

Bias in Local Projections

Edward P. Herbst and Benjamin K. Johannsen*

December 31, 2020

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. Consequently, LP point estimates can be severely biased. We derive simple expressions for this bias and propose a way to bias-correct LPs. Small sample bias can also lead autocorrelation-robust standard errors to dramatically understate sampling uncertainty. We argue they should be avoided in LPs like the ones we study. Using identified monetary policy shocks, we demonstrate that the bias in point estimates can be economically meaningful and the bias in standard errors can affect inference.

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

*edward.p.herbst@frb.gov; benjamin.k.johannsen@frb.gov. This does not reflect the views of the Federal Reserve Board or the Federal Reserve System. We thank 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.

find that, over the past 15 years, LPs have been applied in a variety of settings that are notably different than the setting studied in Jordà (2005). In particular, we find that sample sizes in the time 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.

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. Our Monte Carlo simulations use simple, linear data generating processes. 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 simple scenarios. Notably, we show that the bias in LPs persists even when the shock on the right-hand-side of the regressions in our Monte Carlo analysis is *iid*, as is often the case when researchers have access to a time series of identified structural shocks.

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), Rilstone et al. (1996), Anatolyev (2005), and Bao and Ullah (2007). We show that the bias of the LP estimator at horizon h is a function—specifically, a weighted sum—of the (population) impulse response function at other horizons. As a result, if 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.

We use the higher-order expansion of the LP estimator to develop a simple bias correction that can be added to the ordinary least squared (OLS) estimator. In Monte Carlo simulations, on average our bias-corrected estimators are markedly closer to the true values of the impulse responses.² In addition, using coverage probabilities as a metric, we show that our estimator

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, in some settings researchers

performs relatively well when compared to the OLS estimator.

We extend our analysis to settings using panel data and show that—when using fixed effects—the bias we document persists. We provide formulas for bias correcting the OLS estimator in panel settings with and without controls. Additionally, we show that increasing the number of entities in the panel will not eliminate the bias.

Finally, we also study the downward bias in the standard errors of LP estimators. Our analysis of standard errors is related to recent work by Olea and Plagborg-Møller (2020), who suggest that researchers using LPs should use heteroskedasticity-consistent, but not autocorrelation-robust, standard errors.³ Our focus on finite sample issues leaves us at similar conclusions, but for different reasons.

We show that, like the LPs themselves, commonly-used standard errors for LPs that are heteroskedasticity- and autocorrelation-robust (HAR) are also severely downward biased when sample sizes are similar to those commonly found in the literature. When the regression scores—the regression residuals times the regressors—are autocorrelated, HAR standard errors are necessary for valid inference. Popular HAR estimators, such as the Newey and West (1987) estimator, rely on estimators of the autocorrelation of the regression score. Like the LP point estimators, in finite samples, the estimators of these autocorrelations will be biased. In most empirically-realistic settings, the bias will be downwards. This downward bias is large enough that, in LP settings like the ones we study, estimates of the autocorrelation are, on average, negative in small samples, even when the true autocorrelation is zero or mildly positive!

The possibility of the bias leading to an incorrectly signed estimate of autocorrelations suggests researchers may prefer standard errors that are heteroskedasticity-consistent, but not autocorrelation-robust, such as Huber-White standard errors. In fact, in our empirical examples, switching from Newey-West to Huber-White estimators generally increases the estimates of standard errors, sometimes dramatically so. In the case that researchers need to use autocorrelation-robust standard errors, our small sample analysis suggests that they should be aware that these standard errors may be dramatically biased.

Additionally, when researchers have the shocks in hand, we argue that under the null may prefer methods, such as VARs, that estimate the same impulse responses as LPs (see Plagborg-Møller and Wolf (2019)) and have well-understood, effective methods for bias correction (see Kilian (1998)).

³Olea and Plagborg-Møller (2020) specifically advocate the use of *lag-augmented* LPs. We analyze the effects of lag augmentation in our setup in Appendix G.

hypothesis that is typically of interest, HAR standard errors are not even necessary for valid inference. We provide an analytical characterization of the true autocorrelation function of the regression score in a simple setting. In this case, using HAR standard errors is not conservative; in fact, it will typically understate uncertainty relative to a standard error calculated without attempting to take autocorrelation into account.

We analyze bias in three examples drawn from the empirical monetary economics literature. We show that the bias in point estimates can be economically meaningful. In these examples, consistent with our simulations and analytical results, the use of HAR estimators typically leads to smaller standard error estimates than the use of the Huber-White estimator; in many cases this difference is large enough to affect inference. This is true even when the estimated bias in the point estimate is small.

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 OLS estimate. By contrast, our paper considers the case when a time series of identified shocks is available to the researcher, so right-hand-side variables are *iid*. Our analysis relies on higher-order expansions of the OLS estimator, which illustrate the reasons that the LP estimator is biased and provide a natural bias correction without bootstrapping. In addition, we extend the analysis to panel data settings, which are common settings for LPs in practice. In Appendix F, we compare the bias correction from the block bootstrap used in Kilian and Kim (2011) to our proposed bias correction and use insights from our higher-order expansion of the OLS estimator to explain why the bootstrap performs relatively poorly unless block lengths are relatively long. More generally, our paper is related to work on bias in least-squares estimators of autocorrelation (such as Kendall (1954) and Shaman and Stine (1988)), in dynamic panel data settings (such as Nickell (1981) and Hahn and Kuersteiner (2002)), and in generalized method of moments systems (Rilstone et al. (1996), Anatolyev (2005), and Bao and Ullah (2007)). We extend this work to our LP setting.

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.⁴ Google Scholar’s relevance ranking weights 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 LP).⁵

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 mentioned in the introduction and discussed later, 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 extremely 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 up.

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, and so the time series typically have short 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 $T = 300$

⁴We conducted this search in October 2019. See Appendix I for the list of citations.

⁵If a paper appeared as both a working paper and a published paper, we excluded the working paper version from our analysis.

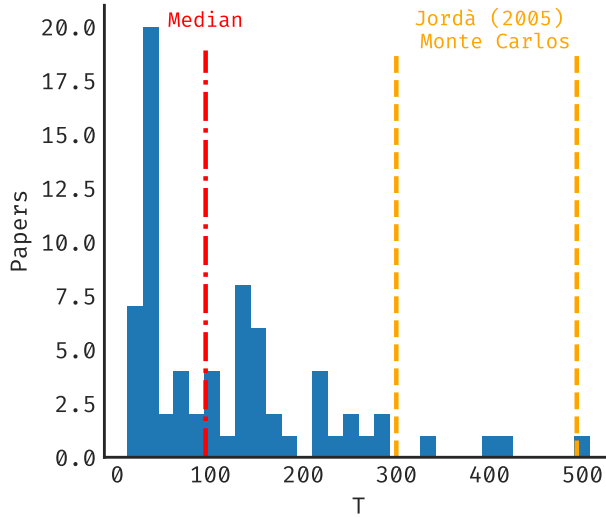


Figure 1: T is small in the literature using LPs.

and $T = 496$ (the orange, dashed lines in Figure 1). Less than 6 percent of the surveyed studies use sample sizes at least that large. Of course, it is difficult to fault Jordà (2005) for not anticipating how researchers would subsequently apply LP methods. 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 completely 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 demonstrate that LPs can be severely biased with sample sizes like those documented in Section 2. We explore this bias using both Monte Carlo evidence and a new analytic approximation of the bias. The analytic approximation yields insights into the bias associated with LPs at different horizons and suggests a correction of the bias.

3.1 Bias in LPs using an AR(1) example

To demonstrate that LPs can be severely biased in small samples, we first consider Monte Carlo analysis using an AR(1) data generating process. Our objective is to study the accuracy of estimated impulse responses via LPs for various sample sizes. For a given T , we simulate $N_{mc} = 10,000$ time series, $\{y_t\}_{t=1}^T$, for the data generating process:

$$y_t = \rho y_{t-1} + \varepsilon_t + \nu_t. \quad (1)$$

Here, ε_t and ν_t are *iid* standard normal random variables. In these Monte Carlo simulations, to be consistent with the high persistence of macroeconomic data, we set $\rho = 0.95$.⁶ We use the AR(1) time series model because LPs were designed to capture the dynamics of a wide range of data generating processes and one would hope that they would perform well in the simplest examples.

We assume that the researcher does not know the true data generating process but is otherwise in a near-ideal setting for estimating the impulse response function of y_t using LPs. The researcher observes $\{y_t, \varepsilon_t\}_{t=1}^T$; that is, the researcher directly observes the shock ε_t which is independent over time and uncorrelated with past values of y_t . In addition, the researcher may control for other variables, denoted by the vector c_t . When we include such controls in the Monte Carlo exercise, we assume $c_t = y_{t-1}$. We stress that our regressions with controls are ideal in the sense that no useful additional information from earlier periods could be added to the regressors and we include the correct number of lags of y_t as controls.

The LP model is the set of regression models indexed by the impulse response horizon h ,

$$y_{t+h} = \alpha_h + \beta_h' x_t + u_{t,h}, \quad h = 0, \dots, H. \quad (2)$$

where $x_t \equiv [\varepsilon_t, c_t]'$.⁷ Thus, 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 $h = 0, \dots, H$. As in the empirical macroeconomics literature, we estimate each β_h using least squares. We denote the estimator of the β_h by $\hat{\beta}_{h,LS}$ and the estimator of the impulse response by $\hat{\theta}_{LS}$.

Using Monte Carlo simulations, we can compute, for any T , the finite sample expectation of the least-squares estimator, $\mathbb{E}[\hat{\theta}_{LS}]$. Figure 2 displays the expectation for the LP estimators

⁶In Appendix C, we provide results for alternative values of ρ . In each simulation, we initialize y_0 at a draw from the unconditional distribution of y_t . The variances of ε_t and ν_t are set to $\sigma_\varepsilon^2 = \sigma_\nu^2 = 1$.

⁷When we do not include controls, $x_t = \varepsilon_t$.

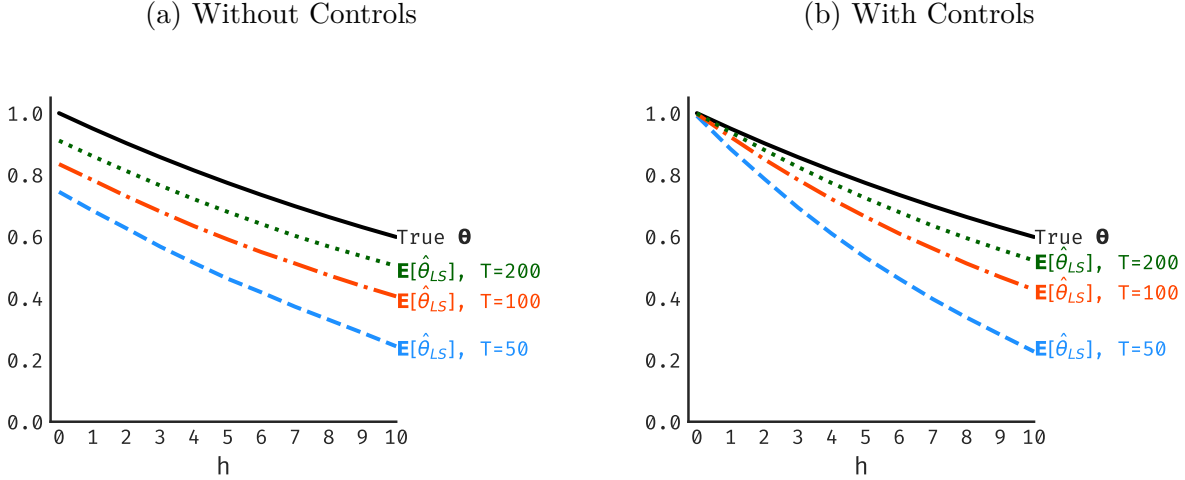


Figure 2: LP estimators are biased in empirically-relevant samples when y_t is an AR(1) with $\rho = 0.95$.

with and without controls for the AR(1) data generating process with $T \in \{50, 100, 200\}$. Recall that about half of the surveyed literature uses T less than 100.

When controls are not included in the LP (the left panel), the estimator is biased even at short horizons. This is true even for moderately long time series—i.e., $T = 200$. As the horizon of the impulse response increases, the bias becomes worse. When y_{t-1} is included as a control (the right panel), the bias diminishes substantially at short horizons. Intuitively, adding controls makes the least-squares error terms less autocorrelated with at short horizons. However, even for the impulse response only 10 periods ahead (2.5 years with quarterly data), the controls alleviate only a small fraction of the bias. The reason that controls are less effective at reducing the bias as h increases is that they are less effective at forecasting y_{t+h} . Note also that, at these longer horizons, the overlapping nature of the left-hand-side variables in the LP implies that the error terms are autocorrelated.⁸

⁸Following Jordà (2005), for each regression of horizon h , researchers typically use all available data, meaning that the regression error term is autocorrelated for at least $h - 1$ periods. With autocorrelated regression error terms, the generalized least-squares estimator asymptotically performs better than the ordinary least-squares estimator. However, researchers typically use the ordinary least-squares estimator because of the small-sample shortcomings of the feasible generalized least-squares estimator. That said, recent work by Lusompa (2019) suggests that well-behaved small sample GLS estimators may be obtained for LPs.

3.2 Approximating the small sample bias in LPs

In this subsection we derive an analytic approximation to the bias of the least-squares estimator, $\widehat{\theta}_{LS}$, given by $\mathbb{E}[\widehat{\theta}_{LS}] - \theta$. The expressions we derive are, to the best of our knowledge, new to the literature and highlight the interdependence of the bias in LPs at different horizons. In order to illustrate the point clearly (and to avoid tedious matrix algebra), we first focus on LPs without controls. We generalize our analytic approximation to the case when controls are included in the LP, and the intuition for the bias is similar (we provide the associated derivations in our Appendix A).

In all of our derivations, we make the following assumptions. Let $w_t = [y_t, \varepsilon_t, c_t]'$.

Assumption 1. *The data $\{w_t\}$ is stationary and ergodic. The demeaned series $\{w_t - \mu_w\}$, where $\mu_w = E[w_t]$, is purely non-deterministic with absolutely summable Wold decomposition coefficients and $\mathbb{E}[|w_t w_t'|^2] < \infty$.*

Under Assumption 1, the Wold representation can be inverted and the data has a VAR(∞) representation. If $\{w_t\}$ is jointly Gaussian, the data satisfies the Assumption 1. The next assumption explicitly characterizes the relationship between y_t and ε_t as well the conditional moments of ε_t .

Assumption 2. *The time series y_t can be decomposed as*

$$y_t = \sum_{j=0}^{\infty} \theta_j \varepsilon_{t-j} + \bar{y}_t$$

where \bar{y}_t is independent of ε_{t-j} for $j = -\infty, \dots, \infty$. Moreover,

$$\begin{aligned} \mathbb{E}[\varepsilon_t | \{w_\tau\}_{\tau < t}, \{\varepsilon_\tau\}_{t < \tau}] &= E[\varepsilon_t] = \mu_\varepsilon \\ \text{and } \mathbb{E}[(\varepsilon_t - \mu_\varepsilon)^2 | \{w_\tau\}_{\tau < t}, \{\varepsilon_\tau\}_{t < \tau}] &= \mathbb{E}[(\varepsilon_t - \mu_\varepsilon)^2] = \sigma_\varepsilon^2 > 0. \end{aligned}$$

The first part of this assumption characterizes y_t as the sum of a casual linear filter in $\{\varepsilon_{t-j}\}_{j=0}^{\infty}$ and an additional component which is independent of all ε_j . The filter coefficients are the impulse response coefficients of interest. This is a stark assumption that is meant to represent an idealized case where “direct causal inference”—as in Nakamura and Steinsson (2018)—is possible. A jointly Gaussian time series trivially satisfies this portion of the assumption. The second part of this assumption precludes first and second moment dependence in ε_t . The idealized case where the researcher has *iid* shocks in hand trivially satisfies

the second part of Assumption 2. Such a stark assumption is consistent with the emerging practice of constructing such shocks with desirable statistical properties. We stress that a violation of Assumption 2 does not eliminate the small sample bias discussed here, it merely makes it difficult to characterize analytically. This assumption allows us to obtain sharp analytical results, a necessary first step in understanding finite sample issues in LPs.

