SVERIGES RIKSBANK
WORKING PAPER SERIES

271

# Un-truncating VARs*

*Ferre De Graeve and Andreas Westermark*

June 2013

WORKING PAPERS ARE OBTAINABLE FROM

Sveriges Riksbank • Information Riksbank • SE-103 37 Stockholm
Fax international: +46 8 787 05 26
Telephone international: +46 8 787 01 00
E-mail: info@riksbank.se

# Un-truncating VARs*

Ferre De Graeve and Andreas Westermark

Sveriges Riksbank Working Paper Series No. 271

June 2013

**Abstract**

Macroeconomic research often relies on structural vector autoregressions to uncover empirical regularities. Critics argue the method goes awry due to lag truncation: short lag-lengths imply a poor approximation to 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. For data generated by frontier DSGE-models long-lag VARs are feasible, reduce bias and variance, and have better coverage. Thus, contrary to conventional wisdom, the trivial solution to the critique actually works.

*Keywords*: VAR, SVAR, Lag-length, Truncation

*JEL*: C18, E37

*This paper has benefited from discussions with Paul Beaudry, Fabio Canova, Yongsung Chang, Martin Eichenbaum, Jesús Fernández-Villaverde, Jonas Fisher, Jordi Galí, Marco Lippi, Jean-Paul L'Huillier, Lutz Kilian, Per Krusell, Ellen McGrattan, Juan Rubio-Ramirez, Federico Ravenna, Barbara Rossi, Chris Sims, Dan Waggoner, Tao Zha and seminar participants at the Atlanta Fed, Cleveland Fed, Dallas Fed, EIEF, Universitat Pompeu Fabra 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.de.graeve@riksbank.se, andreas.westermark@riksbank.se.

# 1 Introduction

Structural Vector Autoregressions (SVARs) 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), Galí (1999), Fisher (2006), Beaudry and Portier (2006) and others provided SVAR-based evidence for a variety of shocks, with each essentially spurring a separate field of research.

Yet the SVAR method is not without its critics. Many critiques of SVARs boil down to the problem of lag truncation. In particular, while DSGE models tend to imply reduced form VAR representations with long lag-length (often infinity), when going to the data, macroeconomists invariably settle on using a very small number of lags (typically one to four quarters). Because lags are truncated, the critics show, impulse response functions (IRFs) computed using the SVAR may not correspond to those of the underlying DSGE model. Chari, Kehoe and McGrattan (2008, henceforth CKM) is the most recent and well-known elicitation of that critique.[1]

The trivial solution to lag truncation, i.e., dramatically increasing lag-length, is unexplored. What keeps macroeconomists from using long lag-lengths is the intuition that uncertainty becomes pervasive. That is, increasing lag-length increases the number of parameters rapidly, thereby reducing the degrees of freedom and making confidence bandwidth explode.

We show that this standard intuition is only part of the story. In the face of misspecification due to lag truncation, increasing lag-length can actually reduce uncertainty. The reason is that as truncation reduces, misspecification reduces. The reduction in misspecifcation not only leads to the well-known bias reduction, but it also reduces variance. This reduction in variance will work 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 sign restrictions.

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

We show that in increasing lag-length in standard SVARs on small 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 almost unequivocally in favor of long-lag VARs: reduced truncation bias, more precise inference, reduced MSE, better coverage rates.

The implication is, contrary to conventional wisdom, that it is possible to estimate SVARs with long lags, and hence reduce truncation bias, and still derive precise structural predictions from them.

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 SVAR impulse responses, where analytics are not tractable. We then assess long-lag VARs on the basis of a series of Monte Carlo experiments. We draw data from a variety of DSGE models, estimate SVARs of different (and possibly very long) lag-length and evaluate their performance. Finally, we assess the implications of our results and discuss some possible avenues for future research.

