

Revealing priors from posteriors

Masako Ikefuji

Faculty of Humanities and Social Sciences, University of Tsukuba
and ISER, 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

October 19, 2021

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 then apply this theory to inflation forecasts by the Bank of England and to (equilibrium) climate sensitivity as reported by the Intergovernmental Panel on Climate Change in an attempt to get some insight into the prior beliefs of central bankers (especially under the uncertainty about Brexit and the corona crisis) and IPCC scientists.

JEL Classification: C11, E31, E58, Q54.

Keywords: Revealed prior; inflation; central bank; climate sensitivity; IPCC.

Corresponding author: Takashi Yamagata, Department of Economics and Related Studies, University of York, York YO10 5DD, UK.

Email addresses: `ikefuji.masako.gn@u.tsukuba.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, and to apply this theory to inflation forecasts and the perception of climate sensitivity.

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 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 present two illustrations of this theory: inflation forecasts by the Bank of England, especially under the uncertainty about Brexit and the corona crisis; and the estimation of the equilibrium climate sensitivity (ECS), which is an important diagnostic in climate modeling. In both cases we are interested to recover the prior beliefs of the decision maker: the Bank of England and the IPCC, respectively.

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. The current paper attempts to fill this gap. Of course, 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 studied by the theory developed in this paper.

In Section 2 we analyze how to recover the prior from the data and the posterior within the framework of the normal distribution. In Section 3 we consider the case when there is only one parameter of interest. Two examples will illustrate the theory: the Bank of England’s forecast of the interest rate (Sections 4 and 5) and the estimation of climate sensitivity (Sections 6 and 7). Section 8 concludes. An Appendix discusses the role of monotonic reparametrizations in general, and the lognormal distribution in particular.

2 From posterior to prior under normality

In order to highlight the issue under simple conditions, we shall first assume normality, which we shall modify later. 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. This has practical consequences as we shall see later.

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 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 $\sigma_2^2 < \sigma_0^2$.

Defining α_m and α_v implicitly by

$$b_2 = \alpha_m b_0, \quad \sigma_2^2 = \alpha_v \sigma_0^2, \quad (12)$$

we can rewrite (11) as

$$b_1 = \kappa_m b_0, \quad \sigma_1^2 = \kappa_v \sigma_0^2, \quad (13)$$

where

$$\kappa_m = \frac{\alpha_m - \alpha_v}{1 - \alpha_v}, \quad \kappa_v = \frac{\alpha_v}{1 - \alpha_v} \quad (14)$$

measure how far the prior is removed from the data 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 data and the variance σ_2^2 in the posterior 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 data 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 rather 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

$$1 - \kappa_m = \frac{1 - \alpha_m}{1 - \alpha_v}, \quad (15)$$

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 data and the mean b_2 in the posterior are approximately equal, but the variance σ_0^2 in the data 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 data 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 $b_1 \approx b_2$.

We shall apply this theory to two cases of interest. In Sections 4 and 5 we try to reveal the prior of the Bank of England in forecasting inflation; then, in Sections 6 and 7 we study the prior of the IPCC in estimating climate sensitivity.

4 Inflation in the UK: posterior and data

Density forecasts provide richer information on forecast uncertainties than point forecasts, and decision makers, professional forecasters, and academic

researchers have increasingly employed this insight to forecast macroeconomic variables. In particular, the Monetary Policy Committee (MPC) of the Bank of England (BoE) has produced quarterly reports on GDP growth and inflation since 1996, and density forecasts are used in these reports to explain the employed monetary policies.

Central banks (including the BoE) and professional forecasters don't follow the model-based forecast densities mechanically; they will add a final touch based on their subjective judgment. Following the approach of McNeese (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". In their study they need to separately identify the adjustment by the forecaster and the unadjusted mechanical density forecasts.

The current paper, rather than aiming at assessing the effects of adjustment, investigates the process in which the decision makers, such as the MPC, finalize the density forecast. Applying the Bayesian framework, in line with the discussion in Winkler (1968), we assume that the decision maker (i.e., the BoE or external forecasters) revises its prior distribution in view of "data" from other experts to form the posterior distribution (i.e., the published density forecast). Applying our method discussed in Sections 2 and 3, we reveal the prior of the MPC, making use of the data and the published density forecast.

In addition to revealing the prior of one institution (the BoE), we also ask a second question, namely what happens if two institutions provide density forecasts based on the same data. Their data are the same but their posteriors are different, which can only mean that their priors are different. How different are these priors, especially when a shock occurs like the Brexit or the Covid-19 lockdown? To investigate this question, we also consider density forecasts by the National Institute of Economic and Social Research (NIESR), an independent and highly-regarded organization in the UK.

