

Bias in Local Projections

Edward P. Herbst and Benjamin K. Johannsen*

June 1, 2023

Abstract

Local projections (LPs) are a popular tool in macroeconomic research. We show that LPs are often used with very small samples in the time dimension and, consequently, that LP point estimates can be severely biased. Under regularity conditions, we derive simple expressions to approximate this bias and propose a way to correct for bias in LPs. Using a medium-scale macroeconomic time-series model, we demonstrate that the bias in point estimates can be economically meaningful. We also show that the same small-sample bias issue can also lead some autocorrelation-robust standard errors to understate sampling uncertainty.

1 Introduction

We show that if a time series is persistent—as is generally the case when researchers are interested in impulse responses—then estimators of impulse responses by local projections (LPs) can be severely biased in sample sizes commonly found in the empirical macroeconomics literature.

Starting with Jordà (2005), LPs have been used by researchers as an alternative to other time series methods, such as vector autoregressions (VARs). We survey the literature and find that, over the past 15 years, LPs have been applied in a variety of settings that are notably different from the setting studied in Jordà (2005). In particular, we find that sample sizes in the time

*edward.p.herbst@frb.gov; benjamin.k.johannsen@frb.gov. The views expressed here are our own and do not necessarily reflect the views of the Federal Reserve Board or anyone else who is associated with the Federal Reserve System. We thank Serena Ng (co-editor), the associate editor, two anonymous referees, Todd Clark, Marco Del Negro, Eric Ghysels, Òscar Jordà, Lutz Kilian, Mikkel Plagborg-Møller, Frank Schorfheide, and Minchul Shin, as well as seminar participants, for helpful comments and discussions. We thank Nicholas Zevanove for excellent research assistance. Any errors are our own.

dimension are typically much smaller than the sample sizes studied in Jordà (2005) and that LPs have become increasingly prevalent when researchers also have a cross section of data (i.e., panel data). Additionally, researchers often approach LPs with identified structural shocks in hand, rather than identifying those shocks as a part of the estimation.¹ We focus on this idealized case in this paper as it is a natural benchmark for understanding the methodology.

We analyze the small-sample bias in LPs using a higher-order expansion of the LP estimator, building on the related work of Kendall (1954), Nagar (1959), Rilstone et al. (1996), and Bao and Ullah (2007). We show that the approximate bias of the LP estimator at horizon h is a function—specifically, a weighted sum—of the (population) impulse response function (IRF) at horizons up to h . As a result, if the data are positively autocorrelated and LP estimators across horizons have the same sign (as is the case for hump-shaped impulse responses), then the least-squares estimators are biased toward zero at every horizon. Additionally, our analysis highlights that the small-sample estimates from LPs are not “local” because the small-sample biases of those estimates depend on the true impulse responses at other horizons.

Using Monte Carlo analysis, we demonstrate that the magnitude of the bias in LPs can be large when sample sizes in the time dimension are similar to those typically found in the empirical macroeconomics literature. We conduct our Monte Carlo simulations using simple, linear data generating processes and an estimated medium-scale macroeconomic VAR. While researchers may be drawn to LPs because they invoke fewer parametric restrictions than other methods, an important standard for this methodology is that it performs well in these simple scenarios.

The expression for the approximate bias that we derive can be used to correct for bias in LP estimators. In our Monte Carlo simulations, our bias-corrected estimators are on average markedly closer to the true values of the impulse responses.² We discuss the tradeoff that researchers face between reducing bias and potentially increasing the mean squared error (MSE) of the estimator. We note that the bias correction does not uniformly increase or decrease MSE in our Monte Carlo simulations.

We extend our analysis to settings using panel data and show that—when using fixed effects—the bias we document persists. We derive formulas to approximate the bias, which could be used

¹In what follows, we always refer to the regressor associated with the LP coefficient as the “shock.”

²Because our bias correction does not completely eliminate small-sample bias, researchers in some settings may prefer methods, such as VARs, that estimate the same impulse responses as LPs (see Plagborg-Møller and Wolf (2021)) and have well understood and effective methods for bias correction (see Kilian (1998)). For a comprehensive comparison among different methods, see Li and Wolf (2022).

to bias-correct LP estimators. We show that increasing the number of entities in the panel will not eliminate the bias.

We also study the downward bias in commonly-used standard errors for LP estimators. Our analysis of standard errors is related to recent work by Montiel Olea and Plagborg-Møller (2021), who suggest that researchers using LPs should use heteroskedasticity-consistent, but not autocorrelation-robust, standard errors. Our focus on finite-sample issues leads us to similar conclusions, but for different reasons. We show that small-sample bias can be an important consideration when using Newey and West (1987) (NW) standard errors. These standard errors rely on estimators of the autocorrelation of the regression score. In finite samples, the estimators of these autocorrelations will be biased for reasons similar to those that cause bias in LP point estimators. In important, empirically realistic settings, the bias will be downward, yielding smaller estimates of standard errors. This result suggests researchers may prefer standard errors that are heteroskedasticity-consistent but not autocorrelation-robust, such as those studied in Huber (1967) and White (1980) (HW). Alternatively, researchers may want to consider alternative heteroskedasticity and autocorrelation robust (HAR) standard errors, like those discussed in Müller (2014), Sun (2014), and Lazarus et al. (2018).

Our paper is related to work by Kilian and Kim (2011), who study the coverage probabilities for confidence intervals for LP estimators using bootstrap methods. Their work focuses on the case when shocks are identified as a part of the LP estimation and uses the block bootstrap to approximate the finite sample distribution of the least-squares estimate. By contrast, our analysis relies on higher-order expansions of the least-squares estimator, which illustrates the reasons that the least-squares estimator is biased and provides a natural bias correction without bootstrapping. More generally, our paper is related to work on bias in least-squares estimators of autocorrelation (such as in Kendall (1954) and Shaman and Stine (1988)), in VARs (such as in Nicholls and Pope (1988) and Pope (1990)), dynamic panel data settings (such as in Nickell (1981) and Hahn and Kuersteiner (2002)), and in generalized method of moments systems (Rilstone et al. (1996), Newey and Smith (2004), Anatolyev (2005), and Bao and Ullah (2007)). We apply this work in our LP setting.

The paper is structured as follows. In Section 2, we present evidence on the ways that LPs are used in the academic literature. In Section 3, we derive an analytic expression for the bias in least-squares estimators of LPs and suggest a bias correction. We extend our analysis to the context of panel data in Section 4. In Section 5, we explore LP bias and bias correction in the context of

a time-series model similar to the one studied in Christiano et al. (2005). We then analyze bias of commonly-used standard errors for least-squares LP estimators in Section 6. Section 7 concludes.

2 Some evidence on the use of LPs

To get a sense of how LPs are used in the literature, we examine the 100 “most relevant” papers citing Jordà (2005) on Google Scholar.³ Google Scholar’s relevance ranking weighs the text of the document, the authors, the source of the publication, and the number of citations. Of these 100 papers, 71 employed LPs in an empirical project (rather than merely citing but not applying LPs).⁴

The focus of this paper is parameter bias associated with short time series, so for each of the studies we recorded the length of the time series, T , in the main LP in each of these papers. About two-thirds of the papers surveyed employed panel data. As we show in Section 4, with entity-specific fixed effects, the time dimension is still the relevant component of the sample size for determining the LP bias. Because many of the panel data sets are unbalanced, constructing a single summary T is challenging. For unbalanced panels, we summarize the size of the time dimension using the mean T across entities, when readily available, or using the largest value of T across entities. In general, our assessment of T is conservative in the sense that it overestimates the time series dimension of the data for many of the LP applications. It is not unusual, for example, to see unbalanced panels that have an average T that is less than half of the time-series dimension of the entire panel or to see robustness exercises that use a small fraction of the data series. In these cases, we use the entire time series dimension of the panel, which biases our estimates of T upward.

Figure 1 displays a histogram of the sample of 71 T s collected in our literature review. The median T (the red dash-dotted line) is around 95. These sample sizes are significantly smaller than those typically used in empirical macroeconometrics papers, as most of the papers surveyed here use the increasingly popular strategy of using observed shocks, such as the monetary policy shocks of Romer and Romer (2004), rather than identified shocks from a VAR, as in Jordà (2005). Constructing these observed shocks is often difficult and costly, so the time series are typically short in length.

The application of LPs to such short time series does not seem to have been anticipated in the early literature on LPs. In fact, the Monte Carlo studies in Jordà (2005) used $T = 300$ and

³We conducted this search in October 2019. See Appendix F for the list of citations.

⁴If a paper appeared as both a working paper and a published paper, we excluded the working paper version from our analysis.




Figure 1: T is small in the literature using LPs.

$T = 496$ (the orange dashed lines in Figure 1). Less than 6 percent of the surveyed studies use sample sizes at least that large. While many studies in our survey use annual or quarterly data, Jordà (2005) used monthly data. In general, however, increasing T by using monthly data rather than quarterly or annual data will not eliminate the issue of small-sample bias in LPs because the monthly series are likely to be more persistent, and the bias in LPs is more severe when the data are more persistent.

3 Bias in LPs

In this section, we analyze the bias in LPs using a Nagar (1959) expansion. We focus on LPs that estimate the impulse response of a macroeconomic variable, y_t , to a structural shock, ε_t . As mentioned in the introduction, the structural shock is observed, and inference is conducted, using linear, least-squares regression. The LP model is the set of regression models indexed by the impulse response horizon, h ,

$$y_{t+h} = \alpha_h + \beta'_h x_t + u_{h,t}, \quad h = 0, \dots, H. \quad (1)$$

where $x_t = [\varepsilon_t, c'_{t-1}]'$ contains the structural shock and (time $t-1$) controls. For ease of exposition, we assume the control vector, c_{t-1} , is not empty—in practice, researchers use c_{t-1} to condition inference on information available at time $t-1$. The first elements of the coefficient vectors $\{\beta_h\}_{h=0}^H$ trace out the impulse response of interest. We denote the $H+1$ vector describing the

impulse response by θ with elements θ_h for $h = 0, \dots, H$. That is, $\theta = [\theta_0, \theta_1, \dots, \theta_H]'$, where $\theta_h = \beta_{h,1}$ (the first element of β_h). As in the empirical macroeconomics literature, we estimate each β_h using ordinary least squares. We denote the estimator of the β_h by $\widehat{\beta}_{h,LS}$ and the estimator of the impulse response by $\widehat{\theta}_{LS}$.⁵

3.1 The approximate bias of the least-squares estimator

To ease notation and without loss of generality, we assume throughout the paper that the data we consider have a mean of zero. To facilitate the derivation of the bias we make the following assumption about the time series properties of y_t , ε_t and c_{t-1} .

Assumption 1. *The series $w_t = [y_t, \varepsilon_t, c_t']'$ is strictly stationary and ergodic. The series has a purely nondeterministic Wold representation with innovations $\omega_t = [\varepsilon_t, \nu_t']'$,*

- ε_t is independent of ε_{t+j} for all $j \neq 0$,
- ε_t is independent of ν_{t+j} for all j ,
- $\mathbb{E}[x_t x_t']$ is invertible, and $\mathbb{E}[\|x_t x_t'\|]$ is finite, where $\|\cdot\|$ denotes the 2-norm.

Our assumption about the properties of ε_t are meant to formalize what is meant by a structural shock. The linearity of w_t in all shocks is a stark assumption that facilitates the derivation of closed-form expressions for the bias in LPs. Linearity in the shock of interest, ε_t , also represents an idealized case where “direct causal inference”—as in Nakamura and Steinsson (2018)—is possible. Notably, in the case where w_t is Gaussian and ε_t is *iid*, the conditions of Assumption 1 are satisfied. We emphasize that violations of Assumption 1 do not imply that LPs are not biased in small samples. Rather, violations of Assumption 1 make it difficult to characterize the small-sample bias. In particular, our analytic expressions derived below depend on homoskedastic shocks. Heteroskedasticity would complicate the analysis. In addition, we make the following assumption about the regression errors from the LP.

Assumption 2. *The regression error, $u_{h,t}$, is independent of ε_{t-j} and ν_{t-j-1} for all $j \geq 0$.*

This assumption formalizes what we mean by conditioning inference on information available at time $t-1$ by including c_{t-1} in the LP. This guarantees that the regression error will be at most an $MA(h)$

⁵It is also common in the LP literature to use $y_{t+h} - y_{t-1}$ as the left-hand side variable. If $y_t - y_{t-1}$ can be included in w_t , as defined in Assumption 1, and so long as our assumptions still hold, then Analytical Result 1 holds. Analogous results to other results in the paper can be derived using straightforward modifications to our analysis.

in ν_t and an $MA(h-1)$ in ε_t , representing an idealized case where the controls perfectly control for time $t-1$ information. This effectively truncates the terms in the approximate bias, simplifying the subsequent analysis considerably. As with Assumption 1, a violation of Assumption 2 does not imply a lack of bias; we discuss implications of departures from this assumption later in the paper. We can now state our main analytical result regarding the approximate bias in LPs.

Analytical Result 1 (Expression for the approximate bias in LPs). *Under Assumptions 1–2, the approximate bias of the LP in (1) is given by*

$$\mathbb{E} \left[\widehat{\theta}_{h,LS} \right] - \theta_h = B_{h,LP} + O \left(T^{-3/2} \right), \quad (2)$$

where

$$B_{h,LP} = -\frac{1}{T-h} \sum_{j=1}^h \left(1 + \text{tr} \left\{ \Sigma_{c,0}^{-1} \Sigma_{c,j} \right\} \right) \theta_{h-j} \text{ for } h > 0, \quad (3)$$

$\Sigma_{c,j} = E [c_{t-j} c_t']$, and $\text{tr}\{\cdot\}$ is the trace operator. Additionally, $B_{0,LP} = 0$.

The derivation of equations (2) and (3) relies on the methodology of Bao and Ullah (2007).⁶ A detailed derivation is in Appendix A. We refer to $B_{h,LP}$ as the approximate bias. Several remarks regarding Analytical Result 1 are in order. First, $|B_{h,LP}|$ is a decreasing function of T ; for fixed h , the least-squares estimator is consistent. Second, $B_{h,LP}$ is a function of the impulse response coefficients at *all horizons up to h* . Intuitively, the data generating process affects the bias at similar horizons in similar ways. Conditioning inference on information available at time $t-1$, however, truncates the terms that contribute to the bias by eliminating the autocorrelation in the regression errors at sufficiently large lags.⁷ Third, the contribution of θ_{h-j} to $B_{h,LP}$ is scaled by

⁶Some additional technical remarks regarding our analytical framework are in order. The least-squares estimator can be cast as a k -class estimator—see Theil (1961). The characteristics of k -class estimators are important for the study of parameter bias (both finite sample and asymptotic) in simultaneous equations models, in particular for the study of (weak) instrumental variables. Sawa (1972) shows that for many k -class estimators the first moment may not even exist, rendering the approximations in this paper inaccurate—see also Srinivasan (1970). We acknowledge this limitation but note that with the additional assumption of normality—following Sawa (1972)—one can guarantee the existence of the first moment of $\widehat{\theta}_{LS}$. Additionally, one can think of the stochastic terms used in the procedure of Bao and Ullah (2007) as offering an approximation of the true finite sample distribution of $\widehat{\theta}_{LS}$. This approximation may have finite moments even when the exact distribution of $\widehat{\theta}_{LS}$ does not. We always assume $\mathbb{E}[\widehat{\theta}_{h,LS}]$ exists.

⁷An earlier version of the paper examined the bias in the alternative polar case: LPs without controls. This LP violates Assumption 2 and requires a slightly different derivation. In the AR(1) model here, the bias is given by the following: $(1 - \rho^{-(T-h)})(\rho^T - \rho^{h+1}) / (1 - \rho)$. This expression can be smaller or larger than the approximate bias in Equation 6, depending on ρ , h , and T . Appendix A.2 and Herbst and Johansen (2021) provide more details on the bias in LPs without controls.

$1 + \text{tr} \left\{ \Sigma_{c,0}^{-1} \Sigma_{c,j} \right\}$. When controls are relatively persistent this scaling factor is relatively large. Fourth, intuitively $B_{0,LP} = 0$ because the structural shock, ε_t , is *iid* and $u_{0,t}$ is independent over time.

It is useful to parse the contribution to $B_{h,LS}$ that is due to estimating α_h . To this end, we derive the approximate bias when α_h is known.

Analytical Result 2 (Expression for the approximate bias in LPs when α_h is known). *Under Assumptions 1–2, and under the additional assumption that α_h is known, the expression for the bias in (1) is given by*

$$\mathbb{E} \left[\widehat{\theta}_{h,LS} \right] - \theta_h = -\frac{1}{T-h} \sum_{j=1}^h \text{tr} \left\{ \Sigma_{c,0}^{-1} \Sigma_{c,j} \right\} \theta_{h-j} + O \left(T^{-3/2} \right), \quad (4)$$

where $\Sigma_{c,j}$ and $\text{tr}\{\cdot\}$ are defined as they were in Analytical Result 1. Additionally, $\mathbb{E} \left[\widehat{\theta}_{0,LS} \right] - \theta_0 = O(T^{-3/2})$.

From equation (4), it is immediately apparent that estimating α_h adds the terms $-(T-h)^{-1} \sum_{j=1}^h \theta_h$ to the approximate bias. When $\theta_h > 0$ and $\text{tr} \left\{ \Sigma_{c,0}^{-1} \Sigma_{c,j} \right\} > 0$, as is the case in many macroeconomic settings, estimating α_h increases the magnitude of the (negative) approximate bias.

3.2 An AR(1) example

Here we examine the approximate bias in LPs using a canonical AR(1) model for y_t . In particular, we assume that

$$y_t = \alpha + \theta_0 \varepsilon_t + \rho y_{t-1} + \nu_t \quad \text{with} \quad \varepsilon_t \stackrel{iid}{\sim} N(0, \sigma_\varepsilon^2) \text{ and } \nu_t \stackrel{iid}{\sim} N(0, \sigma_\nu^2). \quad (5)$$

Further, we assume that $\rho \in (0, 1)$. The researcher observes $\{(y_t, \varepsilon_t)\}_{t=1}^T$ and estimates the LP defined in equation (1) with $c_{t-1} = y_{t-1}$ using ordinary least squares. This example satisfies Assumptions 1 and 2.

3.2.1 Explicit approximate bias

Noting that $\theta_h = \theta_0 \rho^h$, it is easy to see from Analytical Result 1 that the approximate bias of the LP is given by

$$B_{h,LP} = -\frac{1}{T-h} \sum_{j=1}^h (1 + \rho^j) \theta_0 \rho^{h-j} = -\frac{\theta_0}{T-h} \left(\frac{1 - \rho^h}{1 - \rho} + h \rho^h \right). \quad (6)$$

The term $(1-\rho^h)/(1-\rho)$ accounts for the bias arising from estimating the constant, α_h , and the term $h\rho^h$ accounts for the bias arising from the persistence of the control, ρ . If $\theta_0 > 0$ (a normalization), $B_{h,LP} \leq 0$ for all T and $0 < h < T$. That is, the approximate bias is always downward. It is not straightforward to say at which horizon the bias is largest: $\operatorname{argmax}_{0 < h < T} |B_{h,LP}|$ is a complicated function of ρ and T because the component of the bias associated with estimating the constant is generally decreasing in magnitude as ρ increases, whereas the component associated with persistent regressors is increasing in magnitude.

3.2.2 Quality of the approximation

Here we analyze how well $B_{h,LP}$ approximates the small-sample bias of $\hat{\theta}_{h,LS}$ by comparing the approximate bias to the exact bias calculated using Monte Carlo simulations—using equation 5—for various ρ and T . For this exercise, we set $\sigma_\varepsilon = \sigma_\nu = 1$ and $\theta_0 = 1$. The variance parameters do not appear in $B_{h,LP}$.⁸ Figure 2 shows $B_{h,LP}$, along with the bias from the Monte Carlo. The figure shows results for $\rho \in \{0.90, 0.95, 0.99\}$ and $T \in \{50, 100, 200\}$.

For $\rho = 0.9$ and $\rho = 0.95$, $B_{h,LP}$ is a good approximation of the exact small-sample bias in $\hat{\theta}_{h,LS}$ for all h shown. The quality of the approximation improves somewhat as T increases. Clearly, when $\rho = 0.99$, $B_{h,LP}$ is not as good of an approximation as it is for smaller values of ρ . Nevertheless, even with $\rho = 0.99$ and $T = 50$, $B_{h,LP}$ captures salient features of the small-sample bias, including that it is growing in magnitude in h over the values of h shown. We conclude that for empirically relevant sample sizes, $B_{h,LP}$ offers a reasonable approximation of the bias in LPs, though the quality of the approximation is somewhat worse for smaller values of T and larger values of ρ . In Appendix C, we show that similar results hold for an AR(2) model with hump-shaped impulse responses.

3.2.3 Comparison to a parametric approach

LP estimators are often compared to estimates of impulse responses from VARs. These discussions are often centered around the different asymptotic bias and variance tradeoffs associated with the (relatively flexible) LP and (tightly constrained) VAR estimators. A natural question is how the small-sample bias of LP estimators compares to the small-sample bias of VAR estimators.⁹ To

⁸In Appendix C, we show that setting $\sigma_\varepsilon = 10\sigma_\nu = 1$ or $10\sigma_\varepsilon = \sigma_\nu = 1$ has little effect on the quality of the approximation offered by $B_{h,LP}$. In Appendix C, we also analyze approximate bias in the context of the results from a Monte Carlo exercise under the assumption that the means of the data are known.

⁹The small-sample bias of VAR estimators has been explored by Nicholls and Pope (1988), Pope (1990), and others. The non-recursive structure of LPs makes it difficult to extend these result to the settings we consider.

address this, we compare $B_{h,LP}$ to the small-sample bias arising from estimating the coefficients α , θ_0 , and ρ in equation (5) using ordinary least squares and computing the h -period impulse response as $\widehat{\theta}_{0,LS}(\widehat{\rho}_{LS})^h$. This approach, which we call the ‘‘AR approach,’’ is analogous to the estimation of impulse responses using VARs when the researcher has the time series data for the structural shock of interest.

