Frequentist Inference in
Weakly Identified DSGE Models

Pablo Guerron-Quintana   Atsushi Inoue   Lutz Kilian

Federal Reserve Bank   North Carolina   University of Michigan
of Philadelphia   State University   CEPR

October 23, 2009

Abstract

We show that in weakly identified models (1) the posterior mode will not be a consistent estimator of the true parameter vector, (2) the posterior distribution will not be Gaussian even asymptotically, and (3) Bayesian credible sets and frequentist confidence sets will not coincide asymptotically. This means that Bayesian DSGE estimation should not be interpreted merely as a convenient device for obtaining asymptotically valid point estimates and confidence sets from the posterior distribution. As an alternative, we develop new frequentist confidence sets for structural DSGE model parameters that remain asymptotically valid regardless of the strength of the identification.

JEL Classification Codes: C32, C52, E30, E50.

KEYWORDS: DSGE models; Bayesian estimation; Identification; Inference; Confidence sets; Bayes factor; Likelihood ratio.

1 Introduction

In recent years, there has been growing interest in the estimation of dynamic stochastic general equilibrium (DSGE) models by Bayesian methods. One of the chief advantages of the Bayesian approach compared to the frequentist approach is that the use of prior information allows the researcher to estimate structural models that otherwise would be computationally intractable or would produce economically implausible estimates. This feature has made these methods popular even among researchers who think of these methods merely as a convenient device for obtaining model estimates but would not consider themselves Bayesians otherwise.

At the same time, there is growing evidence that many DSGE models used in empirical macroeconomics are only weakly identified (see, e.g., Canova and Sala 2009). Weak identification manifests itself in a likelihood that is nearly flat across the parameter space. For example, Del Negro and Schorfheide (2008) document that DSGE models that have very different policy implications may fit the data equally well. In particular, a New Keynesian model with moderate price rigidities and low wage rigidities is observationally equivalent to a model with high wage rigidities and high price rigidities. As a result, the posterior of the structural parameters of the model becomes highly dependent on the priors used by the researcher. This is a common problem. For example, Smets and Wouters (2007a, p. 594) note that for their main behavioral parameters "the mean of the posterior distribution is typically relatively close to the mean of the prior assumptions." While this fact does not necessarily pose a problem for genuine Bayesians, it is especially troublesome for frequentist users of these methods because it suggests that we learn nothing from the data.

In this paper, we make two contributions. First, we show that in weakly identified models the usual asymptotic equivalence between Bayesian and frequentist estimation and inference breaks down.[1] The problem is that under weak identification the likelihood

---

[1]See Le Cam and Yang (2000, chapter 8) and the references therein for the large sample correspondence between Bayesian and frequentist approaches. For more recent results in the econometrics literature, see Andrews (1994), Chernozhukov and Hong (2003) and Hahn (1997), for example, and Kim (1998) and Phillips and Ploberger (1996) for the nonstationary case in particular.

no longer asymptotically dominates the posterior, which helps explain the sensitivity of Bayesian DSGE estimates to the prior in practice. As a result, one cannot interpret posterior modes (or means or medians) as frequentist point estimates or treat Bayesian credible sets effectively as frequentist confidence sets. In particular, it is not possible to construct confidence intervals from the quantiles of the posterior distribution or by adding multiples of posterior standard deviations to the posterior mean. Specifically, we show that (1) the posterior mode will not be a consistent estimator of the true parameter vector, (2) the posterior distribution will not be Gaussian even asymptotically, and (3) Bayesian credible sets and frequentist confidence sets will not coincide asymptotically. This means that Bayesian DSGE estimation should not be interpreted merely as a convenient device for obtaining asymptotically valid point estimates and confidence sets from the posterior distribution.

Second, as an alternative, we develop a new class of frequentist confidence sets for the structural parameters of DSGE models that remain valid asymptotically regardless of the strength of identification. In general, the strength of identification is a matter of degree and there is no well-defined threshold that separates strongly identified from unidentified models (see, e.g., Canova and Sala 2009, Iskrev 2008). There is little hope of constructing pre-tests for strong identification nor is it clear that pre-testing would be an appropriate strategy in this context. Our approach is instead based on the premise that the structural parameters of the model are weakly identified in the sense that the component of the likelihood function that depends on the structural parameter vector is local to zero. As in the weak instruments literature, we think of this assumption as a device that reflects our inability to determine the strength of the identification from the data. The proposed confidence set is obtained by inverting either the likelihood ratio (LR) statistic or the Bayes factor (BF). The resulting LR confidence set can be constructed using classical estimation methods and has $(1 - \alpha)\%$ coverage probability asymptotically, regardless of the strength of identification. The implied BF confidence set is conservative in that a $(1 - \alpha)\%$ confidence set has at least a $(1 - \alpha)\%$ coverage probability asymptotically. The BF confidence set correctly reflects the uncertainty about the structural parameters even when the likelihood is flat, it protects the researcher from spurious inference, and it is

2

invariant to the prior asymptotically in the case of weak identification. Since the Bayes factor is the ratio of the posterior odds to the prior odds, if the likelihood is flat and hence the prior dominates the posterior, the numerator and the denominator of the ratio will tend to cancel, making the proposed BF confidence set more robust to alternative priors than conventional intervals based on the posterior.

The advantage of the BF interval is that it may be computed easily from existing Bayesian code for DSGE model estimation. No additional numerical estimation is required. This feature facilitates its adoption by the many applied users of Bayesian methods of estimating DSGE models. In contrast, construction of the LR interval dispenses with Bayesian estimation methods altogether. It requires instead the explicit derivation of the state space representation of the unrestricted reduced form of the DSGE model on a case by case basis and numerical estimation of its parameters. The latter task can be computationally challenging in practice, especially when there are many more unobservable state variables than observables. Given the predominance of Bayesian estimation methods in empirical macroeconomics, we focus on applications of the BF interval in this paper. Although our theoretical results show that the LR interval is expected to be tighter asymptotically than the BF interval, we find that in practice even the computationally more straightforward BF interval can help us discriminate between alternative hypotheses of economic interest. We defer to future research a systematic comparison of the coverage accuracy and average length of the two intervals in finite samples.

The work most closely related to our analysis is Moon and Schorfheide's (2009) comparison of frequentist and Bayesian inference in partially identified models. In such models the structural parameter vector of interest can be bounded, but the set of admissible parameter values cannot be narrowed down to a point. Thus, the best a researcher can hope for is to identify the set of parameter values that is consistent with the data. Moon and Schorfheide establish that in partially identified models, Bayesian credible sets tend to be smaller than frequentist confidence sets. This finding is in contrast with the conventional point identified case, in which Bayesian and frequentist sets coincide asymptotically, enabling users to reinterpret Bayesian credible sets as frequentist confidence sets.

Like Moon and Schorfheide we find that Bayesian credible sets and conventional fre-

3

quentist confidence sets need not coincide asymptotically. In particular, the usual Bayesian credible set does not have the correct asymptotic coverage probability in weakly identified models, preventing its interpretation as a frequentist confidence set. Our analysis differs from Moon and Schorfheide's work, first, in that we focus on weakly point identified parameters rather than set identified parameters. The second difference is that we do not stop at documenting these differences but propose an alternative confidence set that remains valid regardless of the strength of the identification. Finally, our study deals with identification in structural macroeconomic models, whereas theirs focuses on microeconomic applications.

Related work also includes Komunjer and Ng (2009), Rubio-Ramirez, Waggoner and Zha (2005, 2009), and Fukač, Waggoner and Zha (2007). Komunjer and Ng (2009) establish conditions for identifying structural parameters in DSGE models from autocovariance structures. Rubio-Ramirez et al. (2005, 2009) develop conditions for identification in structural vector autoregressive models. Fukač et al. (2007) contrast local and global identification. While these procedures are helpful in assessing the identifiability of structural model parameters, they are not informative about the strength of identification, suggesting that there remains a need for approaches such as ours that are robust to weak identification.

The remainder of the paper is organized as follows. In section 2 we investigate the asymptotic behavior of the posterior distribution in weakly identified models. We establish the failure of the conventional frequentist interpretation of Bayesian posterior estimates. We propose the LR and BF confidence sets, establish their asymptotic validity from a frequentist point of view, and discuss their relative merits. In section 3 we investigate the finite-sample performance of traditional pseudo-Bayesian methods by simulation. We focus on a commonly used New Keynesian model consisting of a Phillips curve, an investment-savings equation, and a Taylor rule. We demonstrate that the practice of constructing confidence intervals from the posterior of the structural parameters by adding $+/-1.645$ posterior standard deviations to the posterior mode (or mean) results in intervals with serious coverage deficiencies. In some cases, coverage rates of nominal 90% intervals for commonly used sample sizes may drop as low as 39%. In contrast, the conservative

BF interval proposed in this paper in the simulation has more accurate coverage for all parameters and sample sizes. In section 4, we investigate an empirical example based on a larger scale DSGE model widely used in the DSGE literature (see, e.g., Del Negro and Schorfheide 2008). We focus on the question of the relative importance of wage and price rigidities in the US economy. We also illustrate the robustness of the proposed confidence sets to alternative choices of priors. The concluding remarks are in section 5.

2   Asymptotic Theory

2.1   Asymptotic Behavior of the Posterior Distribution When Parameters Are Weakly Identified

When parameters are strongly identified, the posterior distribution is degenerate about the true parameter value and asymptotically normal after suitable scaling. The latter result is called the Bernstein-von Mises Theorem in the Bayesian literature. We will restate a version of the Bernstein-von Mises Theorem for the multiparameter non-iid case for expository purposes.[2]

Proposition 1 (Bernstein-von Mises Theorem): Denote the log-likelihood function by $\ell_T(\theta) = \ln L_T(X_1, ..., \theta) = \ln f(X_1, X_2, ..., X_T | \theta)$. Suppose that the following conditions hold:

(a) $\theta_0 \in \text{int}(\Theta) \subset \Re^k$.

(b) The prior density $\pi(\theta)$ is continuous on Θ and $\pi(\theta_0) > 0$.

(c) The likelihood function $L_T(\theta)$ is twice continuously differentiable in a neighborhood of θ_0.

(d) For all $\delta > 0$ there exists $\varepsilon(\delta) > 0$ such that

$$\lim_{T \to \infty} P_{\theta_0} \left[\sup_{\theta \in \Theta \cap B_\delta(\theta_0)} [\ell_T(\theta) - \ell_T(\theta_0)] \leq -\varepsilon(\delta)T \right] = 1, \tag{1}$$

where $B_\delta(\theta_0) \equiv \{\theta \in \Theta : \|\theta - \theta_0\| \leq \delta\}$.

---

[2]There are stronger versions of this result. See Bickel and Doksum (2006) and Le Cam and Yang (2000) for more detailed treatments and different versions of this theorem.

(e) There is a matrix valued function $H(\theta)$ such that $\lim_{T\to\infty} \sup_{\theta\in\Theta} |-\frac{1}{T}\nabla_{\theta\theta}\ell_T(\theta) - H(\theta)| \overset{p}{\to} 0$ and $H(\theta_0)$ is positive definite where $\nabla_{\theta\theta}\ell_T(\theta)$ is the Hessian of the log-likelihood function $\ell_T(\theta)$.

(f) The maximum likelihood estimator (MLE) $\hat{\theta}_T$ of θ_0 is strongly consistent, i.e., $\hat{\theta}_T \to \theta_0$ almost surely.

Then for any compact set A

$$\int_{B_T} P(\theta|X_1, X_2, ..., X_T)d\theta \overset{P_{\theta_0}}{\to} P(z \in A). \tag{2}$$

where $B_T = \{\theta \in \Theta : [\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)]^{1/2}(\theta - \hat{\theta}_T) \subset A\}$ and $z \sim N(0_{k\times 1}, I_k)$.

