# Research Statement

Mikkel Plagborg-Møller

February 2, 2025

I seek to develop and analyze practical econometric methods for empirical macroeconomics. My target audience includes applied macroeconomists in academia as well as in policy institutions. In one line of work, I have contributed to the literature on robust estimation of dynamic causal effects. In a second line of work, I have developed methods for estimating the feedback loop between microeconomic heterogeneity and the macroeconomy. In addition, I have worked on a number of applied papers.

#### 1 Dynamic causal effects

In macroeconomics, *impulse response functions* are often used to summarize the dynamic causal effects that surprise changes to policy or fundamentals have on macroeconomic variables. Over the past decade, there has been an active debate about the merits of two popular estimation procedures for impulse responses, Structural Vector Autoregressions (SVARs) and Local Projections (LPs). SVARs extrapolate longer-run impulse responses from short-run correlations in the data through a dynamic model, while LPs directly project future outcomes on current covariates.

In the paper "Local Projections and VARs Estimate the Same Impulse Responses" (Plagborg-Møller and Wolf, 2021, ECMA), we prove that SVARs and LPs in fact estimate exactly the same impulse responses in large samples. More precisely, the population estimands of these two procedures coincide, provided that the number of lagged control variables in the two specifications is infinite. Moreover, if the number of lags is finite, the two procedures have very similar estimands in empirically realistic settings. We conclude that the choice of LP vs. SVAR should not be viewed as a choice been conceptually different procedures, but rather as a choice of small-sample dimension reduction technique. To bolster this interpretation, we give several examples of how commonly used SVAR identification schemes can be equivalently implemented through LPs. The aim of our paper is to allow researchers to freely choose appropriate identification schemes for their applications without feeling constrained by the particular finite-sample estimation procedure they rely on.

The above-mentioned population equivalence result between SVARs and LPs provides a foundation for discussing how/why the methods differ in finite samples. We quantify the bias-variance trade-off between these methods through a comprehensive simulation study in "Local Projections vs. VARs: Lessons From Thousands of DGPs" (Li, Plagborg-Møller, and Wolf, 2024, *JOE*). To ensure the applied relevance of our simulations, we draw

thousands of data generating processes (DGPs) from an empirically estimated encompassing model of hundreds of U.S. macroeconomic time series. We then investigate the average performance across these many DGPs of several impulse response estimators. Our main findings are that (i) use of the conventional least-squares LP estimator can only be justified if the researcher is overwhelmingly concerned with bias (rather than variance); and (ii) shrinkage versions of LP and VAR estimators – which seek to smooth out the estimated impulse response functions – have substantially lower variance than the least-squares estimators while incurring only a moderate cost in terms of bias.

Though LP estimators are frequently outperformed by SVARs when it comes to mean squared error, we show in the working paper "Double Robustness of Local Projections and Some Unpleasant VARithmetic" (Montiel Olea, Plagborg-Møller, Qian, and Wolf, 2024) that LP confidence intervals are much more robust to dynamic misspecification than SVAR confidence intervals. We consider a class of locally misspecified VAR models, thus giving the SVAR procedure the "benefit of the doubt" by modeling any misspecification as asymptotically vanishing. LPs turn out to be surprisingly robust to such misspecification, a property that we link to the machine learning literature on double robustness in partially linear regressions. In contrast, our analytical characterization of the worst-case bias of SVAR estimators reveals that the associated confidence intervals suffer from severe coverage distortions even for misspecification that is quantitatively small, economically plausible, and difficult to detect statistically. Intuitively speaking, since even a small amount of bias is very costly for coverage, researchers who want accurate uncertainty assessments are forced to put high weight on bias when choosing an inference procedure, thus leading to LP being preferred.

