

This is a repository copy of *Revealing priors from posteriors with an application to inflation forecasting in the UK*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/201813/>

Version: Published Version

Article:

Ikefuji, Masako, Magnus, Jan and Yamagata, Takashi orcid.org/0000-0001-5949-8833
(2024) Revealing priors from posteriors with an application to inflation forecasting in the UK. *Econometrics Journal*. pp. 151-170. ISSN 1368-4221

<https://doi.org/10.1093/ectj/utad021>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Revealing priors from posteriors with an application to inflation forecasting in the UK

MASAKO IKEFUJI[†], JAN R. MAGNUS[‡] AND TAKASHI YAMAGATA[§]

[†]*OSIPP, Osaka University, 1-31 Machikaneyama, Toyonaka, Osaka, 560-0043, Japan.*
Email: Ikefuji-osipp@osaka-u.ac.uk

[‡]*Department of Econometrics & Data Science, Vrije Universiteit Amsterdam, School of Business and Economics, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands.*
Email: jan@janmagnus.nl

[§]*DERS, University of York, York YO10 5DD, UK and ISER, Osaka University, 6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan.*
Email: takashi.yamagata@york.ac.uk

First version received: 13 December 2022; final version accepted: 19 July 2023.

Summary: 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 then apply this theory to inflation forecasts by the Bank of England and the National Institute of Economic and Social Research in an attempt to get some insight into the prior beliefs of the policy makers in these two institutions, especially under the uncertainties about the Brexit referendum, the Covid-19 lockdown, and the Russian invasion of Ukraine.

Keywords: *Revealed prior, inflation, central bank.*

JEL codes: *C11, E31, E58.*

1. INTRODUCTION

Everybody has priors, Bayesians and nonBayesians 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, and to highlight some of the difficulties that may arise in this process.

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 knowledge about their discussions, but we do have data—official ‘objective’ statistics and scientific results—and we do have access to their 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 and we shall study the properties of the recovered prior in some detail.

We illustrate the theory by investigating the priors of the Bank of England (BoE) and the National Institute of Economic and Social Research (NIESR) when forecasting inflation, especially under the uncertainties about the Brexit referendum, the Covid-19 lockdown, and the Russian invasion of Ukraine.

The idea of reversing Bayesian thought and—rather than obtaining a posterior from data and prior—recovering the prior from data and posterior does not seem to have received much attention. 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 the revealed priors conform to the party programme? Scientists use data and write papers. The results in these papers may well be influenced by prior beliefs or nonscientific prejudices. Can this influence be quantified? Such and many other questions can, in principle, be studied by the theory developed in this paper.

In Section 2 we analyse how to recover the prior from the data and the posterior, first in general and then within the framework of the normal distribution. In Section 3 we consider the case when there is only one parameter of interest. In the remainder of the paper we present our application to inflation forecasts in the UK: Section 4 provides an introduction, Section 5 presents the posterior, Section 6 discusses the data, and Section 7 reveals the implied prior. Section 8 concludes.

2. FROM POSTERIOR TO PRIOR

Suppose we are given a joint density $f(y, \beta)$ of two sets of random variables: the observations y and the parameters β . From this joint density we can compute the conditional densities $f(y|\beta)$ (the likelihood) and $f(\beta|y)$ (the posterior), and the marginal densities $f(y)$ and $f(\beta)$ (the prior). These densities are related to each other through Bayes' identity

$$f(y|\beta)f(\beta) = f(\beta|y)f(y), \quad (2.1)$$

which can also be written as

$$f(\beta|y) = \frac{f(y|\beta)f(\beta)}{f(y)}, \quad f(y) = \int f(y|\beta)f(\beta) d\beta, \quad (2.2)$$

thus transforming the observations and the prior into the posterior. This is the usual path where we obtain posterior moments based on the data and the prior.

We propose to write Bayes' identity differently, namely as

$$f(\beta) = \frac{f(\beta|y)/f(y|\beta)}{g(y)}, \quad g(y) = \int \frac{f(\beta|y)}{f(y|\beta)} d\beta, \quad (2.3)$$

thus transforming the likelihood and the posterior into the prior. This unusual path allows us to obtain prior moments based on the data and the posterior. The idea is simple, but quite general. In particular, it is not constrained to univariate or normal settings, and it allows both conjugate and nonconjugate priors.

There is, however, a good reason why we may wish to restrict ourselves to conjugate priors. If we observe a process over time, then we would expect to learn more about the process as time passes and we would typically update our prior by using the posterior in one period as the prior in the next period. In such a situation it would not be credible if prior and posterior obey different probability laws.

The class of distributions allowing conjugate priors is quite restrictive. It includes, of course, the normal distribution which, in the case of symmetry, is often a reasonable simplification; see

e.g. Kiss et al. (2023). If there is no symmetry then the normal distribution is not easy to justify, but in our application symmetry seems reasonable.

In order to highlight the issue under simple conditions, we assume normality. Consider a parameter vector of interest β and suppose that the observations are generated from a normal distribution

$$y|\beta \sim N(X\beta, \Omega), \quad (2.4)$$

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

$$b_0 = (X'\Omega^{-1}X)^{-1} X'\Omega^{-1}y, \quad (2.5)$$

with variance

$$\Sigma_0 = (X'\Omega^{-1}X)^{-1}. \quad (2.6)$$

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

$$\beta \sim N(b_1, \Sigma_1), \quad (2.7)$$

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

$$\beta|y \sim N(b_2, \Sigma_2), \quad (2.8)$$

where

$$b_2 = Wb_1 + (I_k - W)b_0, \quad \Sigma_2 = (\Sigma_1^{-1} + \Sigma_0^{-1})^{-1}, \quad (2.9)$$

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

$$\lambda_i(W) = \lambda_i\left(\Sigma_1^{-1/2}\Sigma_2\Sigma_1^{-1/2}\right) = \frac{1}{\lambda_i\left(\Sigma_1^{1/2}\Sigma_2^{-1}\Sigma_1^{1/2}\right)} = \frac{1}{1 + \lambda_i(Z)}. \quad (2.10)$$

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 all 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

$$b_1 = W^{-1}b_2 + (I_k - W^{-1})b_0, \quad \Sigma_1 = (\Sigma_2^{-1} - \Sigma_0^{-1})^{-1}, \quad (2.11)$$

with

$$W^{-1} = \Sigma_1\Sigma_2^{-1} = \Sigma_0(\Sigma_0 - \Sigma_2)^{-1}, \quad (2.12)$$

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. This has practical consequences, as we shall see later.

At this point we have two objects that can be regarded as ‘data’: the observations y and the estimate b_0 . The observations y are what one would typically call the data and we shall not deviate from this custom. But in the process that we are considering, agents do not look at, y but at b_0 as their data. We have a parameter of interest β and the agents will look at b_0 , add their prior b_1 on β , and arrive at a posterior b_2 . To avoid confusion, we shall not refer to b_0 as the data, but as the *input*. In the absence of a prior, the agents accept the input as their only tool: output = input, that is, $b_2 = b_0$. But if the agent’s prior plays a role (which of course it does), then output \neq input and the difference is the prior, which we wish to reveal.

