Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

Bias in Local Projections

Edward P. Herbst and Benjamin K. Johannsen

2020-010

Please cite this paper as:
Herbst, Edward P., and Benjamin K. Johannsen (2020). "Bias in Local Projections," Finance and Economics Discussion Series 2020-010. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2020.010.

Bias in Local Projections

Edward P. Herbst and Benjamin K. Johannsen*

January 13, 2020

Abstract

Local projections (LPs) are a popular tool in applied macroeconomic research. We survey the related literature and find that LPs are often used with very small samples in the time dimension. With small sample sizes, given the high degree of persistence in most macroeconomic data, impulse responses estimated by LPs can be severely biased. This is true even if the right-hand-side variable in the LP is iid, or if the data set includes a large cross-section (i.e., panel data). We derive a simple expression to elucidate the source of the bias. Our expression highlights the interdependence between coefficients of LPs at different horizons. As a byproduct, we propose a way to bias-correct LPs. Using U.S. macroeconomic data and identified monetary policy shocks, we demonstrate that the bias correction can be large.

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 Jorda (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 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 Jorda (2005). In particular, we find that sample sizes in the time dimension are typically much smaller than the sample sizes studied in Jorda (2005) and that LPs have become increasingly prevalent when researchers have a cross section of data (panel data). Additionally, researchers often approach LPs with identified structural shocks in hand, rather than identifying those shocks as a part of the estimation.

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 benchmark for a flexible methodology for estimating impulse responses 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 Taylor-series expansion of the LP estimator. We show that the bias of the LP estimator at horizon h is function—specifically, a weighted average—of the entire (population) impulse response function. 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 expansion 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 all horizons.

We use our Taylor-series expansion to develop a bias-corrected estimator for LPs. In Monte Carlo simulations, our bias-corrected estimator is able to adjust LPs toward their true value, but not completely. These results imply that researchers 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)).

We also analyze LPs applied to panel data or with instrumental variables and show that the bias we document is also found in these settings. Additionally, we give examples where increasing the number of entities in panel data has absolutely no effect on the small sample bias of the LP estimator.

Our paper is related to work by Kilian and Kim (2011), who study the coverage probabilities for confidence intervals for LP estimators. Their work focuses on the case when shocks are identified as a part of the LP estimation. 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 Taylor-series approximation illustrates reasons that the LP estimator is biased, and we extend the analysis to panel data and instrumental variables settings. Additionally, we shed light on why the bootstrap used in Kilian and Kim (2011) performs poorly for the purposes of bias correcting LPs. 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)) and in dynamic panel data settings (such as Nickell (1981)). We extend this work to the LP setting.

Finally, in this paper we focus solely on point estimation, with the goal of documenting and understanding parameter bias in LPs. Of course, proper inference requires also characterizing the uncertainty around point estimators. In the interest of clarity, we omit this important issue. It is worth mentioning, however, that biased parameter estimates will typically be associated with biased estimates of standard errors, even using heteroskedasticity- and autocorrelation-consistent (HAC) corrections.

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 Jorda (2005) on Google Scholar.[1] 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).[2] 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

[1] We conducted this search in October 2019. See Appendix C for the list of citations.

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

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.

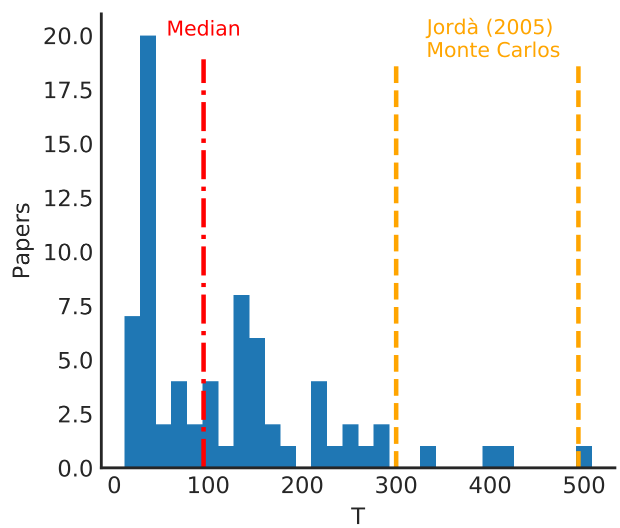

Figure 1: T is small in the literature using LPs.

Figure 1 displays a histogram of the sample of 71 Ts collected in our literature review. The median T (the red dash dotted line) is around 95. These sample sizes are notably less 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 Jorda (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 study in Jorda (2005) used $T = 300$

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 Jorda (2005) for not anticipating how researchers would subsequentyly apply LP meathods. While many studies in our survey use monthly or quarterly data, Jorda (2005) used monthly data. In general, increasing T by using monthly data rather than quarterly or annual data will not avoid the issue of small-sample bias in LPs because the monthly series are likely to be more persistent. In our simulation study, we show that 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 that are similar to 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 (partial) correction of the bias.

3.1 Bias in LPs using two examples

To demonstrate that LPs can be severely biased in small samples, we first consider two simple time series models: an $AR(1)$—the natural benchmark—and an $AR(2)$ that can capture hump-shaped dynamics commonly found in macroeconomics. In the context of these examples, our objective is to study the accuracy of estimated impulse responses via LPs for various sample sizes. We start with a Monte Carlo study. For a given T, we simulate $N_{mc} = 10,000$ time series, $\{y_t\}_{t=1}^{T}$, for each of two data generating processes:

$$y_t = \rho y_{t-1} + \epsilon_t + \nu_t \quad \text{and} \tag{1}$$

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

Here, ϵ_t and $\nu_t \sim iid\ \mathcal{N}(0,1)$. In these Monte Carlo simulations, to be consistent with the high persistence of macroeconomic data, we set $\rho = 0.99$. We set $\psi = 0.4$ to ensure a hump-shaped impulse response.[3] We use these simple linear time series models because LPs

[3]In our Appendix, 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.

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 LP. The researcher observes $\{y_t, \epsilon_t\}_{t=1}^T$; 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 like to control for other variables, denoted by the vector c_t. When we include such controls, we assume $c_t = y_{t-1}$ in the AR(1) case, and $c_t = [y_{t-1}, y_{t-2}]'$ in the AR(2) case. 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,

$$y_{t+h} = \alpha_h + \beta_h' x_t + u_{t,h}, \quad h = 0, \ldots, H. \tag{3}$$

where $x_t \equiv [\epsilon_t, c_t']'$.[4] 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 $h = 0, \ldots, H$. As in the empirical macroeconomics literature, we estimate each β_h using least squares. We denote the estimator of the β_h by $\widehat{\beta}_{h,LS}$ and the estimator of the impulse response by $\widehat{\theta}_{LS}$.

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

[4] When we do not include controls, $x_t = \epsilon_t$.

(a) Without Controls

(b) With Controls

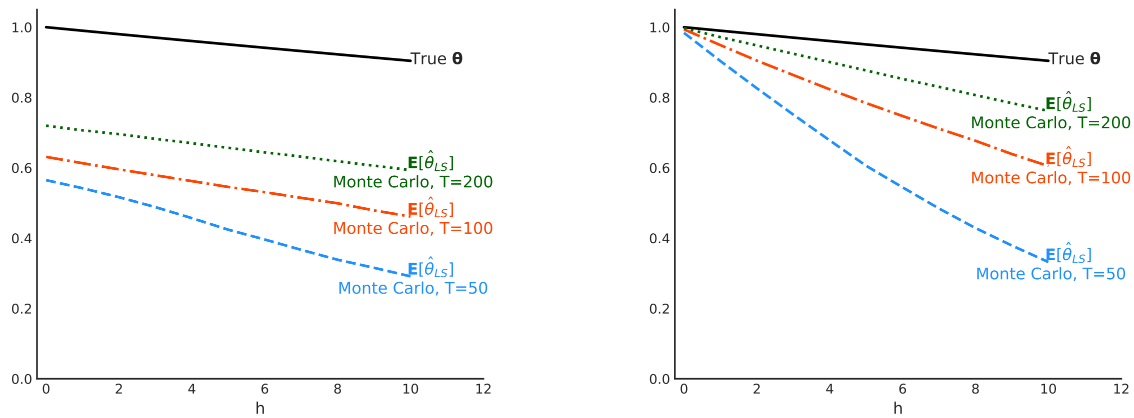

Figure 2: LP estimators without controls are biased in empirically-relevant samples when y_t is an AR(1) with $\rho = 0.99$.

horizons, the overlapping nature of the left-hand-side variables in the LP implies that the error terms are autocorrelated.[5]

Figure 3 displays $E[\widehat{\theta}_{LS}]$ for the $AR(2)$ model. The takeaways are similar to the AR(1) case. With no controls, LPs are severely biased even at short horizons. When controls are added, the bias diminishes are short horizons, but by $h = 10$ it is again severe.

