NBER 40th Annual Conference on Macroeconomics – Discussion Summaries

Local Projections or VARs? A Primer for Macroeconomists

Authors: Jose Luis Montiel Olea, Mikkel Plagborg-Møller, Eric Qian, Christian Wolf

Discussants: Christiane Baumeister, Oscar Jorda

Gabriel Chodorow-Reich opened the discussion with a suggestion to include scatter plots in time series analysis. He argued that these plots can be informative about the relationship between shocks and the dependent variable at a given horizon, thereby promoting transparency in empirical work. He complimented the clarity of the paper but noted a potential inconsistency with respect to the proposed lag selection strategy: The authors discourage using AIC for lag selection in VARs but recommend it in the LP context. Given that one of the advantages of LPs is their ease of implementation, he asked whether a similarly simple lag selection rule could be proposed.

Christian Wolf responded to the issue of lag selection, addressing points raised by both Christiane Baumeister's discussion and Chodorow-Reich. He began by acknowledging the relevance of the paper referenced in Baumeister's discussion (*Local Projections Are VAR Predictions of Increasing Order*, J. Ludwig, 2024), which also evaluates LPs and VARs across a range of lag lengths. He emphasized that, once equivalence between the two methods is established (given the right number of lags), researchers must choose an objective function. He stressed that they adopt an approach centered on reliable uncertainty quantification — ensuring that reported confidence intervals have appropriate coverage. He argued that this is a natural and meaningful objective, though other researchers may reasonably prioritize different criteria.

Wolf also addressed Baumeister's point about long-run monetary policy effects. Even when one has strong priors about stationarity, small uncertainties in the timing of shock effects — even if statistically difficult to detect — can lead to substantial problems in VARs. This motivates their cautious approach to model specification and inference.

Returning to Chodorow-Reich's question, Wolf maintained that there is no internal inconsistency in their proposed strategy. He clarified that while they advise against using AIC to choose the lag length for a VAR used directly for inference, they recommend using an auxiliary VAR to guide lag selection for LPs. Finally, he pointed out that this auxiliary estimation is simple to implement and works well in the class of DGPs considered in the paper.

Mikkel Plagborg-Møller followed up to reinforce the central empirical finding of the paper: Small mistakes in lag selection are relatively harmless in LPs but potentially very costly in VARs. Using AIC to select the number of lags in an auxiliary VAR and applying that to LPs is therefore appropriate, since LPs are more robust to modest misspecification.

James Stock raised a related comment, emphasizing that it would be helpful to provide further guidance on the implementation of local projections with instrumental variables. While instrumental variable (IV) methods are often appealing for identification, they can introduce additional complexity, and users would benefit from more practical advice.

Wolf responded by clarifying that the paper explicitly avoids addressing the issue of weak IVs. However, in terms of general identification, the analysis in the paper closely matches the case of internal IVs, where the instrument is included among the observables. He highlighted that the conclusions of the paper extend to this case.

Next, Mark Watson highlighted a key distinction: If the research goal is to characterize uncertainty, misspecifying the lag length is damaging. However, if the objective is policy analysis or forecasting — that is, maximizing accuracy rather than ensuring proper coverage — then VARs may be preferable.

Christian Wolf agreed and pointed to earlier simulation work that shows how conclusions differ depending on the chosen criterion. For researchers focused on solving an optimal policy problem or minimizing mean squared error, VARs may perform better. He emphasized that the core issue is navigating the bias-variance trade-off, and that the paper offers one natural criterion relevant for applied macroeconomics.

Christiane Baumeister added a technical comment on the robustness of LPs to lag selection. She emphasized that LPs are effectively equivalent to VARs with many lags, which explains why including an additional lag has little effect on the estimates. The authors agreed.

Baumeister then described a numerical example based on the data-generating process (DGP) from the authors' earlier paper (Montiel Olea, Plagborg-Møller, Qian & Wolf, 2024). Using a Smets-Wouters model, she showed that once one maps the LP into its equivalent VAR representation at a given horizon, the VAR can outperform the LP in terms of coverage and even slightly in variance. She argued that, once one accepts that LPs are just composed of many VAR coefficients, the performance gap can disappear — even under the preferred coverage criterion.

The authors agreed with this interpretation and emphasized that their equivalence result hinges on using a sufficiently large number of lags. Only in this case do VARs achieve the same robustness to misspecification and reliable inference properties as LPs.

The authors mentioned that this long-lag equivalence is already discussed in their paper. They emphasized that the performance gap disappears, and VARs achieve the same robustness as LPs, precisely when VARs have enough lags, and only then.