

Cardiff Economics Working Papers



Working Paper No. E2021/22

Estimating macro models and the potentially misleading nature of Bayesian estimation

David Meenagh, Patrick Minford and Michael Wickens

September 2021

ISSN 1749-6010

Cardiff Business School
Cardiff University
Colum Drive
Cardiff CF10 3EU
United Kingdom
t: +44 (0)29 2087 4000
f: +44 (0)29 2087 4419
business.cardiff.ac.uk

This working paper is produced for discussion purpose only. These working papers are expected to be published in due course, in revised form, and should not be quoted or cited without the author's written permission. Cardiff Economics Working Papers are available online from:
<http://econpapers.repec.org/paper/cdfwpaper/> and
business.cardiff.ac.uk/research/academic-sections/economics/working-papers
Enquiries: EconWP@cardiff.ac.uk

Estimating macro models and the potentially misleading nature of Bayesian estimation

David Meenagh*
(Cardiff Business School, Cardiff University)

Patrick Minford†
(Cardiff Business School, Cardiff University, and CEPR)

Michael Wickens‡
(Cardiff Business School, Cardiff University, CEPR, and University of York)

August 2021

Abstract

We ask whether Bayesian estimation creates a potential estimation bias as compared with standard estimation techniques based on the data, such as maximum likelihood or indirect estimation. We investigate this with a Monte Carlo experiment in which the true version of a New Keynesian model may either have high wage/price rigidity or be close to pure flexibility; we treat each in turn as the true model and create Bayesian estimates of it under priors from the true model and its false alternative. The Bayesian estimation of macro models may thus give very misleading results by placing too much weight on prior information compared to observed data; a better method may be Indirect estimation where the bias is found to be low.

Keywords: Bayesian; Maximum Likelihood; Indirect Inference; Estimation Bias

JEL Classification: C11; E12

1 Introduction

This paper addresses the question of how to estimate a macro, DSGE, model. We assume that there is a true model to be discovered and assess the likelihood of discovering it by using various available methods of estimation. In particular we consider Bayesian estimation where the prior distributions may differ from the values of the parameters of the model generating the data; maximum likelihood estimation; and indirect estimation. This question is motivated by two things: first, the increasingly common practice of estimating macroeconomic models by Bayesian methods but not testing the resulting estimated model; second, the findings of Le et al. (2011) that the Bayesian-estimated New Keynesian model of Smets and Wouters (2007) — a seminal paper — is strongly rejected by indirect inference.

In this paper we focus our analysis primarily on the estimation of the degree of price rigidity in the model, since the economy behaves radically differently with and without high rigidity. Under sticky prices, set by producers, demand drives output and employment; and these in turn drive price and wage change, which then impinge on output; under flexible prices, demand at given prices meets output at those prices, and the prices change to equate the two. In macroeconomics, economists divide into two largely non-communicating camps, each of which has pursued its own agenda of policy analysis. On the one hand, a majority of macroeconomists embrace sticky price models and continually estimate them on different country episodes, these days usually via Bayesian methods. On the other, a less numerous but influential group, associated

*Corresponding author: Cardiff Business School, Cardiff University, Colum Drive, Cardiff, UK, CF10 3EU. Tel.: +44 (0)29 2087 5198. Email: MeenaghD@cardiff.ac.uk

†Email: MinfordP@cardiff.ac.uk

‡Email: mike.wickens@york.ac.uk

with Austrian economics (White, 2003) and later Real Business Cycle theory (Kydland and Prescott, 1982), use implicit or calibrated general equilibrium models to predict macro behaviour, rejecting price stickiness as ill-founded; for their calibration they use parameters from pure theory or micro data studies. For example, Chari et al. (2009) write: “Some think New Keynesian models are ready to be used for quarter-to-quarter quantitative policy advice. We do not. Focusing on the state-of-the-art version of these models, we argue that some of its shocks and other features are not structural or consistent with microeconomic evidence. Since an accurate structural model is essential to reliably evaluate the effects of policies, we conclude that New Keynesian models are not yet useful for policy analysis.”

The situation has not been changed by the arrival of micro-founded models since Lucas’ (1976) critique was taken to heart. These DSGE models have simply been specified either within a sticky-price or a flexprice environment, with again no testing in either case. In these models’ specification, this environment is the key driver of behaviour; the parameters of consumption, investment and monetary policy are of secondary importance, it turns out according to the findings of Le et al. (2011) that the Bayesian-estimated New Keynesian model of Smets and Wouters (2007) — a seminal paper — is strongly rejected by a powerful indirect inference test because of its assumptions on price rigidity. Hence it seems to be of great importance to establish good estimates of the degree of price rigidity in macro models. This paper therefore focuses on how this has been and might be done. Examples of macro model estimation include Smets and Wouters(2007)’s use of Bayesian methods, Ireland (2011)’s use of FIML (where he reports finding very high rigidity which he overrules in the final version), and Le et al. (2011) use of Indirect Inference. In this paper , we examine the consequences of using each method, using an extended Monte Carlo experiment.