# 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 $\epsilon$. 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

$$\begin{aligned} Y_t &= B_1Y_{t-1} + ... + B_pY_{t-p} + u_t,\ E(u_tu_t') = \Sigma \\ B(L) &= B_1L + ... + B_pL^p, \end{aligned}$$

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 $\Sigma$. Any SVAR 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 $\Sigma$ they will tend to become less biased if both its arguments become less biased.[2] In other words, reducing truncation reduces bias.

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 $V(B(L)) \uparrow \Longrightarrow V(IRF) \uparrow$ as lag-length increases. But (3) also makes clear that this argument is incomplete. In particular, it neglects that there is a second argument, $\Sigma$. 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 $\Sigma$, 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 $\Sigma$. 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 sample data of length equal to that available in typical macro data samples ($T = 200$).[3] Given one such draw of data, we estimate VARs of differ-

[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).

[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.

ent lag-lengths, calculate impulse response functions and construct confidence bands using standard methods.[4] We repeat that exercise 1000 times for each model and subsequently investigate bias, uncertainty bandwidth, mean-squared error and coverage rates.

## 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, Eichenbaum and Vigfusson (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 SVARs. 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 (or, in the case of SW, 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 DGP).

[4] See Christiano, Eichenbaum and Vigfusson (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.

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

$$r_t = \rho r_{t-1} + (1-\rho)\left\{r_\pi \pi_{t-1} + r_y\left(y_{t-1} - y^p_{t-1}\right)\right\} + \varepsilon^r_t$$

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

Second, invertibility is never a problem; all the models we consider are fundamental. Third, all our experiments 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, and a VAR on GDP-growth and the short term interest rate (in that order).[6] In Section 5.1, we discuss the extent to which restricting attention to two-variable systems matters. Finally, inference is standard. Uncertainty bands are computed as in e.g. Canova (2007) and Uhlig (2005). In particular, given a weak conjugate prior, VARs have a posterior distribution of the Normal-Inverse Wishart form, where the distributions are centered around their OLS estimates.[7]

## 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 significant bias reduction, most notably in CKM and, from intermediate horizons onward, in CEV and SW. To evaluate if such biases are of concern, we now turn to measures of uncertainty.

**Result 1: Uncertainty does not explode for long-lag VARs** Figure 2 plots the median width of the confidence bands across all draws.[8] A first glance at that figure reveals that, contrary to common wisdom, bandwidth does not explode. Instead, even for VARs

[6] Results for different shocks and variables are qualitatively similar.

[7] Our results go through for bootstrap-based confidence bands, used e.g. in CEV. However, bootstrap-based procedures tend to run into non-stationarity problems more frequently. For instance, Kilian's (1998) double bootstrap often implies an unstable bias correction for large lag-lengths.

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

with very long lags uncertainty bands are roughly in the same ballpark as those of short-lag VARs.

**Result 2: Short-lag VARs have maximal uncertainty for horizons where uncertainty is not mechanically low** For short horizons short-lag VARs have maximal bandwidth. 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.

For longer horizons, short-lag VARs trivially attain minimum bandwidth. 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 bandwidth mechanically converges to zero soon after horizon $p$.

**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 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 bandwidth in a different way, by plotting mean-squared errors (MSE) across horizons. The message is very much the same: at short horizons -

[9] The huge swings in coverage for short-lag VARs arise naturally 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.

where uncertainty does not mechanically shrink - short-lag VARs are either comparable or considerably worse than long-lag VARs.

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

Figure 5 plots the Monte Carlo distribution of uncertainty bandwidth for three types of impulse responses. Specifically, for each draw of data from the DSGE model, we measure the bandwidth around the contemporaneous impulse response. For short-lag VARs, the dashed line $(B_4, \Sigma_4)$ plots the distribution of bandwidths across all 1000 draws. Similarly, the solid line plots the distribution of bandwidths for a long-lag VAR $(B_{30}, \Sigma_{30})$. The medians of these two distributions are already contained in Figure 2: the contemporaneous response for CKM and SW, the second horizon for CEV.[10] 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. Such IRFs can be interpreted as isolating the effect of increased parametrization.

[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 the contemporaneous 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 second horizon.