Even under the assumption of a well-specified VAR model, LPs have key advantages over SVARs when it comes to frequentist inference. In the paper "Local Projection Inference Is Simpler and More Robust Than You Think" (Montiel Olea and Plagborg-Møller, 2021, ECMA), we show that confidence intervals for impulse responses based on LPs are robust to issues that applied researchers frequently encounter, namely high persistence of the data and long impulse response horizons. This contrasts with textbook VAR procedures, which are known to be sensitive to these issues. Importantly, we analyze a *lag-augmented* version of LP that controls for lags of the outcome variable and covariates in the regressions. We prove that, in such regressions, the usual confidence interval with heteroskedasticityrobust standard error and normal critical value controls coverage uniformly over both the impulse response horizon (up to a limit that increases with the sample size) and over the DGP (including processes with unit roots and cointegration). Our finding that heteroskedasticityrobust standard errors suffice is perhaps surprising, since applied practice has hitherto been to additionally correct for serial correlation in the regression errors. Such corrections are known to be fragile in realistic sample sizes, and require choosing among a menu of procedures with associated tuning parameters. Thus, our results not only provide a new perspective on the robustness of (lag-augmented) LP, they also simplify the implementation of this popular technique.

A much-studied problem in SVAR analysis is that researchers must exploit a priori economic restrictions to accurately estimate impulse responses, as these are not identified from the data alone. My job market paper "Bayesian Inference on Structural Impulse Re**sponse Functions**" (Plagborg-Møller, 2019, QE) proposes a flexible Bayesian approach to imposing prior information about impulse responses. My contribution is to work directly with a Structural Vector Moving Average (SVMA) model, which not only is more general than the conventional SVAR model, but also has the advantage that the parameters of the model are simply the impulse responses themselves. It is then straight-forward for practitioners to impose all kinds of prior information about impulse response functions, not just about signs or magnitudes, but also about shapes (e.g., smoothness or monotonicity). A second advantage of the SVMA model over the SVAR model is that it does not require the restrictive *invertibility* assumption that the econometrician's information set is the same as that of an economic agent who observes the underlying shocks. To deal with the high dimensionality of the parameter space of the SVMA model, I develop a fast and reliable posterior sampler based on Hamiltonian Monte Carlo and the computationally convenient Whittle approximation to the likelihood function.

In many applications, researchers are interested in the *shape* of an impulse response function. It is well known that conventional pointwise confidence intervals are insufficient for gauging the joint estimation uncertainty across multiple horizons, and various papers in the SVAR literature and elsewhere have proposed methods for constructing simultaneous confidence bands. In "Simultaneous confidence bands: Theory, implementation, and an application to SVARs" (Montiel Olea and Plagborg-Møller, 2019, JAE), we show that almost all these methods (asymptotically) fall into a one-parameter family of bands. In this family there is one band, the *sup-t band*, that is unambiguously superior to all the others. We prove that the sup-t band has a unique minimum regret property in an even larger class of bands, implying that this band is a good default choice for conveying joint uncertainty when the audience's preferences are not fully known. Moreover, we develop a novel Bayesian implementation of the sup-t band, to complement existing frequentist implementations.

Impulse responses, however, are not the only objects of interest to applied researchers. *Variance decompositions*, which quantify the relative importance of different latent shocks that drive the macroeconomy, are another key type of object used to distinguish between disparate theories of the business cycle. In "Instrumental Variable Identification of **Dynamic Variance Decompositions**" (Plagborg-Møller and Wolf, 2022, *JPE*), we provide identification results for such dynamic variance decompositions. We consider the increasingly popular identification approach in which researchers exploit a noisy *proxy* for the shock of interest (also known as an *external instrument*). While such proxies are frequently used in applied work to estimate impulse responses, it is not known how to use them to

identify variance decompositions without imposing the restrictive invertibility assumption that conventional SVAR analysis relies on. We show that, because the signal-to-noise ratio for the proxy is unknown *a priori*, the precise contribution of the latent shock of interest to the variance of macroeconomic observables is not point-identified. However, we are able to derive sharp and informative upper and lower bounds on this variance contribution. Point identification can be achieved under additional assumptions that are substantively weaker than the invertibility assumption. We also prove that the invertibility assumption is testable. In an empirical application, we obtain tight upper bounds on the contribution of monetary policy shocks to inflation dynamics in the U.S.