The Bernstein-von Mises Theorem allows a classical interpretation of Bayesian confidence sets. In other words, Bayesian credible sets for $\hat{\theta}_T$ can be viewed as valid classical confidence sets for θ_0 asymptotically. This fact is important because it allows econometricians who are not Bayesians to use the Bayesian apparatus to estimate DSGE models, taking advantage of its superior convergence properties, while interpreting the results in a classical fashion. However, recent research has shown that DSGE models are often only weakly identified (see, e.g., Del Negro and Schorfheide 2008, Canova and Sala 2009). In that case, the Bernstein-von Mises Theorem does not apply because assumptions (d), (e) and (f) will fail when parameters are not strongly identified. The next result shows formally that the classical interpretation of Bayesian credible sets breaks down when the model is not strongly identified.

Proposition 2 (Posterior Distributions of Exponential Families Under Weak Identification). Consider an exponential family:

$$L_T(x|\theta) = [\Pi_{t=1}^T h(x_t)] \exp\left[\sum_{j=1}^k \eta_j(\theta) \sum_{t=1}^T T_j(x_t) - TB(\theta)\right] \tag{3}$$

where η_j and B are real-valued functions of θ and T_j and h are real-valued functions of the data (see Bickel and Doksum, 2007). Suppose that

(a)

$$\eta_j(\theta) \;=\; \frac{1}{T}q_j(\theta) + o\left(\frac{1}{T}\right) \tag{4}$$

$$B(\theta) \;=\; \frac{1}{T}r(\theta) + o\left(\frac{1}{T}\right). \tag{5}$$

(b) The likelihood function (3) is correctly specified.

(c) $(1/T)\sum_{t=1}^{T} T_j(x_t) \to E(T_j(x_t))$ almost surely for $j = 1, 2, ..., k$.

Then when a conjugate prior is used, the posterior density almost surely converges to

$$\frac{[\exp(\sum_{j=1}^{k} q_j(\theta)E(T_j(x)) - r(\theta))]}{\int_{\Theta} E[\exp(\sum_{j=1}^{k} q_j(\theta)T_j(x) - r(\theta))]d\theta} \tag{6}$$

When a more general, not necessarily conjugate, prior $\pi(\theta)$ is used, the posterior density almost surely converges to

$$\frac{\pi(\theta)[\exp(\sum_{j=1}^{k} q_j(\theta)E(T_j(x)) - r(\theta))]}{\int_{\Theta} \pi(\theta)[\exp(\sum_{j=1}^{k} q_j(\theta)E(T_j(x)) - r(\theta))]d\theta} \tag{7}$$

It can be shown that – even if there is a unique $\theta_0 \in \Theta$ that maximizes $\sum_{j=1}^{k} q_j(\theta)E[T_j(x_i)] - r(\theta)$ – the maximum likelihood estimator is inconsistent and has a nonstandard limiting distribution. Specifically, Proposition 2 shows that (i) the posterior distribution is not degenerate around the true parameter value when the parameter is weakly identified; (ii) that it is not Gaussian; and (iii) that the limit of the posterior distribution depends on the prior. In other words, the effect of the prior on the posterior will not die out asymptotically, invalidating the usual classical interpretation of Bayesian credible sets. This result is intuitive because information does not accumulate even when the sample size grows when parameters are weakly identified. This means first that the posterior mode no longer coincides with the mean or median. Second, this means that, when the econometrician follows the standard procedure for strongly identified DSGE models and computes the mean (or median or mode) of the posterior distribution as the best guess for the parameter value, the resulting estimator will be inconsistent for the true parameter value.

Condition (a) is an extension of Stock and Wright's (2000) concept of weak identification in GMM to exponential families. It is useful to contrast our notion of weak identification in condition (a) of Proposition (2) to the limiting cases of strong identification

and no identification. Strong identification would require that the terms $\eta_j(\theta)$ and $B(\theta)$ in the likelihood function take on values that allow us to solve uniquely for the maximum likelihood estimator. In contrast, lack of identification would correspond to $\eta_j(\theta) = 0$ and $B(\theta) = 0$ such that the likelihood function does not depend on θ. The intermediate case embodied in assumption (a) is that the component of the likelihood function that depends on θ is local to zero. This assumption is designed to represent our inability to determine which of the two limiting cases is a better approximation of reality.[3]

2.2 Likelihood Ratio Tests

As a practical alternative, we propose two frequentist confidence sets for parameters in DSGE models that are valid regardless of the strength of identification. One is based on the likelihood ratio test statistic and the other is based on the Bayes factor.

Suppose that the reduced-form parameters Π are functions of the structural parameters of interest θ:

$$\Pi = g(\theta) \tag{8}$$

where $g : \Theta \to \Re^{dim(\Pi)}$. In DSGE models, Π is the vector of parameters of the state-space model,

$$x_{t+1} = Ax_t + Bw_t, \tag{9}$$

$$y_t = Cx_t + Dw_t, \tag{10}$$

where x_t is a vector of possibly unobserved state variables, y_t is a vector of observed variables, $w_t \overset{iid}{\sim} N(0, I)$. While C is a matrix of zeros and ones, the reduced-form parameters A, B, and D are typically functions of structural parameters θ (see Fernández-Villaverde, Rubio-Ramírez, Sargent, and Watson 2007).

Consider the likelihood ratio statistic for testing

$$H_0 : \Pi = g(\theta) \text{ for some } \theta \in \Theta.$$

---

[3]Our assumptions include as a special case the possibility of no identification, as discussed in Kadane (1975) and Poirier (1998).

and denote it by $LR_T(\theta) = 2(\ell_T(\hat{\Pi}_T) - \ell_T(g(\theta)))$ where $\ell_T(\Pi)$ is the log-likelihood function for Π and $\hat{\Pi}$ is the unconstrained maximum likelihood estimate of Π.

Proposition 3 (Asymptotic Distribution of the LR Test Statistic Under Weak Identification)

(a) $\Theta = A \times B$ is non-empty and compact in $\Re^k$ where $A \subset \Re^{k_1}$, $B \subset \Re^{k_2}$ and $k_1 + k_2 = k$, and θ_0 is in the interior of Θ.

(b) The log-likelihood function $\ell_T(\Pi)$ is correctly specified and twice continuously differentiable in θ.

(c) $g_1 : B \to \Re^{dim(\Pi)}$ and $g_2 : \Theta \to \Re^{dim(\Pi)}$ are continuously differentiable and $g_T(\theta_0) \equiv g_1(\beta_0) + T^{-1/2}g_2(\theta_0) = \Pi_{0,T}$ for all T.

(d) There is a maximum likelihood estimator $\hat{\Pi}_T$ such that $\sqrt{T}(\hat{\Pi}_T - \Pi_{0,T}) \xrightarrow{d} V_\Pi^{1/2}z$ where $V_\Pi = -\text{plim}_{T\to\infty}[(1/T)\nabla_{\Pi\Pi}\ell_T(\Pi_0)]^{-1}$ is positive definite and z is a $dim(\Pi)$-dimensional standard normal random vector.

If $\theta = \theta_0$, then

$$LR_T(\theta_0) \xrightarrow{d} \chi^2_{\dim(\Pi)}. \tag{11}$$

Proposition 3 follows from standard arguments in the weak instrument literature. It shows that the LR test statistic can be used to construct confidence intervals with $(1-\alpha)\%$ coverage probability asymptotically, regardless of the strength of identification. Although this result is not suprising from a technical point of view, it provides a powerful tool for dealing with problems of weak identification of structural parameters in DSGE models. It shows that inference on these parameters may be conducted without ever estimating the structural model. As a result, we can dispense with Bayesian methods of estimating the structural parameters altogether. The construction of the LR confidence set requires instead the explicit derivation of the state space representation of the unrestricted reduced form of the DSGE model on a case by case basis and numerical estimation of its parameters. The latter task can be computationally challenging in practice, however, especially when there are many more unobservable state variables than observables. Given the predominance of Bayesian estimation methods in empirical macroeconomics, below we consider an alternative approach based on the inversion of the Bayes factor.

2.3 Bayes Factors

Consider testing

$$H_0 : \theta \in B_{\delta_T}(\theta_0)$$

against

$$H_1 : \theta \notin B_{\delta_T}(\theta_0)$$

where $B_{\delta_T}(\theta_0) = \{\theta \in \Theta : |\theta - \theta_0| \leq \delta_{T,j} \text{ for } j = 1, 2, ..., p\}$, $\Theta \subset \Re^p$ and $\delta_T = [\delta_{T,1}, ..., \delta_{T,p}]' \to 0_{p \times 1}$ as $T \to \infty$.

We define the *Bayes factor* in favor of H_1 by

$$\text{Bayes Factor}(\theta_0) = \frac{\pi(H_0)p(H_1|X)}{\pi(H_1)p(H_0|X)} \tag{12}$$

where $\pi(H_i)$ and $p(H_i|X)$ are the prior and posterior probabilities of H_i, respectively.

Theorem 1 (Asymptotic Distribution of the Bayes Factor Under Weak Identification)

(a) $\Theta = A \times B$ is non-empty and compact in $\Re^k$ where $A \subset \Re^{k_1}$, $B \subset \Re^{k_2}$ and $k_1 + k_2 = k$, and θ_0 is in the interior of Θ.

(b) $\pi : \Theta \to \Re_+$ is continuous on Θ.

(c) The log-likelihood function $\ell_T(\Pi)$ is correctly specified and twice continuously differentiable in θ.

(d) $g_1 : B \to \Re^{dim(\Pi)}$ and $g_2 : \Theta \to \Re^{dim(\Pi)}$ are continuously differentiable and $g_T(\theta_0) \equiv g_1(\beta_0) + T^{-1/2}g_2(\theta_0) = \Pi_{0,T}$ for all T.

(e) There is a maximum likelihood estimator $\hat{\Pi}_T$ such that $\sqrt{T}(\hat{\Pi}_T - \Pi_{0,T}) \xrightarrow{d} V_\Pi^{1/2}z$ where $V_\Pi = -\text{plim}_{T \to \infty}[(1/T)\nabla_{\Pi\Pi}\ell_T(\Pi_0)]^{-1}$ is positive definite and z is a $dim(\Pi)$-dimensional standard normal random vector.

(f) $\delta_T = [\delta_{T,1}, ..., \delta_{T,p}]'$ satisfies the following condition: If $|\theta_j - \theta_{0,j}| \leq \delta_{T,j}$ for $i = 1, ..., p$ then $Dg(\theta)(\theta - \theta_0) = o(T^{-1/2})$ where $Dg(\theta)$ is the $dim(\Pi) \times k$ Jacobian matrix of $g(\theta)$.

If $\theta = \theta_0$, then

$$\lim_{T \to \infty} P\left(\text{Bayes Factor}(\theta_0) \leq e^{\frac{z'z}{2}}\right) = 1 \tag{13}$$

10

where z is defined in Assumption (e), i.e., $2ln$Bayes Factor(θ_0) is asymptotically bounded by a chi-square random variable with $dim(\Pi)$ degrees of freedom.

Remarks.

1. Extending Stock and Wright's (2000) concept of weak identification in GMM to our context, we model g_T in such a way that the part of g_T that depends on weakly identified parameters vanishes asymptotically. As a result, the rank of the Jacobian of the function g_T can be less than k in the limit. Assumption (d) allows for the case in which the parameters are all weakly identified ($\theta = \alpha$, $g_T(\theta) = T^{-1/2}g_2(\alpha)$, $k = k_1$), the case in which they are partially identified in the sense of Choi and Phillips (1992) ($0 < k_2 < k$ and $g_2(\theta) \equiv 0$ for all θ), and the case in which they are all strongly identified ($\theta = \beta$, $g_T(\theta) = g_2(\beta)$, $k = k_2$). Therefore, (13) holds true regardless of the strength of the identification.