They shut down the effect of misspecification reduction by ignoring the reduced bias in $\Sigma$. The dotted $(B_{30}, \Sigma_4)$ distributions in Figure 5 show the bandwidth 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, there is no effect at all from increased parametrization, since the short-run restriction implies that the contemporaneous IRF only depends on $\Sigma$ and not on $B(L)$. For the CEV and CKM models the right tail of the bandwidth distribution becomes fatter, as standard intuition would suggest. However, two observations stand out. First, the increase in bandwidth 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 $\Sigma$. 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 bandwidth 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 increase in number of parameters, once the misspecification effect through $\Sigma$ is incorporated long-lag VARs appear associated with smaller, not larger uncertainty bandwidth.

Thus, the figures show the uncertainty trade-off: increased parametrization ($B_4 \longrightarrow B_{30}$) which can - but need not - push the distribution outward (from dashed to dotted) vs. reduced misspecification ($\Sigma_4 \longrightarrow \Sigma_{30}$) 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 lag-length can easily imply a reduction in variance.

# 5 Concluding remarks

## 5.1 On the maintained simplifications

Throughout the analysis, the only modification as compared to the standard approach in e.g., Galí (1999), is an increase in lag-length. No additional degree of complexity is introduced, and only standard tools are used. Let us briefly dwell on one of the simplifications, notably that all simulations are based on two variable VARs. Considering small VARs serves to keep the number of parameters limited. As lags increase, the increase in parameters increases

[11] The reduction in uncertainty in the dotted distribution can also be the result of reduced misspecification 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, Eichenbaum and Vigfusson (2004).

faster the more variables in the system.

Since much of the influential SVAR evidence in the literature is based on small VARs, with two, maximum three variables, it seems reasonable to focus on small VARs. Moreover, many developments in empirical macro enable dealing with larger systems. For instance, variants on Minnesota-type priors can allow inclusion of long lags in VARs with many variables. Alternatively, factor dynamics with potentially long lags may well improve structural inference without a large increase in parameters relative to the size of the data. Smoothness priors are yet another available alternative. In short, there are potentially many ways of dealing with larger systems. Irrespective of the particular approach, the variance trade-off we document will be at work in larger systems, too.

## 5.2 On choosing 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 focus on parsimony.[12]

A potential drawback of including longer lags is that it induces overfitting. We have extensively investigated this possibility. For the models and the lag-lengths considered here, we find it not be a major problem. One way to see this is as follows. If present, overfitting should have a first-order effect on bias. In other words, one would expect bias to increase

[12] See Kilian (2001) for a similar argument. He argues that lag criteria which punish parametrization less heavily often have better properties in terms of IRF.

when overfitting sets in. This does not generally occur in our simulations. That said, one avenue of future research lies in the development of information criteria that take into account that the purpose of the model is structural inference, while also avoiding issues of overfitting.

## 5.3 Generality

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 SVARs run on data generated by frequently used DSGE models, longer lag-length tends to imply less bias and more precise inference. Long-lag VARs are therefore a viable instrument in the empirical macroeconomist's toolkit.

In ongoing empirical work we find that the variance trade-off in VARs is not particular to data generated by DSGE models. For long-lag versions of many prominent SVARs in the literature, the balance of uncertainty effects seems to favor misspecification reduction over parametrization concerns. In particular, we find that results can be substantially different from their short-run counterparts and that uncertainty does not explode.

# References


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

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

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

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

[5] 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.

[6] 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.

[7] Christiano, L.J., Eichenbaum, M., Vigfusson, R., 2007. "Assessing structural VARs," In: Acemoglu, D., Rogoff, K.S., Woodford, M. (Eds.), NBER Macroeconomics Annual 2006. MIT Press, Cambridge, 1-106.

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

[9] 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.

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

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

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

[13] Kilian, L., 1998. "Small-sample confidence intervals for impulse response functions," *Review of Economics and Statistics* 80, 218-230.

[14] Kilian, L., 2001. "Impulse response analysis in vector autoregressions with unknown lag order," *Journal of Forecasting* 20, 161-179.