Since LPs and SVARs are motivated by linear models, a natural question is what they estimate when the true DGP is nonlinear. In the lecture "Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly" (Kolesár and Plagborg-Møller, 2024), we argue that LPs and SVARs that project on observed shocks (or proxies) are robust to nonlinearity, in the sense that their asymptotic estimand equals a convex weighted average of the true nonlinear causal effects, and the weights are easy to estimate and report. While this conclusion builds on several other papers, we are able to meaningfully relax their regularity conditions to show that the result applies very generally. Our analysis establishes some new results on identification of average marginal effects that draw connections to the cross-sectional causal literature. In contrast to the case with direct shock measures or proxies, we also prove that common approaches to identifying latent shocks – such as identification via heteroskedasticity or non-Gaussianity – are highly fragile in the face of nonlinearity. Two other shorter papers (Montiel Olea, Plagborg-Møller, and Qian, 2022, AEA P & P; Plagborg-Møller, 2022, JBES) further critique such identification approaches that exploit higher moments of the data or narrative shock restrictions in a likelihood framework.

While the literature on causal effect estimation in macroeconomics has evolved rapidly in the last 10 years, there is an unfortunate dearth of comprehensive review articles or textbooks covering these developments. In the working paper "Local Projections or VARs? A Primer for Macroeconomists" (Montiel Olea, Plagborg-Møller, Qian, and Wolf, 2025), we present a non-technical synthesis of the most important theoretical developments in this area. In addition, we list several concrete take-aways for applied practice that we hope can serve as a recipe book for empirical analyses of the consequences of shocks to economic policies or fundamentals.

#### 2 Heterogeneity and the macroeconomy

The last decade has witnessed an explosion of empirical and theoretical work on the interplay between microeconomic heterogeneity and the macroeconomy. This development was spurred by the increased availability of rich microeconomic data sets that complement traditional aggregate time series data, as well as by growing concern among policy-makers about economic inequality. There is now a need for new econometric methods for combining micro and macro data and for characterizing heterogeneity.

In "Standard Errors for Calibrated Parameters" (Cocci and Plagborg-Møller, 2024, *RESTUD*), we develop a procedure that facilitates inference on the parameters of structural economic models estimated by matching disparate empirical moments (here "moments" should be understood loosely to mean reduced-form parameters). This is a very popular estimation approach in diverse areas of economics, including heterogeneous agent macroeconomics. Unfortunately, traditional inference methods in this setting require the researcher to estimate the full variance-covariance matrix of all the empirical moments, which is often difficult or even impossible in practice if the empirical moments are obtained from different data sources (e.g., micro and macro moments), from different estimation methods, or from previous papers whose underlying data is unavailable. Yet, the marginal standard errors of the individual moments are usually readily available, since these are produced by conventional econometric software or reported in previous papers. We show that these marginal standard errors suffice for doing valid (conservative) inference on structural parameters. Moreover, we derive the optimal weighting of the moments that minimizes the worst-case standard error of the estimator; interestingly, this amounts to selecting a strict subset of the moments. Our methods give researchers the freedom to choose which moments to match based on economic considerations, without giving up on doing simple, valid, and informative inference.

In "Full-Information Estimation of Heterogeneous Agent Models Using Macro and Micro Data" (Liu and Plagborg-Møller, 2023, QE), we propose a method for doing efficient Bayesian inference in heterogeneous agent macro models, exploiting both aggregate time series data as well as repeated cross-sections of micro data. Existing structural estimation procedures fail to fully exploit the information content of such micro data, as they collapse the data to a pre-selected set of cross-sectional moments. We point out that most heterogeneous agent models imply a known functional form for the joint probability distribution of all the macro and micro data. Nevertheless, one encounters a key challenge when computing the likelihood function: according to these models, the cross-sectional micro distribution moves around over time as a function of unobserved macro state variables. Though the latent nature of the state variables renders exact computation of the likelihood infeasible, we develop a computationally convenient numerically unbiased estimate of the likelihood. Based on existing generic theoretical results, we then argue that the likelihood estimate can be incorporated into a standard Markov Chain Monte Carlo procedure, yielding valid and fully efficient Bayesian inference in finite samples. Our numerical illustrations demonstrate that micro data can be essential for doing informative inference on certain parameters in workhorse heterogeneous agent models.

