SIMULTANEOUS CONFIDENCE BANDS: EXPLORING THE DISTRIBUTION OF IMPULSE-RESPONSE FUNCTIONS*

Miroslav Plašil
Czech National Bank

This Version: February, 2017

Preliminary and Incomplete

* First version: 4 December 2017. I would like to thank Michal Andrle for inspiring discussions. All errors are, however, solely my own.

Foreword and Executive Summary

This note presents a collection of exploratory and inferential methods designed for a better understanding of the posterior distribution of impulse-response functions (IRFs) which are generated by structural VARs or other structural models. Importantly, the IRFs identified with sign-restrictions can be tackled by these methods as well. Initial exploration of the set of impulse responses seems to be a natural thing to do and its necessity would be acknowledged by many practitioners. In applied practice, however, natural analytical instincts seem to be somewhat suppressed and both exploration and presentation of the results is often reduced to a set of marginal (pointwise) distributions – a practice which might be insufficient or even misleading. This commonly applied approach can be particularly problematic for VARs identified with sign-restrictions where each sampled IRF corresponds to a different structural model.

One of the potential reasons for the current state of affairs might be that practitioners lack suitable tools for more subtle analysis of the posterior distribution of IRFs. However, such tools are readily available, at least in concept. At this point, it is useful to note that results of the analysis are oftentimes digested visually by inspecting the IRFs plot. This provides a useful format for digesting relevant information as visual cognition of humans is extremely strong. But this also suggests that a pursuit for a better understanding the IRF behaviour can be based on visual exploration. To make visual exploration as effective as possible it might be useful to stick to the philosophy of modern multivariate visualization and draw on the ideas set out by the giants of the exploratory visual analysis including Jacques Bertin or John Tukey.

Modern visualization is an interactive exploratory process where the analysts interrogates the data and receives answers to her questions. It sets out a pursuit for interesting patterns that aid interpretation of the underlying structure of the data. In more formal parlance, the process of „asking questions“ corresponds to formulating some optimization problem of analytical interest and "answering questions" means to find and graphically display its optimum. In this note, I do not have an ambition to codify a set of questions that should be *always* asked, neither to provide much guidance on what plots should be always reported. The note should only serve as an invitation to a certain concept of reasoning.

In the context of IRFs, many empirically relevant questions can be answered within a concept of simultaneous confidence bands, i.e. bands containing entire path of responses over time (up to a fixed horizon). Although simultaneous bands are traditionally related to statistical inference, they might be interpreted in more general terms. In particular, the idea of some constrained area with desired properties where its width, shape or both convey some meaningful information fits well to exploratory needs.

The relevance of questions can be analysis-dependent, but some of the questions will probably appear more frequently. Is the pointwise median a fair representative of the *actual* shape of the IRF and does it truly capture a central tendency of its distribution? Is the path of the majority of sampled IRFs concentrated around the point-wise median? Do the sampled IRFs share the common shape or do they exhibit zig-zag pattern with many crossings? Are there any "dense" regions close to critical values formed by point-wise quantiles (i.e. at the border of traditional confidence bands)? What is the portion of the sampled impulse-response functions whose shape goes against prior views of the analyst? What are the narrowest simultaneous confidence bands with user-specified coverage? These are all questions that commonly arise in the final phase of the analysis.

Although different exploratory questions may emerge in the course of the analysis, they only differ in the formulation of the objective function to be optimized. It follows that technical implementation of the exploratory analysis does not need to change once practitioners are able to formulate their question as an objective function. In general however, the optimization is difficult since one has to optimize over all horizons of the IRF. Since the horizon length will frequently attain dozens of periods, this leads to very complex and computationally highly involved optimization problems that can hardly be efficiently solved in a real time.

To reduce the complexity of the problem, I make an assumption that the shape and the width of the band do not change too frequently and eventual changes are relatively smooth. This assumption enables to optimize the objective function only over few preselected horizons and reconstruct whole path of the band using some smoothly behaved interpolation method. In particular, the reconstruction can be treated as a missing observations problem where the path of the band can be obtained by a convenient filtering (smoothing) algorithm. To keep computational burden low, Hodrick-Prescott filter with observations missing for all but preselected horizons is employed in practical applications.

Due to potentially high complexity and poor tractability the objective function, the optimum is generally found by genetic algorithms. This adds to computational burden, but sensible solutions can usually be obtained after relatively low number of eliminations rounds of the algorithm. Depending on a form of objective function and the number of posterior draws the whole optimization should take no more than few minutes which should not be prohibitive in practical applications. Since genetic algorithms allow for parallel computing, calculations might be further sped up considerably.