2. Assumption (e) requires only the existence of an asymptotically normally distributed maximum likelihood estimator of the reduced-form parameters. We do not need to compute the maximum likelihood estimator of Π to obtain the Bayes factor.

3. Assumption (e) implies strong identification of Π and can be equivalently written as $\sqrt{T}(\hat{\Pi}_T - g_1(\beta)) \xrightarrow{d} g_2(\theta) + V_\pi^{1/2}z$.

4. Because Θ is compact and $g_T(\cdot)$ is continuously differentiable, Assumption (f) is always satisfied if $\delta_T = o(T^{-1/2})$.

5. As an example in which assumptions (d) and (e) are satisfied, consider the wage Phillips curve in Del Negro and Schorfheide (2008):

$$
\begin{aligned}
\widetilde{w}_t &= \zeta_w \beta \mathbb{E}_t \left[\widetilde{w}_{t+1} + \Delta w_{t+1} + \Upsilon_{1,t} \right] + \\
&\quad \frac{1 - \zeta_w}{1 + v_l \left(1 + \lambda_w \right)/\lambda_w} \left(v_l L_t - w_t + \Upsilon_{2,t} \right).
\end{aligned}
$$

Here, $\widetilde{w}_t$ is the optimal real wage relative to the real wage for aggregate services, w, ζ_w is the wage rigidity, v_l is the inverse of the Frisch labor supply elasticity, and $1 + 1/\lambda_w$ is the demand elasticity for labor services (for further discussion see Section 4 and Del Negro and Schorfheide, 2008). The objects $\Upsilon_{i,t}$ relate to terms irrelevant to our discussion. It can be shown that λ_w only enters the systems of equations

11

that define the equilibrium in this DSGE model via the last equation (Del Negro and Schorfheide, 2008). As the inverse labor supply elasticity becomes negligible ($v_l \rightarrow 0$), the parameter λ_w becomes weakly identified in the sense of the zero-information limit condition of Nelson and Startz (2007), whereas the reduced-form parameter $\frac{1-\zeta_w}{1+v_l(1+\lambda_w)/\lambda_w}$ remains strongly identified. Therefore, assumptions (d) and (e) are satisfied.

6. Theorem 1 implies that one can obtain level $(1 - \alpha)$ confidence sets by inverting the Bayes factor:

$$\{\theta \in \Theta : \text{Bayes Factor}(\theta_0) \leq e^{\frac{\chi^2_{\dim(\Pi)}{}^{(1-\alpha)}}{2}}\} \tag{14}$$

which satisfies

$$\lim_{T \rightarrow \infty} P(\Theta_0) \geq 1 - \alpha. \tag{15}$$

7. Note that the fact that we focus on the Bayes factor in favor of the alternative hypothesis (as opposed to the Bayes factor in favor of the null hypothesis) is not innocuous. If we reverse the numerator and denominator in equation (8), under strong identification, an additional $\log(T)$ term will emerge in equation (8) and make it impossible to derive the asymptotic bounds on the distribution of the Bayes factor.

2.4   LR Interval versus BF Interval

It is useful to compare the theoretical properties of the LR and BF intervals. The LR confidence set has $(1-\alpha)\%$ coverage probability asymptotically, regardless of the strength of identification. The BF interval is more conservative in that a $(1-\alpha)\%$ confidence set has at least a $(1 - \alpha)\%$ coverage probability asymptotically. This difference arises because, unlike the LR statistic, $2ln(\text{Bayes Factor}(\theta_0))$ in equation (14) is merely bounded by a random variable with a $\chi^2_{(\dim(\Pi))}$ distribution. Asymptotically, we would expect the LR interval to be tighter because the LR test can be shown to be more powerful asymptotically than the BF approach. Consider the power of the proposed Bayes factor. Suppose that all structural parameters are weakly identified, i.e., $\theta = \alpha$, and that the true parameter value α_1 is different from the hypothetical parameter value α_0.

Theorem 2 (Power of the Bayes Factor Under Weak Identification) Suppose that the As-

sumptions (a)–(f) of Theorem 1 hold with $k = k_2$ and $g_T(\theta) = T^{-1/2}g_2(\alpha)$. Then for $\alpha = \alpha_1$,

$$\text{Bayes Factor}(\alpha_0) \xrightarrow{d} \frac{\int_A \pi(\alpha) \exp\left(-\frac{1}{2}(d(\alpha) + z)'(d(\alpha) + z)\right) d\alpha}{\exp\left(-\frac{1}{2}(d(\alpha_0) + z)'(d(\alpha_0) + z)\right)}, \tag{16}$$

where z is defined in Assumption (d) of Theorem 1 and $d(\alpha) = V_\Pi^{1/2}(g_2(\alpha_1) - g_2(\alpha))$

Remark. Theorem 2 implies

$$2\ln(\text{Bayes Factor}(\alpha_0)) \xrightarrow{d} (d(\alpha_0) + z)'(d(\alpha_0) + z)$$
$$+ 2\ln\left(\int_A \pi(\alpha) \exp\left(-\frac{1}{2}(d(\alpha) + z)'(d(\alpha) + z)\right) d\alpha\right) \tag{17}$$

The first term of (17) is a non-central chi-square random variable whose non-central parameter value increases as the hypothetical value α_0 deviates from the true parameter value α_1, while the second term does not depend on α_0. When the parameter is weakly identified, the test based on the Bayes factor is not consistent, but this result shows that it has nontrivial power against fixed alternatives.

In contrast, Proposition 4 below shows that the LR test has power not only against fixed alternatives, but even against local alternatives provided the structural parameters are strongly identified. This fact suggests that the LR interval will be tighter asymptotically than the BF interval.

Proposition 4 (Local Power of the LR Test) Suppose that the assumption of Proposition 3 hold with $\theta = \beta$. In addition assume that the true structural parameter value is $\beta_T = \beta_0 + T^{-1/2}\delta$ and that $D_\beta g(\beta_0)$ has full rank where δ is a $k_2 \times 1$ vector of constants and $D_\beta g(\beta)$ is the $dim(\Pi) \times k$ Jacobian matrix of $g(\beta)$ with respect to β. Then $LR_T(\beta_0)$ converges in distribution to the noncentral chi-square distribution with $\dim(\Pi)$ degrees of freedom and noncentral parameter $\delta' D_\beta g(\beta_0)' V_\Pi D_\beta g(\beta_0)\delta$.

How large the power advantages of the LR test are in small samples and to what extent we can improve on the length of the BF intervals is a question beyond the scope of this paper. Although our theoretical results show that the LR interval is expected to be tighter than the BF interval asymptotically, we find that in practice even the computationally more straightforward BF interval can help us discriminate between alternative hypotheses of

13

economic interest. The practical advantage of the BF interval is that it may be computed easily from existing Bayesian code for DSGE model estimation. No additional numerical estimation is required. This feature facilitates its adoption by the many applied users of Bayesian methods of estimating DSGE models. Given the predominance of these methods in empirical macroeconomics, we focus on applications of the BF interval in the remainder of this paper. The next subsection discusses how we convert the joint BF confidence set to intervals for individual structural parameters. The same procedure could be used on the LR confidence set.

2.5   The Projection Method

Note that our approach does not allow the construction of point estimates of θ, but the projection method can be used to construct confidence intervals for individual elements of θ (see Dufour and Taamouti, 2005, and Chaudhuri and Zivot, 2008, for the projection method in linear IV and GMM models, respectively). The level $(1-\alpha)$ confidence interval for the ith parameter θ_j is $(\underline{\theta}_j, \bar{\theta}_j)$ where the lower and upper confidence bounds are

$$\underline{\theta}_j = \min\left\{\theta_j \in \Theta_j : \min_{\theta_{-j} \in \Theta_{-j}} \text{Bayes Factor}((\theta_j, \theta_{-j})) \le e^{\frac{\chi_k^2(1-\alpha)}{2}}\right\}, \qquad (18)$$

$$\bar{\theta}_j = \max\left\{\theta_j \in \Theta_j : \min_{\theta_{-j} \in \Theta_{-j}} \text{Bayes Factor}((\theta_j, \theta_{-j})) \le e^{\frac{\chi_k^2(1-\alpha)}{2}}\right\}, \qquad (19)$$

and θ_{-j} is the parameter vector excluding θ_j and Θ_{-j} is the parameter space excluding the parameter space for θ_j. These confidence intervals have confidence level $1 - \alpha$ by construction. Because the Bayes factor is not differentiable in θ when it is computed via simulation and because the number of parameters of a typical DSGE model is large, evaluation of (18) and (19) is computationally challenging. We replace Θ in (18) and (19) by the set of Monte Carlo realizations, which reduces the computational burden. This approach is justified because the set of Monte Carlo realizations becomes dense in the parameter space, as the number of Monte Carlo draws increases.

In practice one has to choose the radius of the neighborhood $B_{\delta_T}(\theta_0)$. We suggest the following data-dependent method for choosing δ_T. Because $\delta_T \to 0_{p \times 1}$, we have

$\pi(H_0) \to 0$, $\pi(H_1) \to 1$, $P(H_0|X) \to 0$ and $P(H_1|X) \to 1$. Thus,

$$\text{Bayes Factor}(\theta_0) \approx \frac{\pi(H_0)}{P(H_0|X)} = \frac{\frac{1}{|\delta_T|}\pi(H_0)}{\frac{1}{|\delta_T|}P(H_0|X)}$$

where $|\delta_T| = \Pi_{i=1}^{p}\delta_{T,i}$. We typically compute $\pi(H_0)$ and $P(H_0|X)$ by Monte Carlo simulation:

$$\hat{\pi}(H_0) = \frac{1}{M}\sum_{j=1}^{M}I(\theta^{(j)} \in B_{\delta_T}(\theta_0)),$$

$$\hat{P}(H_0|X) = \frac{1}{M}\sum_{j=1}^{M}I(\tilde{\theta}^{(j)} \in B_{\delta_T}(\theta_0)),$$

where M is the number of Monte Carlo realizations, $\theta^{(j)}$ is the jth Monte Carlo realization from the prior distribution and $\tilde{\theta}^{(j)}$ is the jth realization from the posterior distribution. Thus

$$\frac{1}{|\delta_T|}\hat{\pi}(H_0) = \frac{1}{|\delta_T|M}\sum_{j=1}^{M}I(\theta^{(j)} \in B_{\delta_T}(\theta_0)), \tag{20}$$

$$\frac{1}{|\delta_T|}\hat{P}(H_0|X) = \frac{1}{|\delta_T|M}\sum_{j=1}^{M}I(\tilde{\theta}^{(j)} \in B_{\delta_T}(\theta_0)), \tag{21}$$

Note that the right-hand sides of (20) and (21) can be interpreted as a multivariate density estimator based on a uniform kernel with δ_T as the bandwidth. Consider a multivariate version of Silverman's rule of thumb:

$$\delta_{T,j} = \hat{\sigma}_j\left(\frac{4}{(p+2)T}\right)^{\frac{1}{p+4}} \tag{22}$$

where $\hat{\sigma}_j$ is the standard deviation of the posterior distribution of θ_j.[4] Because the prior and posterior distributions are not necessarily normal and the kernel is not normal, (22) need not be optimal but satisfies assumption (f). Note that if θ_j is strongly identified, $\hat{\sigma}_j = o_p(1)$ and thus $\delta_{T,j} = o_p(T^{-1/2})$; if θ_j is weakly identified, $\hat{\sigma}_j = O_p(1)$ and $\delta_{T,j} = o_p(1)$. Thus it follows from the discussion in Remark 3 for Theorem 1 that the resulting choice of δ_T satisfies Assumption (f).

---

[4]See Wand and Jones, 1995, p.111, for example.

3   An Illustrative Example

We investigate the accuracy of both traditional pseudo-Bayesian methods and the proposed alternative by Monte Carlo simulation. Given the computational complexity of applying these econometric methods repeatedly, we select as an illustrative example a small-scale New Keynesian model, which is often used as an example in the related literature (see, e.g., Canova and Sala 2009).

