Joint Bayesian Inference for DSGE Models*

Joe Marlow†

This draft: February 2025.

Preliminary, please do not circulate

Abstract

The main criticism of the standard, pointwise approaches to constructing the central estimates and inference in DSGE models is that such approaches make the assumption that responses of macroeconomic variables are independent across time and variables. This paper addresses this by extending the literature on joint Bayesian inference to DSGE models, as well as proposing a way to construct a good measure of the central tendency that is consistent with the model equations. To do so, a Bayes Estimator with Euclidean norm loss is used for the objects of interest. The vector that minimises this loss is shown to be a good measure of the central tendency, while preserving the model structure, and the joint credible set is constructed as the lowest posterior risk region under this loss. An algorithm allowing for joint Bayesian inference of impulse responses and the Forecast Error Variance Decomposition in a DSGE setting is introduced, and the drawbacks of existing approaches are demonstrated using the Smets and Wouters (2007) framework. In addition, some of the practical concerns that researchers may have with using the Bayes Estimator under popular loss functions are discussed, with potential solutions provided.

Keywords: DSGE, Impulse Response Function, Forecast Error Variance Decomposition, Joint Inference, Bayesian.

JEL: C01, C11, C32, E00, E12.

*I would like to thank Alessio Volpicella, Valentina Corradi, Jordi Llorens-Terrazas, Alistair Macaulay and Ricardo Nunes for helpful comments and suggestions.

†University of Surrey, Elizabeth Fry Building (AD), Guildford GU2 7XH, UK. Email: j.marlow@surrey.ac.uk

1 Introduction and Related Literature

Dynamic Stochastic General Equilibrium (DSGE) models are now the workhorse model in macroeconomics, both in academia and policy-making. Increasingly, these models are estimated using data, with most researchers opting for a Bayesian approach due to the relative simplicity of exploring the posterior distribution compared to the likelihood (Fernández-Villaverde, 2009)[1]. The posterior distribution for the parameters are then used to construct estimates and inference for the ultimate objects of interest; usually impulse responses and the forecast error variance decomposition. The most popular approach involves constructing the central estimate, as well as the corresponding inference, using pointwise approaches[2]. Pointwise approaches for constructing the central estimate involve simulating the posterior distribution of the impulse responses and taking the mean (or median) of the distribution at each horizon, for each variable and shock, treating each point independently. For example, central estimates of the impulse responses in Smets and Wouters (2007) and Del Negro, Schorfheide, Smets, and Wouters (2007) are constructed using the posterior mean impulse response, while Justiniano, Primiceri, and Tambalotti (2010) and Mukoyama et al. (2021) use the posterior median, which coincides with standard practice in the VAR literature. However, it is unlikely that the parameter values used to construct the mean response are the same for every variable, horizon and shock. As a result, such measures of central tendency are likely to be inconsistent with the model structure and force practicioners to make the implicit, and unrealistic, assumption that the impulse responses are independent across horizons and variables. Such assumptions are problematic when making statements about co-movement between variables or how responses evolve over time, as it is unclear how much of the changes are being driven by the model structure.

The independent treatment of each point also presents issues for inference, as noted by Sims and Zha (1999). By ignoring the dependence across time and variables, the degree of estimation uncertainty is often understated by the pointwise credible sets that are reported in many applications. These considerations are of high importance to practitioners using DSGE

[1] Bayesian approaches are less exposed to the difficulties of exploring the likelihood of DSGE models which arise due to the high dimensionality of the problem. See, for example, Chernozhukov and Hong (2003).

[2] Dynare 6.1 provides the posterior mean and median impulse response for each variable and horizon, with the credibility regions also calculated for each variable and horizon independently.

models, so an alternative method of constructing impulse responses and the Forecast Error Variance Decomposition (FEVD) that is immune to these criticisms is desirable.

An alternative approach that is widely used and preserves economic interpretability involves collapsing the posterior distribution of the parameters into point estimates, and using these to generate the objects of interest (e.g. Iacoviello and Neri (2010); Christiano, Motto, and Rostagno (2014); Bratsiotis and Pathirage (2024)). While this approach ensures that the objects of interest are generated by the same parameter vector, it does not guarantee a good representation of the central tendency for the impulse responses due to the lack of invariance of such measures of central tendency under non-linear transformations. In addition, by collapsing the posterior distribution of the parameters into a point estimate, researchers ignore the estimation uncertainty involved in the construction of the impulse responses, making the construction of credible sets impossible.

The main contribution of this paper is to address the above concerns in the context of DSGE models. I construct a Bayes Estimator that coincides with a constrained version of the spatial median. I show that the resulting minimiser is the closest point to the true spatial median that preserves the economic structure imposed by the researcher. As a result, the Bayes Estimator allows the researcher to obtain a good measure of the central tendency of the joint distribution of the impulse response functions, while maintaining the economic interpretability that comes from the use of a single structural parameterisation. The treatment of each point independently is also a concern when constructing inference. The joint credible set described in this paper is able to more accurately characterise the estimation uncertainty that is present in the estimates for the ultimate objects of interest.

Bayesian statistics relies on two main approaches when summarising posterior distributions: the Bayes Estimator and the Maximum a Posteriori (MAP), which is also commonly referred to as the posterior mode. Both approaches are capable of achieving the objective of maintaining economic interpretability while taking into account the dependence structure across time and between variables. However, the Bayes Estimator is the more computationally appealing of the two. This is because the MAP estimate requires the researcher to integrate over the joint posterior distribution. In DSGE models this is generally not derived from a known distribution, so would rely on computationally costly numerical algorithms. The Bayes Estimator, on the other hand, is obtained by minimising the expected posterior loss, where the loss function is selected by the researcher. By constructing the loss function to take into account the full vector, rather than just the individual point, we are able to provide a measure of central tendency that maintains economic interpretability. Computationally, this is more straightforward than obtaining the MAP estimator, so is more appealing in this setting.

Inoue and Kilian (2022) propose a similar solution for VARs, which also face the same concerns regarding pointwise approaches. In their paper, the authors generalise earlier work to the situation where there are more impulse responses than structural parameters in the model, as is the case in many applications of interest. In such settings the asymptotic equivalence between the central moment obtained using pointwise and the Bayes estimator breaks down. Although this paper shares the same spirit as Inoue and Kilian (2022), it differs in several dimensions. First, this paper provides a toolkit for empirical practitioners interested in using joint inference in DSGE models, rather than VARs. In particular, I demonstrate that the asymptotic equivalence between pointwise measures of central tendency are almost always infeasible in DSGEs due to the large amount of endogenous variables relative to the number of estimated parameters. This lack of asymptotic equivalence is not always present in VARs, where the larger number of parameters means that asymptotic equivalence between the two estimators is still sometimes possible in scenarios that would be of interest to researchers. That said, the theoretical arguments concerning the choice of loss function and inference are also applicable to the VAR literature.

Second, I demonstrate that the Euclidean norm loss is a sensible choice in this setting, as the Bayes Estimator is, in general, the closest feasible point to the Spatial Median, a well-established measure of central tendency in a multivariate setting. I also demonstrate that additively separable loss functions often produce estimates that are closest to pointwise central tendencies, which are typically not representative of the central tendency of the joint distribution. Third, I extend the methodology to the Forecast Error Variance Decomposition (FEVD), which is another key object of interest for macroeconomists, and is particularly important for business cycle analysis. And fourth, I provide some general advice and suggestions for empirical practitioners using these methods in a time series context, which apply equally to those using DSGE models or VARs. I discuss how standard loss functions may be influenced by scaling or the degree of estimation uncertainty, and provide a solution to address such situations. In addition, I highlight the importance of choosing a policy-relevant horizon length for the impulse responses as the Bayes Estimator is not invariant to the horizon length selected for the impulse responses. Although in practice the effects of this are quantitatively small in the application considered, it may pose a larger concern for researchers using VARs, particularly if the role of estimation uncertainty is not accounted for. A simple solution to this problem is proposed, which involves using a loss function that discounts impulse responses at more distant horizons to prevent a seemingly arbitrary choice of impulse response horizon from altering the results for the policy-relevant horizons. These two features are important to take into consideration in a time series setting to ensure an accurate measure of the central tendency.

To highlight the effectiveness of the approach, I apply the methods to the workhorse DSGE Smets and Wouters (2007) model. I find that ignoring the dependence structure of the impulse responses leads researchers to significantly understate the estimation uncertainty when estimating impulse responses. In the empirical application, only 1 out of the 1,200 posterior impulse response vectors considered is fully contained within the 90% pointwise credible region: the Bayes Estimator under Euclidean norm loss, which lends some support to the idea that the Bayes Estimator under Euclidean norm loss offers a good measure of central tendency.

The paper is organised as follows. Section 2 discusses the issues with current approaches used in the literature. Section 3 introduces joint inference in the context of DSGE models, and sets out the algorithms used in the paper. An application using the Smets and Wouters (2007) model for comparison is shown in Section 4. Section 5 discusses limitations of using the Bayes Estimator in a time series setting, and proposes adjustments to the loss function which may help to mitigate these. Section 6 concludes.

2 Existing Approaches

After estimating the posterior distributions for the model parameters, researchers use these to simulate the posterior distributions for the ultimate objects of interest. There are several existing approaches in the literature for doing so, which can be classified into two subgroups. The first are referred to as pointwise approaches, as they treat each individual point of the impulse response vector independently. These typically involve taking the posterior mean or median of the simulated distribution of the object of interest as the central estimate and constructing pointwise credibility regions designed to reflect the estimation uncertainty. An alternative approach is to collapse the posterior distributions of the parameter estimates into point estimates and then use these fixed parameter values to construct the objects of interest. The Dynare default for this approach is to use the posterior mean for each parameter, but the posterior mode of the joint posterior distribution of the parameters is also commonly used. The following sub-sections outline the issues with both methods, which the approach suggested in this paper is able to address.