3.2 Understanding bias in LPs

In this subsection we derive an analytic approximation to the bias of the least-squares estimators, $\widehat{\theta}_{LS}$, given by $\mathbb{E}[\widehat{\theta}_{LS}] - \theta$. This expression is, to the best of our knowledge, new to the literature, and highlights the interdependence of the bias in LPs at different horizons. In order to illustrate the point clearly (and to avoid tedious matrix algebra), we focus on LPs without controls. Our analytic approximation is easily generalizable to the case when controls are included in the LP, and the intuition for the bias is unchanged (we provide the

[5]Following Jorda (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.

(a) Without Controls

(b) With Controls

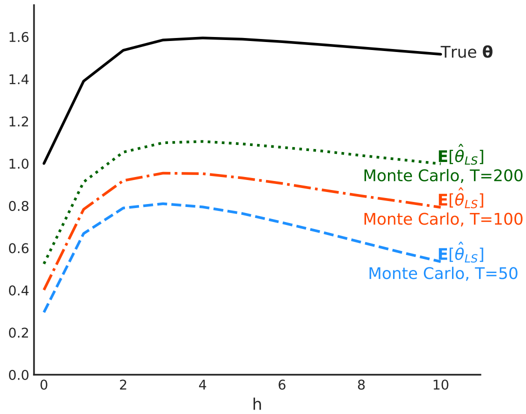

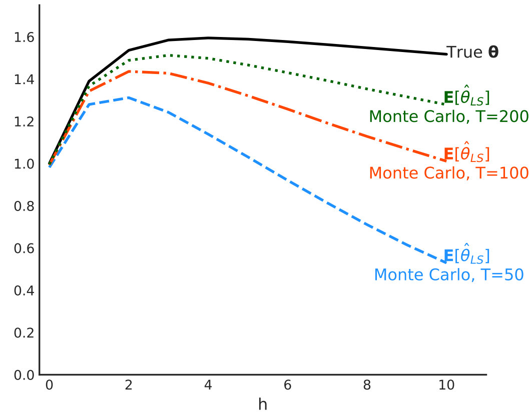

Figure 3: LP estimators without controls are biased in empirically-relevant samples when y_t is an AR(2) with $\rho = 0.99$ and $\psi = 0.4$.

associated derivations in our Appendix A). In all of our derivations, we assume that the data we use in the LP are strictly covariance stationary.

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

$$\widehat{\theta}_{h,LS} = \frac{\frac{1}{T-h}\sum_{t=1}^{T-h}\epsilon_t y_{t+h} - \frac{1}{(T-h)^2}\left(\sum_{t=1}^{T-h}\epsilon_t\right)\left(\sum_{t=1}^{T-h}y_{t+h}\right)}{\frac{1}{T-h}\sum_{t=1}^{T-h}\epsilon_t^2 - \frac{1}{(T-h)^2}\left(\sum_{t=1}^{T-h}\epsilon_t\right)^2} = \frac{\widehat{\text{cov}}[\epsilon_t, y_{t+h}]}{\widehat{\text{var}}[\epsilon_t]}. \tag{4}$$

Here, $\widehat{\text{cov}}$ and $\widehat{\text{var}}$ are the sample covariance and variance, respectively.[6] Equation (4) makes it clear that θ_h in population is the scaled covariance between y_{t+h} and ϵ_t.

To derive an expression for the approximate bias of $\widehat{\theta}_{h,LS}$, we need to compute $\mathbb{E}\left[\widehat{\theta}_{h,LS}\right]$. To do so, we take the expectation of a Taylor-series expansion of (4) around $\mathbb{E}[\widehat{\text{cov}}[\epsilon_t, y_{t+h}]]$ and $\mathbb{E}[\widehat{\text{var}}[\epsilon_t]]$ to yield

$$\mathbb{E}\left[\widehat{\theta}_{h,LS}\right] \approx \frac{\mathbb{E}[\widehat{\text{cov}}[\epsilon_t, y_{t+h}]]}{E[\widehat{\text{var}}[\epsilon_t]]}. \tag{5}$$

All other terms in the Taylor-series expansion are of order $O(T^{-1})$. Noting that $\theta_{h+j} = \text{cov}[\epsilon_t, y_{t+h+j}]/\text{var}[\epsilon_t]$, we calculate the numerator and denominator of equation (5) directly

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

to obtain[7]

$$\mathbb{E}\left[\widehat{\theta}_{h,LS}\right] - \theta_h \approx -\frac{1}{T-h-1} \sum_{j=1}^{T-h-1} \left(1 - \frac{j}{T-h}\right) (\theta_{h+j} + \theta_{h-j}) \,. \tag{6}$$

Equation (6) gives an approximate expression for the bias of $\widehat{\theta}_{h,LS}$ that consists of only the impulse response coefficients at different horizons. Given our maintained assumption that ϵ_t is uncorrelated with past values of y_t, $\theta_{-j} = 0$ for $j > 0$.

Several remarks regarding equation (6) are in order. First, as expected, the bias is a decreasing function of T; for fixed h, the least-squares estimator is asymptotically unbiased. Second, the bias of $\widehat{\theta}_{h,LS}$ is a function of the impulse response *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).

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). However, our approximation ignores some terms of the Taylor-series expansion that are of order $O(T^{-1})$, while the conventional approximation omits $O(T^{-2})$ terms. These $O(T^{-1})$ terms involve the variance and covariance between the numerator and denominator of (4). In principle, we could derive these terms and include them in our expression for the approximate bias of $\widehat{\theta}_{LS}$. In practice, equation (6) performs well without these additional terms.

[7] The expectation of the numerator of (5) is given by

$$\mathbb{E}[\widehat{\text{cov}}[y_{t+h}, \epsilon_t]] = \left(1 - \frac{1}{T-h}\right) \text{cov}[y_{t+h}, \epsilon_t] - \frac{1}{T-h} \sum_{j=1}^{T-h-1} \left(1 - \frac{j}{T-h}\right) (\text{cov}[y_{t+h+j}, \epsilon_t] + \text{cov}[y_{t+h-j}, \epsilon_t]) \,.$$

Here, cov is the true covariance. The expectation of the denominator of of (5) is

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

Here, var is the true variance.