Recall that in our framework, the researcher observes $\{y_t, \varepsilon_t\}_{t=1}^T$. For a given h , the ordinary least-squares estimator of θ_h can be written as

$$\begin{aligned}\widehat{\theta}_{h,LS} &= \frac{\frac{1}{T-h} \sum_{t=1}^{T-h} \varepsilon_t y_{t+h} - \frac{1}{(T-h)^2} \left(\sum_{t=1}^{T-h} \varepsilon_t \right) \left(\sum_{t=1}^{T-h} y_{t+h} \right)}{\frac{1}{T-h} \sum_{t=1}^{T-h} \varepsilon_t^2 - \frac{1}{(T-h)^2} \left(\sum_{t=1}^{T-h} \varepsilon_t \right)^2} \\ &= \frac{\widehat{\text{cov}}[\varepsilon_t, y_{t+h}]}{\widehat{\text{var}}[\varepsilon_t]}.\end{aligned}\tag{3}$$

Here, $\widehat{\text{cov}}$ and $\widehat{\text{var}}$ are the sample covariance and variance, respectively.⁹ Equation (3) makes it clear that θ_h in population is the scaled covariance between y_{t+h} and ε_t .

Analytical Result 1 (Expression for the bias in an LP without controls). *Under Assumptions 1 and 2, the bias of the LP estimator in equation (3) is given by*

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

A proof of this claim is given in Appendix A, but we provide some discussion here. To derive an expression for the approximate bias of $\widehat{\theta}_{h,LS}$, we need to compute $\mathbb{E} \left[\widehat{\theta}_{h,LS} \right]$. In such a simple example, one can calculate higher-order expansions of equation (3). This is a widely adopted approach to these kinds of problems. We apply the methodology of Bao and Ullah (2007), which does not require, for example, that the shocks be normally distributed. Finally, note that equation (4) gives an approximate expression for the bias of $\widehat{\theta}_{h,LS}$ that consists of only the impulse response coefficients at different horizons.¹⁰

To gain intuition for equation (4), it is useful to consider the case when ε_t is symmetric around the origin.¹¹ In this case,

$$\mathbb{E} \left[\widehat{\theta}_{h,LS} \right] = \frac{\mathbb{E}[\widehat{\text{cov}}[\varepsilon_t, y_{t+h}]]}{E[\widehat{\text{var}}[\varepsilon_t]]} + O(T^{-3/2}).\tag{5}$$

⁹Note that for the variance to be consistent with the least-squares estimator, we use only the $T-h$ observations of ε_t .

¹⁰Given our maintained assumption that ε_t is uncorrelated with past values of y_t , $\theta_{-j} = 0$ for $j > 0$.

¹¹In this case, the derivation of our expression for the bias of $\widehat{\theta}_{h,LS}$ is similar to the derivation in Kendall (1954) for the well-known bias of estimators of autocorrelation—see also Shaman and Stine (1988).

The numerator of equation (5) is given by

$$\begin{aligned} \mathbb{E}[\widehat{\text{cov}}[y_{t+h}, \varepsilon_t]] &= \left(1 - \frac{1}{T-h}\right) \text{cov}[y_{t+h}, \varepsilon_t] \\ &\quad - \frac{1}{T-h} \sum_{j=1}^{T-h-1} \left(1 - \frac{j}{T-h}\right) (\text{cov}[y_{t+h+j}, \varepsilon_t] + \text{cov}[y_{t+h-j}, \varepsilon_t]). \end{aligned} \quad (6)$$

Here, cov is the true covariance. The summation in equation (6) comes from the plug-in estimator for the mean in $\widehat{\text{cov}}$. The denominator of equation (5) is

$$\mathbb{E}[\widehat{\text{var}}[\varepsilon_t]] = \left(1 - \frac{1}{T-h}\right) \text{var}[\varepsilon_t]. \quad (7)$$

Here, var is the true variance. Noting that $\theta_{h+j} = \text{cov}[\varepsilon_t, y_{t+h+j}] / \text{var}[\varepsilon_t]$, we obtain equation (4).¹² From this derivation, it is clear that the bias arises because of the need to estimate the means of the variables in the OLS calculations.

Several more remarks regarding equation (4) are in order. First, as expected, the bias is a decreasing function of T ; for fixed h , the least-squares estimator is consistent. Second, the bias of $\widehat{\theta}_{h,LS}$ is a function of the impulse response coefficients at *all* other horizons. Intuitively, the data generating process affects the bias of OLS estimators at similar horizons in similar ways. The interdependence of LP estimates across h highlights that, in finite samples, LPs are not “local.” Third, the contribution of the horizon $h + j$ impulse response coefficients to the bias in the least-squares estimate of the h impulse response coefficient decreases only at linear rate as j increases or decreases. In practice, this means that the bias in the portion of the impulse response of interest—typically, say, the first 20 periods in quarterly macroeconomic applications—can be meaningfully affected by the impulse response at much longer horizons. This is especially true for extremely persistent time series (like many macroeconomic series).

We next consider the case when a researcher includes controls in the LP. We maintain the following assumptions.

Assumption 3. $\alpha_h + x_t' \beta_h$ is an optimal linear forecast for y_{t+h} .

The assumption that the LP produces an optimal linear forecast implies that, using the true parameters, forecast errors are $\text{MA}(h + 1)$ processes.

¹²In this derivation, the expression for the bias would be divided by $T - h - 1$ rather than $T - h$, but that difference is absorbed in the term $O(T^{-3/2})$.

Assumption 4. *The matrix $\mathbb{E}[x_t x_t']$ is full rank.*

This assumption ensures our estimator is consistent and that our setup satisfies the assumptions of Rilstone et al. (1996). With these additional assumptions, we can state our next analytical result.

Analytical Result 2 (Expression for the bias in an LP with controls). *Under Assumptions 1-4, the bias for the LP with controls is given by*

$$\mathbb{E} \left[\widehat{\theta}_{h,LS} \right] - \theta_h = -\frac{1}{T-h} \sum_{j=1}^h \left(1 - \frac{j}{T-h} \right) (1 + \text{tr} \{ \Sigma_{c,0}^{-1} \Sigma_{c,j} \}) \theta_{h-j} + O(T^{-3/2}), \quad (8)$$

where $\Sigma_{c,j} \equiv E[(c_{t-j} - \mu_c)(c_t - \mu_c)']$.

The claim again relies on the results of Bao and Ullah (2007). A detailed derivation is in Appendix A. Several remarks regarding equation (8) are in order. First, as in the case without controls, the bias is a decreasing function of T ; for fixed h , the least-squares estimator is consistent. Second, the bias of $\widehat{\theta}_{h,LS}$ is a function of the impulse response coefficients at *all horizons up to h* . Intuitively, the data generating process affects the bias of OLS estimators at similar horizons in similar ways, however controlling for past data truncates the terms that are important for the bias by truncating the correlation in the regression errors. Third, the contribution of the horizon $h-j$ impulse response coefficients to the bias in the least-squares estimate of the h impulse response coefficient is scaled by $1 + \text{tr} \{ \Sigma_{c,0}^{-1} \Sigma_{c,j} \}$. As a result, when controls are persistent or when unneeded persistent controls are added to the LP, the bias increases.

Using our AR(1) example, figure 3 shows $\mathbb{E} \left[\widehat{\theta}_{LS} \right]$ calculated using the approximation in equation (5), assuming that the true values of θ_h are known. The figure also shows the exact finite sample value from Monte Carlo simulations. Notably, the approximation works quite well in population. For the no-controls case, the analytic approximation is nearly exact for $T \in \{50, 100, 200\}$. With controls, the analytic approximation to the impulse response is somewhat above the true finite-sample expectation, though it still captures most of the bias associated with the least-squares estimator.

3.3 A bias-corrected estimator

Equations 4 and 8 lend themselves to constructing bias corrected estimators for θ_h using plug-in estimators for θ_j and $\Sigma_{c,j}$. In the case of no controls, the bias depends on the values




Figure 3: The bias approximation is accurate in our LPs.

of θ_j for all $j \neq h$ and $|j| \leq T$. Given the inability to estimate all of these parameters with a sample size of T , in practice a researcher could truncate the horizon of the coefficients used in the bias correction. Without theory on how to optimally pick the maximum horizon used in the bias correction, H , we set H to 20, 25, and 50 for T equal to 50, 100, and 200, respectively. In the case when controls are included in the LP, all of the needed values of θ_j and $\Sigma_{c,j}$ are easily computed.

Notably, the researcher could bias correct the coefficients using the OLS estimates of θ_j . When we construct the bias-corrected estimator in this way, we denote the estimator as $\hat{\theta}_{BC,h}$. Alternatively, the researcher could iterate the bias correction on all values of θ_j . When we construct the bias-corrected estimator in this way, we denote the estimator as $\hat{\theta}_{BCC,h}$.

For the case of an LP without controls and an LP with controls, respectively, figures 4 and 5 show the average value of $\hat{\theta}_{LS}$, $\hat{\theta}_{BC}$, and $\hat{\theta}_{BCC}$ over our Monte Carlo simulations when y_t follows an AR(1) with $\rho = 0.95$. Clearly, our bias correction does not completely correct for the bias in $\hat{\theta}_{LS}$ in either case, indicating that our bias-corrected estimator is not a panacea for bias in LPs. Nevertheless, $\hat{\theta}_{BC}$ and $\hat{\theta}_{BCC}$ are markedly closer than $\hat{\theta}_{LS}$ to θ on average.

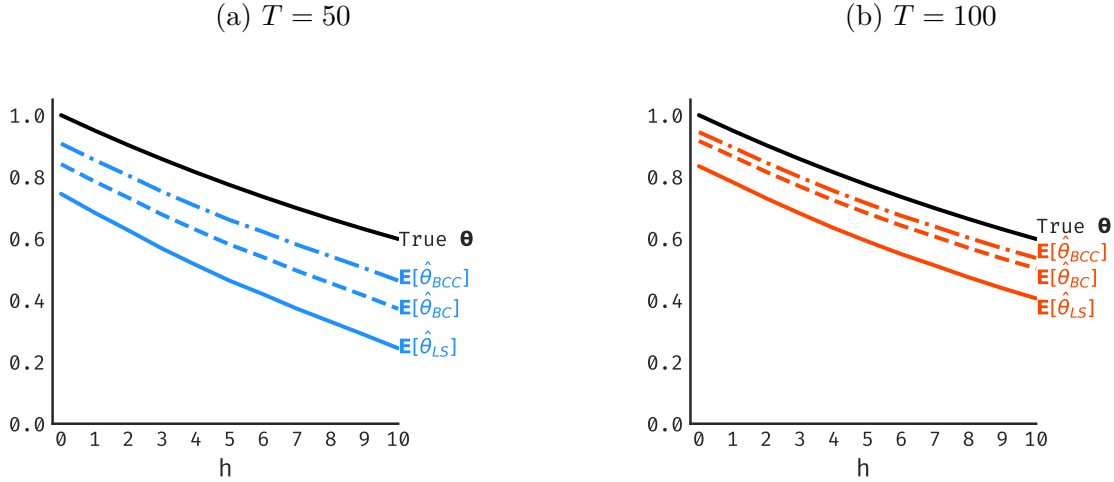


Figure 4: $\hat{\theta}_{BC}$ and $\hat{\theta}_{BCC}$ are closer than $\hat{\theta}_{LS}$ to θ , on average, in our LPs without controls when y_t is an AR(1) with $\rho = 0.95$.

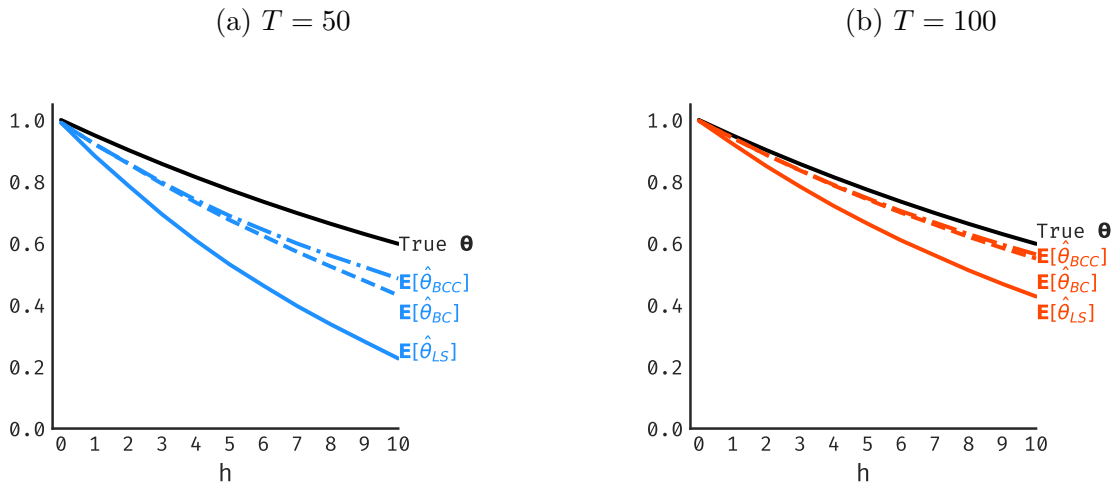


Figure 5: $\hat{\theta}_{BC}$ and $\hat{\theta}_{BCC}$ are closer than $\hat{\theta}_{LS}$ to θ , on average, in our LPs with controls when y_t is an AR(1) with $\rho = 0.95$.

4 Extension to panel data

In this section, we demonstrate that LPs can be severely biased with sample sizes in the time dimension commonly found in the empirical macroeconomic literature even when researchers have access to a large cross-section (i.e. panel data). Of course, parameter bias in dynamic panel data models has been studied since at least Nickell (1981). We illustrate the bias in LPs using the expansion of the OLS estimator that is similar to one in the previous section.

As before, the bias for an LP at horizon h is linked directly to the LP population coefficients at other horizons. In all of our derivations, we maintain the assumptions from the previous section, and, for algebraic simplicity, we assume that the panel is balanced.

4.1 Bias in LPs with panel data using the AR(1) example

To demonstrate that LPs can be severely biased in small samples with panel data, we generate data, $\{y_{i,t}\}_{t=1}^T$ for each entity $i = 1, \dots, I$, using the data generating process specified in equations (1). For simplicity, we assume that all the data are independent across entities, but our derivations for the approximate bias do not depend on this assumption. We show results for panels containing $I = 10, 25$, and 50 entities. As in the previous section, we assume $\rho = 0.95$.