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

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

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

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

[19] Smets, F., Wouters, R., 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review* 97, 586-606.

[20] 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.

Figure 1: Bias

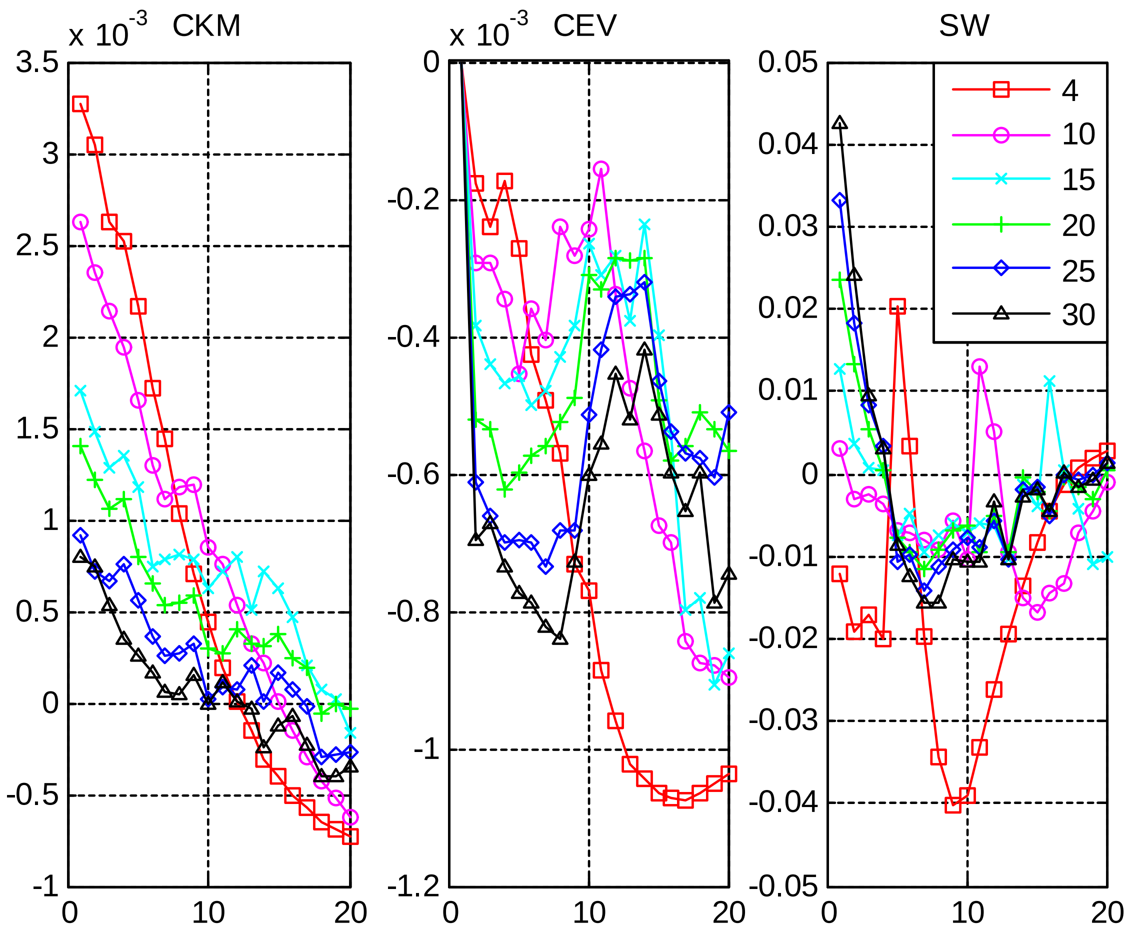

Figure 2: Bandwidth (95th-5th percentile)

