

Smooth Panel Local Projections

Jonathan S. Hartley*

Jackson Mejia[†]

January 8, 2025

This note extends smooth local projections to panel data and uses Monte Carlo methods to explore the bias and variance properties of smooth panel local projections (SPLP). SPLP allows researchers to penalize the impulse response toward a polynomial, while standard local panel projections (PLP) are nonparametric but result in theoretically unappealing IRFs because they are too lumpy. Relative to PLP, SPLP has appealing properties in smaller samples.

*Department of Economics, Stanford University. Email: hartleyj@stanford.edu.

[†]Department of Economics, Massachusetts Institute of Technology. Email: jpmejia@mit.edu. Isaiah Andrews, Tomás Caravello, Pedro Martinez-Bruera, and Vod Vilfort provided useful input. This material is based on work supported by the National Science Foundation under Grant No. 1122374.

1 Introduction

Economists increasingly rely on panel local projections (PLP) to study the dynamic effects of shocks.¹ Indeed, with the advent of PLP in difference-in-difference studies, panel local projections are an important tool for both macroeconomists and microeconomists alike (Dube et al. 2023). However, in a typical example, Arezki, Ramey, and Sheng (2017) note that “the estimated impulse responses [from PLP] are more erratic and often less precise” than other dynamic panel estimators. This problem plagues many studies which estimate local projections because their nonparametric nature leads to unappealingly lumpy impulse responses. Coefficients jump from horizon to horizon in a way that is a result of noise and cannot be meaningfully justified with economic theory. Barnichon and Brownlees (2019) solve that problem with time series data; this note makes several contributions toward solving that problem for panel data.

First, we extend smooth panel local projections from time series to panel data (SPLP). SPLP penalizes the impulse response of a panel local projection (PLP) to a polynomial. The key benefit is the new ability to leverage both the cross-sectional and time series dimensions to appropriately select the penalty parameter and appeal to researchers using both micro and macro data. Second, we give evidence from Monte Carlo simulations that as long as researchers have some confidence in the shape of the impulse response, they should choose smooth panel local projections (SPLP) over standard local panel projections. We provide new evidence of the robustness of SPLP and offer a novel argument that one should use SPLP in connection to economic theory. Theory suggests different kinds of impulse responses for different hypotheses, which can be used to discipline the polynomial choice. Third, we illustrate the method in two applications. Finally, we provide both R and Stata code for practitioners to use.²

2 Smooth Local Panel Projections

This section outlines the procedure for estimating smooth local projections for panel data. If readers are already familiar with Barnichon and Brownlees (2019), they can proceed to the Section 3; much of the following simply rehashes their Section 2, but suitably adapted

1. Because of the extra dimension of variation offered by panel data over time series, Herbst and Johannsen (2024) document that a large share of studies utilizing local projections rely on panel data.

2. The package, called `splp` and `vignette` can be found [here](#). The package is considerably faster than existing PLP estimation packages because it relies on `fixest`.

for panel data. Consider a typical dynamic panel regression

$$y_{i,t+h} = G(i, t) + b_h x_{i,t} + \text{Controls} + \eta_{i,t+h},$$

for cross-sectional units $i = 1, \dots, N$ estimated at time $t + h$ given a shock to $x_{i,t}$ at time t . The sequence of coefficients b_h from t to $t + h$ consist of the impulse response up to h horizons out. In general, our interest is in smoothing the coefficient b_h across horizons.³ $G(i, t)$ is a set of additive fixed effects that depend on unit i at time t but are not necessarily two-way time and individual fixed effects. For example, one may wish to include an industry fixed effect or an interaction between time and industry when studying how firms react to shocks. All that matters in this context is that the dependent and independent variables are demeaned with respect to $G(i, t)$.⁴

As in Barnichon and Brownlees (2019), the goal is to make the coefficient b_h a smooth function of the impulse horizon. To do that, we simply use a B-spline basis function to approximate the coefficient

$$b_h = \sum_{k=1}^K b_k B_k(h)$$

for K sufficiently large and where $B_k : \mathbb{R} \rightarrow \mathbb{R}$. The B-splines link coefficients across horizons. The basis functions are polynomial pieces of order q which are continuous at the kinks by imposing the constraint that all derivatives up to order $q - 1$ are continuous at the kinks. There are $q + 2$ kinks and each one is called an inner knot. See Barnichon and Brownlees (2019) and references therein for a more detailed discussion of basis functions. Following their discussion, we use a cubic basis function throughout.