In various areas of economics, researchers seek to estimate causal effects for many different cross-sectional units. For example, they may be interested in the heterogeneous effects that different neighborhoods in the U.S. have on intergenerational mobility, or the heterogeneous effects of monetary policy on employment in different industrial sectors. Because unrestricted individual effect estimates are typically noisy, it has become standard to report Empirical Bayes (EB) shrinkage estimates of the effects, which have lower variance at the expense of introducing some bias. However, no corresponding method for uncertainty quantification exists, without imposing strong parametric assumptions on the underlying distribution of causal effects. In **"Robust Empirical Bayes Confidence Intervals"** (Armstrong, Kolesár, and Plagborg-Møller, 2022, *ECMA*), we develop a robust Empirical Bayes confidence interval (EBCI) that is centered at the standard EB estimator. Our EBCI is often much narrower than the conventional "un-shrunk" confidence interval, while guaranteeing an "EB coverage" probability of at least  $1 - \alpha$  (i.e., under repeated sampling of both the data and the effects, as in a random effects model). The EBCI is robust in the sense that the EB coverage is controlled regardless of the true distribution of the effects. In addition, we prove that the robust EBCI has a desirable *frequentist* coverage guarantee: with the effect parameters treated as fixed, at least a fraction  $1 - \alpha$  of the many reported confidence intervals contain their respective true parameters, asymptotically.

Slightly further afield, the paper "Consistent factor estimation in dynamic factor models with structural instability" (Bates, Plagborg-Møller, Stock, and Watson, 2013, *JOE*) proves that the popular principal components estimator of latent factors is surprisingly robust to the presence of structural breaks when applied to large, heterogeneous macroeconomic panel data sets, such as multi-sector or multi-country data.

### 3 Applied work

My applied projects span diverse topics in macroeconomics and international economics.

In "Dominant Currency Paradigm" and a companion paper (Gopinath, Boz, Casas, Díez, Gourinchas, and Plagborg-Møller, 2020, AER; Boz, Gopinath, and Plagborg-Møller, 2019,  $AEA \ P \& P$ ), we document the dominance of the U.S. dollar in international trade invoicing and the consequences thereof. Empirically, the dollar is used as invoicing currency for a substantial fraction of trade that does not directly involve the U.S. as a trading partner. Using a sticky-price macro model, aggregate bilateral trade data, as well as micro data from Colombia, we show that dollar dominance fundamentally challenges the textbook Mundell-Fleming view of exchange rate pass-through, expenditure switching, and cross-country policy spillovers. As an example, our regressions imply that a 1% appreciation of the U.S. dollar against every other currency in the world predicts a 0.6% decline within a year in the total trade between countries other than the U.S., controlling for the global business cycle.

In "When is Growth at Risk?" (Plagborg-Møller, Reichlin, Ricco, and Hasenzagl, 2020, *BPEA*), we study the determinants of downside risk to GDP growth. Recent research suggests that various financial market indicators contain valuable advance signals about the potential for severe recessions. In our paper, we first document that these signals are

unreliable once we take into account the real-time information flow of the data that is available to policy-makers. Second, we show through a high-dimensional Bayesian modeling exercise that there are no individual financial indicators that robustly predict movements in the left tail of the GDP growth distribution, especially if one seeks results that generalize to multiple countries. Indeed, we argue that the available time series are simply not long enough to yield a confident assessment of the predictors of tail risks.

Finally, the paper "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve" (Mavroeidis, Plagborg-Møller, and Stock, 2014, JEL) surveys the voluminous literature that has estimated the New Keynesian Phillips Curve (NKPC). This workhorse inflation model is a component of almost every large-scale business cycle modeling framework used in central banks and academia. Unfortunately, empirical papers have reached widely different conclusions about the parameters of the NKPC, despite seemingly minor variations in data and specifications. In our survey, we provide an explanation for these disparities, emphasizing the difficulty of predicting future inflation, which results in a weak instruments problem. This fundamental issue causes different estimation procedures to be biased in systematically different directions. We provide a new, comprehensive set of empirical results by varying different aspects of the regression specifications, and we find that the weak instruments issue is so severe that one can (spuriously) obtain essentially any parameter estimates that one desires with some reasonable-looking specification. Unfortunately, more sophisticated inference procedures that are robust to weak instruments yield very wide confidence sets. We conclude that new data sources are needed to make meaningful progress on estimation of the NKPC.

