SVERIGES RIKSBANK
WORKING PAPER SERIES

451

Long-Lag VARs

Ferre De Graeve and Andreas Westermark

May 2025 (updated September 2025)

WORKING PAPERS ARE OBTAINABLE FROM

www.riksbank.se/en/research
Sveriges Riksbank • SE-103 37 Stockholm
Fax international: +46 8 21 05 31
Telephone international: +46 8 787 00 00

Long-Lag VARs*

Ferre De Graeve and Andreas Westermark

Sveriges Riksbank Working Paper Series No. 451

May 2025 (updated September 2025)

Abstract

Macroeconomic research often relies on structural vector autoregressions, (S)VARs, to uncover empirical regularities. Critics argue the method goes awry due to lag truncation: short lag lengths imply a poor approximation to important data-generating processes (e.g. DSGE models). Empirically, short lag length is deemed necessary as increased parametrization induces excessive uncertainty. The paper shows that this argument is incomplete. Longer lag length simultaneously reduces misspecification, which in turn reduces variance. Contrary to conventional wisdom, the trivial solution to the critique actually works. For data generated by frontier DSGE models long-lag VARs are feasible, reduce bias and variance, and have better mean-squared error. Long-lag VARs are also viable in common macroeconomic data and significantly change structural conclusions about the impact of technology and monetary policy shocks on the economy.

Keywords: VAR, SVAR, Lag length, Lag truncation

JEL: C18, E37

*This paper has benefited from discussions with Christiane Baumeister, Paul Beaudry, Fabio Canova, Yongsung Chang, Martin Eichenbaum, Jesús Fernández-Villaverde, Jonas Fisher, Jordi Galí, Karl Harmenberg, Marco Lippi, Jean-Paul L'Huillier, Per Krusell, Ellen McGrattan, Juan Rubio-Ramirez, Federico Ravenna, Barbara Rossi, Chris Sims, Eric Swanson, Dan Waggoner, Tao Zha and seminar participants at the Atlanta Fed, Cleveland Fed, Dallas Fed, EIEF, Universitat Pompeu Fabra, KU Leuven and NORMAC. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Executive Board of Sveriges Riksbank. Correspondence: ferre.degraeve@kuleuven.be, andreas.westermark@riksbank.se.

1 Introduction

Structural Vector Autoregressions ((S)VARs) have proven to be an important tool for measuring macroeconomic regularities. Following Sims' (1980) seminal contribution Bernanke (1983), Blanchard and Quah (1989), Sims (1989, 1992), Eichenbaum and Evans (1995), Gali (1999), Fisher (2006), Beaudry and Portier (2006) and many others since have provided VAR-based evidence for a variety of shocks and their macroeconomic effects.

Yet the VAR method is not without its critics. Many critiques of VARs boil down to the problem of lag truncation. Empirical macroeconomists invariably use a very small number of lags, typically at most up to four quarters or twelve months. But short-lag VARs are potentially problematic. DSGE models tend to imply reduced form VAR representations with long lag length (often infinity). Because lags are truncated, impulse response functions (IRFs) computed using VARs may not correspond to those of the underlying data-generating process. Chari, Kehoe and McGrattan (2008, henceforth CKM) is perhaps the most well-known elicitation of that critique.[1] Moreover, a recent series of papers (Li et al. (2024a), Montiel Olea et al. (2025)) rather strongly dismisses the use of VARs and advocates the use of Local Projections (LP) instead. In that work, almost every argument against VARs rests on lags being truncated.

The trivial solution to lag truncation, i.e., dramatically increasing the lag length, is unexplored. What keeps macroeconomists from using long lag lengths is the intuition that uncertainty becomes pervasive. That is, increasing the lag length increases the number of

[1] Others include Braun and Mittnik (1993), Faust and Leeper (1997), Cooley and Dwyer (1998) and Ravenna (2007).

parameters rapidly, thereby reducing the degrees of freedom and making confidence interval width prohibitively large.

We show that this standard intuition is only part of the story. In the face of misspecification due to lag truncation, increasing the lag length can actually reduce uncertainty. The reason is that as truncation reduces, misspecification reduces. The reduction in misspecification not only leads to the well-known bias reduction, but it also reduces variance. This reduction in variance works against the imprecision resulting from increased parametrization. This trade-off is general: it applies to all truncated VARs, no matter whether they are identified with short-run, long-run or other restrictions.

When increasing the lag length in standard VARs on realistically sized samples of data generated by standard DSGE models, the variance-effect of misspecification reduction often dominates the increased imprecision due to increased parametrization. The result is then much more in favor of long-lag VARs: reduced truncation bias, more precise inference, reduced MSE, better coverage rates.

We then document that various well-known empirical short-lag VAR studies allow much longer lag representations, without uncertainty becoming prohibitively large. Importantly, increasing the lag length significantly alters their structural economic conclusions.

Our first application considers the short term effect of technology shocks on hours worked. Gali (1999) shows hours drop significantly for the US economy as a whole, a finding that he interprets as supportive of New Keynesian (and starkly opposite to Real Business Cycle) models. Chang and Hong (2004) take issue with that finding and show that for the majority of US industries the impact is positive, thus feeding the disagreement between the two views. The above evidence is all based on short-lag VARs. Our analysis reconciles the two views: for long-lag VARs the aggregate evidence includes a zero response, while the sectoral evidence shows equal support for positive and negative responses. Long-lag VARs thus imply that the support for either view is less strong.

Our second application focuses on the effect of monetary policy shocks on the economy. Miranda-Agrippino and Ricco (2021) construct one of the cleanest measures of monetary policy shocks to date. Using their instrument in a short-lag VAR they find the effects of monetary policy on the economy are both substantial and long lasting. We show that these conclusions are not robust to lag truncation. As soon as one allows for the possibility that the lag length is two years or more, there is essentially no evidence that monetary policy has any significant effect on production, unemployment or prices at horizons beyond one year. While the real and nominal effects within the first year are still significant, they are substantially smaller: a 100bp. increase in the 1-year interest rate leads to fall in industrial production of -0.70% in long-lag VARs, compared to -1.36% in short-lag VARs.

Our work has several implications, both methodological and practical. First, the literature that compares the performance of VARs and local projections (LP) does not consider VARs with long lags a viable approach. While long-lag VARs have excellent theoretical properties (e.g. Plagborg-Møller and Wolf (2021)), they are deemed unworkable in practice due to high variance (e.g. Li et al. (2024a, 2024b)). We explain why the intuition underlying that dismissal is incomplete and establish that long-lag VARs are in fact empirically feasible.

Second, Ludwig (2024) derives the exact finite-sample mathematical mapping between VARs and LPs: LPs are a combination of VARs of increasing lag length. Consequently, appreciating the bias and variance of LPs requires understanding the bias and variance of VARs, also at long lag lengths, which our central result sheds light on. Ludwig's central theorem implies comparing LP and VAR for given lag lengths is misleading because they really cast a large model against a small one. Conditional on model size, there really is no bias-variance tradeoff between LPs and VARs. Our results add to that by showing that increasing the lag length in VARs reduces both bias and variance. In other words, for VARs increasing model size by adding lags may come with benefits on all sides. And by Ludwig's (2024) theorem, which shows how LPs are a combination of VARs of different lag lengths, LPs could benefit from that too, provided they use long lags.

Third, our results imply that a VAR with lag length p, VAR(p), should not be used for structural inference at horizons h larger than p. The reason is that inference is largely uninformative at that point: a VAR(p) cannot propagate much beyond horizon p, which implies that lag truncation leads to uncertainty being mechanically underestimated at those longer horizons. Coupled with potentially significant truncation bias, the short-lag VAR researcher runs the risk of becoming relatively certain about the wrong point estimate. Hence, if the research question centers on long horizon effects the VAR better include lags that extend that far. This helps appreciate the reasons behind the results of Baumeister (2025), who shows that the significant long term real effects of monetary policy found in Jordà et al. (2025) are likely an artefact of lag truncation.

Fourth, all VAR researchers should be concerned about lag truncation, regardless of the horizon of interest dictated by the research question. Truncation affects inference at all horizons, including the very short ones ($h \leqslant p$). Our simulations and applications reinforce that fact. This too sheds light on the performance of LPs. LPs in practice always consider short lag lengths. LP researchers in fact often rely on VAR lag selection criteria to pick LP lag length (see e.g. the discussion in Baumeister (2025)). By Ludwig's (2024) theorem LPs are therefore constructed from short-lag VARs. But because short-lag VARs are biased due to truncation, LPs necessarily inherit that bias. Long-lag VARs, in contrast, hedge against truncation bias at all horizons of interest.

The overarching implication of our investigation is that, contrary to conventional wisdom, long-lag VARs are feasible in practice. Applied research should use them broadly because they safeguard against the significant problem that lag truncation implies for structural economic conclusions. Information criteria focused purely on forecasting are poor guides for accurate structural inference. Forecasting is an out-of-sample exercise that benefits from parsimony (i.e. limited complexity). But structural inference, e.g. IRFs or conditional forecasting, is an in-sample exercise which benefits from added complexity or less parsimony. This trade-off between under vs. overfitting is a well-known fact in statistics/econometrics (e.g. Chatfield (1995)) and also central to model selection questions in e.g. machine learning. While long lags may not be relevant for all data and questions, it is up to the researcher to show that her structural results are robust to significant increases in the lag length, far beyond the lag lengths typically considered in current practice.

The paper is organized as follows. We start by laying out a standard single-equation omitted variables argument. This provides the intuition for the effect of reducing truncation in VAR impulse responses, where analytics are not tractable. We then assess long-lag VARs on the basis of a series of Monte Carlo experiments. We generate data from a variety of DSGE models, estimate VARs of different (and possibly very long) lag length and evaluate their performance. We then turn to the data and re-evaluate some well-known VAR results on the effect of technology and monetary policy shocks. Finally, we provide practical advice for the researcher that wishes to hedge against the negative effects of truncation.

2 Misspecification

We first briefly re-state a textbook omitted variables argument, which facilitates understanding the intuition behind the general VAR results.

