Addressing M. Eichenbaum's Comments

Essays in Dynamic Macroeconomics: from Structural Parameter Estimation to the Evaluation of Central Bank Policies

Juan Castellanos

February 2025

Contents

1. Chapter 1: Local Projections vs. VARs for Structural Parameter Estimation

I have already incorporated these comments and the references into the draft of the paper. Nonetheless, I will try to be exhaustive here on my take on how these papers relate to mine.

1.1. Structural VARs and their lag truncation bias

I really enjoyed reading Christiano, Eichenbaum, Vigfusson, Kehoe, and Watson (2006) and in particular Kehoe's comment. It is reassuring that my findings are consistent with their critique on what they called *the standard approach* (*IRF matching* using my vocabulary).

To start thinking about these issues in relation to my paper, a good place to start is their approximation error decomposition. Using the notation in my paper, we know that if we had an infinite sample of data from a model that satisfies the identification restrictions (e.g. TFP not responding to other variables within the period) or we were to observe the true innovation, then we would have that:

$$\beta^{VAR}(p = \infty, T = \infty|\Theta) = IRF(\Theta) \tag{1}$$

$$\beta^{LP}(T = \infty|\Theta) = IRF(\Theta) \tag{2}$$

where $IRF(\Theta)$ is the structural impulse response function. Then, the error in *IRF matching* relative to *Ind. Inf.* (Sims-Cogely-Nason approach) is simply the difference between $\beta(p, T|\Theta) - IRF(\Theta)$ where $p = \{2, 4, 8, 12\}$ and $T = \{100, 300\}$. As Chari, Kehoe, and McGrattan (2008) show, this error can be decomposed into

$$\left[\beta(p, T|\Theta) - \beta(p, T = \infty|\Theta)\right] + \left[\beta(p, T = \infty|\Theta) - IRF(\Theta)\right] \tag{3}$$

where the first term corresponds to the *small sample bias* and the second term is the *lag truncation bias*. Importantly, Local Projection estimators are independent of the lag-length, as shown by equation (2), and as a result the *lag truncation bias* is only relevant for the VAR approach. In fact, this decomposition helps us further understand why as we increase p the results from the *IRF matching* estimation using VAR get closer and closer to those of the LP. It is just simply that the lag truncation bias of the VARs shrinks as one increases p – see Table A1. Figure 1 illustrates this point. There I am plotting

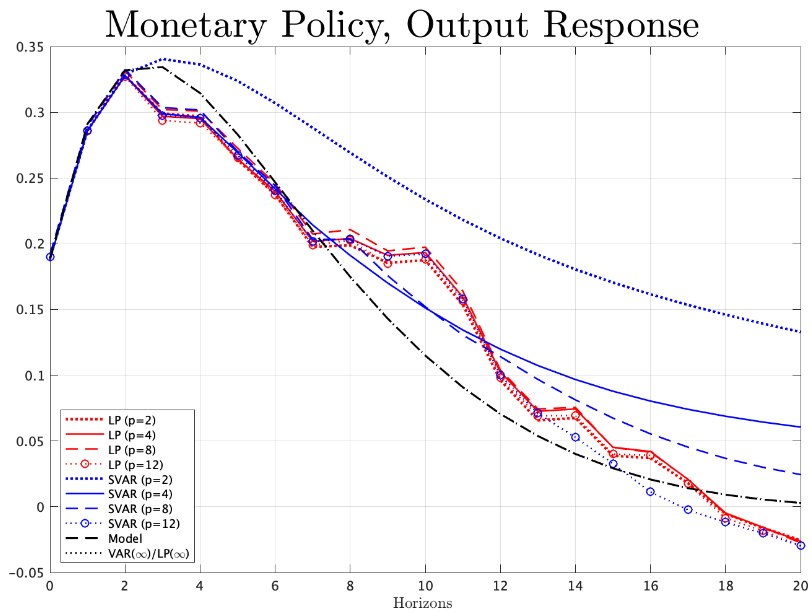

FIGURE 1. An illustration of the *lag truncation bias*

the estimated output response to a monetary policy shock using Local Projections (red lines) or VARs (blue lines) using various lag lengths and under the assumption that the shock is observed. The structural IRFs are depicted in black which indeed coincide with $VAR(\infty)/LP(\infty)$ as stated in (1) and (2). One can clearly see there that lag truncation bias is largest for a VAR(2) and starts decreasing as we increase p. It is not only that LPs and VARs approximately agree up to horizon p (Plagborg-Møller and Wolf 2021), but also that truncation bias shrinks as p increases. Loosely speaking, the difference between SVAR(2) and the SVAR(12) / LP(p) responses is mostly due to *lag truncation bias* while the difference between SVAR(12) / LP(p) estimates and the structural IRF is mostly *small sample bias*.

In Table A1, we also saw that *Indirect Inference* estimates using VARs were getting closer and closer to those that rely on LPs. In that case, we do not have the issue of comparing two statistically different objects when minimizing the distance in the optimization stage. Hence, these biases do not play a role. Instead, the explanation is related to the no-free lunch result in Olea, Plagborg-Møller, Qian, and Wolf (2024). As p gets larger, the bias of VAR responses shrinks but at the cost of getting larger variance. This point is illustrated in Figure A7.