To set notation, let variables with a tilde denote their demeaned version with respect to fixed effects $G(i, t)$. Let H_{max} denote the maximum forecast horizon. Let $\mathbf{y}_{i,t}$ denote the vector of outcomes $(\tilde{y}_{i,t}, \dots, \tilde{y}_{i,t+\min\{T, t+H_{max}\}})^0$ with length d_t . Let $\mathbf{x}_{i,t}$ for $t = 1, \dots, T$ denote the $d_t \times K$ matrix with element (h, k) equal to $B_k(h)x_{i,t}$. Next, let Y denote the stacked vector individual vectors $\mathbf{y}_{i,t}$ and X denote the stacked matrices for individuals $\mathbf{x}_{i,t}$. Finally, let q denote the vector of B-splines coefficients (b_1, \dots, b_K) .

In principle, one could smooth IRFs for multiple variables in SPLP. But in practice, it is hard to think of an application where that would be either practical or relevant be-

3. The method can be readily extended to instrumental variables as in, for example, Jordà, Schularick, and Taylor (2020). We focus on OLS for simplicity, but the computational method is the same and our packages allow for IV estimation.

4. If y is nonstationary, the researcher should include lags of y . Montiel Olea and Plagborg-Møller (2021) show that nonstationary LPs are consistent in time series as long as they include lagged regressors. Of course, this also makes them subject to Nickell bias in a panel data setting.

cause researchers frequently only care about the response of a single outcome variable to a single exogenous shock. Consequently, we focus solely on cases where one wants the smoothed IRF for one variable.⁵ With that notation, a cookbook for estimating SPLP is the following.

1. **Demean data with respect to fixed effects** $G(i, t)$.
2. **Construct matrices** Y and X . Note that maintaining order is crucial for the demeaned data. In particular, demeaned data must be ordered within individual clusters by time and horizon regardless of fixed effects. After this step, what follows is essentially the same as in Barnichon and Brownlees (2019).
3. **Estimate ridge regression and select penalty parameter:**

$$\begin{aligned} \hat{q} &= \arg \min_q f(k) Y' X q^2 + \lambda q' P q \\ &= (X' X + \lambda P)^{-1} X' Y, \end{aligned} \tag{1}$$

where $\lambda > 0$ is a shrinkage parameter and P is a symmetric positive semidefinite penalty matrix. λ determines the bias/variance trade-off in shrinking the impulse response toward the chosen polynomial order. Over a grid of penalty parameters, use generalized cross-validation to select λ .⁶

4. **Construct confidence bands.** In a time series context, Barnichon and Brownlees (2019) propose using Newey-West standard errors. Because panel data researchers are often interested in clustering standard errors, we propose instead using a wild cluster bootstrap. Because the point estimates themselves are biased by construction, the standard errors are too. This is the same inference problem faced by Barnichon and Brownlees (2019).

3 SPLP in Practice

In this section, we proceed in two steps to discuss how researchers should implement SPLP. First, we conduct Monte Carlo simulations to compare SPLP to standard panel

5. To account for the case with controls, suppose there is a vector of controls $C_{i,t}$ of length J . With $J = 1$, we would have a $d_t \times K$ matrix with element (h, K) equal to $c_{i,t}$. That is, rather than multiply by the basis function, we simply multiply by the identity matrix. The same procedure can be extended for multiple controls.