The estimation of macroeconomic models by Bayesian methods has been facilitated by the development of computer programs such as Dynare (Adjemian et al, 2011) which is freely available and requires little knowledge of econometrics. The use of Bayesian methods was initially an attempt to improve on the use of calibration by combining prior beliefs with data instead of relying just on prior beliefs. In calibration the values of parameters are simply imposed on a model derived from theory; often they are based on estimated micro relationships. Validation of calibrated models was by an informal form of indirect inference in which the simulated moments of key variables were roughly compared with their data counterparts. Originally calibration was a response to what Sargent has referred to in an interview with Evans and Honkapohja (2005) as the rejection of too many “good” models using maximum likelihood (ML) estimation. Calibration is now most commonly used to explore the properties of a theoretical model where the calibration is regarded as providing a numerical representation of the model, and not an estimate of the model. The prior distributions in Bayesian estimation provide a constraint on the influence of the data in determining a model’s coefficients. Roughly speaking, the prior beliefs and the data are weighted in proportion to the precision of their information. In calibration the prior beliefs are treated as exact. In Bayesian estimation they are expressed through (non-degenerative) probability distributions and so provide a stochastic constraint on the data.

If a Bayesian estimated model is rejected by a test, it could be because the choice of prior distributions has produced very misleading posterior (i.e. Bayesian) estimates. Another possibility is that the model is mis-specified. In this paper we are concerned with the implications of the choice of prior distribution and model mis-specification. We examine these issues by formulating a ‘true’ model, generate data from the model and then estimate the model’s parameters using different choices of the prior distributions, including choosing priors from a different model specification. While our primary focus is on the effective estimation of the price rigidity parameters, our experiment also has relevance to the estimation of the other parameters of the model as all contribute to the behaviour of the model in response to policy regime choices. However, where previous work has found that certain parameters have a clear range of values, and there is no controversy over this, using Bayesian priors for these parameters, which effectively reduces the role of the data in the posterior estimates, would seem to be justified. Our focus is on parameters that are both important and controversial, such as the rigidity parameters. The fact that they are controversial lends support to using new sample information which might help to settle the controversy. Using strong Bayesian priors could downgrade this information. In this paper we explore by Monte Carlo experiment just how large this downgrading is. We assume there is no controversy on the other parameters, which is broadly the case, and estimate these with prior means that correspond to the true values. However, we allow the true model to have rigidity parameters that lie at the extremes implied by the controversy. We examine how far setting the wrong Bayesian priors for these can bias their posterior estimates away from their true values. Plainly, if Bayesian studies whose priors

favour New Keynesian models predominate in the literature, any such bias will produce posterior estimates that support the use of New Keynesian models, and vice versa. Similarly, priors favouring flexprice models will, if there is a bias, tend to support the use of flexprice models.

One alternative to using fixed priors is to use empirical Bayesian estimates, in which the posterior distribution is used as the new prior distribution. This would provide a more data-based prior but, if this is repeated, the resulting posterior would converge on the ML estimates based just on the data. Consequently, one might as well have used ML estimates in the first place. We focus solely on fixed priors in Bayesian estimation, and consider ML estimation separately.

If a drawback to using Bayesian estimation is having to choose prior distributions, is there a better way to estimate the model? We consider two alternatives: ML and indirect estimation. Whereas Bayesian estimated models tend to be tightly specified with limited dynamics and restricted error processes, models estimated by ML tend either to produce biased estimates of tightly restricted models, or to be weakly identified, having unrestricted time series error processes in order to improve fit. Both may be attributed to model mis-specification. Sims (1980) argued that macroeconomic models tend to be under-identified, not over-identified as implied by their conventional specification and as required for the use of ML estimation. In consequence he doubted the findings from ML estimation. Instead, he proposed the use of unrestricted VAR (or VARMA) models which always provide a valid representation of the data. An over-identified macro model would imply a VAR with coefficient restrictions.

Indirect estimation involves simulating a structural model for given values of its parameters and then using the simulated data to estimate an auxiliary model whose role is to represent characteristics of the data. Sample moments are an example of an auxiliary model, as are sample scores (derivatives of the likelihood function), but neither captures as many characteristics of the data as, for example, an unrestricted VAR. The estimates of the auxiliary model using the simulated data are then compared with estimates of the auxiliary model obtained from observed data. The given values of the structural parameters are revised until the estimates of the auxiliary model based on the simulated data converge on those from the observed data. Even with an auxiliary model with unrestricted parameters, the estimates of its parameters reflect the structural model's restrictions through the simulated data. The indirect estimates are asymptotically equivalent to maximum likelihood estimates; but in the small samples that in practice we are normally faced with, and are our focus here, the Indirect and ML estimators behave quite differently, as we will see.. One reason for using a VAR (or VARMA) as the auxiliary model is that the solution to a linearised DSGE model is a VAR (or VARMA) with coefficient restrictions. Testing these restrictions provides a test of the structural model. This is known as an indirect test. In a series of papers we and other co-authors have proposed the use of indirect testing for Bayesian-estimated models, see Le et al. (2011,2016), and Meenagh et al. (2018) who report that a variety of auxiliary models, including moments, impulse response functions and VARs give results with similar properties.