3. ONE PARAMETER OF INTEREST

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 observations (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

$$b_2 = \frac{\sigma_0^2 b_1 + \sigma_1^2 b_0}{\sigma_0^2 + \sigma_1^2}, \quad \sigma_2^2 = \frac{\sigma_0^2 \sigma_1^2}{\sigma_0^2 + \sigma_1^2}. \quad (3.1)$$

In the reversed case that we are interested in we have an unbiased estimator $b_0 \sim N(\beta, \sigma_0^2)$ based on the observations (our input), 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

$$b_1 = \frac{\sigma_0^2 b_2 - \sigma_2^2 b_0}{\sigma_0^2 - \sigma_2^2}, \quad \sigma_1^2 = \frac{\sigma_0^2 \sigma_2^2}{\sigma_0^2 - \sigma_2^2}, \quad (3.2)$$

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

Defining α_m and α_v implicitly by

$$b_2 = \alpha_m b_0, \quad \sigma_2^2 = \alpha_v \sigma_0^2, \quad (3.3)$$

we can rewrite (3.2) as

$$b_1 = \kappa_m b_0, \quad \sigma_1^2 = \kappa_v \sigma_0^2, \quad (3.4)$$

where

$$\kappa_m = \frac{\alpha_m - \alpha_v}{1 - \alpha_v}, \quad \kappa_v = \frac{\alpha_v}{1 - \alpha_v}, \quad (3.5)$$

measure how far the prior is removed from the input and their effect on the prior mean and variance, respectively. Note that α_m is unrestricted, but that α_v is restricted by $0 < \alpha_v < 1$.

The two fractions κ_m and κ_v capture the essence of our story. First consider κ_v , which relates to the prior variance. What matters here is whether κ_v is small (strong prior information) or large (weak prior information). This depends only on α_v , not on α_m . When α_v is close to one, then the variance σ_0^2 in the input and the variance σ_2^2 in the posterior (the output) are approximately equal,

so that the prior has only a small effect. This is represented by a large value of κ_v and hence a large value of the prior variance σ_1^2 . The prior is then uninformative. But when α_v is close to zero, then the input variance and the posterior variance are not close, and the prior has a big effect. This is represented by a small value of κ_v and hence a small value of the prior variance σ_1^2 . The prior is then informative.

The situation is quite different with κ_m . What matters here is not whether κ_m is small or large, but whether κ_m is close to one or not. This will depend on both α_m and α_v . It is clear that $\kappa_m = 1$ when $\alpha_m = 1$, irrespective of the value of α_v . Writing

$$1 - \kappa_m = \frac{1 - \alpha_m}{1 - \alpha_v}, \quad (3.6)$$

we see that the deviation of κ_m from one depends on the deviation of α_m from one *relative* to the deviation of α_v from one. When α_m is close to one, but α_v is not, then the mean b_0 in the input and the mean b_2 in the posterior are approximately equal, but the variance σ_0^2 in the input and the variance σ_2^2 in the posterior are not approximately equal. In that case $\kappa_m \approx 1$ and the prior mean agrees with the input and the posterior. But when α_v is close to one, but α_m is not, then the variances σ_0^2 and σ_2^2 are approximately equal, but the means b_0 and b_2 are not. In that case κ_m is large (in absolute value). Naturally, for people with a very strong prior ($\sigma_1^2 \approx 0$) we have $\alpha_v \approx 0$, and hence $\kappa_m \approx \alpha_m$ and $b_1 \approx b_2$.

4. THE BANK OF ENGLAND'S INFLATION FORECASTS FOR THE UK

Density forecasts provide richer information on forecast uncertainties than point forecasts, and this simple fact has been increasingly put to practice by decision-makers, professional forecasters, and academic researchers when forecasting macroeconomic variables. For example, the Monetary Policy Committee (MPC) of the BoE has produced quarterly reports on Gross Domestic Product (GDP) growth and inflation since 1996, and density forecasts are typically used in these reports in order to explain the employed monetary policies.

Central banks (including the BoE) and professional forecasters do not follow the model-based forecast densities mechanically—they will add a final touch (possibly more than a mere touch) based on their subjective judgement. Following the approach of McNees (1990) and Turner (1990) for point forecasts, Galvão et al. (2021) investigated whether the subjective adjustment to the (mechanical) density forecast improves the forecast performance and concluded that ‘density forecasts from statistical models prove hard to beat’. Our purpose is not to assess the impact of adjustment on forecast performance, but to examine the process by which decision-makers, such as the MPCs of the BoE, determine density forecasts. Applying the Bayesian framework, in line with the discussion in Winkler (1968), it is assumed that the decision-maker (i.e., the BoE) takes all information available at the time of the decision (i.e., at the MPC meeting) into account. This information includes the latest economic data, internal and external professional forecasts, and the household survey, among others. The process by which the BoE transforms such a large amount of information into published density projections is complex and inaccessible. Even if it were accessible, modelling how this information feeds into the decision-making process would be a huge challenge.

We shall consider the forecast decision process of the BoE in a simplified and stylized manner: the BoE uses the information available at the time of decision-making and produces mechanical forecasts using a representative econometric model. This mechanical (nonBayesian) forecast is

our input b_0 . The committee considers b_0 and adds a final touch (the prior b_1) to arrive at the published forecast (the posterior b_2). The prior comprises the beliefs about the forecast based on all available information other than the mechanical prediction outcome, and it is these beliefs that we want to estimate by applying our proposed methodology.

As a mechanical forecasting model we choose the Phillips curve forecast regression model (Phillips 1958). Central banks make monetary policy decisions based on estimated current and predicted future economic conditions and inflation rates. One may therefore interpret the central bank's role as an attempt to predict how the Phillips curve, which relates the inflation rate to economic conditions, will change.

We investigate two aspects of the BoE's prior forecast distribution. The first aspect concerns changes in the distribution of prior forecasts over time, in particular how the distribution is affected by exogenous shocks. We consider three recent major shocks to the UK economy—the Brexit referendum, the Covid-19 lockdown, and the Russian invasion of Ukraine—and we investigate the changes in the prior distributions by comparing the prior distributions before and during the shocks.

The second aspect concerns the possible difference between the prior distributions of forecasts by two different forecasters, especially after an economic shock. Our second forecaster for this analysis is the NIESR, an independent institution with a high reputation in the UK. We assume that the BoE and the NIESR provide density forecasts based on *the same* input b_0 . If the input is the same then a difference in the outputs (the posteriors) can only be caused by a difference in the priors.