To compute the approximate bias of the AR approach, we again use the methodology of Bao and Ullah (2007). The resulting approximation is

$$\mathbb{E} \left[\widehat{\theta}_{0,LS} \widehat{\rho}_{LS}^h \right] - \theta_0 \rho^h = B_{h,AR} + O(T^{-3/2}) \quad (7)$$

where

$$B_{h,AR} = -\frac{\theta_0 \gamma}{T-1} \left[h \rho^{h-1} (1+3\rho) - \frac{1}{2} h (h-1) \rho^{h-2} (1-\rho^2) \right] \text{ and } \gamma = \frac{\sigma_\nu^2}{\sigma_\nu^2 + \theta_0^2 \sigma_\varepsilon^2}. \quad (8)$$

The details of the derivation are in Appendix B. To understand the expression for $B_{h,AR}$, it is useful to first analyze γ . Note that $\gamma \in [0, 1]$ is the ratio of the variance of $u_{0,t}$ to the one-step-ahead variance of y_t . Larger values of γ increase the magnitude of $|B_{h,AR}|$. That is, when ε_t explains relatively little of the variation in y_t , γ and thus $|B_{h,AR}|$ are larger. Conversely, when instead ν_t explains relatively little of the variation in y_t , γ and thus $|B_{h,AR}|$ are smaller. As the variance of the regression error collapses, the AR approach is less biased.

Several remarks about the relationship between $B_{h,AR}$ and $B_{h,LP}$ are in order. First, $B_{0,AR} = B_{0,LP} = 0$. Second, in general, the values of $B_{h,AR}$ and $B_{h,LP}$ will not coincide for $h > 0$. Third, for $\rho \in (0, 1)$, $B_{h,AR}$ cannot be signed for $h > 0$, unlike $B_{h,LP}$, which is never positive. Fourth, γ affects $B_{h,AR}$ but does not enter $B_{h,LP}$. That is, the approximate bias in the LP is not affected by the amount of variation in y_t that is explained by ε_t . While the magnitude and sign of $B_{h,AR} - B_{h,LP}$ depend on parameters of the AR(1) process and h , it is useful to consider a particular value of $\gamma = \frac{1+\rho}{1+3\rho}$, which sets $B_{1,AR} = B_{1,LP}$. Then, for $\rho \in (0, 1)$ and $h > 1$,

$$B_{h,AR} = -\frac{1}{T-1} \left(h \rho^h + h \rho^{h-1} \right) + \frac{1}{T-1} \frac{1}{2} h (h-1) \rho^{h-2} \frac{1+\rho}{1+3\rho} (1-\rho^2) > B_{h,LP}. \quad (9)$$

So, for any $\gamma < \frac{1+\rho}{1+3\rho}$, $B_{h,AR} - B_{h,LP} > 0$ for all $h > 1$. This result illustrates that there are parameter values for which the AR approach can be uniformly less negatively biased than the LP approach. However, it is difficult to know whether these conditions hold without knowledge of the data generating process. Above all, the AR(1) example illustrates that the small-sample bias in LPs is different from the small-sample bias in VARs.¹⁰

¹⁰In VARs, researchers commonly employ the parametric bootstrap to correct for bias (see, for example, Kilian

3.3 A bias corrected estimator

Analytical Result 1 suggests a bias-corrected estimator of θ_h using plug-in estimators of θ_j and $\Sigma_{c,j}$. With enough data to calculate $\hat{\theta}_{h,LS}$, all of the needed estimates of θ_j and $\Sigma_{c,j}$ are easily computed. Notably, the researcher could bias correct the coefficients using the least-squares estimates of θ_j . We denote the estimator of θ_h constructed in this manner as $\hat{\theta}_{h,BC}$:

$$\hat{\theta}_{h,BC} = \hat{\theta}_{h,LS} + \frac{1}{T-h} \sum_{j=1}^h \left(1 + \text{tr} \left\{ \hat{\Sigma}_{c,0}^{-1} \hat{\Sigma}_{c,j} \right\} \right) \hat{\theta}_{h-j,LS}, \quad h = 1, \dots, H, \quad (10)$$

$$\text{where } \hat{\Sigma}_{c,j} = \frac{1}{T-h-j} \sum_{t=j+1}^{T-h} (c_{t-j} - \bar{c})(c_t - \bar{c})' \text{ for } j = 0, \dots, h \quad \text{and} \quad \bar{c} = \frac{1}{T-h} \sum_{t=1}^{T-h} c_t.$$

The initial condition is given by $\hat{\theta}_{0,BC} = \hat{\theta}_{0,LS}$. Alternatively, the researcher could use bias corrected estimates of θ_h as plug-in estimators in the expression for the approximate bias. We implement this bias corrected estimator by again setting $\hat{\theta}_{0,BCC} = \hat{\theta}_{0,LS}$ and proceeding sequentially to construct:

$$\hat{\theta}_{h,BCC} = \hat{\theta}_{h,LS} + \frac{1}{T-h} \sum_{j=1}^h \left(1 + \text{tr} \left\{ \hat{\Sigma}_{c,0}^{-1} \hat{\Sigma}_{c,j} \right\} \right) \hat{\theta}_{h-j,BCC}, \quad h = 1, \dots, H. \quad (11)$$

If all $\hat{\theta}_{j,LS}$ have the same sign and $\text{tr} \left\{ \hat{\Sigma}_{c,0}^{-1} \hat{\Sigma}_{c,j} \right\} > -1$ for all $j = 0, \dots, h$, then $\hat{\theta}_{h,BCC}$ will be larger in magnitude than $\hat{\theta}_{h,BC}$ for all $h > 1$. Given $\text{tr} \left\{ \hat{\Sigma}_{c,0}^{-1} \hat{\Sigma}_{c,j} \right\} > -1$, both $\hat{\theta}_{h,BC}$ and $\hat{\theta}_{h,BCC}$ are positively weighted sums of $\hat{\theta}_{j,LS}$ for $j = 0, \dots, h$. Because of the recursive definition in equation (11), the weights used to calculate $\hat{\theta}_{h,BCC}$ will be larger than those associated with $\hat{\theta}_{h,BC}$, and so $|\hat{\theta}_{h,BCC}| > |\hat{\theta}_{h,BC}|$ as long as all of the $\hat{\theta}_{j,LS}$ have the same sign. In general, $\hat{\theta}_{h,BCC}$ tends to exhibit superior properties when compared to $\hat{\theta}_{h,BC}$. Therefore, in our subsequent analysis, we will focus on $\hat{\theta}_{h,BCC}$.

Here, we analyze how well $\hat{\theta}_{h,BCC}$ performs in the context of our AR(1) example data generating process. The left column of Figure 3 shows the Monte Carlo estimate $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$, for different values of ρ and T , as a function of h . The bias corrected estimator does not eliminate the bias entirely, particularly for low T and large ρ . That said, it reduces the bias relative to the least-squares estimator. For example, when $\rho = 0.99$ and $T = 50$, $\hat{\theta}_{h,BCC}$ still exhibits a bias of around -0.3 for moderate h , while the corresponding bias of $\hat{\theta}_{h,LS}$ is around -0.6 (see Figure 2). So $\hat{\theta}_{h,BCC}$ represents a substantial improvement.

(1998)). However, in an LP, researchers would need to use, for example, a block bootstrap or the bootstrap discussed in Montiel Olea and Plagborg-Møller (2021). With sample sizes typically seen the literature that uses LPs, block bootstrapping methods are likely to perform poorly.

The right column of Figure 3 shows the MSE ratio of $\widehat{\theta}_{h,BCC}$ relative to $\widehat{\theta}_{h,LS}$. A number larger than unity indicates that the MSE of $\widehat{\theta}_{h,BCC}$ is larger than the MSE of $\widehat{\theta}_{h,LS}$. While $\widehat{\theta}_{h,BCC}$ does a better job, on average, of correcting for the bias when ρ is smaller, the relative MSE is smaller for a larger ρ . Intuitively, when ρ is large, the bias of the LP estimator is also large, and the mean bias correction is enough to reduce the MSE, even though the bias-correction procedure introduces volatility into the estimator. Overall, the MSE metric does not uniformly favor either $\widehat{\theta}_{h,LS}$ or $\widehat{\theta}_{h,BCC}$.

It ought to be mentioned that the potential for asymptotic bias from parametric models may be a key reason that many investigators turn to LPs in the first place. It stands to reason that these researchers may have preferences that give more weight to the bias reduction than to lower-variance estimators. Thus $\widehat{\theta}_{BCC}$, which reduces the finite sample component of this bias, may be attractive even if it exhibits a higher variance (and higher MSE) than its least-squares counterpart.

4 Extension to panel data

Here, we extend our analysis to the setting of panel data with a fixed number of entities, I . We let the subscript i indicate data for a particular panelist (for example, $y_{i,t}$). We consider the setting where the entities have common slope coefficients but entity-specific intercepts, which is a common set-up in the LP literature using panel data. That is,

$$y_{i,t+h} = \alpha_{i,h} + [\varepsilon_{i,t}, c'_{i,t-1}, c'_{t-1}] \beta_h + u_{i,h,t}. \quad (12)$$

The vector c_{t-1} is included to accommodate common regressors. We set up the estimation problem as stacked least squares, which is, by a large margin, the most frequently adopted approach in the LP literature that uses panel data. We refer to the estimator of the impulse response function at horizon h as $\widehat{\theta}_{h,LS,I}$.

4.1 The approximate bias of the least-squares estimator with panel data

To fix notation, let

$$x_{i,t} = [1(i=1), 1(i=2), \dots, 1(i=I), \varepsilon_{i,t}, c'_{i,t-1}, c'_{t-1}]', \quad (13)$$

$$w_t = [\varepsilon_{1,t}, \varepsilon_{2,t}, \dots, \varepsilon_{I,t}, c'_{1,t}, c'_{2,t}, \dots, c'_{I,t}, c'_t]'. \quad (14)$$

Here, $1(i=n)$ denotes the indicator function, which takes the value 1 if the condition $i=n$ is met and 0 otherwise. We make the following assumption, which is analogous to Assumption 1.

Assumption 3. *The series w_t is strictly stationary and ergodic. The series has a purely nonde-terministic Wold representation with innovations $\omega_t = [\varepsilon_{1,t}, \varepsilon_{2,t}, \dots, \varepsilon_{I,t}, \nu_t']'$. Additionally, for all i ,*

- $\varepsilon_{i,t}$ is independent of $\varepsilon_{k,t+j}$ for all $j \neq 0$ and all k .
- $\varepsilon_{i,t}$ is independent of ν_{t+j} for all j .
- $\mathbb{E} \left[\sum_{i=1}^I x_{i,t} x'_{i,t} \right]$ is invertible, and $\mathbb{E} \left[\left\| \sum_{i=1}^I x_{i,t} x'_{i,t} \right\|^2 \right]$ is finite.

Our assumptions about $\varepsilon_{i,t}$ are meant to formalize what we mean by a structural shock.¹¹ Additionally, as was the case without panel data, the linearity of w_t facilitates the derivation of analytic expressions for the approximate bias. One consideration that is specific to panel data is the way that $\varepsilon_{j,t}$ affects $y_{i,t+h}$. We make the following related assumption, which is analogous to Assumption 2.

Assumption 4. *For every i , the regression error $u_{i,h,t}$ is independent of $\varepsilon_{k,t-j}$ and ν_{t-j-1} for all k and $j \geq 0$.*

This assumption formalizes what we mean by conditioning inference on information available at time $t - 1$ by including $c_{i,t-1}$ and c_{t-1} in the LP. Additionally, Assumption 4 requires that if $\varepsilon_{j,t}$ affects $y_{i,t+h}$, it only does so through correlation with $\varepsilon_{i,t}$. In Appendix A.3, we derive the following analytic result.

Analytical Result 3 (Bias in panel LPs with controls). *Under Assumptions 3-4, the approximate bias of $\hat{\theta}_{h,LS,I}$ is given by*

$$\mathbb{E} \left[\hat{\theta}_{h,LS,I} \right] - \theta_h = B_{h,LP,I} + O \left(T^{-3/2} \right), \quad (15)$$

where

$$B_{h,LP,I} = -\frac{1}{T-h} \sum_{j=1}^h [1 + \vartheta_j] \theta_{h-j} \text{ and } \vartheta_j = \frac{1}{I^2} \sum_{i=1}^I \sum_{k=1}^I \frac{\text{tr} \left\{ \Sigma_{c,0,I}^{-1} \Sigma_{c,j,k,i} \right\} \sigma_{\varepsilon,i,k}}{\sigma_{\varepsilon,I}^2}. \quad (16)$$

Here, $\sigma_{\varepsilon,I}^2 = \frac{1}{I} \sum_{i=1}^I \sigma_{\varepsilon,i,i}$, $\sigma_{\varepsilon,i,k} = \mathbb{E}[\varepsilon_{i,t} \varepsilon_{k,t}]$, $\Sigma_{c,0,I} = \frac{1}{I} \sum_{i=1}^I \Sigma_{c,0,i,i}$, $\Sigma_{c,j,i,k} = \mathbb{E}[\tilde{c}_{i,t-j} \tilde{c}'_{k,t}]$, and $\tilde{c}_{i,t} = [c'_{i,t}, c'_t]'$. Additionally, $B_{0,LP,I} = 0$.

¹¹Setting $\varepsilon_{i,t} = \varepsilon_t$ will accommodate a researcher who wants to investigate the response of $y_{i,t+h}$ to a structural shock that is common to all panelists. In our notation, ν_t potentially contains shocks that are panelist specific.

A few comments are in order regarding equation (16). First, as was the case without panel data—and because the LS estimator is consistent—the approximate bias goes to zero as the sample size goes to infinity. Second, the cross-autocovariance of the control variables plays a role in the approximate bias. Third, as was the case without panel data, if the IRF of interest is persistent and the controls are positively autocorrelated, the bias is larger. Fourth, even with controls that are independent across entities or over time, the bias does not go to zero as the number of panelists increases, holding T fixed—similar to the well-known bias discussed in Nickell (1981).

It is instructive to consider the implications of estimating the $\alpha_{i,h}$'s. In Appendix A.3, we show that the contribution to $B_{h,LP,I}$ from estimating the $\alpha_{i,h}$'s is given by

$$-\frac{1}{T-h} \sum_{j=1}^{T-h} \theta_{h-j}. \quad (17)$$

As a result, if the $\alpha_{i,h}$'s (means of the data) are known, increasing the number of panelists, I , can make $B_{h,LP,I}$ smaller, provided that the correlations with existing panelists (reflected in ϑ_j) are not too strong.

4.2 An AR(1) example with panel data

Here, we examine the approximate bias in panel LPs using independent AR(1) data generating processes for $y_{i,t}$. In particular, we assume that

$$y_{i,t} = \alpha_i + \theta_0 \varepsilon_{i,t} + \rho y_{i,t-1} + \nu_{i,t} \quad \text{with} \quad \varepsilon_{i,t} \stackrel{iid}{\sim} N(0, \sigma_\varepsilon^2) \text{ and } \nu_{i,t} \stackrel{iid}{\sim} N(0, \sigma_\nu^2). \quad (18)$$

Note that the independence of the innovations is in both the cross-section and time dimensions. Thus, $y_{i,t}$ is independent across i . Further, we assume that $\rho \in (0, 1)$. The researcher observes $\{(y_{i,t}, \varepsilon_{i,t})\}_{t=1}^T$ for $i = 1, 2, \dots, I$ and estimates the impulse response of $y_{i,t+h}$ to $\varepsilon_{i,t}$ for $h = 0, \dots, H$ using the LP defined in equation (12) with $c_{i,t-1} = y_{i,t-1}$ and $c_{t-1} = \emptyset$. This example satisfies Assumptions 3 and 4.

4.2.1 Explicit approximate bias with panel data

Noting that $\theta_h = \theta_0 \rho^h$, it is easy to see from Analytical Result 3 that the approximate bias of the panel LP is given by

$$B_{h,LP,I} = -\frac{1}{T-h} \sum_{j=1}^h \left(1 + \frac{1}{I} \rho^j\right) \theta_0 \rho^{h-j} = -\frac{\theta_0}{T-h} \left(\frac{1 - \rho^h}{1 - \rho} + \frac{1}{I} h \rho^h\right). \quad (19)$$

Notice that the expression for $B_{h,LP,I}$ in equation (19) is similar to the expression for $B_{h,LP}$ in equation (6), except that the second term in parentheses is divided by I . This result reflects the common AR coefficient across $y_{i,t}$ and our assumption that $y_{i,t}$ is independent of $y_{j,t}$ for $i \neq j$. Equation (19) also illustrates that adding panelists with uncorrelated data can reduce the approximate bias.¹² At the same time, $B_{h,LP,I}$ clearly does not approach zero as I increases. It is worth noting that when ρ is near unity, the first term in parentheses in equation (19) will be relatively large, and increasing I will have little effect on the approximate bias.

4.2.2 Quality of the approximation with panel data

Here we analyze how well $B_{h,LP,I}$ approximates the small-sample bias of $\hat{\theta}_{h,LS,I}$. Figure 4 shows $B_{h,LP,I}$ along with the sample mean of a Monte Carlo exercise conducted using the data generating process in equation (18).¹³ The figure shows results for different values of ρ and T and is constructed under the assumption that $I = 50$. This choice of I is reasonably representative in the context of the LP literature that uses panel data because many studies consider either state-level data or country-level data for advanced economies. Because all of the values of ρ shown imply that the data are persistent, the figure is very similar for smaller values of I .

Notably, Figure 4 is similar to Figure 2. As a result, we conclude that for empirically relevant sample sizes, $B_{h,LP,I}$ offers a reasonable approximation of the bias in LPs, though the quality of the approximation is somewhat worse for smaller values of T and larger values of ρ . Additionally, we conclude that for ρ near unity, adding a panel dimension to the LP does not materially reduce small-sample bias.

4.3 A bias-corrected estimator with panel data

Analytical Result 3 lends itself to constructing bias-corrected estimators of θ_h using plug-in estimators of θ_j and $\Sigma_{c,i,j}$. However, the number of cross-autocovariances that are required to construct the values of $\Sigma_{c,i,j}$ is large. Because the term of the approximate bias involving $\Sigma_{c,i,j}$ is multiplied

¹²With correlation across panelists, increasing I may have less of an effect in reducing the bias. As an extreme example, consider the case where $\varepsilon_{i,t} = \varepsilon_{j,t}$ and $\nu_{i,t} = \nu_{j,t}$ for all i and j . In this case, $y_{i,t} = y_{j,t}$ for all i and j and our expression for the approximate bias given by equation (16) collapses to the expression in equation (3). So, in this case, the expression for the approximate bias in the panel regression is given by equation (6).

¹³For the Monte Carlo exercises, we set $\sigma_\varepsilon = \sigma_\nu = 1$ and $\theta_0 = 1$.

by I^{-1} , in applications with reasonably large I , researchers could instead use

$$B_{h,LP,I} \approx -\frac{1}{T-h} \sum_{j=1}^h \theta_{h-j}. \quad (20)$$

While this expression is not an exact expression for the approximate bias to order $O(T^{-1})$, it could be used to improve upon the least-squares estimator by picking up some (or potentially most if the cross-sectional dependence is not too large) of the approximate bias. We denote the estimator of θ_h constructed in this way as $\hat{\theta}_{h,BC,I}$. Alternatively, similar to the case without panel data, the researcher could construct the approximate bias using bias-corrected estimates of θ_h . We denote this estimator as $\hat{\theta}_{h,BCC,I}$.

Here, we analyze how well $\hat{\theta}_{h,BCC,I}$ performs in the context of our AR(1) example data generating process from equation (18). The left column of Figure 5 shows the Monte Carlo average of $\hat{\theta}_{h,BCC,I} - \theta_h$, for different values of ρ and T , as a function of h . Clearly, the bias-corrected estimator performs better for smaller values of ρ and larger values of T .

The right column of Figure 5 shows the MSEs of $\hat{\theta}_{h,BCC,I}$ relative to the MSEs associated with $\hat{\theta}_{h,LS,I}$.¹⁴ Note that a number larger than unity indicates that the MSE of $\hat{\theta}_{h,BCC,I}$ is larger than the MSE of $\hat{\theta}_{h,LS,I}$. Notably, $\hat{\theta}_{h,BCC,I}$ leads to a smaller MSE than $\hat{\theta}_{h,LS,I}$. Intuitively, the reason is that the values $\hat{\theta}_{h,LS,I}$ have a normal asymptotic limiting distribution when multiplied by \sqrt{TI} , so in small samples this estimator has relatively little variance when used to construct the approximate bias, which is of order $O(T^{-1})$. Of course, estimating the values of $\Sigma_{c,i,j}$ to capture the entire approximate bias would add additional variance.

5 Bias in the Context of a Medium-Scale Monetary Time Series

In this section, we conduct a Monte Carlo study using a larger time series model. To construct an empirically realistic data generating process, we follow Christiano et al. (2005)—subsequently referred to as CEE—and estimate a nine variable VAR(4). We focus on the dynamic effects of a monetary policy shock, identified recursively like in CEE. Figure 6 shows the impulse responses of real GDP, the price level, and the federal funds rate to a monetary policy shock. The estimated VAR(4) serves as the data generating process for our Monte Carlo exercise.¹⁵

Our Monte Carlo study examines the properties of both the standard least-squares estimator and our bias-corrected LP for estimating these three impulse responses. We assume that the

¹⁴Note that the MSE rankings in Figure 5 might differ if we introduced clustered dependence in the simulation

econometrician observes sets of time series of length T of the dependent variable y_t —either the log level of real GDP, the log level of prices, or the federal funds rate—sets of controls c_t , and the monetary policy shock ε_t . In this section, we assume that the vector of controls is “full” in the sense that it contains the four lags of *all* variables like in the VAR. Results under different assumptions about the conditioning set are broadly similar and are available in Appendix E. The sample sizes vary, as in the earlier simulations, with $T \in \{100, 200\}$. We simulate 2000 trajectories from the VAR for each sample size T . Each simulation is initialized from a random point in the stationary distribution.