6. We use generalized cross-validation based on Golub, Heath, and Wahba (1979) because it is faster than k-fold validation for large panel datasets.

local projections (PLP). Second, we discuss how researchers should proceed if they are unsure about the shape of the true impulse response function.

3.1 Monte Carlo Simulations

To evaluate the performance of SPLP vis-à-vis standard PLP, we consider simple regressions of the form

$$y_{i,t+h} = a_i + T_t + b_h x_{i,t} + \#_{i,t} \quad (2)$$

where a_i is an individual fixed effect, T_t is a time fixed effect, and $x_{i,t}$ is the time t shock to y . This is perhaps the most commonly used regression in practice. We evaluate SPLP under three different DGPs, where the shape of the IRF is linear in the first, quadratic in the second, and cubic in the third. Figure 1 plots the targeted impulse responses.




Figure 1: True IRFs used for Monte Carlo simulations. The blue IRF is layered underneath the red one for the first few horizons.

The Monte Carlo simulations test a variety of panel sizes. For each panel size, we generate 5000 replications of sample data for each IRF type. To test the performance of the bootstrap, we draw 5000 bootstrap samples. Throughout, we use a balanced panel. Each replication selects a penalty parameter from a grid of penalty parameters from 10^{-3} to 10^3 distributed logarithmically over 100 grid points.

Point Estimates

To start, suppose we know what the shape of the underlying process is. In general, this will not be true and we will address the contrary case in the following subsection. Following Li, Plagborg-Møller, and Wolf (2021), we compare the performance of SPLP and

PLP at each horizon using the following loss function

$$L_h(\hat{b}_h, b_h, w) = w \mathbb{E} \|\hat{b}_h - b_h\|^2 + (1-w) \mathbf{V} \hat{b}_h, \quad (3)$$

where $w \in [0, 1]$ is the weight the econometrician places on bias. To place a summary number on the comparison, we compute the discounted sum of losses at each horizon as

$$L^i(w) = \sum_{h=0}^H \frac{1}{1+r} L_h,$$

for $i \in \{SPLP, PLP\}$ where $r \geq 0$ is the econometrician's discount rate and PLP refers to standard panel local projections. r may be greater than zero because one cares more about minimizing losses at nearer horizons. For simplicity, we set $r = 0$. Finally, we compare the performance of SPLP and PLP by taking the ratio

$$v(w) = \frac{L^{PLP}(w)}{L^{SPLP}(w)}.$$

If $v(w) > 1$, then the econometrician should prefer SPLP to PLP. Conversely, if $v(w) < 1$, then the econometrician should prefer PLP. In Figure 2, we plot the curve $v(w)$ for $N = 400$ individuals and $T = 100$ time periods each IRF type. Although the sample is large, which means that the variance of the PLP estimator is substantially diminished, an econometrician would prefer SPLP for both cubic and quadratic IRFs unless the weight placed on bias is close to 0.9. In other words, SPLP substantially outperforms PLP. Moreover, an econometrician would never prefer PLP over SPLP if the true IRF is linear.




Figure 2: $v(w)$ curve for $T = 100$ time periods and $N = 400$ individuals.

As a summary measure, consider the threshold variable \bar{w} , where we say that $w > \bar{w}$ means that the researcher should choose SPLP. Table 1 tabulates \bar{w} for linear, quadratic, and cubic IRFs for sample sizes which are empirically relevant in the data. If the true IRF is linear, then the researcher should never choose PLP even if he has an overwhelming preference for bias. As more curvature enters the IRF through either a quadratic or a cubic IRF, then the dominance of SPLP recedes. Even so, threshold values around 0.9 indicate researchers should generally prefer SPLP. Table 4 tabulates the value of the loss function for $w \in \{0, 0.5, 1\}$ to give an idea of the shape of the $v(w)$ function for varying IRF shapes and sample sizes.

