

Revealing priors from posteriors

Masako Ikefuji

OSIPP, Osaka University, Japan

Jan R. Magnus

Department of Econometrics and Data Science,
Vrije Universiteit Amsterdam and Tinbergen Institute, The Netherlands

Takashi Yamagata

Department of Economics and Related Studies, University of York, UK
and ISER, Osaka University, Japan

June 1, 2022

Abstract: A Bayesian typically uses data and a prior to produce a posterior. We shall follow the opposite route, using data and the posterior information to reveal the prior. We apply the theory by inferring the Bank of England's priors when forecasting UK inflation.

JEL Classification: C11, C13, C18

Keywords: Revealed prior, judgment

Corresponding author: Jan R. Magnus, E-mail address: jan@janmagnus.nl
(J.R. Magnus)

Email addresses: ikefuji@osipp.osaka-u.ac.jp (Ikefuji),
jan@janmagnus.nl (Magnus),
takashi.yamagata@york.ac.uk (Yamagata).

1 Introduction

Everybody has priors, Bayesians and non-Bayesians alike. The priors may be vague and difficult to make explicit, but they are there and they may be important. The purpose of this paper is to show that we can make priors explicit from our knowledge of the data and the posterior.

Imagine a group of people (the “committee”) with a collective prior, perhaps based on knowledge and experience, perhaps on political beliefs, perhaps on short-term profit. The committee meets privately and we have no information about their discussions. But we do have scientific data (official “objective” statistics and scientific results) and we do have access to the committee’s published predictions or policy recommendations, which they present to the public. In other words, we have the data and the posterior, but not the prior which the committee does not reveal and possibly may not even be able to formulate or quantify. Can we recover the prior from the data and the posterior? Yes, this is indeed possible as we shall demonstrate.

The idea of reversing Bayesian thought and — rather than obtain a posterior from data and prior — recover the prior from data and posterior, does not seem to have received much attention, with the exception of somewhat related work by Jarociński and Marcet (2019) and Galvão et al. (2021). The current paper attempts to fill this gap.

The list of possible applications is endless. A political party uses scientific data and publishes reports. From these two sources we can recover their priors. Do these conform to the party program? Scientists use data and write papers. The results in these papers may well be influenced by prior beliefs or non-scientific prejudices. Can this influence be quantified? Such questions can, in principle, be answered by the theory developed in this paper.

To illustrate, we report (very) briefly on inflation forecasts by the Bank of England, especially under the uncertainty about Brexit and the Covid crisis. A more thorough investigation of this case is provided in the online supplementary material.

In Section 2 we analyze how to recover the prior from the data and the posterior within the framework of the normal distribution, both for the multi-parameter and the single-parameter case. Section 3 discusses our application. Section 4 concludes and suggests various extensions.

2 From posterior to prior under normality

We consider a parameter vector of interest β and suppose that data are generated from a normal distribution

$$y|\beta \sim N(X\beta, \Omega), \quad (1)$$

where X is a given $n \times k$ matrix of rank k and Ω is a positive definite $n \times n$ matrix. A non-Bayesian frequentist would estimate β using the generalized least-squares (GLS) estimator (which is also the maximum likelihood estimator)

$$b_0 = (X'\Omega^{-1}X)^{-1} X'\Omega^{-1}y \quad (2)$$

with variance

$$\Sigma_0 = (X'\Omega^{-1}X)^{-1}. \quad (3)$$

A Bayesian, on the other hand, would wish to take prior knowledge about β into account. Suppose this prior information is given by

$$\beta \sim N(b_1, \Sigma_1), \quad (4)$$

where Σ_1 is positive definite. Then the posterior distribution of β is

$$\beta|y \sim N(b_2, \Sigma_2), \quad (5)$$

where

$$b_2 = Wb_1 + (I_k - W)b_0, \quad \Sigma_2 = (\Sigma_1^{-1} + \Sigma_0^{-1})^{-1}, \quad (6)$$

and $W = \Sigma_2\Sigma_1^{-1}$ is a $k \times k$ weight matrix.

Although W is, in general, not symmetric, its eigenvalues are real and lie between zero and one. In fact, letting $Z = \Sigma_1^{1/2}\Sigma_0^{-1}\Sigma_1^{1/2}$ with eigenvalues $\lambda_i(Z) > 0$ ($i = 1, \dots, k$), we see that

$$\lambda_i(W) = \lambda_i(\Sigma_1^{-1/2}\Sigma_2\Sigma_1^{-1/2}) = \frac{1}{\lambda_i(\Sigma_1^{1/2}\Sigma_2^{-1}\Sigma_1^{1/2})} = \frac{1}{1 + \lambda_i(Z)}. \quad (7)$$

Note that when the prior becomes uninformative, that is when $\Sigma_1^{-1} \rightarrow 0$, then $b_2 \rightarrow b_0$ and $\Sigma_2 \rightarrow \Sigma_0$. This is well-established basic Bayesian theory.