Figure 7 displays Monte Carlo estimates of the bias of $\hat{\theta}_{h,LS}$ and $\hat{\theta}_{h,BCC}$ (the dashed and solid lines in the left-hand column, respectively) and the MSE ratio (the right-hand column) for output, the price level, and the federal funds. Focusing first on the bias, we see that all of the estimated impulse responses exhibit bias, with the largest bias associated with the impulse response of real GDP. As expected, the bias is decreasing in the sample size T , with the red lines ($T = 200$) generally being closer to zero than the green lines ($T = 100$). Interestingly, unlike for the univariate AR(1) example, this bias is not always—or even typically—downwards. In nearly all cases, the bias of $\hat{\theta}_{h,BCC}$ is smaller in magnitude than the bias of $\hat{\theta}_{h,LS}$, indicating that the bias correction works well. But as with the univariate AR(1) example, the bias corrected estimator does not completely eliminate the bias.

Turning to the relative MSE, which is displayed in the right-hand column of Figure 7, we see that neither $\hat{\theta}_{h,LS}$ nor $\hat{\theta}_{h,BCC}$ is uniformly best. For example, for real GDP the MSE of $\hat{\theta}_{h,BCC}$ is less than that of $\hat{\theta}_{h,LS}$ for $h \in [4, 12]$, as the bias term dominates the MSE calculation at those horizons. When the bias associated with $\hat{\theta}_{h,LS}$ is relatively small (which here is at small or large h), this estimator is preferred in an MSE sense. For the price level the bias switches signs around $h \approx 14$, so for these horizons $\hat{\theta}_{h,BCC}$ has a relatively high MSE. Finally, for the federal funds rate neither estimator dominates across all h . Overall, the MSE criterion does not definitively select either $\hat{\theta}_{h,LS}$ or $\hat{\theta}_{h,BCC}$ as best.¹⁶

innovations.

¹⁵The inflation responses are computed *ex post* as the (annualized) percent change in the price level. Note also that the VAR, at the maximum likelihood estimator, is stationary. The magnitudes of the five largest eigenvalues are $[0.99, 0.97, 0.97, 0.95, 0.95]$. More details about the VAR are available in Appendix E.

¹⁶In the Appendix, we compare the LP estimators to those of VAR(1) and VAR(4) models.

6 Some considerations related to standard errors

LPs are often used to construct confidence intervals for IRFs, which necessitates calculating standard errors that are a function of the asymptotic covariance matrix of the least-squares estimator of the parameters in equation (1), which is given by

$$V_{T-h} = \left(\frac{1}{T-h} \sum_{t=1}^{T-h} \mathbb{E} [\tilde{x}_t \tilde{x}_t'] \right)^{-1} S_{T-h} \left(\frac{1}{T-h} \sum_{t=1}^{T-h} \mathbb{E} [\tilde{x}_t \tilde{x}_t'] \right)^{-1} \quad (21)$$

where $\tilde{x} = [1, x_t']'$ and

$$S_{T-h} = \frac{1}{T-h} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \mathbb{E} [\tilde{x}_t \tilde{x}_s' u_{h,t} u_{h,s}]. \quad (22)$$

The challenging piece of the standard-error calculation is estimating S_{T-h} . Since Jordà (2005), the conventional wisdom has been that HAR standard errors are necessary because the regression residuals of LPs are autocorrelated. A popular choice of estimator of S_{T-h} in the LP literature is the estimator of Newey and West (1987). Because the NW estimator tends to understate sampling uncertainty, work by Müller (2014), Sun (2014), and Lazarus et al. (2018) has suggested alternatives. However, under Assumptions 1 and 2, the *LP regression score*—the product of the ε_t and the LP regression residuals—is serially uncorrelated.¹⁷ Thus, in large samples the (heteroskedasticity-robust) estimator of Huber (1967) and White (1980) (HW) is valid. Here, we consider the small sample implications of using the NW estimator when HW will do.

Because we assume that ε_t is independent of c_{t-1} , the only relevant element of S_{T-h} in computing the standard error of $\hat{\theta}_{h,LS}$ is the diagonal element in the same position as θ_h , which is given by

$$\frac{1}{T-h} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \mathbb{E} [\varepsilon_t (y_{t+h} - \alpha_h - x_t' \beta_h) \varepsilon_s (y_{s+h} - \alpha_h - x_s' \beta_h)] = \frac{1}{T-h} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \gamma_{h,|t-s|}. \quad (23)$$

The NW estimator of this value with bandwidth m is given by

$$\hat{\gamma}_{h,0} + \sum_{\ell=1}^m \left(1 - \frac{\ell}{m+1} \right) \frac{T-h-\ell}{T-h} 2\hat{\gamma}_{h,\ell}. \quad (24)$$

where

$$\hat{\gamma}_{h,\ell} = \frac{1}{T-h-\ell} \sum_{t=\ell+1}^{T-h} \varepsilon_t (y_{t+h} - \alpha_h - x_t' \hat{\beta}_{h,LS}) \varepsilon_{t-\ell} (y_{t-\ell+h} - \alpha_h - x_{t-\ell}' \hat{\beta}_{h,LS}). \quad (25)$$

¹⁷Montiel Olea and Plagborg-Møller (2021) use lag augmentation to achieve (population) residualized regressors, whereas our setup does not require this step because we assume the researcher has access to ε_t .

Note that $\widehat{\gamma}_{h,\ell}$ is the sample autocovariance of the regression score. Under Assumptions 1 and 2, $\gamma_{h,\ell} = 0$ for $\ell > 0$. As a result, setting $m > 0$ is not asymptotically necessary to conduct valid inference in our LP set-up. However, using the NW estimator may have implications for inference in small samples because $\widehat{\gamma}_{h,\ell}$ may have small-sample bias akin to the bias we have documented in LP estimators of impulse response functions.

6.1 Explicit approximate bias of the autocovariance of the regression score

To investigate the small-sample implications of using the NW estimator when HW will do, we focus attention on the case where $\theta_h = 0$ for all h . We choose this setup for two reasons. First, the null hypothesis that $\theta_h = 0$ is often of interest to researchers using LPs, and our derivation would be correct under that hypothesis. Second, the assumption that $\theta_h = 0$ facilitates the derivation of expressions of the approximate bias in $\widehat{\gamma}_{h,\ell}$. In Appendix D we show that, under suitable regularity conditions, when α_h , β_h , and $\gamma_{h,\ell}$ are jointly estimated

$$\mathbb{E}\widehat{\gamma}_{h,\ell} - \gamma_{h,\ell} = -\frac{1}{T-h-\ell}\gamma_{h,0} + O\left(T^{-3/2}\right). \quad (26)$$

It is immediately clear that increasing m in the NW estimator is likely to decrease the value of the element of \widehat{S}_{T-h} that is relevant to constructing the standard errors of the least-squares estimator in our LP setting, which reduces the size of the standard error.¹⁸

6.2 Bias in standard errors in the context of an AR(1) example

To explore the small sample bias of standard errors of LPs, we again use the AR(1) data generating process in equation (5) for Monte Carlo simulations. We consider two specifications. In the first specification, we assume that $\theta_0 = 0$, $\sigma_\varepsilon^2 = 1$, and $\sigma_\nu^2 = 2$. This specification conforms to our assumption in the previous subsection that $\theta_h = 0$ for all h . In the second specification, we assume that $\theta_0 = 1$, $\sigma_\varepsilon^2 = 1$, and $\sigma_\nu^2 = 1$. This specification has $\theta_h \neq 0$ for all h , but maintains the same variance of y_t as in the first specification. We set $c_{t-1} = y_{t-1}$.¹⁹

¹⁸In Appendix D, we show that, for $\ell > 0$, the approximate bias in $\widehat{\gamma}_{h,\ell}$ is the same as in the case where α_h is known. Additionally, we derive the approximate bias in the case when controls are not included in the LP. The expression of the approximate bias in terms of $\gamma_{0,h}$ is unchanged.

¹⁹We focus on HW and NW estimators and standard critical values not only because of their popularity in practice but also because they can be used to isolate the effect of bias in the estimated regression-score autocovariances. In fact, the errors here are homoskedastic, so heteroskedasticity-robust estimators are unnecessary.

We fix the horizon of the LP at $h = 10$ and estimate $\gamma_{10,\ell}$ using the fitted regression scores from the LPs.²⁰ Figure 8 shows the Monte Carlo mean of $\hat{\gamma}_{10,\ell}$ for $\ell > 0$ in each of the two specifications. Several remarks are in order. First, when $\theta_h = 0$, the Monte Carlo mean of $\hat{\gamma}_{10,\ell}$ is roughly constant for all ℓ shown, as implied by the analytical expression of the approximate bias derived in the previous subsection. Second, when $\theta_h > 0$, the Monte Carlo mean of $\hat{\gamma}_{10,\ell}$ is also negative and for some h is more negative than when $\theta_h = 0$. Third, regardless of whether $\theta_h = 0$ or $\theta_h > 0$, the Monte Carlo mean of $\hat{\gamma}_{10,\ell}$ is more negative for smaller values of T or larger values of ρ .

To analyze the effect of using estimates of the autocovariance of the regression score to construct standard errors, we consider coverage probabilities from symmetric 95% confidence intervals constructed using the method of HW and NW. We also consider the equally-weighted cosine (EWC) estimator discussed in Lazarus et al. (2018), a frequency-domain-based alternative standard error estimator.²¹ Figure 9 displays the coverage probabilities for both $\hat{\theta}_{LS}$ (the left-hand column) and $\hat{\theta}_{BCC}$ (the right-hand column) of confidence intervals constructed using the three standard error estimates, with the rows corresponding to different values of ρ . For all ρ and T , the frequency-domain-based EWC delivers the best coverage, except perhaps at $h = 0$. The NW-based confidence interval uniformly provides the worst coverage, but only by a small amount. Finally, the confidence intervals using $\hat{\theta}_{BCC}$ provide better coverage than those using $\hat{\theta}_{LS}$ regardless of the standard errors used. Overall, our results show the favorable performance of $\hat{\theta}_{BCC}$ and the poor performance of NW. While the EWC-based intervals are clearly superior, in practice some users may continue to prefer the conventional time-domain-based estimators, in which case they should use HW.

6.3 Medium-scale model revisited

To extend our analysis of standard errors to a medium-scale time series model, Figure 10 displays the same coverage probabilities shown for the AR(1) but for the CEE VAR variables for real GDP (top), the price level (middle) and the federal funds rate (bottom) for $T = 100$. Once again, the

²⁰In Appendix D, we derive the bias under the assumption that α_h , β_h , and $\gamma_{h,\ell}$ are estimated as a part of the same generalized method of moments system. As a result, those derivations work with only $T - h - \ell$ observations. For our Monte Carlo exercises in this sub-section, we instead calculate $\hat{\gamma}_{h,\ell}$ using the $T - h$ values of the fitted regression score. So, α_h and β_h are estimated using $T - h$ observations for every ℓ .

²¹For EWC we use a bandwidth $\approx 0.41T^{2/3}$. One can also use the critical values for the NW estimator under fixed- b asymptotics with appropriately choosing the bandwidth ($\approx 1.3\sqrt{T-h}$). We found this performed poorly, and so therefore we omit it from the presentation.

EWC-based confidence intervals perform the best, with the NW-based ones performing the worst. The HW-based intervals always perform better than the NW ones, suggesting again that among time-domain based estimators of standard errors, they have the best finite sample performance. The coverage for the price level is particularly poor for large h . Finally, $\hat{\theta}_{h,BCC}$ -based intervals are not uniformly better (or worse) than $\hat{\theta}_{h,LS}$ -based ones.

7 Conclusion

We have shown that LPs can be severely biased in sample sizes commonly found in the related literature. Our analytical expressions of the approximate bias have shown that LPs are intimately linked across horizons in small samples. We have also shown how researchers could use our expression of the approximate bias to correct for bias in LPs. When correcting for bias, researchers face a small-sample tradeoff between bias and MSE. Using Monte Carlo analysis, we have demonstrated that the performance of our bias-corrected estimator depends on the underlying data generating process and on the LP horizon of interest.

In our Monte Carlo examples, our bias correction does not completely correct for the bias in LPs. These results suggest that other time series models with well-understood and effective methods for bias correction (such as VARs) might be better alternatives to estimate impulse responses if researchers have data samples in the time dimension that are similar to those typically found in empirical macroeconomic research. In particular, specifying time series models that are generative for the time series of interest would allow researchers to use likelihood methods. We leave exploration of these ideas, as well as the bias of LP estimators with heteroskedastic shocks, to future research.

We have also analyzed bias in standard errors computed for estimated impulse response functions from LPs. In small samples, standard errors that rely on estimated autocovariances of the regression score, like the NW estimator, typically understate the amount of uncertainty surrounding the estimated impulse response functions. Recent work on standard errors in time series regression has focused on limiting distributions other than the normal distribution (see Sun (2014) and Lazarus et al. (2018)). However, with samples typically found in the LP literature, limiting distributions are not accurate approximations. As a result, if researchers are going to use NW standard errors, they may want to check to see if HW or fixed- b standard errors would lead to different conclusions. If the HW standard errors are larger than the NW standard errors, researchers should

consider what might lead to the apparent negative autocovariance in the regression score. Without another reasonable theory, a plausible explanation might be that the negative estimates of the autocovariance of the regression score are the result of small-sample bias.

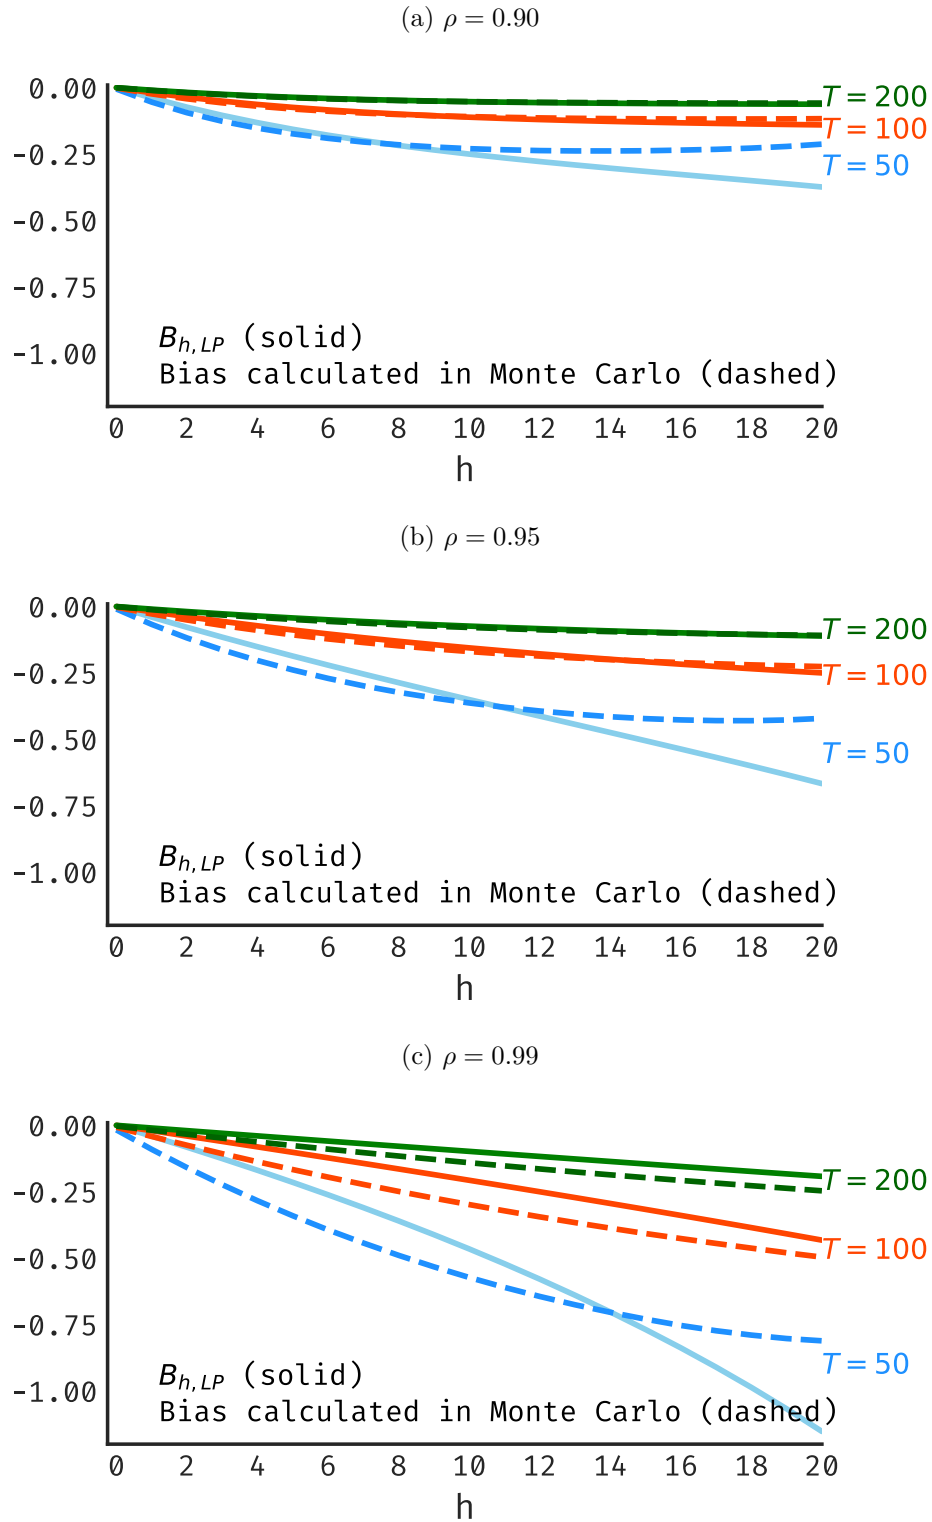
References

- ANATOLYEV, S. (2005): “GMM, GEL, Serial Correlation, and Asymptotic Bias,” *Econometrica*, 73, 983–1002.
- BAO, Y. AND A. ULLAH (2007): “The Second-Order Bias and Mean Squared Error of Estimators in Time-Series Models,” *Journal of Econometrics*, 140, 650 – 669.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal rigidities and the dynamic effects of a shock to monetary policy,” *Journal of Political Economy*, 113, 1–45.
- HAHN, J. AND G. KUERSTEINER (2002): “Asymptotically Unbiased Inference for a Dynamic Panel Model with Fixed Effects When Both n and T are Large,” *Econometrica*, 70, 1639–1657.
- HERBST, E. AND B. K. JOHANSEN (2021): “Bias in local projections,” *FEDS Working Paper No. 2020-010R1*.
- HUBER, P. J. (1967): “The behavior of maximum likelihood estimates under nonstandard conditions,” in *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability: Weather Modification; University of California Press: Berkeley, CA, USA*.
- JORDÀ, Ò. (2005): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- KENDALL, M. G. (1954): “Note on Bias in the Estimation of Autocorrelation,” *Biometrika*, 41, 403–404.
- KILIAN, L. (1998): “Small-Sample Confidence Intervals for Impulse Response Functions,” *Review of economics and statistics*, 80, 218–230.
- KILIAN, L. AND Y. J. KIM (2011): “How Reliable are Local Projection Estimators of Impulse Responses?” *Review of Economics and Statistics*, 93, 1460–1466.
- LAZARUS, E., D. J. LEWIS, J. H. STOCK, AND M. W. WATSON (2018): “HAR Inference: Recommendations for Practice,” *Journal of Business & Economic Statistics*, 36, 541–559.

- LI, DAKE, P.-M. M. AND C. K. WOLF (2022): “Local Projections vs. VARs: Lessons From Thousands of DGPs,” *Mimeo*.
- MONTIEL OLEA, J. L. AND M. PLAGBORG-MØLLER (2021): “Local Projection Inference Is Simpler and More Robust Than You Think,” *Econometrica*, 89, 1789–1823.
- MÜLLER, U. K. (2014): “HAC Corrections for Strongly Autocorrelated Time Series,” *Journal of Business & Economic Statistics*, 32, 311–322.
- NAGAR, A. L. (1959): “The bias and moment matrix of the general k-class estimators of the parameters in simultaneous equations,” *Econometrica: Journal of the Econometric Society*, 575–595.
- NAKAMURA, E. AND J. STEINSSON (2018): “Identification in Macroeconomics,” *Journal of Economic Perspectives*, 32, 59–86.
- NEWKEY, W. K. AND R. J. SMITH (2004): “Higher Order Properties of Gmm and Generalized Empirical Likelihood Estimators,” *Econometrica*, 72, 219–255.
- NEWKEY, W. K. AND K. D. WEST (1987): “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation,” *Econometrica*, 55, 703–708.
- NICHOLLS, D. AND A. L. POPE (1988): “Bias in the estimation of multivariate autoregressions,” *Australian Journal of Statistics*, 30, 296–309.
- NICKELL, S. (1981): “Biases in Dynamic Models with Fixed Effects,” *Econometrica*, 49, 1417–1426.
- PLAGBORG-MØLLER, M. AND C. K. WOLF (2021): “Local Projections and VARs Estimate the Same Impulse Responses,” *Econometrica*, 89, 955–980.
- POPE, A. L. (1990): “Biases of estimators in multivariate non-Gaussian autoregressions,” *Journal of Time Series Analysis*, 11, 249–258.
- RILSTONE, P., V. SRIVASTAVA, AND A. ULLAH (1996): “The Second-Order Bias and Mean Squared Error of Nonlinear Estimators,” *Journal of Econometrics*, 75, 369 – 395.
- ROMER, C. D. AND D. H. ROMER (2004): “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, 94, 1055–1084.
- SAWA, T. (1972): “Finite-sample properties of the k-class estimators,” *Econometrica*, 653–680.