N	T	\bar{w}		
		Linear	Quadratic	Cubic
50	25	1.00	0.97	0.97
100	25	0.99	0.95	0.95
250	25	0.99	0.89	0.90
50	50	0.99	0.97	0.97
100	50	0.99	0.94	0.95
250	50	0.99	0.87	0.88
50	100	0.99	0.96	0.96
100	100	0.99	0.92	0.94
250	100	0.99	0.83	0.86

Table 1: Threshold value \bar{w} for different underlying IRFs and varying sample sizes. If the preference for bias exceeds \bar{w} , then the econometrician should prefer PLP.

Confidence Interval Performance

Standard confidence intervals generally cannot be used with ridge regression, so we instead construct wild cluster bootstrap intervals. We compare the PLP confidence intervals created from clustering at the individual level to a wild cluster bootstrap at the same level of aggregation. Table 2 documents length properties of the wild cluster bootstrap for SPLP compared to clustered standard errors constructed analytically for PLP. Each column reports the percent of the time the 95% confidence interval contains the true parameter. This is value averages over all horizons. Across all sample sizes, the PLP and SPLP confidence intervals perform nearly identically in the linear case. However, introducing some curvature in the true IRF substantially affects the performance of the wild cluster bootstrap when the time dimension is short. When the time dimension is at least fifty, then it per-

forms similarly to the analytically computed cluster standard error for PLP, but a short time dimension is especially problematic for the bootstrap. It is not surprising that the bootstrap struggles in small samples when there is a lot of curvature. SPLP is biased by construction and will tend to underestimate the curvature of the function. For example, if there is a theoretical minimum at $h = 3$, then the coefficient $\hat{b}_3 < b_3$ because it is biased upward by the smoothing with the higher points. One can get around this by using more knots or increasing the polynomial order, but the trade-off is that the estimator just approaches the PLP estimator.

N	T	Linear		Quadratic		Cubic	
		PLP	SPLP	PLP	SPLP	PLP	SPLP
50	25	1.922	0.878	1.941	1.080	1.934	1.216
100	25	0.864	0.402	0.864	0.492	0.864	0.551
250	25	0.683	0.316	0.684	0.391	0.684	0.438
50	50	1.913	0.855	1.932	1.067	1.925	1.206
100	50	0.859	0.389	0.861	0.484	0.857	0.547
250	50	0.680	0.308	0.679	0.384	0.679	0.435
50	100	1.919	0.819	1.911	1.028	1.911	1.166
100	100	0.858	0.373	0.860	0.469	0.856	0.536
250	100	0.680	0.296	0.679	0.371	0.678	0.427

Table 2: Length properties of PLP standard errors versus the wild cluster bootstrap.

The coverage properties of the wild cluster bootstrap for SPLP are also quite promising for applied researchers. Table 3 tabulates the length of the 95% confidence interval for different sample sizes and IRF types. In smaller samples, the length of the SPLP confidence interval is about half that of the equivalently clustered confidence interval and remains significantly smaller even in large samples. Because the coverage properties are nearly as good for the SPLP bootstrap, this means that the SPLP is much better at correctly determining significance as long as there is not too much curvature in the true IRF and the sample size is not too small.

N	T	Linear		Quadratic		Cubic	
		PLP	SPLP	PLP	SPLP	PLP	SPLP
50	25	0.936	0.922	0.931	0.892	0.931	0.930
100	25	0.929	0.877	0.903	0.788	0.931	0.896
250	25	0.915	0.811	0.874	0.745	0.925	0.870
50	50	0.937	0.938	0.939	0.913	0.934	0.932
100	50	0.942	0.937	0.941	0.851	0.950	0.923
250	50	0.942	0.933	0.924	0.801	0.945	0.912
50	100	0.937	0.951	0.939	0.909	0.937	0.921
100	100	0.945	0.945	0.947	0.837	0.953	0.929
250	100	0.949	0.949	0.943	0.784	0.948	0.921

Table 3: Coverage properties of PLP standard errors versus the wild cluster bootstrap.