(a) Without Controls

(b) With Controls

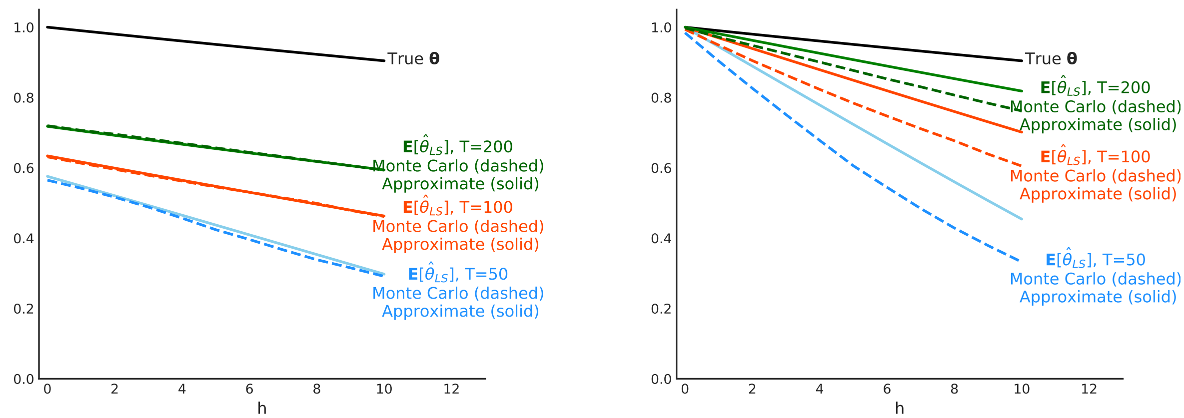

Figure 4: The first-order approximation is accurate in our LPs.

Figure 4 shows $\mathbb{E}\left[\widehat{\theta}_{LS}\right]$ calculated using the first-order 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

To derive a bias-corrected estimator of θ_h, we write the system of equations implied by (6) as

$$\mathbb{E}\left[\widehat{\theta}_{LS}\right] = \theta - M_{T,H}\theta, \tag{7}$$

where

$$M_{T,H} \equiv \begin{bmatrix} 0 & \left(1-\frac{1}{T}\right)\frac{1}{T-1} & \left(1-\frac{2}{T}\right)\frac{1}{T-1} & \dots & \left(1-\frac{H}{T}\right)\frac{1}{T-1} \\ \left(1-\frac{1}{T-1}\right)\frac{1}{T-2} & 0 & \left(1-\frac{1}{T-1}\right)\frac{1}{T-2} & \dots & \left(1-\frac{H+1}{T-1}\right)\frac{1}{T-2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \left(1-\frac{H}{T-H}\right)\frac{1}{T-H-1} & \left(1-\frac{H-1}{T-H}\right)\frac{1}{T-H-1} & \dots & \dots & 0 \end{bmatrix}. \tag{8}$$

The maximum eigenvalue of $M_{T,H}$ is less than one in modulus for any fixed T and $H < T-1$. Additionally, for a given H, $M_{T,H} \longrightarrow \mathcal{O}$, i.e, an $(H+1)\times(H+1)$ matrix of zeros, as $T \longrightarrow \infty$, which is a byproduct of the consistency of $\widehat{\theta}_{LS}$.

Given $\widehat{\theta}_{LS}$, we use (7) to construct our bias-corrected estimator of θ, which we denote $\widehat{\theta}_{BC}$, as

$$\widehat{\theta}_{BC} = (I - M_{T,H})^{-1}\,\widehat{\theta}_{LS}. \tag{9}$$

$\widehat{\theta}_{BC}$ is an easily implementable transformation of the LP estimators regularly reported in the empirical macroeconomics literature.

To implement our estimator, we have to choose a maximum horizon, $H < T-1$, which limits the horizon of the LPs that we estimate. The choice of H clearly has an influence on our estimator, and H must grow at a slower rate than T in order for the bias correction to be asympotitcally consistent. In a particular finite sample, the choice is limited by the length of the time series, rendering the asymptotic growth rate of H unimportant because any H could be consistent with any growth rate. Without theory on how to optimally pick H, we set H to 20, 25, and 50 for T equal to 50, 100, and 200, respectively.

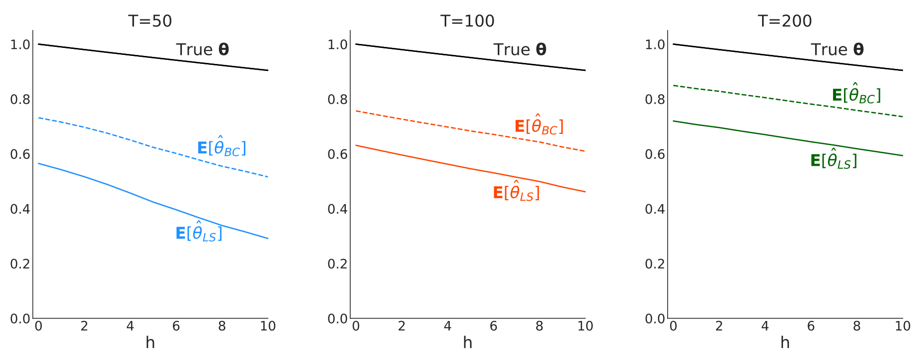

Figure 5: $\widehat{\beta}_{BC}$ is closer than $\widehat{\beta}_{LS}$ to β, on average, in our LPs without controls when y_t is an AR(1) with $\rho = 0.99$.

Figures 5 and 6 show the average value of $\widehat{\theta}_{BC}$ and $\widehat{\theta}_{LS}$ over our Monte Carlo simulations when y_t follows an AR(1) and an AR(2), respectively. Clearly, $\widehat{\theta}_{BC}$ does not completely correct for the bias in $\widehat{\theta}_{LS}$ in either case, indicating that our bias-corrected estimator is not a panacea for bias in LPs. Nevertheless, $\widehat{\theta}_{BC}$ is markedly closer than $\widehat{\theta}_{LS}$ to θ on average.

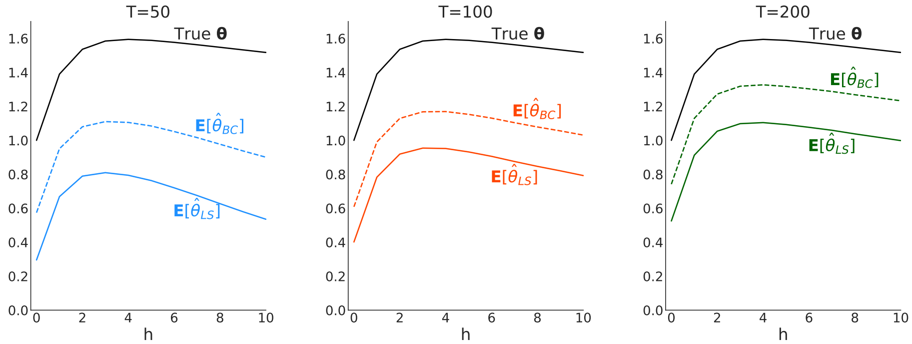

Figure 6: $\widehat{\beta}_{BC}$ is closer than $\widehat{\beta}_{LS}$ to β, on average, in our LPs without controls when y_t is an AR(2) with $\rho = 0.99$ and $\psi = 0.4$.

3.4 Understanding why the block bootstrap performs poorly

An alternative approach to achieve bias correction in LPs is through bootstrapping.[8] Bootstrap methods construct appoximation 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.[9] Figure 7 and 8 display the bias correction from the block bootstrap when y_t is an AR(1) and AR(2), respectively.

Similar to results reported by Kilian and Kim (2011), the block bootstrap offers little in the way of LP bias correction. Given equation (6), this result is not surprising. The block bootstrap works by maintaining the autocorrelation structure of the data in LPs within a give 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.

[8] See Hall (1992) for a textbook treatment.

[9] We use the same block length as Kilian and Kim (2011).