In panel settings, the LP model is the set of regression models indexed by the impulse response horizon h ,

$$y_{i,t+h} = \alpha_{i,h} + \beta'_h x_{i,t} + u_{i,t,h}, \quad h = 0, \dots, H. \quad (9)$$

where $x_{i,t} \equiv [\varepsilon_{i,t}, c'_{i,t}]'$.⁷ The first element of the coefficient vectors $\{\beta_h\}_{h=0}^H$ trace out the impulse response of interest, which we denote $\{\theta_h\}_{h=1}^H$. Note that we incorporate fixed effects in this panel setting.

Figure 6 displays the Monte-Carlo mean of LP estimators with and without controls for $T = 100$ and different numbers of entities in the panels. As was the case without panel data, the LP estimates of the impulse responses are severely biased. Notably, the bias does not approach zero as the number of entities in the panel grows large (see Nickell (1981)).

As in the non-panel setting, the inclusion of controls is less effective at reducing the bias in LP estimators as h increases. The reason that controls are less effective at reducing the bias as h increases is that they are less effective at forecasting $y_{i,t+h}$. Thus, for large h , the bias is similar to the LP estimator without controls. Papers like Acemoglu et al. (2019) have argued that T as small as 40 should make the bias in panel LP estimators relatively small. While the bias documented by Nickell (1981) is small at $h = 1$ when controls are included, that bias can be large at $h = 10$ even when T is relatively large. In general, impulse responses are most of interest at moderate-to-large values of h .




Figure 6: LP estimators without controls are biased in empirically-relevant samples when $y_{i,t}$ is an AR(1) with $\rho = 0.95$.

4.2 Understanding bias in LPs with panel data

In this subsection we derive an approximate bias function for the LP estimator in the context of panel data with entity fixed effects. We do our expansions under the assumption that the time series is growing, but that the number of entities in the panel is constant. This seems like the most relevant setting for macroeconomic applications, where panels generally consists of countries, states, or even counties. We analyze the implications of having a larger number of entities in a panel for the size of the bias.

Throughout this sub-section, we maintain the following assumption.

Assumption 5. *Either $\varepsilon_{i,t} \perp \varepsilon_{j,t}$ for $j \neq i$, or if $\varepsilon_{i,t}$ and $\varepsilon_{j,t}$ are correlated then $E[\varepsilon_{i,t}(y_{j,t+h} - \alpha_{j,h} - \beta_h \varepsilon_{j,t})] = 0$.*

This assumption says that either the shocks are independent or they are valid instruments for one another.

Analytical Result 3 (Bias in Panel LPs without controls.). *Under the maintained assumptions from this and the previous sections, the series expansion of the OLS estimator without controls implies*

$$\mathbb{E}[\hat{\theta}_{h,LS}] - \theta_h = -\frac{1}{T-h} \sum_{j=1}^{T-h-1} \left(1 - \frac{j}{T-h}\right) (\theta_{h+j} + \theta_{h-j}) + O(T^{-3/2}). \quad (10)$$

Equation (10) is identical to equation (4), meaning that the expression for the bias in a panel setting is identical to the expression without panel data.¹³ As a result, without controls the bias-corrected estimator from the previous section could be applied to the panel data setup we analyze here, where T is used as the number of relevant observations rather than $I \times T$.

When controls are included, the bias calculation becomes more complicated, although the intuition about the size and scope of bias is essentially unchanged. Define covariance of the ε s between panelists as

$$\sigma_{j,k} = \mathbb{E}[(\varepsilon_{j,t} - \mu_{j,\varepsilon})(\varepsilon_{k,t} - \mu_{k,\varepsilon})].$$

Next, write the covariance of $t - u$ and t controls of panelists j and k as

$$\Sigma_{j,k,u} = \mathbb{E}[(c_{j,t-u} - \mu_{c,j})(c'_{k,t} - \mu'_{k,z})].$$

Then, the average of the *variance* of the controls is given by

$$\bar{\Sigma}_0 = \frac{1}{I} \sum_{j=1}^I \Sigma_{j,j,c,0}.$$

Analytical Result 4 (Bias in Panel LPs with controls.). *Under the maintained assumptions from this and the previous sections, the series expansion of the OLS estimator with controls implies*

$$\mathbb{E}[\hat{\theta}_{h,LS}] - \theta_h = -\frac{1}{T-h} \sum_{u=1}^h \left(1 - \frac{u}{T-h}\right) [1 + \omega_u] \theta_{h-u} + O(T^{-3/2}) \quad (11)$$

where

$$\omega_u = \frac{1}{I^2} \sum_{i=1}^I \sum_{k=1}^I \frac{\text{tr}\{\bar{\Sigma}_0^{-1} \Sigma_{k,i,c,u}\} \sigma_{i,k,\varepsilon}}{\frac{1}{I} \sum_{j=1}^I \sigma_{j,j,\varepsilon}} \quad \text{and} \quad \Sigma_{j,k,c,u} = \mathbb{E}[(c_{j,t-u} - \mu_{c,j})(c'_{k,t} - \mu'_{k,z})].$$

¹³The moment conditions with and without controls are given by

$$\mathbb{E} \begin{bmatrix} (y_{1,t+h} - \alpha_{1,t} - \beta x_{1,t}) \\ (y_{2,t+h} - \alpha_{2,t} - \beta x_{2,t}) \\ \vdots \\ (y_{I,t+h} - \alpha_{I,t} - \beta x_{I,t}) \\ \frac{1}{I} \sum_{i=1}^I x_{i,t} (y_{i,t+h} - \alpha_{i,t} - \beta x_{i,t}) \end{bmatrix} = 0.$$

A few comments are in order regarding equation (11). First, as was the case without panel data and because the OLS estimator is consistent, the bias goes to zero as the sample size goes to infinity. Second, the cross-autocovariance of the control variables plays a role in the bias. Notably, if the controls are not correlated, the bias is smaller than if they are correlated. Third, as was the case without panel data, if the controls are positively autocorrelated, or if unnecessary positively autocorrelated controls are included in the LP, 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.

5 Beyond point estimation: bias in standard errors

Since Jordà (2005), the conventional wisdom has been that heteroskedasticity and autocorrelation robust (HAR) standard errors are necessary because the regression residuals of LPs are autocorrelated. That is the reason that most practitioners use the HAR standard errors of Newey-West or more recent ones detailed in Sun (2014) and Lazarus et al. (2018). However, under Assumption 2 in the LP with controls, the *regression score*—the product of the ε_t and the regression residuals—is serially uncorrelated.¹⁴ Thus, in large samples HAR standard errors are not necessary; instead, Huber-White (heteroskedasticity-robust) standard errors are valid.

In an LP without controls, under the AR(1) DGP, one can show that the autocovariance function of the regression score, $r_t = \varepsilon_t(y_{t+h} - \theta_h \varepsilon_t)$, is given by

$$\text{cov}[r_t, r_{t-\tau}] = \begin{cases} \rho^{2h} (\sigma_\varepsilon^2)^2 & \tau = 1, \dots, h \\ 0 & \tau > h. \end{cases} \quad (12)$$

Thus, unless ε_t is uncorrelated with y_t , the regression score will be serially correlated, as suggested by the early LP literature. However, in an LP without controls, researchers are generally interested in rejecting the null hypothesis that θ_h is zero. Under the null hypothesis that the $\theta_h = 0$ for all h , the regression score is uncorrelated in population, and Huber-White standard errors are asymptotically valid for the purposes of hypothesis testing.

¹⁴Olea and Plagborg-Møller (2020) use lag augmentation to achieve (population) residualized regressors, whereas our setup does not require this step because of Assumption 2.

HAR standard errors are often justified by researchers as being somehow more conservative than standard errors that do not control for autocorrelation. While HAR standard errors are asymptotically valid even when Huber-White will do, using popular HAR estimators, like the Newey-West estimator, can dramatically *reduce* the size of the estimated standard errors. The reason is that these estimators rely on estimators of the long-run variance (LRV) of the regression score that are functions of the autocovariance of the regression score. In a linear regression like an LP, standard errors are typically a function of the LRV and an estimator of $\mathbb{E}[x_t x_t']$. In fact, when one has $\varepsilon_t \perp c_t$, it is enough to focus on the estimator of the LRV to analyze the size of the standard errors. In finite samples, estimates of the autocovariance of the regression score are typically *downward biased*. Thus, when the actual autocorrelation function is zero, the bias in the autocorrelation function makes it *negative*, on average, in finite samples. We derive a simple expression for this downward bias in the case of an LP without controls in Appendix B.

In the case of LP with controls, where the autocovariance function of the regression score is zero in population, the downward bias means that increasing the bandwidth of a Newey-West estimator, on average, *reduces* the size of the estimated LRV. In the case of LP without controls, the downward bias means increasing the bandwidth reduces the size of the standard errors, on average, even though the regression score is autocorrelated. Thus, HAR standard errors based on estimated autocovariances of the regression score will tend to *underestimate the uncertainty* associated with $\hat{\theta}_{h,LS}$.

To explore the effects of the downward bias in standard errors, we consider Monte Carlo evidence using our AR(1) example. Figure 7 shows the true values of the LRV for the regression score in an LP without controls (panels (a) and (b)) and in an LP with controls (panels (c) and (d)). The horizon of the LP is $h = 5$. The figures also show the Monte Carlo average of Newey-West estimators that use a bandwidth u . The left panels (panels (a) and (c)) show the estimators in the case when $\theta_h = 0$ for all h —that is, the shock is noise uncorrelated with y_t —which is typically the maintained null hypothesis to reject. The right panels (panels (b) and (d)) show the estimators in the case when the shock is correlated with y_t , so $\theta_h = \rho^h$.¹⁵

Several features of these figures are worth noting. First, in all cases the Newey-West estimators dramatically under-estimate the long-run variance. Second, increasing the band-

¹⁵We keep the variance of y_t fixed across all of our specifications.




Figure 7: Estimators of standard errors in an LP with $h = 5$ are biased in empirically-relevant samples when y_t is an AR(1) with $\rho = 0.95$.

width of the Newey-West estimator tends to reduce the estimate of the LRV. Third, even when there is autocorrelation in the population regression score (panel (b)), in empirically relevant sample sizes, increasing the bandwidth of the the Newey-West estimator does little to improve the estimate of the long-run variance and can reduce that estimate. Fourth, the Huber-White estimators, which are given by the Newey-West estimator with $u = 0$, are also downward biased. However, the Huber-White estimator is not further attenuated by the estimates of the autocovariances of the regression score included in the Newey-West estimator.

Often, standard errors are interpreted as providing a credible region around the estimated impulse response without reference to a null hypothesis. To this end, it is also useful to consider the implications of bias in point estimates and standard errors under the maintained assumption that $\theta_h \neq 0$. We analyze this issue using Monte Carlo analysis and an AR(1) generating process for y_t . The OLS estimator of the impulse response function is biased down and typical estimators of standard errors are also biased down. Both sources of bias may make confidence intervals fail to contain the true impulse response function.

Using Huber-White and Newey-West standard errors, Table 1 displays the percentage of confidence sets with 95% nominal coverage probability (based on the asymptotic normal approximation) that contain the true impulse response function at horizon h in a Monte Carlo simulation of an LP without controls where y_t is an AR(1) with $\rho = 0.95$ and a sample size of $T = 50$.¹⁶ Several features of these results are worth noting. First, implementing our proposed bias correction and using controls can increase coverage probabilities. Second, even with our bias correction and controls, the coverage probabilities are markedly lower than 95%. Third, the Huber-White standard errors appear to perform better than the Newey-West standard errors.

Overall, our analysis of standard errors indicates that bias in LPs is an important problem to address when conducting inference. In small samples, Huber-White standard errors are preferable to HAR standard errors when using critical values from the normal asymptotic limiting distribution. When researchers are interested in credible regions around point estimates, implementing our bias adjustment to $\hat{\theta}_{h,LS}$ can improve the coverage probabilities.

¹⁶When constructing the Newey-West estimator, we use a bandwidth of $0.7(T-h)^{1/3}$. We use this bandwidth as the “textbook” choice from Lazarus et al. (2018). When one uses the fixed-b critical values suggested by Lazarus et al. (2018), the performance of the Newey-West estimator improves somewhat (see our Appendix), though the coverage probabilities are not uniformly better than those from the Huber-White standard errors. Fixed-b asymptotics involve using a larger bandwidth for the Newey-West estimator and larger critical values than those implied by the asymptotic normal approximation. It turns out that the bandwidth is not that much larger in sample sizes typically found in the literature. As a result, it is not surprising that using a larger critical value improves the coverage probabilities given that the confidence intervals are too small.

Table 1: Coverage probability of different estimators of standard errors for $\hat{\theta}_h$ in LP when y_t is an AR(1) with $\rho = 0.95$ and $T = 50$

| h | $\hat{\theta}_{h,LS}$, no controls | | $\hat{\theta}_{h,BCC}$, no controls | | $\hat{\theta}_{h,BCC}$, controls | |
|-----|-------------------------------------|------|--------------------------------------|------|-----------------------------------|------|
| | HW | NW | HW | NW | HW | NW |
| 0 | 0.87 | 0.82 | 0.86 | 0.82 | 0.92 | 0.91 |
| 1 | 0.83 | 0.80 | 0.82 | 0.80 | 0.90 | 0.88 |
| 2 | 0.80 | 0.77 | 0.79 | 0.78 | 0.87 | 0.86 |
| 3 | 0.78 | 0.75 | 0.76 | 0.75 | 0.85 | 0.83 |
| 4 | 0.76 | 0.73 | 0.75 | 0.73 | 0.83 | 0.81 |
| 5 | 0.75 | 0.72 | 0.74 | 0.72 | 0.81 | 0.79 |
| 6 | 0.75 | 0.72 | 0.73 | 0.71 | 0.80 | 0.78 |
| 7 | 0.74 | 0.70 | 0.73 | 0.70 | 0.78 | 0.76 |
| 8 | 0.74 | 0.71 | 0.73 | 0.70 | 0.77 | 0.75 |
| 9 | 0.74 | 0.71 | 0.73 | 0.70 | 0.76 | 0.74 |
| 10 | 0.74 | 0.71 | 0.73 | 0.70 | 0.75 | 0.74 |

Note: HW stands for Huber-White standard errors, NW stands for Newey-West standard errors.

6 Application to monetary policy shocks

In this section we provide empirical examples of estimated bias in LPs and also show the sensitivity of inference to the construction of standard errors. We highlight three examples from the literature on the effect of monetary policy shocks on the macroeconomy. These shocks are constructed using either the narrative approach of Romer and Romer (2004) or through asset prices as pioneered by Kuttner (2001).

6.1 The effects of monetary policy shocks on output and inflation

Using a setup similar to Gorodnichenko and Lee (2019), we estimate the effects of Romer and Romer (2004) monetary policy shocks on output and inflation.¹⁷ The data sample runs from 1969:Q1-2008:Q4. We estimate LPs of the form in equation (2) on real output growth

¹⁷The shock series was extended to 2008 by Coibion et al. (2017).



Figure 8: The effect of monetary policy shocks.

and annualized inflation. We include controls consisting of four lags of real output growth, inflation, the federal funds rate, and the monetary policy innovation.¹⁸

The estimated impulse responses of inflation and output to a monetary policy shock are displayed in Figure 8. As in Gorodnichenko and Lee (2019), we cumulate the impulse response of output growth. Figure 8 also shows the bias-corrected estimate of the impulse response. To focus attention on the difference between the two impulse responses in this illustrative example, we omit confidence bands.