The experience so far suggests that the method generates analytically helpful outputs in reasonable time. In cases where some comparison with other methods is possible the method seems to provide very promising results. In particular, the proposed method regularly beats all its known competitors by a great margin when the goal is to construct narrowest confidence bands with desired coverage.

1 Introduction

In many practical situations, a non-linear function of model parameters such as the impulse-response function, multi-period forecasts or some frequency-domain characteristic is used to summarize dynamic behavior of the model. This is because individual coefficients are sometimes hard to interpret or convey only limited information which is not sufficient for full understanding of the model's most relevant features. Functions in question can be readily presented in visual form and thus provide intuitive and appealing tool for both practitioners and wider public to digest the outcomes of the underlying analysis.

To present estimation results or prior-implied behavior of the model through its higher-level features – say an impulse-response function (IRF) – an analyst usually provides the median and supplies a reader with the line of pointwise quantiles as a measure of estimation uncertainty. Unfortunately, such a visual display is not always very informative about the actual behavior implied by the model, or even worse: it can be sometimes highly misleading.

Fig. 1 illustrates the issue. Namely, the panels depict two sets of impulse-responses which share almost the same median and pointwise 67% confidence bands. If the analysis stops here, it might go unnoticed that the models actually produce very different sets of impulse response functions. While the impulse responses in the left panel are relatively similar to each other and share the same hump-shaped pattern, impulse-response functions depicted in the right panel exhibit highly irregular zig-zag behavior. Naturally, this has serious implications for the overall assessment of the model's economic plausibility. By the same token, heedless application of pointwise regions may also blur the soundness of parameter-implied priors if prior predictive analysis is carried out.

It has been recognized for quite a long time that simultaneous confidence/credible intervals may provide desired remedy in situations where distortions from the use of "pointwise" approach arise (see e.g. Jordà, 2009; Inoue and Kilian, 2016; Lűtkepohl et al., 2016 or Montiel Olea and Plagborg-Mőller, 2017). Simultaneous bands may help answer relevant questions about the shape of the IRFs, hypotheses about positive response in selected periods and many others. In this note, I argue that simultaneous confidence/credible[1] bands may also serve as a straightforward and versatile tool for the thorough exploration of the impulse-response functions' distribution.

[1] Despite methodological differences, the terms "confidence bands" and "credible bands" are used interchangeably in this note. Although incorrect from Bayesian perspective, referring to "confidence bands" is sometimes preferred due to the relation to the existing literature.

Figure 1: Example of identical pointwise characteristics for different set of IRFs

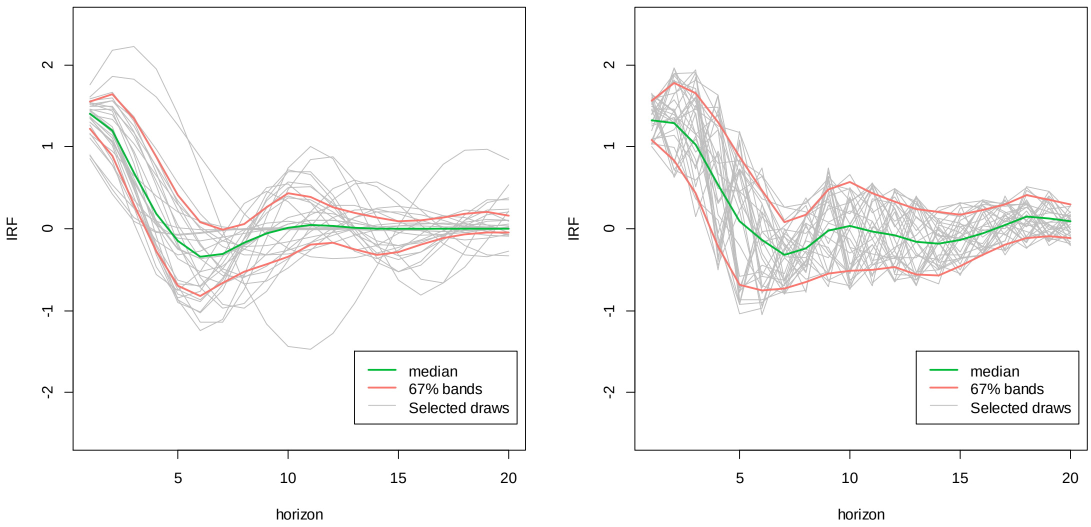

In wider terms, simultaneous confidence bands can be understood as a constrained area showing some desirable feature. User-specified simultaneous coverage would be the most natural requirement for inferential purposes, but the feature to be explored can be of a more general nature. Are there any impulse responses lying close to the pointwise median line and is the calculated median reasonable summary of the set of IRFs? Does the set suggest some common pattern across IRFs or do the individual IRFs show highly irregular behavior? Where is situated the most "dense" region and are there any dense regions lying at the ends of the IRFs distribution? What is the narrowest possible credible band with a predefined coverage? These are all legitimate questions that contribute to a better understanding of the IRFs' characteristics and enhance practitioners' knowledge.