2.1 Pointwise Approaches

Pointwise approaches are the most popular method used in macroeconometrics, with the central estimate constructed as the median (or mean) of the posterior impulse responses for each variable, horizon, and shock. While as a point estimate they provide a useful measure of the central tendency for impulse responses, in general researchers are interested in the evolution of the shocks over time, and the co-movement of different economic variables. These latter points are poorly represented by point wise approaches, as the impulse responses are treated independently across time and variables. As a result, the model structure is often not preserved by the plotted impulse responses, making interpreting them from an economic standpoint problematic.

To illustrate this point, consider a simple AR(1) estimated using Bayesian methods, where $\phi \sim N(\phi_m, \Sigma_\phi)$, and $u_t \sim WN(0,1)$[3].

$$y_t = \phi y_{t-1} + u_t \tag{2.1}$$

[3] A similar example can be found in Inoue and Kilian (2022).

Provided that the process is stationary, Equation 2.1 has an $MA(\infty)$ representation that can be used to obtain the impulse responses. Researchers use the draws from the posterior distribution of ϕ to simulate the posterior distribution of the impulse responses. Conditional on the draw from the posterior distribution of ϕ, the posterior impulse responses associated with that parameterisation follow the model structure imposed by the researcher in Equation 2.1. For example, if ϕ_m is drawn, the first three impulse responses would be $1, \phi_m, \phi_m^2$ for horizons 0, 1, and 2, respectively. Therefore, for each draw of ϕ, the sequence of posterior impulse responses is consistent with the model imposed by the researcher, so can be interpretted from an economic standpoint. The issue with pointwise approaches is that the model-consistency breaks down when constructing the central estimate.

IRFs	$\mathbf{h = 0}$	$\mathbf{h = 1}$	$\mathbf{h = 2}$
Model-consistent	1	ϕ	ϕ^2
Posterior Mean	1	ϕ_m	$\phi_m^2 + \Sigma_\phi$

Table 1: Pointwise vs. Model-Consistent IRFs

As before, to be able to interpret the central estimate from an economic standpoint, we require the impulse responses to be consistent with the model that we impose, i.e. they should follow the path outlined in the first row of Table 1 for the AR(1) model. However, the posterior mean of the impulse responses is inconsistent with the AR(1) structure imposed for $h > 1$. On impact, the simulated posterior distribution for the IRFs is degenerate, as $\phi^0 = 1 \;\; \forall \phi$. At horizon 1, the posterior mean is given by $E[\phi] = \phi_m$. For the posterior mean to be consistent with the AR(1) structure, the impulse response at horizon 2 would need to be ϕ_m^2. This is not the case provided that the posterior distribution of ϕ is not degenerate. Rearranging $Var(\phi) = E[\phi^2] - E[\phi]^2$ gives:

$$E[\phi^2] = Var(\phi) + E[\phi]^2 = \Sigma_\phi + \phi_m^2 \tag{2.2}$$

Where $E[\phi^2]$ is the posterior mean impulse response at horizon 2. Unless $\Sigma_\phi = 0$, there is no ϕ that can generate the impulse responses generated by the posterior mean impulse response vector. As a result, the measure of central tendency obtained by taking the mean of the posterior impulse response at each horizon is not consistent with the model imposed by the researcher, so cannot be used to make statements about the path of the endogenous variable over time following an exogenous shock. This is because it is unclear when interpretting the movements in the response how much of the change is driven by the model equations (the AR(1) structure), and how much is driven by changing the parameter values. For researchers to be able to interpret the impulse responses from an economic perspective, movements in the central tendency for the impulse responses should be entirely driven by the model equations specified by the researcher.

The cause of this discrepancy is the implicit assumption that is made when constructing the posterior mean (or median) impulse responses. By treating each point independently, the researcher assumes that the vector of central estimates for the impulse responses can be obtained from $\mathbb{R}^n$, where n corresponds to the number of impulse responses computed. This allows for the selection of impulse responses that are not consistent with the model structure, as we do not require the central estimate of the impulse response at horizon h to be a function of the central estimates of the impulse responses at horizons $0, ..., h-1$. The question becomes whether, in more complex settings, pointwise approaches are ever able to construct a central estimate for the vector of impulse responses that can be obtained from a single parameterisation, thus preserving economic interpretability. Inoue and Kilian (2022) noted that this only occurs asymptotically in the situation where the number of parameters is greater than or equal to the number of impulse responses to be estimated. In practice, this is highly restrictive in many economic situations and this asymptotic equivalence almost never holds in a DSGE setting. The reason for this is explained below.

The vector of structural parameters is given by $\boldsymbol{\phi} \in \boldsymbol{\Phi}$, where $\boldsymbol{\Phi} \subset \mathbb{R}^{n_p}$, and n_p is the number of parameters. Similarly, the vector of impulse responses (or the Forecast Error Variance Decomposition) is given by $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, where $\boldsymbol{\Theta} \subset \mathbb{R}^{n_{obj}}$, and n_{obj} is the amount of impulse responses of interest to the researcher. In practice, $n_{obj} = n * h * s$, where n is the number of endogenous variables, h is the number of horizons to consider (including the impact horizon) and s is the number of shocks. The mapping between the structural parameters and the impulse responses is given by g(.). This mapping is surjective on the co-domain, $\mathbf{\Theta} = g(\mathbf{\Phi})$, but not on its complement. Intuitively, this means that the set $\mathbf{\Theta}$ contains only the impulse response vectors that are consistent with a single vector of parameters. As a result, the mapping does not necessarily span the full space $\mathbb{R}^{n_{obj}}$. Indeed, as Inoue and Kilian (2022) highlighted, when $n_p < n_{obj}$, the set of impulse responses, $\mathbf{\Theta}$, is of measure zero in the space $\mathbb{R}^{n_{obj}}$, as the mapping cannot span $\mathbb{R}^{n_{obj}}$. While this does occur in the limit when $n_p \geq n_{obj}$, this is almost never the case in DSGE models, so I ignore this scenario.

The reason for this is that DSGE models typically contain many endogenous variables, and the number of structural parameters is not forced to grow as quickly as in a VAR, which has a square structure. For example, adding an additional endogenous variable in a VAR requires the addition of $(2n - 1) * (p + 1)$ structural parameters (for a VAR without a constant term). This is not the case in DSGE models, where there are significantly fewer parameters. For example, in the Smets and Wouters (2007) DSGE model, there are 40 endogenous variables, and 41 parameters. Of these, 36 are estimated. Whether or not $\mathbf{\Theta}$ is able to span $\mathbb{R}^{n_{obj}}$ depends on the dimension of the parameter vector. However, given that in DSGE settings some of the parameters are calibrated, n_p corresponds to the number of estimated parameters, rather than the total number of parameters. To include as many endogenous variables in a VAR, even only accounting for one lag, would require over 3,000 structural parameters. The number of impulse responses to be estimated is $n_{obj} = n * h * s$. In the Smets and Wouters (2007) model, n_{obj} is always greater than n_p. Even if the researcher is only interested in the impact response to a single shock, $n_{obj} = 40 > n_p$. As a result, the central tendency obtained using pointwise methods is never equivalent to the Bayes Estimator, even as the number of posterior draws approaches infinity. More generally, the structure of DSGE models means that the space of impulse responses is almost always of measure zero in the space $\mathbb{R}^{n_{obj}}$.

Taking the posterior mean (or median) impulse responses makes the implicit assumption that $\mathbf{\Theta}$ spans the space $\mathbb{R}^{n_{obj}}$, so pointwise approaches also consider combinations of impulse responses that are impossible to obtain from a single structural parameterisation, as was the case in the AR(1) example described above. In a DSGE setting, the set of impulse responses that can be obtained from a single structural parameterisation are of measure zero in $\mathbb{R}^{n_{obj}}$, so it is never the case that the central estimates obtained using pointwise approaches are obtained from a single parameter vector. As a result, the ability to attach economic interpretation to the resulting impulse responses is restricted, as it is unclear how much of the movement in the IRFs is driven by the model equations, rather than by changing the parameter values. For example, to construct the full set of impulse responses in the Smets and Wouters (2007) framework using the posterior median impulse response, 1,064 different structural parametisations of the model are used, out of the 1,200 posterior impulse responses that are considered.

2.1.1 Pointwise Credible Sets

After computing the central tendency, the uncertainty surrounding the impulse responses is usually displayed by constructing a $(1-\alpha)\%$ credible region:

$$p(\boldsymbol{\theta_i} \in \boldsymbol{\Theta}_{1-\alpha,i}|y) = \int_{\boldsymbol{\Theta}_{1-\alpha,i}} p(\boldsymbol{\theta}_i \mid y) d\boldsymbol{\theta}_i = 1-\alpha \tag{2.3}$$

Where $\boldsymbol{\theta}_i$ is the i^{th} element of the impulse response vector and $\boldsymbol{\Theta}_{1-\alpha,i}$ is chosen such that $p(\boldsymbol{\theta}_i \mid y) \geq c_\alpha$. There are many credible sets that satisfy the above condition. The most popular approach is to construct the Highest Posterior Density Interval (HPDI) which involves selecting c_α to be as large as possible while still satisfying Equation 2.3. As with the central tendency, the key issue is that each element in the vector of impulse responses is considered individually. As a consequence, pointwise approaches to inference, such as the HPDI interval, often misrepresent the degree of estimation uncertainty (Sims and Zha, 1999). This is because treating each point independently does not guarantee that the impulse responses used to construct the HPDI at each point are drawn from the same set of parameter vectors. To construct a credible set that takes these dependence structures into account, the goal should be to obtain a credible set that contains $(1-\alpha)\%$ of the entire impulse response vectors.