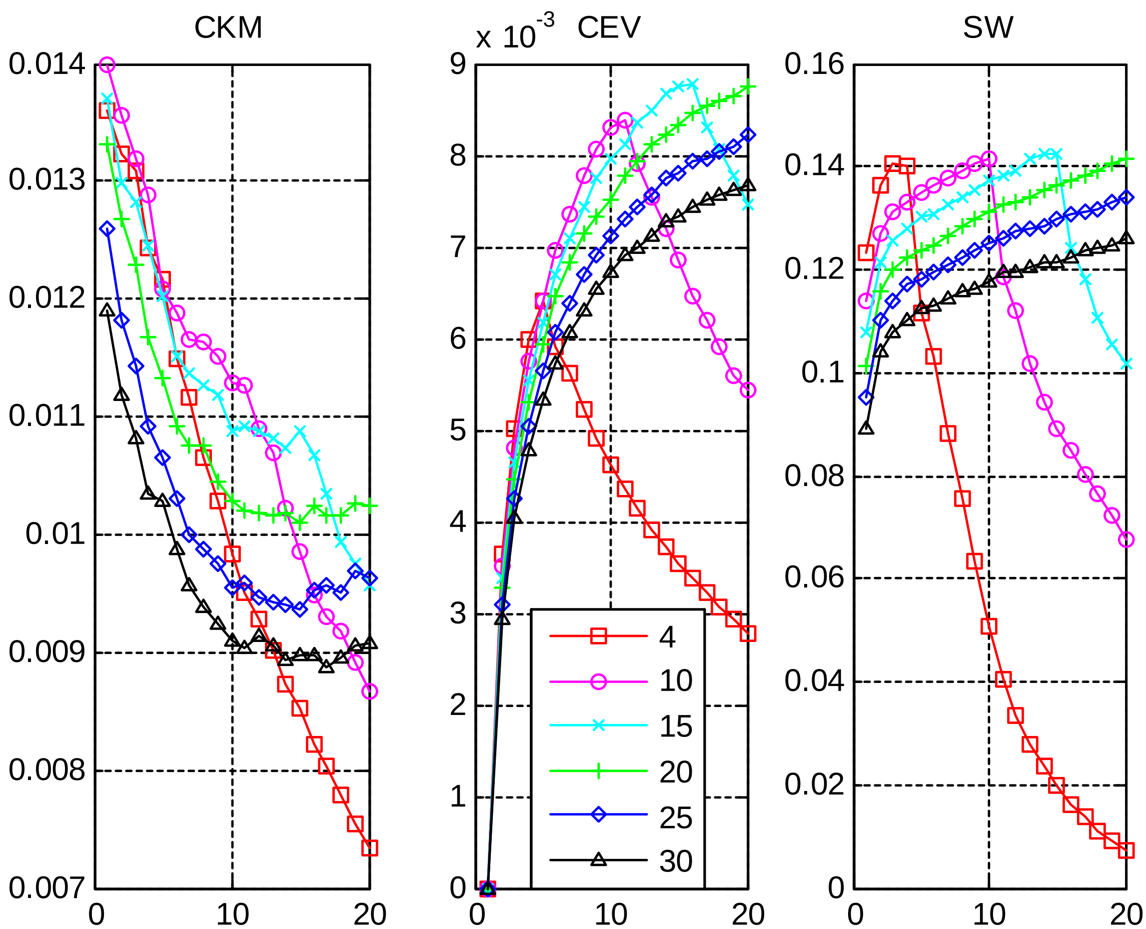

Figure 3: Coverage (90 percent)