4.1 The posterior

The Bank of England's primary responsibility is to keep UK inflation at 2% and the Monetary Policy Committee's task is to decide what monetary policy action to take 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 inflation

rate, economic growth rate, and 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. We focus on the quarters before and after two recent events which significantly 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; see Bank of England (2016a,b; 2020a,b,c), respectively.

The density forecasts by the BoE are based on the so-called two-piece normal distribution.¹ In the quarters under consideration, skewness is absent in three of the five quarters and very mild in the remaining two: 0.1 in 2016/Q3 and -0.26 in 2020/Q3. Hence 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.

The NIESR publishes its economic density forecasts every quarter in *Prospects for the UK Economy*. We consider this density forecast as NIESR’s posterior. The density forecasts are produced using the National Institute Global Econometric Model together with the institution’s judgment (see Source to Figure 7 in Lenoël et al., 2020). The density forecasts for the five quarters under consideration are reported in *Prospects for the UK Economy*; see Kirby et al. (2016a,b), Hantzsche and Young (2020), Lenoël and Young (2020), and Lenoël et al. (2020). The density forecasts are reported as fan-charts. We use these fan-charts to approximate the mean and variance of the appropriate posterior normal distribution.

The posterior means and standard deviations of the BoE and the NIESR for the five quarters are reported in Table 1. The trajectory of the BoE’s posterior mean b_2 is similar to that of the NIESR, but the amplitude of the latter is much wider. The trajectories of the posterior standard deviation σ_2 are very different for the two institutions: BoE’s standard deviation remains the same before and after the referendum, whilst NIESR’s large standard deviation just before the referendum falls to its lowest level right after the referendum.

4.2 The data: BoE’s survey of external forecasters

The data are based on The Bank of England Survey of External Forecasters. In this quarterly survey the BoE asks professional forecasters to indicate

¹The 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.

Table 1: UK inflation: the posterior

		Referendum		Lockdown		
		2016/Q2	2016/Q3	2020/Q1	2020/Q2*	2020/Q3
BoE	b_2	1.52	2.03	1.53	0.50	1.55
	σ_2	1.34	1.34	1.34	2.02	2.02
NIESR	b_2	0.88	3.03	2.10	0.94	2.07
	σ_2	1.72	1.35	1.65	1.49	1.71

* In 2020/Q2 the values for b_2 and σ_2 for the BoE are missing. We take b_2 from “the illustrative scenario” in “conditioning assumptions and scenario summary” (BoE, 2020b), and for σ_2 we employ the value for 2020/Q3.

probabilities they would attach to various possible outcomes in three future quarters. What is special about this survey is that the forecasters can indicate not only their most like range (say, 1.5–2.0 percent) but that they are allowed to attach probabilities to events (say, 25% for the range 1.0–1.5, 50% for the range 1.5–2.0, and 25% for the range 2.0–2.5). The surveyed forecasters cannot be identified by us, but they consist of City firms, academic institutions, and private consultancies, mainly based in London. The BoE then reports a summary of each survey in their quarterly *Monetary Policy Report*. The summary report includes the survey average density forecasts, i.e. the average of the respondents’ probabilities in each interval; see Boero et al. (2015) for more details about the survey.

Table 2: UK inflation: the data

	2016/Q2	2016/Q3	2020/Q1	2020/Q2	2020/Q3
< 0.0	5	2	3	7	5
0.0–1.0	15	4	7	23	20
1.0–1.5	24	9	18	22	19
1.5–2.0	27	18	36	24	25
2.0–2.5	14	23	25	13	16
2.5–3.0	8	21	7	7	10
> 3.0	8	23	3	4	6
b_0	1.63	2.36	1.76	1.42	1.59
σ_*	0.91	0.92	0.67	0.90	0.92

Table 2 gives the data distribution for the five quarters of interest to us. Each column contains percentages adding up to 100 (apart from rounding errors), and each column thus maps out a distribution which we approximate by $N(b_0, \sigma_*^2)$. The parameters b_0 and σ_* are estimated by maximum likelihood for interval censored values, using the R package “fitdistrblus”, provided by Delignette-Muller and Dutang (2015). The estimated values reported in Table 2 fit the survey data well. The trajectory of b_0 is similar to that of b_2 , but the trajectory of σ_* is different from that of σ_2 . Furthermore, σ_* is uniformly smaller than σ_2 .

We shall assume that the NIESR uses the same data as the BoE, and that therefore any difference in the posterior must come from a difference in the prior. This assumption requires some defense. The NIESR forecasters frequently exchange information about economic conditions with economists from major survey organizations and the official sectors.² Professional forecasters also devote significant time and effort to collect information and improve methods for economic forecasting, and their information set can thus be regarded as an upper bound among the forecasters; see Andrade and Le Bihan (2013) and Candia et al. (2021). This implies that, even if the BoE and the NIESR collect density forecasts from different sets of a sufficiently large number of professional forecasters, these two information sets will be highly correlated.