3.1   Simulation Design

Our model setup is taken from Woodford (2003, pp. 246). The economy consists of a Phillips curve, a Taylor rule, an investment-savings relationship, and the exogenous driving processes z_t and ξ_t:

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t \pi_{t+1}, \tag{PC}$$

$$R_t = \rho_r R_{t-1} + (1 - \rho_r)\phi_\pi \pi_t + (1 - \rho_r)\phi_x x_t + \xi_t, \tag{TR}$$

$$x_t = \mathbb{E}_t x_{t+1} - \sigma \left(R_t - \mathbb{E}_t \pi_{t+1} - z_t \right), \tag{IS}$$

$$z_t = \rho_z z_{t-1} + \sigma^z \varepsilon_t^z,$$

$$\xi_t = \sigma^r \varepsilon_t^r.$$

where x_t, π_t and R_t denote the output gap, inflation rate, and interest rate, respectively. The shocks ε_t^z and ε_t^r are assumed to be distributed $\mathcal{NID}\left(0, 1\right)$. The model parameters are the discount factor β, the intertemporal elasticity of substitution σ, the probability α of not adjusting prices for a given firm, the elasticity of substitution across varieties of good, θ, the parameter ω controlling disutility of labor supply; ϕ_π and ϕ_x capture the central bank's reaction to changes in inflation and the output gap, respectively, and $\kappa = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha} \frac{\omega+\sigma}{\sigma(\omega+\theta)}$.

Clearly, the parameters contained in κ are not separately identified. That is, α and θ are at most partially identified. In practice, macroeconomists often fix some parameters such as β, ω and sometimes θ to allow estimation of α based on κ (see, e.g., Eichenbaum and Fisher 2007), but that procedure is not recommended (see Canova and Sala 2009). For related discussion of this approach also see Komunjer and Ng (2009).

Our Monte Carlo experiment consists of the following steps:

1. We generate $1,000$ synthetic data sets of length T for output and inflation using the New Keynesian model as the DGP. In generating the data, we set $\sigma = 1$, $\alpha = 0.75$, $\beta = 0.99$, $\phi_\pi = 1.5$, $\phi_x = 0.125$, $\omega = 1$, $\rho_r = 0.75$, $\rho_z = 0.90$, $\theta = 6$, $\sigma^z = 0.30$, $\sigma^r = 0.20$. These parameter values are standard choices in the macroeconomics literature (see An and Schorfheide 2007, Woodford 2003). We consider two sample sizes: $T = 96$ and $T = 188$. The smaller sample corresponds to the length of quarterly time series starting with the Great Moderation period in 1984 (see Stock and Watson 2002). The later sample reflects the period between 1960 and 2006.

2. For each synthetic data set, we treat output and inflation as our observables and estimate a total of eight parameters: $\Phi = [\alpha \; \phi_\pi \; \phi_x \; \theta \; \rho_r \; \rho_z \; \sigma^r \; \sigma^z]$. The estimation is carried out using Bayesian estimation methods for DSGE models. We characterize the posterior distribution of the parameters of interest using the Random Walk Metropolis-Hasting algorithm documented in An and Schorfheide (2007). We select two types of priors. First, we use uniform priors. Following the literature, we impose boundary restrictions to make the priors proper and to avoid incompatible values (e.g., negative variances, persistence parameters outside the unit circle, and indeterminacy of the model). As an alternative, we use the priors proposed in An and Schorfheide centered around the true values in our DGP and with loose standard deviations (see Table 1). The algorithm involves three steps:

   a. Let $\mathcal{L}(\Phi|Y)$ and $p(\Phi)$ denote the likelihood of the data conditional on the parameters and the prior probability, respectively. Obtain the posterior mode $\widetilde{\Phi}$ $= \arg\max[\ln p(\Phi) + \ln \mathcal{L}(\Phi|Y)]$ using a suitable maximization routine. To ensure that we find the maximum, we provide our maximization procedure with 10 randomly selected starting points, which gives us a set of potential maxima $\left\{\widetilde{\Phi}_i\right\}_{i=1}^{10}$. Then the mode corresponds to the candidate that achieves the highest value among the 10 potential candidates.

   b. Let $\widetilde{\sum}$ be the inverse Hessian evaluated at the posterior mode. Draw $\Phi^{(0)}$ from a normal distribution with mean $\widetilde{\Phi}$ and covariance matrix $\varkappa^2 \widetilde{\sum}$, where $\varkappa^2$ is a scaling parameter.

17

c. For $k = 1, ..., N_s$, draw ϑ from the proposal density $\mathcal{N}\left(\Phi^{(k-1)}, \varkappa^2 \widetilde{\sum}\right)$. The new draw $\Phi^{(k)} = \vartheta$ is accepted with probability $\min\{1, q\}$ and rejected otherwise. The probability r is given by

$$q = \frac{\mathcal{L}\left(\vartheta|Y\right) p\left(\vartheta\right)}{\mathcal{L}\left(\Phi^{(k-1)}|Y\right) p\left(\Phi^{(k-1)}\right)}.$$

The posterior distributions are characterized using $N_s = 100,000$ iterations after discarding an initial burn-in phase of $1,000$ draws. Selecting $\varkappa^2$ is a delicate issue in our experiment. Ideally, one should fine-tune that parameter for each synthetic data set, so that the acceptance rate falls within the values suggested by Roberts et al. (1997). Given the size of our experiment ($5,000$ Monte Carlo replications each consisting of $100,000$ Metropolis-Hasting draws), hand picking $\varkappa^2$ for each synthetic data set is prohibitively expensive. Instead, we set one common scaling parameter for our exercise. To get this value, we fine tune $\varkappa^2$ for 10 separate Monte Carlo replications and then take the average of scaling parameters.

3.2   Simulation Results

The first table reports some statistics of the priors used in our Monte Carlo experiment. Tables 2 and 3 compare the coverage accuracy of alternative confidence sets for the model discussed above. The first table reports the results using uniform priors, while the second table contains the findings with informative priors. The upper panel is for $T = 96$, which corresponds to the sample size of post-Great Moderation quarterly data. The lower panel is for $T = 188$, which corresponds to the standard sample period between 1960 and 2006. Following Gelfand and Smith's (1990) approach, we visually inspected draws from the posterior distribution and discarded data sets in which convergence seems to fail. That left between 600 and 743 synthetic data sets for each sample size and design. The nominal coverage probability is 0.90. The tuning parameter, δ_T, is chosen by the data-dependent method discussed in section 2.2. In light of the computational cost, the results are based on 5,000 draws randomly chosen from 100,000 draws from the posterior distribution.

The first row of the upper panel of Table 2 focuses on the traditional asymptotic

confidence interval that a frequentist user might construct from the posterior mode (or mean or median) by adding $+/-1.645$ posterior standard errors. Some effective coverage rates are well below the nominal rates. The coverage probability may be as low as 52.5%. Alternatively, a frequentist user may focus on the $(1-\alpha)\%$ equal-tailed percentile interval based on the posterior distribution (see, e.g., Balke, Brown, and Yücel 2008). For the percentile interval, the coverage rate may drop as low as 39.1%. If we construct the interval by inverting the Bayes factor (BF interval), in contrast, all intervals for individual parameters have coverage rates of at least 97%, and the joint interval has a coverage probability of 89.4%.

As the sample size is increased in the second panel, the accuracy of the traditional asymptotic interval improves but may remain as low as 71.5%, depending on the parameter. The corresponding percentile intervals have coverage rates as low as 58.8%. The intervals based on inverting the Bayes factor in all cases have at least 90% coverage probability.

Table 3 shows that under informative priors again the effective coverage rates for the traditional confidence intervals may be below 90%. For example, the coverage probability for the scale of the monetary shock, σ^r, can be as low as 62.3% when $T = 96$. In contrast, the proposed BF interval has coverage rates of at least 90% in all cases. An interesting feature of this second exercise is that the use of informative priors benefits the traditional method in that it improves the coverage accuracy relative to the results in Table 2. This result is expected, since these priors are centered around the true parameter values, which forces the posterior mode/median to remain in the neighborhood of the true parameters. This finding highlights the influence that priors have on the construction of traditional confidence intervals. The Monte Carlo experiment shows that the traditional methods are typically least accurate for the parameters describing the stochastic processes of the DSGE model.

The results in Tables 2 and 3 indicate that the accuracy of some traditional intervals for α and θ can be quite good even when those parameters are weakly identified as in our experiment. The reason is as follows: In the weakly identified case, the posterior distribution essentially replicates the prior distribution (see Canova and Sala, 2009). A

natural conjecture is that the symmetry of the priors for α and θ about their true values is responsible for the relatively high accuracy of the traditional methods because it makes it more likely that the credible interval includes the true parameter value. To verify our conjecture, we repeated the Monte Carlo experiment with uniform priors with bounds $[0, 0.8]$ and $[5.5, 15]$ for α and θ, respectively. Under these alternative priors, the true values are close to the boundary of the support of the priors. As Table 4 shows, in that case, the coverage rates for the traditional confidence intervals decline to values as low as 10%. Even under the most optimistic scenario based on the mode, the accuracy for those parameters is only around 50%. On the other hand, our approach remains quite robust to the new priors delivering coverage rates of at least 94% for individual parameters and near 90% for the set.

We conclude that traditional interval estimates for Bayesian DSGE model estimates are not reliable and that the proposed alternative interval has the potential of achieving substantial improvements in accuracy.

4   Empirical Application

To illustrate the usefulness of our methodology, we now construct the BF confidence intervals in a medium-scale DSGE framework. The model for this section builds on the recent literature on dynamic stochastic general equilibrium models (see, e.g., Altig et al. 2005; Smets and Wouters 2007a,b). Our specification follows very closely that of Del Negro et al. (2007) and Del Negro and Schorfheide (2008), who in turn build on Smets and Wouters (2003) and Christiano, Eichenbaum and Evans (2005). Since this type of environment has been extensively discussed in the literature, we provide only a brief discussion. The main features of the model can be summarized as follows: The economy grows along a stochastic path; prices and wages are assumed to be sticky à la Calvo; preferences display internal habit formation; investment is costly; and finally, there are five sources of uncertainty: neutral and capital embodied technology shocks, preference shocks, government expenditure shocks, and monetary shocks. Additional details on the formulation and estimation of DSGE models can be found in Fernandez-Villaverde et al.

(2009).

4.1 Firms

There is a continuum of monopolistically competitive firms indexed by $j \in [0,1]$ each producing an intermediate good from capital services, k_j, and labor services, $L_{j,t}$. The technology function is given by

$$Y_{j,t} = k_{j,t}^{\alpha} \left(Z_t L_{j,t} \right)^{1-\alpha} - Z_t \psi,$$

where ψ makes profits equal to zero in the steady state. The neutral technology shock, Z_t, grows at rate $z_t = \log\left(Z_t / Z_{t-1} \right)$ which is assumed to follow the process[5]

$$z_t = (1 - \rho_z)\gamma + \rho_z z_{t-1} + \sigma_z \epsilon_{z,t},$$

where $\epsilon_{z,t}$ is distributed $\mathcal{NID}(0,1)$. Firms rent capital and labor in perfectly competitive factor markets.

Firms choose prices to maximize the present value of profits; prices are set in Calvo fashion; that is, each period, firms optimally revise their prices with an exogenous probability $1 - \zeta_p$. If, instead, a firm does not re-optimize its price, then the price is updated according to the rule: $P_{j,t} = \left(\pi_{t-1} \right)^{\iota_p} \left(\pi_* \right)^{1-\iota_p} P_{j,t-1}$, where π_{t-1} is the economy-wide inflation in the previous period, π_* is steady-state inflation and $\iota_p \in [0,1]$.