The estimated inflation impulse response roughly accords with Gorodnichenko and Lee (2019): a contractionary 100 basis-point monetary policy shock causes inflation to be little changed for the first few periods after the shock and then eventually decline persistently. The bias-corrected impulse response indicates that inflation responds somewhat more to a monetary policy shock than under the conventional estimates. On average, the response of inflation is about 15 basis points lower in the bias-corrected impulse response, a moderate but nontrivial difference. The bias-corrected estimator is lower than the least-squares LP estimator because the estimated least-squares impulse response is negative for almost hori-

¹⁸The LPs here are slightly different from the ones in Gorodnichenko and Lee (2019) in two ways. First, we use y_{t+h} rather than $y_{t+h} - y_{t-1}$. The bias discussed in this paper is still present under the latter formula. Second, we omit TFP innovations because the objective here is not to study relative variance contributions. Taken together, these differences lead to only minor changes in the estimated LPs.

zons, and the bias-correction is a weighted average of the impulse responses at all previous horizons.

The estimated output impulse response also broadly accords with the results in Gorodnichenko and Lee (2019): a contractionary 100 basis-point monetary policy shock causes the level of output to contract by about 4 percent after two and half years, after which effects of the shock slowly dissipate. Notably, the bias-corrected estimator implies the output decline is about 1/2 percentage point larger. Essentially, this is because the bias-correction at, say, horizon $h = 15$ is influenced by the LP coefficients at previous horizons.

For both inflation and output, the corrections are moderate but economically meaningfully. The corrections are larger at longer horizons, a consequence of the single-sided nature of our bias correction, given in equation (8).¹⁹ It is perhaps not surprising that the bias correction is moderate given that the sample size used here is markedly larger than those typically found in the LP literature.

6.2 State dependence in the effects of monetary policy shocks

One of the key advantages of LPs, relative to other popular methods in macroeconomics, are their ability to handle nonlinearities. Tenreyro and Thwaites (2016) use a “smoothly transitioning local projection model” to investigate whether monetary policy has larger effects during recessions or expansions. The LP specification is given by

$$y_{t+h} = \tau t + F(z_t)(\alpha_h^b + \theta_h^b \varepsilon_t + \gamma^{b'} c_t) + (1 - F(z_t))(\alpha_h^r + \theta_h^r \varepsilon_t + \gamma^{r'} c_t) + u_t. \quad (13)$$

Here y_{t+h} is the endogenous variable of interest (output), c_t is a vector of controls and $F(z_t)$ is a smooth increasing function of an indicator of the state of the economy z_t . The value of $F(z_t)$ is exogenous from the perspective of the regression. The shock ε_t is once again a variant of the Romer and Romer (2004) measure. The coefficients of interest are θ_h^b and θ_h^r , the effects of a monetary policy shock at horizon h in a boom (b) and recession (r), respectively.

¹⁹An earlier working paper version of this paper presented a slightly different formula for the bias correction in the case of LP with controls. While both are valid, the correction in the previous version for horizon h was influenced by estimates of impulse responses at all horizons. The current formula depends only on horizons $j \leq h$. This is much easier to implement in practice.

Since $F(z_t)$ is exogenous, we can estimate this LP model by OLS. We proceed by first removing the trend and estimate (13) and then omit the linear trend in the regression specification.²⁰ We estimate the model on quarterly data using a sample that runs from 1969:Q1 to 2002:Q4. Following, Tenreyro and Thwaites (2016), the control variables c_t are one lag of detrended output and the federal funds rate.

Table 2: State-Dependent Effects of Monetary Policy Shocks

| Horizon | | $\hat{\beta}_{LS}^b$ | $\hat{\beta}_{LS}^r$ | Diff. | $\hat{\beta}_{BC}^b$ | $\hat{\beta}_{BC}^r$ | Diff. |
|----------|------------|----------------------|----------------------|---------|----------------------|----------------------|---------|
| $h = 9$ | Point Est. | -1.27 | -0.70 | -0.58 | -1.36 | -0.80 | -0.56 |
| | NW SE | (0.39) | (0.29) | (0.59) | (0.39) | (0.29) | (0.58) |
| | HW SE | (0.45) | (0.64) | (0.89) | (0.46) | (0.65) | (0.90) |
| $h = 10$ | Point Est. | -1.49 | -0.41 | -1.08 | -1.61 | -0.52 | -1.09 |
| | NW SE | (0.37) | (0.25) | (0.53) | (0.37) | (0.28) | (0.55) |
| | HW SE | (0.50) | (0.74) | (1.01) | (0.51) | (0.75) | (1.02) |
| $h = 11$ | Point Est. | -1.10 | -0.38 | -0.73 | -1.26 | -0.49 | -0.77 |
| | NW SE | (0.38) | (0.34) | (0.64) | (0.39) | (0.39) | (0.67) |
| | HW SE | (0.60) | (0.72) | (1.05) | (0.60) | (0.73) | (1.06) |
| $h = 12$ | Point Est. | -0.64 | -0.47 | -0.17 | -0.81 | -0.58 | -0.23 |
| | NW SE | (0.45) | (0.44) | (0.80) | (0.46) | (0.49) | (0.85) |
| | HW SE | (0.65) | (0.62) | (1.01) | (0.64) | (0.64) | (1.02) |

Notes: Table shows the LS and BC estimates of β^b and β^r with Newey-West (NW) and Huber-White (HW) standard errors.

Table 2 shows the LS and BC estimates of with β^b and β^r , along with with Newey-West and Huber-White standard errors in parenthesis, for $h = 9, \dots, 12$. The coefficient estimates can be interpreted as the percentage point effect on real output in period $t + h$ of a 100 point monetary policy shock at time t when the economy is definitely in either a boom or recession, respectively. The point estimates under both LS and BC indicate that positive shocks occurring during booms are more contractionary than identically-sized shocks occurring during recessions, consistent with Tenreyro and Thwaites (2016). Relative to the LS estimates, the bias-corrected estimates are more negative for both β_h^b and β_h^r , consistent

²⁰Results are similar under alternative ways of detrending output.

with the earlier simulation study.

The difference between the coefficients for booms and recessions is larger when using the bias-corrected estimators. That is, there is moderately more evidence for differential effects of monetary policy shocks depending on the state of the business cycle. As in Tenreyro and Thwaites (2016), the uncertainty surrounding these estimates is large. Notably, the Huber-White standard errors are markedly larger than the Newey-West standard errors, consistent with our analysis in Section 5. For a number of the entries in Table 2, conclusions about statistical significance would be different using the Huber-White standard errors as compared to the Newey-West standard errors.

6.3 Time dependence in the effect of monetary policy shocks

Our final example assesses the evidence for a change in transmission of monetary policy over the early part of the 2000s. Lunsford (2020) identifies two monetary policy shocks from high frequency movements in assets in prices. The first shock is to the level of current federal funds rate and the second shock, the *forward guidance shock*, is to the market’s expected path of federal funds rate beyond the current month. Lunsford (2020) argues that the responses of asset prices and macroeconomic aggregates to forward guidance shocks changed in 2003, as the policy statement issued by the Federal Open Market Committee (FOMC) following monetary policy decisions shifted to emphasize future policy inclinations rather than risks to the economic outlook. Using the identified federal funds rate (ε_t^{FFR}) and forward guidance ε_t^{FG} shocks, Lunsford (2020) effectively estimates LPs of the form

$$y_{t+h} - y_t = \alpha_h + \theta_h^{FFR} \varepsilon_t^{FFR} + \theta_h^{FG} \varepsilon_t^{FG} + u_{t,t+h}, \quad h = 1, \dots, H. \quad (14)$$

where y_{t+h} is a macroeconomic aggregate. The focus here will be on estimates of the coefficient θ_h^{FG} and whether they have changed over two samples of monthly data from February 2000 to June 2003 and August 2003 to May 2006. Following Lunsford (2020), the horizon of interest is $h = 12$, so θ_{12}^{FG} measures the effect of a forward guidance shock one year after its realization (after netting the contemporaneous effect). The size of the two samples deserves emphasis as they are 28 and 23 observations, respectively. These are extremely small—about half of the smallest sample considered in the Monte Carlo simulations in this paper.²¹

²¹For ease of exposition, the paper assumes that the econometrician observes T total observations of both y_t and ε_t (and potentially controls.) So there are $T - h$ observations for the h th horizon LP. In some practical

Before describing the regression results, we note two things about the LP model in (14). First, subtracting y_t from the dependent variables tends to reduce the parameter bias substantially. The model in equation (14) is essentially an LP without controls. One can see, either from the trajectories in Figure 2 or the expression in equation (4), that in an LP without controls a sizable portion of the bias associated with least squares estimate of θ_h at horizons greater than zero is accounted for by the bias in θ_0 . Informally, subtracting y_t from y_{t+h} out removes this portion of the bias.

Table 3 displays $\hat{\theta}_{12,LS}^{FG}$ from the LP model in equation (14) along with Newey-West and Huber-White standard errors over the two samples for three dependent variables: the growth in real personal consumption expenditures (PCE), unemployment rate changes, and the growth in industrial production (IP). The Newey-West standard errors are computed using $u = 10$ lags. The point estimates, as mentioned above, are computed using least squares, as the bias correction would likely be minor when using $y_{t+h} - y_t$ as the dependent variable.

The point estimates indicate that while contractionary—i.e., positive—forward guidance shocks were associated with increases in real PCE growth and a fall in the unemployment rate in the first sample, they were associated with a fall in consumption growth and an increase in unemployment over the second sample. The point estimates associated with IP are both negative. Consistent with standard practice, Lunsford (2020) uses Newey-West standard errors, which are replicated in Table 3; when using the fixed- b asymptotic critical values from Sun (2014), which Table 3 employs, the coefficients associated with PCE growth and the unemployment rate are moderately statistically significant. Table 3 also displays the Huber-White standard errors. With one exception—the unemployment rate in the August to May 2006 sample—these are larger than the Newey-West standard errors, sometimes substantially so. The standard error estimates associated with PCE growth and IP growth in the second sample increase from 3.82 to 6.15 and 4.37 to 10.90, respectively.²²

applications, the econometrician is limited only by the observations of shock. That is, they observe a sample of size $T + H$ of y_t and a sample of size T of ε_t , where H is the maximum horizon considered. In this case, there are T observations for each of the h th horizon LPs—this is the case in Lunsford (2020). Using a constant rather than shrinking sample size results in extremely minor modifications of the analytic expressions and essentially no changes to the Monte Carlo simulations in this paper.

²²In this case, whether the critical values used to assess statistical significance should change depending on the construction of the standard errors is complicated and requires an explicit statement of the null

Table 3: Response to Forward Guidance Shock

| | Feb. 2000 to Jun. 2003 | Aug. 2003 to May 2006 |
|----------------------|------------------------|-----------------------|
| PCE Growth Results | | |
| Point Estimate | 2.88 | -10.51 |
| NW SE | (2.26) | (3.82) |
| HW SE | (2.61) | (6.15) |
| Unemployment Results | | |
| Point Estimate | -4.01 | 3.49 |
| NW SE | (1.13) | (1.24) |
| HW SE | (2.45) | (1.04) |
| IP Growth Results | | |
| Point Estimate | -9.87 | -10.10 |
| NW SE | (8.53) | (4.37) |
| HW SE | (12.03) | (10.90) |

Notes: The Table shows $\hat{\theta}_{12,LS}^{FG}$ from the LP model in (14) along with Newey-West (NW) and Huber-White (HW) standard errors. The NW standard errors are computed using a bandwidth of $u = 10$.

As argued in Section 5, in small samples, even when appropriate, HAC standard errors tend to underestimate the long-run variance relative to the Huber-White estimator. Note that this is true even when the estimate of the coefficient of interest exhibits little bias itself, as we have argued is plausible in this case. The Huber-White standard errors indicate that the uncertainty surrounding these regression coefficients is considerably larger than implied the Newey-West standard errors.

7 Conclusion

We have shown that LPs can be severely biased in sample sizes commonly found in the related literature. We derived an approximate bias function that shows that LPs are intimately hypothesis. In any event, given such small sample sizes, relying on critical values associated with limiting distributions could be problematic.

linked across horizons in small samples. The bias of LPs persists even when researchers have access to a large cross-section (panel data).

We used our approximate bias function to bias correct LPs. In Monte Carlo, analysis our bias correction does not completely correct for the bias in LPs. These results suggest that other time series models with well-understood, effective methods for bias correction (such as VARs) may be better alternatives for estimated 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 also analyzed bias in standard errors computed for estimated impulse response functions from LPs. We showed that, in small samples, standard errors that rely on estimated autocovariances of the regression score, like the Newey-West estimator, typically understate the amount of uncertainty surrounding the estimated impulse response functions. We argued that, in LPs similar to those we consider, researchers should prefer standard errors that are heteroskedasticity consistent, but not autocorrelation robust.

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, it is difficult to appeal to limiting critical values as accurate approximations. As a result, if researchers are going to use HAR standard errors, they may want to check to see if Huber-White standard errors would lead to different conclusions. If the Huber-White standard errors are larger than the HAR standard errors, researchers should consider what might lead to the apparent negative autocovariance in the regression score. Without another reasonable theory, it may be that the negative estimates of the autocovariance of the regression score are the result of small sample bias.

References

- ACEMOGLU, D., S. NAIDU, P. RESTREPO, AND J. A. ROBINSON (2019): “Democracy Does Cause Growth,” *Journal of Political Economy*, 127, 47–100.
- 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.
- COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): “Innocent Bystanders? Monetary Policy and Inequality,” *Journal of Monetary Economics*, 88, 70–89.
- GORODNICHENKO, Y. AND B. LEE (2019): “Forecast Error Variance Decompositions with Local Projections,” *Journal of Business & Economic Statistics*, 0, 1–24.
- 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.
- HALL, P. (1992): *The Bootstrap and Edgeworth Expansion*, Springer New York.
- HALL, P., J. L. HOROWITZ, AND B.-Y. JING (1995): “On Blocking Rules for the Bootstrap with Dependent Data,” *Biometrika*, 82, 561–574.
- 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.
- KUTTNER, K. N. (2001): “Monetary Policy Surprises and Interest Rates: Evidence From the Federal Funds Futures Market,” *Journal of Monetary Economics*, 47, 523–544.
- 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.
- LUNSFORD, K. G. (2020): “Policy Language and Information Effects in the Early Days of Federal Reserve Forward Guidance,” *American Economic Review*, 110, 2899–2934.

- LUSOMPA, A. (2019): “Local Projections, Autocorrelations, and Efficiency,” *Unpublished Manuscript, UC-Irvine*.
- NAKAMURA, E. AND J. STEINSSON (2018): “Identification in Macroeconomics,” *Journal of Economic Perspectives*, 32, 59–86.
- NEWBY, W. K. AND K. D. WEST (1987): “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation,” *Econometrica*, 55, 703–708.
- NICKELL, S. (1981): “Biases in Dynamic Models with Fixed Effects,” *Econometrica*, 49, 1417–1426.
- OLEA, J. L. M. AND M. PLAGBORG-MØLLER (2020): “Local Projection Inference is Simpler and More Robust Than You Think,” .
- PLAGBORG-MØLLER, M. AND C. K. WOLF (2019): “Local Projections and VARs Estimate the Same Impulse Responses,” *Unpublished Manuscript, Princeton University*.
- 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.
- SHAMAN, P. AND R. A. STINE (1988): “The Bias of Autoregressive Coefficient Estimators,” *Journal of the American Statistical Association*, 83, 842–848.
- 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.
- 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.

A Derivations of approximate bias

In this appendix, we derive our expressions for the 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:

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

It can be verified that under Assumption 1, combined with the least squares estimation framework, satisfies the necessary assumptions of Rilstone et al. (1996). Assumptions 2 and 3 allow one to obtain tractable expressions.