Now consider the opposite situation where the data and the posterior are available but not the prior. Can we reveal the prior from the data and the posterior? In general we can, and in the special case of normality we obtain the prior moments as

$$b_1 = W^{-1}b_2 + (I_k - W^{-1})b_0, \quad \Sigma_1 = (\Sigma_2^{-1} - \Sigma_0^{-1})^{-1}, \quad (8)$$

with

$$W^{-1} = \Sigma_1 \Sigma_2^{-1} = \Sigma_0 (\Sigma_0 - \Sigma_2)^{-1}, \quad (9)$$

which assumes implicitly an upper bound to the posterior variance, namely $\Sigma_2 < \Sigma_0$ in the usual sense that $\Sigma_0 - \Sigma_2$ is positive definite. The prior mean is thus a “weighted average” of b_2 and b_0 , but the eigenvalues of W^{-1} do not lie between zero and one. In fact $\lambda_i(W^{-1}) = 1 + \lambda_i(Z) > 1$ and $\lambda_i(I_k - W^{-1}) = -\lambda_i(Z) < 0$ for all $i = 1, \dots, k$.

The restriction $\Sigma_2 < \Sigma_0$ does not play a role in the usual Bayesian framework where we go from data plus prior to posterior, because the underlying variances Σ_0 and Σ_1 are unrestricted (apart from being positive definite) and Σ_2 will automatically satisfy the restriction. But it does play a role when we go from data plus posterior to prior, because now the restriction is not automatically satisfied.

In the special but important case where we have only one parameter β of interest, we write σ_0^2 , σ_1^2 , and σ_2^2 instead of Σ_0 , Σ_1 , and Σ_2 . From the data (without a prior) we obtain an unbiased estimator of β : $b_0 \sim N(\beta, \sigma_0^2)$. If we add a prior $\beta \sim N(b_1, \sigma_1^2)$, then we obtain the posterior $\beta \sim N(b_2, \sigma_2^2)$, where

$$b_2 = \frac{\sigma_0^2 b_1 + \sigma_1^2 b_0}{\sigma_0^2 + \sigma_1^2}, \quad \sigma_2^2 = \frac{\sigma_0^2 \sigma_1^2}{\sigma_0^2 + \sigma_1^2}. \quad (10)$$

In the reversed case that we are interested in, we have an unbiased estimator $b_0 \sim N(\beta, \sigma_0^2)$ from the data and the posterior moments of $\beta \sim N(b_2, \sigma_2^2)$. From these two ingredients we obtain the prior as $\beta \sim N(b_1, \sigma_1^2)$, where

$$b_1 = \frac{\sigma_0^2 b_2 - \sigma_2^2 b_0}{\sigma_0^2 - \sigma_2^2}, \quad \sigma_1^2 = \frac{\sigma_0^2 \sigma_2^2}{\sigma_0^2 - \sigma_2^2}, \quad (11)$$

under the restriction that $\sigma_2^2 < \sigma_0^2$.

3 Inflation forecasts by the Bank of England

The Bank of England’s (BoE) primary responsibility is to keep UK inflation at 2% and the Monetary Policy Committee’s (MPC) task is to decide what monetary policy action it should take to achieve this goal. Every quarter, the BoE publishes its *Monetary Policy Report* (until 2019/Q4 called *Inflation Report*), in which the density forecasts of inflation, economic growth, and the employment rate are provided.

We have chosen the four-quarter (i.e. one-year) ahead density forecast of CPI inflation as the posterior of the BoE and we focus on the quarters before and after two recent events which shocked the UK economy: the referendum

outcome for Brexit in June 2016 and the first lockdown for Covid-19 in March 2020. Accordingly we consider five quarters (2016/Q1 and Q2; and 2020/Q1, Q2, and Q3) in which the density forecasts are published in reports of the BoE.

3.1 Posterior and data

The posterior means b_2 and standard deviations σ_2 for the five quarters are reported in the Table 1.

Table 1: UK inflation: posterior, data, and revealed prior

Year/Quarter	Posterior		Data		Prior	
	b_2	σ_2	b_0	σ_0	b_1	σ_1
<i>Referendum</i>						
2016/Q2	1.52	1.34	1.63	2.6	1.48	1.56
2016/Q3	2.03	1.34	2.36	2.6	1.91	1.56
<i>Lockdown</i>						
2020/Q1	1.53	1.34	1.76	2.6	1.45	1.56
2020/Q2	0.50	2.02	1.42	2.6	-0.90	3.21
2020/Q3	1.55	2.02	1.59	2.6	1.49	3.21

The data are based on The Bank of England Survey of External Forecasters. The BoE reports a summary of each survey in their quarterly *Monetary Policy Report*, and this leads to a mean b_0 and a standard deviation σ_0 .