5. THE POSTERIOR DISTRIBUTION

The BoE's primary responsibility is to keep UK inflation at 2%, and the MPC's task is to decide what monetary policy action to take in order to achieve this goal. Since it will take about two years for monetary policy to have its full effect on the economy, the MPC needs to forecast the development of the economy in general and inflation in particular. Every quarter the BoE publishes its *Monetary Policy Report* (until 2019/Q4 called *Inflation Report*), in which the density forecasts of the inflation rate, economic growth rate, and employment rate are provided. These reports are published in February, May, August, and November of each year.

We have chosen the four-quarter (one-year) ahead density forecast of the Consumer Prices Index (CPI) inflation as the posterior of the BoE.¹ We assume that (almost) all information up to and including the previous quarter is available, so that for example the May report is based on full information of Q1. The information available at the time of publication varies slightly for different variables. For example, the February release contains the monthly CPI up to December of the previous year, while the unemployment rate is only available up to November. We believe it is reasonable to assume that the use of the (nonreal-time) published CPI and unemployment rate provides a sufficiently accurate approximation to the analysis based on real-time data.²

Our objective is to reveal the impact on the BoE's prior on three recent events that have significantly shocked the UK economy: the referendum outcome for Brexit in 2016/Q2 (23 June), the first lockdown for Covid-19 in 2020/Q1 (23 March), and the Russian invasion of Ukraine in

¹ The BoE publishes two types of forecasts: one assumes that interest rates will follow the market expectation, and the other assumes that rates will remain at their current level. We have chosen for the forecast with market interest rates.

² It would have been better to use real-time (vintage) data. However, we have not found a BoE data source that covers the sample period we are considering.

2022/Q1 (24 February). To this end we examine the density forecasts for a given quarter, made in the quarter at which the shock occurred and the two preceding quarters. Thus, in the case of the referendum, we consider density forecasts for 2017/Q2 made in 2016/Q2 (four quarters ahead), 2016/Q1 (five quarters), and 2015/Q4 (six quarters); in the case of the lockdown, we consider density forecasts for 2021/Q1 made in 2020/Q1, 2019/Q4, and 2019/Q3; and in the case of the Russian invasion, we consider density forecasts for 2023/Q1 made in 2022/Q1, 2021/Q4, and 2021/Q3. All forecasts are extracted from the data files associated with Bank of England (2023, February) and are based on the so-called two-piece normal distribution.³ Among the quarters considered, skewness is absent or very mild, with the exception of 2022/Q2. Thus it seems reasonable to assume that a normal approximation $N(b_2, \sigma_2^2)$ of the density forecast (or the ‘fan chart’) with published means and standard deviations is sufficiently accurate.⁴

The NIESR publishes its economic density forecasts every quarter in *Prospects for the UK Economy*, and we consider these forecasts as NIESR’s posteriors. The density forecasts are produced using the National Institute Global Econometric Model (NiGEM) together with the institution’s judgement; see the source to figure 5 in Hantzsche and Young (2020b). The density forecasts are presented as fan charts in the quarterly *National Institute UK Economic Outlook*, published in February, May, August, and November, and we use these charts to approximate the mean and variance of the corresponding posterior normal distribution.⁵ As with the BoE reports, we assume that (almost) all information up to and including the previous quarter is available.

The posterior means and standard deviations of the BoE and the NIESR for the nine quarters of interest are reported in Table 1. The BoE produces on average lower forecasts than the NIESR (about 8% lower) and the standard deviation is also lower (about 6%). But there are large differences, especially in 2016/Q2 (where the NIESR predicted that the outcome of the referendum would have a much bigger impact on inflation than the BoE predicted: 3.06 versus 1.99) and 2022/Q1 (where after the Russian invasion in Ukraine the opposite happened: 6.98 versus 9.31). One would expect that the precision increases as the time horizon decreases, and this is generally the case in each panel with the exception of 2020/Q1 (lockdown), where the BoE produced a much lower inflation prediction than the NIESR (0.50 versus 0.95), but with much larger uncertainty (1.84 versus 1.30). We conclude that there are relatively large differences between the posterior distributions of the two institutions.

6. THE INPUT DISTRIBUTION

Central banks make monetary policy decisions based on current economic conditions and inflation forecasts. Without too much exaggeration one may interpret the central bank’s task as determining how the Phillips curve, which relates the inflation rate to economic conditions, will change. A speech by Catherine L. Mann, one of the BoE’s MPC members, delivered at the September 2022

³ The two-piece normal (split normal, binormal, double-Gaussian) distribution results from joining at the mode the corresponding halves of two normal distributions with the same mode, but different standard deviations. This distribution thus allows skewness; see Wallis (2014) for more details.

⁴ We partially accommodate the possible presence of skewness by taking the median rather than the mean when approximating the normal distribution. This is easy for the BoE because they publish both mean and median. It is more difficult for the NIESR because they publish a fan chart, no numbers, from which we abstracted the median. See the Replication Package for more details on the extraction of medians and standard deviations.

⁵ The reports are called *Prospects for the UK Economy* until November 2020; see Kirby et al. (2016a, b, c), Hantzsche and Young (2019, 2020a, b), Lenoël et al. (2021), Macqueen et al. (2022), and National Institute of Economic and Social Research (2022, Spring).

Table 1. UK inflation: the posterior (published inflation forecasts).

Year/Quarter	BoE		NIESR	
	b_2	σ_2	b_2	σ_2
2017/Q2 inflation forecast distributions made in:				
2015/Q4	1.51	1.37	1.69	1.83
2016/Q1	1.52	1.34	1.00	1.77
2016/Q2 (Referendum)	1.99	1.27	3.06	1.40
2021/Q1 inflation forecasts distributions made in:				
2019/Q3	1.67	1.37	1.99	1.64
2019/Q4	1.53	1.34	2.12	1.59
2020/Q1 (Lockdown)	0.50	1.84	0.95	1.30
2023/Q1 inflation forecasts distributions made in:				
2021/Q3	3.26	1.71	3.58	1.67
2021/Q4	5.29	1.68	5.85	1.76
2022/Q1 (Invasion)	9.31	1.55	6.98	1.53

Note: In 2020/Q1 the values for b_2 and σ_2 for the BoE are missing. For this quarter we take b_2 from table 1.A for the 2021 average in Bank of England (2020, May), and for σ_2 we employ the value for 2020/Q2.

meeting, highlighted the BoE's use of monetary policy instruments to control inflation levels by influencing the inflation expectations of economic agents and economic conditions. As policies do not take effect until one or two years after implementation, it is necessary to make the best possible forecasts. We regard the distribution of inflation forecasts based on the Phillips curve as the input for the MPC when they discuss the forecast distribution to be published in their reports (i.e., the posterior).

To estimate the Phillips curve forecast model we collected data on the quarterly UK consumer price index CPI_t and the unemployment rate x_t . The quarterly CPI series from 1970/Q4 to 2022/Q2 are computed based on the monthly CPI index released by the Office of National Statistics, which is a merge of the 1970–2003 and 1989–2022 series downloaded in March 2023. CPI_t is then computed as the average of the monthly CPI index over the three months of the corresponding quarter t . The quarterly unemployment rate series x_t from 1971/Q1 to 2022/Q2 is downloaded from the Office of National Statistics site, also in March 2023.