In our derivation, we use the notation of Bao and Ullah (2007) where possible. For each derivation, we will cast the OLS estimator as a GMM problem with moment conditions given by $q(\beta; w_t)$, where the data vector $w_t = [y_t, x_t']'$ for the LP models with and without controls. For panel settings, the data vector is expanded to include the additional observables. The objective function is thus given by:

$$\psi_{T-h}(\beta; W_{1:T}) = \frac{1}{T-h} \sum_{t=1}^{T-h} q(\beta; w_t).$$

Let $\nabla^i A(\beta)$ be the matrix of i th order partial derivatives of A with respect to β . In what follows, write $\psi_{T-h}(\beta; W_{1:T})$ as ψ_{T-h} and $q(\beta; w_t)$ as q_t . Define the series of matrices

$$H_i = \nabla^i \psi_{T-h} \text{ and } \bar{H}_i = \mathbb{E}[H_i] \text{ with } Q = \bar{H}_1^{-1}, V = H_1 - \bar{H}_1 \text{ and } W = H_2 - \bar{H}_2.$$

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

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

We are interested in computed in the bias, that is $\mathbb{E}[\hat{\beta} - \beta]$. It is obvious that $\mathbb{E}[a_{-1/2}] = 0$. Moreover, because the moment conditions associated with the LP estimator are linear in β , we have $H_2 = \bar{H}_2 = 0$. It can be verified that the expression for the bias to $O(T^{-1})$ thus simplifies considerably to

$$\mathbb{E}[\hat{\beta} - \beta] \approx \mathbb{E}[QH_1Q\psi_{T-h}] := B. \quad (\text{A.2})$$

Our objective is to derive the element of the vector B that is associated with the shock of interest ε for the LP model with and without controls. Before proceeding with introduce notation to define first and second moments of the shocks, ε_t , and controls, c_t :

$$\mu_\varepsilon = \mathbb{E}[\varepsilon_t], \quad \sigma_\varepsilon^2 = \mathbb{E}[(\varepsilon_t - \mu_\varepsilon)^2], \quad \mu_c = \mathbb{E}[c_t], \quad \text{and } \Sigma_c = \mathbb{E}[(c_t - \mu_c)(c_t - \mu_c)'].$$

A.1 LP without controls

For the LP without controls, the moment conditions associated with the OLS estimator are defined as

$$q_t \equiv \begin{bmatrix} y_{t+h} - \alpha_h - \theta_h \varepsilon_t \\ \varepsilon_t (y_{t+h} - \alpha_h - \theta_h \varepsilon_t) \end{bmatrix},$$

where we have departed slightly from the notation of Bao and Ullah (2007) in defining the parameter vector. Before constructing the matrices of derivatives, note that we can rearrange the moment conditions to deduce:

$$\theta_h = \frac{\mathbb{E}[(\varepsilon_t - \mu_\varepsilon)(y_{t+h} - \alpha_h)]}{\mathbb{E}[(\varepsilon_t - \mu_\varepsilon)^2]} = \frac{\mathbb{E}[(\varepsilon_t - \mu_\varepsilon)(y_{t+h} - \alpha_h)]}{\sigma_\varepsilon^2}$$

Additionally for $s < t$, we have

$$\theta_h = \frac{\mathbb{E}[(\varepsilon_t - \mu_\varepsilon)(y_{t+h} - \mathbb{E}_s[y_{t+h}] + \mathbb{E}_s[y_{t+h}] - \alpha_h)]}{\sigma^2} = \frac{\mathbb{E}[(\varepsilon_t - \mu_\varepsilon)(y_{t+h} - \mathbb{E}_s[y_{t+h}] - \alpha_h)]}{\sigma^2} \quad (\text{A.3})$$

Where the second equality follows from the fact that $\mathbb{E}_s[\varepsilon_t] = \mu$ for all $s < t$. It is easy to see that:

$$H_1 = \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} -1 & -\varepsilon_t \\ -\varepsilon_t & -\varepsilon_t^2 \end{bmatrix}, \quad \bar{H}_1 = \begin{bmatrix} -1 & -\mu_\varepsilon \\ -\mu_\varepsilon & -(\sigma_\varepsilon^2 + \mu_\varepsilon^2) \end{bmatrix}, \quad \text{and } Q = \begin{bmatrix} -\left(1 + \frac{\mu_\varepsilon^2}{\sigma_\varepsilon^2}\right) & \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} \\ \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} & -\frac{1}{\sigma_\varepsilon^2} \end{bmatrix}.$$

Tedious algebra yields:

$$QH_1Q = \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} -\left(1 - 2\frac{\mu_\varepsilon}{\sigma_\varepsilon^2}(\varepsilon_t - \mu_\varepsilon) + \left(\frac{\mu_\varepsilon}{\sigma_\varepsilon^2}\right)^2(\varepsilon_t - \mu_\varepsilon)^2\right) & -\left(\frac{1}{\sigma_\varepsilon^2}(\varepsilon_t - \mu_\varepsilon) - \frac{\mu_\varepsilon}{(\sigma_\varepsilon^2)^2}(\varepsilon_t - \mu_\varepsilon)^2\right) \\ -\frac{1}{\sigma_\varepsilon^2}(\varepsilon_t - \mu_\varepsilon) + \frac{\mu_\varepsilon}{(\sigma_\varepsilon^2)^2}(\varepsilon_t - \mu_\varepsilon)^2 & -\frac{1}{(\sigma_\varepsilon^2)^2}(\varepsilon_t - \mu_\varepsilon)^2 \end{bmatrix}.$$

Recall that we only need to obtain the second element of B which corresponds with the bias associated with $\hat{\theta}_h$. Thus, we only need to calculate $[QH_1Q]_{2,\psi_{T-h}}$. This is given by:

$$\begin{aligned} [QH_1Q]_{2,\psi_{T-h}} &= \frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \left(-\frac{1}{\sigma_\varepsilon^2}(\varepsilon_t - \mu_\varepsilon) + \frac{\mu_\varepsilon}{(\sigma_\varepsilon^2)^2}(\varepsilon_t - \mu_\varepsilon)^2 \right) \sum_{t=1}^{T-h} (y_{t+h} - \alpha_h - \theta_h \varepsilon_t) \\ &\quad - \frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \frac{1}{(\sigma_\varepsilon^2)^2}(\varepsilon_t - \mu_\varepsilon)^2 \sum_{t=1}^{T-h} \varepsilon_t (y_{t+h} - \alpha_h - \theta_h \varepsilon_t) \\ &= \frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} -\frac{1}{\sigma_\varepsilon^2}(\varepsilon_t - \mu_\varepsilon)(y_{s+h} - \alpha_h - \theta_h \varepsilon_s) - \frac{1}{(\sigma_\varepsilon^2)^2}(\varepsilon_t - \mu_\varepsilon)^2(\varepsilon_s - \mu_\varepsilon)(y_{s+h} - \alpha_h - \theta_h \varepsilon_s) \\ &= \frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} (\phi_I(t, s) + \phi_{II}(t, s)) \end{aligned}$$

Consider the expectation of the term ϕ_I . We have:

$$\mathbb{E}[\phi_I(t, s)] = \begin{cases} 0 & \text{if } t = s \\ -\theta_{h+s-t} & \text{otherwise.} \end{cases}$$

Now consider the expectation of the term ϕ_{II} ,

$$\begin{aligned} \mathbb{E}[\phi_{II}(t, s)] &= \frac{1}{(\sigma_\varepsilon^2)^2} \mathbb{E}[(\varepsilon_t - \mu_\varepsilon)^2(\varepsilon_s - \mu_\varepsilon)(y_{s+h} - \alpha_h - \theta_h \varepsilon_s)] \\ &= \frac{1}{(\sigma_\varepsilon^2)^2} \mathbb{E}[(\varepsilon_t - \mu_\varepsilon)^2 \mathbb{E}[(\varepsilon_s - \mu_\varepsilon)(y_{s+h} - \alpha_h - \theta_h \varepsilon_s) | \varepsilon_t]] \\ &= \frac{1}{(\sigma_\varepsilon^2)^2} \mathbb{E} \left[(\varepsilon_t - \mu_\varepsilon)^2 \mathbb{E} \left[(\varepsilon_s - \mu_\varepsilon) \left(\underbrace{y_{s+h} - \alpha_{h-(t-s)} - \theta_{h-(t-s)} \varepsilon_t}_{\delta_I(t,s)} + \underbrace{\alpha_{h-(t-s)} + \theta_{h-(t-s)} \varepsilon_t - \alpha_h - \theta_h \varepsilon_s}_{\delta_{II}(t,s)} \right) | \varepsilon_t \right] \right]. \end{aligned}$$

By Assumption 2, ε_t does not enter into $\delta_I(t, s)$. Moreover, under Assumption 2, conditioning on it does not convey any useful information about ε_s or other components of y_{t+s} .

$$\mathbb{E}[(\varepsilon_s - \mu_\varepsilon)\delta_I(t, s)|\varepsilon_t] = \mathbb{E}[(\varepsilon_s - \mu_\varepsilon)\delta_I(t, s)]. \quad (\text{A.4})$$

Direct calculation yields:

$$\mathbb{E}[(\varepsilon_s - \mu_\varepsilon)\delta_I(t, s)] = \begin{cases} 0 & \text{if } t = s \\ \theta_h \sigma_\varepsilon^2 & \text{otherwise.} \end{cases}$$

Consider now the term involving $\delta_{II}(t, s)$. By direct calculation:

$$\mathbb{E}[(\varepsilon_s - \mu_\varepsilon)\delta_{II}(t, s)|\varepsilon_t] = \begin{cases} 0 & \text{if } t = s \\ -\theta_h \sigma_\varepsilon^2 & \text{otherwise.} \end{cases}$$

Thus $\mathbb{E}[\phi_{II}(t, s)] = 0$ for all t and s . Combining these results, we have:

$$\mathbb{E}[[QH_1Q]_{2,\cdot}\psi_{T-h}] = \frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} -\theta_{h+s-t} \mathbf{1}_{\{t \neq s\}}.$$

Tedious arithmetic confirms that:

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

This delivers equation (4).

A.2 LP with controls

In the case of controls, we have $x_t = [\varepsilon_t, c_t]'$, so the moment conditions associated with the OLS estimator are defined so that

$$q_t \equiv \begin{bmatrix} y_{t+h} - \alpha_h - x_t' \beta_h \\ x_t (y_{t+h} - \alpha_h - x_t' \beta_h) \end{bmatrix} = \begin{bmatrix} y_{t+h} - \alpha_h - x_t' \beta_h \\ \varepsilon_t (y_{t+h} - \alpha_h - x_t' \beta_h) \\ c_t (y_{t+h} - \alpha_h - x_t' \beta_h) \end{bmatrix}.$$

Recall the object of interest, θ_h , is the first element of β_h . We have that

$$H_1 = \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} -1 & -\varepsilon_t & -c_{t-1}' \\ -\varepsilon_t & -\varepsilon_t^2 & -\varepsilon_t c_{t-1}' \\ -c_{t-1} & -\varepsilon_t c_{t-1} & -c_{t-1} c_{t-1}' \end{bmatrix}, \quad \bar{H}_1 = \begin{bmatrix} -1 & -\mu_\varepsilon & -\mu_c' \\ -\mu_\varepsilon & -(\sigma_\varepsilon^2 + \mu_\varepsilon^2) & -\mu_\varepsilon \mu_c' \\ -\mu_c & -\mu_\varepsilon \mu_c & -(\Sigma_c + \mu_c \mu_c') \end{bmatrix},$$

$$\text{and } Q = \begin{bmatrix} -\left(1 + \frac{\mu_\varepsilon^2}{\sigma_\varepsilon^2} + \mu_c' \Sigma_c^{-1} \mu_c\right) & \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} & \mu_c' \Sigma_c^{-1} \\ \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} & -\frac{1}{\sigma_\varepsilon^2} & 0 \\ \Sigma_c^{-1} \mu_c & 0 & -\Sigma_c^{-1} \end{bmatrix}.$$

As with previous derivation, we need only calculate the second row of QH_1Q . Direct calculation yields

$$[QH_1Q]_{2,\cdot} = \frac{1}{T-h} \sum_{t=1}^{T-h} \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) \begin{bmatrix} 1 - \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) - \mu'_c \Sigma_c^{-1} (c_{t-1} - \mu_c) \\ \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) \\ \Sigma_c^{-1} (c_{t-1} - \mu_c) \end{bmatrix}'.$$

Then the second element of $QH_1Q\psi_{T-h}$ is given by

$$\begin{aligned} [QH_1Q\psi_{T-h}]_2 &= \frac{1}{T-h} \sum_{t=1}^{T-h} \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) \begin{bmatrix} 1 - \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) - \mu'_c \Sigma_c^{-1} (c_{t-1} - \mu_c) \\ \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) \\ \Sigma_c^{-1} (c_{t-1} - \mu_c) \end{bmatrix}' \\ &\quad \times \frac{1}{T-h} \sum_{s=1}^{T-h} \begin{bmatrix} y_{s+h} - \alpha_h - x'_s \beta_h \\ \varepsilon_s (y_{s+h} - \alpha_h - x'_s \beta_h) \\ c_s (y_{s+h} - \alpha_h - x'_s \beta_h) \end{bmatrix}. \end{aligned}$$

Explicit calculation of this object yields:

$$\begin{aligned} [QH_1Q\psi_{T-h}]_2 &= -\frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) (y_{s+h} - \alpha_h - x'_s \beta_h) \\ &\quad - \frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \left[\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) \right]^2 (\varepsilon_s - \mu_\varepsilon) (y_{s+h} - \alpha_h - x'_s \beta_h) \\ &\quad - \frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) (c_{t-1} - \mu_c)' \Sigma_c^{-1} (c_{s-1} - \mu_c) (y_{s+h} - \alpha_h - x'_s \beta_h) \\ &= -\frac{1}{(T-h)^2} \sum_{t=1}^{T-h} \sum_{s=1}^{T-h} \phi_I(t, s) + \phi_{II}(t, s) + \phi_{III}(t, s). \end{aligned}$$

Consider first $\mathbb{E}[\phi_I(t, s)]$. Direct calculation yields:

$$\mathbb{E}[\phi_I(t, s)] = \begin{cases} \theta_{h-(t-s)} & \text{if } s < t \leq s+h \\ 0 & \text{otherwise.} \end{cases}$$

Note, unlike the no controls case, for $t < s$, $\mathbb{E}[\phi_I(t, s)] = 0$ as a direct consequence of Assumption 3. By similar argument to the previous section, $\mathbb{E}[\phi_{II}(t, s)] = 0$ for all t and s . Consider the expectation of $\phi_{III}(t, s)$:

$$\mathbb{E}[\phi_{III}(t, s)] = \mathbb{E} \left[\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) (c_{t-1} - \mu_c)' \Sigma_c^{-1} (c_{s-1} - \mu_c) (y_{s+h} - \alpha_h - x'_s \beta_h) \right]$$

By direct calculation, we have:

$$\mathbb{E}[\phi_{III}(t, s)] = \begin{cases} \theta_{h-(t-s)} \mathbb{E} [(c_{t-1} - \mu_c)' \Sigma_c^{-1} (c_{s-1} - \mu_c)] & \text{if } s < t \leq s+h \\ 0 & \text{otherwise.} \end{cases}$$

Plugging these expectations into the expression for B delivers equation (8).

