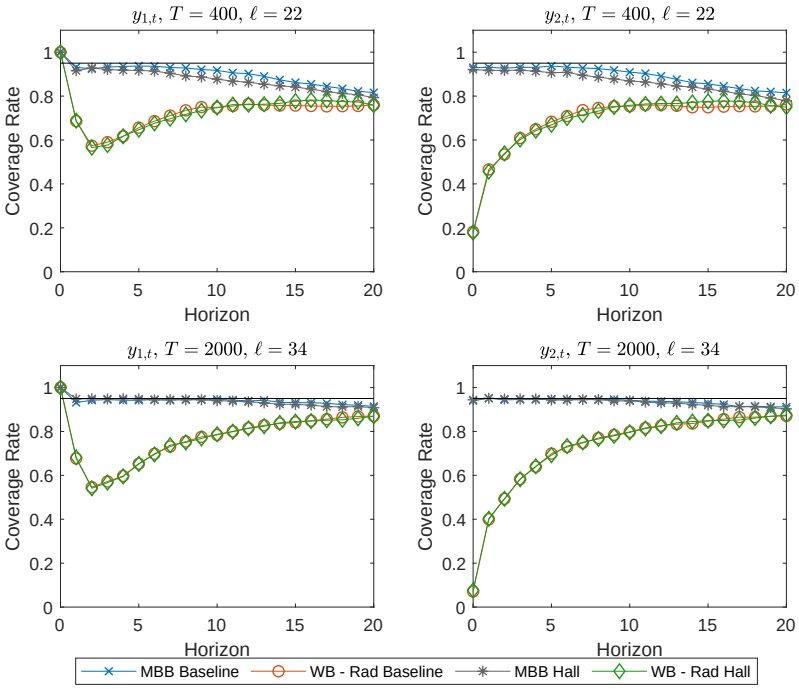

Figure C.61: Coverage rates of 95% confidence intervals for normalized IRFs under DGP3. The solid horizontal line shows the 0.95 target level.

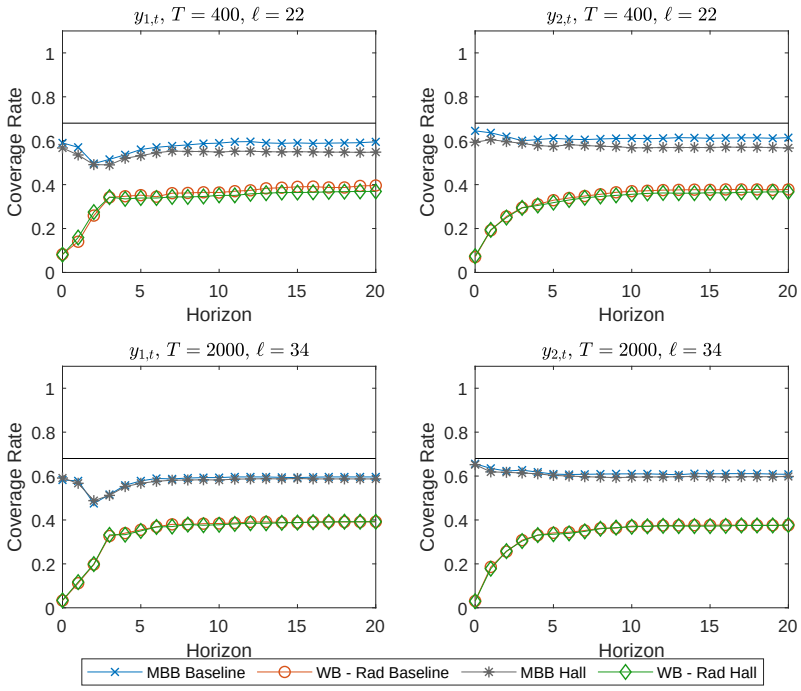

Figure C.62: Coverage rates of 68% confidence intervals for FEVDs under DGP3. The solid horizontal line shows the 0.68 target level.

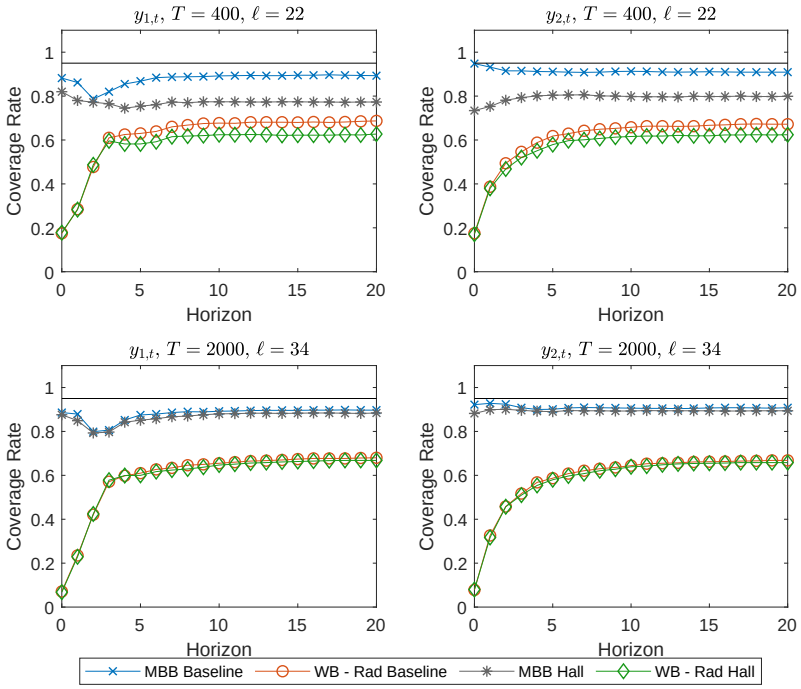

Figure C.63: Coverage rates of 95% confidence intervals for FEVDs under DGP3. The solid horizontal line shows the 0.95 target level.

References

Brüggemann, Ralf, Carsten Jentsch, and Carsten Trenkler. 2016. "Inference in VARs with Conditional Heteroskedasticity of Unknown Form." *Journal of Econometrics* 191 (1):69–85. URL https://doi.org/10.1016/j.jeconom.2015.10.004.

Gertler, Mark and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7 (1):44–76. URL https://doi.org/10.1257/mac.20130329.

Hall, Peter. 1992. *The Bootstrap and Edgeworth Expansion*. Springer, New York. URL https://doi.org/10.1007/978-1-4612-4384-7.

Jentsch, Carsten and Kurt G. Lunsford. 2019. "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States: Comment." *American Economic Review*. Forthcoming. URL https://www.aeaweb.org/articles?id=10.1257/aer.20162011.

Kilian, Lutz. 1999. "Finite-Sample Properties of Percentile and Percentile-t Bootstrap Confidence Intervals for Impulse Responses." *Review of Economics and Statistics* 81 (4):652–660. URL https://doi.org/10.1162/003465399558517.

Lütkepohl, Helmut. 2005. *New Introduction to Multiple Time Series Analysis*. Springer-Verlag. URL https://doi.org/10.1007/3-540-27752-8.

Mertens, Karel and Morten O. Ravn. 2013. "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States." *American Economic Review* 103 (4):1212–1247. URL https://doi.org/10.1257/aer.103.4.1212.