Following Stock and Watson (2009, 2020), we employ the autoregressive distributed-lag model version of the Phillips curve. Defining the difference operator as $\Delta a_t = a_t - a_{t-1}$, the h -quarter-ahead forecast regression is given by

$$\pi_{t+h}^h - \pi_t = \alpha + \sum_{\ell=0}^q \gamma_{\ell} \Delta \pi_{t-\ell} + \sum_{\ell=0}^q \beta_{\ell} x_{t-\ell} + \epsilon_{t+h}, \quad (6.1)$$

where $\pi_t = 400 \ln(CPI_t/CPI_{t-1})$ and $\pi_{t+h}^h = h^{-1} \sum_{j=1}^h \pi_{t+j}$ is the log-approximation of the h -quarter inflation at an annual rate. Observe that π_{t+4}^4 is the four-quarter (year-on-year) inflation rate. The unemployment rate x_t is a slack variable. We assume that the ϵ_{t+h} are independent and identically distributed as $N(0, \sigma_{\epsilon}^2)$. The lag-order q is empirically determined by information

criteria. The forecast of π_{t+h}^h at time t is then given by

$$\pi_{t+h|t}^h = \alpha + \pi_t + \sum_{\ell=0}^q \gamma_{\ell} \Delta \pi_{t-\ell} + \sum_{\ell=0}^q \beta_{\ell} x_{t-\ell}, \quad (6.2)$$

and, assuming that the model also holds in the forecasting period, we have

$$\pi_{t+h}^h - \pi_{t+h|t}^h \sim N(0, \sigma_h^2), \quad (6.3)$$

conditional on the observations up to time t .

Using the estimation sub-sample $t = 1971/Q1 + q + h, \dots, \tau - h$ we now run the forecast regression (6.1) for the forecast quarter of interest $\tau + h$ and obtain parameter estimates

$$\hat{\alpha}_{h,\tau}, (\hat{\gamma}_{0,h,\tau}, \dots, \hat{\gamma}_{q,h,\tau}), (\hat{\beta}_{0,h,\tau}, \dots, \hat{\beta}_{q,h,\tau}), \hat{\sigma}_{h,\tau}^2, \quad (6.4)$$

where $\hat{\sigma}_{h,\tau}^2$ is the estimate of the error variance based on the regression residuals. Notice that the last observation of the dependent variable for the estimation is π_{τ}^h , where τ indexes the last quarter used in the estimation. The forecast distribution for the input is normal with mean

$$b_0 = \hat{\pi}_{\tau+h|\tau}^h = \hat{\alpha}_{h,\tau} + \pi_{\tau} + \sum_{\ell=0}^q \hat{\gamma}_{\ell,h,\tau} \Delta \pi_{\tau-\ell} + \sum_{\ell=0}^q \hat{\beta}_{\ell,h,\tau} x_{\tau-\ell}, \quad (6.5)$$

and variance $\sigma_0^2 = \hat{\sigma}_{h,\tau}^2$. The Bayesian Information Criterion (BIC) selected $q = 2$ for all forecast regressions and subsequent results are based on these results. The Akaike Information Criterion (AIC) selected less parsimonious models, but the associated results are qualitatively similar to those based on the BIC. Detailed forecast regression results based on the BIC and AIC are reported in the Replication Package. The results suggest that contemporaneous changes of the inflation and unemployment rates are highly significant, while the time-lagged unemployment rates are not as strongly significant as the time-lagged changes in inflation rates.

The forecast distributions based on the forecast regression results for the three events of interest ($\tau = 2016/Q2, 2020/Q1$, and $2022/Q1$) are reported in Table 2. The forecast mean b_0 appears to be responsive to changes in the unemployment rate and inflation, as expected; and it reacts negatively to changes in the contemporaneous unemployment rate.

The 2017/Q2 (referendum) inflation forecasts show stable uncertainty around $\sigma_0 = 2.45$, while the mean b_0 falls sharply from 1.47% to 1.09% and then jumps back up to 1.55%. For the 2021/Q1 (lockdown) forecasts we see a decrease in both mean and variance from 2019/Q3 to 2020/Q1. The 2023/Q1 (invasion) forecasts show a steep increase in b_0 from 2021/Q3 to 2022/Q1. The uncertainty of the forecasts is stable during these quarters.

When comparing Tables 1 and 2 we see that $\sigma_2 < \sigma_0$ in all cases: the variance in the input is larger than the variance in the output (the posterior). This important theoretical restriction, discussed at the end of Section 2, is thus satisfied in our case. Since data are noisy and stubborn, this may not always be the case in applications. A particular challenge arises when we have several inputs, a group of expert advisors to the committee rather than one advisor (Professor William Phillips, MBE) as in our case. With a group of advisors the consensus b_0 could have a misleadingly high precision (low variance), caused by the fact that these expert advisors tend to be highly correlated. This correlation needs to be neutralized in order to achieve a credible input; see Magnus and Vasnev (2023) for an investigation of this problem.

Table 2. UK inflation: the inputs (estimated forecast distributions and variables).

Year/Quarter (τ)	Forecasts (inputs)		Variables at τ			at $\tau + h$
	b_0	σ_0	x_τ	π_τ^4	π_τ	$\pi_{\tau+h}^4$
Forecast for 2017/Q2 ($\tau + h$) made in:						
2015/Q4	1.55	2.47	5.10	0.10	0.53	2.15
2016/Q1	1.09	2.47	5.10	0.33	-1.87	2.15
2016/Q2 (Referendum)	1.47	2.43	4.90	0.33	2.26	2.15
Forecast for 2021/Q1 ($\tau + h$) made in:						
2019/Q3	2.66	2.41	3.80	1.83	1.73	0.64
2019/Q4	1.75	2.39	3.80	1.39	0.62	0.64
2020/Q1 (Lockdown)	1.65	2.35	4.00	1.64	0.12	0.64
Forecast for 2023/Q1 ($\tau + h$) made in:						
2021/Q3	4.55	2.38	4.30	2.72	4.31	9.67
2021/Q4	6.53	2.38	4.00	4.77	8.72	9.67
2022/Q1 (Invasion)	7.61	2.39	3.70	6.01	5.33	9.67

Note: The input parameters are the estimates from the forecast distribution $N(b_0, \sigma_0^2)$, where b_0 is given in (6.5), $\sigma_0 = \hat{\sigma}_{h,\tau}$, x_τ is the unemployment rate, π_τ^4 is the four-quarter inflation rate, and π_τ is the annualized inflation rate at quarter τ .