2.2 Using a Point Estimate for the Parameters

An alternative approach that is commonly used in the literature is to estimate the parameters, and then collapse the posterior distributions for each parameter into a point estimate for use in constructing the objects of interest. Indeed, this is how Smets and Wouters (2007) construct the Forecast Error Variance Decomposition (FEVD), and it has also been used in popular applications such as Iacoviello and Neri (2010) and Christiano et al. (2014). The advantage of this approach is that the resulting objects of interest are derived from a single parameterisation, so maintain interpretability, unlike in the pointwise approaches described above.

However, the method is not without flaws. First, by collapsing the posterior distributions down to point estimates, researchers are unable to demonstrate the degree of estimation uncertainty. Second, the central moment of the posterior distribution is not invariant to nonlinear transformations, so does not guarantee that the impulse responses generated by using the central moment of the parameters are a good representation of the central moment of the impulse responses or forecast error variance decompostions themselves. The default for this approach in Dynare is to use the mean of the marginal posterior distribution for each parameter, while the mode of the joint posterior distribution is more frequently used to correspond with the Bayesian statistics literature, where it is commonly referred to as the Maximum a Posteriori (MAP). However, it is well known that both the mean and the MAP are not invariant to nonlinear transformations[4]. As the mapping between the structural parameters and the impulse responses is highly nonlinear, taking the mean (or mode) of each structural parameter will not, in general, generate the mean (or mode) of the posterior distribution for the impulse responses. As a result, there is no reason to believe *a priori* that such an approach will provide a good estimate of the central estimate for the ultimate object of interest.

3 Bayes Estimator and Joint Inference for DSGE models

Standard practice when plotting impulse response functions for DSGE models is to use the posterior mean (or median) impulse response for each variable, horizon, and shock, and conduct pointwise inference (e.g Smets and Wouters (2007)). This requires the researcher to implicitly assume that each response is obtained independently of previous responses, the responses of other variables to the same shock, and the response of the variable to other structural shocks.

[4]For the mean, this is easily demonstrated by invoking the same arguments used in Jensen's inequality. For the MAP, see, for example, Druilhet and Marin (2007).

Clearly this is an undesirable assumption to make in the vast majority of macroeconomic situations. A further consequence of this assumption is that the responses are likely to be obtained from different structural parametisations of the DSGE model, which also limits the economic interpretability of the results. This concern has been raised by Fry and Pagan (2011), among others, in a VAR context, but is equally problematic in a DSGE framework. In fact, as demonstrated in Section 2.1, pointwise IRFS in DSGE models are always derived from a combination of different structural models in any realistic application. The fact that pointwise approaches do not take into account the dependence between variables and time periods is one of the main criticisms of using such methods, as doing so leads us to misrepresent the estimation uncertainty associated with the object of interest (Sims and Zha, 1999).

There is a growing literature attempting to construct methods for estimation and inference that take this dependence into account for impulse responses estimated using VARs. Inoue and Kilian (2013) select the vector of impulse responses that maximises the joint density of the admissible structural VAR models as the central tendency, and then construct a joint credible set by ranking the admissible models based on the value of their joint density. The drawback of their algorithm is that it is possible only when there is a one-to-one mapping between the structural parameters and the impulse responses, which limits its use in empirical settings. To address this, Inoue and Kilian (2022) suggest a Bayes estimator, with inference obtained by constructing lowest posterior risk regions as advocated by Bernardo (2010) in a generic setting. Montiel Olea and Plagborg-Møller (2019) propose a sup-t confidence band that has exact finite-sample simultaneous credibility, but do not comment on how to construct an estimate for the central tendency. In a frequentist setting various methodologies have been considered to obtain inference that is able to capture the dependence across variables and time. For example, Lütkepohl et al. (2015) consider approaches that use the Bonferroni principle and the Wald statistic, Inoue and Kilian (2016) use the inverted joint Wald statistic, while Lütkepohl et al. (2018) and Bruder and Wolf (2018) use bootstrap methods.

However, these concerns have not yet been addressed in a DSGE setting, which is the main contribution of this paper. This paper provides an algorithm to construct a Bayes Estimator for the impulse responses and the forecast error variance decomposition in DSGE models, which takes these concerns into account by construction, ensuring that the central estimates are generated by the same structural parameterisation of the model and that the inference properly takes into account the dependence structures that are prevalent in macroeconomic data. Namely the response to a shock at horizon h will necessarily depend on the response to the same variable at horizon h-1. Likewise, the responses of the macroeconomic variables are also linked through the parametisation of the model, allowing the researcher to make statements about the co-movement of economic variables in response to particular shocks. The Bayes Estimator is selected by choosing the vector of impulse responses that minimise a given expected posterior loss function:

$$\tilde{\boldsymbol{\theta}} = \underset{\tilde{\boldsymbol{\theta}} \in \boldsymbol{\Theta}}{argmin} \ E_{\theta}\left[L(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}})\right] \tag{3.1}$$

Where $L(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}})$ represents a generic loss function and $\boldsymbol{\theta}$ the object of interest. In practice, there is usually not an analytical solution to Equation 3.1, so it must be solved numerically:

$$\tilde{\boldsymbol{\theta}}_M = \underset{\tilde{\boldsymbol{\theta}} \in \boldsymbol{\Theta}_M}{argmin} \ \frac{1}{M} \sum_{i=1}^{M} L(\boldsymbol{\theta}^{(i)}, \tilde{\boldsymbol{\theta}}) \tag{3.2}$$

Where M represents the total number of draws from the posterior distribution and $\boldsymbol{\theta}^{(i)}$ is the vector of impulse responses obtained using the i^{th} draw from the posterior distribution. The resulting Bayes Estimator is the vector of impulse responses that minimises the expected posterior loss, conditional on the impulse responses being generated by the same set of structural parameters, thus ensuring economic interpretability. The requirement that a researcher must select a loss function may appear to be a drawback of the approach, as the measure of central tendency will, in general, differ under different loss functions. However, pointwise approaches are also sensitive to the choice of loss function, so this should not be thought of as a significant drawback of using the Bayes Estimator. In a pointwise setting, the impulse response functions selected are the ones that minimise the following generic expected posterior loss:

$$E_{\theta_{PW}}[\mathcal{L}(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}})] = \int L_1\left(\theta_1, \tilde{\theta}_1\right) f(\theta_1|\boldsymbol{x})d\theta_1 + \cdots \int L_{n_{obj}}\left(\theta_{n_{obj}}, \tilde{\theta}_{n_{irf}}\right) f(\theta_{n_{obj}}|\boldsymbol{x})d\theta_{n_{obj}} \tag{3.3}$$

Where θ_i represents the i^{th} element of the vector of impulse responses, $f(\theta_i|\boldsymbol{x})$ denotes the posterior distribution of θ_i, and $\tilde{\boldsymbol{\theta}}$ is the vector of impulse responses selected that minimise the expected posterior loss. As a result, each of the individual components of the expected posterior loss in Equation 3.3 is minimised without taking into account the other impulse responses. If the researcher is interested in the posterior median impulse response, this is equivalent to minimising Equation 3.3 using the absolute loss function, $L_i(\theta_i, \tilde{\theta}_i) = |\theta_i - \tilde{\theta}_i|$, whereas for the posterior mean, the quadratic loss function, $L_i(\theta_i, \tilde{\theta}_i) = (\theta_i - \tilde{\theta}_i)^2$, would be used instead. To illustrate this point for the absolute loss, define the loss function over a single element of θ as:

$$L_i(\theta_i, \tilde{\theta}_i) = \begin{cases} (\theta_i - \tilde{\theta}_i) & \text{for } \theta_i \geq \tilde{\theta}_i \\ (\tilde{\theta}_i - \theta_i) & \text{for } \tilde{\theta}_i > \theta_i \end{cases} \tag{3.4}$$

The expected posterior loss at each point can be written as:

$$\int L_i\left(\theta_i, \bar{\theta}_i\right) f(\theta_i|\boldsymbol{x}) d\theta_i = \int_{-\infty}^{\tilde{\theta}_i} (\tilde{\theta}_i - \theta_i) f(\theta_i|\boldsymbol{x}) d\theta_i + \int_{\tilde{\theta}_i}^{\infty} (\theta_i - \tilde{\theta}_i) f(\theta_i|\boldsymbol{x}) d\theta_i \tag{3.5}$$

The Bayes Estimator is the $\tilde{\theta}_i$ that minimises the expected posterior loss. As the loss function considered is convex, the minimiser can be obtained by differentiating with respect to $\tilde{\theta}_i$ and setting the result equal to zero.

$$\int_{-\infty}^{\tilde{\theta}_i} f(\theta_i|\boldsymbol{x}) d\theta_i - \int_{\tilde{\theta}_i}^{\infty} f(\theta_i|\boldsymbol{x}) d\theta_i = 0 \tag{3.6}$$

At the loss-minimising point, $\int_{-\infty}^{\tilde{\theta}_i} f(\theta_i|\boldsymbol{x}) d\theta_i = \int_{\tilde{\theta}_i}^{\infty} f(\theta_i|\boldsymbol{x}) d\theta_i$, which implies:

$$2\int_{-\infty}^{\tilde{\theta}_i} f(\theta_i|\boldsymbol{x}) d\theta_i = \int_{-\infty}^{\infty} f(\theta_i|\boldsymbol{x}) d\theta_i = 1 \tag{3.7}$$

Therefore:

$$\int_{-\infty}^{\tilde{\theta}_i} f(\theta_i|\boldsymbol{x}) d\theta_i = \frac{1}{2} \tag{3.8}$$

As a result, the Bayes Estimator under the absolute loss function in a scalar setting is the posterior median. Equation 3.3 is then minimised by taking the posterior median of each impulse response separately. Therefore, the key difference between pointwise and joint approaches is not the inclusion of a loss function, but rather the dimension of the object being minimised. While pointwise approaches minimise over a scalar, the Bayes Estimator instead minimises the loss over the whole vector of impulse responses. More formally, the Bayes Estimator used in this paper chooses the vector of impulse responses that minimises $E_\theta \left[||\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}||_2 |\boldsymbol{x} \right]$, pointwise methods would instead choose $\tilde{\boldsymbol{\theta}}$ to minimise $E_{\theta_1} \left[|\theta_1 - \tilde{\theta}_1||\boldsymbol{x} \right] + ... + E_{\theta_{n_{obj}}} \left[|\theta_{n_{obj}} - \tilde{\theta}_{n_{obs}}||\boldsymbol{x} \right]$ if the researcher opts for the pointwise median impulse response function. Characterising the loss in vector form ensures that the collection of impulse responses that minimise the expected loss function of choice are compatible with the same structural model, and so retain the economic interpretability that is of paramount importance to most researchers utilising DSGE models. The following subsection outlines how to achieve this.