The model we use to make our comparisons is the New Keynesian model of the US constructed by Christiano, Eichenbaum and Evans (2005), which was estimated by Bayesian methods by Smets and Wouters (2007). In this model the US is treated as a closed continental economy. In essence it is a standard Real Business Cycle model but with the addition of sticky wages and prices which allows monetary policy feedback to affect the real economy. Although Smets and Wouters found that their estimated model forecasted more accurately than unrestricted VAR models, we note that such forecast tests have been shown to have little power (Minford et al., 2015).

The degree of wage/price rigidity in the economy is a central issue in macroeconomics. In our first set of comparisons we specify a New Keynesian (NK) model with high rigidity. A second set assumes an NK model with virtually full wage/price flexibility — where the Calvo chances of resetting wages and prices is a shade short of 100%; we label this a “flexprice” (FP) model. In all other respects, for maximum simplicity and transparency, the two models are the same. This allows us to focus narrowly on the implications of these various estimation methods for determining the extent of rigidity in macro models.

In all of our experiments we take the NK model as the “true” model (or DGM) and generate 1000 samples from it. Two versions of the model are considered, one with wage/price rigidities and the other with flexible wages and prices. In our first set of experiments we examine the effects of the choice of prior. We obtain Bayesian estimates of each model for each sample using two different priors: a prior with wage/price rigidities (a high rigidity, or HR prior) and a prior with flexible wages and prices (an FP prior). We obtain the very striking result that the choice of prior distribution strongly biases the posterior estimates towards the prior

whatever version of the model generates the simulated data; this bias can be reduced by ‘flattening’ the prior distribution (raising its variance) and centering it closer to the true mean, but it cannot be eliminated. We also compare ML and Indirect estimates of the NK version of the model. It would seem from these findings that the reliance on Bayesian estimation in support of the dominant NK model of the US post-war economy is highly vulnerable to the choice of prior distributions. This might help to explain why these models are rejected by Indirect Inference tests; the tests might be implicitly rejecting the NK priors instead of (or as well as) the specification of the model.

We compare the use of Bayesian estimation with two classical estimators: ML and Indirect estimation. We find that the ML estimates are also highly biased while the Indirect estimates have low bias. This suggests that a better general strategy than using Bayesian estimation might be the use of Indirect estimation.

In Section 2 we show how the choice of prior and biases in the maximum likelihood estimator may affect the posterior estimates in Bayesian estimation. In Section 3 we discuss the choice of the New Keynesian model for our Monte Carlo experiments. The consequences for the Bayesian estimates of the New Keynesian model of alternative choices of the prior distributions are reported in Section 4. We also report the biases when using instead ML and Indirect estimators. A brief summary of our results and their broader implications are reported in Section 5.

2 Bias in Bayesian estimation — the role of priors and data

The effect on the posterior distribution of the choice of prior distribution and of biases in the ML estimator can be illustrated as follows. In classical estimation with data x and T observations we choose θ to maximise the log likelihood function $\ln L(x/\theta)$; i.e.

$$\arg \max_{\theta} \ln L(x/\theta)$$

The ML estimator $\hat{\theta}$ is obtained by solving

$$\frac{\partial \ln L(x/\theta)}{\partial \theta} \Big|_{\theta=\hat{\theta}} = 0.$$

In Bayesian estimation either we estimate θ using the mean of the posterior distribution, or we use the mode of the posterior distribution $\tilde{\theta}$. For a symmetric posterior distribution the mean and the mode are the same. In general, computationally, it is easier to find the mode. To obtain the mode we maximise the posterior distribution; i.e.

$$\arg \max_{\theta} p(\theta/x) \equiv \arg \max_{\theta} \ln p(\theta/x)$$

As

$$\ln p(\theta/x) = \ln L(x/\theta) + \ln p(\theta) - \ln f(x)$$

and the last term doesn’t contain θ , we can ignore it. Hence

$$\arg \max_{\theta} \ln p(\theta/x) \equiv \arg \max_{\theta} [\ln L(x/\theta) + \ln p(\theta)].$$

The mode of the posterior distribution is obtained from

$$\left[\frac{\partial \ln L(x/\theta)}{\partial \theta} + \frac{\partial \ln p(\theta)}{\partial \theta} \right]_{\theta=\tilde{\theta}} = 0. \tag{1}$$

We note that solving $\frac{\partial \ln L(x/\theta)}{\partial \theta} = 0$ for θ gives the mode of the likelihood function (i.e. the ML estimator), and solving $\frac{\partial \ln p(\theta)}{\partial \theta} = 0$ for θ gives the mode of the prior distribution. The posterior mode is obtained by solving the sum of the two.