We now have the posteriors (different for BoE and NIESR) and the data (same for BoE and NIESR), so we can reveal the priors. We shall do so in the next section, but before applying our theory we need to discuss one remaining issue which occurs here and also in the next application, and in fact is a common problem. The problem is that the variation in the data is too small or, in other words, the precision of the data is overestimated. We have noted above that σ_* (from the data) is uniformly smaller than σ_2 (from the posterior). But the theory prescribes that the posterior variance must be smaller than the data variance (and also than the prior variance). After all, we add information to the data (in the form of a prior), so the precision should increase. A comparison of σ_* and σ_2 shows that this condition is violated, and hence we cannot choose $\sigma_0 = \sigma_*$ for the data.

This important aspect of the data seems to be largely ignored in the literature. We deal with this problem in a pragmatic way, by choosing three different values for σ_0 , namely 2.2, 2.6, and 3.0. These values are larger (but not much larger) than the values of σ_2 in Table 1, and they also yield a reasonable spectrum of revealed priors, as we shall see in the next section.

²See <https://www.niesr.ac.uk/business-conditions-forum>.

5 Inflation in the UK: revealed prior

Both the BoE and the NIESR revise their data in view of their priors before they publish their forecasts (their posteriors). We observe the data and the posterior but not the priors, but we can reveal the priors from the data and the posterior. The estimated priors of the BoE and the NIESR are reported in Table 3. We discuss and interpret these priors for each of the five quarters below.

Table 3: UK inflation: data and priors

Year/Quarter	Data		Priors			
	b_0	σ_0	BoE		NIESR	
			b_1	σ_1	b_1	σ_1
<i>Referendum</i>						
2016/Q2	1.63	2.2	1.46	1.69	-0.30	2.76
		2.6	1.48	1.56	0.30	2.29
		3.0	1.49	1.50	0.51	2.10
2016/Q3	2.36	2.2	1.84	1.69	3.43	1.71
		2.6	1.91	1.56	3.28	1.58
		3.0	1.95	1.50	3.20	1.51
<i>Lockdown</i>						
2020/Q1	1.76	2.2	1.39	1.69	2.54	2.49
		2.6	1.45	1.56	2.33	2.14
		3.0	1.47	1.50	2.25	1.98
2020/Q2	1.42	2.2	-4.44	5.10	0.53	2.03
		2.6	-0.90	3.21	0.71	1.82
		3.0	-0.26	2.73	0.78	1.72
2020/Q3	1.59	2.2	1.34	5.10	2.80	2.72
		2.6	1.49	3.21	2.44	2.27
		3.0	1.52	2.73	2.30	2.08

5.1 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 and the NIESR were made under the assumption that the “remainers” would win.

The quarter labeled 2016/Q2 is associated with the reports by the BoE and the NIESR which came out in May 2016, one month before the referendum, containing forecasts made in 2016/Q2 about the inflation one year later in 2017/Q2. We see that the BoE and the NIESR have very different priors. The BoE’s prior mean b_1 is about 1.47, irrespective of the assumption on the standard deviation σ_0 in the data. But the prior mean of the NIESR is quite sensitive to σ_0 and could even have been negative. The prior standard deviation σ_1 is relatively small (around 1.6) for the BoE and large (around 2.4) for the NIESR. The NIESR appears to have been quite uncertain about their prior beliefs, and the standard deviations in this quarter are the highest among the five quarters.

What is the reason for the large difference in the prior means between the BoE and the NIESR? The most important reason is probably that the two organizations differ in their hypothesized monetary policy paths. The BoE considered that inflation would rise automatically to the 2% target by mid-2018, and was reluctant to place monetary policy which might bring too rapid inflation. In contrast, the NIESR considered that there could be substantial inflation in two years’ time, and they 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 per cent by the end of 2017.”

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.

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. Both the NIESR and the BoE appear to have been rather confident about their prior beliefs given the small (and stable) standard deviations.

5.2 The lockdown

On March 16, 2020 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 March 26. 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 labor 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. This may explain, in part, why the NIESR has a higher prior mean than the BoE in the 2020/Q1 forecast.

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.

In 2020/Q2, the BoE did not provide density inflation forecasts, which was unprecedented. They 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 . Regarding the posterior standard deviation, we employ $\sigma_2 = 2.02$, the published value in the next quarter. This is much higher than $\sigma_2 = 1.34$ in the previous quarters, thus reflecting the increased uncertainty.

The prior for the BoE reflects its pessimistic feelings and its inability to make accurate forecasts. In contrast, the NIESR considered that the inflation rate would rise to around 1% in 2021, thus higher than BoE’s inflation expectation, because the NIESR considered that the Covid-19 shock would reduce both demand and supply, which would have a broadly neutral effect on inflation. Their prior is also pessimistic, but less so than the BoE, and their confidence in this prior is much higher.