There is a competitive firm that produces the final good using intermediate goods according to the technology

$$Y_t = \left[\int_0^1 Y_{j,t}^{1/\left(1+\lambda_{f,t}\right)} dj \right]^{1+\lambda_{f,t}}.$$

Here $\lambda_{f,t}$ is the degree of monopoly power and evolves according to the process $\log \lambda_{f,t} = \left(1 - \rho_{\lambda_f}\right) \log \lambda_f + \rho_{\lambda_f} \log \lambda_{f,t-1} + \sigma_{\lambda_f} \epsilon_{\lambda,t}$. The shock $\epsilon_{\lambda,t}$ is assumed to be $\mathcal{NID}(0,1)$.

4.2 Households

The economy is populated by a continuum of households indexed by i. Every period households must decide how much to consume, work, and invest. In addition, they must

---

[5]The growth term is needed to have a well-defined steady state around which we can solve the model.

choose the amount of money to be sent to a financial intermediary. Agents in the economy have access to complete markets; such an assumption is needed to eliminate wealth differentials arising from wage heterogeneity. Households maximize the expected present discounted value of utility

$$\mathbb{E}_0^i \sum_{t=0}^{\infty} \beta^t \left[\log(C_{i,t} - hC_{i,t-1}) - \phi_t \frac{L_{i,t}^{1+v_l}}{1+v_l} \right] \tag{23}$$

subject to

$$P_t C_{i,t} + P_t \left(I_{i,t} + a(u_{i,t}) \overline{K}_{i,t} \right) + B_{i,t+1} = R_t^K u_{i,t} \overline{K}_{i,t} + W_{i,t} L_{i,t} + R_{t-1} B_{i,t} + A_{i,t} + \Pi_t + T_{i,t},$$

and

$$\overline{K}_{i,t+1} = (1-\delta)\overline{K}_{i,t} + I_{i,t} \left(1 - \Gamma(\frac{I_{i,t}}{I_{i,t-1}}) \right).$$

Here, $\mathbb{E}_t^i$ is the time t expectation operator conditional on the information set of household i; ϕ_t is a preference shifter that follows the process $\log \phi_t = (1 - \rho_\phi) \log \phi + \rho_\phi \log \phi_{t-1} + \sigma_\phi \epsilon_{\phi,t}$ with $\varepsilon_{\phi,t}$ distributed $\mathcal{N}(0,1)$; preferences display internal habit formation measured by $h \in (0,1)$; and Γ is a function reflecting the costs associated with adjusting the investment portfolio. This function is assumed to be increasing and convex satisfying $\Gamma(e^\gamma) = \Gamma'(e^\gamma) = 0$ and $\Gamma''(e^\gamma) > 0$ in the steady state. $T_{j,t}$ corresponds to lump-sum transfers from the government to household i. $B_{i,t}$ is the individual demand for one-period government bonds, which pay the gross nominal interest rate R_t. As in the related literature, it is assumed that physical capital can be used at different intensities (see, e.g., Altig et al. 2005). Furthermore, using the capital with intensity $u_{i,t}$ yields the return $R_t^K u_{i,t} \overline{K}_{i,t}$ but entails the cost $a(u_{i,t})$, which satisfies $a(1) = 0; a''(1) > 0; a'(1) > 0$. Finally, the term $A_{i,t}$ captures net payments from complete markets, while Π_t corresponds to profits from producers.

4.3 Wage Setting

Following Erceg et al. (2000), we assume that each household is a monopolistic supplier of a differentiated labor service, $L_{i,t}$. Households sell these labor services to a competitive firm that aggregates labor and sells it to final firms. The technology used by the aggregator

is

$$L_t = \left[\int_0^1 L_{i,t}^{1/(1+\lambda_w)} dj \right]^{1+\lambda_w}, \quad 0 < \lambda_w < \infty.$$

It is straightforward to show that the relationship between the labor aggregate and the aggregate wage, W_t, is given by

$$L_{i,t} = \left[\frac{W_{i,t}}{W_t} \right]^{-(1+\lambda_w)/\lambda_w} L_t.$$

To induce wage sluggishness, we assume that households set their wages in Calvo fashion. In particular, with exogenous probability ζ_w a household does not re-optimize wages each period. Hence, wages are set according to the rule of thumb $W_{i,t} = (\pi_* e^\gamma)^{1-\iota} (\pi_{t-1} e^{z_{t-1}})^{\iota_w} W_{i,t-1}$.

4.4 Government

As in most of the recent New Keynesian literature, we assume a cashless economy (Woodford, 2003). The monetary authority sets the short-term interest rate according to a Taylor rule. In particular, the central bank smoothes interest rates and responds to deviations of actual inflation from steady-state inflation, π_*, and deviations of output from its target level, Y_t^*.

$$\frac{R_t}{R^*} = \left(\frac{R_{t-1}}{R^*} \right)^{\rho_r} \left[\left(\frac{\pi_t}{\pi_*} \right)^{\psi_1} \left(\frac{Y_t}{Y_t^*} \right)^{\psi_2} \right]^{1-\rho_r} \exp(\sigma_r \epsilon_{r,t}). \tag{24}$$

The term $\epsilon_{r,t}$ is a random shock to the systematic component of monetary policy and is assumed to be standard normal; σ_r is the size of the monetary shock. This is the same Taylor rule used in Del Negro et al. (2007) and Del Negro and Schorfheide (2008). R^* corresponds to the steady-state gross nominal interest rate.

Finally, we assume that government spending is given by $G_t = (1 - 1/g_t) Y_t$ where g_t follows the exogenous process $\log g_t = (1 - \rho_g) \log g + \rho_g \log g_{t-1} + \sigma_g \epsilon_{g,t}$, where $\epsilon_{g,t} \sim \mathcal{NID}(0,1)$. The government uses taxes and one-period bonds to finance its purchases.

4.5 Data and Estimation

We follow Del Negro and Schorfheide (2008) in estimating the model using five observables: real output growth, per capita hours worked, labor share, inflation (annualized),

23

and nominal interest rates (annualized). We use their quarterly data set for the period $1982.Q1 - 2005.Q4$. We set our priors alternatively to the non-dogmatic agnostic, low-rigidities, and high-rigidities priors employed in Del Negro and Schorfheide (see Tables 1 through 3 in their paper).

The parameter space is divided into two sets: $\Theta_1 = [\alpha\ \delta\ g\ L^*\ \psi]$, which is not estimated, and $\Theta_2 = [r_*\ \gamma\ \lambda_f\ \pi_*\ \zeta_p\ \iota_p\ \zeta_w\ \iota_w\ \lambda_w\ \Gamma''\ h\ a''\ v_l\ \psi_1\ \psi_2\ \rho_r\ \rho_z\ \sigma_z\ \rho_\phi\ \sigma_\phi\ \rho_\lambda\ \sigma_\lambda\ \rho_g\ \sigma_g\ \sigma_r\ L_{adj}]$, which is. The following values are used for the first set of parameters: $\alpha = 0.33$, $\delta = 0.025$, $g = 0.22$, $L^* = 1$, $\psi = 0$. Although these values are standard choices in the DSGE literature, some clarifications are in order. As in Del Negro and Schorfheide (2008), our parametrization imposes the constraint that firms make zero profits in the steady state. We also assume that households work one unit of time in steady state. This assumption in turn has two implications: First, the parameter ϕ is endogenously determined by the optimality conditions in the model. Second, because hours worked have a mean different from that in the data, the measurement equation in the state space representation is

$$\log L_t(data) = \log L_t(model) + \log L_{adj}.$$

Here, the term L_{adj} is required to match the mean observed in the data. Finally, rather than imposing priors on the great ratios as in Del Negro and Schorfheide, we follow the standard practice (Christiano et al. 2005) of fixing the capital share, α, the depreciation rate, δ, and the share of government expenditure on production, g.

The posterior distributions of the parameters in the set Θ_2 are characterized using the Random Walk-Metropolis-Hasting algorithm outlined in Section 3. A total of three independent chains, each of length $100,000$ were run. We conducted standard tests to check the convergence of each chain (see Gelman et al. 2004).

4.6 What Do the Data Tell Us about the Relative Importance of Wage and Price Rigidities?

Table 5 summarizes the posterior means, medians, and modes as well as the posterior standard deviations, as shown in Table 6 of Del Negro and Schorfheide (2008).[6] For each structural parameter, we also show the 90% Bayesian credible interval and the proposed 90% confidence interval based on inverting the Bayes factor (*BF interval*).

For our purposes, the parameters of greatest interest are ζ_p and ζ_w, which quantify the degree of price and wage rigidities, respectively. Del Negro and Schorfheide found that the posterior of these parameters was heavily influenced by their prior, so a researcher entering a prior favoring one of these rigidities would inevitably arrive at a posterior favoring that same rigidity. This finding suggests that a properly constructed confidence band should be wider. Our BF based interval delivers wider confidence intervals, which shields the econometrician from making unduly strong statements about the degree of stickiness in the data. In contrast, a researcher naïvely interpreting the credible sets as frequentist confidence sets would have concluded that these same parameters are fairly tightly estimated. Although the BF intervals are wider, they are not so wide as to make the exercise useless, indicating that even under weak identification there is some information in the data about the structural parameters.

There is an active literature on measuring the degree of price rigidity at the micro level (see, e.g., Klenow and Kryvtsov 2008; Nakamura and Steinsson 2008). For example, Klenow and Kryvtsov (2008) find that price contracts last, on average, about 2.3 quarters. Based on the credible intervals, a researcher would conclude that the length of those price spells is incompatible with the macro evidence in Table 5. In contrast, a researcher relying on the *BF interval* would view Klenow and Kryvtsov's findings as perfectly consistent with the results from the Bayesian estimation exercise (the lower bound of the interval implies that prices are reset every 1.9 quarters).[7] When we turn to wage stickiness, the credible

---

[6] The attentive reader may notice that our posteriors differ somewhat from those in Del Negro and Schorfheide (2008). This is because, as previously explained, we opt not to use priors on the great ratios. For the discussion below, these differences are immaterial.

[7] The length of price contracts is defined as $\frac{1}{1-\zeta_p}$, where ζ_p is the probability of not re-optimizing prices today.

interval favors a model with a fairly flexible wage setting (the longest wage contract lasts only for 1.5 quarters). Our BF approach, however, suggests that the data are compatible with a model displaying wage contracts of up to 2.3 quarters.

Tables 6 and 7 provide evidence that the BF interval is not very sensitive to the choice of prior. We compare the low-rigidity and high-rigidity priors explored by Del Negro and Schorfheide (2008). The BF interval suggests that Klenow and Kryvstov's findings are plausible even under priors that assume substantial price rigidity (Table 7) or low price rigidity (Table 6). To summarize, the BF interval is designed to help protect researchers from overly optimistic inferences. It allows applied users who are merely Bayesians of convenience to compute asymptotically valid confidence sets from DSGE models estimated by Bayesian methods, even when conventional methods relying on the asymptotic equivalence of Bayesian and frequentist estimation and inference would be invalid.

5 Concluding Remarks

An attractive feature of Bayesian DSGE estimation methods is that they facilitate the estimation of structural models that are too large to be estimated reliably by conventional maximum likelihood methods. This feature has made these methods popular even among researchers who think of these methods merely as a convenient device for obtaining model estimates but would not consider themselves Bayesians otherwise. If the DSGE model is only weakly identified, however, Bayesian posterior estimates tend to be dominated by the prior and the usual asymptotic equivalence between frequentist and Bayesian methods of estimation and inference breaks down. We showed that attempts to construct classical confidence sets from the posterior, whether based on percentile intervals or by adding multiples of posterior standard deviations to posterior modes, are invalid if the model is weakly identified. Moreover, the posterior mode, mean, and median are inconsistent estimators of the true structural parameter values. Given mounting evidence that many DSGE models used in the literature suffer from weak identification problems, this finding suggests caution in interpreting posterior modes as traditional point estimates and

highlights the limitations of traditional confidence sets constructed from the posterior.