7. THE REVEALED PRIOR

We now have the posteriors and the input, so we can reveal the priors. The posteriors, discussed in Section 5, are different for the BoE and the NIESR, but the input is the same: both institutions are assumed to use the Phillips curve inflation forecasts, discussed in Section 6. Hence, any difference in the posterior must be due to a difference in the prior. We shall see that the implied priors reveal significant differences between the two institutions, and we shall discuss the likely sources of these differences.

The estimated moments of the prior distribution for both the BoE and the NIESR are reported in Table 3, where we use $\hat{\pi}_{\tau+h|\tau}^h$ ($h = 4, 5, 6$) as an approximation to b_0 and $\sigma_0 = \hat{\sigma}_{h,\tau}$ as an approximation to σ_0 . We discuss and interpret these priors for each of the three exogenous shocks below.

7.1. The Brexit referendum

The Brexit referendum was held on 23 June 2016 (2016/Q2) and the outcome was a surprise. The general expectation before the referendum was that the UK population would vote to remain in the European Union, and the forecasts by the BoE and the NIESR were made under the assumption that the ‘remainers’ would win. The BoE and NIESR prior distributions reported in Table 2 are shown in Figures 1 and 2.

Consider first the prior uncertainties σ_1 . The BoE appears to be more confident about their forecasts than the NIESR, as indicated by the lower standard deviations, but the trend over the quarters is roughly the same. The value of σ_1 falls slightly from 2015/Q4 to 2016/Q1, then falls significantly in the referendum quarter. This significant fall may be due to the fact that after the referendum the uncertainty about its outcome had been removed, although other kinds of

Table 3. UK inflation: the priors.

Year/Quarter	BoE		NIESR	
	b_1	σ_1	b_1	σ_1
2017/Q2 inflation forecast distributions made in:				
2015/Q4	1.49	1.65	1.85	2.72
2016/Q1	1.70	1.60	0.91	2.54
2016/Q2 (Referendum)	2.19	1.49	3.86	1.72
2021/Q1 inflation forecast distributions made in:				
2019/Q3	1.20	1.67	1.41	2.24
2019/Q4	1.43	1.62	2.41	2.12
2020/Q1 (Lockdown)	-1.31	2.96	0.65	1.56
2023/Q1 inflation forecast distributions made in:				
2021/Q3	1.89	2.46	2.65	2.34
2021/Q4	4.07	2.37	5.03	2.61
2022/Q1 (Invasion)	10.55	2.04	6.54	1.98




Figure 1. BoE: prior forecast for 2017/Q2.

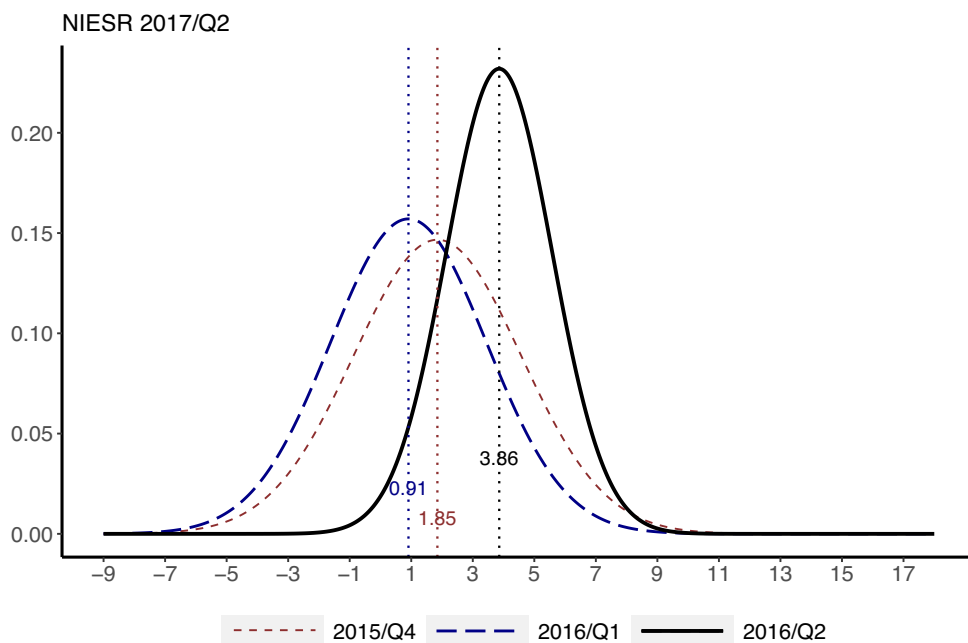


Figure 2. NIESR: prior forecast for 2017/Q2.

uncertainty remained or emerged, for example uncertainty about the final trading arrangements. The MPC assumed (Bank of England 2016, May) that such uncertainties would not persist, which may have contributed to a smaller σ_1 than the NIESR.

In contrast, the trends of the BoE and NIESR prior means b_1 are very different. The BoE's prior inflation forecast rises rapidly across quarters, accelerating towards the referendum quarter, while the NIESR prior mean falls from 2015/Q4 to 2016/Q1 and then rises significantly in 2016/Q2. What could lie behind this large difference in prior beliefs between the two institutions?

The difference in 2016/Q1 can perhaps be explained by the fact that the two institutions worked under different assumptions about future monetary policy. The BoE assumed that inflation would rise automatically to the 2% target by mid-2018, and was reluctant to initiate monetary policy that might bring too rapid inflation (see Bank of England 2016, February and May). The NIESR, however, assumed substantial inflation in two years' time, and expected 'the Monetary Policy Committee to move to raise rates by the end of this year and then follow a policy of gradually tightening to 1.5% by the end of 2017' (Kirby et al. 2016a, b).

Concerning the difference in prior means in 2016/Q2, both the BoE and the NIESR forecasts were made given that the base rate would be cut from 0.50% to 0.25% in August 2016. The prior mean b_1 is much higher for the NIESR than for the BoE, possibly because the NIESR assumed a further cut of the base rate to 0.1% within 2016, while the BoE did not make such an assumption (see Bank of England 2016, August; Kirby et al. 2016c). Nonetheless, both the NIESR and the BoE appear to have been rather confident about their prior beliefs in 2016/Q2 as evidenced by the small (and stable) standard deviations.