1.2. Bayesian Limited Information Estimation

It sounds like an interesting extension given that as you point out part of the literature has moved towards Bayesian estimation. After digging a bit deeper into this I've seen that Bayesian Limited Information was introduced in Kim (2002), and that Christiano, Trabandt, and Walentin (2010) have provided a textbook type explanation of these methods, which then have been used not only by you in Christiano, Eichenbaum, and Trabandt (2016), but also in other influential papers like Bianchi, Ilut, and Saijo (2024). This already puts in perspective my limited knowledge on the topic. Hence, before attempting to address your comment, I would like to make a disclaimer: I am not a statistician and I also have very limited training in Bayesian estimation.

Having said that, my understanding is that the appeal from the Kim-CTW approach comes from combining Bayesian estimation with limited information techniques. In other words, maximum likelihood estimation is the frequentist counterpart of full bayesian estimation, while IRF matching / simulated method of moments / indirect inference are the frequentist counterparts of Bayesian Limited Information techniques. My conjecture then would be that the results that rely on the trade-offs between Local Projections and VARs as well as on the distinction between using structural or estimated IRFs – $\beta(\Theta)$ in my notation, $\psi(\Theta)$ in equation (38) of your paper – would still apply to this set-up and differences would arise mostly due to the frequentist vs. the bayesian approach.

2. Chapter 2: The Aggregate and Distributional Implications of Credit Shocks on Housing and Rental Markets

2.1. Rent stickinesses

Anecdotal evidence from Ireland – Andrew is Irish – suggests that the typical duration of rental contracts is about a year long. Part of the rent stickiness at lower frequencies is due to the prevalence of these one-year leases. In our model, one period is equivalent to one year. As a result, this type of rent stickiness wouldn't be present in our framework. Hence, stickiness would arise when landlords and tenants agree on the same price for the following year in a new contract.

Empirical evidence on rent stickiness is scatter due to data limitations. Gallin and Verbrugge (2019), based on their previous work, offer some empirical evidence on the matter for the U.S. during the period 1999-2008 and show that *"about half of rents in the US were unchanged after 12 months, and about one third of all rents were unchanged after two years"*. Indeed, nominal rents are sticky in the US! Unfortunately, similar data is not available to us for the case of Ireland.

Gallin and Verbrugge (2019) also show that rent stickiness arises through the landlord's choice of renegotiation every twelve months. Hence, simplifications like time dependent pricing or state dependent pricing through menu costs do not really apply to this market. We would have to model additional discrete choices into an already computationally costly household optimization problem.

Despite my slightly defensive position on this matter, it is spot on to worry about the reaction of rental prices as it drives most of our welfare implications. We are aware of that and in future versions of the paper we aim at replicating the "distance" regressions using model simulated data. This exercise will allow us to bring the model closer to the data and better capture the variation in both house and rental prices.

On a different note, thinking about these issues has made me wonder if introducing rent stickiness, in a way that respects the theoretical result in Gallin and Verbrugge (2019) which states that single-unit landlords tend to offer the same contract while multiple unit landlords tend to take the risk of renegotiating as the cost of vacancy averages out, may be a good way to think about co-existing large-scale investors and individual landlords within the model.

3. Chapter 3: The Role of Mortgage Fixation Periods for Macro-Prudential and Monetary Policies

3.1. Habit formation and sticky wages

We know from your work with Christiano and Evans that introducing habit formation in consumption into these class of models helps in getting the delayed, hump shaped response of consumption to a monetary policy shock. Moreover, we also know from your work that wage stickiness is the crucial nominal friction to generate persistent movements in output (Christiano, Eichenbaum, and Evans 2005). Hence, of course, including these features will help improve the fit of our model along those dimensions. Nonetheless, I would argue that as long as habit formation and wage stickiness do not interact with the interest fixation period, which I believe is the case, then it should affect all three economies considered – ARM, HRM and FRM – in a similar way. As a result, the qualitative implications from our analysis, e.g. the similar consumption responses to a transitory shock across the three economies, should still be present even in the case these elements were added into the model.

References

Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo. 2024. "Diagnostic Business Cycles." *Review of Economic Studies* 91 (1): 129–162.

Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan. 2008. "Are structural VARs with long-run restrictions useful in developing business cycle theory?" *Journal of Monetary Economics* 55 (8): 1337–1352.

Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt. 2016. "Unemployment and Business Cycles." *Econometrica* 84 (4): 1523–1569.

Christiano, Lawrence J., Martin Eichenbaum, Robert Vigfusson, Patrick J. Kehoe, and Mark W. Watson. 2006. "Assessing Structural VARs [with Comments and Discussion]." *NBER Macroeconomics Annual* 21: 1–105.

Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin. 2010. "DSGE Models for Monetary Policy Analysis." In *Handbook of Monetary Economics,* vol. 3, 285–367: Elsevier.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.

Gallin, Joshua, and Randal J. Verbrugge. 2019. "A theory of sticky rents: Search and bargaining with incomplete information." *Journal of Economic Theory* 183: 478–519.

Kim, Jae-Young. 2002. "Limited information likelihood and Bayesian analysis." *Journal of Econometrics* 107 (1): 175–193.

Olea, José Luis Montiel, Mikkel Plagborg-Møller, Eric Qian, and Christian K Wolf. 2024. "Double Robustness of Local Projections and Some Unpleasant VARithmetic." NBER Working Paper No. 32495.

Plagborg-Møller, Mikkel, and Christian K. Wolf. 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica* 89 (2): 955–980.