3.1 Central Tendency

The question of what forms a good measure of central tendency for a multi-dimensional object is still debated among researchers. One popular approach is the Spatial Median, as proposed by Haldane (1948). The Spatial Median is the multivariate extension of the univariate median, and is defined as the vector that minimises the sum of the Euclidean distances to all of the other points within the set of interest. For a given set of draws from the posterior distribution, $\boldsymbol{\theta}^{(1)}$... $\boldsymbol{\theta}^{(M)}$, where $\boldsymbol{\theta}^{(i)} \in \boldsymbol{\Theta_M}$, the spatial median is defined as the vector that solves:

$$\tilde{\boldsymbol{\theta}}_M^{sm} = \underset{\tilde{\boldsymbol{\theta}} \in \mathbb{R}^{n_{obj}}}{\text{argmin}} \sum_{i=1}^{M} ||\boldsymbol{\theta}^{(i)} - \tilde{\boldsymbol{\theta}}||_2 \tag{3.9}$$

The multivariate mean can be similarly obtained by selecting the vector that minimises the squared Euclidean distance. In practice, the Spatial Median is obtained using numerical methods [5]. While the Spatial Median is able to solve one of the issues with the pointwise posterior median, namely that the vector of medians is not a good measure of the central tendency of the multivariate object (Small, 1990), it does not solve the issue of model inconsistency. The

[5] The most popular method, Weiszfeld's algorithm, was introduced by Weiszfeld (1937). A translation of the original paper into English was provided by Weiszfeld and Plastria (2009). Cohen et al. (2016) recently proposed a more computationally efficient algorithm for approximating the Spatial Median.

Spatial Median is infeasible if we wish to maintain model consistency for our central estimate for similar reasons to those outlined in the argument against using the pointwise posterior median. To address this, I define the Bayes Estimator to minimise the sum of the Euclidean distances to the other posterior draws, constraining the minimiser to lie within $\boldsymbol{\Theta_M}$:

$$\tilde{\boldsymbol{\theta}}_M^{BE} = \underset{\tilde{\boldsymbol{\theta}} \in \boldsymbol{\Theta}_M}{\operatorname{argmin}} \sum_{i=1}^{M} ||\boldsymbol{\theta}^{(i)} - \tilde{\boldsymbol{\theta}}||_2 \tag{3.10}$$

The Bayes Estimator obtained in Equation 3.10 is consistent with the model equations and can be described as a constrained version of the Spatial Median. This ensures that the resulting vector of impulse responses is a good measure of the central tendency of the impulse response vector, while only considering the vectors that are possible to obtain from a single parameterisation of the model. The two coincide in the case where $\tilde{\boldsymbol{\theta}}_M^{sm} \in \Theta_M$. When $\tilde{\boldsymbol{\theta}}_M^{sm} \notin \Theta_M$, simulation evidence suggests that the Bayes estimator is the vector in Θ_M that minimises the Euclidean distance to the unconstrained spatial median.

As a consequence, the method of obtaining the central tendency proposed in this paper can be considered similar in spirit to the suggestion in Fry and Pagan (2011), which was to select the structural model closest to the median. However, Fry and Pagan (2011) suggested selecting the set of feasible IRFs that minimise the squared Euclidean distance to the the vector of pointwise medians, which is not, in general, representative of the central tendency for the entire vector. One of the key drawbacks of using an additively separable loss function, such as the ones proposed in Inoue and Kilian (2022), is that the Bayes Estimator asymptotically coincides with the vector of pointwise medians when $n_{obj} = n_p$, as demonstrated in Appendix A. The approach in this paper avoids such criticisms, as the Spatial Median is recognised as a multivariate measure of the central tendency.

A key consideration that needs to be addressed to obtain an accurate measure of the central tendency is the scaling of the elements contained within $\boldsymbol{\theta}$. Solving the above equation using the standard impulse response vector can leave the researcher vulnerable to over-weighting impulse responses for variables measured in smaller scales, or where there is more estimation uncertainty present, as the contribution to the overall loss will be greater. This is also a concern if the researcher wants to use the same parameterisation for all shocks, as the standard deviation of each shock will necessarily influence the magnitude of the impulse responses and therefore their contribution to the loss function. As a result, not accounting for these differences would result in the estimate of the spatial median being skewed toward the objects of interest corresponding with larger estimated shocks, smaller scales, or a larger degree of estimation uncertainty. Admittedly, in a log-linearised DSGE framework, the scales are percentage-point deviations from the steady state, so concerns around the scale of the variable are usually not a major concern, but this may be more problematic in a VAR. Transforming the data beforehand allows us to alleviate some of these concerns.

$$\nu_{n,h,s}^{(i)}(\theta^{(i)}) = \frac{\theta_{n,h,s}^{(i)}}{\sigma_{n,h,s}} \tag{3.11}$$

Where σ denotes the standard deviation and n corresponds to the variable, h the horizon, and s the shock. The transformation ensures that the elements are measured in the same scale, standard deviation units, so the relative weights will be more equal across variables, horizons, and shocks. The resulting minimisation problem then becomes:

$$\tilde{\boldsymbol{\theta}}_M^{BE_2} = \underset{\tilde{\boldsymbol{\theta}} \in \boldsymbol{\Theta}_M}{\text{argmin}} \frac{1}{M} \sum_{i=1}^{M} ||\boldsymbol{\nu}^{(i)}(\boldsymbol{\theta}^{(i)}) - \tilde{\boldsymbol{\nu}}(\tilde{\boldsymbol{\theta}})||_2 \tag{3.12}$$

Applying this adjustment helps to overcome the criticisms leveled at the standard loss functions by Inoue and Kilian (2022), and offers a more computationally efficient alternative to the scale-invariant Dirac-Delta loss proposed in their paper. However, it is unable to control for another important consideration for researchers, which is the horizon of interest for the object of interest. This particular point is discussed in more detail in Section 5.

3.2 Joint Credible Sets

There is a growing literature on how to construct joint inference in a time series setting. Most approaches rely on the asymptotic normality of the central estimate and construct the credible set or confidence bands around this point. The general form of the Bayes Estimator falls under the umbrella of M-Estimators, which have been shown to be asymptotically normal under certain conditions (Huber, 1964; Arcones, 1998). These conditions typically include convexity of the set of structural impulse responses. While this holds for the spatial median if we allow the vector of impulse responses to be selected from $\mathbb{R}^{n_{obj}}$ (Möttönen, Nordhausen, and Oja, 2010), it does not hold when we restrict the choice set to Θ, as is the case when we wish to preserve economic interpretability. As a result, many of the approaches to joint inference currently available in the literature cannot be applied in this setting. One approach that does not require normality is the joint credible set proposed by Bernardo (2010) in a general setting, and by Inoue and Kilian (2022) in a time series setting. The joint credible set for the vector of impulse responses can be constructed as the lowest posterior risk region. The $(1-\alpha)100\%$ joint credible set can be defined as:

$$\boldsymbol{\Theta}_{1-\alpha,L} = \left\{ \tilde{\boldsymbol{\theta}} \in \Theta : E_\theta(L(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}})) \leq c_{1-\alpha,L} \right\} \tag{3.13}$$

Where $E_\theta(L(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}))$ represents a generic loss function to be selected by the researcher. The critical value, $c_{1-\alpha,L}$, is set to be the smallest value that is able to achieve a posterior probability for $\boldsymbol{\Theta}_{1-\alpha,L}$ of $1-\alpha$. In this paper, the expected posterior loss is given by $E_\theta(L(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}})) = E_\theta(||\boldsymbol{\nu}(\boldsymbol{\theta}) - \tilde{\boldsymbol{\nu}}(\tilde{\boldsymbol{\theta}})||_2|\boldsymbol{x})$. In finite samples, the set can be visually approximated by sorting the losses given by the numerical approximation of the Bayes Estimator from smallest to largest, and selecting the sample critical value such that $(1-\alpha)100\%$ of the draws are contained.

$$\boldsymbol{\Theta}_{1-\alpha,L,M} = \left\{ \tilde{\boldsymbol{\theta}} \in \boldsymbol{\Theta}_M : \frac{1}{M}\sum_{i=1}^{M} ||\boldsymbol{\nu}^{(i)}(\boldsymbol{\theta}^{(i)}) - \tilde{\boldsymbol{\nu}}(\tilde{\boldsymbol{\theta}})||_2 \leq c_{1-\alpha,L,M} \right\} \tag{3.14}$$

As the number of posterior draws approaches infinity, the numerical approximation of the posterior distribution of the IRFs, $\boldsymbol{\Theta}_M$, approaches the true, continuous, posterior distribution $\boldsymbol{\Theta}$. In addition, the numerical approximation of the expected posterior loss approaches the expected posterior loss:

$$\sup_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \left| \frac{1}{M}\sum_{i=1}^{M} ||\boldsymbol{\nu}^{(i)}(\boldsymbol{\theta}^{(i)}) - \tilde{\boldsymbol{\nu}}(\tilde{\boldsymbol{\theta}})||_2 - E_\theta(L(\boldsymbol{\theta}, \tilde{\boldsymbol{\theta}}) \right| = o_{as}(1) \tag{3.15}$$

In addition, as each draw from the posterior distribution, $\boldsymbol{\theta}^{(i)}$ is iid, $c_{1-\alpha,L,M} - c_{1-\alpha,L} = o_{as}(1)$ as $M \rightarrow \infty$. As a result, Equation 3.14 can be viewed as a numeric approximation of the true Joint Credible Set given by Equation 3.13. One concern with this visual approximation is that the posterior probability of the set given by Equation 3.14 is zero. This is because the set contains finitely many elements, while the joint posterior distribution is continuous. However, given that the numerical approximation of the expected posterior loss will converge to its true value as the number of posterior draws approaches infinity, the visual approximation of the joint credible set can be used to assess the degree of estimation uncertainty. An advantage of using this approach is that the joint credible set takes the form of a shot-gun plot, which allows researchers to see the individual paths of the impulse responses under different parameterisations.