(a) Without Controls

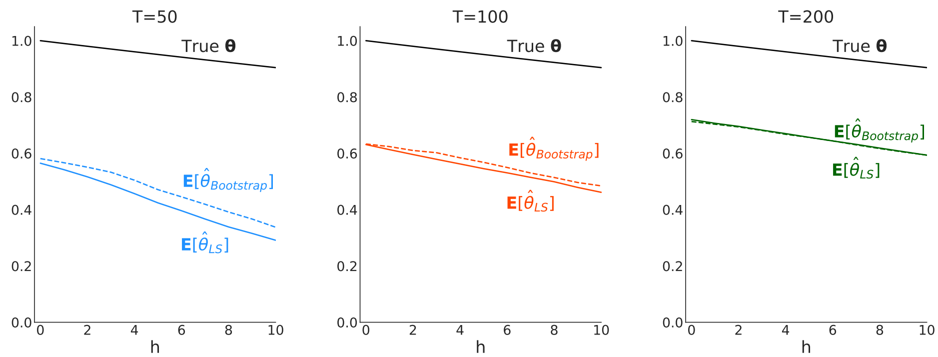

(b) With Controls

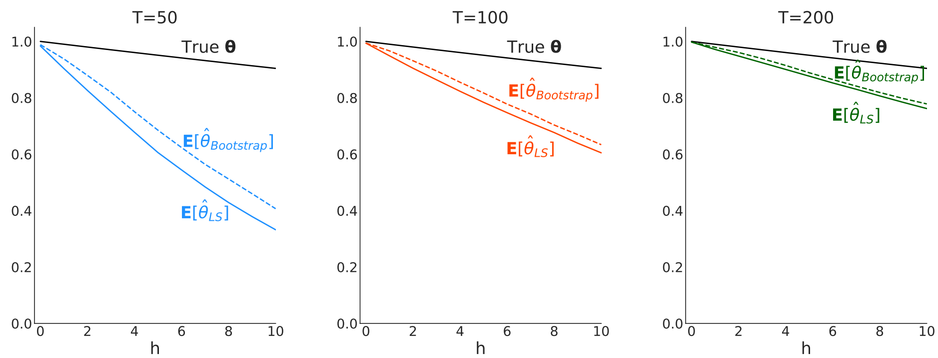

Figure 7: $\widehat{\beta}_{Bootstrap}$ provides little bias correction in our LPs y_t is an AR(1) with $\rho = 0.99$.

4 Bias in LPs with panel data

In this section, we demonstrate that LPs can be severely biased with sample sizes 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). Here, we illustrate the bias in LPs using the Taylor-series approximation of the ordinarly least-squares estimator that is similar to one in the previous section. As with the univariate case, the bias for an LP at horizon h is linked directly to the entire sequence of LP population coefficients. In all of our derivations, we assume that the data we use in the LP are strictly covariance stationary, and, for algebraic simplicity, we assume that the panel is balanced.

(a) Without Controls

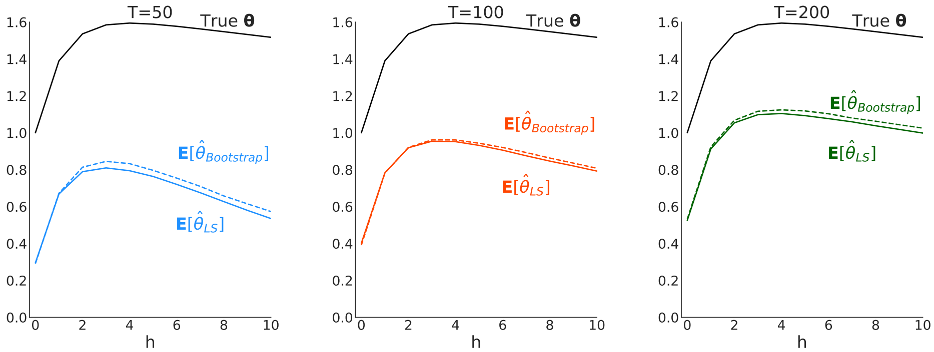

(b) With Controls

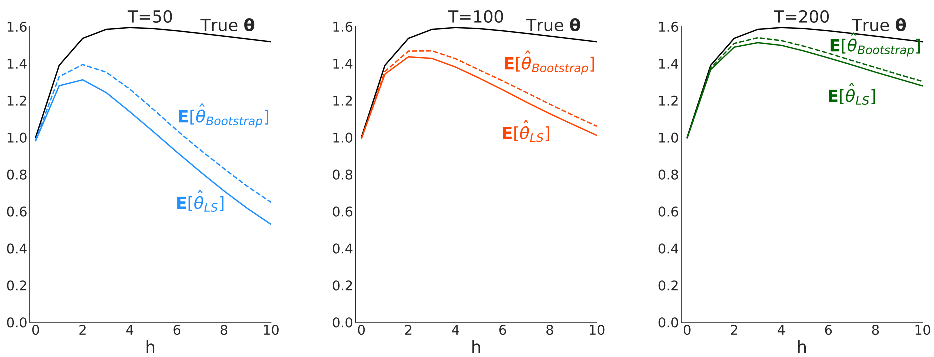

Figure 8: $\widehat{\beta}_{Bootstrap}$ provides little bias correction in our LPs y_t is an AR(2) with $\rho = 0.99$ and $\psi = 0.4$.

4.1 Bias in LPs with panel data using two examples

To demonstrate that LPs can be severely biased in small samples with panel data, we generate data, $\{y_{i,t}\}_{t=1}^{T}$ for eachentity $i = 1, \ldots, I$, using the data generating processes specified in equations (1) and (2). For simplicity, we assume that all the data are independent across entities, but the results generalize to a setting where the data are correlated across entities. We show results for panels containing $I = 10, 25,$ and 50 entities. As in the previous section we assume $\rho = 0.99$ and $\psi = 0.4$.

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

$$y_{i,t+h} = \alpha_{i,h} + \beta_h' x_{i,t} + u_{i,t,h}, \quad h = 0, \ldots, H. \tag{10}$$

where $x_{i,t} \equiv [\epsilon_{i,t}, c_{i,t}']'$.[4] As in the previous section, the first element of the coefficient vectors

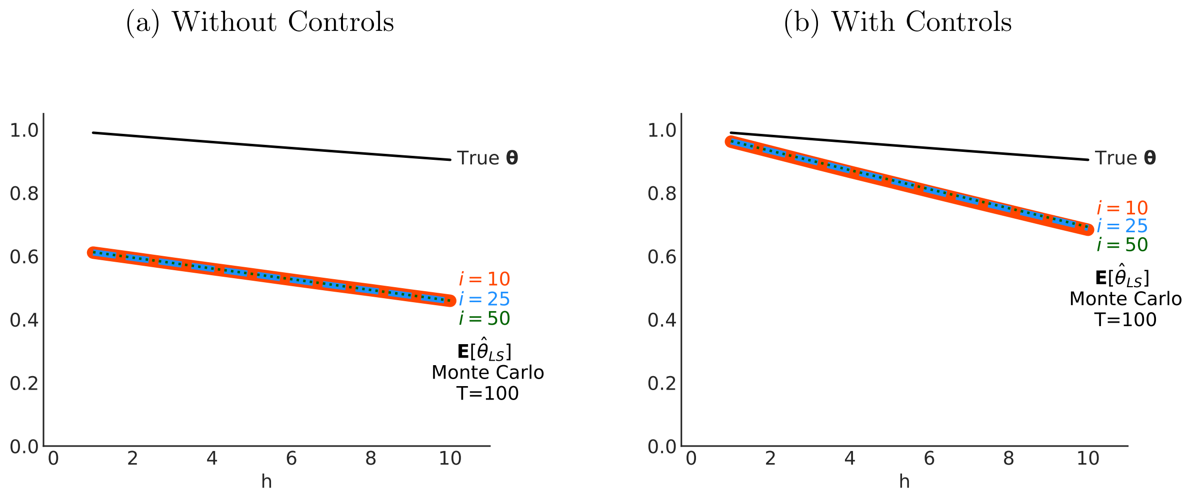

Figure 9: LP estimators without controls are biased in empirically-relevant samples when $y_{i,t}$ is an AR(1) with $\rho = 0.99$.

$\{\beta_h\}_{h=0}^{H}$ trace out the impulse response of interest, which we denote $\{\theta_h\}_{h=1}^{H}$.