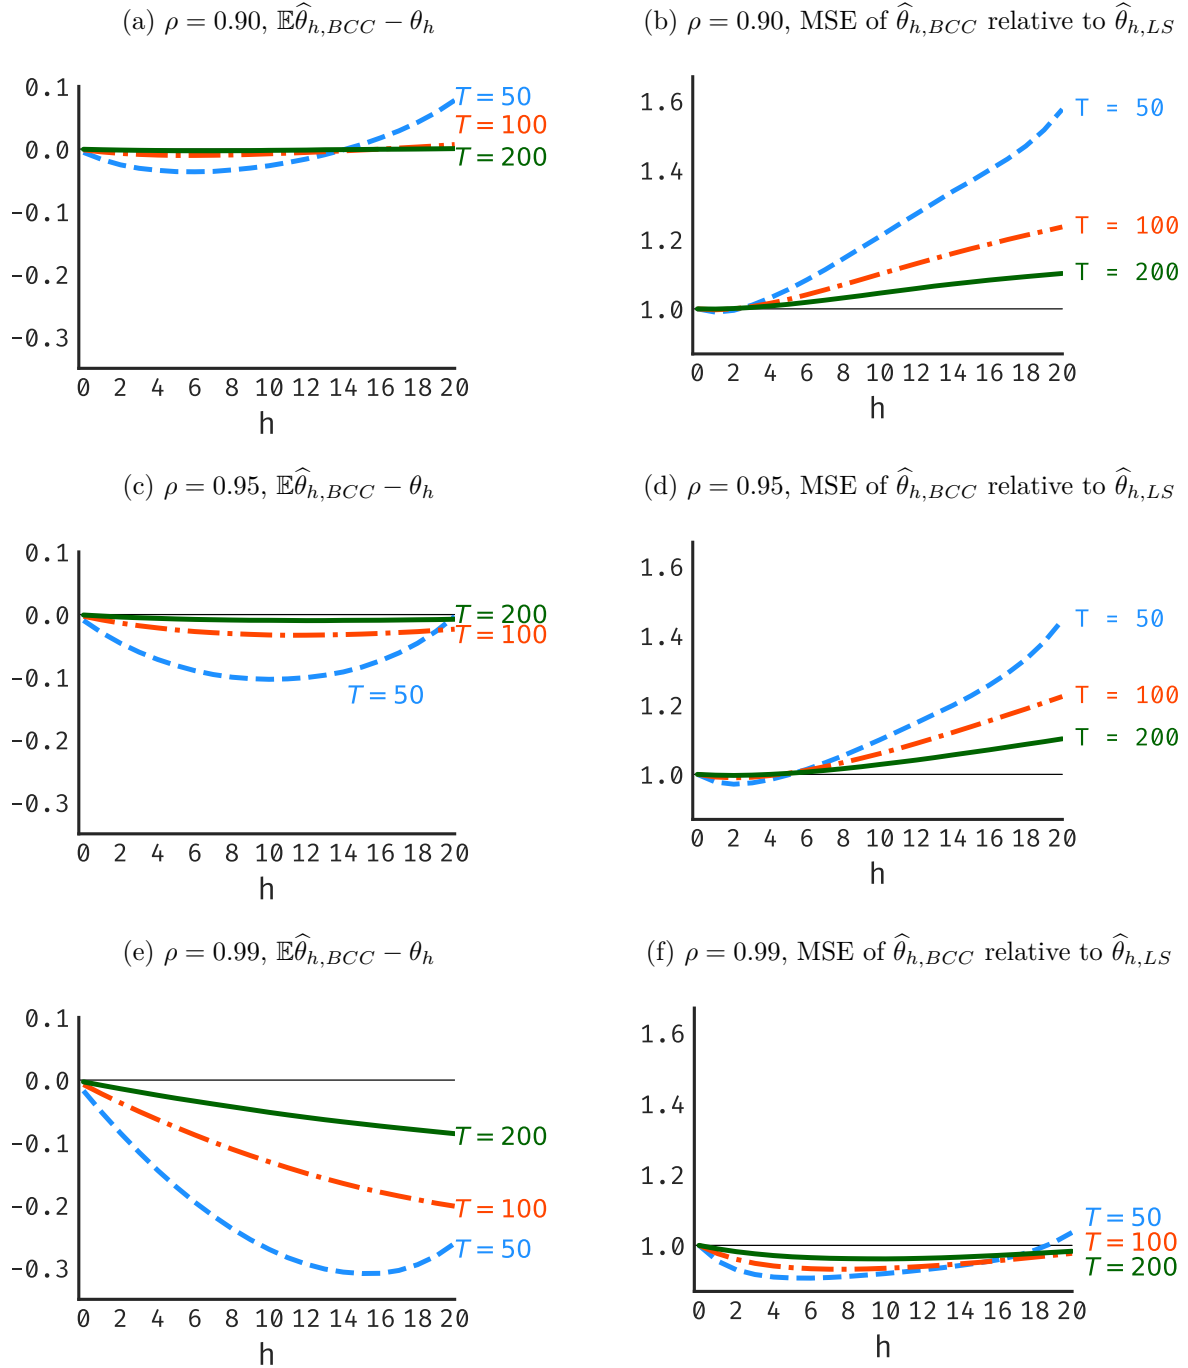
- SHAMAN, P. AND R. A. STINE (1988): “The Bias of Autoregressive Coefficient Estimators,” *Journal of the American Statistical Association*, 83, 842–848.
- SIMS, C. A. AND T. ZHA (1999): “Error Bands for Impulse Responses,” *Econometrica*, 67, 1113–1155.
- SRINIVASAN, T. N. (1970): “Approximations to Finite Sample Moments of Estimators Whose Exact Sampling Distributions are Unknown,” *Econometrica*, 38, 533–541.
- SUN, Y. (2014): “Let’s Fix It: Fixed-b Asymptotics Versus Small-b Asymptotics in Heteroskedasticity and Autocorrelation Robust Inference,” *Journal of Econometrics*, 178, 659–677.
- THEIL, H. (1961): *Economic forecasts and policy*, North-Holland Pub. Co.
- WHITE, H. (1980): “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity,” *Econometrica*, 48, 817–838.

Figure 2: $B_{h,LP}$ performs well in empirically-relevant samples when y_t is an AR(1).




Note: the sub-figures show the value of $B_{h,LP}$ and the Monte Carlo means of $\widehat{\theta}_{h,LS} - \theta_h$ estimated on data simulated from equation (5). We use 1,000,000 Monte Carlo simulations. We set $\sigma_\varepsilon = \sigma_\nu = 1$ and $\theta_0 = 1$.

Figure 3: Performance of $B_{h,BCC}$ in an AR(1) example.



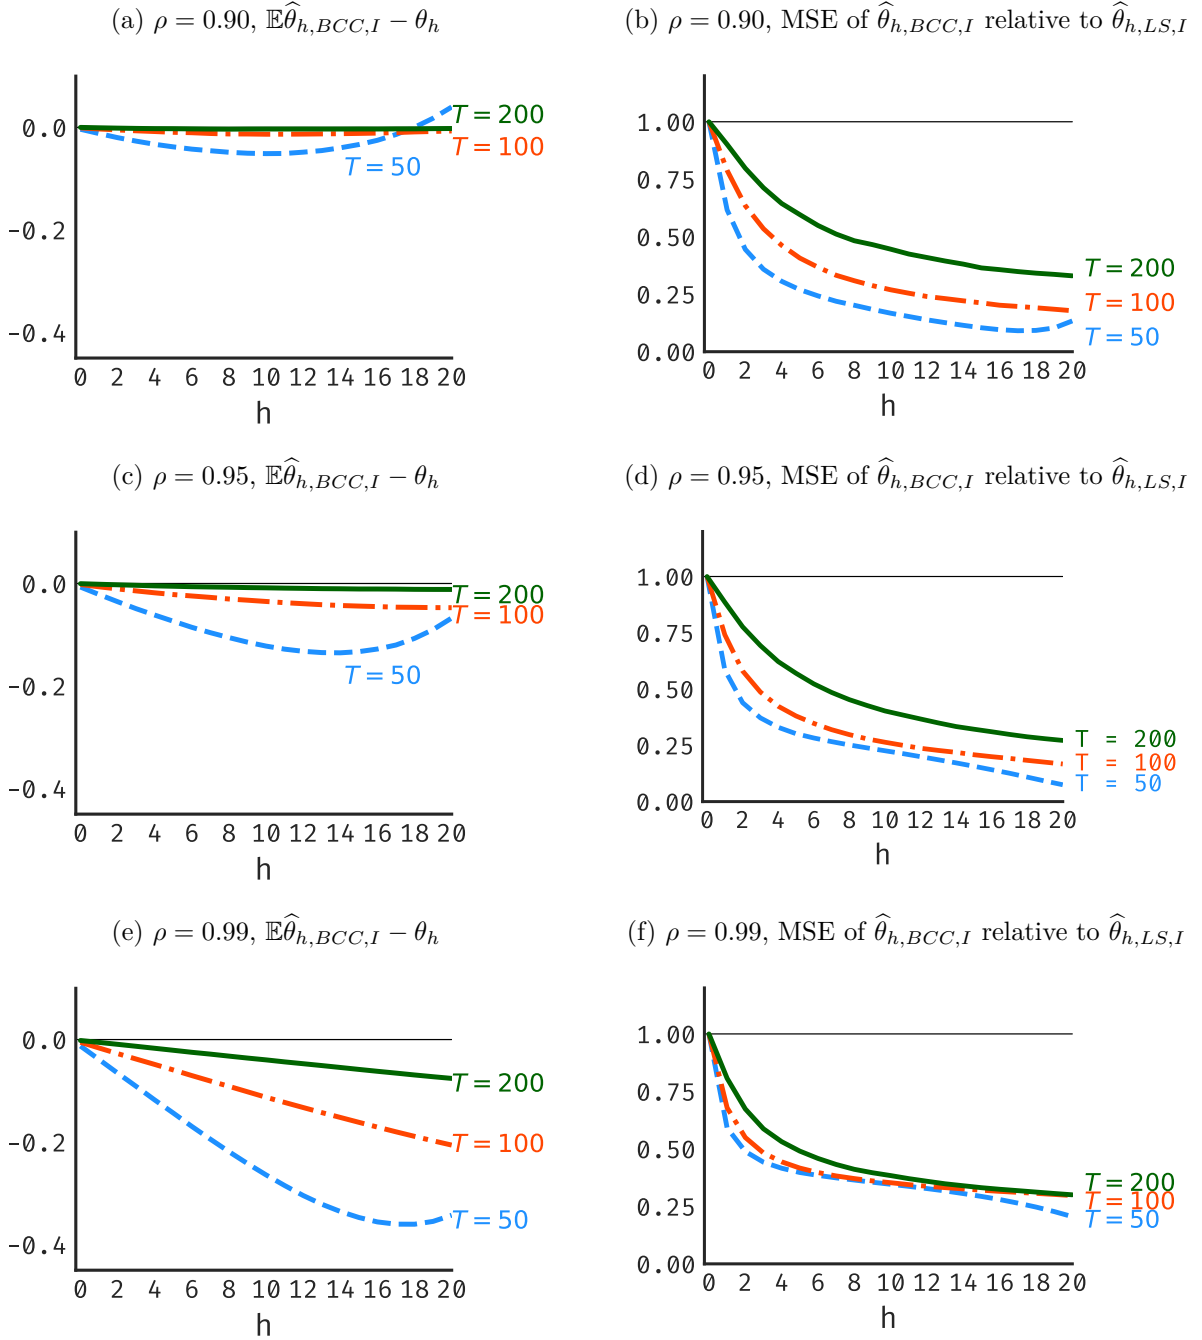
Note: the sub-figures on the left show the Monte Carlo means of $\hat{\theta}_{h,BCC} - \theta_h$ estimated on data simulated from equation (5). The sub-figures on the right show the Monte Carlo value of the MSE of $\hat{\theta}_{h,BCC}$ relative to the MSE of $\hat{\theta}_{h,LS}$. A number larger than unity indicates that the MSE of $\hat{\theta}_{h,BCC}$ is larger than $\hat{\theta}_{h,LS}$. We use 1,000,000 Monte Carlo simulations. We set $\sigma_\varepsilon = \sigma_\nu = 1$ and $\theta_0 = 1$.

Figure 4: $B_{h,LP,I}$ performs well in empirically-relevant samples when $y_{i,t}$ is an AR(1) and $I = 50$.




Note: the sub-figures show the value of $B_{h,LP,I}$ and the Monte Carlo means of $\widehat{\theta}_{h,LS,I} - \theta_h$ estimated on data simulated from equation (18). We use 1,000,000 Monte Carlo simulations. We set $\sigma_\varepsilon = \sigma_\nu = 1$ and $\theta_0 = 1$.

Figure 5: Performance of $B_{h,BCC,I}$ in an AR(1) example.



Note: the sub-figures on the left show the Monte Carlo means of $\hat{\theta}_{h,BCC,I} - \theta_h$ estimated on data simulated from equation (18). The sub-figures on the right show the Monte Carlo value of the MSE of $\hat{\theta}_{h,BCC,I}$ relative to the MSE of $\hat{\theta}_{h,LS,I}$. A number larger than unity indicates that the MSE of $\hat{\theta}_{h,BCC,I}$ is larger than $\hat{\theta}_{h,LS,I}$. We use 1,000,000 Monte Carlo simulations. We set $\sigma_\varepsilon = \sigma_\nu = 1$ and $\theta_0 = 1$.

Figure 6: Impulse Response to Monetary Policy Shock in a CEE-style VAR




Note: The figure displays impulse response of output (in percent), the price level (in percent), and the federal funds rate (in percentage points) to a one standard deviation increase in the monetary policy shock (identified recursively via the Cholesky factorization). The solid lines display the median impulse responses and the dashed lines 90 percent confidence intervals computed using the method of Sims and Zha (1999).


Figure 7: Bias and MSE under a CEE-type VAR Data Generating Process

Real GDP

(a) $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid)




(b) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$




Price Level

(c) $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid)




(d) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$




Federal Funds Rate

(e) $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid)




(f) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$



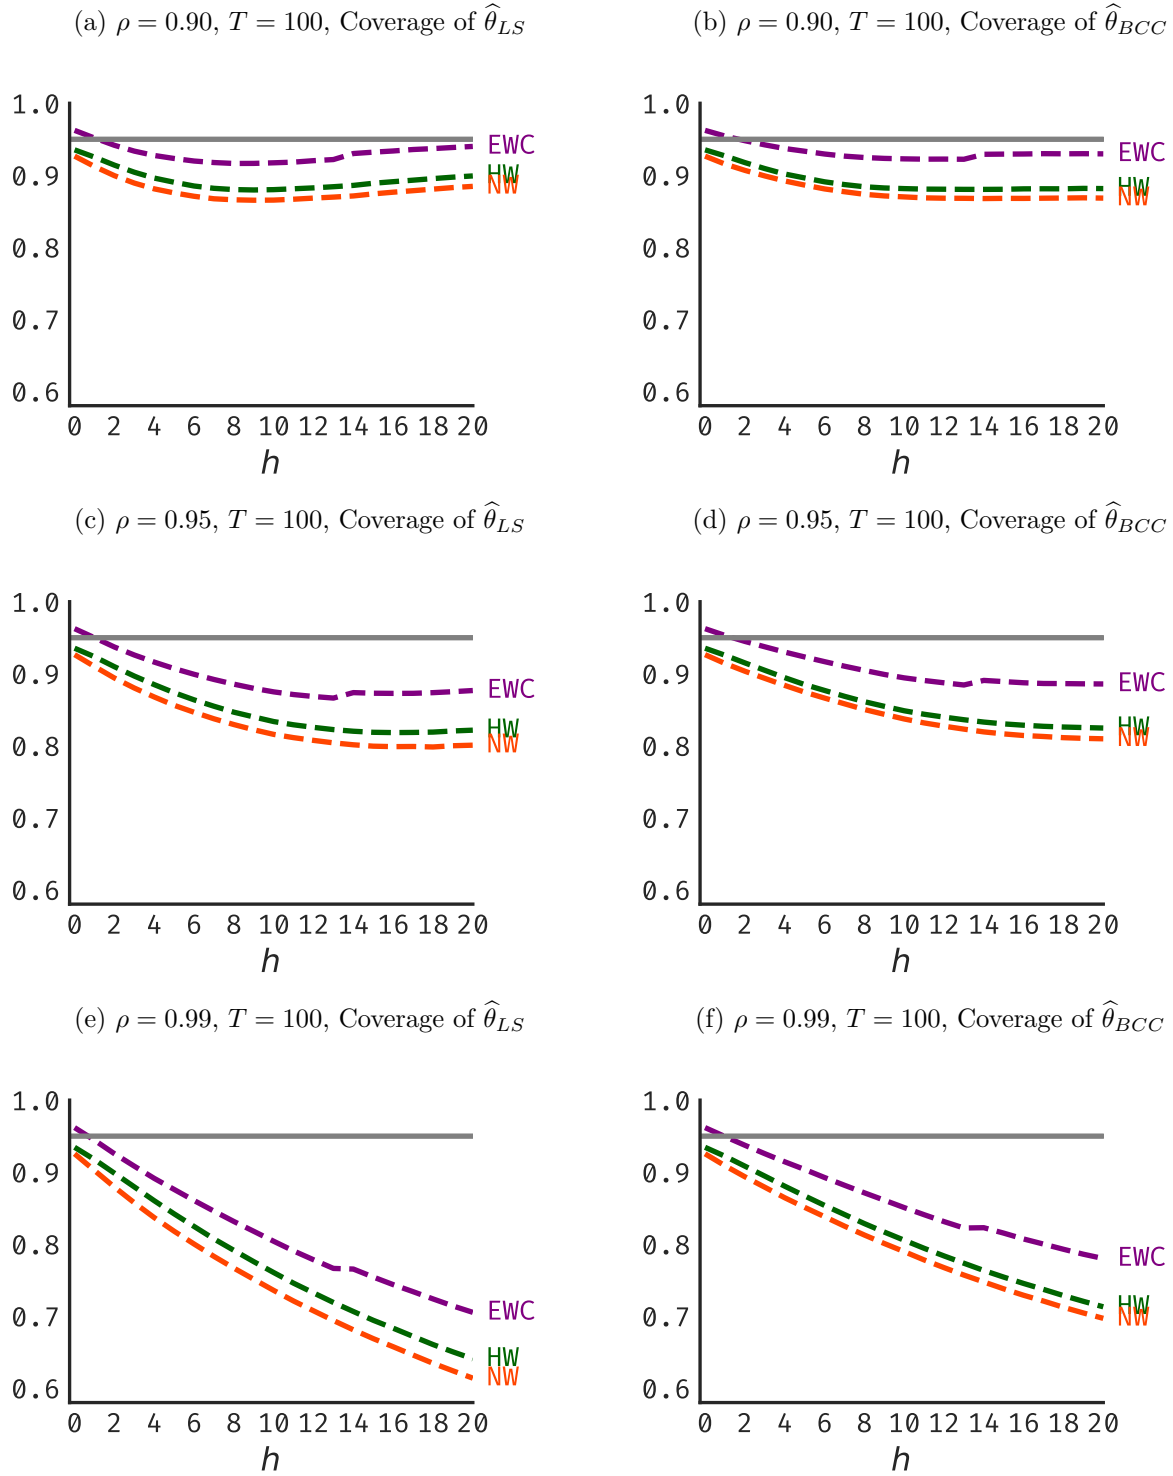
Note: the sub-figures in the left column show the Monte Carlo means of $\hat{\theta}_{h,LS} - \theta_h$ (dashed lines) and $\hat{\theta}_{h,BCC} - \theta_h$ (solid lines) for $T = 100$ (red) and $T = 200$ (green). The sub-figures in the right column show the ratio of the MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$ for $T = 100$ (red) and $T = 200$ (green).

Figure 8: $\widehat{\gamma}_{10,\ell}$ is biased down in small samples when y_t is an AR(1).



Note: the sub-figures on the left show the Monte Carlo means of $\widehat{\gamma}_{10,\ell}$ estimated on data simulated from equation (5) when $\theta_0 = 0$. The sub-figures on the right show analogous figures when $\theta_0 = 1$. When $\theta_0 = 1$, we set $\sigma_\varepsilon = \sigma_\nu = 1$. When $\theta_0 = 0$, we increase σ_ν so that the variance of y_t is unchanged. We use 1,000,000 Monte Carlo simulations.

Figure 9: Empirical Coverage of 95% Confidence Intervals




Note: the sub-figures on the left show the empirical coverage of symmetric 95% confidence intervals constructed using HW (green), NW (orange), and EWC (purple) based standard errors for $\hat{\theta}_{LS}$ and $\hat{\theta}_{BCC}$.


Figure 10: Coverage of 95% Confidence Intervals under a CEE-type VAR DGP

Real GDP

(a) Coverage of $\hat{\theta}_{h,LS}$




(b) Coverage of $\hat{\theta}_{h,BCC}$




Price Level

(c) Coverage of $\hat{\theta}_{h,LS}$




(d) Coverage of $\hat{\theta}_{h,BCC}$




Federal Funds Rate

(e) Coverage of $\hat{\theta}_{h,LS}$



(f) Coverage of $\hat{\theta}_{h,BCC}$



Note: the sub-figures on the left show the empirical coverage of symmetric 95% confidence intervals constructed using HW (green), NW (orange), and EWC (purple) based standard errors for $\hat{\theta}_{LS}$ and $\hat{\theta}_{BCC}$ for $T = 100$.

Appendix for “Bias in Local Projections”

A Derivations of approximate bias

In this Appendix, we derive our expressions for the approximate bias of the LP estimators studied in our paper. To do so, we employ the framework proposed by Rilstone et al. (1996) and extended to time series models by Bao and Ullah (2007). These papers derive expressions for finite-sample moments for a wide class of estimators via an approximation of an estimator $\hat{\beta}$ of the form:

$$\hat{\beta} - \beta = a_{-1/2} + a_{-1} + O_p(T^{-3/2}). \quad (\text{A.1})$$

Assumption 1, combined with the least-squares estimation framework, fulfills the necessary assumptions of Rilstone et al. (1996). Meanwhile, Assumption 2 facilitates the derivation of tractable expressions for the approximate bias.