We proposed two alternative frequentist confidence sets that remain asymptotically valid regardless of the strength of identification. One is constructed by inverting a LR statistic; the other involves inverting the Bayes factor. We showed that the proposed confidence sets tends to have higher coverage accuracy than the alternative methods we showed to be theoretically invalid. The proposed confidence sets are designed to help applied users separate the information conveyed by the data from the information conveyed by the prior. This is an especially useful feature for non-Bayesian users of Bayesian DSGE estimation methods, given recent evidence that DSGE models with very different policy implications may be observationally equivalent. This means that the posterior tends to move nearly one for one with the prior. In such cases, one would like a frequentist confidence set to reflect the fact that there is little information about the structural parameter in the data. This is indeed what we found in several examples based on the recent literature. While the proposed intervals tend to be appropriately wide, we also showed that it is not necessarily the case that they are completely uninformative. We showed that even under weak identification the data contained enough information to discriminate between competing hypotheses of economic interest. Moreover, we found that the strength of the identification and hence the width of the confidence set may differ from one structural parameter to the next.

We contrasted the relative merits of the LR and BF approach from a theoretical and a practical point of view. Although LR intervals are expected to be tighter asymptotically than BF intervals, their implementation can be computationally challenging and requires fundamental changes in the estimation of DSGE models. In contrast, the BF interval is straightforward to implement using existing Bayesian code for estimating DSGE models. Given the predominance of Bayesian methods in practice, we focused on the BF interval in the empirical applications. The proposed BF confidence set takes full advantage of Bayesian estimation methods in that it is based on the inversion of the Bayes factor. Our BF method has two attractive features. One is that it circumvents the problems of estimating structural DSGE models by classical maximum likelihood methods by using the Bayesian estimation framework in constructing the Bayes factor. The other is that by

construction the proposed confidence set is asymptotically invariant to the choice of prior in the case of weak identification. Since the Bayes factor is the ratio of the posterior odds to the prior odds, if the likelihood is flat and hence the prior dominates the posterior, the numerator and the denominator of the ratio will tend to cancel, making the proposed confidence set more robust to alternative priors than conventional intervals. We illustrated this point in the context of the question of the relative importance of wage and price rigidities in a New Keynesian model.

It is of course not necessary to use our method to reveal weak identification. For example, inspection of the likelihood or a comparison of the posterior under different priors may suffice to diagnose problems of weak identification, as illustrated in Del Negro and Schorfheide (2008). Weak identification does not mean that there is no information in the data about the structural parameters of interest, however. The value added of the BF and LR intervals relative to other methods is that they allow one to quantify how informative the data are about the structural parameter of interest.

Appendix

Proof of Proposition 1: Let $B_T = \{\theta \in \Theta : [\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)]^{1/2}(\theta - \hat{\theta}_T) \in A\}$. Note that $B_T \overset{P_{\theta_0}}{\to} \{\theta_0\}$ because A is compact, $\hat{\theta}_T$ is strongly consistent and $\nabla_{\theta\theta}\ell_T(\theta)$ is diverging. Recall that $B_\delta(\theta_0) = \{\theta \in \Theta : \|\theta - \theta_0\| \leq \delta\}$ as defined in Assumption (d) of Proposition 1. Define $I_{1,T}$ and $I_{2,T}$ by

$$
\frac{|\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)|^{\frac{1}{2}}}{\pi(\theta_0)L_T(X_1,...,X_T|\hat{\theta}_T)} \int_\Theta \pi(\theta)L_T(X_1,...,X_T|\theta)d\theta
$$

$$
= \frac{|\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)|^{\frac{1}{2}}}{\pi(\theta_0)L_T(X_1,...,X_T|\hat{\theta}_T)} \int_{\Theta \cap B_\delta(\theta_0)^c} \pi(\theta)L_T(X_1,...,X_T|\theta)d\theta
$$

$$
+ \frac{|\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)|^{\frac{1}{2}}}{\pi(\theta_0)L_T(X_1,...,X_T|\hat{\theta}_T)} \int_{\Theta \cap B_\delta(\theta_0)} \pi(\theta)L_T(X_1,...,X_T|\theta)d\theta
$$

$$
= I_{1,T} + I_{2,T}. \tag{25}
$$

$$
\begin{aligned}
I_{1,T} &= \frac{1}{\pi(\theta_0)} \exp(\ell_T(\theta_0) - \ell_T(\hat{\theta}_T))|\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)|^{\frac{1}{2}} \\
&\quad \times \int_{\Theta \cap B_\delta(\theta_0)^c} \pi(\theta)\exp(\ell_T(\theta) - \ell_T(\theta_0))d\theta \\
&\leq \frac{1}{\pi(\theta_0)}|\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)|^{\frac{1}{2}}\exp(-\varepsilon(\delta)T) \\
&\to 0, \tag{26}
\end{aligned}
$$

where the first inequality follows from $\exp(\ell_T(\theta_0) - \ell_T(\hat{\theta}_T)) \leq 1$ and Assumption (f) and the last convergence follows from Assumptions (d) and (f).

$$
\begin{aligned}
I_{2,T} &= |\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)|^{\frac{1}{2}} \int_{B_\delta(\theta_0)} \frac{\pi(\theta)}{\pi(\theta_0)}\exp(\ell_T(\theta) - \ell_T(\hat{\theta}_T))d\theta \\
&= |\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)|^{\frac{1}{2}} \int_{B_\delta(\theta_0)} \exp(\ell_T(\theta) - \ell_T(\hat{\theta}_T))d\theta + O(\delta) \\
&= |\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)|^{\frac{1}{2}} \int_{B_\delta(\theta_0)} \exp\left[\frac{1}{2}(\theta - \hat{\theta}_T)'\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)(\theta - \hat{\theta}_T)\right]\exp(R_T(\theta))d\theta \\
&\quad + O(\delta) \ P_{\theta_0}\text{-a.s.} \\
&= |\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)|^{\frac{1}{2}} \int_{B_\delta(\theta_0)} \exp\left[\frac{1}{2}(\theta - \hat{\theta}_T)'\nabla_{\theta\theta}\ell_T(\hat{\theta}_T)(\theta - \hat{\theta}_T)\right]d\theta \\
&\quad + O(\delta) + o(1) \ P_{\theta_0}\text{-a.s.} \\
&\to (2\pi)^{\frac{1}{2}}, \tag{27}
\end{aligned}
$$

as $\delta \to 0$, where $R_T(\theta) = (\theta - \hat{\theta}_T)' \nabla_{\theta\theta} \ell_T(\bar{\theta}_T)(\theta - \hat{\theta}_T) - (\theta - \hat{\theta}_T)' \nabla_{\theta\theta} \ell_T(\hat{\theta}_T)(\theta - \hat{\theta}_T)$,
$\bar{\theta}_T$ is a point between $\hat{\theta}_T$ and θ_0, the second equality follows from Assumption (b) and $\exp(\ell_T(\theta) - \ell_T(\hat{\theta}_T)) \leq 1$, the third follows from the Taylor's theorem, the fourth from Assumption (c) and the last from Assumption (f). It follows from (25), (26) and (27) that

$$\frac{|\nabla_{\theta\theta} \ell_T(\hat{\theta}_T)|^{\frac{1}{2}}}{\pi(\theta_0) L_T(X_1, ..., X_T | \hat{\theta}_T)} \int_\Theta \pi(\theta) L_T(X_1, ..., X_T | \theta) d\theta \to (2\pi)^{\frac{1}{2}} \qquad (28)$$

For sufficiently large T, $B_T(\hat{\theta}_T) \subset B_\delta(\theta_0)$. By repeating arguments we obtain

$$\frac{|\nabla_{\theta\theta} \ell_T(\hat{\theta}_T)|^{\frac{1}{2}}}{\pi(\theta_0) L_T(X_1, ..., X_T | \hat{\theta}_T)} \int_{\Theta \cap B_T(\hat{\theta}_T)} \pi(\theta) L_T(X_1, ..., X_T | \theta) d\theta \to (2\pi)^{\frac{1}{2}} P(z \in A) \qquad (29)$$

where $z \sim N(0, I(\theta_0))$. The desired result follows from (28) and (29).

Proof of Theorem 1: It follows from Assumption (c), the Taylor theorem and the first order condition for MLE that

$$
\begin{aligned}
I_3 &\equiv \int_{B_{\delta_T}(\theta_0)} \pi(\theta) \exp(\ell_T(g_T(\theta)) - \ell_T(\hat{\Pi}_T)) d\theta \\
&= \int_{B_{\delta_T}(\theta_0)} \pi(\theta) \exp\left(\frac{1}{2}(g_T(\theta) - \hat{\Pi}_T)' \nabla_{\Pi\Pi} \ell_T(\bar{\Pi}_T(\theta))(g_T(\theta) - \hat{\Pi}_T)\right) d\theta, \quad (30)
\end{aligned}
$$

where $\bar{\Pi}_T(\theta)$ is between $g_T(\theta)$ and $\hat{\Pi}_T$. It follows from Assumptions (d) and (f) that

$$
\begin{aligned}
g_T(\theta) &= g_T(\theta_0) + Dg_T(\bar{\theta}(\theta))(\theta - \theta_0) \\
&= g_T(\theta_0) + o(T^{-1/2}), \qquad (31)
\end{aligned}
$$

where $\bar{\theta}(\theta)$ is a point between θ and θ_0. It follows from Assumptions (a) and (b), (30) and (31) that

$$
\begin{aligned}
I_3 &= \int_{B_{\delta_T}(\theta_0)} \pi(\theta) \exp\left(\frac{1}{2}(g_T(\theta_0) - \hat{\Pi}_T)' \nabla_{\Pi\Pi} \ell_T(g_T(\theta_0))(g_T(\theta_0) - \hat{\Pi}_T)\right) d\theta + o_p(1) \\
&= \int_{B_{\delta_T}(\theta_0)} \pi(\theta) d\theta \exp\left(\frac{1}{2}(g_T(\theta_0) - \hat{\Pi}_T)' \nabla_{\Pi\Pi} \ell_T(g_T(\theta_0))(\theta_0) - \hat{\Pi}_T)\right) + o_p(1)
\end{aligned}
$$
$$\qquad (32)$$

It follows from Assumption (e) and (39) that

$$I_3 = \int_{B_{\delta_T}(\theta_0)} \pi(\theta) d\theta \exp\left(-\frac{1}{2} z'z\right) + o_p(1) \qquad (33)$$

30

where z is the standard normal random vector defined in Assumption (e).

Let

$$I_4 = \int_{\Theta \backslash B_{\delta_T}(\theta_0)} \pi(\theta) \exp(\ell_T(g_T(\theta)) - \ell_T(\hat{\Pi}_T)) d\theta, \tag{34}$$

where $A \backslash B = \{x : x \in A \text{ and } x \notin B\}$. Since $\ell_T(g_T(\theta)) \leq \ell_T(\hat{\Pi}_T)$ by the definition of MLE, it follows from (34) that

$$I_4 \leq \int_{\Theta \backslash B_{\delta_T}} \pi(\theta) d\theta. \tag{35}$$

Because $\int_\Theta \pi(\theta) \exp(\ell_T(g(\theta))) d\theta$ cancels out, the Bayes factor in favor of H_1 can be written as

$$\begin{aligned}
\text{Bayes Factor } (\theta_0) &= \frac{\int_{B_{\delta_T}(\theta_0)} \pi(\theta) d\theta}{\int_{\Theta \backslash B_{\delta_T}(\theta_0)} \pi(\theta) d\theta} \frac{I_4}{I_3} \\
&\leq \exp\left(\frac{1}{2} z' z\right) + o_p(1) \tag{36}
\end{aligned}$$

where the inequality follows from (35) and the last equality from (33). Therefore it follows from (36) that

$$2\ln(\text{Bayes Factor } (\theta_0)) \leq z' z + o_p(1) \tag{37}$$

from which we obtain the desired result.