Figure 9 displays the Monte-Carlo mean of LP estimators with and without controls for the AR(1) data generating process, $T = 100$, and different numbers of entitites in the panels. Figures 10 displays analogous results for the AR(2) data generating process. As was the case without panel data, for both the AR(1) and the AR(2) specifications, the LP estimates of the impulse responses are severely biased. Notably, the bias is independent of the number of entitites in the panel data set (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. Thus, our results indicate that LP bias can be very large for moderate values of h and relatively large values of T. In general, impulse responses are most of interest at moderate values of h.

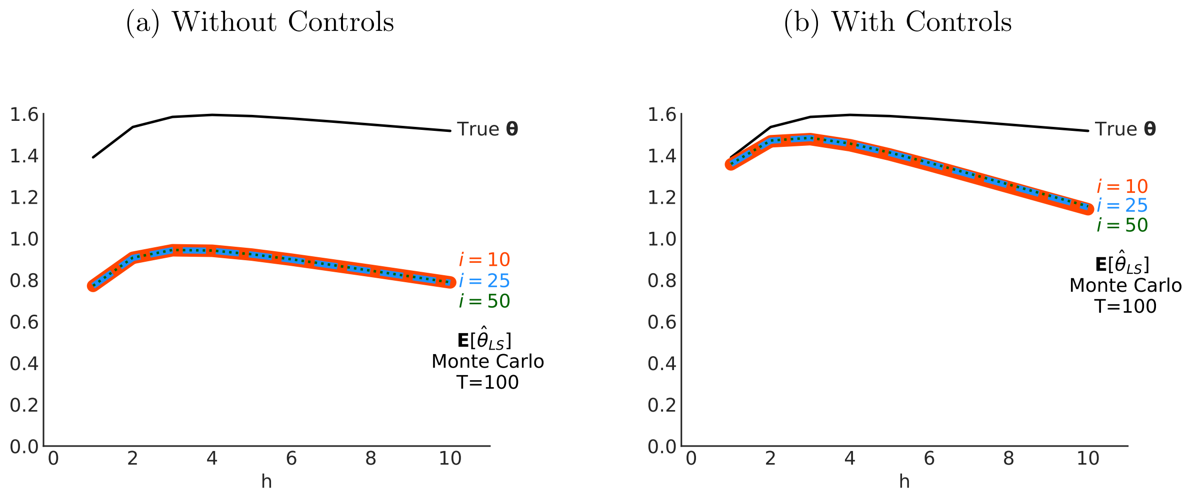

Figure 10: LP estimators without controls are biased in empirically-relevant samples when $y_{i,t}$ is an AR(2) with $\rho = 0.99$ and $\psi = 0.4$.

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. As in the univariate setting, for simplicity we focus on the case without controls so that the LP estimator coincides with the impulse response, given in population by $\{\theta_h\}_{h=0}^H$. The case with controls is analogous, though the notation becomes tedious when each panelist uses different controls.

For a given h, the ordinary least-squares estimator of θ_h with entity fixed effects is

$$\begin{aligned}\widehat{\theta}_{h,LS} &= \frac{\frac{1}{I}\sum_{i=1}^{I}\left[\frac{1}{T-h}\sum_{t=1}^{T-h}\epsilon_{i,t}y_{i,t+h} - \frac{1}{(T-h)^2}\left(\sum_{t=1}^{T-h}\epsilon_{i,t}\right)\left(\sum_{t=1}^{T-h}y_{i,t+h}\right)\right]}{\frac{1}{I}\sum_{i=1}^{I}\left[\frac{1}{T-h}\sum_{t=1}^{T-h}\epsilon_{i,t}^2 - \frac{1}{(T-h)^2}\left(\sum_{t=1}^{T-h}\epsilon_{i,t}\right)^2\right]} \\ &= \frac{\frac{1}{I}\sum_{i=1}^{I}\widehat{\text{cov}}[\epsilon_{i,t}, y_{i,t+h}]}{\frac{1}{I}\sum_{i=1}^{I}\widehat{\text{var}}[\epsilon_{i,t}]}. \qquad (11)\end{aligned}$$

As before, $\widehat{\text{cov}}$ and $\widehat{\text{var}}$ are the sample covariance and variance, respectively.

The expressions we derived for the sample variance and covariance in Section 3.2 are valid for $\widehat{\text{cov}}[\epsilon_{i,t}, y_{i,t+h}]$ and $\widehat{\text{var}}[\epsilon_{i,t}]$. The first-order Taylor series expansion from the previous section then implies

$$\mathbb{E}\left[\widehat{\theta}_{h,LS}\right]-\theta_h \approx -\frac{1}{I}\sum_{i=1}^{I}\frac{1}{T-h-1}\sum_{j=1}^{T-h-1}\left(1-\frac{j}{T-h}\right)(\theta_{h+j}+\theta_{h-j})$$
$$\approx -\frac{1}{T-h-1}\sum_{j=1}^{T-h-1}\left(1-\frac{j}{T-h}\right)(\theta_{h+j}+\theta_{h-j}). \quad (12)$$

Equation (12) is identical to (6), meaning that the expression for the bias in a panel setting is identical to the expression without panel data, which helps explain why the number of entities in the panel has no effect on the bias in our Monte Carlo simulations. As a result, the bias-corrected estimator from the previous section could be applied to panel data setup we analyze here, where T is used as the number of relevant observations rather than $I \times T$.

5 Bias in LPs with instrumental variables

A number of recent papers have employed LPs with instrumental variables (see, for example, Jordà et al. (2015), Ramey and Zubairy (2018), and Stock and Watson (2018)). In this section, we demonstrate that LPs with instrumental variables can also be biased with sample sizes commonly found in the empirical macroeconomic literature. As before, we assume that the data we use in the LP are strictly covariance stationary.

The LP model with instrumental variables is the set of regression models, indexed by the impulse response horizon h, given by equation (3). We assume that the researcher has access to a vector of valid instruments z_t and that z_t is the same length as x_t. For algebraic simplicity, we assume that $x_t = \epsilon_t$, z_t is some other variable that is a valid instrument, and that each variables is independent over time. As before, our results generalize to including additional variables and to allowing for general autocovariance structures.

For a given h, the instrumental-variables estimator of θ_h is

$$\widehat{\theta}_{h,IV} = \frac{\frac{1}{T-h}\sum_{t=1}^{T-h} z_t y_{t+h} - \frac{1}{(T-h)^2}\left(\sum_{t=1}^{T-h} z_t\right)\left(\sum_{t=1}^{T-h} y_{t+h}\right)}{\frac{1}{T-h}\sum_{t=1}^{T-h}\epsilon_t z_t - \frac{1}{(T-h)^2}\left(\sum_{t=1}^{T-h}\epsilon_t\right)\left(\sum_{t=1}^{T-h} z_t\right)} = \frac{\widehat{\text{cov}}[z_t, y_{t+h}]}{\widehat{\text{cov}}[z_t,\epsilon_t]}. \quad (13)$$

We take the expectation of a Taylor-series expansion of (13) around $\mathbb{E}[\widehat{\text{cov}}[z_t, y_{t+h}]]$ and $\mathbb{E}[\widehat{\text{cov}}[z_t,\epsilon_t]]$ to yield

$$\mathbb{E}\left[\widehat{\theta}_{h,IV}\right] \approx \frac{\mathbb{E}[\widehat{\text{cov}}[z_t, y_{t+h}]]}{E[\widehat{\text{cov}}[z_t,\epsilon_t]]}. \quad (14)$$

All other terms in the Taylor-series expansion are of order $O(T^{-1})$. Noting that $\theta_{h+j} = \text{cov}[z_t, y_{t+h+j}]/\text{cov}[z_t, \epsilon_t]$, we calculate the numerator and denominator of (14) directly to obtain

$$\mathbb{E}\left[\widehat{\theta}_{h,IV}\right] - \theta_h \approx -\frac{1}{T-h-1} \sum_{j=1}^{T-h-1} \left(1 - \frac{j}{T-h}\right)(\theta_{h+j} + \theta_{h-j}), \tag{15}$$

which is identical to equation (6). As a result, the bias of LPs that we analyzed earlier persists in settings with instrumental variables.