Importantly, all these questions can be answered within the context of simultaneous bands. Deeper exploration of the behavior might be particularly needed in structural VAR models identified with sign restrictions where a set of IRFs does not describe behavior of a single model but summarizes information coming from all admissible (and potentially very different) structural models. At this point, it is also worth noting that although further discussion is limited to impulse-response functions (IRFs), the methodology proposed below can be equally employed for other functions of individual coefficients where similar needs for simultaneous inference emerge. An example would be a summary of multiple-horizon forecasts.

To facilitate *exploratory* process of the posterior distribution of the IRFs, I propose a suite of very simple methods for the construction of simultaneous confidence bands with user-specified properties. All methods are based on the same principle and only differ in a desired feature to be exhibited via the IRF plot. The approach is primarily designed for Bayesian inference where IRFs are derived from posterior draws of parameters but it may also be employed for the construction of the bands from bootstrap samples. It provides enough flexibility to account for varying exploratory strategies pursued by practitioners and its algorithmic implementation is fairly straightforward.

It should be noted, however, that the flexibility comes at the price of higher computational burden since the method relies on explicit optimization methods. Nevertheless under some reasonable assumptions which will be discussed below the complexity of the optimization problem can be considerably reduced. The experience so far suggests that reasonable results can be obtained in order of minutes using standard personal computer (and only one core of the processor), which should not be prohibitive in real-world applications. Moreover, optimization algorithm applied below allows for parallel computations which can greatly speed up the searching process and reduce the computation time to only few minutes or dozens of seconds.

Even though exploratory motives prevail somewhat over inferential purposes in this note; the method can be readily used for traditional inference. If the ultimate goal is to produce simultaneous credible intervals with desired coverage then the proposed method seems to always provide the narrowest bands among all its known competitors. Although the application of the method to the real-world problems is still limited, the performance observed so far is encouraging.

The remainder of the note is organized as follows. In Section 2, I sketch out the method of constructing simultaneous confidence bands and discuss its merits and limitations. In Section 3, the proposed methodology is illustrated on a simple example of an IRF generated by the univariate AR(2) process. Some exploratory strategies with a potential to provide useful insights to practitioners are presented. Section 4 documents the performance of the proposed approach in real world applications and the final section concludes. Computational details are relegated to Appendices.

2 Construction of simultaneous credible intervals

Recently, the interest in simultaneous credible intervals for impulse response functions has gained some momentum and a bulk of methods has been proposed in the ("VAR-oriented") literature. Their thorough description is beyond the scope of this note and interested readers are referred to Lűtkepohl et al. (2016) or Montiel Olea and Plagborg-Mőller (2017) for a comprehensive review. Note that these methods were all designed with inferential objectives in mind. This is not the ultimate objective pursued here, but it may serve as a good starting point as the goal to create simultaneous confidence bands with desired coverage is well understood by the practitioners. Later, other objectives will be pursued using the same methodology.

If we disregard inferential methods drawing on asymptotic considerations which usually show disappointing performance in real-world applications (cf. Lűtkepohl et al., 2016), the majority of methods for simultaneous bands relies on classical techniques of simultaneous parameter inference applied to bootstrap or posterior distribution of IRFs. Common problem with these methods is that they tend to be too conservative and provide unnecessarily wide bands. Moreover, their relative performance may vary across applications and it is thus not clear which method should be preferred in the particular case. To provide a way out, Montiel Olea and Plagborg-Mőller (2017) introduced a one-parameter class of confidence bands and showed that it includes most of the popular choices in applied work.[2] The one-parameter band, $\widehat{\mathcal{B}}(c)$, can be defined as:

$$\widehat{\mathcal{B}}(c) \equiv \left[\hat{\theta}_1 - \hat{\sigma}_1 c, \hat{\theta}_1 + \hat{\sigma}_1 c\right] \times \left[\hat{\theta}_2 - \hat{\sigma}_2 c, \hat{\theta}_2 + \hat{\sigma}_2 c\right] \times \ldots \times \left[\hat{\theta}_k - \hat{\sigma}_k c, \hat{\theta}_k + \hat{\sigma}_k c\right],$$