3.2 Uncertainty about the IRF

Economists often have in mind the shape of an IRF prior to engaging with the data. Typically, a model or a class of models maps into a linear, quadratic, or cubic IRF. However, one may wish to use the IRF to distinguish between models. For example, capital adjustment costs yield approximately linear IRFs, while investment adjustment costs produce hump-shaped IRFs. In such situations, when the econometrician has a weak prior on the shape of the IRF, standard LP will be practically useless because it is too noisy to determine with any certainty the shape of the IRF. In contrast, a tool like SPLP can be quite useful. Consider a simple panel local projection of the form

$$y_{i,t+h} = a_i + T_t + b_h x_{i,t} + \#_{i,t+h}.$$

Since a significant part of the procedure for SPLP is selecting a polynomial order to discipline the IRF, one may worry that about naively choosing the wrong polynomial order will lead to incorrect results. For example, one may select a linear polynomial when the true IRF is quadratic. However, SPLP is robust to this concern. Through appropriate selection of the penalty parameter, the resulting IRFs will be quite similar for lower polynomial orders. To show that, we generate sample data with $N = 250$ cross-sectional units and $T = 50$ time periods. In Figure 5, we plot the impulse responses for several different polynomial orders when the true DGP is either linear, quadratic, or cubic. It turns out that generalized cross-validation appropriately selects a penalty parameter for all three

cases such that it does not matter very much in practice whether the selected polynomial order is less than four. However, as the polynomial order increases, the SPLP estimator increasingly resembles the PLP estimator.

The practical takeaway is that it will generally lead to robust results if the econometrician selects a polynomial of order one, two, or three because the penalty selection process will ensure that they are all similar. However, one should generally not select a polynomial order much higher than three and there is not a good econometric reason to do that anyway. Moreover, unless one thinks that the IRF is a jump process with a unit root, then it likewise is not recommended to select a polynomial of order zero.

4 Applications to Oil and Democracy

This section applies SPLP to oil shocks from Arezki, Ramey, and Sheng (2017) and democracy shocks from Acemoglu et al. (2019).

4.1 Oil Shocks

Using a novel data source, Arezki, Ramey, and Sheng (2017) construct a panel of oil news shocks and show that the impulse responses of macro variables to these shocks conforms to a simple open economy model. The paper uses a variety of dynamic panel data models to estimate impulse responses, including PLP methods which result in erratic IRFs. As an example of SPLP, we smooth out those those IRFs. That procedure is aided by the fact that the paper draws a tight connection between theoretical and empirical IRFs, which allows us to use theory to discipline the polynomial order of impulse response functions.

In our illustration of SPLP, we use their baseline measure of the oil news shock, which is the net present value of oil discovery scaled by time t GDP and discounts the future using country-specific discount rates. As a first application, consider the response of the current account scaled by GDP to an oil shock in their full sample of data.⁷ Since the current account is stationary, this comes from the regression