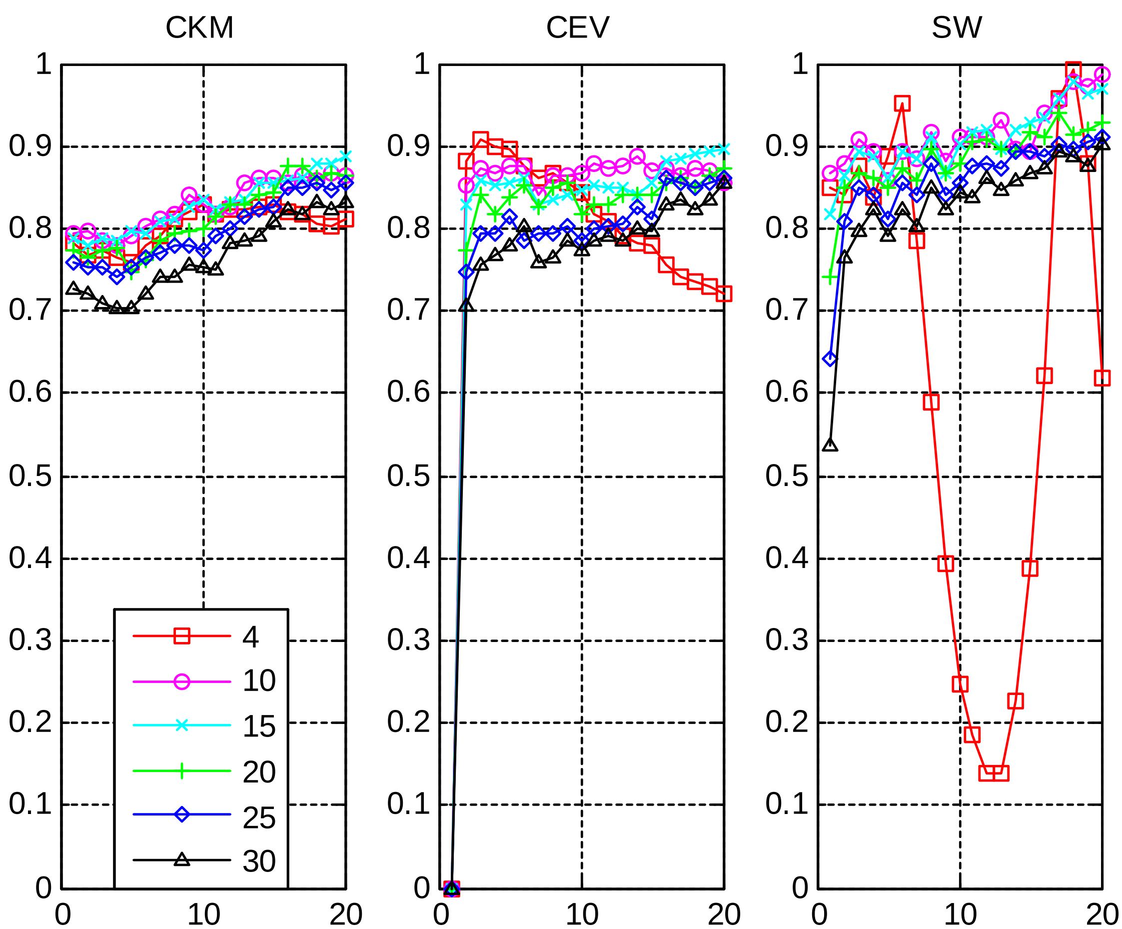

Figure 4: Mean-squared error

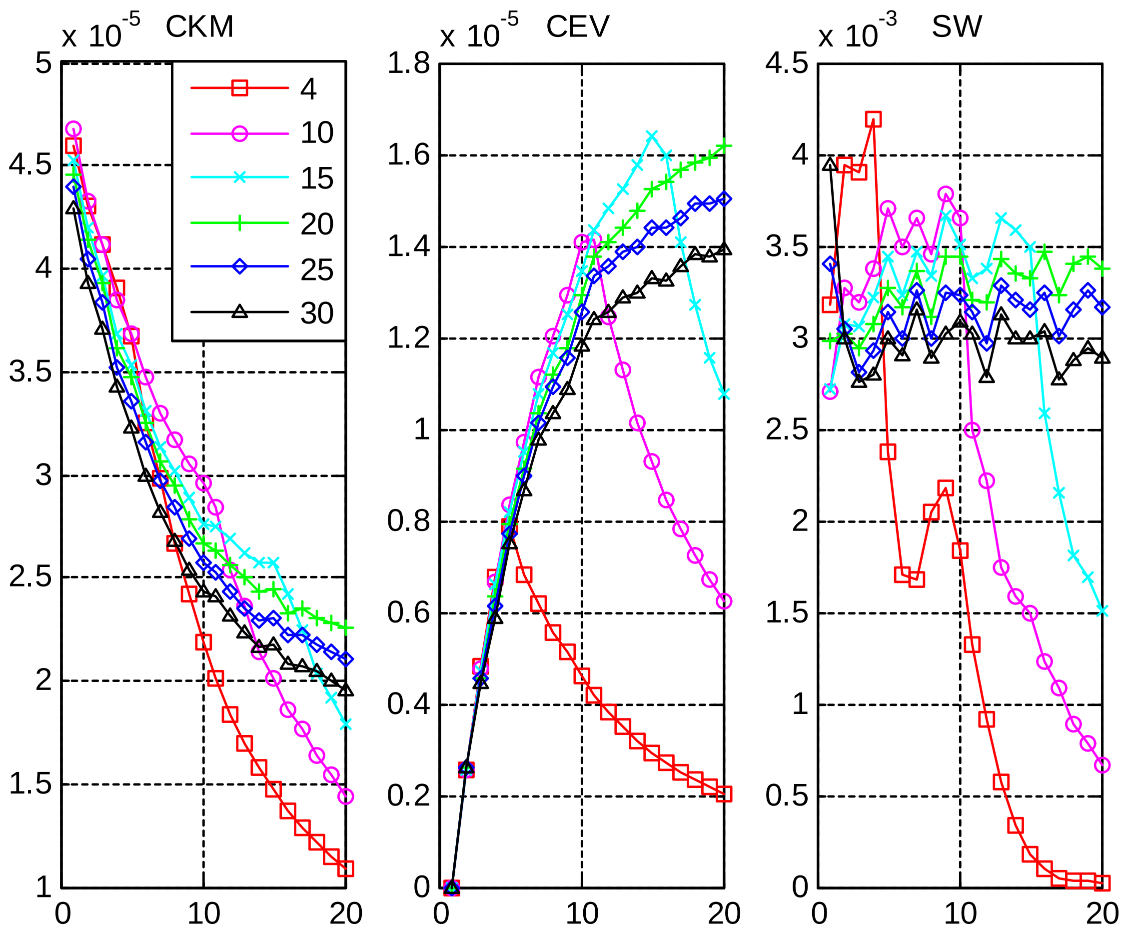

Figure 5: Bandwidth distribution

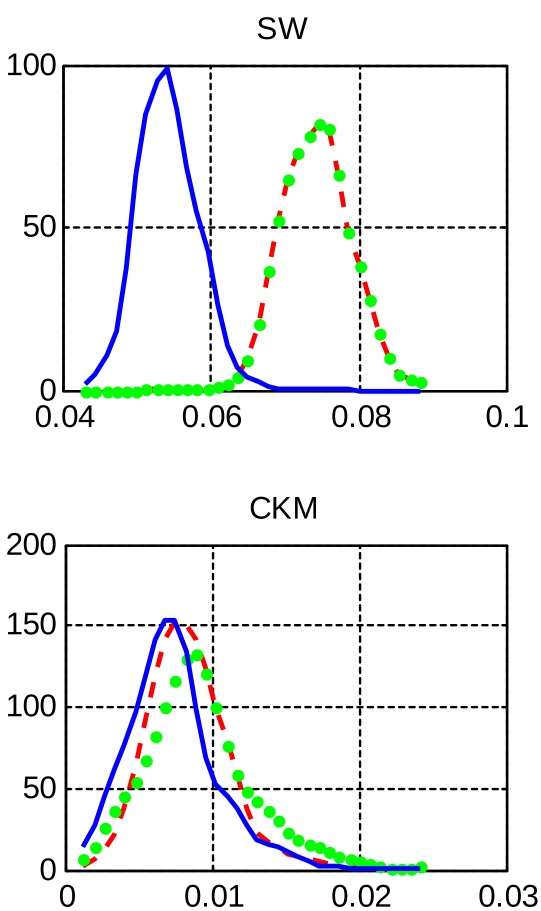

# Earlier Working Papers:

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

| | |
|---|---|
| Estimation of an Adaptive Stock Market Model with Heterogeneous Agents<br>*by Henrik Amilon* | 2005:177 |
| Some Further Evidence on Interest-Rate Smoothing: The Role of Measurement Errors in the Output Gap<br>*by Mikael Apel and Per Jansson* | 2005:178 |
| Bayesian Estimation of an Open Economy DSGE Model with Incomplete Pass-Through<br>by Malin Adolfson, Stefan Laséen, Jesper Lindé and Mattias Villani | 2005:179 |
| Are Constant Interest Rate Forecasts Modest Interventions? Evidence from an Estimated Open Economy DSGE Model of the Euro Area<br>*by Malin Adolfson, Stefan Laséen, Jesper Lindé and Mattias Villani* | 2005:180 |
| Inference in Vector Autoregressive Models with an Informative Prior on the Steady State<br>*by Mattias Villani* | 2005:181 |
| Bank Mergers, Competition and Liquidity<br>*by Elena Carletti, Philipp Hartmann and Giancarlo Spagnolo* | 2005:182 |