2.1 Some useful single-equation intuition

Consider a data-generating process

$$y_t = X_{1t}\beta_1 + X_{2t}\beta_2 + \epsilon_t,\ V(\epsilon_t) = \sigma^2 \tag{1}$$

where a variable y is determined by two (sets of) exogenous variables, X_1 and X_2 and a shock ϵ. Now run the regression

$$y_t = X_{1t}b_1 + e_t,\ V(e_t) = s^2. \tag{2}$$

It is well-known that omission of the relevant variable X_2 leads to biased point estimates (unless $X_1 \perp X_2$):

$$E(b_1) \neq \beta_1$$

as well as an upwardly biased variance estimate (always):

$$s^2 > \sigma^2.$$

2.2 Omitted variables and truncation in VARs

The single-equation textbook result straightforwardly generalizes to VARs. It suffices to think of y as a vector of variables, X_1 as the lags the researcher includes, and X_2 as the lags not included, or truncated.

It is then immediate that a VAR, denoted by

$$Y_t = B_1 Y_{t-1} + ... + B_p Y_{t-p} + u_t,\ \ E(u_t u_t') = \Sigma$$

$$B(L) = B_1 L + ... + B_p L^p,$$

which has $p \ll p^*$ (where p^* denotes the true lag length) will suffer from truncation bias. The omitted variables argument above highlights why: lag truncation (or omitting relevant variables) results in a bias in the reduced form coefficients $B(L)$ *and* in the reduced form covariance matrix Σ. Any VAR analysis has impulse responses as a function of both these reduced form objects; let

$$IRF = f(B(L), \Sigma). \tag{3}$$

Because impulse responses are a function of both $B(L)$ and Σ they will tend to become less biased if both arguments become less biased. In other words, reducing truncation reduces bias.[2]

But what do we know about variance? Recall that the intuition that keeps macroeconomics from considering long lag lengths is that the increased parametrization (dimension of $B(L)$) leads to increased imprecision.

Though conceptually simple, equation (3) helps formalize that standard intuition. Essentially, recalling that $V(.)$ denotes variance, the intuition simply states that $V(B(L)) \uparrow \Longrightarrow$

[2] We merely refer to a documented tendency in DSGE models analyzed in the literature (see, for instance, CKM). From a theoretical perspective, this reduction in bias is not a certitude. Generally, bias reduction in its arguments does not guarantee bias reduction in the impulse response function. See Sims (1972) for an elicitation of a related point in terms of reduced form objects: convergence in individual point estimates (i.e. function arguments) may imply divergence of the sum of coefficients (i.e. the function itself).

$V(IRF) \uparrow$ as the lag length increases. But (3) also makes clear that this argument is incomplete. In particular, it neglects that there is a second argument, Σ. Therefore, any claims about $V(IRF)$ solely based on $V(B(L))$ are only partial. Importantly, the omitted variables argument suggests a reduction in bias of the estimate of Σ, which may well contribute to a reduction in variance of impulse responses.

Equation (3) also makes clear why general statements about $V(IRF)$ are hard to make: the non-linearity of f (also across horizons) interacts with the multi-dimensionality of both its arguments, $B(L)$ and Σ. Therefore, we ascertain the balance of this trade-off by means of a series of Monte Carlo experiments based on frequently studied models in macroeconomics.

3 Monte Carlo evidence

For each DSGE model considered, we generate data of length equal to that available in typical macro data samples ($T = 200$ quarters).[3] Given one such sample of data, we estimate VARs of different lag lengths, calculate impulse response functions and construct confidence bands.[4] We repeat that exercise 1000 times for each model and subsequently investigate bias, confidence interval (CI) width, mean-squared error and coverage rates.

[3] When comparing VARs of different lag-length, we ensure each VAR has the same number of effective observations, equal to $T = 170$. That is, lag initialization does not affect sample size. That said, our results do not hinge on this implementation detail.

[4] See Christiano et al. (2007) for a discussion of why this is the appropriate way to evaluate SVARs. Essentially, one takes an econometrician's perspective - who has only one draw of data and faces a question of inference on the basis of just that data.

3.1 Setup

We consider a range of models, both real and nominal, and identified with both short and long-run restrictions. More precisely, we consider estimating IRFs using long-run restrictions on data generated from CKM's RBC model as well as the short-run restriction version in Christiano et al. (2007), henceforth CEV, of that same model (in which agents do not observe the productivity shock at the time of making the labor decision). We consider both these models because they have taken center stage in much of the debate on the use of VARs. In addition, we also consider the Smets and Wouters (2007) model, henceforth SW, because it nests many shocks and frictions frequently discussed in macro and arguably captures dynamics deemed important in the data. As a simple way of building in a short-run restriction in that model, we assume that monetary policy responds only to lagged macroeconomic aggregates. The identifying restriction is then that only the monetary policy shock affects the interest rate contemporaneously. Because each of these models is well-known, we refer the reader to the respective papers for a precise description of model equations and parameter calibration/estimation.[5]

We work under a number of maintained simplifications. First, the identification assumptions are invariably correct (i.e., the long or short-run restrictions hold true in the data generating process (DGP)). Second, invertibility is never a problem; all the models we consider are fundamental (e.g. Fernandez-Villaverde et al. (2007)). Third, all our experiments

[5] For CKM and CEV, we follow the CKM baseline calibration. For SW we modify the policy rule to

$$r_t = \rho r_{t-1} + (1-\rho)\left\{r_\pi \pi_{t-1} + r_y\left(y_{t-1} - y_{t-1}^p\right)\right\} + \varepsilon_t^r$$

and calibrate the model at the median of SW's posterior distribution.

are based on two-shock models and two-variable VARs. Both RBC models fit that framework by construction, but the SW model does not. For the latter, we consider the model with only monetary policy and preference shocks. For the CKM and CEV model we run a VAR on labor productivity growth and hours in log-levels. The moment of interest is the response of hours to a productivity shock. For SW we run a VAR on GDP-growth and the interest rate, and the moment of interest is the IRF of GDP to a monetary policy shock. Finally, uncertainty bands are computed as in Sims and Zha (1999), Canova (2007) and Uhlig (2005). In particular, given a weak conjugate prior centered around their OLS estimates VARs have a posterior distribution of the Normal-Inverse Wishart form. Sims and Zha (1999) show that these provide good approximations to frequentist bootstrapped intervals.[6]

3.2 Results

Figure 1 contains, for each model, the median bias across all replications for VARs of different lag length. The figure resembles those found in the literature and shows how short lag length can imply substantial bias. Particularly, the short-lag VAR ($p = 4$) frequently exhibits the maximum bias at multiple horizons for the different models considered. Long lag length, or reduced truncation, can induce substantial bias reduction, most notably in CKM and, from intermediate horizons onward, in CEV and SW.[7] To evaluate if such biases are of concern,

[6]Our results also go through for bootstrapped confidence bands. Note furthermore that the priors we use do not put different weight on short vs. long lags, as one would in e.g. a Minnesota prior. We discuss alternative priors and inference procedures in Section 5.2.

[7]Note that the bias in CKM reduces fairly monotonically over horizons and lag length, but it does not in CEV and SW. This happens because in the latter transmission is more complicated, both in terms of dynamics (e.g. humps) and more endogenous model variables. Our investigation purely focuses on the effect

we now turn to measures of uncertainty.

Figure 1: Median bias in IRF across 1000 synthetic data sets from three DSGE models

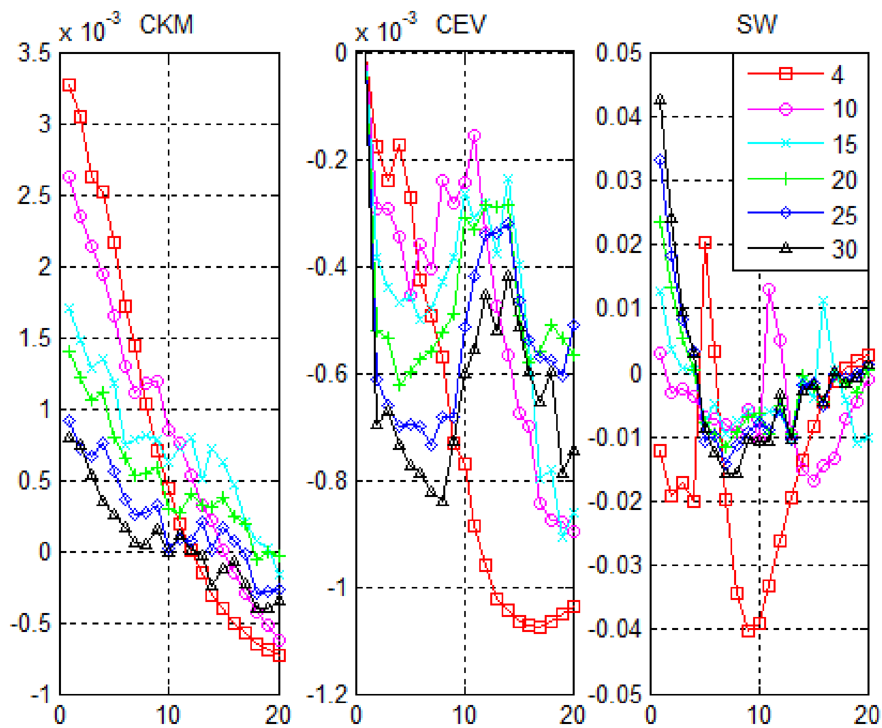

Note: Bias calculated as $IRF(VAR(p)) - IRF(DSGE)$. Horizontal axis is horizon in quarters. The IRF under investigation is the response of hours to a technology shock in CKM and CEV, and the response of GDP to a monetary policy shock for SW.

of lag length and its impact on bias and variance. The VARs could further benefit in terms of bias reduction from including omitted endogenous variables to fully capture the more complex transmission. See Braun and Mittnik (1993) for a discussion of how omitting relevant endogenous variables affects bias.

Result 1: Uncertainty does not explode with the lag length p Figure 2 plots a standard measure of uncertainty about IRF: the median width of the confidence bands.[8] A first glance at that figure reveals that, contrary to common wisdom, CI width does not explode as the lag length p increases. Instead, even for VARs with very long lags uncertainty bands are roughly in the same ballpark as those of short-lag VARs.

For longer horizons, short-lag VARs trivially attain minimum CI width. The reason is that a VAR(p) cannot propagate much beyond horizon p. As a result, uncertainty cannot propagate much beyond that horizon either. The consequence is, as apparent from Figure 2, that CI width mechanically converges to zero soon after horizon p.