B Derivation of bias in standard errors

We are going to consider a case where a researcher uses an LP without controls and wants to estimate the standard error of the estimator. To do so, the researcher needs to compute the variance of the regression score, which is composed of terms of the form

$$\gamma_{h,u} \equiv \mathbb{E} \left[\varepsilon_t (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) (y_{t+h-u} - \alpha_h - \varepsilon_{t-u} \beta_h) \varepsilon_{t-u} \right].$$

The researcher wants to report standard errors under the null hypothesis that the coefficient on the shock of interest is zero. If this hypothesis is maintained at all horizons, then the researcher's maintained assumption is that

$$\varepsilon_t \perp y_s$$

for all t and s . We will use this maintained assumption. In addition, we assume that $\mathbb{E}[\varepsilon_t^8] < \infty$ in order to satisfy the assumptions in Rilstone et al. (1996). In this section, for ease of exposition, we assume $E[\varepsilon_t] = 0$, but this is without loss of generality. Notably, typical HAC estimators of the standard error of $\widehat{\beta}_{h,LS}$ are functions only of $\gamma_{h,u}$, the sample size, and a bandwidth parameter.

When researchers calculate $\gamma_{h,u}$ (for example, when computing Newey-West standard errors), they typically use plug-in estimators derived from the empirical regression scores. To understand how this procedure affects the small-sample properties of $\widehat{\gamma}_{h,u}$, it is useful to think about estimating the regression coefficients and $\gamma_{h,u}$ jointly. In this case

$$q_t = \begin{bmatrix} (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) \\ \varepsilon_t (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) \\ \varepsilon_t (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) \varepsilon_{t-u} (y_{t+h-u} - \alpha_h - \varepsilon_{t-u} \beta_h) - \gamma_{h,u} \end{bmatrix}.$$

Unlike our previous derivations, because of the $\beta_h \beta_h'$ term in q_t , the matrix H_2 is not a matrix of zeros.

$$\mathbb{E}(\nabla q_t) = \begin{bmatrix} -1 & 0 & 0 \\ 0 & -\sigma_\varepsilon^2 & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

$$Q = \begin{bmatrix} -1 & 0 & 0 \\ 0 & -\frac{1}{\sigma_\varepsilon^2} & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

$$Q\psi_{T-h-u} = -\frac{1}{T-h-u} \sum_{t=u+1}^{T-h} \begin{bmatrix} (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) \\ \frac{\varepsilon_t}{\sigma_\varepsilon^2} (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) \\ \varepsilon_t (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) \varepsilon_{t-u} (y_{t+h-u} - \alpha_h - \varepsilon_{t-u} \beta_h) - \gamma_{h,u} \end{bmatrix}$$

$$\begin{aligned}
QVQ\psi_{T-h-u} &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sigma_\varepsilon^2} & 0 \\ 0 & 0 & 1 \end{bmatrix} \times \frac{1}{T-h-u} \sum_{t=1+u}^{T-h} \nabla q_t \\
&\times \frac{1}{T-h-u} \sum_{s=1+u}^{T-h} \begin{bmatrix} (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \\ \frac{\varepsilon_s}{\sigma_\varepsilon^2} (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \\ \varepsilon_s (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \varepsilon_{s-u} (y_{s+h-u} - \alpha_h - \varepsilon_{s-u} \beta_h) - \gamma_{h,u} \end{bmatrix}
\end{aligned}$$

We care about the third row of $QVQ\psi_{T-h-u}$. Fix s and t , the terms in the third row are of the form

$$\begin{aligned}
&[-\varepsilon_t \varepsilon_{t-u} (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) - \varepsilon_t \varepsilon_{t-u} (y_{t+h-u} - \alpha_h - \varepsilon_{t-u} \beta_h)] (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \\
&+ [-\varepsilon_t \varepsilon_{t-u}^2 (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) - \varepsilon_t^2 \varepsilon_{t-u} (y_{t+h-u} - \alpha_h - \varepsilon_{t-u} \beta_h)] \frac{\varepsilon_s}{\sigma_\varepsilon^2} (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \\
&- \varepsilon_s (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \varepsilon_{s-u} (y_{s+h-u} - \alpha_h - \varepsilon_{s-u} \beta_h) + \gamma_{h,u}
\end{aligned}$$

Clearly, $\mathbb{E}[-\varepsilon_s (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \varepsilon_{s-u} (y_{s+h-u} - \alpha_h - \varepsilon_{s-u} \beta_h) + \gamma_{h,u}] = 0$ for all s . Consider

$$-[\varepsilon_t \varepsilon_{t-u} (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) + \varepsilon_t \varepsilon_{t-u} (y_{t+h-u} - \alpha_h - \varepsilon_{t-u} \beta_h)] (y_{s+h} - \alpha_h - \varepsilon_s \beta_h)$$

If $u > 0$ and $\varepsilon_t \perp y_s$, then this is zero in expectation for all t and s . If $u = 0$, then the expectation is

$$-2\sigma_\varepsilon^2 \mathbb{E}[(y_t - \mu_y)^2]$$

Consider

$$-[\varepsilon_t \varepsilon_{t-u}^2 (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) + \varepsilon_t^2 \varepsilon_{t-u} (y_{t+h-u} - \alpha_h - \varepsilon_{t-u} \beta_h)] \frac{\varepsilon_s}{\sigma_\varepsilon^2} (y_{s+h} - \alpha_h - \varepsilon_s \beta_h)$$

If $u > 0$ and $\varepsilon_t \perp y_s$, then this is zero in expectation if $s \neq t$ and $s \neq t - u$. If $s = t$ or $s = t - u$, then this equals

$$-\sigma_\varepsilon^2 \mathbb{E}(y_t - \mu_y).$$

If $u = 0$, then the expectation is

$$-\mathbb{E} \left\{ [2\varepsilon_t^3 (y_{t+h} - \alpha_h - \varepsilon_t \beta_h)] \frac{\varepsilon_s}{\sigma_\varepsilon^2} (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \right\}.$$

If $s \neq t$ and $\varepsilon_t \perp y_s$, then the expectation is zero. If $s = t$, the the expectation is

$$-2 \frac{\mathbb{E}(\varepsilon_t^4)}{\sigma_\varepsilon^2} \mathbb{E}(y_t - \mu_y).$$

Note that

$$\nabla^2 q_t = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 2\varepsilon_t \varepsilon_{t-u} & \varepsilon_t \varepsilon_{t-u}^2 + \varepsilon_t^2 \varepsilon_{t-u} & 0 & \varepsilon_t^2 \varepsilon_{t-u} + \varepsilon_{t-u}^2 \varepsilon_t & 2\varepsilon_t^2 \varepsilon_{t-u}^2 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$\bar{H}_2 = E(\nabla^2 q_t) = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2(\sigma_\varepsilon^2)^2 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$-\frac{1}{2}Q\bar{H}_2 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & (\sigma_\varepsilon^2)^2 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

To compute

$$-\frac{1}{2}Q\bar{H}_2E[Q\psi_{T-h-u} \otimes Q\psi_{T-h-u}]$$

we need to evaluate

$$\mathbb{E} \left[\frac{1}{T-h-u} \sum_{t=1+u}^{T-h} \frac{\varepsilon_t}{\sigma_\varepsilon^2} (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) \frac{1}{T-h-u} \sum_{s=1+u}^{T-h} \frac{\varepsilon_s}{\sigma_\varepsilon^2} (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \right].$$

If $t \neq s$ and $\varepsilon_t \perp y_s$, then

$$\mathbb{E} \left[\frac{\varepsilon_t}{\sigma_\varepsilon^2} (y_{t+h} - \alpha_h - \varepsilon_t \beta_h) \frac{\varepsilon_s}{\sigma_\varepsilon^2} (y_{s+h} - \alpha_h - \varepsilon_s \beta_h) \right] = 0.$$

If $t = s$ and $\varepsilon_t \perp y_s$, then

$$\mathbb{E} \left[\frac{\varepsilon_t^2}{(\sigma_\varepsilon^2)^2} (y_{t+h} - \alpha_h)^2 \right] = E \left[\frac{1}{\sigma_\varepsilon^2} (y_{t+h} - \alpha_h)^2 \right].$$

Putting all of this together, if $u > 0$, the bias is

$$-\frac{1}{T-h-u} \frac{T-h-2u}{T-h-u} \sigma_\varepsilon^2 E \left[(y_t - \mu_y)^2 \right].$$

If $u = 0$, the bias is

$$-\frac{1}{T-h-u} 2 \frac{E(\varepsilon_t^4)}{\sigma_\varepsilon^2} E(y_t - \mu_y) - \frac{1}{T-h-u} \sigma_\varepsilon^2 E \left[(y_t - \mu_y)^2 \right]$$

In the case of normal variation, this is

$$-\frac{1}{T-h-u} 7\sigma_\varepsilon^2 E(y_t - \mu_y)$$

Notice that under the null hypothesis,

$$\gamma_{h,0} = \sigma_\varepsilon^2 \mathbb{E} \left[(y_t - \mu_y)^2 \right].$$

Under the null hypothesis that $\varepsilon_t \perp y_s$ and under normal variation,

$$\mathbb{E}(\hat{\gamma}_{h,0}) = \gamma_{h,0} \left(1 - \frac{7}{T-h} \right) + O(T^{-3/2})$$

$$\mathbb{E}(\hat{\gamma}_{h,u}) = \gamma_{h,u} - \frac{1}{T-h} \gamma_{h,0} + O(T^{-3/2}).$$

Clearly, when T is small, these distortions can be substantial, which explains why increasing the bandwidth of a Newey-West estimator makes the standard errors even smaller in expectation.

C Additional figures for the AR(1) example

In this appendix, we show figures analogous to those in the text for different values of ρ in the AR(1) example. In the main text, we set $\rho = 0.95$. Here, we consider $\rho = 0.9$ and $\rho = 0.99$.

C.1 $\rho = 0.9$

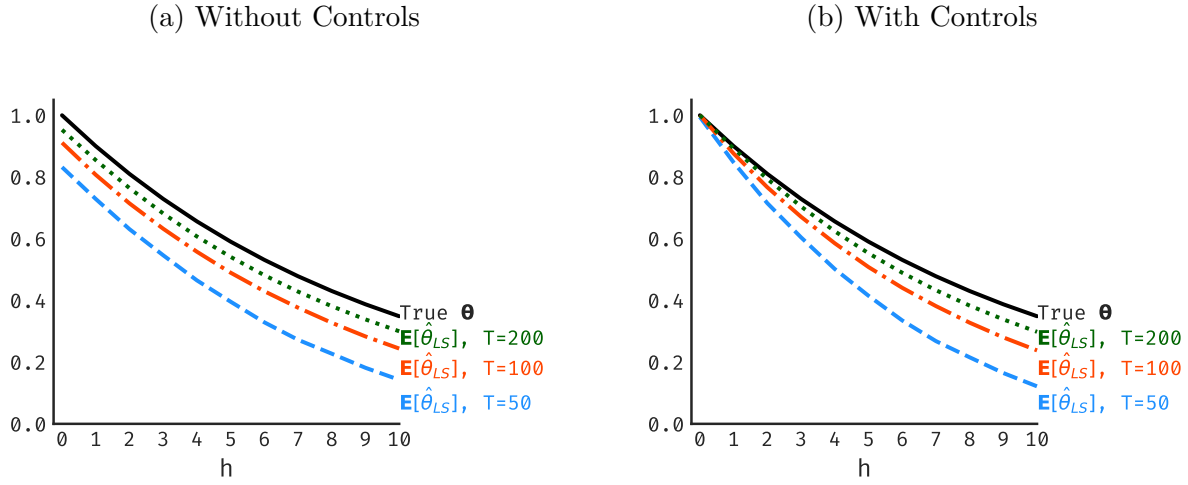


Figure C.1: LP estimators are biased in empirically-relevant samples when y_t is an AR(1) with $\rho = 0.90$.

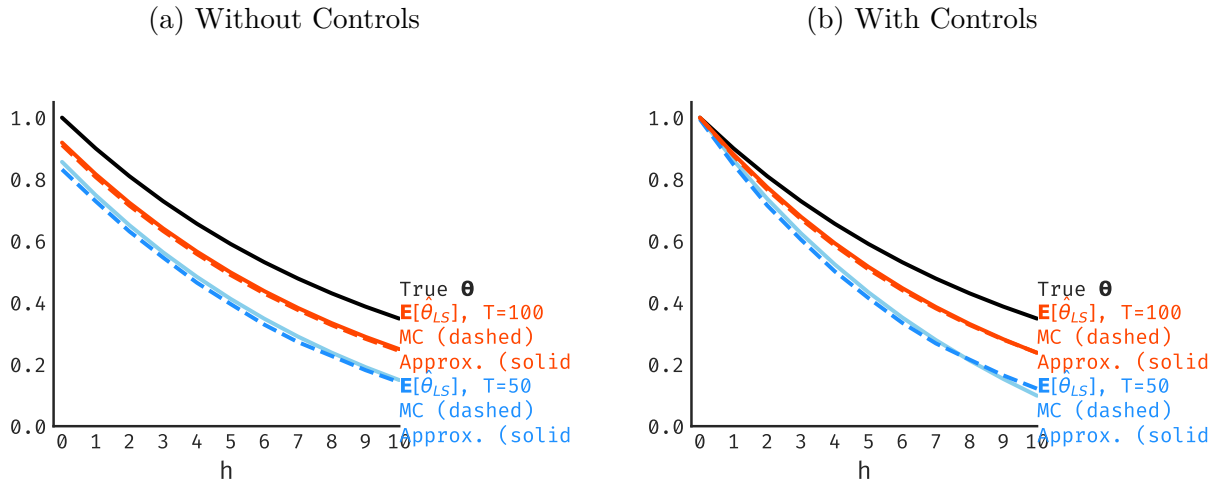


Figure C.2: The bias approximation is accurate in our LPs.




Figure C.3: $\hat{\theta}_{BC}$ and $\hat{\theta}_{BCC}$ are closer than $\hat{\theta}_{LS}$ to θ , on average, in our LPs without controls when y_t is an AR(1) with $\rho = 0.90$.




Figure C.4: $\hat{\theta}_{BC}$ and $\hat{\theta}_{BCC}$ are closer than $\hat{\theta}_{LS}$ to θ , on average, in our LPs with controls when y_t is an AR(1) with $\rho = 0.90$.

C.2 $\rho = 0.99$




Figure C.5: LP estimators are biased in empirically-relevant samples when y_t is an AR(1) with $\rho = 0.99$.

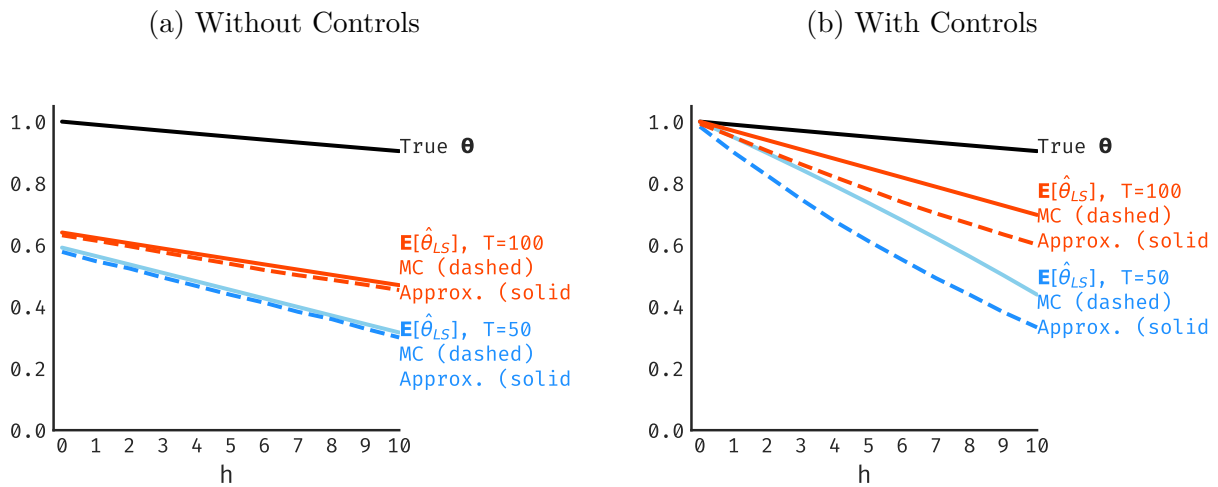


Figure C.6: The bias approximation is accurate in our LPs.