In the next quarter, 2020/Q3, both the BoE and the NIESR substantially adjusted their 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. In the NIESR scenario, inflation would fall to -0.1% in 2020/Q3 but then would recover to about 2% in 2021. The higher prior of the NIESR shows that the NIESR considered deflation risk exaggerated.

5.3 Strength of the prior

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

Table 4: UK inflation: ratios of parameters
 $\alpha_m = b_2/b_0$, $\alpha_v = (\sigma_2/\sigma_0)^2$, $\kappa_m = b_1/b_0$, $\kappa_v = (\sigma_1/\sigma_0)^2$

Year/Quarter	Data		BoE				NIESR			
	b_0	σ_0	α_m	α_v	κ_m	κ_v	α_m	α_v	κ_m	κ_v
<i>Referendum</i>										
2016/Q2	1.63	2.2	0.93	0.37	0.89	0.59	0.54	0.61	-0.18	1.57
	1.63	2.6	0.93	0.27	0.91	0.36	0.54	0.44	0.18	0.78
	1.63	3.0	0.93	0.20	0.92	0.25	0.54	0.33	0.31	0.49
2016/Q3	2.36	2.2	0.86	0.37	0.78	0.59	1.28	0.38	1.46	0.60
	2.36	2.6	0.86	0.27	0.81	0.36	1.28	0.27	1.39	0.37
	2.36	3.0	0.86	0.20	0.83	0.25	1.28	0.20	1.36	0.25
<i>Lockdown</i>										
2020/Q1	1.76	2.2	0.87	0.37	0.79	0.59	1.19	0.56	1.44	1.29
	1.76	2.6	0.87	0.27	0.82	0.36	1.19	0.40	1.32	0.67
	1.76	3.0	0.87	0.20	0.84	0.25	1.19	0.30	1.28	0.43
2020/Q2	1.42	2.2	0.35	0.84	-3.13	5.37	0.66	0.46	0.38	0.85
	1.42	2.6	0.35	0.60	-0.63	1.52	0.66	0.33	0.50	0.49
	1.42	3.0	0.35	0.45	-0.19	0.83	0.66	0.25	0.55	0.33
2020/Q3	1.59	2.2	0.97	0.84	0.84	5.37	1.30	0.60	1.76	1.53
	1.59	2.6	0.97	0.60	0.94	1.52	1.30	0.43	1.53	0.76
	1.59	3.0	0.97	0.45	0.95	0.83	1.30	0.32	1.45	0.48

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 smaller (resp. larger) is the value of κ_v , the stronger (resp. weaker) is the prior information. Throughout the referendum period and in the quarter just before the first lockdown, BoE's κ_v ranges from 0.25 to 0.59, uniformly smaller than NIESR's. This shows the BoE holds stronger prior views than the NIESR. After the lockdown, the situation is reversed and the NIESR holds stronger prior views than the BoE, since κ_v ranges from 0.83 to 5.37 (BoE) and from 0.33 to 1.53 (NIESR).

Concerning κ_m we recall that the closer κ_m is 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 data. Except for the first lockdown quarter, BoE's κ_m is much closer to unity than NIESR's, which shows again that the BoE holds strong priors (except during the first lockdown) compared to the NIESR.

In conclusion, 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 1.34 and 1.95, 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 NIESR is more flexible and its prior views vary over the quarters. The referendum in particular dramatically affected the strength of the NIESR’s prior. Except in 2020/Q2, the NIESR has weaker prior views than the BoE.

6 Climate sensitivity: posterior and data

When the radiation balance of the Earth is perturbed, the temperature will change. By how much is measured by the *equilibrium climate sensitivity* (ECS): the long-term temperature rise that is expected to result from a doubling of the atmospheric CO₂ concentration, usually relative to the pre-industrial level (around 1750). It is a prediction of the new global mean near-surface air temperature once the CO₂ concentration has stopped increasing and most of the feedbacks have had time to have their full effect. The ECS is an important diagnostic in climate modelling, but it cannot be measured directly and forms a large source of uncertainty. CO₂ levels rose from 280 parts per million (ppm) in the eighteenth century (IPCC, 2013, p. 100) to about 416 ppm by 2020, an increase of almost 50%. In the same period, the Earth’s temperature rose by a little over one degree Celsius. In this and the next section, the ECS will be our parameter of interest, and we shall denote it by β .

In contrast to Sections 4 and 5 where we assumed normality (hence symmetry) of the distributions, this assumption is not plausible here. So we need to search for a credible alternative, and the simplest alternative is normality of a monotonic transformation of β , say $h(\beta)$. Given this transformation the same theory shows that if $b_0 \sim N(h(\beta), \sigma_0^2)$ is an unbiased estimator of $h(\beta)$ from the data and $h(\beta) \sim N(b_2, \sigma_2^2)$ is the posterior, then the prior can be recovered as

$$h(\beta) \sim N(b_1, \sigma_1^2), \tag{16}$$

where b_1 and σ_1^2 are given in (11). From the prior distribution of $h(\beta)$ we then obtain the prior distribution of β . The most important example is $h(\beta) = \log \beta$ in which case the posterior and the prior are both lognormally distributed; see the Appendix for details. This is the route that we shall follow.