where $\hat{\theta}_j$ is the value of impulse-response at time j, $\hat{\sigma}_1$is the pointwise standard error for $\hat{\theta}_j$, $c > 0$ is a positive scaling constant and k is the length of the IRF vector. In other words, one-parameter class is defined as a Cartesian product of scaled-up versions of the traditional pointwise confidence intervals where the same scaling factor is used for all elements of the IRF vector. Montiel Olea and Plagborg-Mőller, 2017 then suggest using sup-t bands as an optimal representative of the one-parameter class and show that these are indeed the narrowest possible bands within the class. The bands are obtained by picking the smallest c that guarantees predefined simultaneous coverage equal to $1 - \alpha$. The implementation of sup-t bands for posterior or bootstrap draws is straightforward and computationally convenient – simply calculate empirical pointwise α quantiles and scale α up or down with a single factor to obtain critical values that guarantee the exact simultaneous coverage of $1 - \alpha$. Since the fraction of IRFs contained in the simultaneous bands is a monotonic function in the scaling factor, the univariate optimization is fairly easy and optimal value of the factor can be obtained very fast (even by a trial-and-error approach).

Unlike other traditional methods falling into the one-parameter class, construction of sup-t confidence bands requires explicit optimization; although an uninvolved one.

[2] The class includes Bonferroni or Šidák simultaneous bands, for example.

As such, sup-t bands may serve as a useful benchmark for more complex optimization-based methods. In principle, the paths of critical values can be found by using different scaling factors for each horizon so as to recover the narrowest bands with nominal coverage of $1-\alpha$. However given the horizon length, this approach leads to multivariate optimization problem of alarming complexity. In practice, heuristic searches were proposed to tackle the issue (see Staszewska, 2007 and Staszewska-Bystrova and Winker, 2013), but the proposed heuristics are computationally demanding and do not even guarantee desired coverage. In particular, they usually tend to produce bands with actual coverage below nominal level (Lűtkepohl et al., 2016).

The method proposed in this note can be seen as a compromise between excessively complex multivariate search of critical values over each horizon and a univariate optimization where the shape and width of bands are regulated through a single tuning parameter. The method does not rely on explicit scaling of the pointwise quantiles and can be applied directly to actual "values" of the IRF defining the band (however, individual scaling of the quantiles is possible, if needed).

Reduction in complexity is based on the idea that the intervals' shape and width do not exhibit too frequent abrupt changes from one horizon to another. Under this assumption, the overall shape can be reasonably approximated by only selecting critical values for a small number of horizons. Remaining critical values are then set equal to values of some well-behaved (and relatively smooth) function. There are many options how to implement this general idea in practice.

In this note, I treat the interpolation of the bands' path between predefined points as a traditional *missing observations* problem with the observations missing for all but predefined horizons. Specifically, I apply Hodrick-Prescott filter principles to derive a smooth path for the bands (see e.g. Schlicht, 2008 or technical Appendix A for computational details). This approach guarantees that desired critical values are exactly attained in predefined horizons and smooth transition of critical values is observed otherwise. Fig. 2 provides some illustrative examples of the confidence bands constructed by this approach for a varying number of predefined time-points and different values of the smoothing parameter lambda in Hodrick-Prescott filter.

I will return to practical questions of how to select the initial number of time-points (number of preselected horizons) and the value of the smoothing parameter for Hodrick-Prescott filter in Section 2.1. At the current stage it is only important to realize that the goal of finding simultaneous confidence intervals with desired features (e.g. narrowest bands with predefined coverage) can be thought of as an optimization problem where the optimum of the objective function is only optimized over the arguments related to preselected time-points. In general, the objective function can take various forms depending on the exploratory strategy pursued by a researcher. The most obvious choice for inferential purposes would be to minimize the overall width of confidence bands guaranteeing the desired coverage but other – more exploration-driven – choices are possible and scientifically relevant. Some of them are discussed in the following sections.

Figure 2: Example of the bands based on a different number of time-points

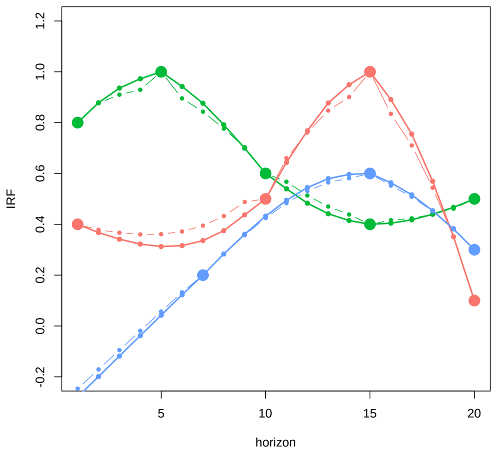

Since the objective function can be quite complex to be solved analytically, I find the (sub)optimum by means of real-valued genetic algorithms (see e.g. Wright, 2001 or Herrera et al. 1998).[3] Optimization by genetic algorithms naturally adds to computational burden but sensible solutions can usually be found in reasonable time if the initial population is not too far from the optimum.