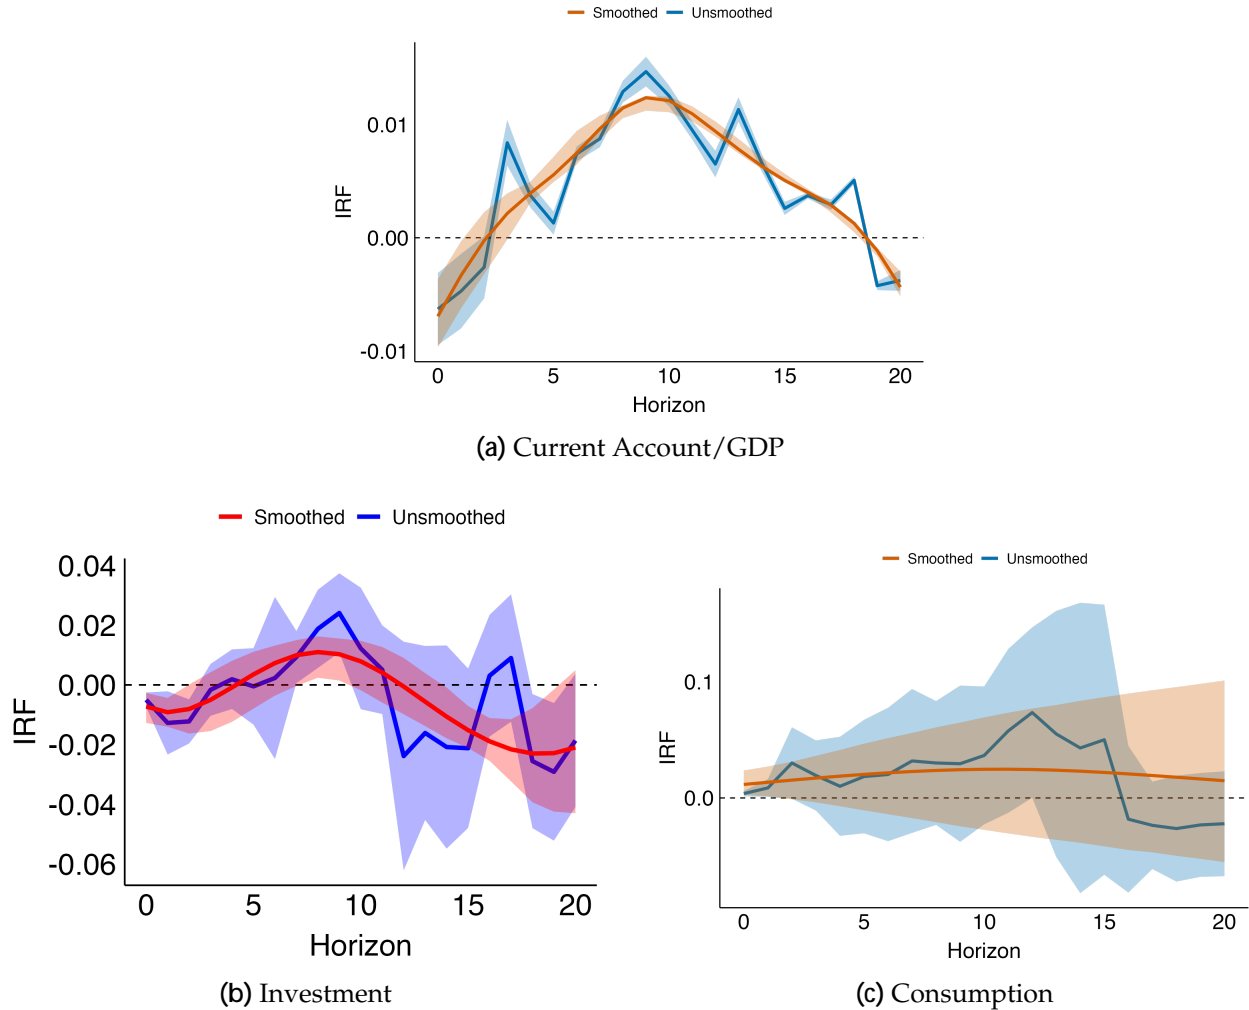
$$y_{i,t+h} = a_i + T_t + b_h x_{i,t} + \#_{i,t+h}, \quad (4)$$

where a_i is a country i fixed effect and T_t is a time fixed effect. $x_{i,t}$ is the oil news shock for country i . Figure 3a plots the impulse response of the current account for up to twenty years out penalized to a quadratic along with a 95% confidence interval in orange. In blue,

7. I use their replication data and simply filter out any observations without a current account or an oil shock.

we also plot the PLP estimator. The PLP estimator is lumpy by comparison, although there is very little difference in the confidence interval.

Figure 3: Responses of aggregate variables to oil shocks



Notes: Impulse responses of macro variables to oil shocks. The PLP impulse response is in blue and the SPLP IRF is in orange. The associated 95% confidence interval is a wild cluster bootstrap with 2000 samples.