If $\ln L(x/\theta)$ is flat then the data are uninformative about θ and $\frac{\partial \ln L(x/\theta)}{\partial \theta}$ is close to zero for a range of values of θ . It then follows that the Bayesian estimator is dominated by the prior. If $p(\theta)$ is flat (i.e. the prior is a uniform distribution) then $\frac{\partial \ln p(\theta)}{\partial \theta} = \frac{1}{p(\theta)} \frac{\partial p(\theta)}{\partial \theta} = 0$ and so the data dominate.

To find the posterior mode $\tilde{\theta}$ consider an expansion of (1) about θ_0 which gives

$$\frac{\partial \ln L(x/\theta)}{\partial \theta} + \frac{\partial \ln p(\theta)}{\partial \theta} \simeq \left[\frac{\partial \ln L(x/\theta)}{\partial \theta} + \frac{\partial \ln p(\theta)}{\partial \theta} \right]_{\theta=\theta_0} + \left[\frac{\partial^2 \ln L(x/\theta)}{\partial \theta \partial \theta'} + \frac{\partial^2 \ln p(\theta)}{\partial \theta \partial \theta'} \right]_{\theta=\theta_0} (\theta - \theta_0).$$

Setting this to zero and solving for $\theta = \tilde{\theta}$ gives

$$\tilde{\theta} - \theta_0 = - \left[\frac{\partial^2 \ln L(x/\theta)}{\partial \theta \partial \theta'} + \frac{\partial^2 \ln p(\theta)}{\partial \theta \partial \theta'} \right]_{\theta=\theta_0}^{-1} \left[\frac{\partial \ln L(x/\theta)}{\partial \theta} + \frac{\partial \ln p(\theta)}{\partial \theta} \right]_{\theta=\theta_0}.$$

We can obtain $\tilde{\theta}$ through an iterative process. For iteration r we have $\theta = \tilde{\theta}_{(r)}$, $\theta_0 = \tilde{\theta}_{(r-1)}$. As

$$\begin{aligned} p \lim T^{-1} \frac{\partial^2 \ln L(x/\theta)}{\partial \theta \partial \theta'} &= -p \lim T^{-1} \frac{\partial \ln L(x/\theta)}{\partial \theta} \frac{\partial \ln L(x/\theta)}{\partial \theta'} \\ p \lim T^{-1} \frac{\partial^2 \ln p(\theta)}{\partial \theta \partial \theta'} &= -p \lim T^{-1} \frac{\partial \ln p(\theta)}{\partial \theta} \frac{\partial \ln p(\theta)}{\partial \theta'} \end{aligned}$$

it follows that

$$\tilde{\theta} - \theta_0 \simeq \left[\frac{\partial \ln L(x/\theta)}{\partial \theta} \frac{\partial \ln L(x/\theta)}{\partial \theta'} + \frac{\partial \ln p(\theta)}{\partial \theta} \frac{\partial \ln p(\theta)}{\partial \theta'} \right]_{\theta=\theta_0}^{-1} \left[\frac{\partial \ln L(x/\theta)}{\partial \theta} + \frac{\partial \ln p(\theta)}{\partial \theta} \right]_{\theta=\theta_0} \quad (2)$$

If θ_0 is the true value of θ then asymptotically the mode of the posterior distribution has the distribution

$$\begin{aligned} \tilde{\theta} &\sim N(\theta, \tilde{\Sigma}) \\ \tilde{\Sigma} &= p \lim T \left[\frac{\partial \ln L(x/\theta)}{\partial \theta} \frac{\partial \ln L(x/\theta)}{\partial \theta'} + \frac{\partial \ln p(\theta)}{\partial \theta} \frac{\partial \ln p(\theta)}{\partial \theta'} \right]_{\theta=\theta_0}^{-1} \end{aligned} \quad (3)$$

It follows from (1), (2) and (3) that the posterior mode is approximately a weighted average of the score $\frac{\partial \ln L(x/\theta)}{\partial \theta}$ and $\frac{\partial \ln p(\theta)}{\partial \theta}$. The weights are proportional to the precision of the ML estimator $p \lim T \left[\frac{\partial \ln L(x/\theta)}{\partial \theta} \frac{\partial \ln L(x/\theta)}{\partial \theta'} \right]_{\theta=\theta_0}^{-1}$ and the variance of the prior distribution $p \lim T \left[\frac{\partial \ln p(\theta)}{\partial \theta} \frac{\partial \ln p(\theta)}{\partial \theta'} \right]_{\theta=\theta_0}^{-1}$. The more precise these estimators the more they determine the posterior mode.