In many situations of practical interest the exploration strategy may result in a constrained optimization problem (e.g. finding the narrowest confidence bands given the user-specified coverage) and genetic algorithm must be adapted to handle these constraints. Many constraint-handling techniques were proposed in the literature – the penalty-function approach being arguably the most popular one (see for example Ponsich et al., 2008 for a review). This is also the strategy pursued here. Namely, infeasible solutions are assigned a prohibitively high penalty to obtain low probability to survive in the next elimination round.[4] This method seems to work well for the problems at hand, if the initial population contains some feasible solutions. If these solutions are hard to find by generating random populations, one may obtain initial feasible solutions by first applying genetic algorithm to auxiliary

[3] These are all-purpose optimization routines inspired by natural evolution. Unfamiliar readers are encouraged to consult existing literature.

[4] This is very close in spirit to "death penalty method" where infeasible solutions are given zero chance to pass the selection step.

objective function minimizing the violation of the constraints (e.g. using objective function that minimizes the difference between observed and user-specified coverage).

Another option is to follow the suggestion of Chehouri et al. (2016) who propose a genetic algorithm operating simultaneously on the population separated into two families. These are formed by feasible solutions and solutions that violate the constraints, respectively. In each family different objective function is optimized and solutions may travel from one family to another one. This approach may lead to additional computational cost, however it might be worth a try if it is difficult to build sensible initial population and many rounds of the genetic algorithm would be needed to reach the optimum. For traditional inference, feasible solutions can be generated by other inferential methods that are known to guarantee the predefined simultaneous coverage (e.g. sup-t bands).

In practice, I used a population of 30 individuals and 20 to 100 elimination rounds to report all the simultaneous bands in this note. Further improvements in the optimum were usually quite small and analytically insignificant. Technical details on the exact implementation of the genetic algorithm; in particular the choice of cross-over and mutation operators and survivor strategies are given in technical Appendix B.[5]

2.1 The Choice of Arbitrary Parameters

The proposed method relies on some arbitrary decisions which are fully in hands of the practitioners. In particular, the researcher needs to select a number of time-points and their location to obtain a reasonably-behaved path of the confidence bands. The issue of somewhat lesser importance is the choice of the parameter lambda used to smoothly connect all arbitrarily chosen time-points. This subchapter discusses practical hints that may guide this selection. It also highlights some possibilities how to minimize the arbitrariness of the choices.

No doubt: the concrete choice of a number of time-points and their position is analysis-dependent and subject to some initial exploration. One has to face a common tradeoff: the number must be relatively small to reduce the complexity of the problem but not excessively small to preserve the flexibility of the method. It seems natural to select the very first and the very last time-point of the confidence interval and complement those with few points lying somewhere in between. In

[5] It should be stressed that technical implementation adopted here should not be seen as golden standard but rather as a proof of concept. Other options may work equally well or even considerably better. Practitioners are advised to stick to their preferred implementation of genetic algorithms or other flexible optimization routines.

general, it is useful to pick those horizons where a break in the shape of confidence bands is to be expected.

This information can sometimes be inferred from the shape of the pointwise median or based on some other methods for simultaneous confidence bands which are fast to compute. In cases where no such information is available (for example, this can be complicated for IRFs identified with sign restrictions) the choice of five or six equidistantly placed time-points should work reasonably well in the vast majority of practical applications. Some experimentation with possible alternatives is a useful part of the exploratory process that contributes to a better understanding of the analysis outcome.

Contrary to the usual practice in HP filtering applications, the smoothing parameter lambda needs to be set to a quite small value. Setting the value to 1 should work well in general and using values above 10 does not usually make much sense unless one has a specific reason for doing so. For higher values of lambda, the penalty on smoothness of the band's path is excessively high which produces (usually undesirable) spikes in predefined time-points where specific values have to be attained – see Fig. 3. The plausibility of the smoothing parameter can be checked by visual inspection and compared to prior views on the shape of the IRF.