| Testing Near-Rationality using Detailed Survey Data<br>*by Michael F. Bryan and Stefan Palmqvist* | 2005:183 |
| Exploring Interactions between Real Activity and the Financial Stance<br>*by Tor Jacobson, Jesper Lindé and Kasper Roszbach* | 2005:184 |
| Two-Sided Network Effects, Bank Interchange Fees, and the Allocation of Fixed Costs<br>*by Mats A. Bergman* | 2005:185 |
| Trade Deficits in the Baltic States: How Long Will the Party Last?<br>*by Rudolfs Bems and Kristian Jönsson* | 2005:186 |
| Real Exchange Rate and Consumption Fluctuations follwing Trade Liberalization<br>*by Kristian Jönsson* | 2005:187 |
| Modern Forecasting Models in Action: Improving Macroeconomic Analyses at Central Banks<br>*by Malin Adolfson, Michael K. Andersson, Jesper Lindé, Mattias Villani and Anders Vredin* | 2005:188 |
| Bayesian Inference of General Linear Restrictions on the Cointegration Space<br>*by Mattias Villani* | 2005:189 |
| Forecasting Performance of an Open Economy Dynamic Stochastic General Equilibrium Model<br>*by Malin Adolfson, Stefan Laséen, Jesper Lindé and Mattias Villani* | 2005:190 |
| Forecast Combination and Model Averaging using Predictive Measures<br>*by Jana Eklund and Sune Karlsson* | 2005:191 |