6.1 The posterior

In estimating β we rely exclusively on the Intergovernmental Panel on Climate Change (IPCC) reports. So far, five so-called Assessment Reports have appeared, the first in 1990, the fifth in 2013.³ In these reports we find estimates (and precisions) of studies on the ECS (our data) and the IPCC’s own estimates (our posterior). From this information we will attempt to recover the IPCC’s priors.

In the fifth report, more precisely the Working Group I contribution (IPCC, 2013), hereafter IPCC5, the authors state that “no best estimate for equilibrium climate sensitivity can now be given because of a lack of agreement on values across assessed lines of evidence and studies” (IPCC5, p. 16, footnote). But later in the same report they do provide estimates, as follows:

“...ECS is likely in the range 1.5°C to 4.5°C with high confidence. ECS is positive, extremely unlikely less than 1°C (high confidence), and very unlikely greater than 6°C (medium confidence)” (IPCC5, pp. 83–84)

The IPCC also provides a precise interpretation of terms like “extremely unlikely” and “medium confidence” (IPCC5, p. 36), which differs slightly from the interpretation in the previous Assessment Report (IPCC, 2007, p. 22) by explicitly taking into account the degree of “agreement” in the team about the evidence provided by each study. Given this interpretation, IPCC5 concludes that

$$\begin{aligned}\Pr(1.5 < \text{ECS} < 4.5) &= 0.67, \\ \Pr(\text{ECS} < 1.0) &< 0.05, \text{ and} \\ \Pr(\text{ECS} > 6.0) &< 0.10.\end{aligned}$$

In addition (pp. 75 and 817), they summarize information of experiments by the Coupled Model Intercomparison Project Phase 5 (CMIP5) who report a range 2.1–4.7 for the ECS, without however stating the likelihood of this range.

Assuming the ECS β to be lognormally distributed, so that $\log \beta \sim N(b_2, \sigma_2^2)$, we seek combinations (b_2, σ_2) such that the posterior probabilities closely match the probabilities in the IPCC report. There is no unique lognormal distribution that fits our data, but $b_2 = 1.07$ and $\sigma_2 = 0.53$ seems a reasonable approximation and is also in line with Hwang et al. (2013, Figure 4) where $b_2 = 1.071$ and $\sigma_2 = 0.527$.

³The sixth report (IPCC, 2021) is about to appear. So far, only the *Summary for Policy Makers* is publicly available.




Figure 1: Posterior distribution (in red) and the fifteen studies, lognormal distributions

The selected posterior distribution, plotted in Figure 1 (bold line, in red), satisfies

$$\begin{aligned} \Pr(1.5 < \beta < 4.5) &= 68.9\% \text{ (about 67\%),} \\ \Pr(0 < \beta < 1.0) &= 2.2\% \text{ (less than 5\%), and} \\ \Pr(\beta > 6.0) &= 8.7\% \text{ (less than 10\%),} \end{aligned}$$

in accordance with the IPCC report's conclusions. In addition, the interquartile range is $\Pr(2.04 < \beta < 4.17) = 50\%$ and there is a 1% probability of $\beta > 10.0$. The skewness of the distribution is well illustrated by the fact that the mode and median of β are quite different: The mode is $e^{b_2 - \sigma_2^2} = 2.20$, while the median is $e^{b_2} = 2.92$.

6.2 The data

In addition to the posterior we need information on the data. Our data consist of $n = 15$ studies mentioned in the IPCC5 report. Fourteen of these are contained in Figure 10.20 (p. 925) and Box 12.2 (p. 1110) of the report, and one study (Huber et al., 2011) is included in Figure 2 of Knutti et al. (2017) and referred to in various places of IPCC5.

The i th study produces a range (l_i, u_i) (lower and upper bound) with an associated probability p_i (typically 90% or 95%) for the ECS β , our parameter

Table 5: Data on climate sensitivity, Fifth Assessment Report

Study	p_i	Bounds		Lognormal approx.	
		l_i	u_i	y_i	σ_{0i}
1 Lindzen and Choi, 2011	95%	0.5	1.1	-0.30	0.20
2 Schmittner et al., 2011	90%	1.4	2.8	0.68	0.21
3 Aldrin et al., 2012	90%	1.2	3.5	0.72	0.33
4 Hargreaves et al., 2012	90%	1.0	4.2	0.72	0.44
5 Lewis, 2013	90%	2.0	3.6	0.99	0.18
6 Bender et al., 2010	95%	1.7	4.1	0.97	0.22
7 Otto et al., 2013	90%	0.9	5.0	0.75	0.52
8 Schwartz, 2012	90%	1.2	4.9	0.89	0.43
9 Lin et al., 2010	90%	2.8	3.7	1.17	0.08
10 Libardoni and Forest, 2011	90%	1.2	5.3	0.93	0.45
11 Köhler et al., 2010	90%	1.4	5.2	0.99	0.40
12 Olson et al., 2012	95%	1.8	4.9	1.09	0.26
13 Huber et al. 2011	67%	2.9	4.0	1.23	0.17
14 Holden et al., 2010	90%	2.0	5.0	1.15	0.28
15 Palaeosens, 2012	95%	1.1	7.0	1.02	0.47