6 Application to monetary policy shocks

We close our analysis with an empirical illustration of our bias correction using a setup similar to Gorodnichenko and Lee (2019). Using LPs, we estimate the effects of Romer and Romer (2004) monetary policy shocks on output and inflation.[10] Using quarterly data from 1969:Q1-2008:Q4, we estimate LPs of the form in (3) on real output growth and annualized inflation. We include controls consisting of four lags of real output growth, inflation, the federal funds rate, and the monetary policy innovation.[11]

The estimated impulse responses of inflation and output to a monetary policy shock are displayed in Figure 11. As in Gorodnichenko and Lee (2019), we cumulate the impulse response of output growth. Figure 11 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

[10]The shock series was extended to 2008 by Coibion et al. (2017).

[11]The 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.

(a) Inflation (b) Output

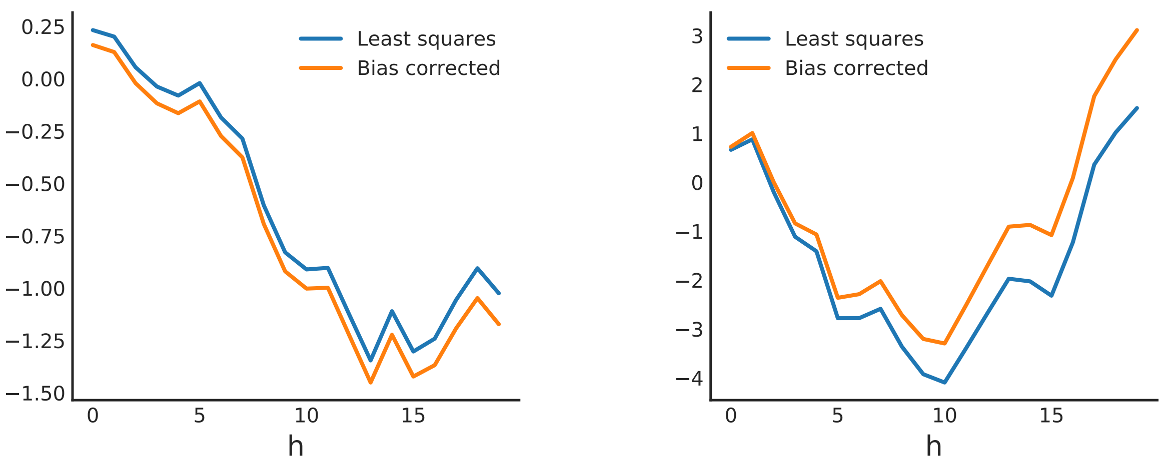

Figure 11: The effect of monetary policy shocks.

but nontrivial difference. The bias-corrected estimator is lower than the least-squares LP estimator because the estimated impulse response at long horizons is negative, and the bias-correction is a weighted average of the impulse responses at all 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 nearly 1 percentage point smaller. Essentially, this is because the bias-correction at say, horizon $h = 10$ is influenced by the LP coefficients at later horizons (which are themselves influenced by coefficients at even later horizons, et cetera). The LP coefficients at very long horizons (beyond the horizons show) are positive. The result of these positive values at long horizons is that the negative LP estimates at short horizons are adjusted up.

The bias corrections for both inflation and output were computed using $H = 50$, and by design this choice affects the results. As discussed in Section 3, parameter bias in LPs is a function of *all* the true LP coefficients, including those beyond the horizons of interest. Even if longer-horizon responses are not of interest in and of themselves, these responses affect the finite sample behavior of the LP estimator at shorter (and presumably more interesting) horizons. While it is true that the longer-horizon responses are, in general, more imprecisely estimated than those at shorter horizons, our Monte Carlo evidence suggests it is beneficial to use the estimates in bias correction.

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

We used our approximate bias function to bias correct LPs. However, in Monte Carlo analysis our bias correction does not completely correct for the bias in LPs, especially when the data are as persistent as most macroeconomic time series of interest. Our 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.

References

ACEMOGLU, D., S. NAIDU, P. RESTREPO, AND J. A. ROBINSON (2019): "Democracy does cause growth," *Journal of Political Economy*, 127, 47–100.

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.

HALL, P. (1992): *The Bootstrap and Edgeworth Expansion*, Springer New York.

JORDA, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161?182.

JORDÀ, Ò., M. SCHULARICK, AND A. M. TAYLOR (2015): "Betting the house," *Journal of International Economics*, 96, S2–S18.

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.

Lusompa, A. (2019): "Local Projections, Autocorrelations, and Efficiency," *Unpublished Manuscript, UC-Irvine.*

Nickell, S. (1981): "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 49, 1417–1426.

Plagborg-Møller, M. and C. K. Wolf (2019): "Local Projections and VARs Estimate the Same Impulse Responses," *Unpublished Manuscript, Princeton University.*

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.

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.

Stock, J. H. and M. W. Watson (2018): "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *Economic Journal*, 128, 917–948.

A Analytic approximation for bias with controls

In this section, we derive an analytic expression for LP bias in the case with controls. Consider the set of LPs in equation (3), reproduced below.

$$y_{t+h} = \alpha_h + \beta_h' x_t + u_{t,h}, \quad h = 0, \ldots, H. \tag{16}$$

Assume that x_t is stationary with mean μ_x and autocovariances $\Gamma_{x,j}, j = -\infty, \ldots, \infty$. Denote the cross-autocovariances $\Gamma_{xy,j}, j = -\infty, \ldots, \infty$.. Let $\tilde{x}_t = [1, x_t']'$. The least-squares estimator of β_h is given by:

$$\begin{aligned}
\hat{\beta}_{h,LS} &= \left(\frac{1}{T-h}\sum_{t=1}^{T-h}\tilde{x}_t\tilde{x}_t'\right)^{-1}\left(\frac{1}{T-h}\sum_{t=1}^{T-h}\tilde{x}_t y_{t+h}\right) \\
&= \left(\frac{1}{T-h}\sum_{t=1}^{T-h}x_t x_t' - \left(\frac{1}{T-h}\sum_{t=1}^{T-h}x_t\right)^2\right)^{-1} \\
&\quad \times \left(\frac{1}{T-h}\sum_{t=1}^{T-h}x_t y_{t+h} - \left(\frac{1}{T-h}\sum_{t=1}^{T-h}x_t\right)\left(\frac{1}{T-h}\sum_{t=1}^{T-h}y_t\right)\right) \\
&= (\widehat{\text{var}}[x_t])^{-1}\widehat{\text{cov}}[x_t, y_t].
\end{aligned} \tag{17}$$

The second equality follows from the Frisch-Waugh-Lovell theorem and $\widehat{\text{cov}}$ and $\widehat{\text{var}}$ are the sample covariance and variance. Taking a first-order Taylor expansion of (17) around $\mathbb{E}\left[\widehat{\text{var}}[x_t]\right]$ and $\mathbb{E}\left[\widehat{\text{cov}}[x_t, y_t]\right]$ yields

$$\mathbb{E}\left[\hat{\beta}_{h,LS}\right] \approx \mathbb{E}\left[\widehat{\text{var}}[x_t]\right]^{-1}\mathbb{E}\left[\widehat{\text{cov}}[x_t, y_t]\right]. \tag{18}$$

Consider first $\mathbb{E}\left[\widehat{\text{var}}[x_t]\right]$. Tedious algebra shows that:

$$E\left[\widehat{\text{var}}[x_t]\right] = \frac{1}{T-h}\sum_{t=1}^{T-h}\mathbb{E}\left[(x_t-\mu_x)(x_t-\mu_x)'\right] + \frac{1}{(T-h)^2}\mathbb{E}\left[\sum_{t=1}^{T-h}(x_t-\mu_x)\sum_{t=1}^{T-h}(x_t-\mu_x)'\right]. \tag{19}$$