Figure 3: Spikes produced by setting too high smoothing parameter

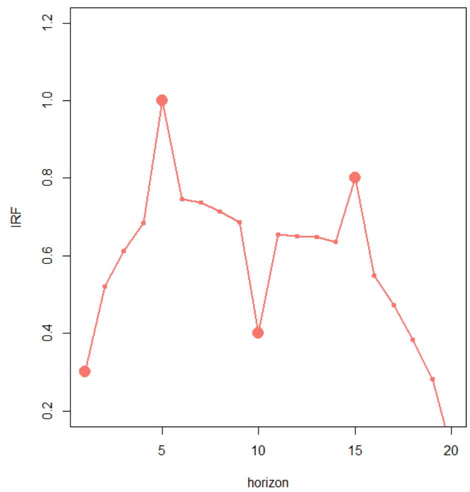

If desired, the arbitrariness of the presented choices can be substituted by explicit optimization over additional arguments. Note that it is straightforward to express the objective function as a function of additional parameters, say, the smoothing parameter or time-points' location. These would simply represent additional genes of the individual (chromosome) in the genetic algorithm. For example, if the number of the time-points is set to d, one simply adds d genes into each chromosome to encode their position. One additional gene would be necessary to encode the value of the smoothing parameter.

2.2 Illustrative example: IRF of a simple AR(2) process

To illustrate the method proposed above I use an impulse-response function describing behavior of the univariate autoregressive process of order 2. Its marginal parameter priors were updated with a system prior that the AR(2) model was designed for capturing business-cycle frequencies ranging between 2-8 years (see Andrle and Plašil, 2017 for details). As a part of the prior predictive analysis practitioners may ask whether such system restrictions on parameters also discipline the course of the IRF and if its shape is fully in line with their prior views. To shed some lights on this issue, 5000 posterior parameter draws were stored to calculate implied path of IFRs. Horizon length was set to 20 periods. Readers unfamiliar with system priors or prior predictive analysis do not need to worry and may look at the following example as if it were an outcome of some traditional Bayesian estimation.

Fig. 4a presents a traditional summary with the point-wise median and 67% point-wise credible intervals (using 16.5% and 83.5% quantiles). Fig. 4b compares pointwise credible bands with simultaneous credible bands obtained by a benchmark sup-t algorithm. As expected, pointwise credible intervals are considerably narrower as their simultaneous coverage is considerably below 67%. Despite different width of the bands, both methods seem to deliver similar shape of the IRFs, with diminishing but still conspicuous oscillatory waves up to horizon 20.

Recall that sup-t bands are the narrowest simultaneous credible bands within the one-parameter class and can obtained very fast. A natural question is if we can do any better by sticking to computationally more intensive methods outside this class – in particular, by exploiting the method outlined above.

To reduce the complexity of the problem I make a use of a "symmetry" principle where I only optimize over the time-points defining the "upper" band. The lower band is then obtained automatically – as its mirror image around the pointwise median. The assumption of the "central" role of the pointwise median is not necessary in general but it saves computational cost and makes the comparison with sup-t bands more straightforward (note that sup-t bands are actually based on the very same assumption). The bands were obtained using 6 time points placed at the horizons 1, 3, 7, 10, 15 and 20. Smoothing parameter was set to 1. Placement of the time-points was chosen with respect to the shape of pointwise median.

To start optimization routine, two options can still be considered: namely, one can minimize the overall width of the bands either *i*) indirectly over the set of pointwise quantiles defining the critical values, or *ii*) directly over the set of actual values of the IRF.[6] First approach can be seen as a multivariate generalization of the sup-t band algorithm where different scaling factors are used for each horizon. The bands obtained through indirect and direct approach are presented in Fig. 4c and Fig. 4d respectively; along with the basic sup-t band to facilitate their mutual comparison.