of interest. Given the range (l_i, u_i) and the associated probability $\Pr(l_i < \beta_i < u_i) = p_i$ we can identify the parameters of the associated lognormal distributions, as follows. Since $\log \beta_i \sim N(y_i, \sigma_{0i}^2)$ in the i th study, we have

$$\Pr\left(\frac{\log l_i - y_i}{\sigma_{0i}} < z_i < \frac{\log u_i - y_i}{\sigma_{0i}}\right) = p_i, \quad z_i \sim N(0, 1), \quad (17)$$

and hence

$$y_i = \log(l_i u_i)^{1/2}, \quad \sigma_{0i} = \frac{\log(u_i/l_i)^{1/2}}{q_i}, \quad \Phi(q_i) = \frac{p_i + 1}{2}, \quad (18)$$

where Φ denotes the c.d.f. of the standard-normal distribution. In this way we end up with fifteen observations y_i with associated standard deviations σ_{0i} ; see Table 5. Our task is to find an estimator $b_0 \sim N(\log \beta, \sigma_0^2)$ based on the fifteen data $y_i \sim N(\log \beta, \sigma_{0i}^2)$. This task is not trivial as we shall see.

The IPCC5 report distinguishes between an instrumental period and a palaeoclimatic period. The instrumental period is the short period in the Earth's long history where direct instrumental records on climate are available, while the palaeoclimatic period is the long period preceding such

records. Of our fifteen studies some only use data from the instrumental period (studies 1, 3, 5–9, 12, 13), some only from the palaeoclimatic period (2, 4, 11, 14, 15), and some combine different lines of evidence (3, 10, 12). Studies 3 and 12 appear in both the groups instrumental and combination. This highlights the first of several problems: comparability.

The second problem is that the first study (Lindzen and Choi, 2011) has a big impact and is in essence an outlier. The IPCC raises doubts about the reliability of this study (IPCC5, pp. 923–924), but it has not removed the study from their report. Not only is the revealed value of y_i much lower than in the other studies, but the effect is much strengthened by the fact that the reported precision is high.

Third, we should take into account that the studies are correlated with each other. They all estimate the same parameter, probably using similar highly correlated data sets.

To gain further insight, Figure 1 presents the fifteen lognormal curves corresponding to the fifteen studies. From the figure several things become clear. First, that the first study (with the highest peak) is an outlier. And second, that what the data tell us is ambiguous. The modes range from 0.71 (1.89 if we exclude the first study) to 3.31, and the medians from 0.74 (1.98 if we exclude the first study) to 3.41.

Table 6: Estimates of b_0 and σ_0 from the data

b_0	σ_0	mode	quantiles		
			5%	50%	95%
	0.6	1.55	0.83	2.23	5.97
0.8	0.7	1.36	0.70	2.23	7.04
	0.8	1.17	0.60	2.23	8.30
	0.6	1.72	0.92	2.46	6.60
0.9	0.7	1.51	0.78	2.46	7.78
	0.8	1.30	0.66	2.46	9.17
	0.6	1.90	1.01	2.72	7.29
1.0	0.7	1.67	0.86	2.72	8.60
	0.8	1.43	0.73	2.72	10.13
	0.6	2.10	1.12	3.00	8.06
1.1	0.7	1.84	0.95	3.00	9.50
	0.8	1.58	0.81	3.00	11.20

Our task is to find a lognormal distribution which can be thought of as a reasonable representation of the data. This distribution must have a *larger* variance than the red curve in Figure 1 because theory prescribes that the posterior variance is smaller than the data variance (and also than the prior variance); see also the discussion at the end of Section 4.2.

Based on the graphical information and the constraint $\sigma_0 > 0.53$, we consider three values of σ_0 : 0.6, 0.7, and 0.8; and four values of b_0 : 0.8, 0.9, 1.0, and 1.1. Thus we shall consider twelve scenarios, where each scenario obeys the theoretical restrictions and the combination of scenarios covers what we believe is a credible range of the data parameters. Table 6 presents the mode and three quantiles (5%, 50%, and 95%) for the implied distributions. We see that the median ranges from 2.23 to 3.00, which is a little narrower than in the fourteen studies where the median ranges from 1.98 to 3.41 (if we exclude the first study where the median is 0.74). The scenarios also allow for a wide range of (right) tail behavior. The value β^* of β which defines the right tail (that is, where 5% of the distribution lies to the right of β^*) ranges from 5.97 to 11.20.