Result 2: Short-lag VARs have maximal uncertainty for horizons where uncertainty is not mechanically low For short horizons short-lag VARs have maximal CI width. This holds true for each of the models considered. A possible reason for that to occur is that misspecification error is maximal for short-lag VARs. Individual reduced form coefficients may be estimated more precisely for a given draw, but across draws short-lag VARs have increased variance due to the misspecification of the VAR. Long-lag VARs, by contrast, may have individually imprecise reduced form coefficients, but they suffer much less from misspecification.

Result 3: Long-lag VARs have comparable coverage and comparable or better MSE than short-lag VARs Combined with a tendency to produce smaller biases, long-lag VARs have favorable properties compared to more standard short-lag VARs. Figure 3

[8] That is, for each generated sample we subtract the 5th percentile from the 95th, and then take the median across all draws. Results are similar for 68% credible intervals.

Figure 2: Median CI width across 1000 synthetic data sets from three DSGE models

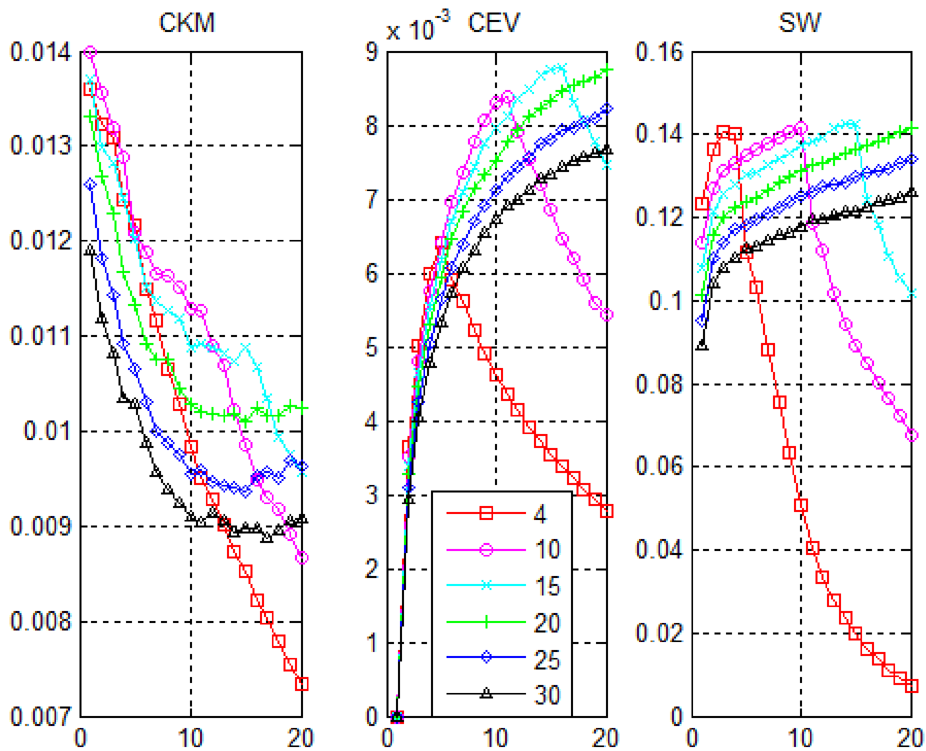

Note: 90% confidence interval width. The IRF under investigation is the response of hours to a technology shock in CKM and CEV, and the response of GDP to a monetary policy shock for SW.

documents how long-lag VARs attain coverage rates that are 1) reasonably good overall, 2) comparable to those for short-lag VARs for the CKM and CEV models, 3) much better for the SW model, where short-lag VARs with short run restrictions go astray entirely.[9]

Figure 4 combines bias and variance in a different way, by plotting mean-squared errors

[9]The huge swings in coverage for short-lag VARs in SW arise as the combination of substantial bias and mechanically low uncertainty. As a result, from intermediate horizons onward, the econometrician becomes relatively certain about the wrong point.

Figure 3: Coverage across 1000 synthetic data sets from three DSGE models

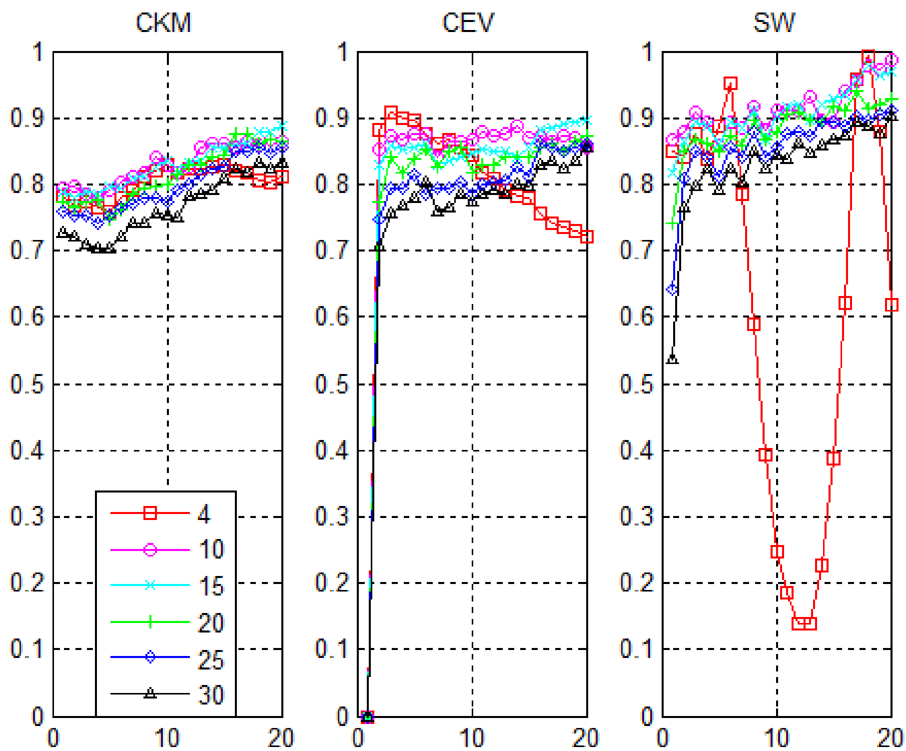

Note: 90% coverage. The IRF under investigation is the response of hours to a technology shock in CKM and CEV, and the response of GDP to a monetary policy shock for SW.

(MSE) across horizons. The message is very much the same: at short horizons - where uncertainty does not mechanically shrink - short-lag VARs are either comparable or considerably worse than long-lag VARs.

3.3 Decomposing uncertainty effects

From the above results it may not be obvious that standard intuition - increased parametrization leading to increased uncertainty - holds at all. We here provide a decomposition to

Figure 4: Median MSE across 1000 synthetic data sets from three DSGE models

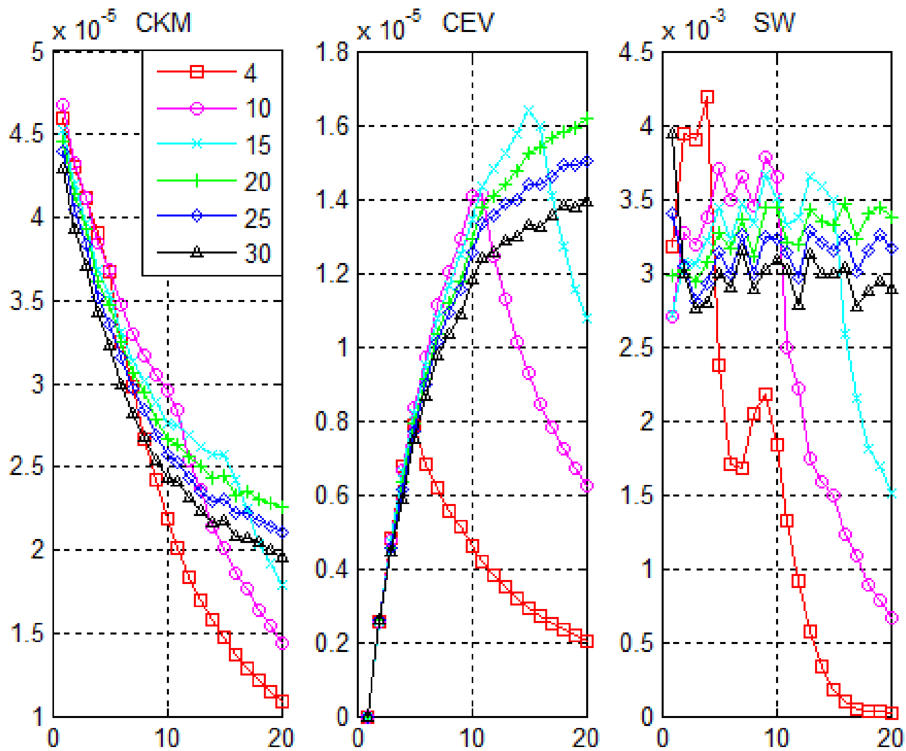

Note: Mean-squared error (MSE). The IRF under investigation is the response of hours to a technology shock in CKM and CEV, and the response of GDP to a monetary policy shock for SW.

measure the impact of the standard intuition on the total variance effect.

Figure 5 plots the Monte Carlo distribution of CI width for three types of impulse responses. Specifically, for each sample of data from the DSGE model, we measure the CI width around the contemporaneous impulse response for CKM and SW, and the second horizon for CEV.[10] The medians of these two distributions are already contained in Figure

[10]While similar effects are at work at longer horizons for all models considered, they are harder to disentangle due to the mechanical reduction in uncertainty for short-lag VARs, as apparent in Figure 2. For CEV

2. For short-lag VARs, here illustrated by a VAR(4), the dashed line (B_4, Σ_4) plots the distribution of CI widths across all 1000 draws, based on the lag polynomial $B(L)$ and covariance matrix Σ of a short-lag VAR(4), denoted by B_4 and Σ_4. Similarly, the solid line (B_{30}, Σ_{30}) plots the distribution of CI widths for a long-lag VAR, here illustrated by a VAR(30), based on its lag polynomial $B(L)$ and covariance matrix Σ, denoted by B_{30} and Σ_{30}. Comparing these two distributions confirms the earlier results: long-lag VARs do not necessarily imply overwhelmingly dispersed uncertainty bands.