Figure 4: Comparison of pointwise and simultaneous credible bands

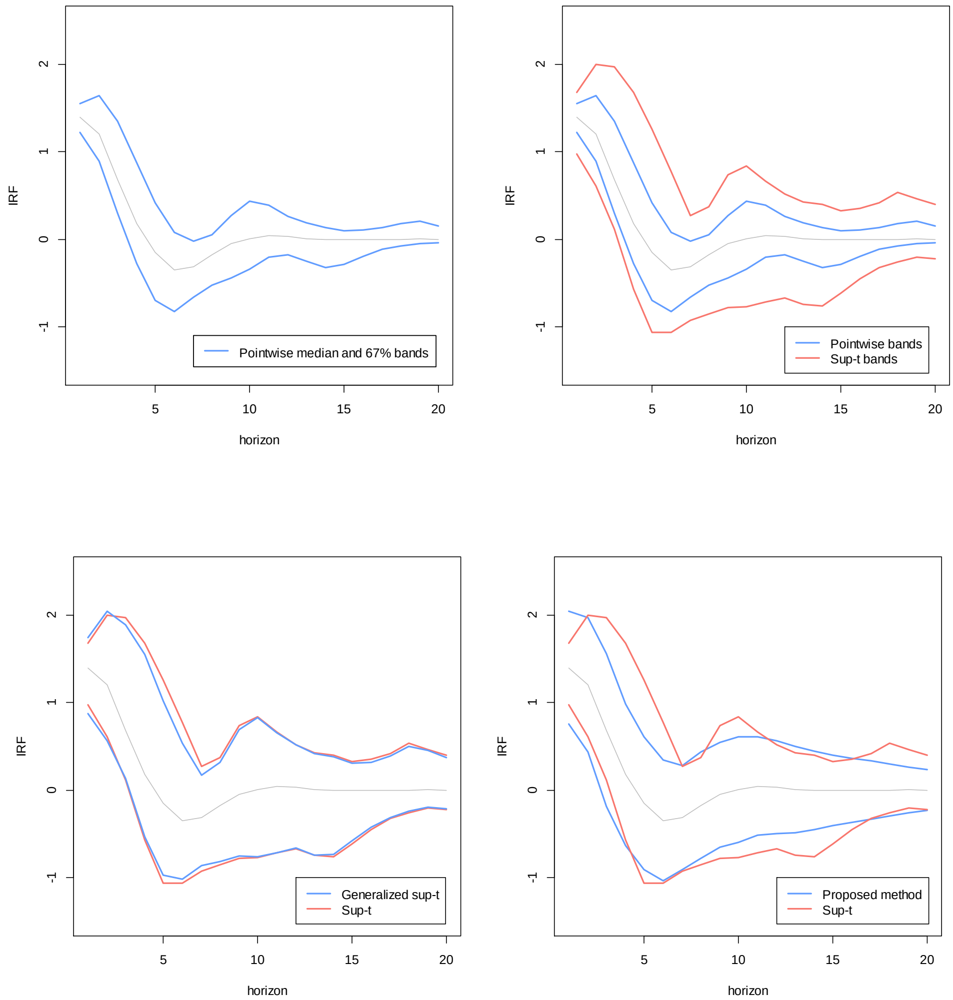

The comparison suggests that the basic sup-t bands which are based on univariate optimization seem to work quite well within the methods based on the scaling of quantiles since the observed improvements attained by employing different scaling factors are quite minor overall. Recall however that the "quantile-scaling" methods are generally ineffective if they are used for inference in sign-identified VARs since

[6] These two options will differ due to the assumption of symmetry. Note that the bands based on indirect method are not necessarily symmetric in this case, since the principle of symmetry is only applied to quantiles. This does not guarantee symmetric bands for asymmetric pointwise distributions.

point-wise quantiles make little sense in this case. Luckily, the application of our method directly to actual values of the IRF can still be used in the context of sign restrictions and quite importantly it also seems to provide sizeable gains vis-à-vis quantile-based methods in situations where both approaches are viable (see Fig. 4d). One can clearly see that obtained confidence bands are narrower than all its competitors and exhibit less wriggling pattern. This is naturally relevant for practical applications if simultaneous credible bands are to be used for inference.

Now suppose that a practitioner wants to learn more about how representative the pointwise median is. In many situations doubts about its usefulness as a summary statistics are warranted since dense regions of IRFs' posterior distribution may lie elsewhere and may exhibit substantially different shape. It is thus worth exploring this aspect more thoroughly. Underlying strategy that corresponds with the goal of the exploration would be to find dense simultaneous regions and inspect their shape. This can be again achieved by using strategy of narrowest simultaneous bands with user-specified coverage, however with a difference that this time the bands may lie anywhere in the IRF space and are not by construction tied down to pointwise median. If the shape and magnitude of the pointwise median coincide with bands' path for all horizons, it can be considered a useful summary.

Simultaneous bands complying with the outlined strategy can be constructed in the following way: upper band is obtained as before while the lower band is calculated as the upper band minus interval width. The width is assumed to differ across horizons but only its relatively smooth changes are allowed. This again reduces the complexity of the problem and allows for the optimization over a smaller number of arguments. More concretely, the objective function is optimized over 9 parameters. First six parameters encode the shape of the (upper) simultaneous band. Identically to previous exercise relevant time-points are placed at horizons 1, 3, 7, 10, 15 and 20. Last three parameters then encode the interval width using the horizons 1, 10 and 20. To maintain the complexity of the problem manageable, the width is only encoded using three parameters. In practical applications the interval width seems to be considerably less volatile than the bands themselves, thus this assumption should be quite innocuous. However, experimentation with more complex encoding is of course possible.