In this Appendix, we use the notation of Bao and Ullah (2007) where possible. For each derivation, we will cast the least-squares estimator as a generalized method of moments (GMM) estimator with moment conditions given by $q_h(\beta; \xi_t)$, where $\beta = [\alpha_h, \beta'_h]'$ and $\xi_t = [y_{t+h}, x'_t]'$ is the data vector for the LP model. The GMM empirical moments are given by

$$\psi_{h,T-h}(\beta; \{\xi_t\}_{t=1}^{T-h}) = \frac{1}{T-h} \sum_{t=1}^{T-h} q_h(\beta; \xi_t). \quad (\text{A.2})$$

Let $\nabla^i A(\beta)$ be the matrix of i th order partial derivatives of A with respect to β . In what follows, we abbreviate $\psi_{h,T-h}(\beta; \{\xi_t\}_{t=1}^{T-h})$ as $\psi_{h,T-h}$ and $q_h(\beta; \xi_t)$ as $q_{h,t}$. We also introduce the series of matrices

$$H_i = \nabla^i \psi_{h,T-h} \text{ and } \bar{H}_i = \mathbb{E}[H_i] \text{ with } Q = \bar{H}_1^{-1}, V = H_1 - \bar{H}_1. \quad (\text{A.3})$$

Bao and Ullah (2007) show that the expressions for the terms in (A.1) are given by:

$$a_{-1/2} = -Q\psi_{h,T-h} \quad \text{and} \quad a_{-1} = -QVa_{-1/2} - \frac{1}{2}Q\bar{H}_2 [a_{-1/2} \otimes a_{-1/2}]. \quad (\text{A.4})$$

We are interested in computing

$$B = \mathbb{E}[a_{-1/2} + a_{-1}] = \mathbb{E}[QVQ\psi_{h,T-h}] - \mathbb{E}\left[\frac{1}{2}Q\bar{H}_2 [(Q\psi_{h,T-h}) \otimes (Q\psi_{h,T-h})]\right], \quad (\text{A.5})$$

which we refer to as the approximate bias. Throughout this Appendix, and without loss of generality, we assume all data have mean zero. Before proceeding, we introduce notation for second moments of the data

$$\sigma_\varepsilon^2 = \mathbb{E}[\varepsilon_t^2], \quad \sigma_y^2 = \mathbb{E}[y_t^2], \quad \Sigma_{c,j} = \mathbb{E}[c_{t-j}c_t']. \quad (\text{A.6})$$

A.1 LP when means are estimated

In the LP model when means are estimated, the moment conditions for the least-squares estimator are

$$\mathbb{E} \begin{bmatrix} y_{t+h} - \alpha_h - x_t'\beta_h \\ x_t(y_{t+h} - \alpha_h - x_t'\beta_h) \end{bmatrix} = 0. \quad (\text{A.7})$$

Because, in the notation of Bao and Ullah (2007), $\bar{H}_2 = \mathbf{0}$, to calculate $B_{h,LP}$ we only need to calculate the second element of $\mathbb{E}[QH_1Q\psi_{h,T-h}]$, which is given by

$$\mathbb{E}[QH_1Q\psi_{h,T-h}]_2 = -\frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \mathbb{E}[\phi_1(t,s) + \phi_2(t,s) + \phi_3(t,s)], \quad (\text{A.8})$$

where

$$\phi_1(t,s) = \frac{\varepsilon_t}{\sigma_\varepsilon^2} (y_{s+h} - \alpha_h - x_s'\beta_h), \quad (\text{A.9})$$

$$\phi_2(t,s) = \left(\frac{\varepsilon_t}{\sigma_\varepsilon^2}\right)^2 \varepsilon_s (y_{s+h} - \alpha_h - x_s'\beta_h), \quad (\text{A.10})$$

$$\phi_3(t,s) = \frac{\varepsilon_t}{\sigma_\varepsilon^2} c_{t-1}' \Sigma_{c,0}^{-1} c_{s-1} (y_{s+h} - \alpha_h - x_s'\beta_h). \quad (\text{A.11})$$

Consider first $\mathbb{E}[\phi_1(t,s)]$. When $t \leq s$ this expectation equals zero by Assumption 2. When $t > s+h$ this expectation is zero by Assumption 1. Finally, when $s < t \leq s+h$, then direct calculation yields $\mathbb{E}[\phi_1(t,s)] = \theta_{s+h-t}$. Next consider $\mathbb{E}[\phi_2(t,s)]$. If $t = s$, then under Assumption 2, this expectation is zero. If $t \neq s$, then the expectation is again zero owing to Assumptions 1 and 2. Finally, consider $\mathbb{E}[\phi_3(t,s)]$. When $t > s+h$, or $t \leq s$ $\mathbb{E}[\phi_3(t,s)] = 0$ by Assumptions 1 and 2. When $s < t \leq s+h$

$$\mathbb{E}[\phi_3(t,s)] = \mathbb{E} \left[\frac{\varepsilon_t}{\sigma_\varepsilon^2} c_{t-1}' \Sigma_{c,0}^{-1} c_{s-1} (u_{h-(t-s),t} + x_t'\beta_{h-(t-s)} - x_s'\beta_h) \right] \quad (\text{A.12})$$

$$= \mathbb{E} \left[\frac{\varepsilon_t}{\sigma_\varepsilon^2} c_{t-1}' \Sigma_{c,0}^{-1} c_{s-1} (u_{t,h-(t-s)} + x_t'\beta_{h-(t-s)}) \right] \quad (\text{A.13})$$

$$= \mathbb{E} \left[\frac{\varepsilon_t}{\sigma_\varepsilon^2} c_{t-1}' \Sigma_{c,0}^{-1} c_{s-1} (x_t'\beta_{h-(t-s)}) \right] \quad (\text{A.14})$$

$$= \mathbb{E} \left[\frac{1}{\sigma_\varepsilon^2} \mathbb{E}_{t-1} [\varepsilon_t (x_t'\beta_{h-(t-s)})] \text{tr} \left\{ c_{t-1}' \Sigma_{c,0}^{-1} c_{s-1} \right\} \right] \quad (\text{A.15})$$

$$= \theta_{h-(t-s)} \text{tr} \left\{ \Sigma_{c,0}^{-1} \Sigma_{c,t-s} \right\}. \quad (\text{A.16})$$

Note that equation (A.12) follows from the definition of $u_{h,t}$, equation (A.13) follows from Assumption 1, equation (A.14) follows from Assumption 2, equation (A.15) follows from the law of iterated expectations, and equation (A.16) follows from Assumption 1. Combining these resulting gives expression for $B_{h,LP}$ given in equation (3).

It is instructive to understand the effect of estimating α_h . To do so, we assume the mean is known—so that the intercepts of the regressions α_h are known—and compare the bias of the LP estimator to our benchmark specification. The moment conditions for the LP are

$$\mathbb{E} [x_t (y_{t+h} - x_t' \beta_h)] = 0 \quad (\text{A.17})$$

We are interested in the first element of β_h . Similar calculation to the case when α_h is estimated yields

$$\mathbb{E}[\widehat{\theta}_{h,LS} - \theta_h] = \mathbb{E} \frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} [\phi_2(t,s) + \phi_3(t,s)] + O(T^{-3/2}), \quad (\text{A.18})$$

where ϕ_2 and ϕ_3 are define in A.10 and A.11. It is immediate that the contribution to $B_{h,LP}$ from estimating the mean is given by

$$-\frac{1}{T-h} \sum_{j=1}^{T-h} \theta_{h-j}. \quad (\text{A.19})$$

Additionally, the approximate bias takes the form

$$B_{h,LP} = -\frac{1}{T-h} \sum_{j=1}^h \text{tr} \left\{ \Sigma_{c,0}^{-1} \Sigma_{c,j} \right\} \theta_{h-j} \text{ for } h > 0, \quad (\text{A.20})$$

A.2 LP with no controls

It is not necessary to include controls in equation (1) to have a consistent estimator of θ_h , and a sub-set of the related LP literature does not include controls. In this sub-section, we derive the approximate bias when controls are not included.

To do this, we maintain Assumption 1, but drop Assumption 2, and set c_{t-1} to the empty vector. We assume that the following moment conditions hold

$$\mathbb{E} \begin{bmatrix} y_{t+h} - \alpha_h - \theta_h \varepsilon_t \\ \varepsilon_t (y_{t+h} - \alpha_h - \theta_h \varepsilon_t) \end{bmatrix} = 0. \quad (\text{A.21})$$

To get an expression for the approximate bias, the arguments in A.1 go through under the obvious modifications. The resulting approximate bias is

$$-\frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \mathbb{E} [\phi_1(t,s) + \phi_2(t,s)], \quad (\text{A.22})$$

where $\phi_3(t, s)$ does not appear because no controls are included.

Consider first $\mathbb{E}[\phi_1(t, s)]$. When $t > s + h$ this expectation is zero by Assumption 1. When $t = s$, this expectation is zero by the moment condition. When $s < t \leq s + h$ or $t < s$, direct calculation yields $\mathbb{E}[\phi_1(t, s)] = \theta_{s+h-t}$. Next consider $\mathbb{E}[\phi_2(t, s)]$, which is zero. We can now state the approximate bias for the case without controls

Analytical Result 4 (Bias in LPs without controls.). *Under Assumption 1 and under the assumption that equation A.21 holds, the approximate bias of the least-squares estimator of θ_h*

$$-\frac{1}{T-h} \sum_{j=1}^{T-h-1} \left(1 - \frac{j}{T-h}\right) (\theta_{h+j} + \theta_{h-j}). \quad (\text{A.23})$$

A few comments are in order. First, the approximate bias is a function of the impulse response at all horizons up to $T-h-1$. Relative to the case with controls, this means that the impulse response at many more horizons, including those beyond the horizon of interest, affect the approximate bias. Second, it is not feasible to estimate all of the terms that enter the expression for the approximate bias with a finite sample. Any attempt to bias correct would need to truncate j . Third, it is not the case that the approximate bias without controls is necessarily smaller or larger than the approximate bias with controls. Instead, the magnitude of the bias depends on the particular regression models.

It is instructive to understand the effect of estimating α_h . To do this, we assume that the means of the data (and thus α_h) are known. For the LS estimator, we assume that the following moment condition holds

$$\mathbb{E} \left[\varepsilon_t (y_{t+h} - \theta_h \varepsilon_t) \right] = 0. \quad (\text{A.24})$$

The approximate bias is

$$-\frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \mathbb{E} [\phi_2(t, s)], \quad (\text{A.25})$$

where ϕ_2 is defined in equation (A.10). Note that $\phi_1(t, s)$ does not appear in this expression because the means are known, and $\phi_3(t, s)$ does not appear because no controls are included. It is immediate that from the arguments in the previous sub-section that when controls are not included, all of the approximate bias is due to estimating the mean.

A.3 LP with panel data

In this Appendix, we use the notation defined in Section 4. The moment conditions for the estimator are

$$0 = \mathbb{E} \left[\frac{1}{I} \sum_{i=1}^I x_{i,t} \left(y_{i,t+h} - x'_{i,t} \beta_h \right) \right]. \quad (\text{A.26})$$

Applying the results of Bao and Ullah (2007), the approximate bias of $\widehat{\theta}_{h,LS,I}$ is given by

$$-\frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \mathbb{E} [\Phi_1(t, s) + \Phi_2(t, s) + \Phi_3(t, s)], \quad (\text{A.27})$$

where

$$\Phi_1(t, s) = \mathbb{E} \left[\frac{1}{I} \sum_{i=1}^I \frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,I}^2} \left(y_{i,s+h} - x'_{i,s} \beta_h \right) \right] \quad (\text{A.28})$$

$$\Phi_2(t, s) = \mathbb{E} \left[\frac{1}{I} \sum_{i=1}^I \frac{\varepsilon_{i,t}^2}{\left(\sigma_{\varepsilon,I}^2 \right)^2} \frac{1}{I} \sum_{k=1}^I \varepsilon_{k,s} \left(y_{k,s+h} - x'_{k,s} \beta_h \right) \right] \quad (\text{A.29})$$

$$\Phi_3(t, s) = \mathbb{E} \left[\frac{1}{I} \sum_{i=1}^I \frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,I}^2} \frac{1}{I} \sum_{k=1}^I \tilde{c}'_{i,t-1} \Sigma_{c,0,I}^{-1} \tilde{c}_{k,s-1} \left(y_{k,s+h} - x'_{k,s} \beta_h \right) \right] \quad (\text{A.30})$$

Consider $\Phi_1(t, s)$. Under Assumptions 3 and 4, $\Phi_1(t, s) = 0$ when $t \leq s$ or $t > s+h$. If $s < t \leq s+h$, $\Phi_1(t, s) = \theta_{s+h-t}$. Consider $\Phi_2(t, s)$. Under Assumptions 3 and 4, $\Phi_2(t, s) = 0$ for all t and s . Consider $\Phi_3(t, s)$. Under Assumptions 3 and 4, $\Phi_3(t, s) = 0$ if $t \leq s$ or $t > s+h$. When $s < t \leq s+h$

$$\Phi_3(t, s) = \mathbb{E} \left[\frac{1}{I} \sum_{i=1}^I \frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,I}^2} \frac{1}{I} \sum_{k=1}^I \tilde{c}'_{i,t-1} \Sigma_{c,0,I}^{-1} \tilde{c}_{k,s-1} \left(u_{k,h-(t-s),t} + x'_{k,t} \beta_{h-(t-s)} \right) \right] \quad (\text{A.31})$$

$$= \mathbb{E} \left[\frac{1}{I} \sum_{i=1}^I \frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,I}^2} \frac{1}{I} \sum_{k=1}^I \tilde{c}'_{i,t-1} \Sigma_{c,0,I}^{-1} \tilde{c}_{k,s-1} x'_{k,t} \beta_{h-(t-s)} \right] \quad (\text{A.32})$$

$$= \mathbb{E} \left[\frac{1}{I} \sum_{i=1}^I \frac{1}{\sigma_{\varepsilon,I}^2} \frac{1}{I} \sum_{k=1}^I \mathbb{E}_{t-1} \left[\varepsilon_{i,t} x'_{k,t} \beta_{h-(t-s)} \right] \tilde{c}'_{i,t-1} \Sigma_{c,0,I}^{-1} \tilde{c}_{k,s-1} \right] \quad (\text{A.33})$$

$$= \frac{1}{I} \sum_{i=1}^I \frac{1}{I} \sum_{k=1}^I \text{tr} \left\{ \Sigma_{c,0,I}^{-1} \mathbb{E} \left[\tilde{c}_{k,s-1} \tilde{c}'_{i,t-1} \right] \right\} \frac{\sigma_{\varepsilon,i,k}}{\sigma_{\varepsilon,I}^2} \theta_{s+h-t}. \quad (\text{A.34})$$

Note that A.31 follows from the definition of $u_{k,h-(t-s),t}$ and Assumption 3, A.32 follows from Assumption 4, A.33 follows from the law of iterated expectations, and A.34 follows from Assumptions 3 and 4. Combining these observations delivers the results reported in Analytic Result 3.

It is instructive to understand the effect of estimating $\alpha_{i,h}$. In this case, the moment conditions are

$$0 = \mathbb{E} \left[\frac{1}{I} \sum_{i=1}^I [\varepsilon_{i,t}, c'_{i,t-1}, c'_{t-1}]' (y_{i,t+h} - x'_{i,t} \beta_h) \right], \quad (\text{A.35})$$

where the first I elements of β_h are known. Algebra similar to the previous sub-section yields that the approximate bias is given by

$$-\frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \mathbb{E} [\Phi_2(t, s) + \Phi_3(t, s)], \quad (\text{A.36})$$

where $\Phi_2(t, s)$ and $\Phi_3(t, s)$ are as defined in equations (A.29) and (A.30). It is immediate that the contribution to $B_{h,LP,I}$ from estimating the mean is given by

$$-\frac{1}{T-h} \sum_{j=1}^{T-h} \theta_{h-j}. \quad (\text{A.37})$$

B Analytic comparison between LP and parametric model

Here, we provide details of the derivation of the bias in the AR estimator $\widehat{\theta}_0 \widehat{\rho}^h$ discussed in Section 3.2.3 and under the data generating process given in equation 5. To derive the bias, we work with the following moment conditions

$$\mathbb{E} \begin{bmatrix} y_t - \alpha - \varepsilon_t \theta_0 - \rho y_{t-1} \\ \varepsilon_t (y_t - \alpha - \varepsilon_t \theta_0 - \rho y_{t-1}) \\ y_{t-1} (y_t - \alpha - \varepsilon_t \theta_0 - \rho y_{t-1}) \\ \theta_h - \theta_0 \rho^h \end{bmatrix} = \mathbf{0}. \quad (\text{B.1})$$

Note that Assumptions 1 and 2 are satisfied. For notational ease, and without loss of generality we assume that $\mathbb{E}[y_t] = 0$, meaning $\alpha = 0$.

We proceed by calculating the two terms defined in equation (A.5). With appropriately defined vectors and matrices analogous to those discussed in Appendix A, $\mathbb{E}[QVQ\psi_{h,T-1}]_4$ is given by

$$\left(\frac{1}{T-1} \right)^2 \sum_{t=2}^T \sum_{s=2}^T \mathbb{E} \left[- \left(\frac{\rho^h}{\sigma_\varepsilon^2} \varepsilon_t + \frac{\theta_0 h \rho^{h-1}}{\sigma_y^2} y_{t-1} \right) \left(1 + \frac{\varepsilon_t \varepsilon_s}{\sigma_\varepsilon^2} + \frac{y_{t-1} y_{s-1}}{\sigma_y^2} \right) \nu_s \right]. \quad (\text{B.2})$$

It can be shown that this quantity can be expressed as

$$\mathbb{E}[QVQ\psi_{h,T-1}]_4 = -\frac{1}{T-1} \theta_0 h \rho^{h-1} \gamma (1 + 3\rho) + O(T^{-3/2}), \quad (\text{B.3})$$

where $\gamma = \frac{\sigma_\nu^2}{\sigma_\nu^2 + \theta_0^2 \sigma_\varepsilon^2}$. The second term in equation (A.5) is given by:

$$\begin{aligned} \mathbb{E} \left[\frac{1}{2} Q \bar{H}_2 [(Q\psi_{h,T-h}) \otimes (Q\psi_{h,T-h})] \right]_4 &= -\frac{1}{2} h \rho^{h-1} \left(\frac{1}{T-1} \right)^2 \sum_{t=2}^T \sum_{s=2}^T \mathbb{E} \left[\frac{\varepsilon_s}{\sigma_\varepsilon^2} \nu_s \frac{y_{t-1}}{\sigma_y^2} \nu_t \right] \\ &\quad - \frac{1}{2} h \rho^{h-1} \left(\frac{1}{T-1} \right)^2 \sum_{t=2}^T \sum_{s=2}^T \mathbb{E} \left[\frac{\varepsilon_t}{\sigma_\varepsilon^2} \nu_t \frac{y_{s-1}}{\sigma_y^2} \nu_s \right] \\ &\quad - \frac{1}{2} \theta_0 h (h-1) \rho^{h-2} \left(\frac{1}{T-1} \right)^2 \sum_{t=2}^T \sum_{s=2}^T \mathbb{E} \left[\frac{y_{s-1}}{\sigma_y^2} \nu_s \frac{y_{t-1}}{\sigma_y^2} \nu_t \right] \end{aligned} \tag{B.4}$$

$$= -\frac{1}{2} \theta_0 h (h-1) \rho^{h-2} \frac{1}{T-1} \gamma (1-\rho^2). \tag{B.5}$$

Then

$$B_{h,AR} = -\frac{1}{T-1} \gamma \theta_0 h \rho^{h-1} \left[(1+3\rho) - \frac{1}{2} (h-1) (1-\rho^2) \rho^{-1} \right], \tag{B.6}$$

which is the expression given in Section 3.2.3.

C Additional analysis in the context of simple data generating processes

In section 3.2 we analyzed $B_{h,LP}$ in the context of an AR(1) data generating process given by equation (5). In this Appendix, we analyze the quality of the approximate bias when means are known, different parameterizations of equation (5) and an AR(2) model that generates hump-shaped impulse response functions.

C.1 The quality of the approximation in an AR(1) when the means are known

Here, we analyze the quality of the approximation to the bias in the context of the AR(1) data generating process in equation (5) when the means are known, as derived in Appendix A.1. Under this assumption, figure C.1 shows analogous results to those reported in figure 2.

As in the case when α_h is estimated, for $\rho = 0.9$ and $\rho = 0.95$, our analytic expression is a good approximation to the exact finite-sample bias in $\hat{\theta}_{h,LS}$ for all h shown. The quality of the approximation improves somewhat as T increases. Clearly, when $\rho = 0.99$, our analytic expression is not as good of an approximation as it is for smaller values of ρ . We conclude that when the means are known and for empirically relevant sample sizes, our analytic expression offers a reasonable

approximation to the bias in LPs, though the quality of the approximation is somewhat worse for smaller values of T and larger values of ρ .

C.2 Alternative parameterizations of the AR(1) data generating process

In the main text, we assumed $\sigma_\varepsilon = \sigma_\nu = 1$. Here, we consider different settings.

First, consider the case where $\sigma_\varepsilon = 1$ and $\sigma_\nu = 0.1$. In this case, the structural shock ε_t explains almost all of the variation in y_t . Using this parameterization, figure C.2 shows analogous results to those reported in figure 2. Note that the values of σ_ε and σ_ν do not appear in $B_{h,LP}$. So, the line representing $B_{h,LP}$ are identical in figures 2 and C.2. Additionally, it is apparent that the Monte Carlo means of $\hat{\theta}_{LS}$ are also very similar when comparing figures 2 and C.2. Figure C.3 shows analogous results to those reported in Figure 3. While the average bias correction (left panels) is little changed by the change in σ_ν , the performance of $\hat{\theta}_{h,BCC}$ improves relative to $\hat{\theta}_{h,LS}$ as measured by MSE.

Next, consider the case where $\sigma_\varepsilon = 0.1$ and $\sigma_\nu = 1$. In this case, the structural shock ε explains almost none of the variation in y_t . Using this parameterization, figure C.4 shows analogous results to those reported in figure 2. Note that the values of σ_ε and σ_ν do not appear in $B_{h,LP}$. So, the line representing $B_{h,LP}$ are identical in figures 2 and C.4. Additionally, the Monte Carlo means of $\hat{\theta}_{LS}$ are also similar when comparing figures 2 and C.4. Figure C.5 shows the performance of the bias corrected estimator under this data generating process. The bias correction is again unaffected by the relative sizes of σ_ε and σ_ν , but the MSE of $\hat{\theta}_{BCC}$ is worse, relative to $\hat{\theta}_{LS}$, because the bias term contributes little to the MSE.