3.2 The revealed prior

Given the posterior and the data, we can reveal the unknown prior.

3.2.1 The referendum

The Brexit referendum was held on 23 June, 2016. Before the referendum, the general expectation was that the UK population would vote to remain in the European Union, and the forecasts by the BoE were made under the assumption that the “remainers” would win. The quarter labeled 2016/Q2 is associated with the report by the BoE of May 2016, one month before the referendum, containing forecasts made in 2016/Q2 about the inflation one year later in 2017/Q2. The BoE’s prior mean b_1 is about 1.48 and the prior

standard deviation σ_1 is relatively small (around 1.6). The BoE considered that inflation would rise automatically to the 2% target by mid-2018, and was reluctant to impose monetary policy which might cause too rapid inflation. The priors changed dramatically after the referendum, as we can see in the quarter labeled 2016/Q3. The outcome of the referendum was unexpected and not assumed in the previous report published in 2016/Q2. So, the priors had to adjust. In addition, after the referendum the exchange rate fell sharply and the outlook for growth in the short to medium term weakened markedly.

3.2.2 The lockdown

Lockdown Covid measures in the UK came legally into force on March 26, 2020, ten days after being announced by Boris Johnson. One month before the lockdown, in the February 2020 *Monetary Policy Report*, the MPC decided to maintain the base rate at 0.75%. After the national lockdown, domestic and world economic conditions deteriorated sharply. To respond to the new situation, the MPC reduced the base rate to 0.1% on 19 March 2020, right after the first cut to 0.25% on 11 March 2020. The inflation rate declined to 1.5% in March. The prior for the BoE reflects its pessimistic feelings and its inability to make accurate forecasts. In the next quarter, 2020/Q3, the BoE substantially adjusted its prior inflation forecast upwards. After declining sharply to 0.6% in 2020/Q2, the BoE expected that inflation would fall further due to the low energy prices and the temporary cut in value-added tax for the hospitality industry, and that inflation would rise during 2021, as the impacts of low energy price and the value-added tax cut would fade.

Summarizing, we have found that BoE's prior is highly stable and inflexible over the five quarters, except the outlier during the initial quarter of the national lockdown. Except in 2020/Q2, BoE's b_1 ranges between 1.45 and 1.91, which is just below the institutional inflation target level of 2%. The BoE apparently has great confidence in its ability to achieve the target, except during the Covid-19 lockdown. The online appendix provides a much more detailed analysis.

4 Conclusions and further work

In this paper we have tried to reveal the prior, given information about the data and the posterior. We do so in the context of the normal distribution. The normality assumption can be relaxed without changing the underlying theory, but the expressions would become less transparent.

There are at least two further closely related questions worth investigating. First, suppose that our observations stretch over several, say two, periods. In both periods we have data and posteriors, and we can recover the priors in periods 1 and 2. Dynamic consistency requires that the prior in period 2 equals the posterior in period 1. But is it? If it is, then the agent is consistent or rational in this Bayesian updating scheme. But if it isn't, then the agent is not consistent. One can easily imagine a situation where the agent remains too loyal to their original prior, which one may call “prior stubbornness” or “bunching”. This stubbornness may be politically motivated and is related to the theory of learning. It may continue until some bound has been reached (a tipping point), after which the prior is adjusted and moves to a new level. In fact, it should be possible to derive a measure for such stubbornness.

Second, we may consider the situation where there is not one but several, say two, agents over several, say two, periods. The data available to the two agents are the same, but their posteriors are not, and hence their priors are also different. We may then ask: under what conditions would their priors converge?

Acknowledgements

We are grateful to Peter Boswijk, Siem Jan Koopman, Gideon Magnus, Mark Steel, and Rick van der Ploeg for helpful discussions. Earlier versions of this paper were presented at the Vrije Universiteit Amsterdam (April 2020), the NESG Conference (May 2020), and at ISER, Osaka University (December 2020). This work was financially supported by the JSPS KAKENHI Grants Numbers JP20H01477 and JP21H00700 and the program of the Joint Usage/Research Center for Behavioral Economics at ISER, Osaka University.

Appendix: Supplementary data

Supplementary material related to this article can be found online at [TO FOLLOW].

References

- Galvão, A.B., A. Garratt, and J. Mitchell (2021). Does judgment improve macroeconomic density forecasts? *International Journal of Forecasting*, 37, 1247–1260.

Jarociński, M. and A. Marcet (2019). Priors about observables in vector autoregressions. *Journal of Econometrics*, 209, 238–255.