To understand why, and to relate our results to the standard intuition, we construct the following counterfactual impulse responses:

$$IRF = f(B_{30}(L), \Sigma_4).$$

These hypothetical IRFs are constructed using the (many) reduced form coefficients of a long-lag VAR, $B_{30}(L)$, combined with the reduced form covariance matrix of a short-lag VAR Σ_4. Such IRFs can be interpreted as isolating the effect of increased parametrization. They shut down the effect of misspecification reduction by ignoring the reduced bias in Σ. The dotted (B_{30}, Σ_4) distributions in Figure 5 show the CI width associated with these counterfactual impulse responses. Standard intuition dictates that long lag length makes the entire distribution shift outward, through the additional uncertainty created by the strong increase in number of parameters.

It is immediately apparent that, across models, the dotted distribution does not unequivocally lie to the right of the dashed distribution. In other words, the strong increase in number of parameters need not imply an increase in uncertainty. For the SW model,

the contemporaneous ($h = 0$) response of hours to technology shocks is subject to a zero restriction and is thus uninformative. The figure therefore contains the IRF uncertainty distribution for the horizon $h = 1$.

Figure 5: CI width distribution across 1000 synthetic data sets from three DSGE models

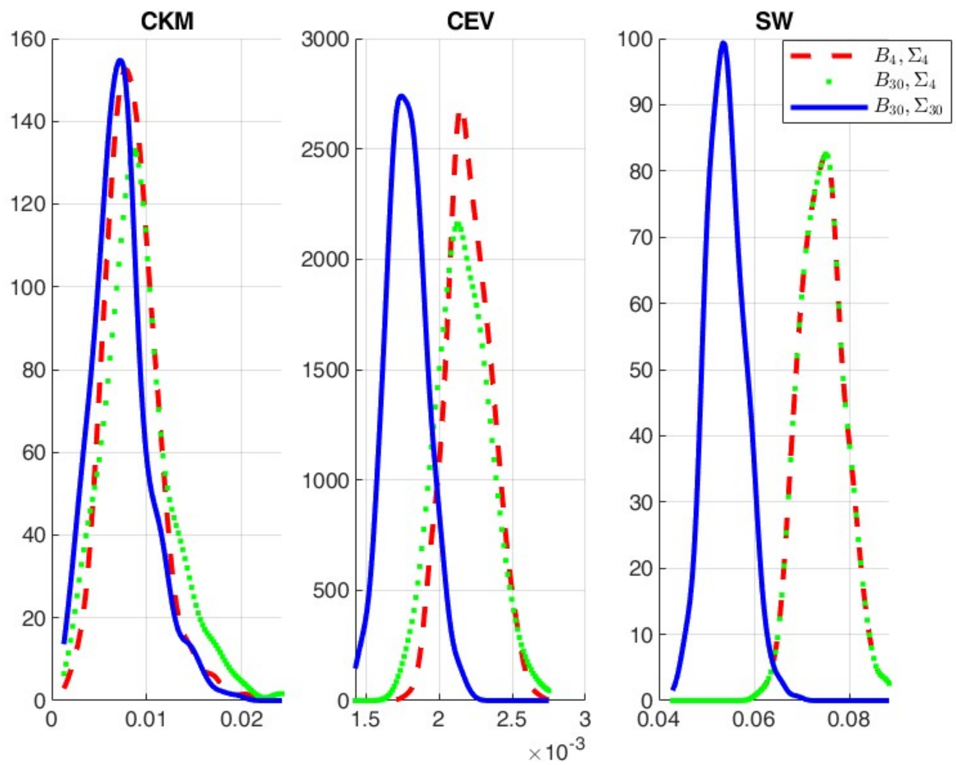

Note: CI width (90%). The IRF under investigation is the response of hours to a technology shock in CKM and CEV, and the response of GDP to a monetary policy shock for SW. For CKM and SW horizon $h = 0$. For CEV horizon $h = 1$.

there is no effect at all from increased parametrization, since the short-run restriction implies that the contemporaneous IRF only depends on Σ and not on $B(L)$. For the CEV and CKM models the right tail of the CI width distribution becomes fatter, as standard intuition would suggest. However, two observations stand out. First, the increase in CI width is not overwhelming. Second, a significant portion of the mass is shifting to the left of the dashed, short-lag distribution, indicating reduced uncertainty.

The fact that increased parametrization does not invariably increase uncertainty is at odds both with standard intuition (less degrees of freedom) and with a well-known omitted variables result. Particularly, coefficient estimates b_1 in (2) are not only biased, but also have too low variance. Intuitively, to the extent that omitted variables correlate with included ones, the explanatory power of those included will appear to be larger than it really is. Analytically, if we denote the coefficients on X_1 in the correct regression (which does include X_2) by $b_{1.2}$, then

$$Var(b_1) < Var(b_{1.2}). \tag{4}$$

This suggests that by including additional relevant variables one increases the variance of coefficients. We now provide detail on the effects in each of the individual models, which will lay bare the reasons for these seemingly counterintuitive results.

Let us start with the SW model in Figure 5. As mentioned above, since identification is based on short run restrictions, contemporaneous IRFs are not a function of $B(L)$, only of Σ. Hence, the dashed and dotted lines overlap. The effect of misspecification reduction, on the other hand, substantially reduces uncertainty, as can be seen by the shift to the solid distribution.

Now consider the CI width distribution for the CEV model. Here, taking into account the long-lag polynomial clearly only partially results in an increase in uncertainty measures. To see the reason for this, note that IRFs are functions involving multiple coefficients. As a result, covariance between coefficients becomes an issue. For the sake of argument, consider the simplest possible function involving two parameters in (1), their sum. Let $X_1 = [X_{1a}, X_{1b}]$ and denote the corresponding point estimates by b_{1a} and b_{1b}. Then the variance of the sum of the two coefficients in b_1 in the equation that omits X_2 is

$$V(b_{1a} + b_{1b}) = V(b_{1a}) + V(b_{1b}) + 2Cov(b_{1a}, b_{1b}). \tag{5}$$

Similarly, the variance of the sum in the correct regression (which includes X_2) is

$$V(b_{1a.2} + b_{1b.2}) = V(b_{1a.2}) + V(b_{1b.2}) + 2Cov(b_{1a.2}, b_{1b.2}). \tag{6}$$

While we know that each of the first two terms is smaller in (5) than the corresponding terms in (6), the presence of the covariances prevents any automatic conclusion on whether $V(b_{1a} + b_{1b}) \lesseqgtr V(b_{1a.2} + b_{1b.2})$.

Thus, as soon as one considers functions that combine coefficients of a regression subject to omitted variables, the usual variance relation in (4) can break down. This explains the shift from the dashed to the dotted distribution in the CEV model, and particularly why there can be significant mass shifting towards lower uncertainty despite having a big increase in the number of parameters.

The quantitatively more important effect on uncertainty is not due to the big increase in parametrization, however, but rather the effect of the reduction in misspecification. This is illustrated by the shift from the dotted to the solid distribution.

Finally, consider the CKM model in Figure 5. The dotted line in the figure shows how increased parametrization, along the lines of standard intuition, tends to shift the distribution of uncertainty outward compared to the short-lag VAR. Here, too, there is some mass that shifts leftward. As in the case of the CEV model, this can occur because IRFs involve a combination of parameters.[11] Despite the push toward increased uncertainty following the

[11] The reduction in uncertainty in the dotted distribution can also be the result of reduced misspecification

increase in number of parameters, once the misspecification effect through Σ is incorporated long-lag VARs are associated with smaller, not larger uncertainty.

Thus, the figures show the uncertainty trade-off: increased parametrization ($B_4 \rightarrow B_{30}$) which can - but need not - push the distribution outward (from dashed to dotted) vs. reduced misspecification ($\Sigma_4 \rightarrow \Sigma_{30}$) which shrinks uncertainty and thus pulls the distribution to the left (from dotted to solid). In sum, while standard intuition on increased parametrization is partially correct and clearly part of the story, misspecification reduction tends to have more substantial variance effects. As a result, for VARs on data generated by standard DSGE models, the total effect of increasing the lag length can easily imply a reduction in variance. While the simulation setup provides ample insight into VAR inference in typical macroeconomic samples, the empirical macroeconomist does not have the luxury of knowing the DGP. We now ask if long-lag VARs are also viable in the data by revisiting a few well-known empirical VAR studies.

4 Empirical applications

4.1 Application 1: Technology shocks and hours

Much of the debate about the use and pitfalls of VARs centered on the question of the hours response to technology shocks. Gali's (1999) finding that hours fall after a positive technology shock was met with severe backlash. On the theoretical front, CKM stressed that

in $B(1)$, documented by Sims (1972), in combination with long-run identifying restrictions. This effect exists because $B(1)$ enters the identification procedure in the case of long-run restrictions. For more on the importance of $B(1)$, see Christiano et al. (2004).

lag truncation bias makes short-lag VARs with long-run identification restrictions unable to recover true IRF. On the empirical front, Chang and Hong (2006) applied Gali's approach to 458 US manufacturing sectors and found that most sectors exhibit a positive rather than a negative hours response. With the knowledge that long-lag VARs overcome the truncation problem and may still allow precise inference in realistic samples we now revisit that debate. Specifically, we ask whether long-lag versions of the VARs of Gali (1999) and Chang and Hong (2006) imply excessive uncertainty and whether their conclusions are different from the original short-lag versions. We use the same samples as the original studies: quarterly US data 1948:1-1994:4 for Gali (1999; 4 quarterly lags) and annual sectoral US data 1958-1996 for Chang and Hong (2006; 1 annual lag). We use the same Bayesian methods as in the Monte Carlo evaluations.

Figure 6 shows the Gali (1999) result for VARs of increasing lag length. The estimated impact effect of technology shocks on hours is invariably negative, regardless of the lag length. This effect is significant for VARs with (quarterly) lag lengths up to $p = 20$ (5 years). At longer lag lengths (e.g. $p = 28, 36$) the effect turns insignificant. The bottom row shows that CI width does not obviously increase with the lag length. While the VAR(4) has minimal variance at horizon $h = 0$, the long-lag VARs often have smaller variance at horizons $h = 2, 3, 4$. At longer horizons $h > p$ mechanical effects start kicking in for short-lag VARs. Broadly speaking, however, for horizons $h < p$ VARs of different lag length have largely comparable CI width.

Figure 7 (top panel, dashed line) contains our replication of Chang and Hong's (2006) result: the majority of sectors see an increase in hours in the wake of a positive technology shock. The bottom panel shows the associated cross-sectional distribution of CI width.