4 Empirical Application

To illustrate the advantages of the approaches proposed in this paper compared to the standard approaches, I use the Smets and Wouters (2007) framework[6]. First, I construct the Bayes Estimator as the measure of central tendency using the loss function described in Equation 3.12. Ranking the expected posterior losses from smallest to largest allows for the construction of a visual approximation for the joint credible set. While the construction of the central estimate requires the same parameterisation to be used for all endogenous variables, only a subset of the full set responses, corresponding to the observable variables, are displayed in the paper.

[6]For a full description of the model and the estimation procedures, refer to Smets and Wouters (2007).

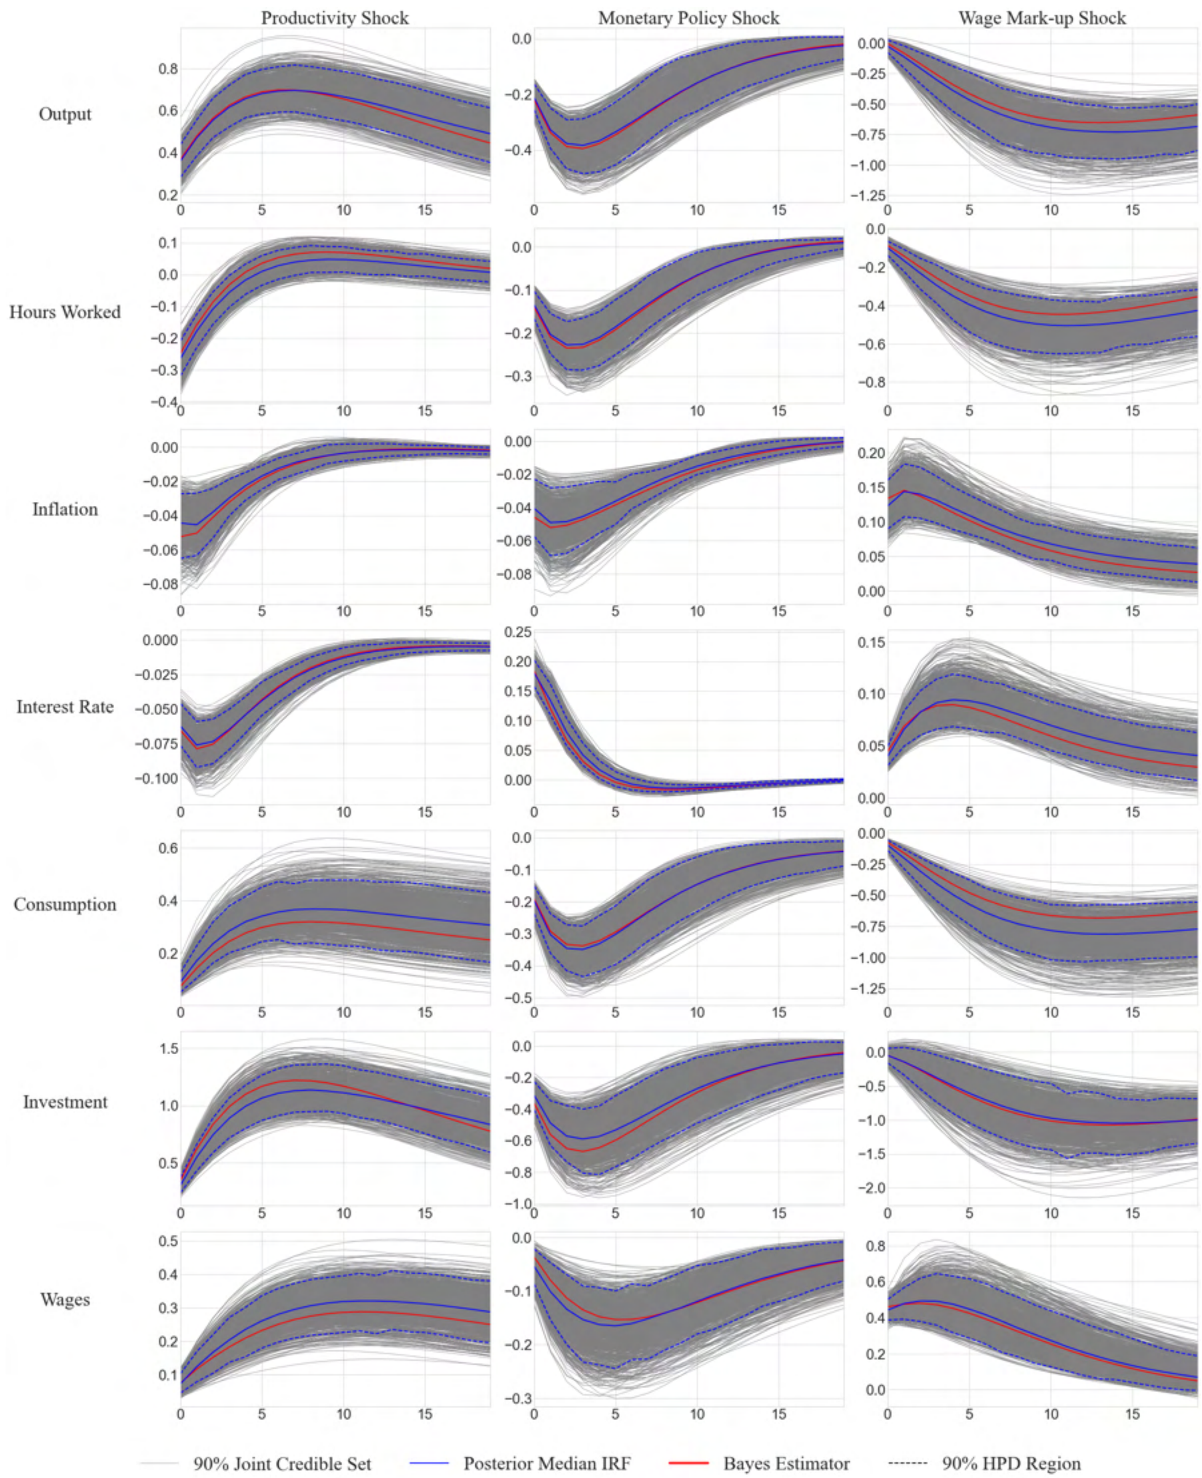

Figure 1: Bayes Estimator and Joint Credible Set vs. Posterior Median and HPD Region

Figure 1 plots the impulse responses for the observable variables in the Smets and Wouters (2007) model following three different shocks. The red lines correspond to the Bayes Estimators for the impulse responses obtained by minimising the Euclidean distance, while the solid blue lines are the posterior median impulse responses and the dashed blue lines represent the 90% pointwise credible regions. The blue lines are obtained directly from the Smets and Wouters (2007) replication file. The different approaches are shown to produce different measures of the central tendency, although in most cases the quantitative differences between the two sets of impulse responses are relatively small. The Bayes Estimator implies a smaller rise in consumption following a productivity shock, and a smaller decline in consumption following a wage mark-up shock compared to its pointwise counterpart.

The main difference between the two approaches is the corresponding inference. The pointwise HPDI, like the central estimate, also treats each point of the impulse response vector independently. As a result, the degree of estimation uncertainty implied by the HPD region understates the true estimation uncertainty that is present as it does not take into account the dependence across variables and horizons, reflecting the criticism of such approaches in Sims and Zha (1999). This finding is also consistent with the literature on joint inference in VARs, which found that pointwise approaches to inference often understate the degree of estimation uncertainty (Lütkepohl et al., 2015). The 90% HPD region underrepresents the degree of estimation uncertainty compared to the joint credible set. By construction, 90% of the impulse responses considered are contained within the joint credible set. In the Smets and Wouters (2007) example considered in this paper, only 1 out of the 1,200 posterior impulse response vectors considered is fully contained within the pointwise credible region for all variables and horizons: the Bayes Estimator. Allowing for a small tolerance to account for numerical precision[7] results in just 1.6% of the 1,200 impulse response vectors being completely contained within the pointwise credible region for all variables, demonstrating that conventional approaches severely understate the degree of estimation uncertainty.

[7] The check allows for a tolerance of e^{-5} to account for rounding errors and machine precision.