The first term on the right-hand side of (19) is clearly equal to $\Gamma_{x,0}$. The second term can obtained through more tedious algebra:

$$\frac{1}{(T-h)^2}\mathbb{E}\left[\sum_{t=1}^{T-h}(x_t-\mu_x)\sum_{t=1}^{T-h}(x_t-\mu_x)'\right] = \frac{1}{T-h}\Gamma_{x,0} + \frac{1}{(T-h)^2}\sum_{j=1}^{T-h}(T-h-j)\left(\Gamma_{x,j}+\Gamma_{x,-j}\right).$$

Thus:

$$E\left[\widehat{\text{var}}[x_t]\right] = \left(1-\frac{1}{T-h}\right)\Gamma_{x,0} - \frac{1}{T-h}\sum_{j=1}^{T-h}\left(1-\frac{j}{T-h}\right)(\Gamma_{x,j}+\Gamma_{x,-j}). \tag{20}$$

Note that this generalizes the IID case. The covariance term can be similarly obtained as:

$$\begin{aligned}\mathbb{E}\left[\widehat{\text{cov}}[x_t, y_t]\right] &= \frac{1}{T-h}\sum_{t=1}^{T-h}\mathbb{E}\left[(x_t-\mu_x)(y_{t+h}-\mu_y)'\right] \\ &\quad -\frac{1}{(T-h)^2}\mathbb{E}\left[\sum_{t=1}^{T-h}(x_t-\mu_x)\sum_{t=1}^{T-h}(y_{t+h}-\mu_y)\right]. \qquad (21)\end{aligned}$$

Clearly the first term on the right-hand side of (21) is equal to $\Gamma_{xy,h}$. Tedious algebra shows that the second term is equal to:

$$\begin{aligned}\frac{1}{(T-h)^2}\mathbb{E}\left[\sum_{t=1}^{T-h}(x_t-\mu_x)\sum_{t=1}^{T-h}(y_{t+h}-\mu_y)\right] &= \frac{1}{T-h}\Gamma_{xy,h} \\ &+ \frac{1}{(T-h)^2}\sum_{j=1}^{T-h-1}(T-h-j)\left(\Gamma_{xy,h+j}+\Gamma_{xy,h-j}\right).\end{aligned}$$

Thus,

$$E\left[\widehat{\text{cov}}[x_t, y_t]\right] = \left(1-\frac{1}{T-h}\right)\Gamma_{xy,h} - \frac{1}{T-h}\sum_{j=1}^{T-h}\left(1-\frac{j}{T-h}\right)(\Gamma_{xy,h+j}+\Gamma_{xy,h-j}). \qquad (22)$$

Next, note that:

$$\begin{aligned}\mathbb{E}\left[\widehat{\text{var}}[x_t]\right]^{-1}\mathbb{E}\left[\widehat{\text{cov}}[x_t, y_t]\right] &= \mathbb{E}\left[\widehat{\text{var}}[x_t]\right]^{-1}\left(\left(1-\frac{1}{T-h}\right)\Gamma_{x,0}\right)\left(\left(1-\frac{1}{T-h}\right)\Gamma_{x,0}\right)^{-1}\mathbb{E}\left[\widehat{\text{cov}}[x_t, y_t]\right] \\ &= \left[\left(\left(1-\frac{1}{T-h}\right)\Gamma_{x,0}\right)^{-1}\mathbb{E}\left[\widehat{\text{var}}[x_t]\right]\right]^{-1} \\ &\qquad \left(\left(1-\frac{1}{T-h}\right)\Gamma_{x,0}\right)^{-1}\mathbb{E}\left[\widehat{\text{cov}}[x_t, y_t]\right]. \qquad (23)\end{aligned}$$

For the first term of the right hand side of (23) we have:

$$\begin{aligned}\left(\left(1-\frac{1}{T-h}\right)\Gamma_{x,0}\right)^{-1}\mathbb{E}\left[\widehat{\text{var}}[x_t]\right] &= I + \frac{1}{T-h-1}\sum_{j=1}^{T-h}\left(1-\frac{j}{T-h}\right)\Gamma_{x,0}^{-1}(\Gamma_{x,j}+\Gamma_{x,-j}) \\ &:= \Psi_h. \qquad (24)\end{aligned}$$

The second term is:

$$\begin{aligned}\left(\left(1-\frac{1}{T-h}\right)\Gamma_{x,0}\right)^{-1}\mathbb{E}\left[\widehat{\text{cov}}[x_t, y_t]\right] &= \\ \beta_h + \frac{1}{T-h-1}\sum_{j=1}^{T-h-1}\left(1-\frac{j}{T-h}\right)(\beta_{h+j}+\beta_{h-j}). \qquad (25)\end{aligned}$$

Thus, we deduce that:

$$\mathbb{E}\left[\hat{\beta}_{h,LS}\right] \approx \Psi_h \left(\beta_h + \frac{1}{T-h-1}\sum_{j=1}^{T-h-1}\left(1-\frac{j}{T-h}\right)(\beta_{h+j}+\beta_{h-j})\right). \tag{26}$$

This expression nests the one in (6). In that case, $\Psi_h = 1$, because ϵ_t is iid. In the controls case, Ψ is a weighted average of the population regression coefficients associated with the direct forecast of x_t for horizons $j = -(T-h), \ldots, (T-h)$.

B Additional Monte Carlo simulations

(a) Without Controls (b) With Controls

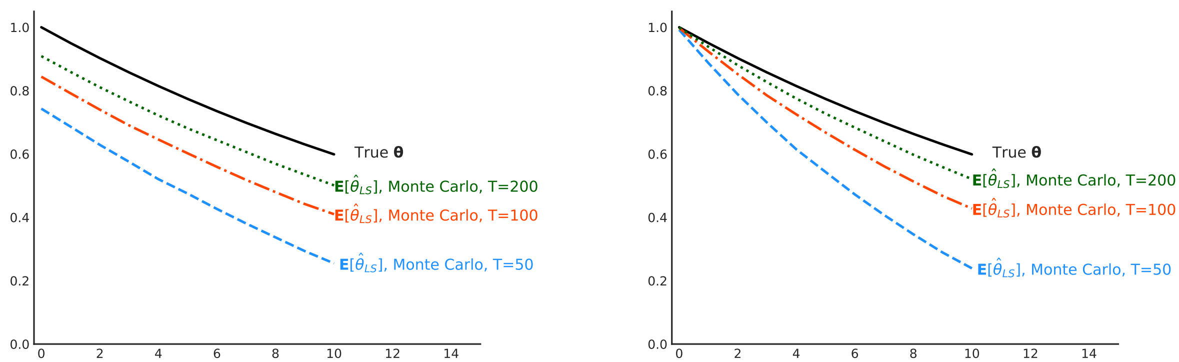

Figure B.1: LP estimators without controls are biased in empirically-relevant samples when y_t is an AR(1) with $\rho = 0.95$.

(a) Without Controls

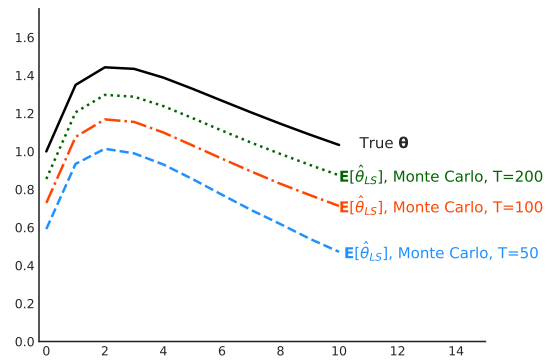

(b) With Controls

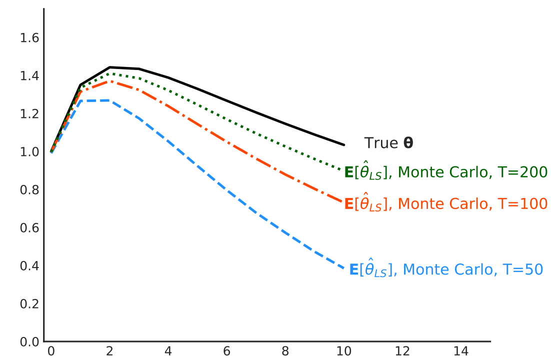

Figure B.2: LP estimators without controls are biased in empirically-relevant samples when y_t is an AR(2) with $\rho = 0.95$ and $\psi = 0.4$.