Proof of Theorem 2: Because the log-likelihood function is twice continuously differentiable by Assumption (c) and because $\hat{\Pi}_T - T^{-1/2} g_2(\alpha) = o_p(1)$ uniformly in $\alpha \in A$ where the uniform convergence follows from Assumption (c) and the compactness of A by Assumption (a),

$$\frac{1}{T} \nabla_{\Pi\Pi} \ell_T(\bar{\Pi}_T(\alpha)) \xrightarrow{p} -V_\Pi^{-1} \tag{38}$$

where $\bar{\Pi}_T$ is any point between $\hat{\Pi}_T$ and $T^{-1/2} g_2(\alpha)$ and the convergence is uniform in α.

Define I_3 and I_4 as in the proof of Theorem 1. It follows from Assumptions (a), (b) and

(e), (30), (31) and (38) that

$$
\begin{aligned}
I_3 &= \int_{B_{\delta_T}(\alpha_0)} \pi(\alpha) \exp\left(\frac{1}{2}(g_T(\alpha_0) - \hat{\Pi}_T)'\nabla_{\Pi\Pi}\ell_T(\bar{\Pi}_T)(g_T(\alpha_0) - \hat{\Pi}_T)\right) d\alpha + o_p(1) \\
&= \int_{B_{\delta_T}(\alpha_0)} \pi(\alpha) d\alpha \exp\left(\frac{1}{2}(g_2(\alpha_0) - \hat{\Pi}_T)'\nabla_{\Pi\Pi}\ell_T(\bar{\Pi}_T)(g_2(\alpha_0) - \hat{\Pi}_T)\right) + o_p(1) \\
&= \int_{B_{\delta_T}(\alpha_0)} \pi(\alpha) d\alpha \\
&\quad \times \exp\left(\frac{1}{2}(g_2(\alpha_0) - g_2(\alpha_1) + g_2(\alpha_1) - \hat{\Pi}_T)'\nabla_{\Pi\Pi}\ell_T(\bar{\Pi}_T)\right. \\
&\quad \times \left.(g_2(\alpha_0) - g_2(\alpha_1) + g_2(\alpha_1) - \hat{\Pi}_T)\right) + o_p(1) \\
&= \int_{B_{\delta_T}(\alpha_0)} \pi(\alpha) d\alpha \exp\left(-\frac{1}{2}(d(\alpha_0) + z)'(d(\alpha_0) + z)\right), \tag{39}
\end{aligned}
$$

where $\bar{\Pi}_T$ is a point between $\hat{\Pi}_T$ and $g_T(\alpha_0)$.

It follows from (38) that

$$
\begin{aligned}
I_4 &= \int_{A \backslash B_{\delta_T}(\alpha_0)} \pi(\alpha) \exp(\ell_T(g_2(\alpha)) - \ell_T(\hat{\Pi}_T)) d\alpha \tag{40} \\
&= \int_{A \backslash B_{\delta_T}(\alpha_0)} \pi(\alpha) \exp\left(\frac{1}{2}(g_2(\alpha) - g_2(\alpha_1) + g_2(\alpha_1) - \hat{\Pi}_T)'\nabla_{\Pi\Pi}\ell_T(\bar{\Pi}_T(\alpha))\right. \\
&\quad \left. \times(g_2(\alpha) - g_2(\alpha_1) + g_2(\alpha_1) - \hat{\Pi}_T)\right) \tag{41} \\
&= \int_{A \backslash B_{\delta_T}(\alpha_0)} \pi(\alpha) \exp\left(-\frac{1}{2}(d(\alpha) + V_{\Pi}^{\frac{1}{2}}z)'(d(\alpha) + V_{\Pi}^{\frac{1}{2}}z)\right) + o_p(1) \tag{42} \\
&= \int_{A} \pi(\alpha) \exp\left(-\frac{1}{2}(d(\alpha) + V_{\Pi}^{\frac{1}{2}}z)'(d(\alpha) + V_{\Pi}^{\frac{1}{2}}z)\right) + o_p(1), \tag{43}
\end{aligned}
$$

where $\bar{\Pi}_T(\alpha)$ is between $\hat{\Pi}_T$ and $g_2(\alpha)$ and the last equality follows since $B_\delta(\alpha_0)$ is compact and $d(\alpha)$ is continuous in α.

Combining (39) and (43), the Bayes factor in favor of H_1 can be written as

$$
\begin{aligned}
\text{Bayes Factor}(\alpha_0) &= \frac{\int_{B_{\delta_T}(\theta_0)} \pi(\theta) d\theta}{\int_{\Theta \backslash B_{\delta_T}(\theta_0)} \pi(\theta) d\theta} \frac{I_4}{I_3} \\
&= \frac{\int \pi(\alpha) \exp\left(-\frac{1}{2}(d(\alpha) + z)'(d(\alpha) + z)\right) d\alpha}{\exp\left(-\frac{1}{2}(d(\alpha_0) + z)'(d(\alpha_0) + z)\right)} + o_p(1), \tag{44}
\end{aligned}
$$

which completes the proof.

Acknowledgments

References

1. Altig, David, Lawrence Christiano, Martin Eichenbaum, and Jesper Linde (2005), "Firm-Specific Capital, Nominal Rigidities and the Business Cycles," NBER working paper 11034.

2. An, Sungbae, and Frank Schorfheide (2007), "Bayesian Analysis of DSGE Models," *Econometric Reviews*, 26, 113-172.

3. Andrews, Donald W.K. (1994), "The Large Sample Correspondence Between Classical Hypothesis Tests and Bayesian Posterior Odds Tests," *Econometrica* 62, 1207–1232.

4. Balke, Nathan S., Stephen P.A. Brown, and Mine K. Yücel (2008), "Globalization, Oil Price Shocks and U.S. Economic Activity," unpublished manuscript, Federal Reserve Bank of Dallas.

5. Bickel, Peter J. and Kjell Doksum (2007), *Mathematical Statistics: Basic Concepts and Selected Ideas,* Vol. I, Second Edition, Prentice Hall: Upper Saddle River, NJ.

6. Canova, Fabio, and Luca Sala (2009), "Back to Square One: Identification Issues in DSGE Models," *Journal of Monetary Economics*, 56, 431-449.

7. Chaudhuri, Saraswata, and Eric Zivot (2008), "A New Method of Projection-Based Inference in GMM with Weakly Identified Nuisance Parameters," unpublished manuscript, University of North Carolina, Chapel Hill, and University of Washington, Seattle.

8. Chernozhukov, Victor and Han Hong (2003), "An MCMC Approach to Classical Estimation," *Journal of Econometrics,* 115, 293–346.

9. Choi, In, and Peter C.B. Phillips (1992), "Asymptotic and Finite Sample Distribution Theory for IV Estimators and Tests in Partially Identified Structural Equations," *Journal of Econometrics,* 51, 113–150.

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

11. Del Negro, Marco, and Frank Schorfheide (2008), "Forming Priors for DSGE models (and How It Affects the Assessment of Nominal Rigidities)," *Journal of Monetary Economics*, 55, 1191-1208.

12. Del Negro, Marco, Frank Schorfheide, Frank Smets and Rafael Wouters (2007), "On the Fit and Forecasting Performance of New Keynesian Models," *Journal of Business and Economic Statistics*, 25, 123–162.

13. Dufour, Jean-Marie, and Mohamed Taamouti, (2005), "Projection-Based Statistical Inference in Linear Structural Models with Possibly Weak Instruments," *Econometrica,* 73, 1351–1365.

14. Eichenbaum, Martin, and Jones Fisher (2007), "Estimating the Frequency of Price Re-optimization in Calvo-style Models," *Journal of Monetary Economics*, 54, 2023-2047.

15. Erceg, C., D. Henderson, and A. Levin (2000), "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics,* 46, 281-313.

16. Fernández-Villaverde, Jesús, Pablo A. Guerron-Quintana, Juan F. Rubio-Ramírez (2009), "The New Macroeconometrics: A Bayesian Approach," *Handbook of Applied Bayesian Analysis*, forthcoming.

17. Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent and Mark W. Watson (2007), "A, B, C's (and D's) for Understanding VARs," *American Economic Review,* 97, 1021-1026.

18. Fukač, Martin, Daniel F. Waggoner and Tao Zha (2007), "Local and Global Identification of DSGE Models: A Simultaneous-Equation Approach," unpublished manuscript, Reserve Bank of New Zealand and Federal Reserve Bank of Atlanta.

19. Gelfand, Allan and Adrian Smith (1990), "Sampling Based Approaches to Calculating Marginal Densities," *Journal of American Statistical Association,* 85, 398–409.

20. Gelman, Andrew, John B. Carlin, Hal S. Stern and Donald B. Rubin (2004), *Bayesian Data Analysis,* second edition, Chapman & Hall/CRC: Boca Raton, FL.

21. Hahn, Jinyong (1997), "Bayesian Bootstrap of the Quantile Regression Estimator: A Large Sample Study," *International Economic Review,* 38, 795–808.

22. Iskrev, Nikolay (2008), "How Much Do We Learn from the Estimation of DSGE Models? A Case Study of Identification Issues in a New Keynesian Business Cycle Model," unpublished manuscript, University of Michigan.

23. Kadane, Joseph B. (1975), "The Role of Identification in Bayesian Theory," in: Stephen E. Fienberg and Arnold Zellner (eds.), *Studies in Bayesian Econometrics and Statistics: In Honor of Leonard J. Savage*, North-Holland: Amsterdam, 175-191.

24. Klenow, Peter J. and Oleksiy Kryvtsov (2008). "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics*, 123, 863-904.

25. Kim, Jae-Young (1998), "Large Sample Properties of Posterior Densities, Bayesian Information Criterion and the Likelihood Principle in Nonstationary Time Series Models," *Econometrica,* 66, 359–380.

26. Komunjer, Ivana, and Serena Ng (2009), "Dynamic Identification of DSGE Models," unpublished manuscript, University of California, San Diego, and Columbia University.

27. Le Cam, Lucien, and Grace Lo Yang (2000), *Asymptotics in Statistics: Some Basic Concepts*, Second Edition, Springer-Verlag: New York, NY.

28. Moon, Hyungsik Roger, and Frank Schorfheide (2009), "Bayesian and Frequentist Inference in Partially Identified Models," unpublished manuscript, University of Southern California and University of Pennsylvania.

29. Nakamura, Emi, and Jón Steinsson (2008), "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics,* 123, 1415–1464.

30. Nelson, Charles R. and Richard Startz (2007), "The Zero-Information-Limit Condition and Spurious Inference in Weakly Identified Models," *Journal of Econometrics,* 138, 47–62.

31. Phillips, Peter C.B. and Werner Ploberger (1996), "An Asymptotic Theory of Bayesian Inference For Time Series," *Econometrica,* 64, 381–412.

32. Poirier, Dale J. (1998), "Revising Beliefs in Nonidentified Models," *Econometric Theory*, 14, 483-509.

33. Roberts, Gareth, Andrew Gelman, and Walter Gilks (1997). "Weak Convergence and Optimal Scaling of Random Walk Metropolis Algorithms." *Annals of Applied Probability*, 7, 110-120.

34. Rubio-Ramirez, Juan F., Daniel Waggoner and Tao Zha (2005), "Markov-Switching Structural Vector Autoregressions: Theory and Application," unpublished manuscript, Federal Reserve Board of Atlanta.

35. Rubio-Ramirez, Juan F., Daniel Waggoner and Tao Zha (2009), "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *Review of Economic Studies,* forthcoming.

36. Smets, F., and R. Wouters (2003), "An Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association*, 1, 1123-1175.