C.3 An AR(2) example

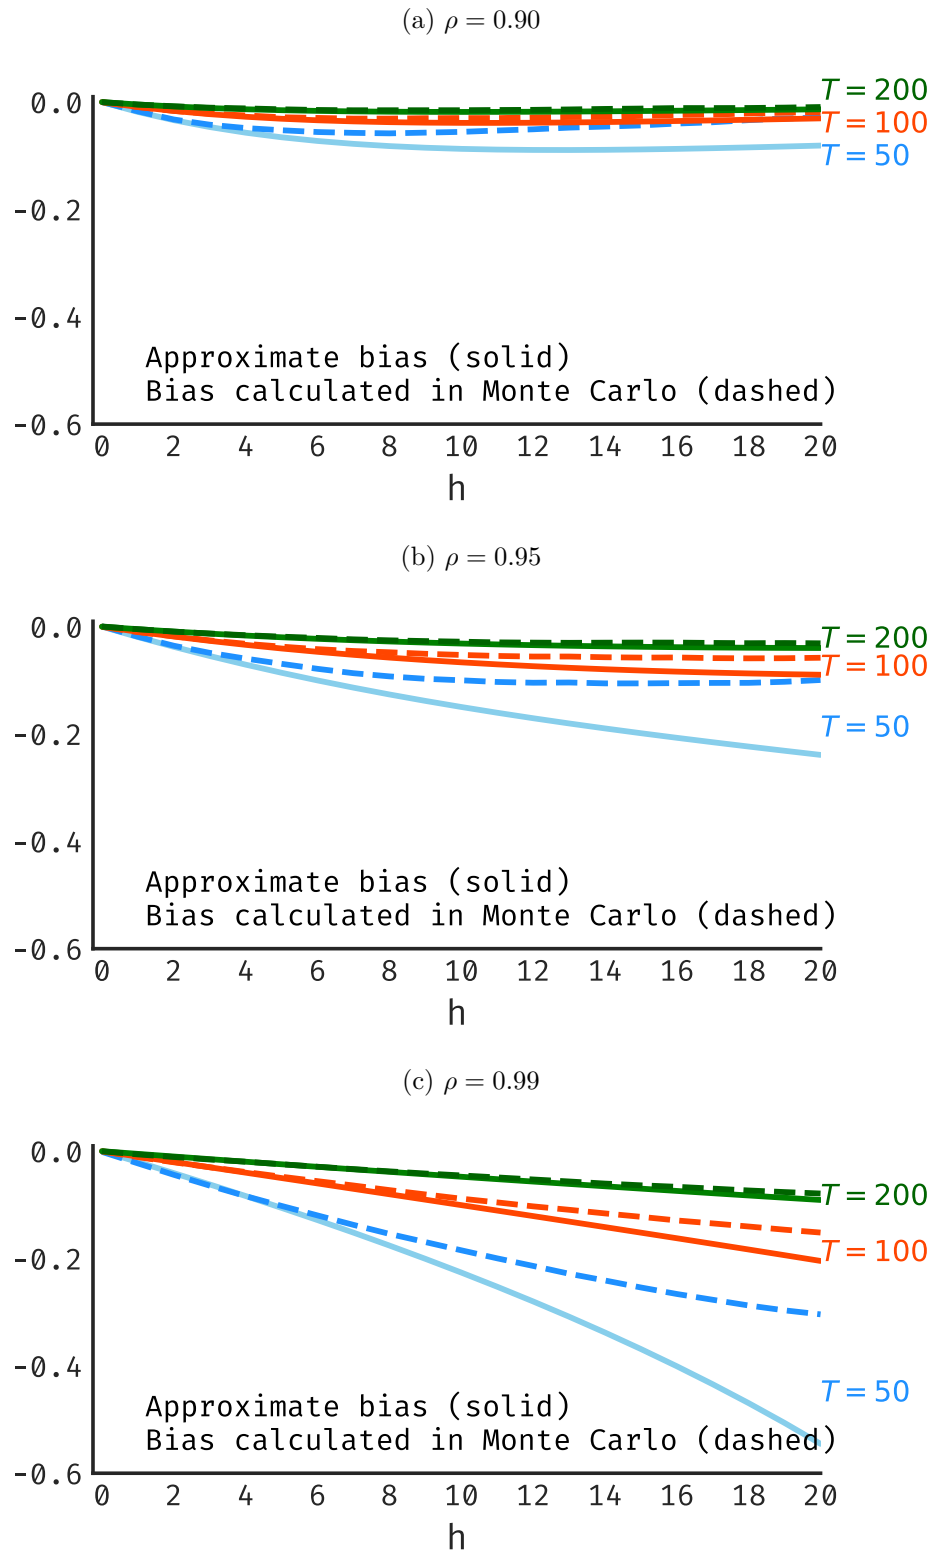
Here, we analyze the performance of the approximation $B_{h,LP}$ using a simple AR(2) model. We specify the data generating process so that

$$y_t = (\rho + \psi)y_{t-1} - \psi\rho y_{t-2} + \theta_0\varepsilon_t + \nu_t \quad (\text{C.1})$$

where $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$, $\nu_t \sim N(0, \sigma_\nu^2)$, and each is independent of each other and over time. We set $\psi = 0.4$ and $\theta_0 = 1$. For the values of ρ that we consider, this data generating process delivers hump-shaped impulse response functions. This particular formulation of the AR(2) is useful because ρ plays a similar role to the AR(1) in that it determines the persistence of y_t . We consider different values for ρ within each figure. The LP is specified so that $x_t = [\varepsilon_t, y_{t-1}, y_{t-2}]'$.


As was the case for the AR(1), For $\rho = 0.9$ and $\rho = 0.95$, $B_{h,LP}$ is a good approximation to the exact finite-sample bias in $\hat{\theta}_{h,LS}$ for all h shown. The quality of the approximation improves somewhat as T increases. Clearly, when $\rho = 0.99$, $B_{h,LP}$ is not as good of an approximation as it is for smaller values of ρ . Nevertheless, even with $\rho = 0.99$ and $T = 50$, $B_{h,LP}$ captures salient features of the finite-sample bias, including that it is growing in magnitude with h over the values of h shown. We conclude that the quality of the approximation is somewhat worse for smaller values of T and larger values of ρ .

Figure C.1: The approximate bias performs well in empirically-relevant samples when y_t is an AR(1) and the means are known.



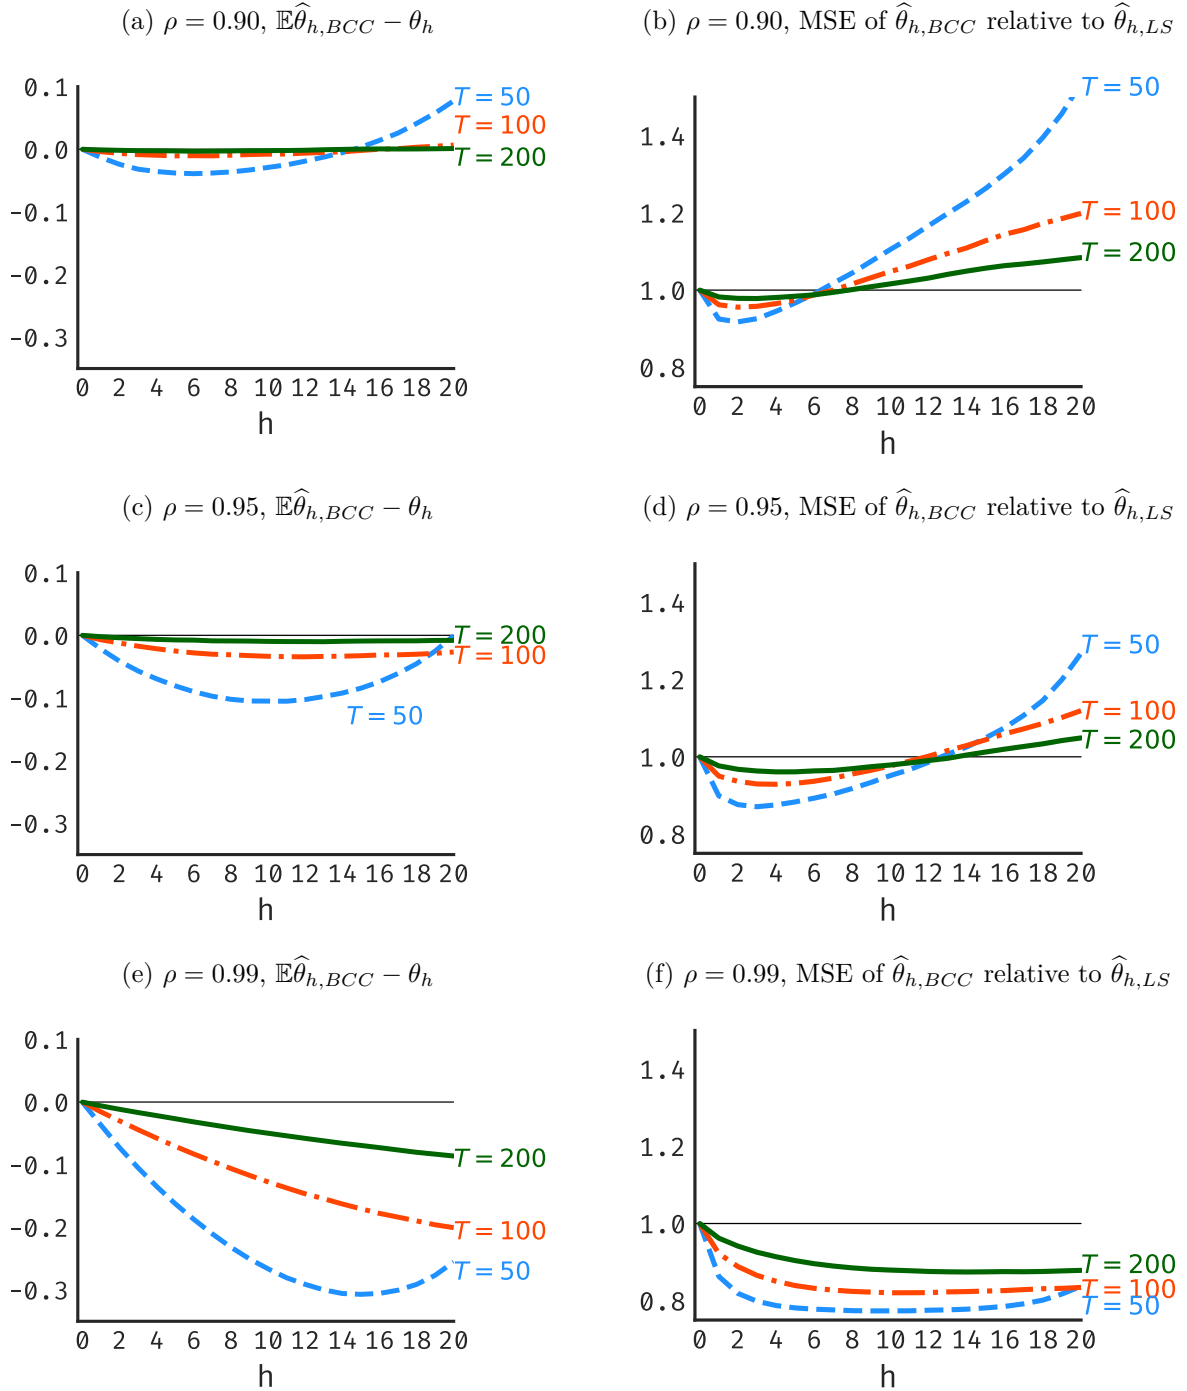
Note: This figure is analogous to Figure 2, but the estimation is done under the assumption that the means of the data are known. The sub-figures show the value of the approximate bias and the Monte Carlo means of $\hat{\theta}_{h,LS} - \theta_h$ estimated on data simulated from equation (5) under the assumption that α_h is known. We use 100,000 Monte Carlo simulations. We set $\sigma_\varepsilon = \sigma_\nu = 1$ and $\theta_0 = 1$.

Figure C.2: $B_{h,LP}$ performs well in empirically-relevant samples when y_t is an AR(1) and σ_ν is small.




Note: This figure is analogous to Figure 2, but with $\sigma_\varepsilon = 1$ and $\sigma_\nu = 0.1$ (and $\theta_0 = 1$). The sub-figures show the value of $B_{h,LP}$ and the Monte Carlo means of $\hat{\theta}_{h,LS} - \theta_h$ estimated on data simulated from equation (5). We use 100,000 Monte Carlo simulations.

Figure C.3: Performance of $B_{h,BCC}$ in the AR(1) example when σ_ν is small.



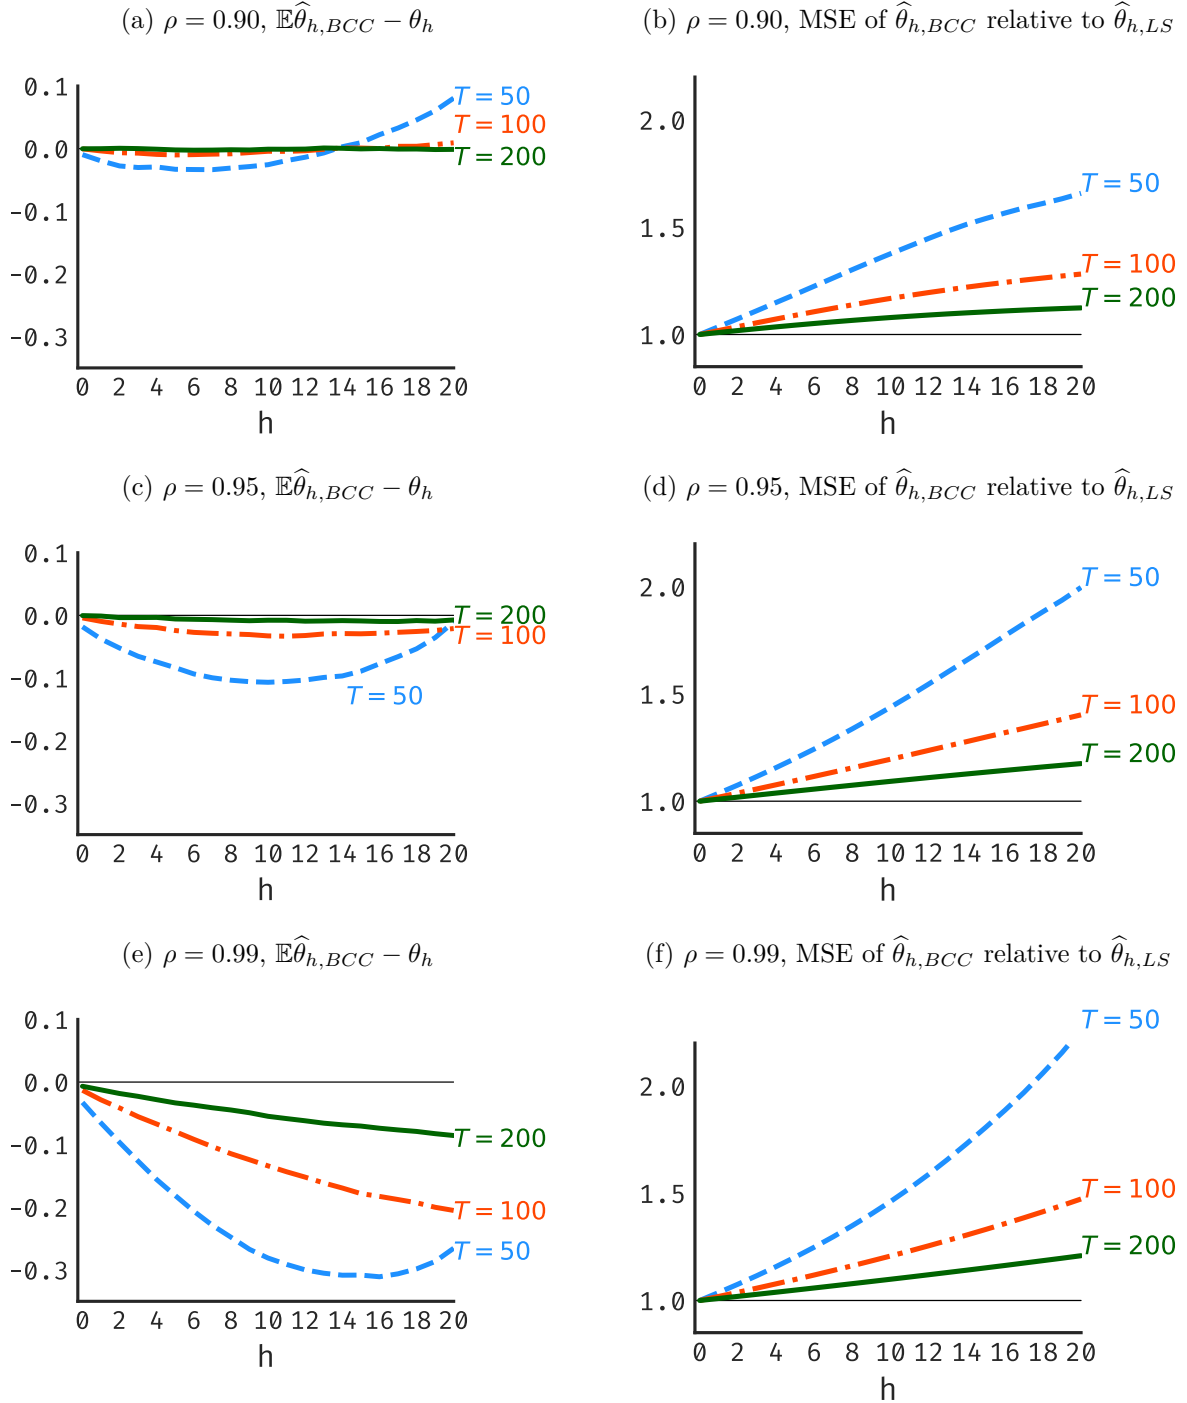
Note: This figure is analogous to Figure 3, but with $\sigma_\varepsilon = 1$ and $\sigma_\nu = 0.1$ (and $\theta_0 = 1$). The sub-figures on the left show the Monte Carlo means of $\hat{\theta}_{h,BCC} - \theta_h$ estimated on data simulated from equation (5). The sub-figures on the right show the Monte Carlo value of the MSE of $\hat{\theta}_{h,BCC}$ relative to the MSE of $\hat{\theta}_{h,LS}$. A number larger than unity indicates that the MSE of $\hat{\theta}_{h,BCC}$ is larger than $\hat{\theta}_{h,LS}$. We use 100,000 Monte Carlo simulations.

Figure C.4: $B_{h,LP}$ performs well in empirically-relevant samples when y_t is an AR(1) and σ_ε is small.



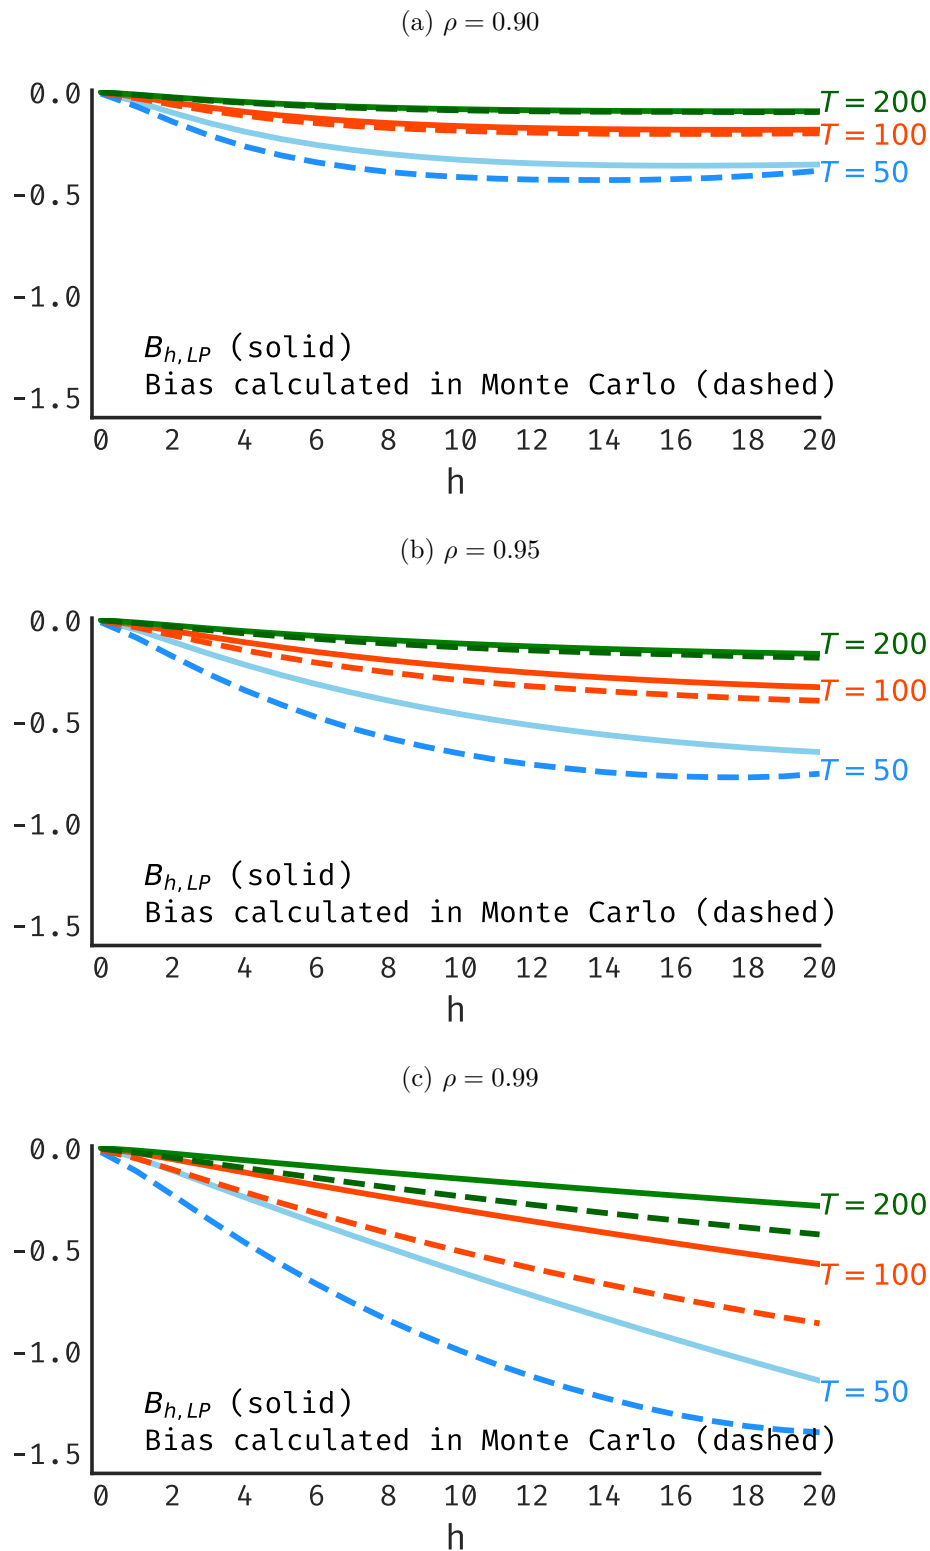
Note: This figure is analogous to Figure 2, but with $\sigma_\varepsilon = 0.1$ and $\sigma_\nu = 1$ (and $\theta_0 = 1$). The sub-figures show the value of $B_{h,LP}$ and the Monte Carlo means of $\hat{\theta}_{h,LS} - \theta_h$ estimated on data simulated from equation (5). We use 100,000 Monte Carlo simulations.