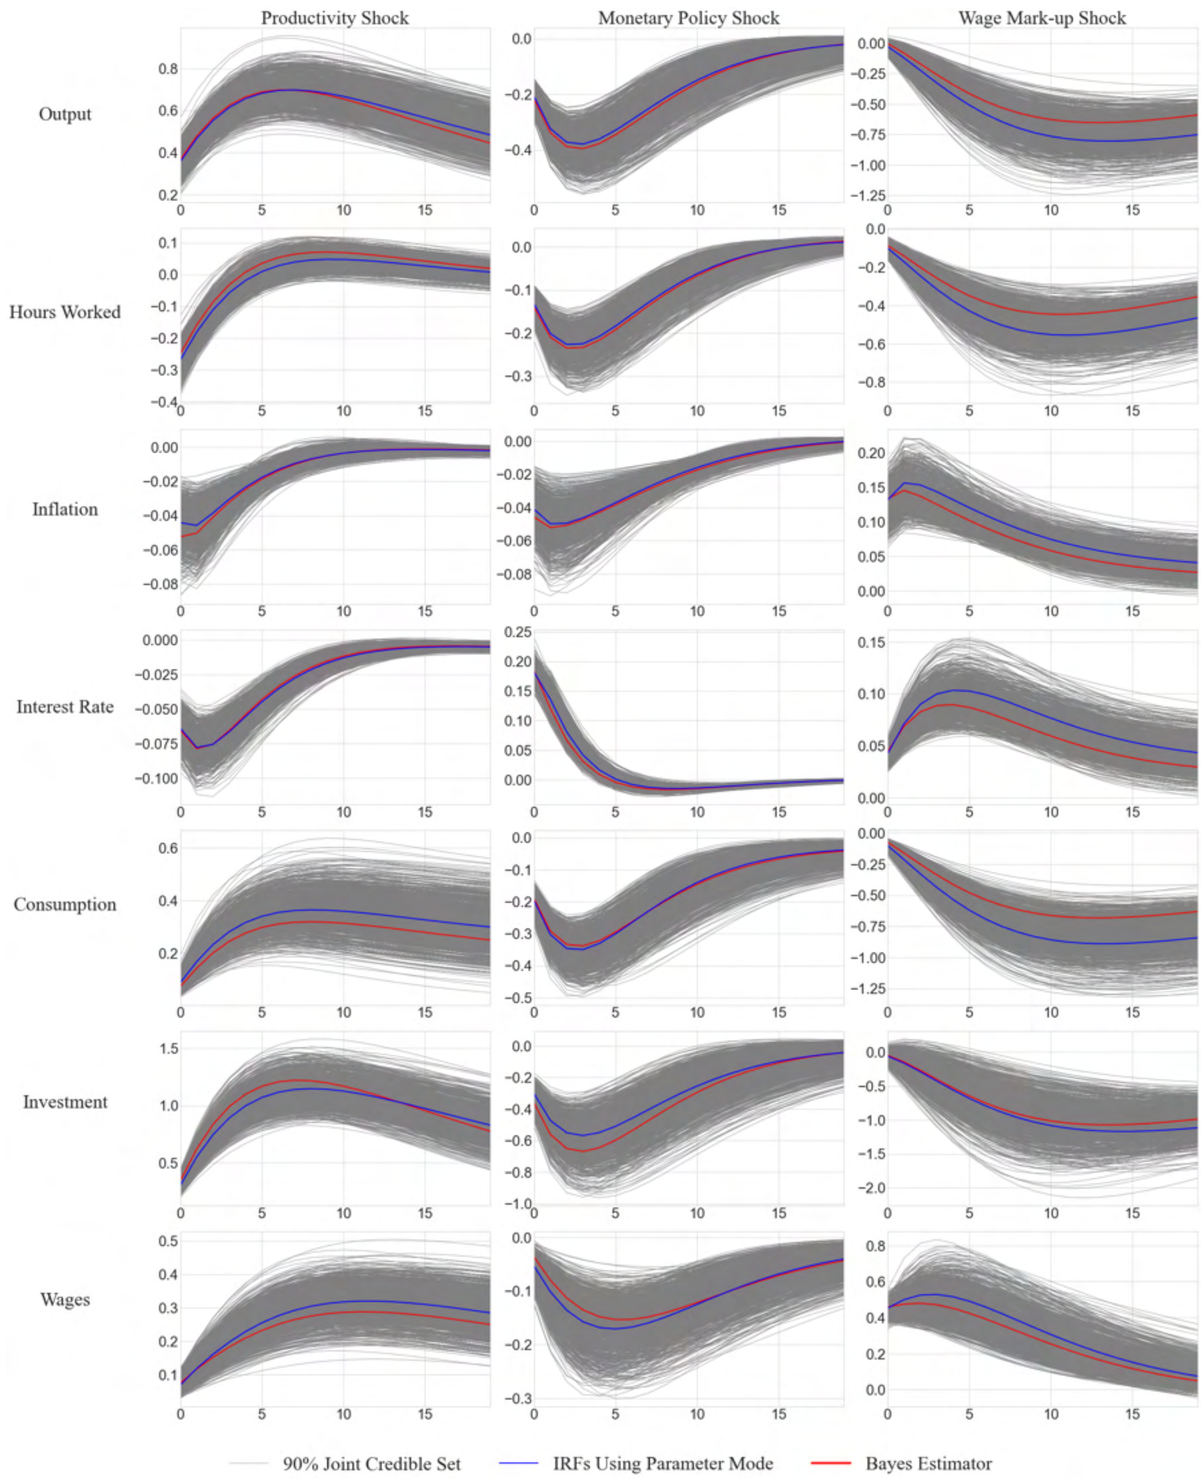

Figure 2: IRFs: Mode of Each Parameter vs. Bayes estimator

Figure 3 demonstrates the differences between the Bayes Estimator and the impulse responses generated by fixing the parameters at their posterior mode in the Smets and Wouters (2007) application. Again, although in general the differences between the Bayes Estimator and those obtained using the parameter modes are small, there are noticeable differences in the response of consumption, investment, and wages to a productivity shock, and the responses of several variables to a wage mark-up shock. However, while both of these estimates are drawn from a single structural model, so are able to maintain economic interpretability, only the approach using the Bayes Estimator is able to provide both the reassurance that the impulse responses generated will provide a good central estimate of the joint distribution and also allow researchers to characterise the estimation uncertainty that is present.

There are also implications for the Forecast Error Variance Decomposition (FEVD), which is another important object of interest for researchers when estimating DSGE models, particularly for business cycle analysis. The Bayes Estimator for the FEVD can be constructed in the same way as for the impulse responses, changing $\boldsymbol{\theta}$ to represent the vector of FEVDs that are compatible with the same structural parameterisation. Fixing the parameter values at their mode is the most popular method for computing the FEVD for DSGEs. This is because pointwise approaches do not guarantee that the FEVD sums to one, as different parameterisations are combined. As the FEVD is a non-linear function of the impulse responses, it is even less likely that the FEVD obtained from the parameter modes will provide a good measure of the central tendency for the FEVD. The Bayes Estimator is also able to contribute here, and produces noticeable differences compared to the FEVDs obtained using the parameter modes. In particular, the Bayes Estimator implies that productivity and monetary policy shocks are more important for explaining the variation in most of the observed variables. However, as demonstrated by the joint credible set, there is a large degree of estimation uncertainty surrounding these estimates, with a productivity shock explaining somewhere between 10-45% of the variation in output at a business cycle frequency, and similar degrees of uncertainty in many other cases. This highlights the advantage of using the approach outlined in this paper for the FEVD, as the standard approach of fixing the parameters at their posterior modes is unable to capture the uncertainty present in the estimation of the FEVD.