7 Climate sensitivity: revealed prior

Given the moments (b_0, σ_0) from the data and (b_2, σ_2) from the posterior, we can now discuss the revealed prior moments. The prior mean is given by (11),

$$b_1 = \frac{b_2 - \alpha_v b_0}{1 - \alpha_v}, \quad \alpha_v = \sigma_2^2 / \sigma_0^2, \quad (19)$$

and thus depends only on the posterior mean b_2 (which we set at 1.07), the data mean b_0 (which ranges from 0.8 to 1.1), and on the ratio of $\sigma_2 = 0.53$ and σ_0 (which ranges from 0.6 to 0.8).

In contrast, the prior standard deviation, also given in (11), is

$$\sigma_1 = \frac{\sigma_0 \sigma_2}{\sqrt{\sigma_0^2 - \sigma_2^2}} = \frac{\sigma_2}{\sqrt{1 - \alpha_v}}, \quad (20)$$

which, given $\sigma_2 = 0.53$, depends only on σ_0 or alternatively on α_v .

Table 7 shows the revealed prior means and standards deviations for each of the twelve scenarios (four values of b_0 , three values of σ_0). Given $(b_2, \sigma_2) = (1.07, 0.53)$ and the twelve selected pairs (b_0, σ_0) , we present the induced values of κ_m , κ_v , b_1 , and σ_1 .

Since b_2 must lie in-between b_0 and b_1 , it follows that $b_1 > 1.07$ if $b_0 < 1.07$ and $b_1 < 1.07$ if $b_0 > 1.07$. The lower is b_0 , the higher must be b_1 . If $b_0 = 0.8$ then we need a high prior mean, leading to a prior median of 3.6°C or more.

Table 7: Estimates of prior moments b_1 and σ_1 for given posterior moments $(b_2, \sigma_2) = (1.07, 0.53)$

data		κ_m	κ_v	prior		mode	quantiles		
b_0	σ_0			b_1	σ_1		5%	50%	95%
0.8	0.6	2.54	3.55	2.03	1.13	2.12	1.18	7.61	48.84
	0.7	1.79	1.34	1.43	0.81	2.17	1.10	4.19	15.91
	0.8	1.60	0.78	1.28	0.71	2.18	1.12	3.60	11.53
0.9	0.6	1.86	3.55	1.67	1.13	1.48	0.83	5.33	34.24
	0.7	1.44	1.34	1.30	0.81	1.90	0.96	3.66	13.91
	0.8	1.34	0.78	1.20	0.71	2.02	1.04	3.33	10.66
1.0	0.6	1.32	3.55	1.32	1.13	1.04	0.58	3.74	24.01
	0.7	1.16	1.34	1.16	0.81	1.66	0.84	3.20	12.16
	0.8	1.12	0.78	1.12	0.71	1.87	0.96	3.08	9.86
1.1	0.6	0.88	3.55	0.96	1.13	0.73	0.41	2.62	16.83
	0.7	0.94	1.34	1.03	0.81	1.45	0.74	2.80	10.64
	0.8	0.95	0.78	1.05	0.71	1.73	0.89	2.85	9.12

If the standard deviation in the data is relatively large ($\sigma_0 = 0.8$), then the standard deviation in the prior will be relatively small. But even then the right tail is substantial with a 5% probability that the ECS will exceed 10°C or worse. In our bolded scenario there is a 5% probability that the ECS will exceed 12°C.

As discussed in Section 3, the story can be summarized in terms on the key parameters κ_m and κ_v . What matters regarding κ_m is whether or not it is close to one. When the mean b_0 in the data and the mean b_2 in the posterior are approximately equal, but the variance σ_0^2 in the data and the variance σ_2^2 in the posterior are not approximately equal, then $\kappa_m \approx 1$ and the prior mean agrees with the data and the posterior. But when the variances σ_0^2 and σ_2^2 are approximately equal, but the means b_0 and b_2 are not, then κ_m is large (in absolute value). This is well illustrated in Table 7 where $\kappa_m = 1.12$ is close to one when $b_0 = 1.0$ is close to $b_2 = 1.07$, while $\sigma_0 = 0.8$ is not close to $\sigma_2 = 0.53$. At the other extreme, $\kappa_m = 2.54$ is not close to one, because $b_0 = 0.8$ is not close to $b_2 = 1.07$, while $\sigma_0 = 0.6$ is close to $\sigma_2 = 0.53$.