Figure 6: Hours response to technology shock: Gali (1999)

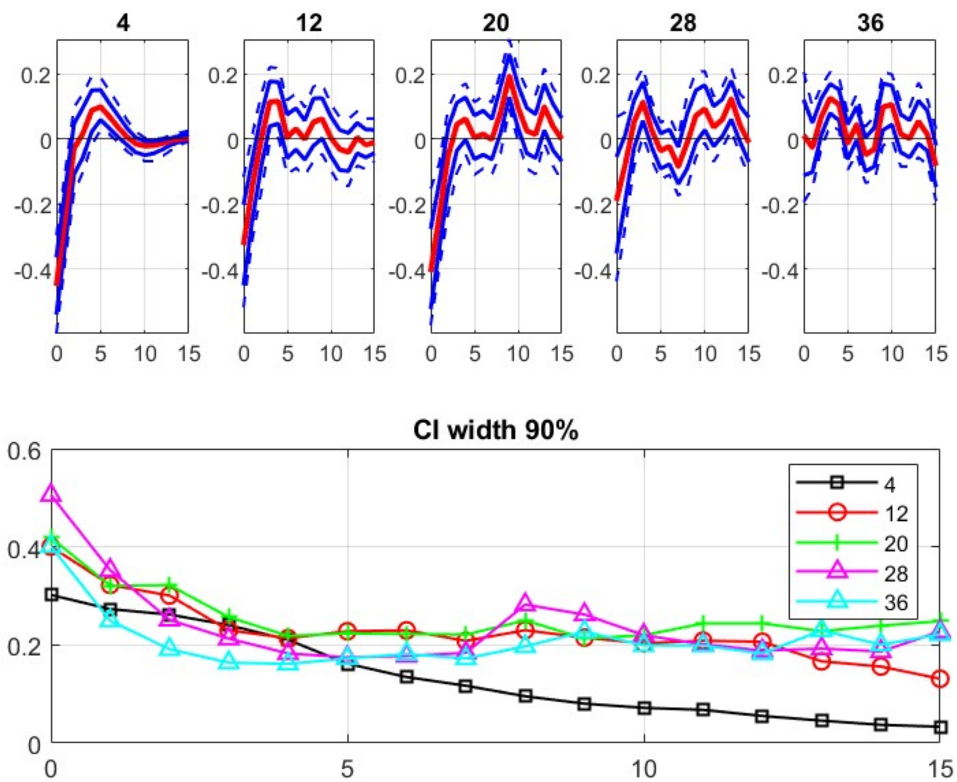

Note: The top row shows IRF and CI (68% and 90%) for VARs of increasing lag length. The bottom row overlays CI width (90%) for the different VARs.

Combined, the short-lag VAR finds 203 significantly positive sectors and only 46 significantly negative (Table 1).

The results for the long-lag VAR are striking. First, the point estimates differ dramatically. The cross-sectional distribution of IRFs shifts substantially to the left in Figure 7 (top panel, solid line). While short-lag VARs estimate 70% of sectors (Table 1: 317/458) respond positively, long-lag VARs find 50% of sectors respond positively and the other half

Figure 7: Hours response to technology shock: Chang and Hong (2006)

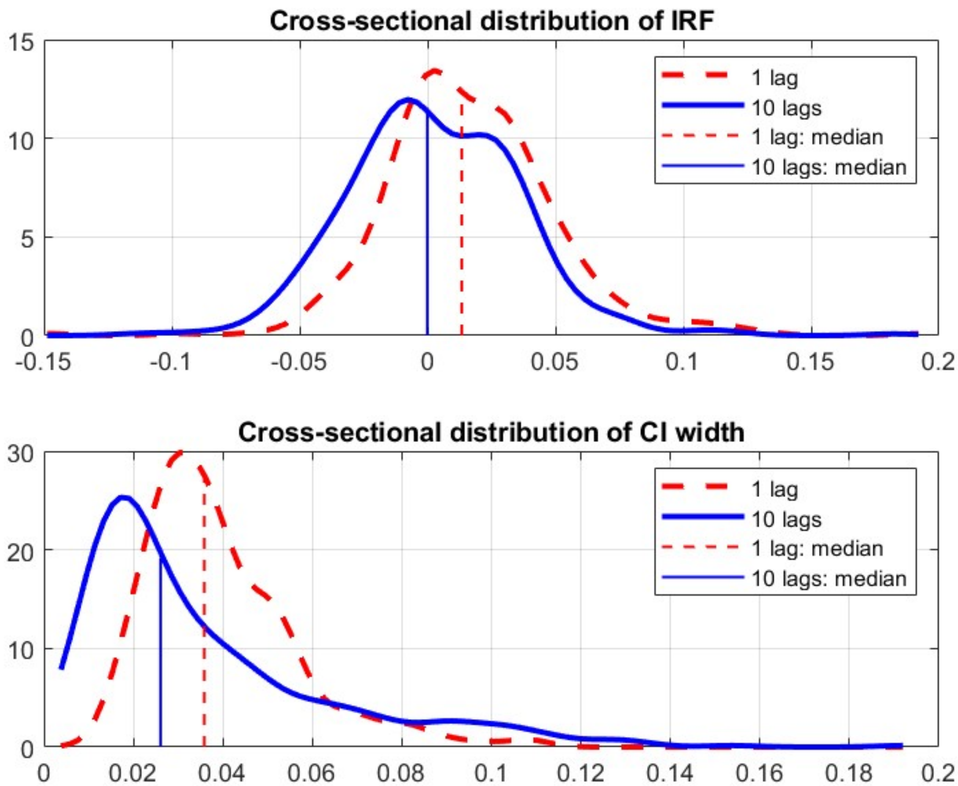

negatively. Second, variance does not become excessively large. The bottom panel of Figure 7 shows that long-lag VARs have mostly lower variance than their short-lag counterpart: the cross-sectional distribution of CI width predominantly shifts to the left for long lags. Combined, the number of industries that differ significantly from zero is almost the same for short and long-lag VARs (249 and 244 respectively). Hence, once again, long-lag VARs do not exhibit excessive variance, rather the opposite. While there is a stark contrast between the short-lag results of Gali (1999) and Chang and Hong (2006), long-lag results seem more in line with one another. In contrast to the conflicting results from short-lag VARs, long-lag

VARs support both salt and freshwater camps.

Table 1: Short-run hours response to technology shock - Chang and Hong (2006)

	1 lag			10 lags		
	Number of industries			Number of industries		
	Positive	Negative	Total	Positive	Negative	Total
Estimate (median)	317	141	458	229	229	458
Significant at 10%	203	46	249	128	116	244

We take from this evidence that long-lag VARs are a viable, feasible tool in the macro empiricist's toolkit. Long-lag VARs do not imply excessive variance even in realistic samples. The evidence also suggests that there are substantial effects on point estimates, confirming the importance of truncation bias also in practice.

4.2 Application 2: Monetary policy shocks

Perhaps most of the development of VAR methods involves the study of monetary policy shocks. We now revisit a state-of-the-art monetary policy VAR and again ask to what extent long-lag VARs change inference. Our analysis revisits the influential work of Miranda-Agrippino and Ricco (2021), who construct one of the most convincing instruments for monetary policy shocks to date. Through the lens of a short-lag VAR (i.e. a VAR(12) with monthly data), these shocks have significant and long-lasting effects on real and nominal variables. The thick line with shaded confidence band in Figure 8 reproduces that result (i.e. Figure 3 in Miranda-Agrippino and Ricco (2021)). The negative effects on industrial production and CPI are immediate and remain significant for at least two years. Point estimates at the two year horizon are even more negative than the immediate impacts and are invariably significant.

Figure 8: Monetary policy shock: Miranda-Agrippino and Ricco (2021)

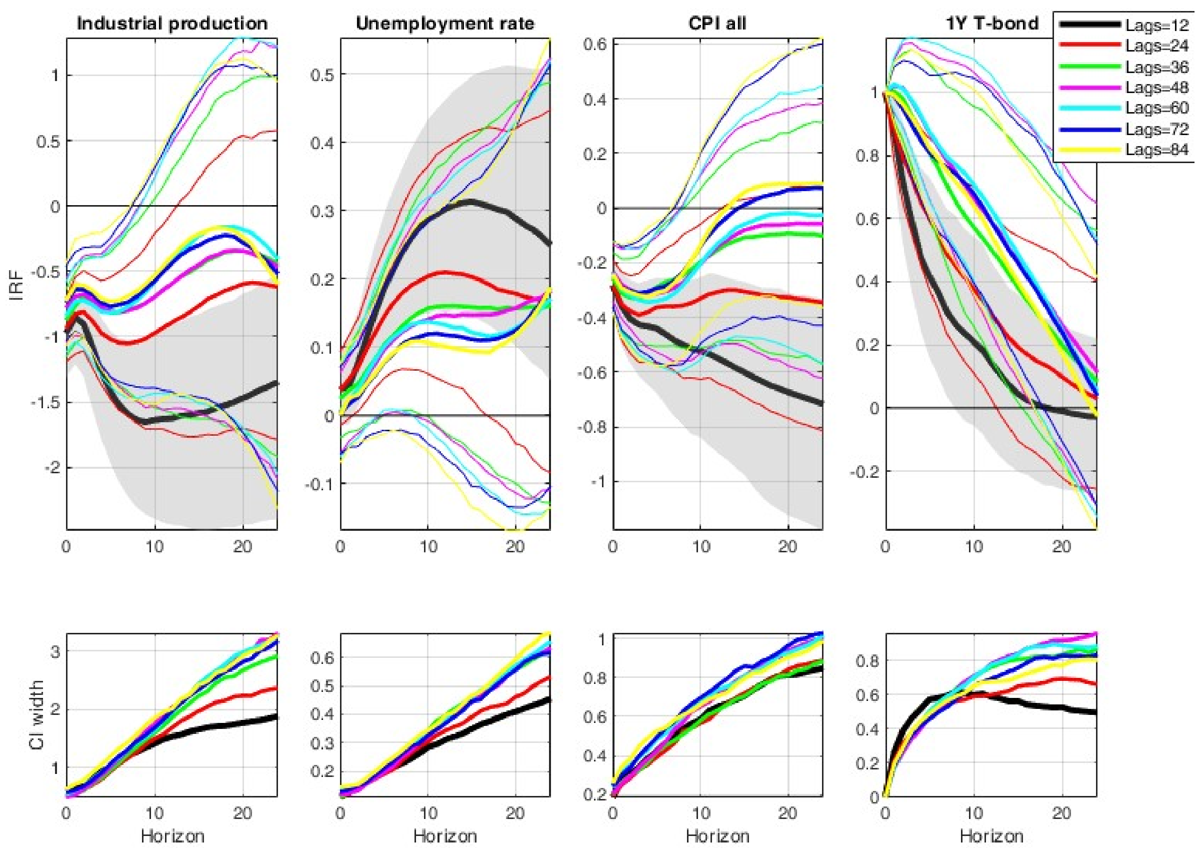

Note: Thick black line and grey band: VAR(12). Longer lag VARs in coloured lines: thick median, thin confidence bands.