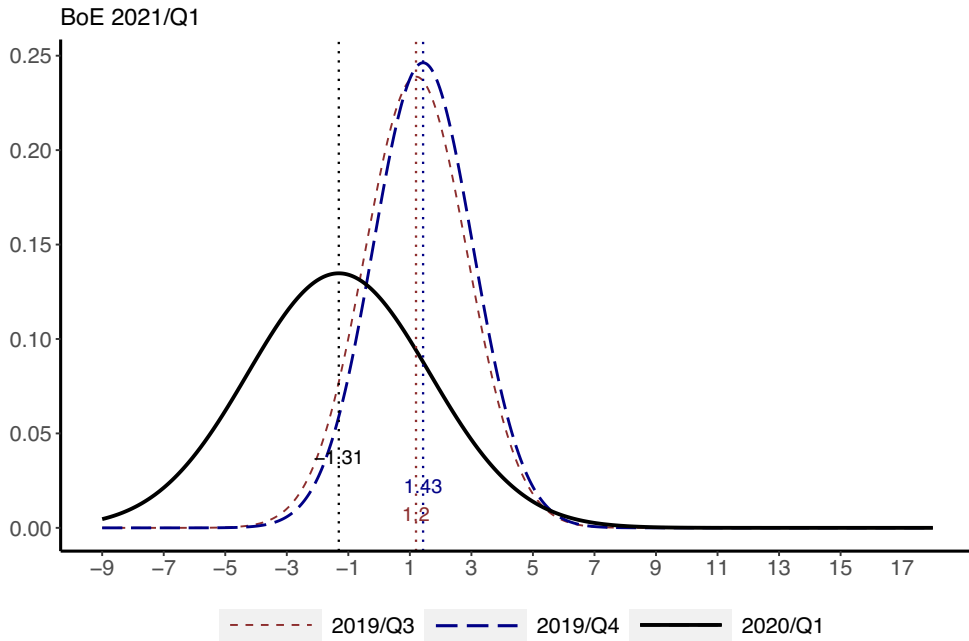


Figure 3. BoE: prior forecast for 2021/Q1.

7.2. The Covid-19 lockdown

On 16 March 2020 (2020/Q1) UK Prime Minister Boris Johnson announced, in response to the Covid-19 threat, that ‘now is the time for everyone to stop non-essential contact and travel’, and lockdown measures came legally into force on 26 March. Following the lockdown, the BoE decided not to provide a density inflation forecast in their quarterly report (Bank of England 2020, May), an unprecedented step showing large uncertainty. The BoE did however describe a scenario in which the annual inflation rate in 2021 would be 0.50% (see the note to Table 1). We use this value as the posterior mean b_2 .

The BoE and NIESR prior distributions reported in Table 2 are shown in Figures 3 and 4. The trends in the BoE and NIESR prior means are similar, rising from 2019/Q3 to 2019/Q4 and then falling dramatically in 2020/Q1. The NIESR prior means are consistently higher than those of the BoE.

One month before the lockdown, in the February 2020 *Monetary Policy Report*, the MPC decided to maintain the base rate at 0.75%. This was because the growth in regular pay fell back to around 3.5%, although unit labour costs continued to grow at rates above those consistent with meeting the inflation target in the medium term. The NIESR had a similar view, but it assumed that the base rate would be cut by 0.25% at the end of March and then remain at 0.5% until the end of 2021 (Hantzsche and Young 2020a). This may explain, in part, why the NIESR had a higher prior mean than the BoE in 2019/Q4.

After the national lockdown, domestic and world economic conditions deteriorated sharply. To respond to this 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. The inflation rate declined to 1.5% in March (see Bank

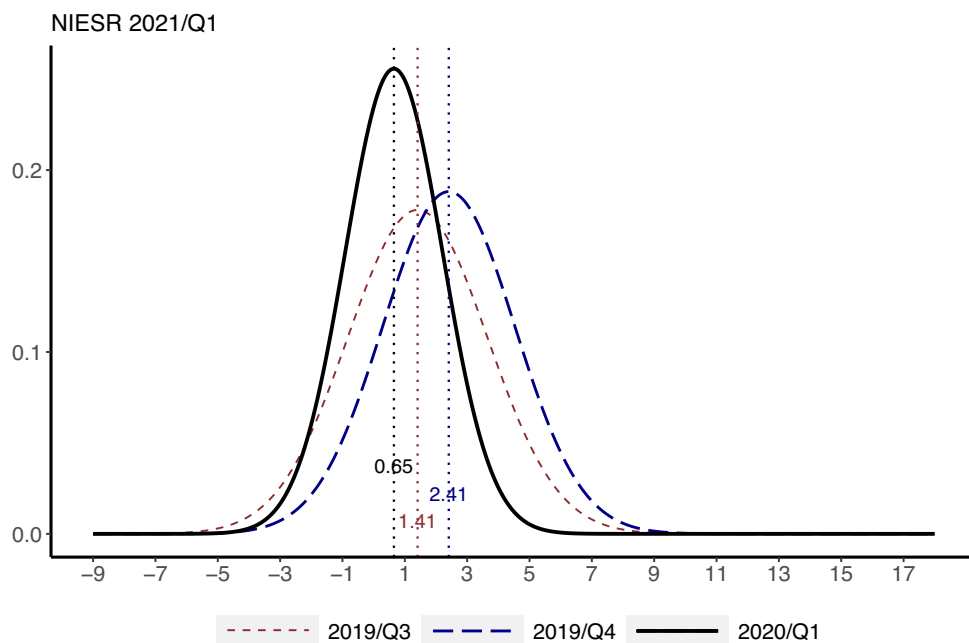


Figure 4. NIESR: prior forecast for 2021/Q1.

of England 2020, May) and the priors of the BoE and the NIESR in 2020/Q1 had to be adjusted after this unexpected shock. The difference in the 2020/Q1 prior means is partially explained by the fact that the NIESR argued that the Covid-19 shock would reduce both demand and supply, thus having a broadly neutral effect on inflation (Hantzsche and Young 2020b).

The trends of the prior uncertainties of the BoE and the NIESR are very different, especially towards the lockdown quarter. From 2019/Q3 to 2019/Q4 the prior uncertainty of the BoE is much smaller than that of the NIESR, but in the lockdown quarter the prior uncertainty of the BoE increased unprecedentedly, while that of the NIESR sharply decreased. The BoE's prior distribution, with a huge standard deviation and a hard-to-believe mean of -1.31% , reflects the massive uncertainty caused by the pandemic and the lockdown. The confidence of the NIESR in the prior in 2020/Q1 is much higher than the confidence of the BoE, even higher than their own confidence in the previous quarters. The NIESR must have made their forecast expecting that the economic shock due to the lockdown would disappear well before 2021/Q1.

7.3. The Russian invasion of Ukraine

On 24 February 2022 (2022/Q1) Russia invaded Ukraine, and the resulting war was expected to drive international energy and food prices higher for a long period to come.

The BoE and NIESR prior distributions reported in Table 2 are shown in Figures 5 and 6. In 2021/Q3–2022/Q1 energy prices rose faster than expected and there were bottlenecks in supplies, especially, but not only after the invasion. During periods of sharp rises in energy prices, the UK inflation forecast is heavily influenced by the price control facilities regulated by Ofgem (Office of Gas and Electricity Markets), known as the Energy Price Cap (January 2019–September 2022)




Figure 5. BoE: prior forecast for 2023/Q1.

followed by the Energy Price Guarantee (October 2022–April 2024 (plan)), both of which set caps on household gas and electricity prices. As a result, the forecasts had to be significantly revised each quarter.