We can now see the effect on the posterior mode of the choice of prior and biases in the ML estimator. If the mode of the prior distribution differs from the true value of θ then this will affect $\frac{\partial \ln p(\theta)}{\partial \theta}$ and hence $\tilde{\theta}$. If the ML estimator is biased then this will affect $\frac{\partial \ln L(x/\theta)}{\partial \theta}$ and hence $\tilde{\theta}$. Replacing $\frac{\partial \ln p(\theta)}{\partial \theta}$ and $\frac{\partial \ln L(x/\theta)}{\partial \theta}$ in (2) by these differences gives an approximate idea of their effects. The biases will be weighted by the relevant measure of precision. We conclude that the greater the biases and the measure of precision, the larger will be the effect of these two biases on $\tilde{\theta}$.

3 Choice of model for Monte Carlo experiment — the central role of wage/price rigidity

Our focus on the New Keynesian model and its assumption of widespread rigidity in wage/price setting largely reflects its widespread use by central banks in setting monetary policy. The priors commonly used in the model make monetary policy very powerful. There have, however, been warnings against the uncritical use of the New Keynesian model in policy analysis. For example, Chari et al. (2009) as quoted earlier concluded that New Keynesian models should not be used for policy analysis.

This concern was borne out in the investigation by Le et al. (2011) who questioned the findings of Smets and Wouters on the degree of nominal rigidity in the posterior model. Le et al. applied the indirect inference test to the Smets-Wouters model, first investigating their New Keynesian version and then investigating a New Classical version with no rigidity. They rejected both versions based on the full post-war sample used by Smets and Wouters. With a three-variable VAR(1) (in output, inflation and interest rates) as the auxiliary model they obtained a Wald test equivalent t-value of around 2.5. They noted that the power of this test, though considerable, was lower than that of a Wald test based on a VAR with all 7 variables in the model; the t-value was also very much higher.

They also found that there were two highly significant break-points in the sample, in the mid-1960s and the mid-1980s. They argued that this suggested that there were parts of the economy where prices and wages were flexible. To improve the match to the data they therefore proposed a ‘hybrid’ model and estimated this, not by Bayesian methods, but by indirect estimation. They found that this mixed model better matched the data from the mid-1980s until 2004, a period known as “the great moderation”. However, no such version of the model could match the data for two earlier sub-periods in which there were very low shares for the “flexible sectors”. But when the sample was extended to include the period of financial crisis up to 2012, these shares rose dramatically and became dominant. These findings offer at least partial support for the critics of nominal rigidity.

Using micro-data, Zhou and Dixon (2019) show that matters may be even more complicated. They found that firms normally set prices for a period of time but when shocks are large they change them frequently, implying that there is time-dependence and also shock dependence in the length of pricing periods. In the great moderation period there was a lack of large shocks which could explain the finding of high rigidity. Once the large shocks of the financial crisis hit, this rigidity mostly disappears. Normally, however, there is some rigidity.

This discussion illustrates the two concerns made before about the Bayesian estimation of DSGE models, and especially the ubiquitous New Keynesian model. First, the significance of indirect inference tests of the Smets-Wouters model indicates model mis-specification — effectively that the priors are wrong. Second, the tests passed by the hybrid version, with the extent of rigidity varying with shocks, indicate that the mis-specification lies in the imposition of fixed price/wage rigidity across the whole economy.

4 Monte Carlo experiments

4.1 Bayesian estimation

In this section, using Monte Carlo experiments, we explore the consequences for Bayesian estimation of the New Keynesian model of alternative choices of the prior distributions. We take the Smets-Wouters model with their estimated parameters to be the true model and generate 1000 samples of data from it. These are treated as the observed data in the Bayesian estimation. We perform two experiments. In the first we set the true model so that the degree of wage and price stickiness parameters (ξ_w and ξ_p) are equal to the values estimated by Smets-Wouters (approximately 0.7), which we refer to as the high-rigidity (HR), typical New Keynesian, version. In the second, we set the true model so that both ξ_w and ξ_p are set to 0.05, and call it the flexible price (FP) version, which implies that the probability that prices and wages are fixed is close to zero — thereby eliminating its typical New Keynesian properties. In each experiment we use two sets of priors: high-rigidity priors (HR) and flexible price (FP) priors. For the HR priors we set the mean to be 0.5 and the standard deviation to be 0.1, the same distribution as Smets-Wouters. For the FP priors the mean and standard deviation are set to 0.05 and 0.1 with a minimum lower bound of 0.001 to ensure the model solves. In each case one of these is the false set. For all the other parameters whose values are not critical to whether the model is HR or FP, we used the same priors as in SW.

The results for the first experiment (HR true) are reported in columns 2 and 3 of Table 1. We show the average estimates for the 1000 samples of the key parameters ξ_w and ξ_p for each prior distribution together with the standard deviations of these 1000 estimates.