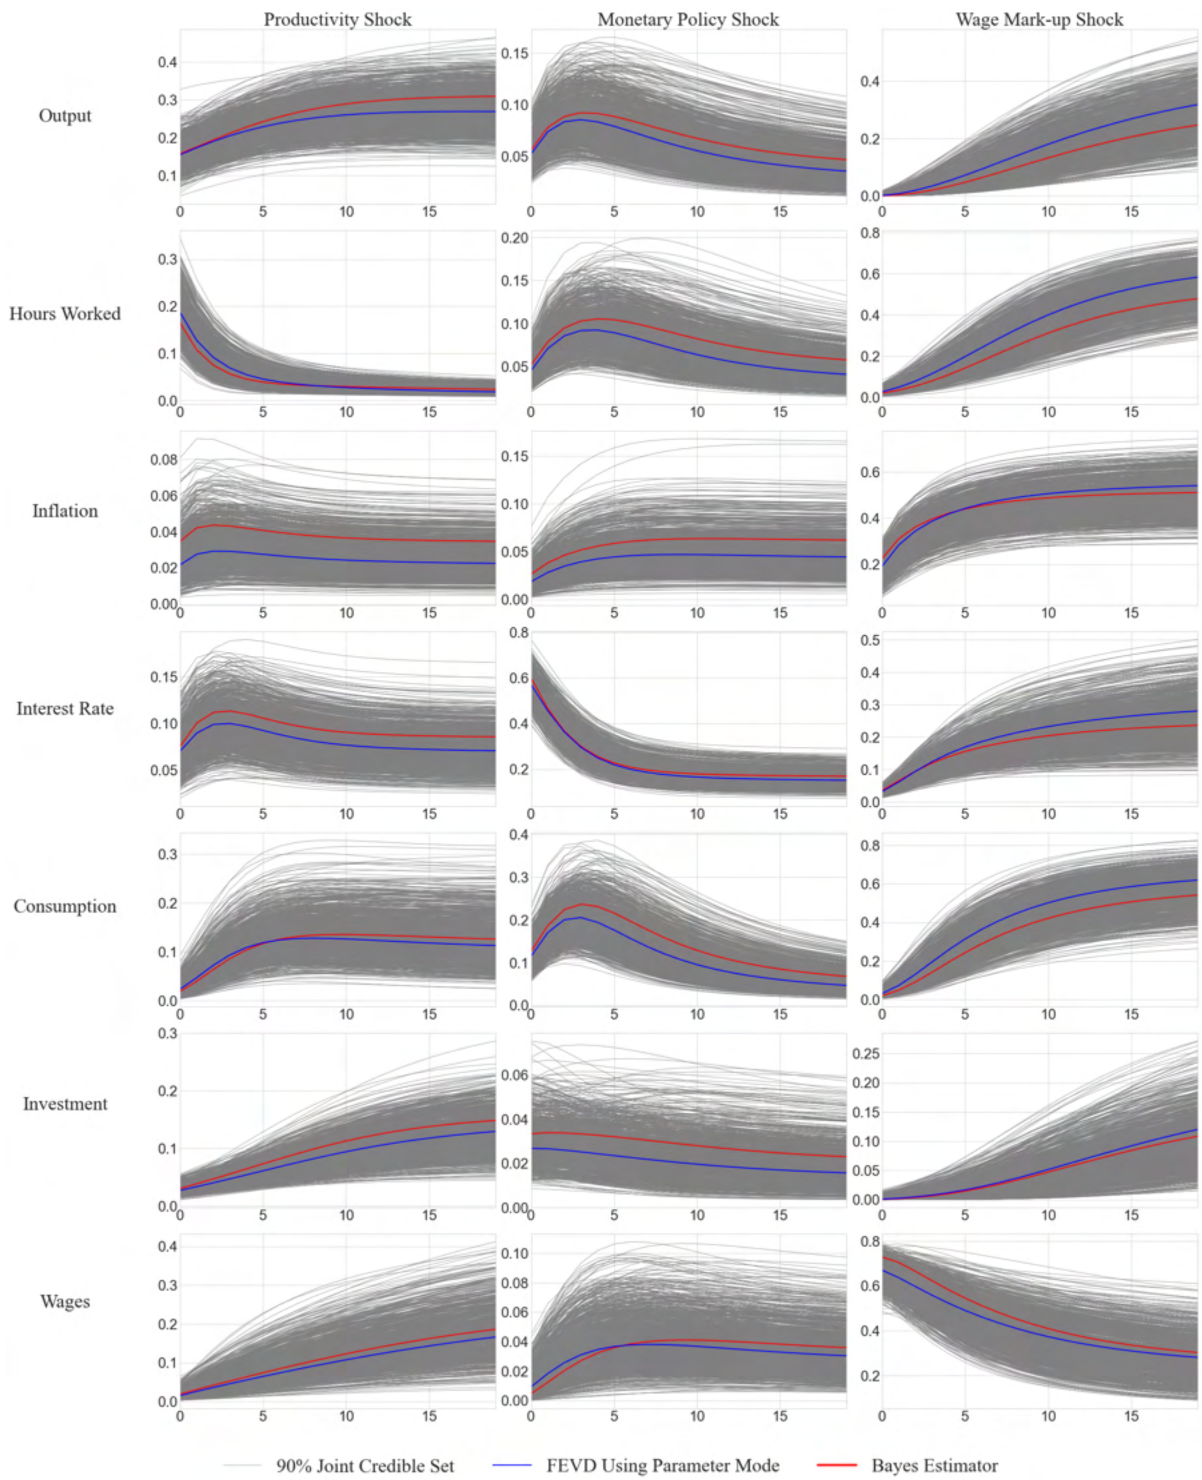

Figure 3: FEVD: Mode of Each Parameter vs. Bayes estimator

5 Non-invariance to Horizon length

The previous sections outlined the advantages of taking the entire vector of impulse responses into account when constructing the central tendency and inference. However, this more holistic approach is not without a cost. By definition, the structural model that is selected using the Bayes Estimator will depend on how many horizons the researcher is interested in computing impulse responses for, as the expected posterior loss will change when more horizons are added to the vector of impulse responses. While this is not a concern when using standard approaches such as the pointwise approaches described in section 2.1, it should not be thought of as an advantage of using such methods, but rather an artifact of the unrealistic assumptions (independence of responses across variables and time) that we make when doing so. That said, the lack of invariance may still be cause for concern if the findings are particularly sensitive to the choice of horizon. While the focus of this paper is on DSGE models, the point raised in this section is equally applicable to any application of joint inference in a time series setting, including VARs.

This lack of invariance to the horizon of interest requires researchers using joint inference approaches to make a conscious decision about the appropriate number of horizons to consider. In DSGE settings with quarterly data, researchers are usually most interested in the first 20 horizons, corresponding to business cycle frequency. An alternative approach is to reduce the sensitivity of the results to changes in the horizon of interest by discounting horizons that are further out. This would help to ensure that the estimate obtained is an accurate measure of the central tendency for horizons that are most relevant for policymakers. For example, one could consider an adjustment to Equation 3.11 of the form:

$$\nu^{(i)}_{n,h,s}(\theta^{(i)}) = \beta^{\lambda} \frac{\theta^{(i)}_{n,h,s}}{\sigma_{n,h,s}} \tag{5.1}$$

With $\beta \in [0, 1]$ and $\lambda = h - min(h, \delta)$, where δ denotes the period from which discounting begins. This approach allows some flexibility to ensure that sufficient weight is placed on the policy relevant horizons, but also nests the cases of no discounting ($\beta = 1$) and full discounting beyond horizon δ ($\beta = 0$ and $\delta > 0$). The selection of β and δ then become tuning parameters for the analysis that can be tailored to the specific application. As a rule of thumb, they should be selected to ensure that relevant policymaking horizons carry a non-trivial weight. Setting $\beta = 0.9$ and $\delta = 20$ would imply no discounting for the first 20 periods, and then progressively discounting the subsequent horizons, resulting in a weight of 0.12 in period 40. Repeating the analysis in Section 2.1 with a maximum horizon of 40 and the above values for β and δ, we obtain the results in Figure 4.

The central estimate is shown to vary depending on the horizon of interest. The red solid line shows the measure of the central tendency corresponding to the adjusted Bayes Estimator using Equation 5.1 with $\beta = 0.9$ and $\delta = 20$, and the joint credible set is also constructed under this adjustment. The solid black line shows the central tendency if only the first 20 periods were considered when selecting the parameterisation ($\beta = 0, \delta = 20$), i.e. the black line uses the same structural parameters as those used for the central tendency in Figure 1. Finally, the solid Blue line applies no discounting ($\beta = 1$). The figure shows that changing the horizon of interest from 20 to 40 causes the central tendency to change. However, in this case, progressively discounting beyond horizon 20 ensures that the selected structural parameterization remains the same. While researchers could choose to assign zero weight beyond horizon 20, assigning some weight ensures that more distant horizons still feature in the expected posterior loss, and so guards against selecting a parameterisation that diverges at more distant horizons. This is perhaps not a major concern in most DSGE settings, where the responses converge to zero as the variable moves back to its steady state, but may be more of a concern in VARs, where this is not always the case.