The non-black lines in Figure 8 show IRFs based on VARs with increasing lag length.[12]

[12] We use the replication codes of Miranda-Aggripino and Ricco (2021) and change only the lag length.

As soon as the VAR contains lags of two years and more, there is no evidence that monetary policy has effects on industrial production or CPI that extend further than a year out. All the confidence bands include zero for any horizon $h > 13$. Even a VAR(24) finds IRFs that disagree with those of the VAR(12): their point estimates lie outside of each others' confidence bands for multiple horizons. VARs with lags of three years and more strengthen that conclusion, and are also very close to one another in terms of point estimates.

It is worth emphasizing that the finding that monetary policy has no significant effects two years out is not the result of prohibitively large uncertainty in long-lag VARs. As the lower panel in Figure 8 shows, CI width is comparable for short and long-lag VARs at all horizons where uncertainty is not mechanically low for short-lag VARs.[13] The main change in conclusions comes from substantial differences in point estimates. The within year one average point estimate of a 100bp. rate increase on industrial production for long-lag VARs robustly hovers around -0.7%, while the short-lag VAR puts it at -1.36%.

The long-lag VAR researcher can find comfort in the fact that her conclusions are robust to perturbations in the lag length. The short-lag VAR researcher, by contrast, must admit her results are fragile: even changing lags from one to two years significantly affects conclusions. Our results complement the findings of Baumeister (2025) in an almost identical VAR setup. She argues the very long term (around 8-10 years) significant effects of monetary policy

[13] It is worth pointing out that the mechanical breakdown in uncertainty documented in the simulations in Section 3 pertained to stationary VARs, where e.g. labor productivity and GDP are specified in growth rates. In the present application the VAR is specified in (log) levels and thus has roots close to one. The bottom row of Figure 8 illustrates that the mechanical breakdown in uncertainty of short-lag VARs *relative* to long-lag VARs also occurs in such cases.

in Jordà et al. (2025) are purely due to lag truncation. Her argument as well as our simulations suggest it is unwise to trust long-horizon implications of short-lag VARs, as inference becomes mechanical. On top of that, the present application shows that important short term conclusions are also subject to lag truncation. Long-lag VARs suggest there are no discernible effects of monetary policy farther than one year out. And even within the first year, long-lag VARs estimate real and nominal effects of monetary policy that are substantially smaller than what short-lag VARs find.

5 Discussion and practical advice

5.1 On choosing the lag length

All of the above results are in terms of structural inference. None of our results imply that long-lag VARs ought to be used for matters such as forecasting. For instance, the large dimensionality of the lag polynomial in long-lag VARs prohibits any success in forecasting due to the lack of parsimony. While one can certainly envisage ways to reduce the dimensionality, that is not the issue here. Rather, if one wants to draw structural conclusions, e.g. by means of IRFs, then misspecification concerns are essential. Therefore, if forecasting is not the main purpose of the model, it may be ill-advised to trust lag-selection criteria which purely focus on forecasting/parsimony. Long-lag VARs should be in the set of models considered and not ruled out based on forecasting criteria. Recent work by González-Casasús and Schorfheide (2025) takes significant steps in that direction by proposing criteria that acknowledge the model will also be used for structural inference.

Yet the applied researcher does not need to wait for the field to settle on ideal novel criteria. A simple heuristic approach researchers can take is to plot the structural moment(s) central to the research question against lag length. If the structural conclusion dramatically changes as the lag length increases, that should raise a red flag and is a clear indication truncation is important for the question at hand. Settling on a longer lag length where structural conclusions are robust to perturbations in the lag length hedges against truncation effects in short-lag VARs. The researcher can lean on the traditional econometric adage that adding potentially irrelevant variables (here: lags) is better than omitting relevant variables. The presumption that adding more lags necessarily leads to prohibitively large uncertainty is hereby disproven. This practice would easily constitute an improvement over the current widespread approach that either leaves lag choice up to forecasting criteria, or where researchers just state they use one year's worth of lags and we all accept that as appropriate for macroeconomic data. Our study of monetary policy shocks exemplifies how this heuristic approach is directly useful in practice and meaningfully affects structural inference.

5.2 On model size and inference procedures

The variance trade-off we document applies to all VAR approaches that rely on a correct (untruncated) reduced form. It does not depend on the number of variables in the VAR, nor on whether the researcher uses frequentist or Bayesian methods, nor on the identification assumptions.

Yet a frequentist VAR researcher will, as the lag length increases, ultimately bump into a degrees of freedom problem. An n-variable VAR(p) has $n + n^2p$ coefficients and $\frac{n(n+1)}{2}$ elements in the variance-covariance matrix. As lags grow, the number of parameters to be estimated increases rapidly, and more so when there are more variables in the VAR. At some point there just are not enough observations to estimate all these parameters. While this is a relevant theoretical consideration, short-lag VARs with one year of lags are far removed from the theoretical limit for typical macroeconomic data. To see that, let us consider Sims (1980) as an indicative benchmark - it is, after all, the study that brought VARs to the fore. In a six-variable VAR with up to two year lags he has roughly two observations per estimated parameter. That is similar to the ratio one obtains in both our technology 2-variable VARs with ten year lags. Entertaining larger systems poses no practical difficulty either. The short-lag VAR evaluation of Li et al. (2024a), for instance, focuses on 5-variable VARs for a quarterly dataset spanning 64 years. It takes more than six year lags before the number of parameters reaches half the number of data points. They thus easily allow investigating the effect of adding several years of lags. This comparison of data points to number of parameters is not meant as a target for how far to extend lags. Instead, it illustrates that with realistic macroeconomic data samples the VAR researcher already has ample flexibility to incorporate much longer lag lengths.

Importantly, four and a half decades of empirical macro research has jumped on the suggestion of Sims (1980) and developed various methods for dimension reduction in VAR models. All of those methods enable incoporating longer lags easily.

On the frequentist side, those working with many variables often adopt dynamic factor models or FAVARs (see e.g. Stock and Watson (2010) for an overview). And adding lots of lags to the dynamic factor equation comes at low cost in terms of increased parametrization.

The advent of big data has also seen ample recent research in economics and beyond on Lasso and Ridge techniques to facilitate the estimation of large VARs. Here too, including more lags is easy to encompass.

Bayesian techniques are incredibly popular in applied VAR research. Adding lags in Bayesian VARs is particularly straightforward since Bayesian approaches inherently avoid over-parametrization. One can think of priors generally as adding data points, which automatically counteracts potential degrees of freedom issues.[14] Therein lies the power of Bayesian methods in dealing with big data, where dealing with lots of parameters is the central issue (e.g. Giannone et al. (2021)). As with short-lag BVARs, the researcher has a wide menu of priors to choose from (see e.g. Hauzenberger et al. (2024)). Any of those priors can accommodate long lags. Whether the prior weights short and long lags differently is up to the researcher and should depend on the data and problem at hand. Our general result is not a call to favor certain priors over others. It is a call to incorporate long lags in the model such that posterior estimation can attach weight to them where relevant and thus overcome the problem of truncation. Our monetary application is an example where the prior heavily favors short over long lags (i.e. random walk prior), yet despite that estimation clearly finds long lags important in the data, as evidenced by the meaningful change in posterior structural inference.[15]

In sum, long-lag VARs should universally be considered in applied macro research. There

[14] Conjugate priors (e.g. Minnesota prior), for instance, are often implemented by adding dummy observations. Yet the point stands more generally, as priors serve the purpose of "regularization" in estimation.

[15] Antolin-Diaz and Surico (2025) is another example where long lags matter despite priors tilting towards short lags. Their sample is however much longer than the data span available for most macroeconomic questions.

is no need to adapt one's favorite modeling approach or inference method. All that is required is to allow for significantly longer lag lengths than current practice does.

5.3 On the horizon of interest and lag length: h vs. p

Both our simulations and empirical applications document the mechanical breakdown in inference when the IRF horizon h extends beyond the number of lags p. Long-lag VARs hedge against that breakdown and provide a fairer representation of the actual uncertainty at long horizons faced by the researcher. These results help in understanding how a short-lag researcher may find spurious significant long term effects. Baumeister (2025) provides an example of such a situation. Jordà et al. (2025) find significant long term (8 to 10 years) effects of monetary policy on the economy. Part of their evidence is based on a series of prominent short-lag VARs in the literature. Baumeister (2025) shows that these long term effects vanish when long-lag VARs are considered. This contrast can be understood as the consequence of the short-lag VAR researcher facing a significant bias in IRF and significantly underestimating uncertainty around those IRFs at longer horizons. Long-lag VARs address both problems.

But our results say more. Lag truncation is really a concern for every VAR researcher, as inference about horizons $h \leqslant p$ is also importantly affected. As evident from the literature and throughout our analysis, lag truncation induces bias at every horizon, including the very short ones. And again, long-lag VARs effectively hedge against that bias. That this matters for empirical work is clear from both our applications, where short term ($h \leqslant p$) structural conclusions change dramatically as the lag length increases.

In sum, every VAR researcher should be concerned with lag truncation, regardless of the horizon of interest that the structural question centers on. And long-lag VARs provide a practical tool to do so.