(a) Without Controls

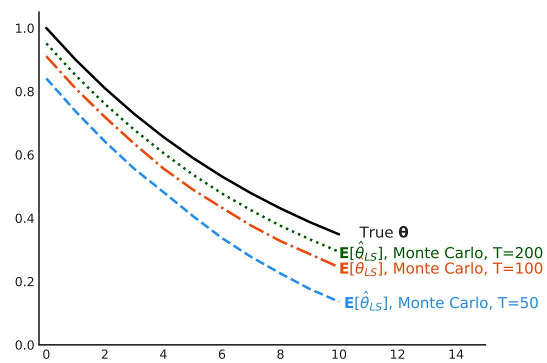

(b) With Controls

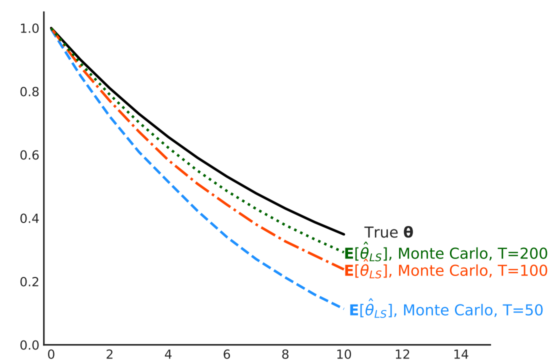

Figure B.3: LP estimators without controls are biased in empirically-relevant samples when y_t is an AR(1) with $\rho = 0.90$.

C 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

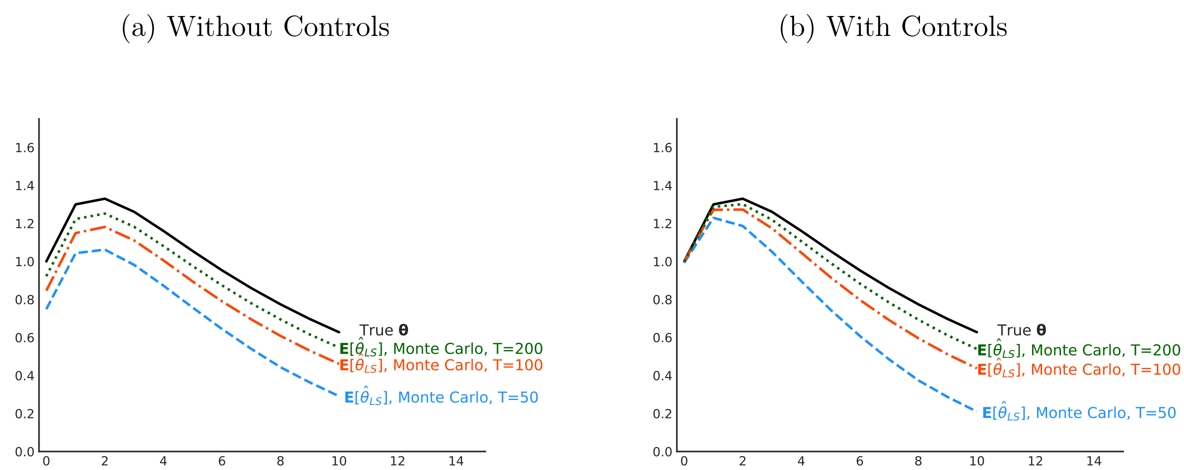

Figure B.4: LP estimators without controls are biased in empirically-relevant samples when y_t is an AR(2) with $\rho = 0.90$ and $\psi = 0.4$.

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-Villaverde, 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. Krznar, and S. Mitra (2013): "Evaluating the net benefits of macroprudential policy: A cookbook," .

Auerbach, A. J. and Y. Gorodnichenko (2012): "Fiscal multipliers in recession and expansion," in *Fiscal policy after the financial crisis*, University of Chicago Press, 63–98.

——— (2017): "Fiscal multipliers in Japan," *Research in Economics*, 71, 411–421.

Ball, L. M., D. Furceri, M. D. Leigh, and M. P. Loungani (2013): *The distributional effects of fiscal consolidation*, 13-151, International Monetary Fund.

Banerjee, R., M. B. Devereux, and G. Lombardo (2016): "Self-oriented monetary policy, global financial markets and excess volatility of international capital flows," *Journal of International Money and Finance*, 68, 275–297.

Banerjee, R. N. and H. Mio (2018): "The impact of liquidity regulation on banks," *Journal of Financial intermediation*, 35, 30–44.

Barnichon, R. and C. Brownlees (2019): "Impulse response estimation by smooth local projections," *Review of Economics and Statistics*, 101, 522–530.

Basher, S. A., A. A. Haug, and P. Sadorsky (2012): "Oil prices, exchange rates and emerging stock markets," *Energy Economics*, 34, 227–240.

Baumeister, C. and L. Kilian (2016): "Lower oil prices and the US economy: Is this time different?" *Brookings Papers on Economic Activity*, 2016, 287–357.

Bayer, C., R. Lütticke, L. Pham-Dao, and V. Tjaden (2019): "Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk," *Econometrica*, 87, 255–290.

Ben Zeev, N. and E. Pappa (2017): "Chronicle of a war foretold: The macroeconomic effects of anticipated defence spending shocks," *The Economic Journal*, 127, 1568–1597.

Berton, F., S. Mocetti, A. F. Presbitero, and M. Richiardi (2018): "Banks, firms, and jobs," *The Review of Financial Studies*, 31, 2113–2156.

Bhattacharya, U., N. Galpin, R. Ray, and X. Yu (2009): "The role of the media in the internet IPO bubble," *Journal of Financial and Quantitative Analysis*, 44, 657–682.

Bluedorn, J. C. and C. Bowdler (2011): "The open economy consequences of US monetary policy," *Journal of International Money and Finance*, 30, 309–336.

Borio, C., M. Drehmann, and K. Tsatsaronis (2014): "Stress-testing macro stress testing: does it live up to expectations?" *Journal of Financial Stability*, 12, 3–15.

Borio, C. E., E. Kharroubi, C. Upper, and F. Zampolli (2016): "Labour reallocation and productivity dynamics: financial causes, real consequences," .

Born, B., G. J. Müller, and J. Pfeifer (2014): "Does austerity pay off?" *Review of Economics and Statistics*, 1–45.

Boss, M., G. Fenz, J. Pann, C. Puhr, M. Schneider, E. Ubl, et al. (2009): "modeling credit risk through the Austrian business cycle: An update of the OeNb model," *Financial Stability Report*, 17, 85–101.

Brady, R. R. (2011): "Measuring the diffusion of housing prices across space and over time," *Journal of Applied Econometrics*, 26, 213–231.

Butt, N., R. Churm, M. F. McMahon, A. Morotz, and J. F. Schanz (2014): "QE and the bank lending channel in the United Kingdom," .

Caggiano, G., E. Castelnuovo, V. Colombo, and G. Nodari (2015): "Estimating fiscal multipliers: News from a non-linear world," *The Economic Journal*, 125, 746–776.

Caldara, D. and E. Herbst (2019): "Monetary policy, real activity, and credit spreads: Evidence from bayesian proxy svars," *American Economic Journal: Macroeconomics*, 11, 157–92.

Caselli, F. G. and A. Roitman (2016): "Nonlinear exchange-rate pass-through in emerging markets," *International Finance*.

Chiţu, 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 (2012): "Innocent bystanders? Monetary policy and inequality in the US," Tech. rep., National Bureau of Economic Research.

——— (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-Crichton, 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.