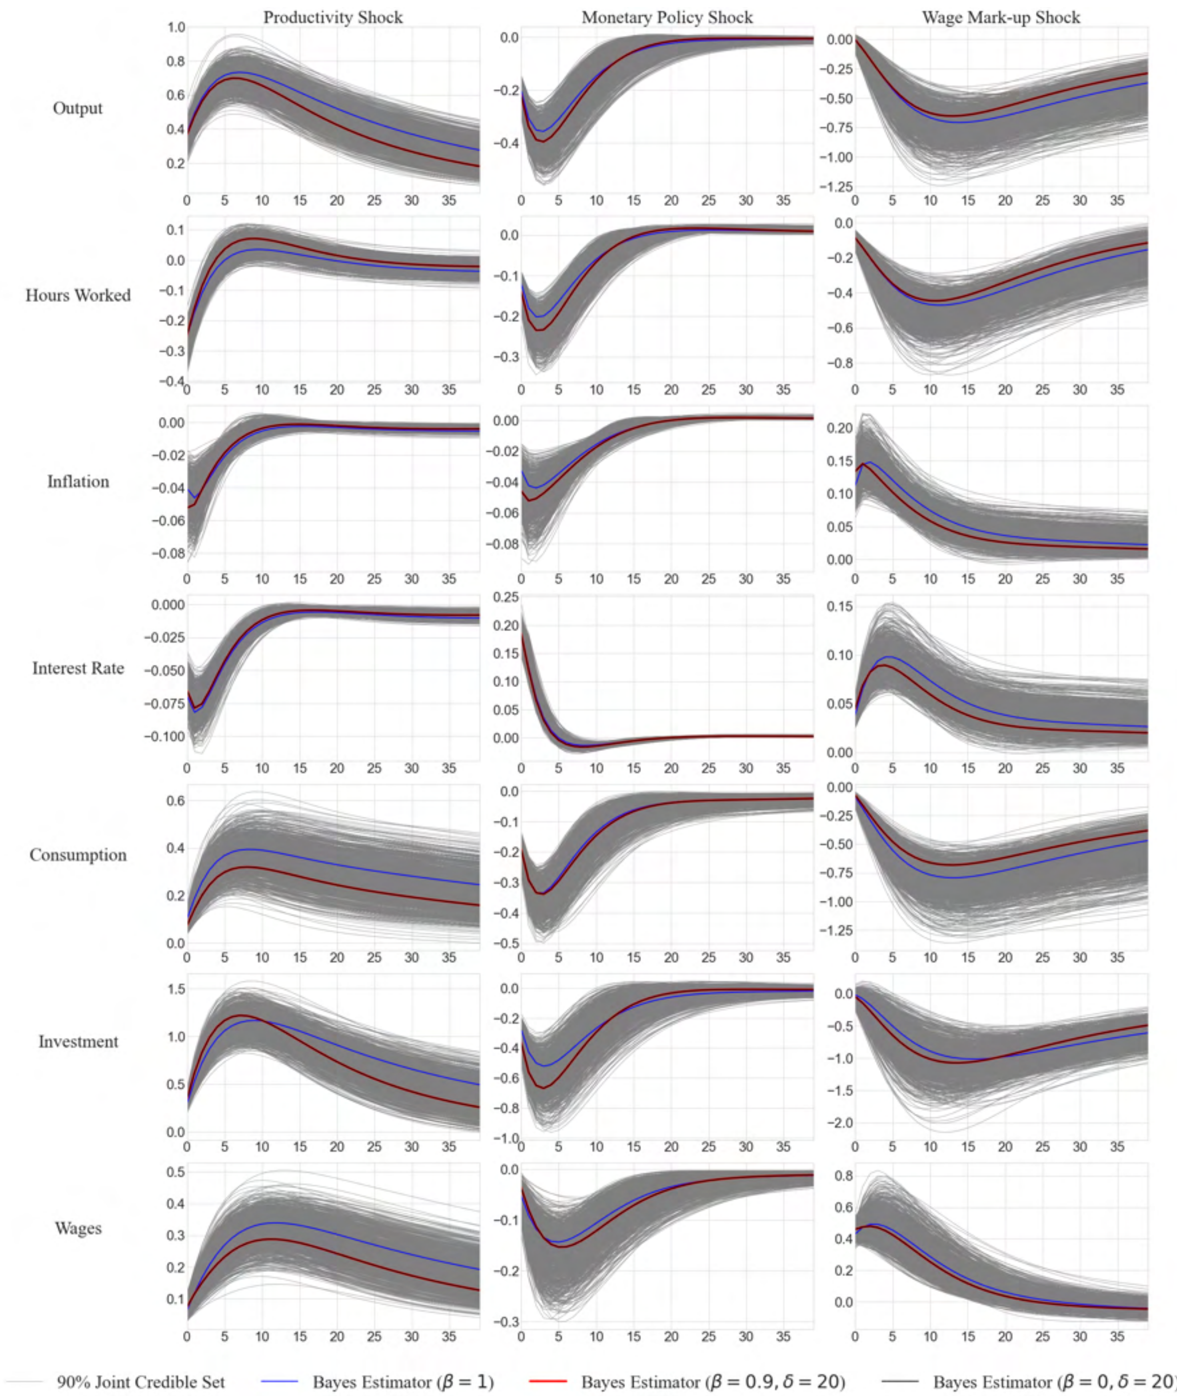

Figure 4: Bayes Estimator Under Different Values of β

6 Conclusion

This paper provides a toolkit to conduct joint inference for DSGE models estimated using Bayesian methods. Using the Bayes Estimator has two key advantages over the current approaches used in the literature. First, it provides an accurate measure of the central tendency for the object of interest that can be obtained from a single structural model. Under the loss function considered in this paper, the central estimate coincides with a constrained version of the spatial median, where only economically interpretable impulse responses or forecast error variance decompositions are considered. Second, it allows researchers to characterise the uncertainty that is present in the estimation of impulse responses. Current approaches either ignore the estimation uncertainty (using the mean/mode of the posterior distribution of the parameters) or understate it due to the assumption of independence across horizons and variables. In the empirical example considered in the paper, the Bayes Estimator is the only vector of impulse responses, out of the 1,200 considered, that is fully contained within the standard pointwise 90% HPD region for all variables, horizons and shocks. The paper also suggests adjustments to the loss function that are shown to mitigate some of the concerns that practicioners may have when using the Bayes Estimator.

References

Arcones, M. A. (1998): "Asymptotic Theory For M-Estimators Over A Convex Kernel," *Econometric Theory*, 14, 387–422.

Bernardo, J. M. (2010): "Integrated Objective Bayesian Estimation and Hypothesis Testing," *Bayesian Statistics 9*, 1–25.

Bratsiotis, G. J. and K. D. Pathirage (2024): "Monetary and Macroprudential Policy and Welfare in an Estimated Four-Agent New Keynesian Model," *Journal of Money, Credit and Banking*, n/a.

Bruder, S. and M. Wolf (2018): "Balanced Bootstrap Joint Confidence Bands for Structural Impulse Response Functions," *Journal of Time Series Analysis*, 39, 641–664.

Chernozhukov, V. and H. Hong (2003): "An MCMC approach to classical estimation," *Journal of Econometrics*, 115, 293–346.

Christiano, L. J., R. Motto, and M. Rostagno (2014): "Risk Shocks," *American Economic Review*, 104, 27–65.

Cohen, M. B., Y. T. Lee, G. Miller, J. Pachocki, and A. Sidford (2016): "Geometric median in nearly linear time," *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, 9–21.

Del Negro, M., F. Schorfheide, F. Smets, and R. Wouters (2007): "On the fit of new Keynesian models," *Journal of Business and Economic Statistics*, 25, 123–143.

Druilhet, P. and J.-M. Marin (2007): "Invariant HPD credible sets and MAP estimators," .

Fernández-Villaverde, J. (2009): "The Econometrics of DSGE Models," *NBER Working Paper 14677*.

Fry, R. and A. Pagan (2011): "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 49, 938–60.

Haldane, J. B. S. (1948): "Note on the median of a multivariate distribution," *Biometrika*, 35, 414–417.

Huber, P. J. (1964): "Robust Estimation of a Location Parameter," *Annals of Mathematical Statistics*, 35, 73–101.

Iacoviello, M. and S. Neri (2010): "Housing Market Spillovers: Evidence from an Estimated DSGE Model," *American Economic Journal: Macroeconomics*, 2, 125–64.

Inoue, A. and L. Kilian (2013): "Inference on impulse response functions in structural VAR models," *Journal of Econometrics*, 177, 1–13.

——— (2016): "Joint confidence sets for structural impulse responses," *Journal of Econometrics*, 192, 421–432, innovations in Multiple Time Series Analysis.

——— (2022): "Joint Bayesian inference about impulse responses in VAR models," *Journal of Econometrics*, 231, 457–476, special Issue: The Econometrics of Macroeconomic and Financial Data.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010): "Investment shocks and business cycles," *Journal of Monetary Economics*, 57, 132–145.

Lütkepohl, H., A. Staszewska-Bystrova, and P. Winker (2015): "Confidence Bands for Impulse Responses: Bonferroni vs. Wald," *Oxford Bulletin of Economics and Statistics*, 77, 800–821.

——— (2018): "Calculating joint confidence bands for impulse response functions using highest density regions," *Empirical Economics*, 55, 1389–1411.

Montiel Olea, J. L. and M. Plagborg-Møller (2019): "Simultaneous confidence bands: Theory, implementation, and an application to SVARs," *Journal of Applied Econometrics*, 34, 1–17.

Mukoyama, T., M. Shintani, and K. Teramoto (2021): "Cyclical Part-Time Employment in an Estimated New Keynesian Model with Search Frictions," *Journal of Money, Credit and Banking*, 53, 1929–1968.

Möttönen, J., K. Nordhausen, and H. Oja (2010): "Asymptotic theory of the spatial median," *Institute of Mathematical Statistics Collections*, 182–193.

Sims, C. A. and T. Zha (1999): "Error Bands for Impulse Responses," *Econometrica*, 67, 1113–1155.

Small, C. G. (1990): "A Survey of Multidimensional Medians," *International Statistical Review / Revue Internationale de Statistique*, 58, 263–277.

Smets, F. and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586–606.

Weiszfeld, E. (1937): "Sur le point pour lequel la Somme des distances de ¡i¿n¡/i¿ points donneacute;s est minimum," *Tohoku Mathematical Journal, First Series*, 43, 355–386.

Weiszfeld, E. and F. Plastria (2009): "On the point for which the sum of the distances to n given points is minimum," *Annals of Operations Research*, 167, 7–41.

A Appendix A - Alternative Loss Functions

This section compares the loss function of choice in this paper with other alternative loss functions, such as those proposed by Inoue and Kilian (2022). The main focus is on whether these loss functions are able to provide a good measure of the central tendency for the set of posterior draws. To make the arguments more tractable, I focus on the case where $n_{obj} = n_p$. In this case, as the number of posterior draws approaches infinity, the mapping $\boldsymbol{\theta} = f(\boldsymbol{\phi})$ spans $\mathbb{R}^{n_{obj}}$. As a result, the optimisation problem becomes unconstrained, which simplifies the arguments.

A.1 Additively Separable Loss Functions

The simplest loss functions to implement are additively separable. These loss functions can be generically defined as those taking on the following form:

$$\tilde{\boldsymbol{\theta}}^{AS} = \underset{\tilde{\boldsymbol{\theta}} \in \boldsymbol{\Theta}}{\operatorname{argmin}} \sum_{j=1}^{n_{obj}} E\left[L\left(\theta_j, \tilde{\theta}_j\right)\right] \tag{A.1}$$

This nests two of the loss functions proposed in Inoue and Kilian (2022), the absolute and quadratic losses. The sample analogues of which are defined as follows:

$$\tilde{\boldsymbol{\theta}}^{AL}_M = \underset{\tilde{\boldsymbol{\theta}} \in \boldsymbol{\Theta}_M}{\operatorname{argmin}} \frac{1}{M} \sum_{i=1}^{M} \sum_{j=1}^{n_{obj}} \left|\theta_j^{(i)} - \tilde{\theta}_j\right|, \tag{A.2}$$

$$\tilde{\boldsymbol{\theta}}^{QL}_M = \underset{\tilde{\boldsymbol{\theta}} \in \boldsymbol{\Theta}_M}{\operatorname{argmin}} \frac{1}{M} \sum_{i=1}^{M} \sum_{j=1}^{n_{obj}} \left(\theta_j^{(i)} - \tilde{\theta}_j\right)^2, \tag{A.3}$$

In the setting when $M \to \infty$ and $n_p = n_{obj}$, $\widehat{\Theta}_M \to \mathbb{R}^{n_{obj}}$. Focusing on the absolute loss function, Equation A.2 becomes:

$$\tilde{\boldsymbol{\theta}}^{AL} = \underset{\tilde{\boldsymbol{\theta}} \in \mathbb{R}^{n_{obj}}}{\operatorname{argmin}} \sum_{j=1}^{n_{obj}} E\left[\left|\theta_j - \tilde{\theta}_j\right|\right] \tag{A.4}$$

In this setting, it is clear to see that the minimiser of Equation A.4 is the vector of point-wise medians. To demonstrate this, note that Equation A.4 is identical to Equation 3.3, the minimisation problem to obtain the pointwise posterior medians. It is well known that the vector of pointwise medians is, in general, not a good measure of the central tendency in a multivariate setting (Small, 1990). Similar arguments can be applied to any additively separable loss function, which raises concerns about their usefulness for the task at hand; finding a good measure of the central tendency for the joint distribution.