5.4 On local projections

The reader may wonder whether local projections are a better safeguard against lag truncation, in light of many of the advantages claimed in the literature comparing LPs and VARs. They are not. To see why, Ludwig's (2024) central theorem establishes that LP estimates combine VARs of increasing lag lengths. Specifically, an LP(p) at horizon h combines estimates of a VAR(p), VAR($p+1$), ..., VAR($p+h-1$). In practice LPs invariably use short lags. In fact, a common motivation for lag choice in LP is to use the VAR lag length dictated by information criteria (see e.g. the discussion in Baumeister (2025)). Yet short-lag VARs suffer from truncation bias. Hence, whenever lags are truncated, this bias necessarily carries over into LP estimates. Long-lag VARs, by contrast, guard against truncation bias. And as our results show, despite their larger model size, long-lag VARs may simultaneously reduce variance. In principle, long-lag LPs can also be unbiased because they build on long-lag VARs. But if the VAR(p) has long enough lags to address lag truncation, then while the LP also becomes unbiased it simultaneously becomes less efficient since it combines multiple VARs and therefore has many more coefficients (Ludwig (2024)). Hence, from both a bias and a variance perspective, long-lag VARs are preferable over LPs. Baumeister (2025) provides simulation evidence that further reinforces the point that long-lag VARs outperform LPs.

6 Conclusion

We document a general trade-off. Of course, it is possible to design models or find data for which the balance of the trade-off leans toward short lag lengths. However, contrary to common wisdom, long lag length need not imply prohibitively large imprecision. While increased parametrization in itself may increase uncertainty, this effect is counteracted by a reduction in misspecification. For VARs estimated on data generated by frequently used DSGE models, longer lag length tends to imply less bias and more precise inference. In empirical applications we find that the variance trade-off in VARs is not particular to data generated by DSGE models. For long-lag versions of prominent VARs in the literature, the balance of uncertainty effects seems to favor misspecification reduction over parametrization concerns. In particular, we find that structural results can be substantially different from their short-lag counterparts and that uncertainty does not necessarily become excessively large. Long-lag VARs are therefore an important instrument in the empirical macroeconomist's toolkit.

References

[1] Antolin-Diaz, J., Surico, P., 2025. "The long-run effects of government spending", *American Economic Review*, forthcoming.

[2] Baumeister, C., 2025. "Discussion of: Local projections or VARs? A primer for macroeconomists", In: Leahy, J.V., Ramey, V.A. (eds.), *NBER Macroeconomics Annual* 40, forthcoming.

[3] Beaudry, P., Portier, F., 2006. "Stock prices, news, and economic fluctuations", *American Economic Review* 96, 1293-1307.

[4] Bernanke, B.S., 1983. "Nonmonetary effects of the financial crisis in the propagation of the Great Depression", *American Economic Review* 73, 257-76.

[5] Blanchard, O.J., Quah, D., 1989. "The dynamic effects of aggregate demand and supply disturbances", *American Economic Review* 79, 655-73.

[6] Braun, P.A., Mittnik, S., 1993. "Misspecifications in vector autoregressions and their effects on impulse responses and variance decompositions", *Journal of Econometrics* 59, 319-341.

[7] Canova, F., 2007. Applied Macroeconomic Research, Princeton University Press, Princeton, New Jersey.

[8] Chang, Y., Hong, J.H., 2006. "Do technological improvements in the manufacturing sector raise or lower employment?", *American Economic Review* 96, 352-368.

[9] Chari, V.V., Kehoe, P.J., McGrattan, E.R., 2008. “Are structural VARs with long-run restrictions useful in developing business cycle theory?”, *Journal of Monetary Economics* 55, 1337-52.

[10] Chatfield, C., 1995. “Model uncertainty, data mining and statistical inference”, *Journal of the Royal Statistical Society, Series A* 158 (3), 419-44.

[11] Christiano, L.J., Eichenbaum, M., Vigfusson, R., 2004. “The response of hours to a technology shock: Evidence based on direct measures of technology”, *Journal of the European Economic Association* 2, 381-95.

[12] Christiano, L.J., Eichenbaum, M., Vigfusson, R., 2007. “Assessing structural VARs”, In: Acemoglu, D., Rogoff, K.S., Woodford, M. (eds.), NBER Macroeconomics Annual 2006.

[13] Cooley, T.F., Dwyer, M., 1998. “Business cycle analysis without much theory. A look at structural VARs”, *Journal of Econometrics* 83, 57-88.

[14] Eichenbaum, M., Evans, C., 1995. “Some empirical evidence on the effects of shocks to monetary policy on exchange rates”, *Quarterly Journal of Economics* 110, 975-1009.

[15] Faust, J., Leeper, E.M., 1997. “When do long-run identifying restrictions give reliable results?”, *Journal of Business and Economic Statistics* 15, 345-53.

[16] Fernandez-Villaverde, J., Rubio-Ramirez, J.F., Sargent, T.J., Watson, M.W., 2007. “ABCs (and Ds) of understanding VARs”, *American Economic Review* 97 (3), 1021–26.

[17] Fisher, J.D.M., 2006. “The dynamic effects of neutral and investment-specific technology shocks”, *Journal of Political Economy* 114, 413-51.

[18] Gali, J., 1999. “Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations?”, *American Economic Review* 89, 249-71.

[19] Giannone, D, Lenza, M., Primiceri, G.E., 2021. “Economic predictions with big data: The illusion of sparsity”, *Econometrica* 89 (5), 2409-37.

[20] González-Casasús, O., Schorfheide, F., 2025. “Misspecification-robust shrinkage and selection for VAR forecasts and IRFs”, NBER Working Paper 33474.

[21] Hauzenberger, N., Huber, F., Koop, G., 2024. “Macroeconomic forecasting using BVARs”, In: Clements, M., Galvão, A., *Handbook of Research Methods and Applications in Macroeconomic Forecasting.*

[22] Jordà, Ò., Singh, S.R., Taylor, A.M., 2025. “The long-run effects of monetary policy”, *Review of Economics and Statistics*, forthcoming.

[23] Li, D., Plagborg-Møller, M., Wolf, C.K., 2024a. “Local projections vs. VARs: Lessons from thousands of DGPs”, *Journal of Econometrics* 244, 105722, 1-21.

[24] Li, D., Plagborg-Møller, M., Wolf, C.K., 2024b. “Local projections vs. VARs: Lessons from thousands of DGPs”, slides, available at www.mikkelpm.com/research/.

[25] Ludwig, J., 2024. “Local projections are VAR predictions of increasing order”, mimeo.

[26] Miranda-Agrippino, S., Ricco, G., 2021. “The Transmission of Monetary Policy Shocks”, *American Economic Journal: Macroeconomics* 13 (3), 74–107.

[27] Montiel Olea, J.L., Qian, E., Plagborg-Møller, M., Wolf, C.K., 2025. "Local projections or VARs? A primer for macroeconomists", In: Leahy, J.V., Ramey, V.A. (eds.), *NBER Macroeconomics Annual* 40, forthcoming.

[28] Plagborg-Møller, M., Wolf, C.K., 2021. "Local Projections and VARs Estimate the Same Impulse Responses", *Econometrica* 89(2), 955-80.

[29] Ravenna, F., 2007. "Vector autoregressions and reduced form representations of DSGE models", *Journal of Monetary Economics* 54, 2048-64.

[30] Sims, C., 1972. "The role of approximate prior restrictions in distributed lag estimation", *Journal of the American Statistical Association* 67, 169-75.

[31] Sims, C., 1980. "Macroeconomics and reality", *Econometrica* 48, 1-48.

[32] Sims, C., 1989. "Models and their uses", *American Journal of Agricultural Economics* 71, 489-94.

[33] Sims, C., Zha, T., 1999. "Error bands for impulse responses", *Econometrica* 67, 1113-55.

[34] Smets, F., Wouters, R., 2007. "Shocks and frictions in US business cycles: a Bayesian DSGE approach", *American Economic Review* 97, 586-606.

[35] Stock, J.H., Watson, M.W., 2001. "Vector autoregressions", *Journal of Economic Perspectives* 15, 101-15.

[36] Stock, J.H., Watson, M.W., 2010. "Dynamic Factor Models", Oxford Handbook of Economic Forecasting, M.P. Clements and D.F. Hendry (eds), Oxford University Press.

[37] Uhlig, H., 2005. "What are the effects of monetary policy on output? Results from an agnostic identification procedure", *Journal of Monetary Economics* 52, 381-419.

Recent Working Papers:

For a complete list of Working Papers published by Sveriges Riksbank, see www.riksbank.se

The Macroeconomic Effects of Trade Tariffs: Revisiting the Lerner Symmetry Result *by Jesper Lindé and Andrea Pescatori*	2019:363
Biased Forecasts to Affect Voting Decisions? The Brexit Case *by Davide Cipullo and André Reslow*	2019:364
The Interaction Between Fiscal and Monetary Policies: Evidence from Sweden *by Sebastian Ankargren and Hovick Shahnazarian*	2019:365
Designing a Simple Loss Function for Central Banks: Does a Dual Mandate Make Sense? *by Davide Debortoli, Jinill Kim and Jesper Lindé*	2019:366
Gains from Wage Flexibility and the Zero Lower Bound *by Roberto M. Billi and Jordi Galí*	2019:367
Fixed Wage Contracts and Monetary Non-Neutrality *by Maria Björklund, Mikael Carlsson and Oskar Nordström Skans*	2019:368
The Consequences of Uncertainty: Climate Sensitivity and Economic Sensitivity to the Climate *by John Hassler, Per Krusell and Conny Olovsson*	2019:369
Does Inflation Targeting Reduce the Dispersion of Price Setters' Inflation Expectations? *by Charlotte Paulie*	2019:370
Subsampling Sequential Monte Carlo for Static Bayesian Models *by David Gunawan, Khue-Dung Dang, Matias Quiroz, Robert Kohn and Minh-Ngoc Tran*	2019:371
Hamiltonian Monte Carlo with Energy Conserving Subsampling *by Khue-Dung Dang, Matias Quiroz, Robert Kohn, Minh-Ngoc Tran and Mattias Villani*	2019:372
Institutional Investors and Corporate Investment *by Cristina Cella*	2019:373
The Impact of Local Taxes and Public Services on Property Values *by Anna Grodecka and Isaiah Hull*	2019:374
Directed technical change as a response to natural-resource scarcity *by John Hassler, Per Krusell and Conny Olovsson*	2019:375
A Tale of Two Countries: Cash Demand in Canada and Sweden *by Walter Engert, Ben Fung and Björn Segendorf*	2019:376
Tax and spending shocks in the open economy: are the deficits twins? *by Mathias Klein and Ludger Linnemann*	2019:377
Mind the gap! Stylized dynamic facts and structural models *by Fabio Canova and Filippo Ferroni*	2019:378
Financial Buffers, Unemployment Duration and Replacement Labor Income *by Mats Levander*	2019:379
Inefficient Use of Competitors' Forecasts? *by André Reslow*	2019:380
How Much Information Do Monetary Policy Committees Disclose? Evidence from the FOMC's Minutes and Transcripts *by Mikael Apel, Marianna Blix Grimaldi and Isaiah Hull*	2019:381
Risk endogeneity at the lender/investor-of-last-resort *by Diego Caballero, André Lucas, Bernd Schwaab and Xin Zhang*	2019:382
Heterogeneity in Households' Expectations of Housing Prices – Evidence from Micro Data *by Erik Hjalmarsson and Pär Österholm*	2019:383
Big Broad Banks: How Does Cross-Selling A Affect Lending? *by Yingjie Qi*	2020:384
Unemployment Fluctuations and Nominal GDP Targeting *by Roberto Billi*	2020:385
FAQ: How do I extract the output gap? *by Fabio Canova*	2020:386