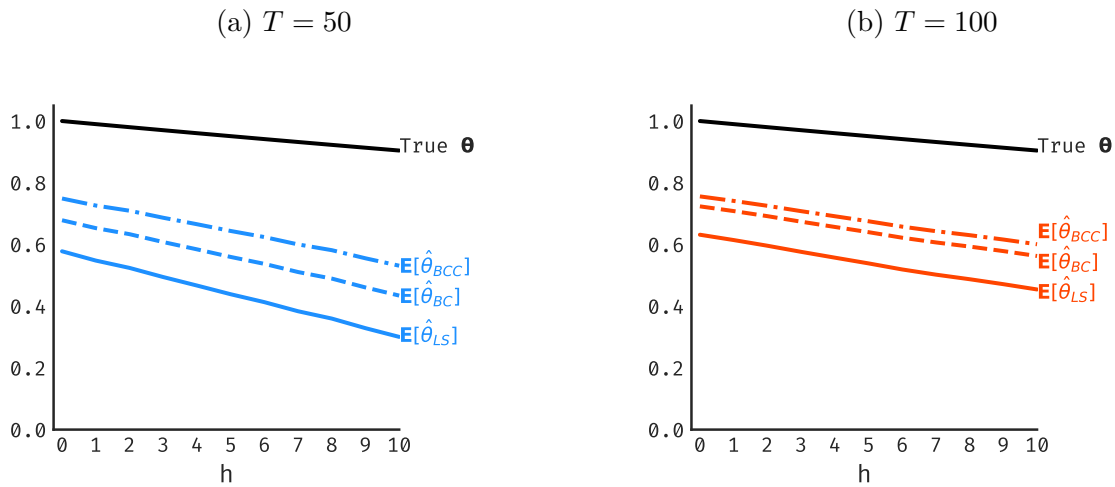


Figure C.7: $\hat{\theta}_{BC}$ and $\hat{\theta}_{BCC}$ are closer than $\hat{\theta}_{LS}$ to θ , on average, in our LPs without controls when y_t is an AR(1) with $\rho = 0.99$.

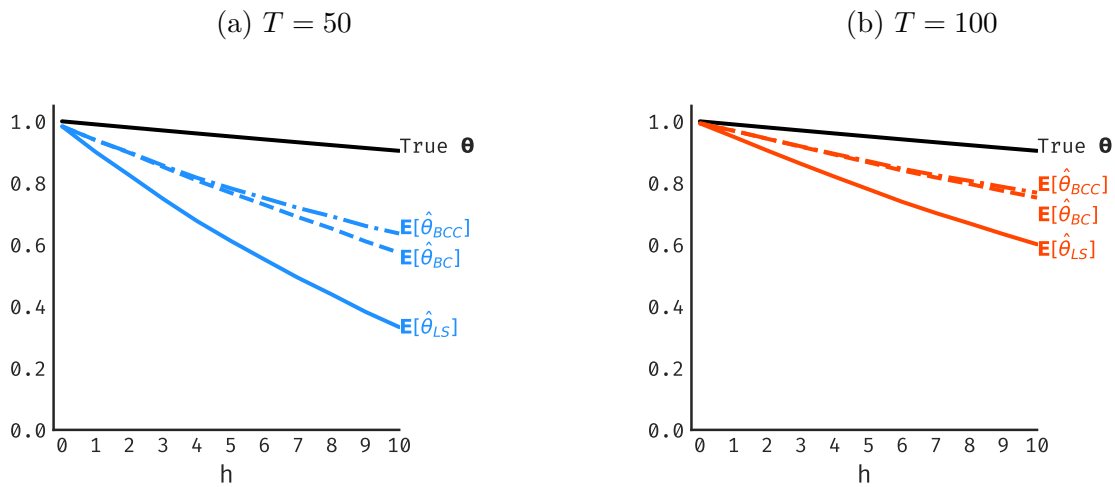


Figure C.8: $\hat{\theta}_{BC}$ and $\hat{\theta}_{BCC}$ are closer than $\hat{\theta}_{LS}$ to θ , on average, in our LPs with controls when y_t is an AR(1) with $\rho = 0.99$.

D An AR(2) example

In this appendix, we show figures analogous to those in the text for an AR(2) example. We specify the example so that

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

This process delivers hump-shaped impulse response functions. We set $\rho = 0.95$ and $\psi = 0.4$. When we include controls, we set $c_t = [y_{t-1}, y_{t-2}]'$.




Figure D.1: LP estimators are biased in empirically-relevant samples when y_t is an AR(2) with $\rho = 0.95$ and $\psi = 0.4$.

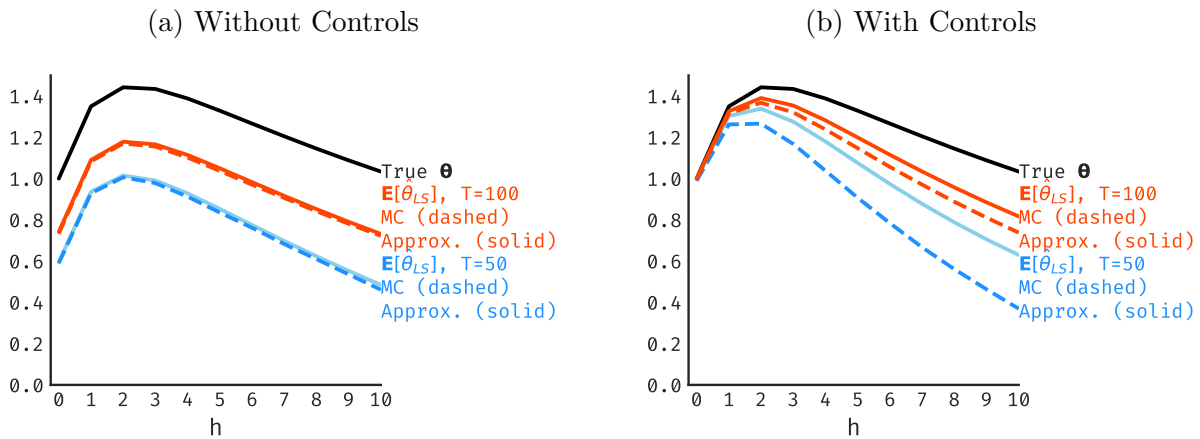


Figure D.2: The bias approximation is accurate in our LPs.

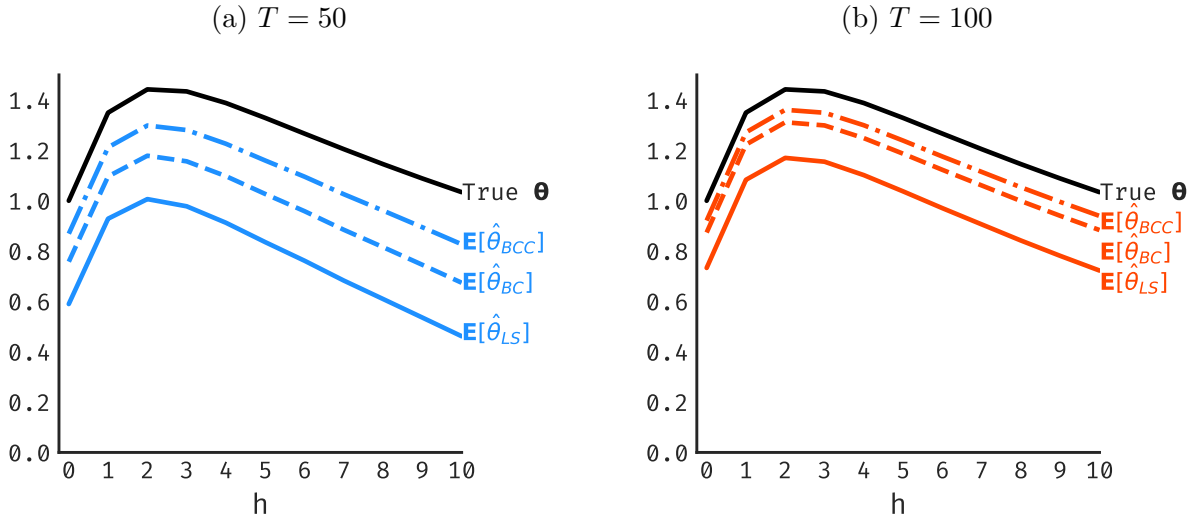


Figure D.3: $\hat{\theta}_{BC}$ and $\hat{\theta}_{BCC}$ are closer than $\hat{\theta}_{LS}$ to θ , on average, in our LPs without controls when y_t is an AR(1) with $\rho = 0.95$ and $\psi = 0.4$.

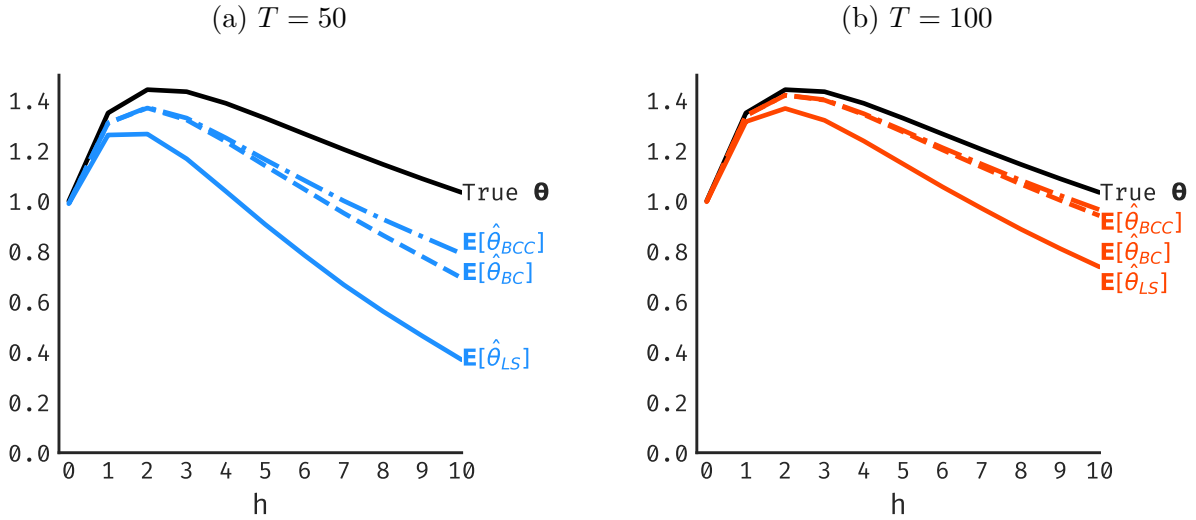


Figure D.4: $\hat{\theta}_{BC}$ and $\hat{\theta}_{BCC}$ are closer than $\hat{\theta}_{LS}$ to θ , on average, in our LPs with controls when y_t is an AR(2) with $\rho = 0.90$ and $\psi = 0.4$.

E Newey-West standard errors and fixed-b asymptotics

Sun (2014) and Lazarus et al. (2018) suggest that researchers should use fixed-b asymptotics when conducting inference using HAR estimators. For the Newey-West estimator, they suggest using a bandwidth of $1.3\sqrt{T-h}$. The asymptotic limiting distribution of the test statistic is not standard.

Table 1 shows results that are analogous to the results shown in Table 1, but using the bandwidth for the Newey-West estimator suggested by Lazarus et al. (2018) and the non-standard critical values to construct confidence sets using the Newey-West standard error. When one uses the fixed-b critical values, the performance of the Newey-West estimator improves somewhat, though the coverage probabilities are not uniformly better than those from the Huber-White standard errors.

Fixed-b asymptotics involve using a larger bandwidth for the Newey-West estimator and larger critical values than those implied by the asymptotic normal approximation. It turns out that the bandwidth is not that much larger in sample sizes typically found in the literature. As a result, the bias is not that much larger when the larger bandwidth is used. It is then not surprising that using a larger critical value improves the coverage probabilities given that the confidence intervals are too small.

Table 1: Coverage probability of different estimators of standard errors for $\hat{\theta}_h$ in LP without controls when y_t is an AR(1) with $\rho = 0.95$ and $T = 50$ using fixed-b critical values

| h | $\hat{\theta}_{h,LS}$, no controls | | $\hat{\theta}_{h,BCC}$, no controls | | $\hat{\theta}_{h,BCC}$, controls | |
|----|-------------------------------------|------------|--------------------------------------|------------|-----------------------------------|------------|
| | Huber-White | Newey-West | Huber-White | Newey-West | Huber-White | Newey-West |
| 0 | 0.87 | 0.82 | 0.86 | 0.85 | 0.92 | 0.93 |
| 1 | 0.83 | 0.81 | 0.82 | 0.83 | 0.90 | 0.90 |
| 2 | 0.80 | 0.80 | 0.79 | 0.82 | 0.87 | 0.89 |
| 3 | 0.78 | 0.77 | 0.76 | 0.79 | 0.85 | 0.86 |
| 4 | 0.76 | 0.76 | 0.75 | 0.78 | 0.83 | 0.85 |
| 5 | 0.75 | 0.75 | 0.74 | 0.78 | 0.81 | 0.83 |
| 6 | 0.75 | 0.75 | 0.73 | 0.77 | 0.80 | 0.82 |
| 7 | 0.74 | 0.73 | 0.73 | 0.76 | 0.78 | 0.82 |
| 8 | 0.74 | 0.74 | 0.73 | 0.77 | 0.77 | 0.82 |
| 9 | 0.74 | 0.75 | 0.73 | 0.77 | 0.76 | 0.81 |
| 10 | 0.74 | 0.75 | 0.73 | 0.78 | 0.75 | 0.81 |

F Bias and the block bootstrap

An alternative approach to achieve bias correction in LPs is through bootstrapping.¹ Bootstrap methods construct approximation to the distribution of an estimator (for example), by resampling observables or the errors from a parametric model. In a time series context, where the dynamic relationship between observables or errors is important to preserve, block bootstrapping techniques—in which the resampling scheme seeks to preserve some of the correlation in the original data set—are typically used, as in Kilian and Kim (2011). We revisit the Monte Carlo simulations in Section 3.1 and attempt to correct for the finite-sample bias in LPs using the block bootstrap.² Figure F.1 displays the bias correction from the block bootstrap when y_t is our AR(1) example.

Similar to results reported by Kilian and Kim (2011), the block bootstrap offers little in the way of LP bias correction. As in Kilian and Kim (2011), we set the block length to 4.³ Given equation (4), it is not surprising that with a short block length the block bootstrap does little to bias correct. The block bootstrap works by maintaining the autocorrelation structure of the data in LPs within a given block, but by destroying the autocorrelation across blocks. By destroying the autocorrelation across blocks, the block bootstrap destroys some of the autocorrelation information needed to adjust the estimates. As a result, the block bootstrap underestimates the bias in LPs, rendering bias correction based on the block bootstrap relatively ineffective.