‘True’ Model	HR		FP	
	HR	FP	HR	FP
Degree of Wage Stickiness (ξ_w)	0.6873 (0.0452)	0.4113 (0.1631)	0.6482 (0.1327)	0.1246 (0.0892)
Degree of Price Stickiness (ξ_p)	0.7082 (0.0527)	0.1653 (0.1281)	0.6934 (0.0943)	0.0481 (0.0213)

Table 1: Average Estimates and their standard deviations (over 1000 samples) of the Wage and Price Stickiness Parameters for the NK and FP models with NK and FP Priors

In the HR case, the true parameter values for both ξ_w and ξ_p are approximately 0.7. The average Bayesian estimates based on the HR prior distribution are close to, and not significantly different from, their

true values. For the FP priors centred on 0.05 they are a long way below, and highly significantly different from, the true values of 0.7. Figure 1 shows the histograms of the ξ_w and ξ_p parameters for this HR case under both HR and FP priors. Under HR priors the parameters are centred approximately around the true value of 0.7. Under FP priors the parameters are centred approximately around 0.1; but a large number of the estimates are spread above this.

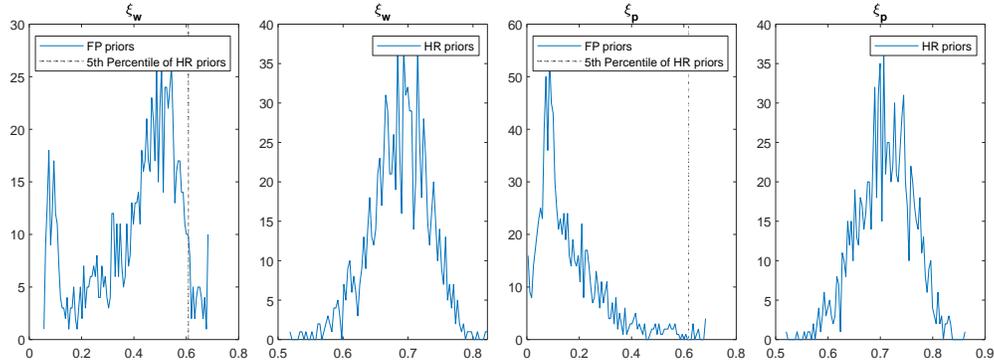


Figure 1: Histograms of the rigidity of wage (ξ_w) and price (ξ_p) coefficients and under HR and FP priors — where HR is the "true" model

The corresponding results for the second experiment where FP is the true model are reported in columns 4 and 5 of Table 1. For the FP priors the estimate of ξ_p is close to, and not significantly different from its "true" value of 0.05. The estimate of ξ_w is further from 0.05, but still not significantly different. For the HR priors the estimates of both parameters are close to their prior means of 0.7, but they are significantly far from their true values of 0.05.

Figure 2 shows the histograms of the ξ_w and ξ_p parameters for this FP model version under both HR and FP priors. With FP priors the histograms are centred close to 0.05. With HR priors the distributions of both ξ_w and ξ_p are centred around 0.7, far from the true values of 0.05.

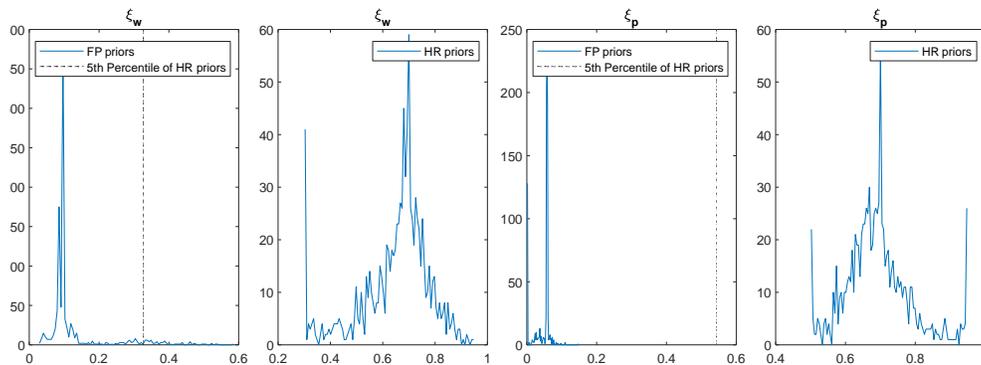


Figure 2: Histograms of the rigidity of wage (ξ_w) and price (ξ_p) coefficients and under HR and FP priors — where FP is the "true" model

One response to these results might be that they are due to excessively tight prior distributions. Perhaps greater flexibility could allow Bayesian estimation to give less biased results? Accordingly we show how the results remain largely unaffected as we allow the mean of the prior distributions to vary while still differentiating the priors meaningfully from the truth. Essentially what we find is that the priors act powerfully to distort the posterior estimates of price rigidity away from the true values. Nor do we find that substituting the mean for the mode alters our findings; estimating the mean rather than the mode is far more time-consuming but gives essentially the same results (as widely noted, e.g. by SW, 2007).