## **Bibliography**

- ARMSTRONG, T. B., M. KOLESÁR, AND M. PLAGBORG-MØLLER (2022): "Robust Empirical Bayes Confidence Intervals," *Econometrica*, 90(6), 2567–2602.
- BATES, B. J., M. PLAGBORG-MØLLER, J. H. STOCK, AND M. W. WATSON (2013): "Consistent factor estimation in dynamic factor models with structural instability," *Journal of Econometrics*, 177(2), 289–304.
- BOZ, E., G. GOPINATH, AND M. PLAGBORG-MØLLER (2019): "Dollar Invoicing and the Heterogeneity of Exchange Rate Pass-Through," *AEA Papers and Proceedings*, 109, 527–532.
- COCCI, M. D., AND M. PLAGBORG-MØLLER (2024): "Standard Errors for Calibrated Parameters," *Review of Economic Studies*, forthcoming.
- GOPINATH, G., E. BOZ, C. CASAS, F. DÍEZ, P.-O. GOURINCHAS, AND M. PLAGBORG-MØLLER (2020): "Dominant Currency Paradigm," *American Economic Review*, 110(3), 677–719.
- KOLESÁR, M., AND M. PLAGBORG-MØLLER (2024): "Dynamic Causal Effects in a Nonlinear World: the Good, the Bad, and the Ugly," arXiv:2411.10415. Invited JBES Lecture at the 2025 North American Winter Meeting of the Econometric Society.
- LI, D., M. PLAGBORG-MØLLER, AND C. K. WOLF (2024): "Local projections vs. VARs: Lessons from thousands of DGPs," *Journal of Econometrics*, 244(2), 105722.
- LIU, L., AND M. PLAGBORG-MØLLER (2023): "Full-Information Estimation of Heterogeneous Agent Models Using Macro and Micro Data," *Quantitative Economics*, 14(1), 1–35.
- MAVROEIDIS, S., M. PLAGBORG-MØLLER, AND J. H. STOCK (2014): "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 52(1), 124–188.
- MONTIEL OLEA, J. L., AND M. PLAGBORG-MØLLER (2019): "Simultaneous confidence bands: Theory, implementation, and an application to SVARs," *Journal of Applied Econometrics*, 34(1), 1–17.
- (2021): "Local Projection Inference Is Simpler and More Robust Than You Think," *Econometrica*, 89(4), 1789–1823.
- MONTIEL OLEA, J. L., M. PLAGBORG-MØLLER, AND E. QIAN (2022): "SVAR Identification from Higher Moments: Has the Simultaneous Causality Problem Been Solved?," *AEA Papers and Proceedings*, 112, 481–485.

MONTIEL OLEA, J. L., M. PLAGBORG-MØLLER, E. QIAN, AND C. K. WOLF (2024): "Double Robustness of Local Projections and Some Unpleasant VARithmetic," arXiv:2405.09509.

(2025): "Local Projections or VARs? A Primer for Macroeconomists," prepared for the 2025 NBER Macroeconomics Annual.

- PLAGBORG-MØLLER, M. (2019): "Bayesian Inference on Structural Impulse Response Functions," Quantitative Economics, 10(1), 145–184.
- (2022): "Discussion of "Narrative Restrictions and Proxies" by Raffaella Giacomini, Toru Kitagawa, and Matthew Read," *Journal of Business & Economic Statistics*, 40(4), 1434–1437.
- PLAGBORG-MØLLER, M., L. REICHLIN, G. RICCO, AND T. HASENZAGL (2020): "When is Growth at Risk?," *Brookings Papers on Economic Activity*, 2020(Spring), 167–229.
- PLAGBORG-MØLLER, M., AND C. K. WOLF (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89(2), 955–980.
  - (2022): "Instrumental Variable Identification of Dynamic Variance Decompositions," *Journal of Political Economy*, 130(8), 2164–2202.