Figure C.5: Performance of $B_{h,BCC}$ in the AR(1) example when σ_ε is small.



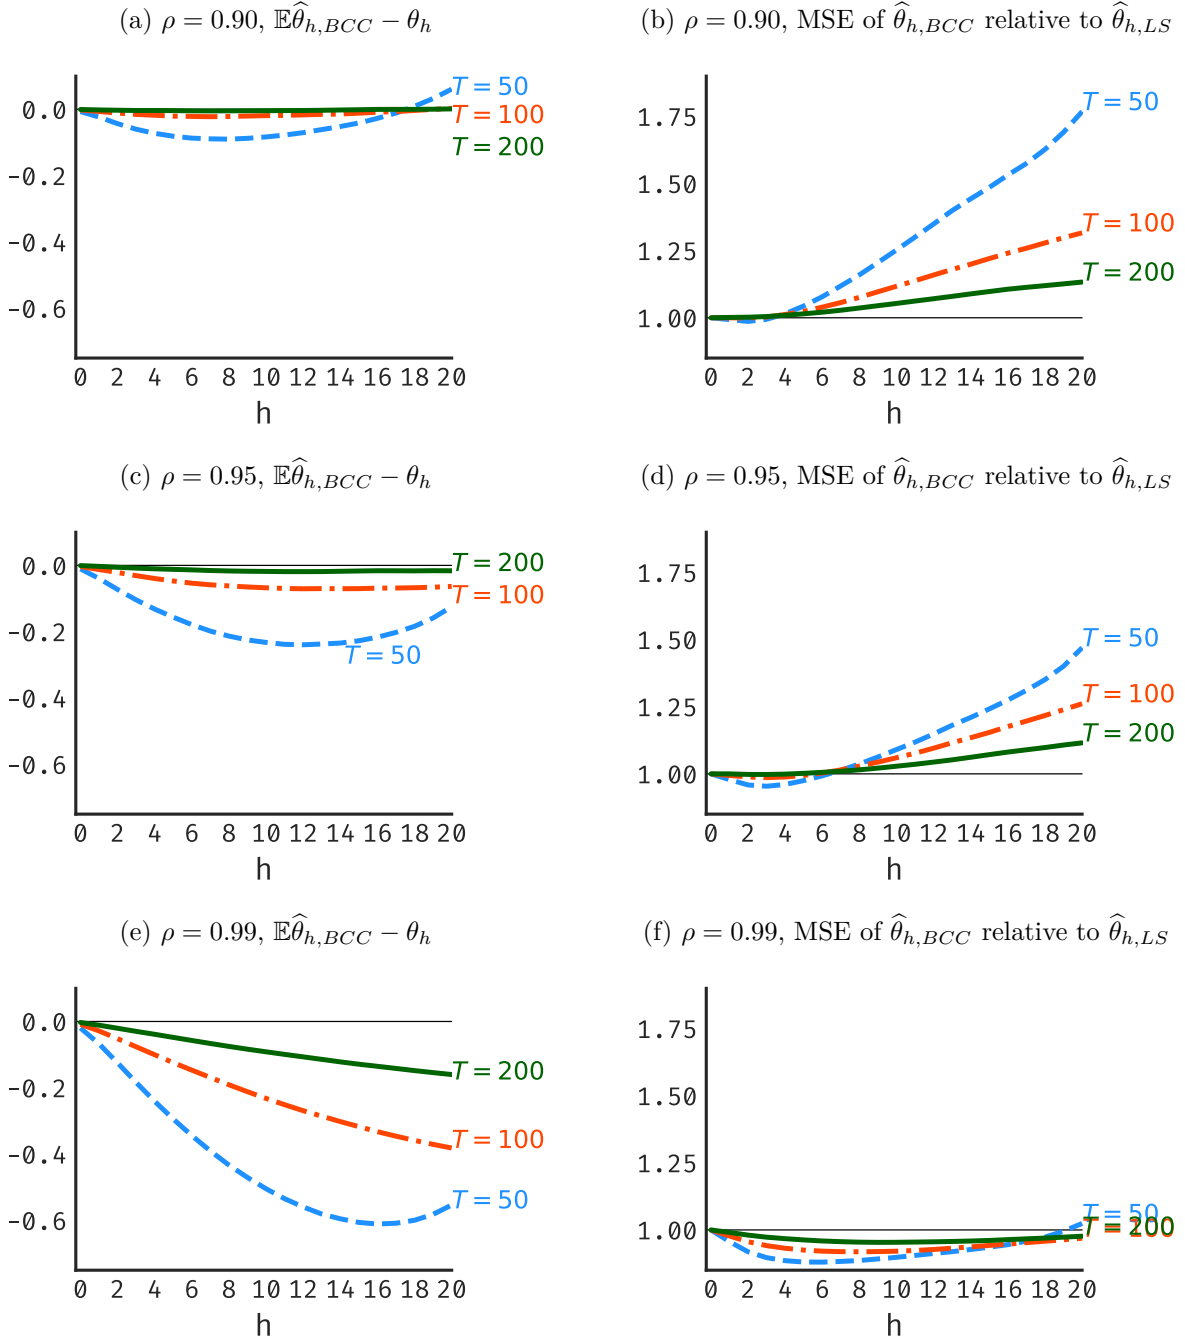
Note: This figure is analogous to Figure 3, but with $\sigma_\varepsilon = 0.1$ and $\sigma_\nu = 1$ (and $\theta_0 = 1$). The sub-figures on the left show the Monte Carlo means of $\hat{\theta}_{h,BCC} - \theta_h$ estimated on data simulated from equation (5). The sub-figures on the right show the Monte Carlo value of the MSE of $\hat{\theta}_{h,BCC}$ relative to the MSE of $\hat{\theta}_{h,LS}$. A number larger than unity indicates that the MSE of $\hat{\theta}_{h,BCC}$ is larger than $\hat{\theta}_{h,LS}$. We use 100,000 Monte Carlo simulations.

Figure C.6: $B_{h,LP}$ performs well in empirically-relevant samples when y_t is an AR(2).



Note: the sub-figures show the value of $B_{h,LP}$ and the Monte Carlo means of $\hat{\theta}_{LS}$ estimated on data simulated from equation (C.1). We use 100,000 Monte Carlo simulations. We set $\sigma_\varepsilon = \sigma_\nu = 1$ and $\theta_0 = 1$.

Figure C.7: Performance of $B_{h,BCC}$ in an AR(2) example.



Note: the sub-figures on the left show the Monte Carlo means of $\hat{\theta}_{h,BCC} - \theta_h$ estimated on data simulated from equation (5). The sub-figures on the right show the Monte Carlo value of the MSE of $\hat{\theta}_{h,BCC}$ relative to the MSE of $\hat{\theta}_{h,LS}$. A number larger than unity indicates that the MSE of $\hat{\theta}_{h,BCC}$ is larger than $\hat{\theta}_{h,LS}$. We use 100,000 Monte Carlo simulations. We set $\sigma_\varepsilon = \sigma_\nu = 1$ and $\theta_0 = 1$.

D Derivations related to standard errors

Here, we provide details of the derivation of the bias in the GMM estimator of the autocovariance of the regression score under the assumption that $\theta_h = 0$ for all h . Our notation mirrors that of Appendix A.

D.1 Derivations related to SEs in LP with controls

To derive the approximate bias in $\widehat{\gamma}_{h,\ell}$, in an LP with controls, we work with the following moment conditions

$$\mathbb{E} \begin{bmatrix} y_{t+h} - \alpha - x'_t \beta_h \\ \varepsilon_t (y_{t+h} - \alpha - x'_t \beta_h) \\ c_{t-1} (y_{t+h} - \alpha - x'_t \beta_h) \\ \varepsilon_t (y_{t+h} - \alpha - x'_t \beta_h) \varepsilon_{t-\ell} (y_{t-\ell+h} - \alpha - x'_{t-\ell} \beta_h) - \gamma_{h,\ell} \end{bmatrix} = \mathbf{0}, \quad (\text{D.1})$$

where $\ell > 0$. We maintain Assumptions 1 and 2 and assume that Assumptions A-C in Rilstone et al. (1996) are satisfied. For notational ease, and without loss of generality, we assume that all data have zero mean. Note that the GMM empirical moment function associated with equation (D.1) will only have $T - h - \ell$ observations owing to the additional autocovariance term.

With appropriately defined vectors and matrices analogous to those discussed in Appendix A that also are in the notation of Bao and Ullah (2007), the final element of $\mathbb{E}[QVQ\psi_{h,T-h-\ell}]$ is given by

$$\left(\frac{1}{T-h-\ell} \right)^2 \sum_{t=1+\ell}^{T-h} \sum_{s=1+\ell}^{T-h} \mathbb{E}[\Upsilon_1(t,s) + \Upsilon_2(t,s) + \Upsilon_3(t,s) + \Upsilon_4(t,s)], \quad (\text{D.2})$$

where

$$\Upsilon_1(t,s) = -[\varepsilon_{t-\ell} \varepsilon_t u_{h,t} + \varepsilon_t \varepsilon_{t-\ell} u_{h,t-\ell}] u_{h,s}, \quad (\text{D.3})$$

$$\Upsilon_2(t,s) = -[\varepsilon_t \varepsilon_{t-\ell}^2 u_{h,t} + \varepsilon_t^2 \varepsilon_{t-\ell} u_{h,t-\ell}] \frac{\varepsilon_s}{\sigma_\varepsilon^2} u_{h,s}, \quad (\text{D.4})$$

$$\Upsilon_3(t,s) = -[c'_{t-\ell-1} \varepsilon_{t-\ell} \varepsilon_t u_{h,t} + \varepsilon_t c'_{t-1} \varepsilon_{t-\ell} u_{h,t-\ell}] \Sigma_c^{-1} c_{s-1} u_{h,s}, \quad (\text{D.5})$$

$$\Upsilon_4(t,s) = -\varepsilon_t u_{h,t} \varepsilon_{s-\ell} u_{h,s-\ell} + \gamma_{h,\ell}. \quad (\text{D.6})$$

It can be shown that for all t and s , $\mathbb{E}\Upsilon_1(t,s) = 0$, $\mathbb{E}\Upsilon_3(t,s) = 0$, $\mathbb{E}\Upsilon_4(t,s) = 0$. If $t = s$ or $t = s + \ell$, then

$$\mathbb{E}\Upsilon_2(t,s) = -\sigma_\varepsilon^2 \mathbb{E} \left[(y_{t+h} - \alpha_h - x'_t \beta_h)^2 \right]. \quad (\text{D.7})$$

If $t \neq s$ and $t \neq s + \ell$, then $\mathbb{E}\Upsilon_2(t, s) = 0$. So, the final element of $\mathbb{E} [Q (H_1 - \bar{H}_1) Q\psi_{h, T-h-\ell}]$ is given by

$$-2 \left(\frac{1}{T-h-\ell} \right) \sigma_\varepsilon^2 \mathbb{E} \left[(y_{t+h} - \alpha_h - x'_t \beta_h)^2 \right] + O \left(T^{-3/2} \right). \quad (\text{D.8})$$

Additionally, the final element of $\mathbb{E} \left[\frac{1}{2} Q \bar{H}_2 [a_{-1/2} \otimes a_{-1/2}] \right]$ is given by

$$\begin{aligned} & - \left(\frac{1}{T-h-\ell} \right)^2 \sum_{t=1+\ell}^{T-h} \sum_{s=1+\ell}^{T-h} \mathbb{E} \left[\varepsilon_t (y_{t+h} - \alpha_h - x'_t \beta_h) \varepsilon_s (y_{s+h} - \alpha_h - x'_s \beta_h) \right] \\ & = - \left(\frac{1}{T-h-\ell} \right) \sigma_\varepsilon^2 \mathbb{E} \left[(y_{s+h} - \alpha_h - x'_s \beta_h)^2 \right] \end{aligned} \quad (\text{D.9})$$

Noting that $\gamma_{h,0} = \sigma_\varepsilon^2 \mathbb{E} \left[(y_{s+h} - \alpha_h - x'_s \beta_h)^2 \right]$, the approximate bias of $\hat{\gamma}_{h,\ell}$ is

$$-\frac{1}{T-h-\ell} \gamma_{h,0}. \quad (\text{D.10})$$

It is useful to parse the effect of estimating α_h . If the means of the data are known (α_h is known), the moment conditions are

$$\mathbb{E} \begin{bmatrix} \varepsilon_t (y_{t+h} - \alpha - x'_t \beta_h) \\ c_{t-1} (y_{t+h} - \alpha - x'_t \beta_h) \\ \varepsilon_t (y_{t+h} - \alpha - x'_t \beta_h) \varepsilon_{t-\ell} (y_{t-\ell+h} - \alpha - x'_{t-\ell} \beta_h) - \gamma_{h,\ell} \end{bmatrix} = \mathbf{0}, \quad (\text{D.11})$$

where $\ell > 0$. It can be shown that the final element of $\mathbb{E} [Q (H_1 - \bar{H}_1) Q\psi_{h, T-h-\ell}]$ is given by

$$\left(\frac{1}{T-h-\ell} \right)^2 \sum_{t=1+\ell}^{T-h} \sum_{s=1+\ell}^{T-h} \mathbb{E} [\Upsilon_2(t, s) + \Upsilon_3(t, s) + \Upsilon_4(t, s)], \quad (\text{D.12})$$

and that the final element of $\frac{1}{2} Q \bar{H}_2 [a_{-1/2} \otimes a_{-1/2}]$ is unchanged from the case when controls are included. Because $\mathbb{E}\Upsilon_1(t, s) = 0$ for all t and s , the approximate bias of $\hat{\gamma}_{h,\ell}$ is the same as in the case when α_h is estimated.

D.2 Derivations related to SEs in LP with no controls

Here, we consider the approximate bias of $\hat{\gamma}_{h,\ell}$ in an LP without controls. We maintain Assumption 1 and assume that Assumptions A-C in Rilstone et al. (1996) are satisfied. The moment conditions are the same as in the LP with controls, but with $c_{t-1} = \emptyset$. It can be shown that the final element of $\mathbb{E} [Q (H_1 - \bar{H}_1) Q\psi_{h, T-h-\ell}]$ is given by

$$\left(\frac{1}{T-h-\ell} \right)^2 \sum_{t=1+\ell}^{T-h} \sum_{s=1+\ell}^{T-h} \mathbb{E} [\Upsilon_1(t, s) + \Upsilon_2(t, s) + \Upsilon_4(t, s)], \quad (\text{D.13})$$

and the final element of $\frac{1}{2}Q\bar{H}_2 [a_{-1/2} \otimes a_{-1/2}]$ is unchanged in form from when controls are included. Because $\theta_h = 0$, it is immediate that the approximate bias of $\hat{\gamma}_{h,\ell}$ is

$$-\frac{1}{T-h-\ell}\sigma_\varepsilon^2\mathbb{E}[y_t^2] = -\frac{1}{T-h-\ell}\gamma_{h,0}. \quad (\text{D.14})$$

Following similar arguments to those made in the case of controls, it is easily shown that the approximate bias of $\hat{\gamma}_{h,\ell}$ is unchanged in the case without controls if α_h is known.

E CEE VAR

Christiano et al. (2005) estimate a 9 variable VAR(4) using quarterly U.S. data on real GDP, real consumption, real investment, GDP deflator prices, real wages, labor productivity, federal funds rate, real profits, and the growth rate of M2. All variables except the federal funds rate and M2 growth rate enter the VAR in log levels.

E.1 Data Construction

Following Christiano et al. (2005), our sample begins in 1965Q3 and ends in 1995Q3. Details on the variable construction are given below. Variable names from FRED (<https://fred.stlouisfed.org>) are provided. In instances where the data has higher than quarterly frequency, we use quarterly averages.

1. **Real GDP.** Take the level of real gross domestic product (FRED mnemonic `GDPC1`). The VAR observable is constructed as:

$$\text{Real GDP}_t = \log(GDPC1_t).$$

Citation: U.S. Bureau of Economic Analysis, Real Gross Domestic Product [GDPC1], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/GDPC1>.

2. **Real Consumption.** Take the level of real consumption (FRED mnemonic `PCECC96`). The VAR observable is constructed as:

$$\text{Real Consumption}_t = \log(PCECC96_t).$$

Citation: U.S. Bureau of Economic Analysis, Real Personal Consumption Expenditures [PCECC96], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/PCECC96>.

3. **GDP Deflator Prices.** Take the level of GDP deflator prices (FRED mnemonic `GDPDEF`). The VAR observable is constructed as:

$$\text{Price Level}_t = \log(GDPDEF_t).$$

Citation: U.S. Bureau of Economic Analysis, Gross Domestic Product: Implicit Price Deflator [`GDPDEF`], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/GDPDEF>.

4. **Real Investment.** Take the level of real gross domestic private investment (FRED mnemonic `GPDI1`). The VAR observable is constructed as:

$$\text{Real Investment}_t = \log(GPDI1_t).$$

Citation: U.S. Bureau of Economic Analysis, Real Gross Private Domestic Investment [`GPDI1`], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/GPDI1>.

5. **Real Wages.** Take hourly compensation for all employed persons (FRED mnemonic `COMPNFB`), the level of the GDP deflator prices, and the average weekly hours worked (FRED mnemonic `PRS85006023`). The VAR observable is constructed as:

$$\text{Real Wages}_t = \log\left(\frac{COMPNFB_t}{GDPDEF_t \times PRS85006023_t}\right).$$

Citations: U.S. Bureau of Labor Statistics, Nonfarm Business Sector: Hourly Compensation for All Workers [`COMPNFB`], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/COMPNFB>. U.S. Bureau of Labor Statistics, Nonfarm Business Sector: Average Weekly Hours for All Workers [`PRS85006023`], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/PRS85006023>.

6. **Labor Productivity.** Take labor productivity (FRED mnemonic `OPHNFB`). The VAR observable is constructed as:

$$\text{Labor Productivity}_t = \log(OPHNFB_t).$$

Citation: U.S. Bureau of Labor Statistics, Nonfarm Business Sector: Labor Productivity (Output per Hour) for All Workers [`OPHNFB`], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/OPHNFB>.

7. **Federal Funds Rate.** Take the quarterly average of the monthly federal funds rate (FRED mnemonic FEDFUNDS). The VAR observable is constructed as:

$$\text{Federal Funds Rate}_t = FEDFUNDS_t.$$

Citation: Board of Governors of the Federal Reserve System (US), Federal Funds Effective Rate [FEDFUNDS], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/FEDFUNDS>.

8. **Real Profits.** Take corporate profits after tax (FRED mnemonic CP). The VAR observable is constructed as:

$$\text{Real Profits}_t = \log \left(\frac{CP_t}{GDPDEF_t} \right).$$

Citations: U.S. Bureau of Economic Analysis, Corporate Profits After Tax (without IVA and CCAdj) [CP], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/CP>. U.S. Bureau of Economic Analysis, Gross Domestic Product: Implicit Price Deflator [GDPDEF], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/GDPDEF>.

9. **The growth rate of M2.** Take the M2 monetary aggregate (FRED mnemonic M2). The VAR observable is constructed as:

$$\Delta M2_t = 100 \times \left(\frac{M2_t - M2_{t-1}}{M2_{t-1}} \right).$$

Citation: Board of Governors of the Federal Reserve System (US), M2 (DISCONTINUED) [M2], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/M2>.

E.2 Identification of the Monetary Policy Shock in the VAR

The structural shocks underpinning the VAR are identified using the Cholesky factorization of the estimated covariance matrix of the one step ahead forecast errors. The observables enter the VAR in the following order: real GDP, real consumption, real investment, GDP deflator prices, real wages, labor productivity, federal funds rate, real profits, and the growth rate of M2. The monetary policy shock is assumed to be the one associated with the federal funds “equation” (i.e., the 7th one). The Cholesky identification scheme can thus be interpreted as a set of timing assumptions. The monetary policy shock cannot contemporaneously affect GDP, consumption, investment, the price level, wages, or labor productivity.

E.3 Local Projections and VARs, Redux

Here we put finite sample issues in LPs in the context a VAR-type estimator. Unless, the estimated VAR coincides with the data generating process, there will be an *asymptotic* bias and variance trade-off between the LP and VAR. In this section, we gauge the extent of this issue in the context of finite sample approximation for our data generating process. We compare the impulse response estimator $\hat{\theta}_{BCC}$ as well as the estimators constructed by iterating the $h = 0$ impact estimate using either an estimated VAR(1) or estimated VAR(4). We denote these estimates by $\hat{\theta}_{VAR(1)}$ and $\hat{\theta}_{VAR(4)}$, respectively. The VAR(4) represents the idealized case where the specification coincides with the data generating process. The VAR(1) stands in as a misspecified alternative that might be attractive to an investigator for its parsimony.