| Swedish Intervention and the Krona Float, 1993-2002<br>*by Owen F. Humpage and Javiera Ragnartz* | 2006:192 |
| A Simultaneous Model of the Swedish Krona, the US Dollar and the Euro<br>*by Hans Lindblad and Peter Sellin* | 2006:193 |
| Testing Theories of Job Creation: Does Supply Create Its Own Demand?<br>*by Mikael Carlsson, Stefan Eriksson and Nils Gottfries* | 2006:194 |
| Down or Out: Assessing The Welfare Costs of Household Investment Mistakes<br>*by Laurent E. Calvet, John Y. Campbell and Paolo Sodini* | 2006:195 |
| Efficient Bayesian Inference for Multiple Change-Point and Mixture Innovation Models<br>*by Paolo Giordani and Robert Kohn* | 2006:196 |
| Derivation and Estimation of a New Keynesian Phillips Curve in a Small Open Economy<br>*by Karolina Holmberg* | 2006:197 |
| Technology Shocks and the Labour-Input Response: Evidence from Firm-Level Data<br>*by Mikael Carlsson and Jon Smedsaas* | 2006:198 |
| Monetary Policy and Staggered Wage Bargaining when Prices are Sticky<br>*by Mikael Carlsson and Andreas Westermark* | 2006:199 |
| The Swedish External Position and the Krona<br>*by Philip R. Lane* | 2006:200 |

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