Drivers of consumer prices and exchange rates in small open economies *by Vesna Corbo and Paola Di Casola*	2020:387
TFP news, stock market booms and the business cycle: Revisiting the evidence with VEC models *by Paola Di Casola and Spyridon Sichlimiris*	2020:388
The costs of macroprudential deleveraging in a liquidity trap *by Jiaqian Chen, Daria Finocchiaro, Jesper Lindé and Karl Walentin*	2020:389
The Role of Money in Monetary Policy at the Lower Bound *by Roberto M. Billi, Ulf Söderström and Carl E. Walsh*	2020:390
MAJA: A two-region DSGE model for Sweden and its main trading partners *by Vesna Corbo and Ingvar Strid*	2020:391
The interaction between macroprudential and monetary policies: The cases of Norway and Sweden *by Jin Cao, Valeriya Dinger, Anna Grodecka-Messi, Ragnar Juelsrud and Xin Zhang*	2020:392
Withering Cash: Is Sweden ahead of the curve or just special? *by Hanna Armelius, Carl Andreas Claussen and André Reslow*	2020:393
Labor shortages and wage growth *by Erik Frohm*	2020:394
Macro Uncertainty and Unemployment Risk *by Joonseok Oh and Anna Rogantini Picco*	2020:395
Monetary Policy Surprises, Central Bank Information Shocks, and Economic Activity in a Small Open Economy *by Stefan Laséen*	2020:396
Econometric issues with Laubach and Williams' estimates of the natural rate of interest *by Daniel Buncic*	2020:397
Quantum Technology for Economists *by Isaiah Hull, Or Sattath, Eleni Diamanti and Göran Wendin*	2020:398
Modeling extreme events: time-varying extreme tail shape *by Bernd Schwaab, Xin Zhang and André Lucas*	2020:399
The Effects of Government Spending in the Eurozone *by Ricardo Duque Gabriel, Mathias Klein and Ana Sofia Pessoa*	2020:400
Narrative Fragmentation and the Business Cycle *by Christoph Bertsch, Isaiah Hull and Xin Zhang*	2021:401
The Liquidity of the Government Bond Market – What Impact Does Quantitative Easing Have? Evidence from Sweden *by Marianna Blix Grimaldi, Alberto Crosta and Dong Zhang*	2021:402
Five Facts about the Distributional Income Effects of Monetary Policy *by Niklas Amberg, Thomas Jansson, Mathias Klein and Anna Rogantini Picco*	2021:403
When domestic and foreign QE overlap: evidence from Sweden *by Paola Di Casola and Pär Stockhammar*	2021:404
Dynamic Macroeconomic Implications of Immigration *by Conny Olovsson, Karl Walentin, and Andreas Westermark*	2021:405
Revisiting the Properties of Money *by Isaiah Hull and Or Sattath*	2021:406
The cost of disinflation in a small open economy vis-à-vis a closed economy *by Oleksandr Faryna, Magnus Jonsson and Nadiia Shapovalenko*	2021:407
On the Performance of Cryptocurrency Funds *by Daniele Bianchi and Mykola Babiak*	2021:408
The low-carbon transition, climate commitments and firm credit risk *by Sante Carbone, Margherita Giuzio, Sujit Kapadia, Johannes Sebastian Krämer, Ken Nyholm and Katia Vozian*	2022:409
Seemingly Irresponsible but Welfare Improving Fiscal Policy at the Lower Bound *by Roberto M. Billi and Carl E. Walsh*	2022:410
Pension Reform and Wealth Inequality: Evidence from Denmark *by Torben M. Andersen, Joydeep Bhattacharya, Anna Grodecka-Messi and Katja Mann*	2022:411

Inflation Targeting or Fiscal Activism? *by Roberto M. Billi*	2022:412
Trading volume and liquidity provision in cryptocurrency markets *by Daniele Bianchi, Mykola Babiak and Alexander Dickerson*	2022:413
DISPERSION OVER THE BUSINESS CYCLE: PASSTHROUGH, PRODUCTIVITY, AND DEMAND *by Mikael Carlsson, Alex Clymo and Knut-Eric Joslin*	2022:414
Electoral Cycles in Macroeconomic Forecasts *by Davide Cipullo and André Reslow*	2022:415
The Curious Incidence of Monetary Policy Across the Income Distribution *by Tobias Broer, John Kramer and Kurt Mitman*	2022:416
Central Bank Mandates and Monetary Policy Stances: through the Lens of Federal Reserve Speeches *by Christoph Bertsch, Isaiah Hull, Robin L. Lumsdaine, and Xin Zhang*	2022:417
The Political Costs of Austerity *by Ricardo Duque Gabriel, Mathias Klein and Ana Sofia Pessoa*	2022:418
Central bank asset purchases: Insights from quantitative easing auctions of government bonds *by Stefan Laséen*	2023:419
Greenflation? *by Conny Olovsson and David Vestin*	2023:420
Effects of foreign and domestic central bank government bond purchases in a small open economy DSGE model: Evidence from Sweden before and during the coronavirus pandemic *by Yildiz Akkaya, Carl-Johan Belfrage, Paola Di Casola and Ingvar Strid*	2023:421
Dynamic Credit Constraints: Theory and Evidence from Credit Lines* *by Niklas Amberg, Tor Jacobson, Vincenzo Quadrini and Anna Rogantini Picco*	2023:422
Stablecoins: Adoption and Fragility *by Christoph Bertsch*	2023:423
CBDC: Lesson from a Historical Experience *by Anna Grodecka-Messi and Xin Zhang*	2023:424
Do Credit Lines Provide Reliable Liquidity Insurance? Evidence from Commercial-Paper Backup Lines *by Niklas Amberg*	2023:425
Price Pass-Through Along the Supply Chain: Evidence from PPI and CPI Microdata *by Edvin Ahlander, Mikael Carlsson and Mathias Klein*	2023:426
Cash for Transactions or Store-of-Value? A comparative study on Sweden and peer countries *by Carl Andreas Claussen, Björn Segendorf and Franz Seitz*	2023:427
Fed QE and bank lending behaviour: a heterogeneity analysis of asset purchases *by Marianna Blix Grimaldi and Supriya Kapoor*	2023:428
Monetary policy in Sweden after the end of Bretton Woods *by Emma Bylund, Jens Iversen and Anders Vredin*	2023:429
Banking Without Branches *by Niklas Amberg and Bo Becker*	2024:430
Climate impact assessment of retail payment services *by Niklas Arvidsson, Fumi Harahap, Frauke Urban and Anissa Nurdiawati*	2024:431
Four Facts about International Central Bank Communication by *Christoph Bertsch, Isaiah Hull, Robin L. Lumsdaine, and Xin Zhang*	2024:432
Optimal Monetary Policy with r* < 0 *by Roberto Billi, Jordi Galí, and Anton Nakov*	2024:433
Quantitative Easing, Bond Risk Premia and the Exchange Rate in a Small Open Economy *by Jens H. E. Christensen and Xin Zhang*	2024:434
Supply-Chain Finance: An Empirical Evaluation of Supplier Outcomes *by Niklas Amberg, Tor Jacobson and Yingjie Qi*	2024:435
Optimal Contracts and Inflation Targets Revisited *by Torsten Persson and Guido Tabellini*	2024:436
Potential Climate Impact of Retail CBDC Models *by Niklas Arvidsson, Fumi Harahap, Frauke Urban and Anissa Nurdiawati*	2024:437

Do we need firm data to understand macroeconomic dynamics? *by Michele Lenza and Ettore Savoia*	2024:438
Inflation-Dependent Exchange Rate Pass-Through in Sweden: Insights from a Logistic Smooth Transition VAR Model *by Gabriella Linderoth and Malte Meuller*	2024:439
Quantitative Easing and the Supply of Safe Assets: Evidence from International Bond Safety Premia *by Jens H. E. Christensen, Nikola N. Mirkov and Xin Zhang*	2024:440
Bank fragility and the incentives to manage risk *by Toni Ahnert, Christoph Bertsch, Agnese Leonello and Robert Marquez*	2024:441
A Traffic-Jam Theory of Growth *by Daria Finocchiaro and Philippe Weil*	2024:442
Intertemporal MPC and Shock Size *by Tullio Jappelli, Ettore Savoia and Alessandro Sciacchetano*	2024:443
Plundered or profitably pumped-up? The effects of private equity takeover *by Anders Kärnä and Samantha Myers*	2024:444
Measuring Riksbank Monetary Policy: Shocks and Macroeconomic Transmission *by Jakob Almerud, Dominika Krygier, Henrik Lundvall and Mambuna Njie*	2024:445
Joint extreme Value-at-Risk and Expected Shortfall dynamics with a single integrated tail shape parameter *by Enzo D'Innocenzo, André Lucas, Bernd Schwaab and Xin Zhang*	2025:446
The Inflationary Effects of Quantitative Easing *by Mathias Klein and Xin Zhang*	2025:447
Shadow banks or just not banks? Growth of the Swedish non-bank sector *by Jieying Li and Samantha Myers*	2025:448
Sveriges Riksbank´s Foreign Exchange Reserve, 1823–2023 *by Gustav Ingman*	2025:449
Unconventional Monetary Policies in Small Open Economies *by Marcin Kolasa, Stefan Laséen and Jesper Lindé*	2025:450

Sveriges Riksbank
Visiting address: Brunkebergs torg 11
Mail address: se-103 37 Stockholm

Website: www.riksbank.se
Telephone: +46 8 787 00 00, Fax: +46 8 21 05 31
E-mail: registratorn@riksbank.se