The left column of figure E.1 displays Monte Carlo estimates of $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid lines), $\mathbb{E}\hat{\theta}_{h,VAR(1)} - \theta_h$ (dotted lines), and $\mathbb{E}\hat{\theta}_{h,VAR(4)} - \theta_h$ (dash-dotted lines) for output, the price level, and the federal funds (rows) for $T = 100$. The VAR(4), the dash-dotted lines, being correctly specified, exhibits the least amount of bias in general. Note that while $\hat{\theta}_{LS}$ and the VAR(4) estimate are both unbiased asymptotically in this setting, the finite sample bias associated with the (even bias-corrected) LP estimator is substantially larger. While this is not a general statement—our theoretical results confirm that this difference cannot be signed in even a simple AR(1) setting—it is instructive that in empirically realistic settings this difference is sizable, as the LP bias and VAR bias are generated by the statistical considerations. Thus, the common practice of ignoring issues of finite sample bias in VARs may not be appropriate for LPs. The VAR(1), being misspecified, exhibits bias coming from both finite sample and asymptotic considerations. For real GDP and the price level, the bias associated with the VAR(1) is substantial, sometimes much worse than even the non-corrected LP (not shown), as their responses display more complex dynamics.


The right column of Figure E.1 displays the MSE ratios of $\hat{\theta}_{BCC}$ and $\hat{\theta}_{VAR(1)}$ relative to the $\hat{\theta}_{VAR(4)}$ for $T = 100$. The VAR(1) exhibits lower MSE than $\hat{\theta}_{BCC}$ for essentially all variables and horizons—indeed, it is typically preferred to the (correctly specified) VAR(4) estimator owing to finite sample considerations. One way to interpret this result, in the context of increasing popularity of LPs, is that investigators may care more about the bias than the variance of estimators—otherwise, we would see much use of misspecified VARs.

E.4 Additional Results


Figure E.1: Bias and MSE of $\hat{\theta}_{BCC}$ compared to VAR-based estimates under a CEE-type VAR Data Generating Process

Real GDP

(a) $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid), $\mathbb{E}\hat{\theta}_{h,VAR(1)} - \theta_h$ (dotted) and $\mathbb{E}\hat{\theta}_{h,VAR(4)} - \theta_h$ (dash-dotted)




(b) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,VAR(4)}$ (solid) and $\hat{\theta}_{h,VAR(1)}$ relative to $\hat{\theta}_{h,VAR(4)}$ (dotted)




Price Level

(c) $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid), $\mathbb{E}\hat{\theta}_{h,VAR(1)} - \theta_h$ (dotted) and $\mathbb{E}\hat{\theta}_{h,VAR(4)} - \theta_h$ (dash-dotted)




(d) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,VAR(4)}$ (solid) and $\hat{\theta}_{h,VAR(1)}$ relative to $\hat{\theta}_{h,VAR(4)}$ (dotted)




Federal Funds Rate

(e) $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid), $\mathbb{E}\hat{\theta}_{h,VAR(1)} - \theta_h$ (dotted) and $\mathbb{E}\hat{\theta}_{h,VAR(4)} - \theta_h$ (dash-dotted)



(f) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,VAR(4)}$ (solid) and $\hat{\theta}_{h,VAR(1)}$ relative to $\hat{\theta}_{h,VAR(4)}$ (dotted)




Note: the sub-figures show Monte Carlo means of $\hat{\theta}_{h,BCC} - \theta_h$ (solid lines), $\hat{\theta}_{h,VAR(1)} - \theta_h$ (dotted lines) and $\hat{\theta}_{h,VAR(4)} - \theta_h$ (dash-dotted lines) for $T = 100$. The sub-figures on the show the ratio of the MSEs of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,VAR(4)}$ (solid lines) and for $\hat{\theta}_{h,VAR(1)}$ relative to $\hat{\theta}_{h,VAR(4)}$ (dotted lines) for $T = 100$. We use 2,000 Monte Carlo simulations.


Figure E.2: Bias and MSE under a CEE-type VAR DGP (No controls)

Output

(a) $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid)




(b) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$




Price Level

(c) $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid)




(d) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$




Federal Funds Rate

(e) $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid)



(f) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$




Note: the sub-figures in the left column show the $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed lines) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid lines) for $T = 100$ (red) and $T = 200$ (green). The sub-figures in the right column show the ratio of the MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$ for $T = 100$ (red) and $T = 200$ (green).


Figure E.3: Bias and MSE under a CEE-type VAR DGP (Partial Controls)

Output

(a) $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid)




(b) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$




Price Level

(c) $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid)




(d) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$




Federal Funds Rate

(e) $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid)



(f) MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$



Note: the sub-figures in the left column show the $\mathbb{E}\hat{\theta}_{h,LS} - \theta_h$ (dashed lines) and $\mathbb{E}\hat{\theta}_{h,BCC} - \theta_h$ (solid lines) for $T = 100$ (red) and $T = 200$ (green). The sub-figures in the right column show the ratio of the MSE of $\hat{\theta}_{h,BCC}$ relative to $\hat{\theta}_{h,LS}$ for $T = 100$ (red) and $T = 200$ (green).

F Sources for meta study

- ABIAD, M. A., D. FURCERI, AND P. TOPALOVA (2015): *The macroeconomic effects of public investment: evidence from advanced economies*, 15-95, International Monetary Fund.
- ACEMOGLU, D., S. NAIDU, P. RESTREPO, AND J. A. ROBINSON (2019): “Democracy Does Cause Growth,” *Journal of Political Economy*, 127, 47–100.
- ADB, A. A., D. FURCERI, AND P. T. IMF (2016): “The macroeconomic effects of public investment: Evidence from advanced economies,” *Journal of Macroeconomics*, 50, 224–240.
- ADLER, G., M. R. A. DUVAL, D. FURCERI, K. SINEM, K. KOLOSKOVA, M. POPLAWSKI-RIBEIRO, ET AL. (2017): *Gone with the headwinds: Global productivity*, International Monetary Fund.
- ALESINA, A., O. BARBIERO, C. FAVERO, F. GIAVAZZI, AND M. PARADISI (2015a): “Austerity in 2009–13,” *Economic Policy*, 30, 383–437.
- (2017): “The effects of fiscal consolidations: Theory and evidence,” Tech. rep., National Bureau of Economic Research.
- ALESINA, A., C. FAVERO, AND F. GIAVAZZI (2015b): “The output effect of fiscal consolidation plans,” *Journal of International Economics*, 96, S19–S42.
- AMIOR, M. AND A. MANNING (2018): “The persistence of local joblessness,” *American Economic Review*, 108, 1942–70.
- ANDREASEN, M. M., J. FERNÁNDEZ-VILLAYERDE, AND J. F. RUBIO-RAMÍREZ (2017): “The pruned state-space system for non-linear DSGE models: Theory and empirical applications,” *The Review of Economic Studies*, 85, 1–49.
- ANGRIST, J. D., Ò. JORDÀ, AND G. M. KUERSTEINER (2018): “Semiparametric estimates of monetary policy effects: string theory revisited,” *Journal of Business & Economic Statistics*, 36, 371–387.
- AREZKI, R., V. A. RAMEY, AND L. SHENG (2017): “News shocks in open economies: Evidence from giant oil discoveries,” *The quarterly journal of economics*, 132, 103–155.

- ARREGUI, N., J. BENEŠ, I. KRZMAR, AND S. MITRA (2013): “Evaluating the net benefits of macroprudential policy: A cookbook,” .
- AUERBACH, A. J. AND Y. GORODNICHENKO (2012): “Fiscal multipliers in recession and expansion,” in *Fiscal policy after the financial crisis*, University of Chicago Press, 63–98.
- (2017): “Fiscal multipliers in Japan,” *Research in Economics*, 71, 411–421.
- BALL, L. M., D. FURCERI, M. D. LEIGH, AND M. P. LOUNGANI (2013): *The distributional effects of fiscal consolidation*, 13-151, International Monetary Fund.
- BANERJEE, R., M. B. DEVEREUX, AND G. LOMBARDO (2016): “Self-oriented monetary policy, global financial markets and excess volatility of international capital flows,” *Journal of International Money and Finance*, 68, 275–297.
- BANERJEE, R. N. AND H. MIO (2018): “The impact of liquidity regulation on banks,” *Journal of Financial intermediation*, 35, 30–44.
- BARNICHON, R. AND C. BROWNLEES (2019): “Impulse response estimation by smooth local projections,” *Review of Economics and Statistics*, 101, 522–530.
- BASHER, S. A., A. A. HAUG, AND P. SADORSKY (2012): “Oil prices, exchange rates and emerging stock markets,” *Energy Economics*, 34, 227–240.
- BAUMEISTER, C. AND L. KILIAN (2016): “Lower oil prices and the US economy: Is this time different?” *Brookings Papers on Economic Activity*, 2016, 287–357.
- BAYER, C., R. LÜTTICKE, L. PHAM-DAO, AND V. TJADEN (2019): “Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk,” *Econometrica*, 87, 255–290.
- BEN ZEEV, N. AND E. PAPPA (2017): “Chronicle of a war foretold: The macroeconomic effects of anticipated defence spending shocks,” *The Economic Journal*, 127, 1568–1597.
- BERTON, F., S. MOCETTI, A. F. PRESBITERO, AND M. RICHIARDI (2018): “Banks, firms, and jobs,” *The Review of Financial Studies*, 31, 2113–2156.
- BHATTACHARYA, U., N. GALPIN, R. RAY, AND X. YU (2009): “The role of the media in the internet IPO bubble,” *Journal of Financial and Quantitative Analysis*, 44, 657–682.

- BLUEDORN, J. C. AND C. BOWDLER (2011): “The open economy consequences of US monetary policy,” *Journal of International Money and Finance*, 30, 309–336.
- BORIO, C., M. DREHMANN, AND K. TSATSARONIS (2014): “Stress-testing macro stress testing: does it live up to expectations?” *Journal of Financial Stability*, 12, 3–15.
- BORIO, C. E., E. KHARROUBI, C. UPPER, AND F. ZAMPOLLI (2016): “Labour reallocation and productivity dynamics: financial causes, real consequences,” .
- BORN, B., G. J. MÜLLER, AND J. PFEIFER (2014): “Does austerity pay off?” *Review of Economics and Statistics*, 1–45.
- BOSS, M., G. FENZ, J. PANN, C. PUHR, M. SCHNEIDER, E. UBL, ET AL. (2009): “modeling credit risk through the Austrian business cycle: An update of the OeNb model,” *Financial Stability Report*, 17, 85–101.
- BRADY, R. R. (2011): “Measuring the diffusion of housing prices across space and over time,” *Journal of Applied Econometrics*, 26, 213–231.
- BUTT, N., R. CHURM, M. F. MCMAHON, A. MOROTZ, AND J. F. SCHANZ (2014): “QE and the bank lending channel in the United Kingdom,” .
- CAGGIANO, G., E. CASTELNUOVO, V. COLOMBO, AND G. NODARI (2015): “Estimating fiscal multipliers: News from a non-linear world,” *The Economic Journal*, 125, 746–776.
- CALDARA, D. AND E. HERBST (2019): “Monetary policy, real activity, and credit spreads: Evidence from bayesian proxy svars,” *American Economic Journal: Macroeconomics*, 11, 157–92.
- CASELLI, F. G. AND A. ROITMAN (2016): “Nonlinear exchange-rate pass-through in emerging markets,” *International Finance*.
- CHIȚU, L., B. EICHENGREEN, AND A. MEHL (2014): “When did the dollar overtake sterling as the leading international currency? Evidence from the bond markets,” *Journal of Development Economics*, 111, 225–245.
- CHODOROW-REICH, G. AND L. KARABARBOUNIS (2016): “The limited macroeconomic effects of unemployment benefit extensions,” Tech. rep., National Bureau of Economic Research.

- CHONG, Y., Ò. JORDÀ, AND A. M. TAYLOR (2012): “The Harrod–Balassa–Samuelson Hypothesis: Real Exchange Rates And Their Long-Run Equilibrium,” *International Economic Review*, 53, 609–634.
- CLOYNE, J. AND P. HÜRTGEN (2016): “The macroeconomic effects of monetary policy: a new measure for the United Kingdom,” *American Economic Journal: Macroeconomics*, 8, 75–102.
- COIBION, O. AND Y. GORODNICHENKO (2012): “What can survey forecasts tell us about information rigidities?” *Journal of Political Economy*, 120, 116–159.
- COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): “Innocent Bystanders? Monetary Policy and Inequality,” *Journal of Monetary Economics*, 88, 70–89.
- DABLA-NORRIS, M. E., M. S. GUO, M. V. HAKSAR, M. KIM, M. K. KOCHHAR, K. WISEMAN, AND A. ZDZIENICKA (2015): *The new normal: A sector-level perspective on productivity trends in advanced economies*, International Monetary Fund.
- DE COS, P. H. AND E. MORAL-BENITO (2016): “Fiscal multipliers in turbulent times: the case of Spain,” *Empirical Economics*, 50, 1589–1625.
- DESSAINT, O. AND A. MATRAY (2017): “Do managers overreact to salient risks? Evidence from hurricane strikes,” *Journal of Financial Economics*, 126, 97–121.
- DUPOR, B., J. HAN, AND Y.-C. TSAI (2009): “What do technology shocks tell us about the New Keynesian paradigm?” *Journal of Monetary Economics*, 56, 560–569.
- EICHLER, M. (2013): “Causal inference with multiple time series: principles and problems,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371, 20110613.
- FAVARA, G. AND J. IMBS (2015): “Credit supply and the price of housing,” *American Economic Review*, 105, 958–92.
- FAZZARI, S. M., J. MORLEY, AND I. PANOVSKA (2015): “State-dependent effects of fiscal policy,” *Studies in Nonlinear Dynamics & Econometrics*, 19, 285–315.
- FRANCIS, N., M. T. OWYANG, J. E. ROUSH, AND R. DICECIO (2014): “A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks,” *Review of Economics and Statistics*, 96, 638–647.

- FUNKE, M., M. SCHULARICK, AND C. TREBESCH (2016): “Going to extremes: Politics after financial crises, 1870–2014,” *European Economic Review*, 88, 227–260.
- FURCERI, D., M. L. E. BERNAL-VERDUGO, AND M. D. M. GUILLAUME (2012a): *Crises, labor market policy, and unemployment*, 12-65, International Monetary Fund.
- FURCERI, D., S. GUICHARD, AND E. RUSTICELLI (2012b): “The effect of episodes of large capital inflows on domestic credit,” *The North American Journal of Economics and Finance*, 23, 325–344.
- FURCERI, D., P. LOUNGANI, AND A. ZDZIENICKA (2018): “The effects of monetary policy shocks on inequality,” *Journal of International Money and Finance*, 85, 168–186.
- FURCERI, D. AND A. ZDZIENICKA (2011): “How costly are debt crises?” *IMF working papers*, 1–29.
- (2012): “How costly are debt crises?” *Journal of International Money and Finance*, 31, 726–742.
- GAL, P. N. AND A. HIJZEN (2016): *The short-term impact of product market reforms: A cross-country firm-level analysis*, International Monetary Fund.
- GERTLER, M. AND S. GILCHRIST (2018): “What happened: Financial factors in the great recession,” *Journal of Economic Perspectives*, 32, 3–30.
- HALL, A. R., A. INOUE, J. M. NASON, AND B. ROSSI (2012): “Information criteria for impulse response function matching estimation of DSGE models,” *Journal of Econometrics*, 170, 499–518.
- HAMILTON, J. D. (2011): “Nonlinearities and the macroeconomic effects of oil prices,” *Macroeconomic dynamics*, 15, 364–378.
- HAUTSCH, N. AND R. HUANG (2012): “The market impact of a limit order,” *Journal of Economic Dynamics and Control*, 36, 501–522.
- HOLLO, D., M. KREMER, AND M. LO DUCA (2012): “CISS-a composite indicator of systemic stress in the financial system,” .
- JORDÀ, Ò. (2009): “Simultaneous confidence regions for impulse responses,” *The Review of Economics and Statistics*, 91, 629–647.

- JORDÀ, Ò. AND M. MARCELLINO (2010): “Path forecast evaluation,” *Journal of Applied Econometrics*, 25, 635–662.
- JORDÀ, Ò., M. SCHULARICK, AND A. M. TAYLOR (2013): “When credit bites back,” *Journal of Money, Credit and Banking*, 45, 3–28.
- (2015a): “Betting the House,” *Journal of International Economics*, 96, S2–S18.
- (2015b): “Leveraged bubbles,” *Journal of Monetary Economics*, 76, S1–S20.
- (2016): “Sovereigns versus banks: credit, crises, and consequences,” *Journal of the European Economic Association*, 14, 45–79.
- JORDÀ, Ò. AND A. M. TAYLOR (2016): “The time for austerity: estimating the average treatment effect of fiscal policy,” *The Economic Journal*, 126, 219–255.
- KILIAN, L. AND Y. J. KIM (2011): “How Reliable are Local Projection Estimators of Impulse Responses?” *Review of Economics and Statistics*, 93, 1460–1466.
- KILIAN, L. AND R. J. VIGFUSSON (2011): “Nonlinearities in the oil price–output relationship,” *Macroeconomic Dynamics*, 15, 337–363.
- (2017): “The role of oil price shocks in causing US recessions,” *Journal of Money, Credit and Banking*, 49, 1747–1776.
- KRAAY, A. (2014): “Government spending multipliers in developing countries: evidence from lending by official creditors,” *American Economic Journal: Macroeconomics*, 6, 170–208.
- KRISHNAMURTHY, A. AND T. MUIR (2017): “How credit cycles across a financial crisis,” Tech. rep., National Bureau of Economic Research.
- LEDUC, S. AND D. WILSON (2013): “Roads to prosperity or bridges to nowhere? Theory and evidence on the impact of public infrastructure investment,” *NBER Macroeconomics Annual*, 27, 89–142.
- LEIGH, M. D., W. LIAN, M. POPLAWSKI-RIBEIRO, R. SZYMANSKI, V. TSYRENNIKOV, AND H. YANG (2017): *Exchange rates and trade: A disconnect?*, International Monetary Fund.

- LISTORTI, G. AND R. ESPOSTI (2012): “Horizontal price transmission in agricultural markets: fundamental concepts and open empirical issues,” *Bio-based and Applied Economics Journal*, 1, 81–108.
- LUETTICKE, R. (2018): “Transmission of monetary policy with heterogeneity in household portfolios,” .
- MENKHOFF, L., L. SARNO, M. SCHMELING, AND A. SCHRIMPF (2016): “Currency value,” *The Review of Financial Studies*, 30, 416–441.
- MERTENS, K. AND J. L. MONTIEL OLEA (2018): “Marginal tax rates and income: New time series evidence,” *The Quarterly Journal of Economics*, 133, 1803–1884.
- MIAN, A., A. SUFI, AND E. VERNER (2017): “Household debt and business cycles worldwide,” *The Quarterly Journal of Economics*, 132, 1755–1817.
- MIYAMOTO, W., T. L. NGUYEN, AND D. SERGEYEV (2018): “Government spending multipliers under the zero lower bound: Evidence from Japan,” *American Economic Journal: Macroeconomics*, 10, 247–77.
- MORETTI, E. AND D. J. WILSON (2017): “The effect of state taxes on the geographical location of top earners: evidence from star scientists,” *American Economic Review*, 107, 1858–1903.
- NAKAMURA, E. AND J. STEINSSON (2018): “Identification in Macroeconomics,” *Journal of Economic Perspectives*, 32, 59–86.
- OTTONELLO, P. AND T. WINBERRY (2018): “Financial heterogeneity and the investment channel of monetary policy,” Tech. rep., National Bureau of Economic Research.
- OWYANG, M. T., V. A. RAMEY, AND S. ZUBAIRY (2013): “Are government spending multipliers greater during periods of slack? Evidence from twentieth-century historical data,” *American Economic Review*, 103, 129–34.
- RAMEY, V. A. (2016): “Macroeconomic shocks and their propagation,” in *Handbook of macroeconomics*, Elsevier, vol. 2, 71–162.
- RAMEY, V. A. AND S. ZUBAIRY (2018): “Government spending multipliers in good times and in bad: evidence from US historical data,” *Journal of Political Economy*, 126, 850–901.

- RIERA-CRITCHTON, D., C. A. VEGH, AND G. VULETIN (2015): “Procyclical and countercyclical fiscal multipliers: Evidence from OECD countries,” *Journal of International Money and Finance*, 52, 15–31.
- ROMER, C. D. AND D. H. ROMER (2015): “New evidence on the impact of financial crises in advanced countries,” Tech. rep., National Bureau of Economic Research.
- (2017): “New evidence on the aftermath of financial crises in advanced countries,” *American Economic Review*, 107, 3072–3118.
- SALOMONS, A. ET AL. (2018): “Is automation labor-displacing? Productivity growth, employment, and the labor share,” Tech. rep., National Bureau of Economic Research.
- SANTORO, E., I. PETRELLA, D. PFAJFAR, AND E. GAFFEO (2014): “Loss aversion and the asymmetric transmission of monetary policy,” *Journal of Monetary Economics*, 68, 19–36.
- SCHALLER, J. (2013): “For richer, if not for poorer? Marriage and divorce over the business cycle,” *Journal of Population Economics*, 26, 1007–1033.
- STOCK, J. H. AND M. W. WATSON (2018): “Identification and estimation of dynamic causal effects in macroeconomics using external instruments,” *The Economic Journal*, 128, 917–948.
- STOYANOV, A. AND N. ZUBANOV (2012): “Productivity spillovers across firms through worker mobility,” *American Economic Journal: Applied Economics*, 4, 168–98.
- TAYLOR, A. M. (2015): “Credit, financial stability, and the macroeconomy,” *Annu. Rev. Econ.*, 7, 309–339.
- TENREYRO, S. AND G. THWAITES (2016): “Pushing on a String: US Monetary Policy is Less Powerful in Recessions,” *American Economic Journal: Macroeconomics*, 8, 43–74.
- TEULINGS, C. N. AND N. ZUBANOV (2014): “Is economic recovery a myth? Robust estimation of impulse responses,” *Journal of Applied Econometrics*, 29, 497–514.
- YANG, J., H. GUO, AND Z. WANG (2006): “International transmission of inflation among G-7 countries: A data-determined VAR analysis,” *Journal of Banking & Finance*, 30, 2681–2700.
- ZIDAR, O. (2019): “Tax cuts for whom? Heterogeneous effects of income tax changes on growth and employment,” *Journal of Political Economy*, 127, 1437–1472.