In the following exercise we simulated the HR and FP models, then treated each simulation as the data and estimated the parameters with a set prior mean. Table 2 shows the results when the HR model is treated as the true model and Table 3 when the FP model is treated as the true model. What is clear is that as the prior mean changes from low to high values the mean of the posterior mode of the rigidity parameters also increases.

Prior Mean	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Prior Stdev	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
mean(ξ_w)	0.4351	0.5490	0.6797	0.6685	0.6714	0.6986	0.7392	0.7708	0.7917
std(ξ_w)	0.1422	0.1558	0.0961	0.0574	0.0398	0.0405	0.0517	0.0515	0.0524
mean(ξ_p)	0.2666	0.3859	0.5049	0.6064	0.6763	0.7265	0.7681	0.7673	0.8264
std(ξ_p)	0.1551	0.1326	0.0762	0.0523	0.0464	0.0467	0.0532	0.0620	0.0382

Table 2: Estimating the mode with shifting priors on HR data

Prior Mean	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Prior Stdev	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
mean(ξ_w)	0.1063	0.1541	0.1617	0.2331	0.4575	0.5484	0.7060	0.7145	0.7687
std(ξ_w)	0.0162	0.0695	0.0626	0.1008	0.1117	0.1315	0.1349	0.1177	0.0705
mean(ξ_p)	0.0979	0.1334	0.2889	0.4270	0.6478	0.6039	0.7128	0.7071	0.8229
std(ξ_p)	0.0146	0.1201	0.1354	0.1175	0.0686	0.1086	0.1067	0.0854	0.0671

Table 3: Estimating the mode with shifting priors on FP data

These two experiments show with startling clarity how the choice of prior distribution affects the posterior estimates. The most striking result, which holds in both experiments, is that the posterior estimates are completely dominated by the prior distributions. Whether the data are generated by an HR or an FP model is immaterial as here the data play little role. It might be argued that this is what Bayesian econometrics aims to achieve, i.e. incorporate prior beliefs. The danger, of course, is that it will be inferred that the model is correct no matter how flawed it may be. This is why we have urged in several papers that Bayesian estimated models be tested.

4.2 ML and Indirect estimation

If the use of Bayesian estimation is suspect, what other method of estimation might be preferable? We compare two classical estimators: ML and Indirect estimation. As noted above, the use of Bayesian estimation was in part a response to the deficiencies of ML estimation. ML estimation — which can also be interpreted as Bayesian estimation with uninformative, uniform priors — seeks to choose parameter values that give the best in-sample forecasting performance by the model. This can produce highly biased parameter estimates, especially if the model is mis-specified; the estimator compensates for the mis-specification by distorting the parameters, thereby improving the forecasts.

In contrast, Indirect estimation chooses the model parameters to generate data from the structural model that gives estimates of an auxiliary model closest to those using the observed data. In a recent paper Le et al. (2016) carried out small sample Monte Carlo experiments which showed that the Indirect estimator has low bias and the associated Indirect test — based on the significance of differences between estimates of the parameters of the auxiliary model from data simulated from the structural estimates and the observed data

— has very high power against a mis-specified model such as the FP version of the NK model. The ML estimator by contrast was highly biased and had no power against a mis-specified model.

We now go on to apply ML and II here to the estimation of the price rigidity parameters in our Monte Carlo experiment. We produce 1000 simulations from both the HR and FP models, then estimate only the price rigidity parameters (keeping all other parameters fixed) treating each simulation as the data. What we find is consistent with those of Le et al. (2016). Table 4 shows that the ML estimate is seriously biased, while the II bias is very small. When we treat the HR model as the true model we find II has very low bias of approximately 1-2%, compared to 4-13% for ML. When the FP model is treated as the truth the difference in bias is much greater. The bias for II is approximately 1-5% compared to the massive bias of 1200-1700% for the ML estimates.¹

	True Model							
	HR				FP			
	ML		II		ML		II	
	ξ_w	ξ_p	ξ_w	ξ_p	ξ_w	ξ_p	ξ_w	ξ_p
Actual	0.6881	0.6978	0.6881	0.6978	0.0500	0.0500	0.0500	0.0500
Mean	0.7781	0.6687	0.6958	0.7107	0.8806	0.6509	0.0506	0.0524
St. dev	0.0555	0.2140	0.1140	0.1192	0.1113	0.2669	0.0286	0.0296
Bias (%)	13.0874	4.1733	1.1261	1.8454	1661.2000	1201.8000	1.2000	4.8000

Table 4: Bias of ML and II estimators

5 Conclusions

In this paper we have explored the estimation of price rigidity in macro models; the degree of rigidity is central to the behaviour of the economy in these models and there has been a long-running debate in macroeconomics about the appropriateness of assuming price rigidity as opposed to flexible prices. Yet we observe the widespread use of Bayesian estimation of macro models in which price rigidity is treated as a prior; we have investigated the consequences of this method for the resulting estimates of the price rigidity parameter. Our central finding is that in Bayesian estimation of the New Keynesian model the choice of prior distribution tends to distort the posterior estimates of this parameter, causing it to be substantially biased. A further result is that Maximum Likelihood estimates of it are also highly biased and that Indirect estimates have much lower bias.