As a second example, consider the response of investment/GDP, plotted in Figure 3c. Under King-Plosser-Rebeloe preferences, the saving rate should rise, decline, and then rise again in response to an oil shock. This suggests disciplining the IRF to a cubic. The smoothed IRF in Figure 3c supports the theory, but the PLP estimator IRF is ambiguous. As a third example, we consider the response of consumption. Theory implies that the consumption response to an oil news shock should be flat (Arezki, Ramey, and Sheng 2017, p. 115) to reflect consumption smoothing. Figure 3b plots the consumption response to an oil news shock. As predicted, it is essentially flat and positive, albeit not significantly

so.

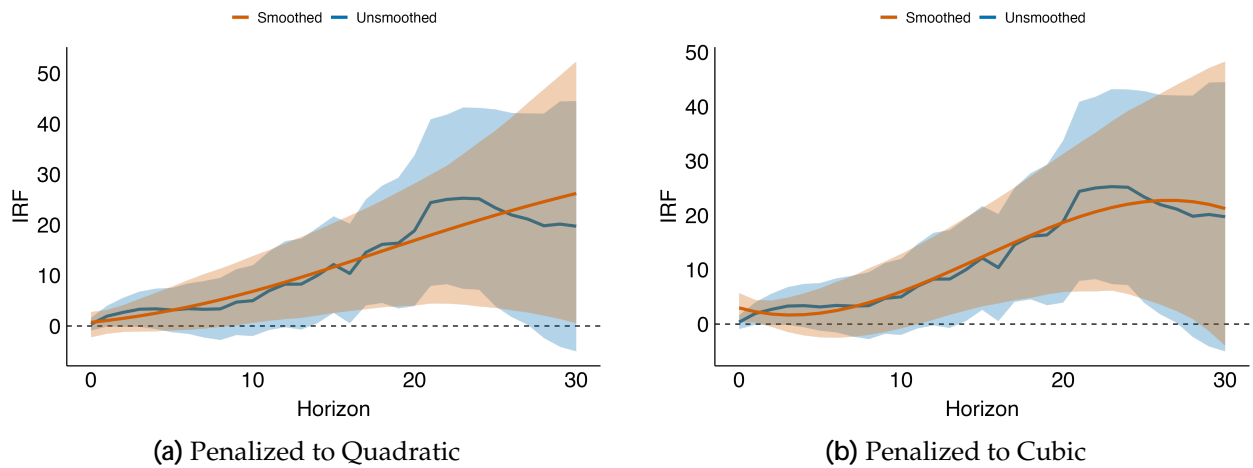
4.2 Democracy Shocks

As another example, consider Acemoglu et al. (2019), which studies the causal effect of democracy on aggregate outcomes. One of their main results (Figure 3) estimates

$$y_{i,t+h} - y_{i,t-1} = T_t + b_h D_{i,t} + \sum_{j=1}^p g_{h,j} y_{i,t-j} + \epsilon_{i,t+h},$$

where y is log output and D is an indicator for democratic transition. Figure 4 plots the SPLP and PLP dynamic responses of output to a democratic transition for both quadratic and cubic polynomials.

Figure 4: Response of output to democratic transitions



Notes: Impulse response of output to a transition into democracy. The PLP impulse response is in blue and the SPLP IRF is in orange. The associated 95% confidence interval is a wild cluster bootstrap with 2000 samples.

Although the SPLP response is not dramatically different from the PLP response, it offers another interpretation for how to understand such regressions. When the polynomial is lower-order, it ends up looking similar to the dynamic response of a dynamic panel like

$$y_{i,t} = T_t + a_i + b D_{i,t} + g y_{i,t-1} + \epsilon_{i,t+h},$$

where the IRF is extrapolated from $b/(1-g)$. That is, a near-linear polynomial approximation to the IRF ends up looking similar to what one would estimate from simply run-

ning a dynamic regression with a lagged dependent variable and constructing the IRF based on that. Indeed, one can see that from looking at Figure 2 in Acemoglu et al. (2019).