The main difference between the BoE and NIESR prior distributions is the prior mean for 2022/Q1: a staggering 10.55% for the BoE and 6.54% for the NIESR. The BoE's prior uncertainty is similar to that of the NIESR. Apparently, the BoE and the NIESR have different prior beliefs about the impact and persistence of energy price increases on CPI inflation. The BoE assumes (Bank of England 2022, May) that energy prices will persist near their peak for several months, highlighting the strong inflationary pressure from rising energy prices, while the NIESR's report (National Institute of Economic and Social Research 2022, May) highlights the impact of significant energy and food price inflation on the real economy by reducing real wages and household disposable income, which will pull inflation downwards.

The BoE analysed (Bank of England 2022, May) the contribution of the energy price hike to inflation, and it estimated that the 1.7 percentage points of the 7% CPI inflation in March 2022 was directly accounted for by energy prices and that the energy contribution to inflation was expected to increase by a further 1.6 percentage points in 2022/Q4. In contrast, the NIESR seems to somewhat underestimate the contribution of energy prices to inflation, stating that 'electricity, gas and other fuels account for only 3.3% of the UK CPI basket, compared with 7.7% of the US CPI basket' (National Institute of Economic and Social Research 2022, May), while emphasizing the negative impact of inflation on real disposable income, consumption, and GDP.

The NIESR's forecasts for unemployment in 2022–23 are also more pessimistic than those of the BoE. The NIESR forecast an average unemployment rate of 4.4% in 2022 and a slight

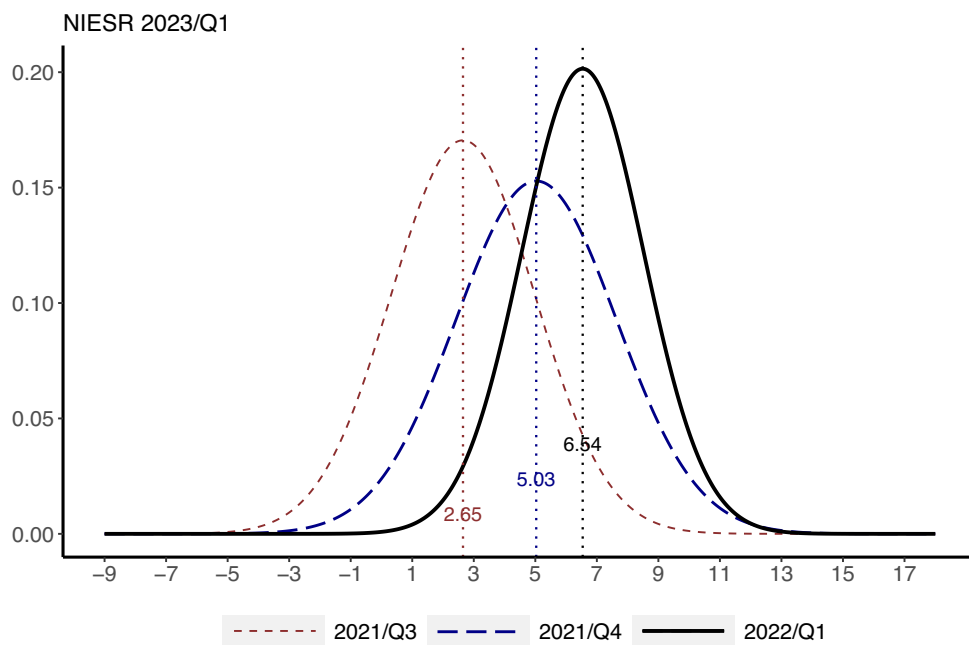


Figure 6. NIESR: prior forecast for 2023/Q1.

increase to an average of 5% in 2023, while the BoE forecast an average rate of slightly below 4% in 2022 and slightly above 4% in 2023 (Bank of England 2022, May; National Institute of Economic and Social Research 2022, May).

7.4. Strength of the prior

So far we have analysed the priors of the BoE and the NIESR for each quarter trying to identify the source of their differences. Now we analyse the strengths of their priors.

In Table 4 we present the key parameter ratios, as discussed in Section 3. We focus on the parameters related to the prior: κ_m and κ_v . Recall that the closer is κ_m to unity, the stronger is the prior in the sense that b_2 (the posterior mean) is largely determined by b_1 (the prior mean) and not by the input b_0 . For the lockdown quarter and invasion period (2020/Q1 and 2021/Q3–2022/Q1), NIESR's κ_m is uniformly closer to unity than BoE's, showing that the NIESR holds strong views concerning the prior mean compared to the BoE in that period.

Concerning κ_v we recall that the smaller the value of κ_v , the stronger the prior information. Throughout the referendum quarters (2015/Q4–2016/Q2) and in the quarters before the first lockdown (2019/Q3–Q4), BoE's κ_v ranges from 0.38 to 0.48, uniformly smaller than the NIESR's by a large margin. This shows that the BoE holds stronger prior views than the NIESR during this period. At the lockdown quarter the situation is reversed and the NIESR holds much stronger prior views than the BoE: $\kappa_v = 1.58$ for BoE and 0.44 for the NIESR. During the invasion quarters (2021/Q3–2022/Q1), the strengths of the priors held by the BoE and the NIESR are similar. Apparently, the BoE had great confidence in its ability to achieve the inflation target, but

Table 4. UK inflation: $\alpha_m = b_2/b_0$, $\alpha_v = \sigma_2^2/\sigma_0^2$, $\kappa_m = b_1/b_0$, and $\kappa_v = \sigma_1^2/\sigma_0^2$.

Year/Quarter	BoE				NIESR			
	α_m	α_v	κ_m	κ_v	α_m	α_v	κ_m	κ_v
The forecast for 2017/Q2 made in:								
2015/Q4	0.98	0.31	0.96	0.44	1.09	0.55	1.20	1.21
2016/Q1	1.40	0.29	1.56	0.42	0.92	0.51	0.83	1.06
2016/Q2 (Referendum)	1.35	0.27	1.49	0.38	2.08	0.33	2.63	0.50
The forecast for 2021/Q1 made in:								
2019/Q3	0.63	0.32	0.45	0.48	0.75	0.46	0.53	0.87
2019/Q4	0.87	0.31	0.81	0.46	1.21	0.44	1.37	0.79
2020/Q1 (Lockdown)	0.30	0.61	-0.80	1.58	0.58	0.31	0.39	0.44
The forecast for 2023/Q1 made in:								
2021/Q3	0.72	0.52	0.42	1.06	0.79	0.49	0.58	0.96
2021/Q4	0.81	0.50	0.62	0.99	0.90	0.55	0.77	1.20
2022/Q1 (Invasion)	1.22	0.42	1.39	0.73	0.92	0.41	0.86	0.69

this confidence was undermined by the lockdown and the energy crisis caused by the Russian invasion.