The broader significance of these findings is that the Bayesian estimation of macro models may give very misleading results by placing too much weight on prior information compared to observed data and that a better method may be Indirect estimation. While our extended example has focused narrowly on the price/wage rigidity parameters which are of central importance in DSGE models, the conclusion applies to all the parameters of the model, which all have some bearing on the welfare assessment of policy regimes within the economy; all should be estimated with minimum bias, which can be achieved in small samples by indirect estimation — as found already by Le et al. (2016) across all the parameters of a DSGE model. The reason this is an important finding is the widespread use of Bayesian estimation in macroeconomics which has been facilitated by Dynare. This has resulted in an implicit consensus in favour of the New Keynesian model with highly sticky wages and prices, in spite of its rejection by indirect inference tests in favour of a hybrid model whose rigidity varies with the evolution of shocks — as found recently on US data by Le et al. (2019). The danger for macroeconomics is that this consensus becomes an orthodox opinion that is not supported by scientific evidence. Eventually, of course, theories not supported by the evidence will be rejected, much as the Great Depression overturned classical macroeconomics. Such overturning is however bad for the reputation of economics as a science. Rather than protect a theory by biasing estimation results in its favour — for example, through using strong priors — it is better to submit theories to ongoing tests of their consistency with the data.

¹To check whether this large ML bias could be due to the shocks being highly persistent (and so possibly mimicking the effects of price/wage rigidity) we redid the experiment setting the shock persistence to 0 for the FP model. This reduced the ML bias to 861% for ξ_w and 196% for ξ_p , still much higher than II.

6 References

Adjemian, S., Bastani, H., Juillard, M., Karamé, F., Maih, J., Mihoubi, F., Mutschler, W., Perendia, G., Pfeifer, J., Ratto, M., and Villemot, S. (2011), "Dynare: Reference Manual, Version 4." Dynare Working Papers, 1, CEPREMAP

Chari, V. V., Kehoe, P. J. and McGrattan, E.R. (2009) "New Keynesian Models: Not Yet Useful for Policy Analysis." *American Economic Journal: Macroeconomics*, 1(1): 242-66.

Christiano, L., Eichenbaum, M. and Evans, C. (2005) "Nominal Rigidities and the Dynamic effects of a shock to Monetary Policy." *Journal of Political Economy*, 113: 1-45.

Evans, G. W. and Honkapohja, S. (2005) "Interview with Thomas J. Sargent." *Macroeconomic Dynamics*, 9(4): 561-583.

Ireland, P. (2011) "A New Keynesian Perspective on the Great Recession." *Journal of Money, Credit and Banking*, 43(1): 31-54

Kydland, F., and Prescott, E.(1982) "Time to build and aggregate fluctuations." *Econometrica*, 50: 1345-70.

Le, M., Meenagh, D. and Minford, P. (2019) "State-dependent pricing turns money into a two-edged sword." Cardiff Economics Working Paper no. E2019/15, Cardiff University, Cardiff Business School, Economics Section.

Le, M., Meenagh, D., Minford, P., and Wickens, M. (2011) "How much nominal rigidity is there in the US economy? Testing a new Keynesian DSGE model using indirect inference." *Journal of Economic Dynamics and Control*, 35 (12): 2078-2104.

Le, V., Meenagh, D., Minford, P., Wickens, W. and Xu, Y. (2016) "Testing Macro Models by Indirect Inference: A Survey for Users." *Open Economies Review*, 27: 1-38.

Lucas, R.E. (1976) "Econometric Policy Evaluation: a Critique." *Carnegie-Rochester Conference Series on Public Policy*, 1: 19-46.

Meenagh, D., Minford, P., Wickens, M. and Xu, Y.(2019) "Testing DSGE Models by Indirect Inference: a Survey of Recent Findings", *Open Economies Review*, 30: 593-620.

Minford, P., Xu, Y. and Zhou, P. (2015) "How Good are Out of Sample Forecasting Tests on DSGE Models?" *Italian Economic Journal: A Continuation of Rivista Italiana degli Economisti and Giornale degli Economisti*, 1(3): 333-351.

Sims, C..A. (1980) "Macroeconomics and reality." *Econometrica*, 48: 1-48.

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

White, L. H. (2003). "The Methodology of the Austrian School Economists (revised ed.)". Ludwig von Mises Institute.

Zhou, P. and Dixon, H. (2019) "The Determinants of Price Rigidity in the UK: Analysis of the CPI and PPI Microdata and Application to Macrodata Modelling." *The Manchester School*, 87(5): 640-677.