5 Concluding Remarks

This paper extends smooth local projections to panel data. We show that SPLP is preferable to standard panel local projections except in rather extreme circumstances. Additionally, we show that the estimator is generally robust to selecting the wrong polynomial order. Finally, we demonstrate the utility of the estimator through two applications.

References

- Acemoglu, Daron, Suresh Naidu, Pascual Restrepo, and James A. Robinson. 2019. "Democracy Does Cause Growth." *Journal of Political Economy* 127, no. 1 (February): 47–100. ISSN: 0022-3808. <https://doi.org/10.1086/700936>.
- Arezki, Rabah, Valerie A. Ramey, and Liugang Sheng. 2017. "News Shocks in Open Economies: Evidence from Giant Oil Discoveries*." *The Quarterly Journal of Economics* 132, no. 1 (February): 103–155. ISSN: 0033-5533. <https://doi.org/10.1093/qje/qjw030>.
- Barnichon, Regis, and Christian Brownlees. 2019. "Impulse Response Estimation by Smooth Local Projections." *The Review of Economics and Statistics* 101, no. 3 (July): 522–530. ISSN: 0034-6535. https://doi.org/10.1162/restfn_ga_fn_g00778.
- Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan Taylor. 2023. *A Local Projections Approach to Difference-in-Differences*. Technical report. Cambridge, MA: National Bureau of Economic Research, April. <https://doi.org/10.3386/w31184>.
- Golub, Gene H., Michael Heath, and Grace Wahba. 1979. "Generalized Cross-Validation as a Method for Choosing a Good Ridge Parameter." *Technometrics* 21, no. 2 (May): 215–223.
- Herbst, Edward P., and Benjamin K. Johansson. 2024. "Bias in local projections." *Journal of Econometrics* 240, no. 1 (March): 105655. ISSN: 03044076. <https://doi.org/10.1016/j.jeconom.2024.105655>.
- Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2020. "The effects of quasi-random monetary experiments." *Journal of Monetary Economics* 112 (June): 22–40. ISSN: 03043932. <https://doi.org/10.1016/j.jmoneco.2019.01.021>.
- Li, Dake, Mikkel Plagborg-Møller, and Christian K. Wolf. 2021. "Local Projections vs. VARs: Lessons From Thousands of DGPs" (April).
- Montiel Olea, José Luis, and Mikkel Plagborg-Møller. 2021. "Local Projection Inference Is Simpler and More Robust Than You Think." *Econometrica* 89 (4): 1789–1823. ISSN: 0012-9682. <https://doi.org/10.3982/ECTA18756>.

N	T	Linear			Quadratic			Cubic		
		$w = 0$	$w = 0.5$	$w = 1$	$w = 0$	$w = 0.5$	$w = 1$	$w = 0$	$w = 0.5$	$w = 1$
50	25	4.66	4.13	1.00	3.18	2.67	0.84	2.55	2.38	0.67
100	25	4.61	3.68	0.98	3.03	2.34	0.81	2.61	2.29	0.64
250	25	4.85	3.06	0.98	3.00	1.78	0.82	2.51	1.92	0.67
50	50	5.02	4.86	0.93	3.31	3.04	0.39	2.57	2.44	0.34
100	50	4.79	4.52	0.87	3.21	2.69	0.44	2.56	2.33	0.27
250	50	5.20	4.33	0.88	3.19	2.18	0.45	2.48	2.02	0.27
50	100	5.40	5.32	0.50	3.46	3.13	0.14	2.75	2.59	0.12
100	100	5.45	5.25	0.62	3.52	2.88	0.17	2.68	2.41	0.10
250	100	5.42	4.96	0.60	3.47	2.25	0.13	2.67	2.10	0.07

Table 4: $v(w)$ for $w \in \{0, 0.5, 1\}$.