37. Smets, Frank, and Rafael Wouters (2007a), "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586-606.

38. Smets, Frank, and Rafael Wouters (2007b), "On the Fit and Forecasting Performance of New Keynesian Models," *Journal of Business and Economic Statistics*, 25, 123-162.

39. Stock, James H. and Mark Watson (2002), "Has the Business Cycle Changed and Why?" *NBER Macroeconomics Annual 2002*, 159-218.

40. Stock, James H. and Jonathan H. Wright (2000), "GMM with Weak Identification," *Econometrica,* 68, 1055–1096.

41. Wand, M.P. and M.C. Jones (1995), *Kernel Smoothing*, Chapman & Hall: London, UK.

42. Woodford, Michael M. (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy,* Princeton University Press: Princeton, NJ.

Table 1: Prior Specification for Parameters of the Small-Scale New Keynesian Model

Uniform Priors

Parameters	Distributions	Lower Bound	Upper Bound
ϕ_p	Uniform	1	5
ϕ_x	Uniform	0	5
α	Uniform	0	1
θ	Uniform	1	15
ρ_z	Uniform	0	1
ρ_r	Uniform	0	1
σ_z	Uniform	0	1
σ_r	Uniform	0	1

Informative Priors

Parameters	Distributions	Mean	Standard Deviations
ϕ_p	Gamma	1.5	0.25
ϕ_x	Gamma	0.125	0.1
α	Beta	0.75	0.2
θ	Normal	6	2
ρ_z	Beta	0.9	0.2
ρ_r	Beta	0.75	0.2
σ_z	Inverse Gamma	0.3	0.2
σ_r	Inverse Gamma	0.2	0.2

Table 2: Effective Coverage Rates of Nominal 90% Confidence Intervals Based on the Posterior and Bayes Factor Intervals: Small-Scale New Keynesian Model with Uniform Priors

	ϕ_π	ϕ_x	α	θ	ρ_z	ρ_r	σ_z	σ_r	Joint
				$T = 96$					
Mean$\pm1.645SD$	0.996	0.971	0.923	1.000	0.793	0.858	0.593	0.525	
Median$\pm1.645SD$	0.990	0.971	0.793	1.000	0.790	0.858	0.623	0.642	
Mode$\pm1.645SD$	0.926	0.983	0.846	0.772	0.864	0.861	0.891	0.946	
Percentile	1.000	0.952	0.993	1.000	0.794	0.855	0.535	0.391	
BF Interval	1.000	0.994	1.000	1.000	0.970	0.972	0.991	0.972	0.894
				$T = 188$					
	ϕ_π	ϕ_x	α	θ	ρ_z	ρ_r	σ_z	σ_r	Joint
Mean$\pm1.645SD$	0.995	0.985	0.977	0.998	0.832	0.900	0.715	0.733	
Median$\pm1.645SD$	0.995	0.987	0.943	0.998	0.832	0.898	0.748	0.778	
Mode$\pm1.645SD$	0.973	0.977	0.880	0.782	0.880	0.893	0.925	0.962	
Percentile	0.997	0.982	0.993	0.998	0.832	0.895	0.643	0.588	
BF Interval	0.998	0.993	0.998	0.998	0.992	0.993	0.985	0.977	0.932

Table 3: Effective Coverage Rates of Nominal 90% Confidence Intervals Based on the Posterior and Bayes Factor Intervals: Small-Scale New Keynesian Model with Informative Priors

	ϕ_π	ϕ_x	α	θ	ρ_z	ρ_r	σ_z	σ_r	Joint
				$T = 96$					
Mean$\pm1.645SD$	0.997	0.996	0.995	0.997	0.880	0.915	0.890	0.945	
Median$\pm1.645SD$	0.997	0.977	0.995	0.997	0.875	0.911	0.869	0.934	
Mode$\pm1.645SD$	0.997	0.840	0.974	0.997	0.774	0.902	0.623	0.727	
Percentile	0.997	0.996	0.995	0.997	0.887	0.915	0.921	0.964	
BF Interval	0.997	0.997	0.997	0.997	0.996	0.995	0.995	0.997	0.969
				$T = 188$					
	ϕ_π	ϕ_x	α	θ	ρ_z	ρ_r	σ_z	σ_r	Joint
Mean$\pm1.645SD$	0.999	0.965	0.999	0.999	0.908	0.921	0.913	0.911	
Median$\pm1.645SD$	0.999	0.938	0.999	0.999	0.905	0.919	0.899	0.894	
Mode$\pm1.645SD$	0.999	0.761	0.987	0.999	0.761	0.915	0.636	0.699	
Percentile	0.999	0.975	0.997	0.999	0.911	0.919	0.930	0.921	
BF Interval	0.999	0.999	0.999	0.999	0.994	0.982	0.988	0.991	0.934

Table 4: Effective Coverage Rates of Nominal 90% Confidence Intervals Based on the Posterior and Bayes Factor Intervals: Small-Scale New Keynesian Model with Modified Uniform Priors

	ϕ_π	ϕ_x	α	θ	ρ_z	ρ_τ	σ_z	σ_r	joint
				$T = 188$					
Mean±1.645SD	0.996	0.988	0.104	0.490	0.808	0.931	0.635	0.679	
Median±1.645SD	0.996	0.988	0.102	0.406	0.815	0.927	0.671	0.750	
Mode±1.645SD	0.948	0.992	0.529	0.515	0.913	0.902	0.937	0.983	
Percentile	0.996	0.981	0.102	0.263	0.804	0.923	0.552	0.506	
BF Interval	0.996	0.994	0.942	0.994	0.994	0.990	0.973	0.973	0.885

Table 5: Nominal 90% Confidence Intervals Based on the Posterior and Bayes Factor Interval: Medium-Scale New Keynesian Model with Agnostic Priors

	Rigidity Parameters					
	Posterior				Credible	BF confidence
	means	medians	modes	SD	intervals	intervals
ζ_p	0.694	0.695	0.695	0.046	[0.615,0.769]	[0.479, 0.840]
ζ_w	0.219	0.214	0.164	0.071	[0.112,0.341]	[0.036, 0.563]
	Other Endogenous Propagation Parameters					
	Posterior				Credible	BF confidence
	means	medians	modes	SD	intervals	intervals
ν_l	1.767	1.699	1.388	0.530	[1.026,2.749]	[0.458, 4.661]
ψ_1	2.382	2.376	2.353	0.253	[1.983,2.806]	[1.278, 3.390]
ψ_2	0.074	0.073	0.074	0.024	[0.038,0.116]	[0.007, 0.180]
ρ_r	0.722	0.724	0.715	0.038	[0.657,0.781]	[0.547, 0.840]
$\imath_p$	0.096	0.074	0.014	0.083	[0.007,0.260]	[0.000, 0.562]
$\imath_w$	0.279	0.269	0.281	0.121	[0.102,0.489]	[0.000, 0.947]
s'	9.146	8.986	8.986	1.973	[6.177,12.639]	[3.036, 17.459]
h	0.753	0.756	0.762	0.056	[0.657,0.840]	[0.512, 0.905]
a''	0.230	0.214	0.168	0.101	[0.096,0.418]	[0.035, 0.888]
	Exogenous Propagation Parameters					
	Posterior				Credible	BF confidence
	means	medians	modes	SD	intervals	intervals
ρ_z	0.232	0.222	0.185	0.117	[0.052,0.445]	[0.001, 0.676]
ρ_ϕ	0.958	0.960	0.964	0.018	[0.926,0.984]	[0.867, 0.999]
ρ_{λ_f}	0.942	0.949	0.960	0.033	[0.875,0.982]	[0.771, 0.998]
ρ_g	0.916	0.916	0.918	0.027	[0.870,0.959]	[0.781, 0.992]

Table 6: Nominal 90% Confidence Intervals Based on the Posterior and Bayes Factor Interval: Medium-Scale New Keynesian Model with Low-Rigidity Priors

	Rigidity Parameters					
		Posterior			Credible	BF confidence
	means	medians	modes	SD	intervals	intervals
ζ_p	0.658	0.660	0.681	0.045	[0.580,0.730]	[0.466, 0.793]
ζ_w	0.266	0.263	0.260	0.057	[0.177,0.364]	[0.089, 0.551]
	Other Endogenous Propagation Parameters					
		Posterior			Credible	BF confidence
	means	medians	modes	SD	intervals	intervals
ν_l	1.856	1.776	1.404	0.597	[1.018,2.961]	[0.426, 5.016]
ψ_1	2.417	2.412	2.275	0.249	[2.014,2.838]	[1.611, 3.692]
ψ_2	0.074	0.072	0.062	0.023	[0.038,0.115]	[0.007, 0.178]
ρ_r	0.724	0.726	0.692	0.037	[0.659,0.780]	[0.556, 0.836]
$\imath_p$	0.154	0.128	0.015	0.116	[0.014,0.381]	[0.000, 0.684]
$\imath_w$	0.270	0.261	0.293	0.112	[0.101,0.467]	[0.001, 0.770]
s'	8.955	8.788	8.851	2.007	[5.966,12.497]	[3.799, 20.434]
h	0.740	0.742	0.736	0.055	[0.648,0.826]	[0.535, 0.914]
a''	0.242	0.226	0.157	0.103	[0.103,0.430]	[0.035, 0.858]
	Exogenous Propagation Parameters					
		Posterior			Credible	BF confidence
	means	medians	modes	SD	intervals	intervals
ρ_z	0.217	0.207	0.211	0.115	[0.042,0.421]	[0.000, 0.704]
ρ_ϕ	0.956	0.957	0.960	0.018	[0.924,0.982]	[0.853, 0.998]
ρ_{λ_f}	0.950	0.955	0.972	0.027	[0.897,0.984]	[0.782, 1.000]
ρ_g	0.911	0.912	0.911	0.027	[0.866,0.955]	[0.791, 0.996]

Table 7: Nominal 90% Confidence Intervals Based on the Posterior and Bayes Factor
Interval: Medium-Scale New Keynesian Model with High-Rigidity Priors

		Posterior			Credible	BF confidence
	means	medians	modes	SD	intervals	intervals
		Rigidity Parameters				
ζ_p	0.773	0.774	0.739	0.056	[0.678,0.855]	[0.580, 0.905]
ζ_w	0.446	0.427	0.312	0.115	[0.288,0.670]	[0.155, 0.848]
		Other Endogenous Propagation Parameters				
	means	medians	modes	SD	intervals	intervals
ν_l	1.520	1.466	1.745	0.667	[0.493,2.709]	[0.133, 4.794]
ψ_1	2.322	2.316	2.326	0.241	[1.944,2.733]	[1.521, 3.374]
ψ_2	0.057	0.055	0.061	0.021	[0.025,0.095]	[0.007, 0.171]
ρ_r	0.744	0.746	0.707	0.036	[0.682,0.800]	[0.577, 0.857]
$\imath_p$	0.082	0.063	0.025	0.072	[0.006,0.224]	[0.000, 0.492]
$\imath_w$	0.188	0.181	0.235	0.091	[0.047,0.347]	[0.001, 0.611]
s'	9.996	9.860	7.798	2.073	[6.857,13.592]	[3.983, 20.630]
h	0.795	0.801	0.771	0.052	[0.701,0.870]	[0.562, 0.935]
a''	0.217	0.202	0.228	0.103	[0.082,0.407]	[0.018, 0.747]
		Exogenous Propagation Parameters				
	means	medians	modes	SD	intervals	intervals
ρ_z	0.235	0.223	0.240	0.125	[0.052,0.463]	[0.000, 0.730]
ρ_ϕ	0.900	0.918	0.945	0.066	[0.749,0.966]	[0.607, 0.996]
ρ_{λ_f}	0.852	0.888	0.967	0.106	[0.647,0.970]	[0.492, 0.999]
ρ_g	0.927	0.926	0.910	0.029	[0.880,0.978]	[0.822, 0.999]