Quite naturally, finding optimal solution is now computationally more involved than in former case. This is partially because one has to optimize over larger number of parameters but also some additional complications arise. In particular, it is more difficult to generate initial population of feasible solutions as one usually needs to start quite far from the optimum to cover the entire space of the IRF's posterior distribution. This increases number of algorithm rounds necessary to reach final solution – in particular for bands with low coverage. In practice, I started with generating very wide confidence bands that covered entire space and used auxiliary objective function to minimize distance between observed and user-specified coverage. Once desired coverage was attained, I minimized the width of the bands to find the densest regions complying with desired coverage.[7]

Obtained simultaneous credible bands are presented in Fig. 5. Left panel depicts simultaneous credible intervals with 30%, 50% and 67% coverage and the pointwise median. The results indicate that the pointwise median seems to be a relatively fair summary of the IRFs in this particular case as it lies entirely within the 30% simultaneous credible region. This is also corroborated by direct comparison of the bands with those based on "symmetry" assumption – see Fig. 5, right panel. This is not surprising since the behavior of the IRF was a priori disciplined by – economically-motivated – restrictions (see above) to exhibit reasonable performance.

Figure 5: Simultaneous credible bands and the shape of the posterior distribution

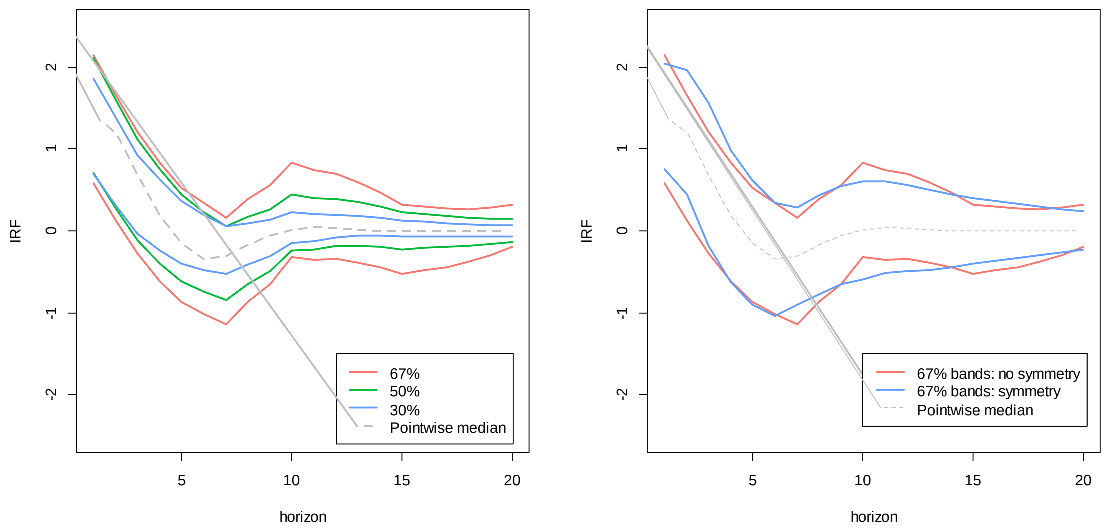

In addition, the results also confirm that posterior mass of the IRF distribution is mainly concentrated around some "central" IRF with simultaneous confidence intervals being quite symmetrical and lying inside each other. This may provide additional indication that the course of all IRFs shows quite similar and well-behaved pattern. As the bands (their coverage) widen, IRFs exhibit increasingly

[7] In theory genetic algorithms do not guarantee attaining global optimum and the algorithm may get stuck in local optima. This may lead to different simultaneous credible bands each time the optimization is performed. While different solutions may support exploratory goals of the analysis, some users may find this feature undesirable. To inspect a danger of occurrence of locally-optimal solutions in greater detail, I run the optimization couple of times with different starting populations and all solutions were almost identical. This provides some positive indication of a convergence, although it does not need to hold in other settings.

discernable downs (at about horizon 7) and ups (horizon 10) which suggests a presence of higher and longer-lasting oscillations at the "ends" of the distribution. This may be explained by the form of employed system prior which places maximum weight on business cycle frequencies but also passes through some portion of longer oscillations. If such oscillatory behavior in IRF is at odds with practitioner's prior views, she may consider elicitation of tighter priors on business-cycle frequencies.

3 Practical applications

3.1 Sign-identified monetary shock

To do…

4 Conclusions

To do…

Statistical inference and formal statistical tests are only valuable when one fully understands what the data tell us. However, once you know your data thoroughly, all statistical inference becomes redundant.

Literature

TBC