Longer block lengths make the block bootstrap more effective at bias correcting, but they reduce the number of non-overlapping blocks in a dataset. Kilian and Kim (2011) report that longer block lengths lead to worse coverage probabilities.

¹See Hall (1992) for a textbook treatment.

²We use the same block length as Kilian and Kim (2011).

³Results are similar when we set the block length to $(T - h)^{1/3}$, as suggested by Hall et al. (1995).




Figure F.1: $\hat{\beta}_{Bootstrap}$ provides little bias correction in our LPs y_t is an AR(1) with $\rho = 0.95$.

G Lag augmentation in the AR(1) example

In this section, we consider the effects of lag augmentation in our AR(1) example. Olea and Plagborg-Møller (2020) argue that lag augmentation improves the performance of LPs. There are some important differences between the LP setting we consider and the setting in Olea and Plagborg-Møller (2020). First, we consider an LP in which a constant must be estimated, while Olea and Plagborg-Møller (2020) do not. Second, we consider LP estimates of impulse responses to identified shocks that that the researcher brings to the LP, rather than identifying the shock as a part of the LP system. Here, we consider the effects of lag augmentation in our setup. The main example in Olea and Plagborg-Møller (2020) is an AR(1) without a constant, so the example from our paper has few other differences than the two important differences identified above.

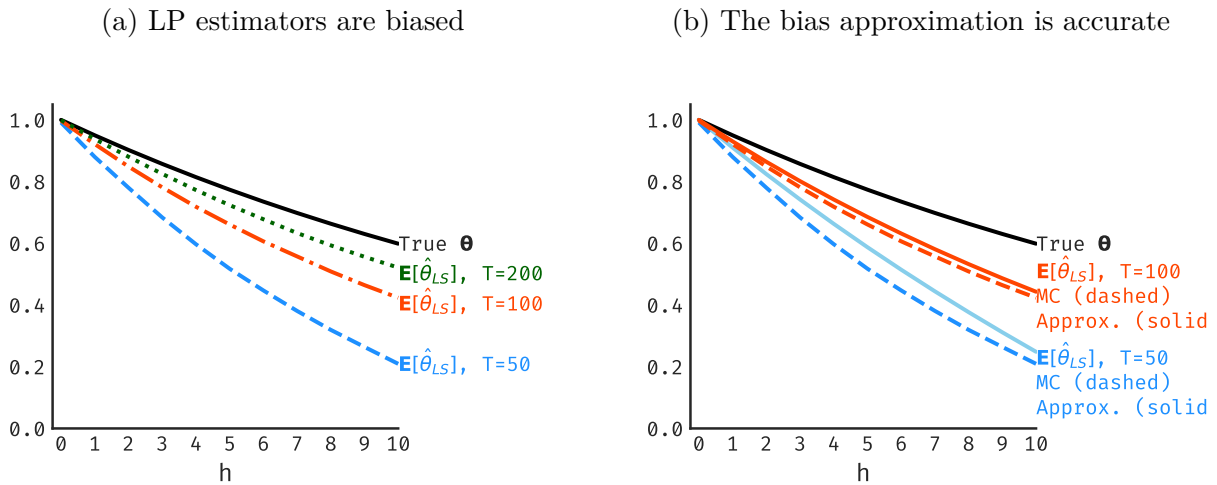


Figure G.1: LP with lag augmentation when y_t is an AR(1) with $\rho = 0.95$.

Figure G.1 shows the Monte Carlo mean of $\hat{\theta}_{LS}$ in our AR(1) example when $\rho = 0.95$. The bias in the impulse response estimator is essentially the same as without lag augmentation.

Figure G.2 shows the Monte Carlo mean of $\hat{\theta}_{LS}$, $\hat{\theta}_{BC}$, and $\hat{\theta}_{BCC}$ in our AR(1) example when $\rho = 0.95$. The mean bias correction is essentially the same as without lag augmentation.

Table 2 shows the coverage probabilities for confidence intervals constructed using Huber-White and Newey-West standard errors in our AR(1) example with lag augmentation when $\rho = 0.95$. The coverage probabilities are similar to their values without lag augmentation.

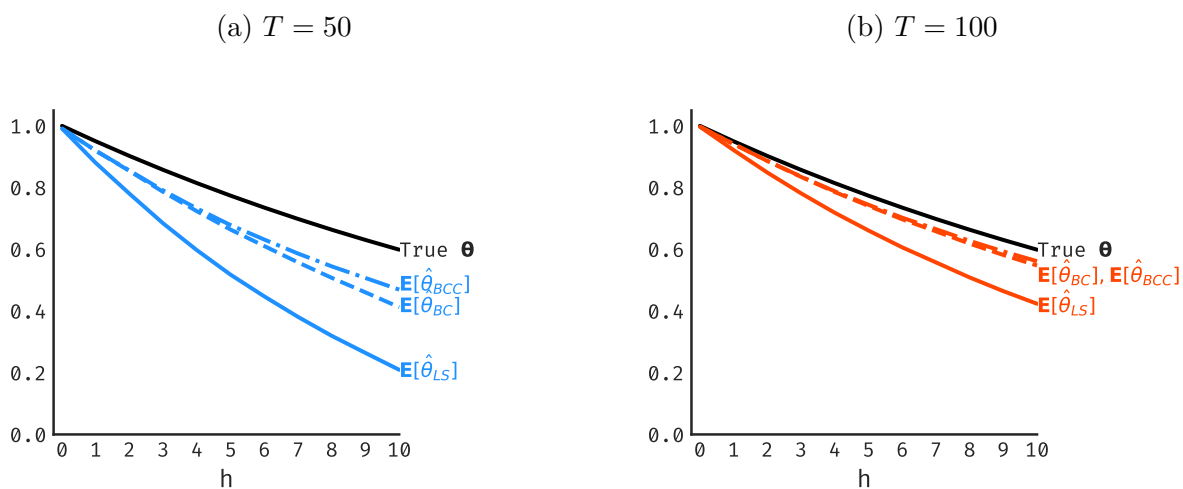


Figure G.2: Bias correction in an LP with lag augmentation when y_t is an AR(1) with $\rho = 0.95$.

Table 2: Coverage probability of different estimators of standard errors for $\hat{\theta}_h$ in LP with lag augmentation when y_t is an AR(1) with $\rho = 0.95$ and $T = 50$

| h | $\hat{\theta}_{h,LS}$ | | $\hat{\theta}_{h,BCC}$ | |
|----|-----------------------|------------|------------------------|------------|
| | Huber-White | Newey-West | Huber-White | Newey-West |
| 0 | 0.92 | 0.91 | 0.92 | 0.91 |
| 1 | 0.89 | 0.87 | 0.90 | 0.88 |
| 2 | 0.85 | 0.83 | 0.87 | 0.85 |
| 3 | 0.83 | 0.80 | 0.84 | 0.82 |
| 4 | 0.80 | 0.77 | 0.82 | 0.80 |
| 5 | 0.77 | 0.74 | 0.80 | 0.78 |
| 6 | 0.76 | 0.73 | 0.78 | 0.76 |
| 7 | 0.74 | 0.71 | 0.77 | 0.75 |
| 8 | 0.73 | 0.69 | 0.76 | 0.74 |
| 9 | 0.71 | 0.68 | 0.75 | 0.73 |
| 10 | 0.70 | 0.67 | 0.74 | 0.72 |

H Derivation of bias in an LP without controls using differences in y_t

Some authors use $y_{t+h} - y_{t-1}$ or $y_{t+h} - y_t$ as the left hand side variable for their LP. If controls are included and the researcher uses $y_{t+h} - y_{t-1}$ as the left-hand side variable, it has no effect on our earlier derivations. If $\theta_0 = 0$ and controls are included, it will also have no effect on our earlier derivation if the researcher uses $y_{t+h} - y_t$ as the left-hand-side variable. If controls are not included in the regression, then using the difference of y_t reduces the bias. Intuitively, using the difference of y_t removes much of the persistence in the dependent variable, and for near-unit-root processes is almost well-specified so that the regression errors are nearly MA(h+1) processes. Lunsford (2020) uses an LP without controls and $y_{t+h} - y_t$, so we focus on this setup. Given that Lunsford (2020) works with high-frequency monetary policy and forward guidance shocks, along with monthly macroeconomic data, it is reasonable to assume that $\theta_0 = 0$. The algebra for the case when $y_{t+h} - y_{t-1}$ is the left-hand-side variable is similar.

To show that using $y_{t+h} - y_t$ as the left-hand-side variable reduces bias, consider that the moment conditions associated with the OLS estimator are

$$\mathbb{E}[q_t] = 0$$

where

$$q_t \equiv \begin{bmatrix} y_{t+h} - y_t - \alpha_h - \beta_h \varepsilon_t \\ \varepsilon_t (y_{t+h} - y_t - \alpha_h - \beta_h \varepsilon_t) \end{bmatrix}.$$

We have that

$$\begin{aligned} \nabla^1 q_t &= \begin{bmatrix} -1 & -\varepsilon_t \\ -\varepsilon_t & -\varepsilon_t^2 \end{bmatrix} \\ H_1 &= \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} -1 & -\varepsilon_t \\ -\varepsilon_t & -\varepsilon_t^2 \end{bmatrix} \\ \bar{H}_1 &= \begin{bmatrix} -1 & -\mu_\varepsilon \\ -\mu_\varepsilon & -(\sigma_\varepsilon^2 + \mu_\varepsilon^2) \end{bmatrix} \\ \nabla^2 q_t = H_2 = \bar{H}_2 &= \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \\ V &= \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} 0 & -(\varepsilon_t - \mu_\varepsilon) \\ -(\varepsilon_t - \mu_\varepsilon) & -(\varepsilon_t^2 - (\sigma_\varepsilon^2 + \mu_\varepsilon^2)) \end{bmatrix} \\ Q &= \begin{bmatrix} -\left(1 + \frac{\mu_\varepsilon^2}{\sigma_\varepsilon^2}\right) & \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} \\ \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} & -\frac{1}{\sigma_\varepsilon^2} \end{bmatrix} \end{aligned}$$

Because \bar{H}_2 is a matrix of zeros,

$$B = \mathbb{E}\{QVQ\psi_{T-h}\}.$$

Note that

$$B = \mathbb{E} \{ Q (H_1 - \bar{H}_1) Q \psi_{T-h} \} = \mathbb{E} \{ Q H_1 Q \psi_{T-h} \}.$$

Also,

$$\begin{aligned} QH_1 &= \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} -\left(1 + \frac{\mu_\varepsilon^2}{\sigma_\varepsilon^2}\right) & \frac{\mu_\varepsilon}{\sigma_\varepsilon} \\ \frac{\mu_\varepsilon}{\sigma_\varepsilon} & -\frac{1}{\sigma_\varepsilon^2} \end{bmatrix} \begin{bmatrix} -1 & -\varepsilon_t \\ -\varepsilon_t & -\varepsilon_t^2 \end{bmatrix} \\ &= \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} \left(1 - \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon)\right) & \left(1 - \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon)\right) \varepsilon_t \\ \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) & \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) \varepsilon_t \end{bmatrix} \end{aligned}$$

and

$$\begin{aligned} QH_1Q &= \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} \left(1 - \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon)\right) & \left(1 - \frac{\mu_\varepsilon}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon)\right) \varepsilon_t \\ \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) & \frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) \varepsilon_t \end{bmatrix} \begin{bmatrix} -\left(1 + \frac{\mu_\varepsilon^2}{\sigma_\varepsilon^2}\right) & \frac{\mu_\varepsilon}{\sigma_\varepsilon} \\ \frac{\mu_\varepsilon}{\sigma_\varepsilon} & -\frac{1}{\sigma_\varepsilon^2} \end{bmatrix} \\ &= \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} -\left(1 - 2\frac{\mu_\varepsilon}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) + \left(\frac{\mu_\varepsilon}{\sigma_\varepsilon^2}\right)^2 (\varepsilon_t - \mu_\varepsilon)^2\right) & -\left(\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) - \frac{\mu_\varepsilon}{(\sigma_\varepsilon^2)^2} (\varepsilon_t - \mu_\varepsilon)^2\right) \\ -\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) + \frac{\mu_\varepsilon}{(\sigma_\varepsilon^2)^2} (\varepsilon_t - \mu_\varepsilon)^2 & -\frac{1}{(\sigma_\varepsilon^2)^2} (\varepsilon_t - \mu_\varepsilon)^2 \end{bmatrix} \end{aligned}$$

Then

$$\begin{aligned} QH_1Q\psi_{T-h} &= \frac{1}{T-h} \sum_{t=1}^{T-h} \begin{bmatrix} -\left(1 - 2\frac{\mu_\varepsilon}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) + \left(\frac{\mu_\varepsilon}{\sigma_\varepsilon^2}\right)^2 (\varepsilon_t - \mu_\varepsilon)^2\right) & -\left(\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) - \frac{\mu_\varepsilon}{(\sigma_\varepsilon^2)^2} (\varepsilon_t - \mu_\varepsilon)^2\right) \\ -\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) + \frac{\mu_\varepsilon}{(\sigma_\varepsilon^2)^2} (\varepsilon_t - \mu_\varepsilon)^2 & -\frac{1}{(\sigma_\varepsilon^2)^2} (\varepsilon_t - \mu_\varepsilon)^2 \end{bmatrix} \\ &\quad \times \frac{1}{T-h} \sum_{s=1}^{T-h} \begin{bmatrix} y_{s+h} - y_s - \alpha_h - \beta_h \varepsilon_s \\ \varepsilon_s (y_{s+h} - y_s - \alpha_h - \beta_h \varepsilon_s) \end{bmatrix} \end{aligned}$$

We only need the expectation of the second row of this expression. To calculate it, fix t and consider

$$\begin{aligned} \mathbb{E} \left\{ -\left(\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) - \frac{\mu_\varepsilon}{(\sigma_\varepsilon^2)^2} (\varepsilon_t - \mu_\varepsilon)^2\right) (y_{s+h} - y_s - \alpha_h - \beta_h \varepsilon_s) - \frac{1}{(\sigma_\varepsilon^2)^2} (\varepsilon_t - \mu_\varepsilon)^2 \varepsilon_s (y_{s+h} - y_s - \alpha_h - \beta_h \varepsilon_s) \right\} &= \\ \mathbb{E} \left\{ -\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) (y_{s+h} - y_s - \alpha_h - \beta_h \varepsilon_s) - \frac{1}{(\sigma_\varepsilon^2)^2} (\varepsilon_t - \mu_\varepsilon)^2 (\varepsilon_s - \mu_\varepsilon) (y_{s+h} - y_s - \alpha_h - \beta_h \varepsilon_s) \right\} &= \\ \mathbb{E} \left\{ -\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) (y_{s+h} - y_s - \alpha_h - \beta_h \varepsilon_s) \right\} &= \\ \mathbb{E} \left\{ -\frac{1}{\sigma_\varepsilon^2} (\varepsilon_t - \mu_\varepsilon) (y_{s+h} - y_t + y_t - y_s - \alpha_h - \beta_h \varepsilon_s) \right\} & \end{aligned}$$

If $s = t$, or $s < t - h$, then this expression is zero. If $t - h \leq s < t$, then this expression is equal to $-\beta_{s+h-t}$.

If $t < s$, this expression is equal to

$$-\beta_{s+h-t} + \beta_{s-t}$$

The existence of the second term mitigates the bias relative to the case where the difference is not taken.

I 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. PAPP (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.