8. SUMMARY AND SOME OPEN ISSUES

In this paper we have tried to reveal the prior, given information about the data (which we call the input) and the posterior, thus reversing the Bayesian method. We have assumed normality in the development and the application, although normality can be relaxed at the cost of more cumbersome and less transparent expressions. We then applied the theory to inflation density forecasts by the BoE, where we assumed a simple model of how the BoE operates in the process leading to its forecast publication. The BoE uses data, but indirectly. The data are used to form an estimate of the forecast and its distribution (in our case, through the Phillips curve) and this estimate is then used as input into the Bayesian framework. The committee deciding on the report takes this input while adding their own ideas, knowledge, and prejudices, which we call the prior. Input and prior together give the posterior which we observe as the published report. In practice, we observe the input (generated by data) and the posterior, but not the prior, which is what we are trying to reveal.

Two aspects of the BoE's prior forecast distribution are considered. The first aspect analyses changes in the distribution of prior forecasts over time, in particular how exogenous shocks affect it. We consider three recent major shocks to the UK economy: the Brexit referendum, the Covid-19 lockdown, and the Russian invasion of Ukraine. This is investigated by comparing the prior distributions before and during the shocks.

In addition, as a second aspect, we analyse how the prior distribution of the forecasts of two different forecasters may differ, especially when an economic shock occurs. For this analysis, the density forecasts of the BoE were compared with those of the NIESR, an independent UK institution of high repute.

Analysis of the prior distributions during the three exogenous shocks reveals that:

- at 2016/Q2, when the referendum took place, the means of the prior distribution for the BoE and NIESR are very different, possibly due to different assertions on the economic situation, but the prior uncertainties are both small and similar, indicating that the Brexit outcome did not add uncertainty to prior distribution;
- in 2020/Q1, the first lockdown quarter, prior uncertainty for the BoE rose sharply, while prior uncertainty for the NIESR fell, indicating that prior uncertainty for the NIESR was not affected by the lockdown; and
- in 2022/Q1, during the Russian invasion quarter, the BoE and NIESR prior means rose, but the prior mean of the BoE (10.55%) was much higher than that of the NIESR (6.54%). This may have been due to the fact that the BoE estimated that the contribution of higher energy prices on CPI inflation was much higher than the contribution estimated by the NIESR.

Comparing the parameters which indicate the degree of confidence in the prior belief, we found that the BoE had great confidence in its ability to achieve the inflation target, which was undermined by the lockdown and by the energy crisis and Russian invasion.

There are many further issues worth investigating, of which we mention two. First, the way in which the input is constructed (the Phillips curve) is highly stylized. If we had more information then more realistic input data could be constructed. For example, the NIESR has an in-house forecasting model called NiGEM and it would be of interest to use this information and refine the process from input to prior. Second, one may consider a situation where there are several agents, as in the case of the BoE and NIESR, but with different inputs. In such a situation, the selection of the data source may and probably will be subject to prior beliefs as well.

ACKNOWLEDGEMENTS

We are grateful to Peter Boswijk, Siem Jan Koopman, Gideon Magnus, Francesca Monti, Mark Steel, Takayuki Tsuruga, and Rick van der Ploeg for helpful discussions, and to the editor and referees for their constructive comments. Earlier versions of this paper were presented at the Vrije Universiteit Amsterdam (April 2020), the Netherlands Econometric Study Group (May 2020), ISER, Osaka University (December 2020), the University of Palermo (March 2023), Kobe University (May 2023), and Tohoku University (July 2023). This work was financially supported by the JSPS KAKENHI Grants Number JP20H01477, JP21H00700, JP21H04397, and the programme of the Joint Usage/Research Center for Behavioral Economics at ISER, Osaka University.

REFERENCES

- Bank of England (2016). *Inflation Report*, May 2016, London, UK. <https://www.bankofengland.co.uk/-/media/boe/files/inflation-report/2016/may-2016.pdf>.
- Bank of England (2019–2023). *Monetary Policy Report*, published quarterly, London, UK. <https://www.bankofengland.co.uk/monetary-policy-report/monetary-policy-report>.
- Galvão, A. B., A. Garratt and J. Mitchell (2021). Does judgment improve macroeconomic density forecasts? *International Journal of Forecasting* 37, 1247–60.

- Hantzsche, A. and G. Young (2019). Prospects for the UK economy. *National Institute Economic Review* 250, F4–F33.
- Hantzsche, A. and G. Young (2020a). Prospects for the UK economy. *National Institute Economic Review* 251, F4–F34.
- Hantzsche, A. and G. Young (2020b). Prospects for the UK economy. *National Institute Economic Review* 252, F10–F43.
- Kirby, S., O. Carreras, J. Meaning, R. Piggott and J. Warren (2016a). Prospects for the UK Economy. *National Institute Economic Review* 235, F47–F76.
- Kirby, S., O. Carreras, J. Meaning and R. Piggott (2016b). Prospects for the UK economy. *National Institute Economic Review* 236, 82–107.
- Kirby, S., O. Carreras, J. Meaning, R. Piggott and J. Warren (2016c). Prospects for the UK economy. *National Institute Economic Review* 237, F42–F71.
- Kiss, T., S. Mazur, H. Nguyen and P. Österholm (2023). Modeling the relation between the US real economy and the corporate bond-yield spread in Bayesian VARs with non-Gaussian innovations. *Journal of Forecasting* 42, 347–68.
- Lenoël, C., R. Macqueen and P. Mortimer-Lee (2021). UK economic outlook: shortages and fiscal tightening threaten economic recovery from Covid-19. *National Institute Economic Review* 258, E1.
- Macqueen, R., S. Millard, U. Patel and K. Whyte (2022). UK economic outlook. *National Institute Economic Review* 260, E1.
- Magnus, J. R. and A. Vasnev (2023). On the uncertainty of a combined forecast: the critical role of correlation. *International Journal of Forecasting* 39, 1895–908.
- McNees, S. K. (1990). The role of judgement in macroeconomic forecasting accuracy. *International Journal of Forecasting* 6, 287–99.
- National Institute of Economic and Social Research (2022). *National Institute UK Economic Outlook—Spring 2022*, London, UK.
- Phillips, A. W. (1958). The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957. *Economica* 25, 283–99.
- Stock, J. H. and M. W. Watson (2009). Phillips curve inflation forecasts. In J. Fuhrer and others (Eds.), *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, MIT Press, Cambridge, MA. <https://www.amazon.co.uk/Understanding-Inflation-Implications-Monetary-Policy/dp/0262013630>.
- Stock, J. H. and M. W. Watson, (2020). Slack and cyclically sensitive inflation. *Journal of Money, Credit and Banking* 52, 393–428.
- Turner, D. S. (1990). The role of judgement in macroeconomic forecasting. *Journal of Forecasting* 9, 315–45.
- Wallis, K. F. (2014). The two-piece normal, binormal, or double Gaussian distribution: its origin and rediscoveries. *Statistical Science* 29, 106–12.
- Winkler, R. (1968). The consensus of subjective probability distributions. *Management Science* 15, B61–B75.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at the publisher's website:

Replication Package

Co-editor Raffaella Giacomini handled this manuscript.