The prior standard deviation satisfies $\sigma_1 > \sigma_2 = 0.53$ (as it must), but σ_1 can be larger or smaller than σ_0 . In our set up, σ_1 can only take three values: 0.71 (small), 0.81 (medium), or 1.13 (large), depending on whether σ_0 equals

0.8 (large), 0.7 (medium), or 0.6 (small). Regarding κ_v what matters is whether it is small ($\kappa_v = 0.78$, strong prior information) or large ($\kappa_v = 3.55$, weak prior information). When the standard deviations σ_0 in the data and σ_2 in the posterior are approximately equal, then we see from Table 7 that the prior has only a small effect. This is represented by a large value of κ_v (3.55) and hence a large value of the prior variance σ_1 (1.13). The prior is then uninformative. But when the standard deviation in the data and the posterior are not close, then the prior has a big effect. This is represented by a small value of κ_v (0.78) and hence a small value of the prior standard deviation σ_1 (0.71). The prior is then informative.

So, what do these results tell us about the unknown prior of the IPCC? And in particular, how strong is this prior in view of the fact that a strong prior seems to be required to make any sense of the data.

Let us compare our prior with the conclusions (the posterior) of the previous IPCC report (IPCC, 2007). One would expect that the posterior of the previous study serves as the prior in the next study, at least approximately. The fourth report concludes about the ECS:

“It is likely to be in the range 2.0°C to 4.5°C with a best estimate of about 3.0°C, and is very unlikely to be less than 1.5°C. Values substantially higher than 4.5°C cannot be excluded, but agreement of models with observations is not as good for those values”.
(IPCC, 2007, p. 12)

If we fit a lognormal distribution with mean 1.1 and standard deviation 0.4, then these statements translate to a median of 3.0 and

$$\begin{aligned}\Pr(2.0 < \text{ECS} < 4.5) &= 0.69 > 0.67 \text{ (likely),} \\ \Pr(\text{ECS} < 1.5) &= 0.04 < 0.10 \text{ (very unlikely), and} \\ \Pr(\text{ECS} > 4.5) &= 0.16,\end{aligned}$$

which seems about right.

The posterior mean 1.10 from the previous report is a little lower than the prior mean in the current report (say, 1.16), but it is comparable. The same is true for the implied median, which is 3.0 in the previous report, and now a little higher, say 3.2. The main difference is in the standard deviation, which is 0.4 in the posterior distribution of the previous report and about 0.7 in the prior distribution of the current report. It appears that the IPCC scientists have agreed *a priori* on a value for the ECS between 3°C and 4.0°C, while judging the occurrence of a real disaster much more likely than the posterior in the previous report predicts. For example, while $\Pr(\text{ECS} > 6.0)$ equals 4.2% for the posterior in the previous report and 8.7% for the posterior in the

current report, the revealed prior gives a probability range between 14.8% ($b_0 = 1.05$, $\sigma_0 = 0.71$) and 58.3% ($b_0 = 2.03$, $\sigma_0 = 1.13$). These high numbers reflect the *a priori* view of the IPCC scientists; they are not based on new information becoming available in the current report.




Figure 2: Distributions of β based on data b_0 , prior b_1 , posterior b_2 , and previous report b_2^*

In Figure 2 we plot the four representative lognormal distributions of β for the data ($b_0 = 1.00$, $\sigma_0 = 0.70$), the prior ($b_1 = 1.16$, $\sigma_1 = 0.81$), the posterior ($b_2 = 1.07$, $\sigma_2 = 0.53$), and the posterior from the previous report ($b_2^* = 1.10$, $\sigma_2^* = 0.40$). These four graphs summarize our story. The smaller the variance of the lognormal distribution, the higher the top. The posterior b_2^* from the previous reports has the smallest variance, hence the highest top and the smallest probability of a really large temperature. There is a large discrepancy between the posterior b_2^* of the previous report and the prior b_1 of the current report, showing dynamic inconsistency. The main difference between the conclusions in the previous and the current report is that uncertainty has increased. Why has uncertainty increased in the presence of many new scientific studies, often claiming rather precise predictions? This can only be because the IPCC's prior judgement has become more cautious and less optimistic, perhaps rightly so, than the previous posterior justifies.

Since the Sixth Report IPCC6 (IPCC, 2021) is about to appear, let's briefly compare the posteriors of IPCC5 with the latest IPCC6. One would expect that improved knowledge about climate processes would lead to a

narrower range for the climate sensitivity parameter β . This is indeed the case. In IPCC6 (Section 7.5.5) the “likely” range for the climate sensitivity parameter is given as 2.5°C to 4.0°C, as compared to 1.5°C to 4.5°C in IPCC5. This corresponds to a slightly higher mode and a substantially smaller variance. The IPCC6 authors claim that the tighter range is caused by improved data (rather than by a sharper prior). It would be of interest to investigate if this is indeed the case, that is, if the impact of the prior has decreased and if so by how much.

8 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 and an extension (reparametrization) based on a monotonic transformation of the normal distribution, in our case the lognormal distribution. The (log)normality assumption can be relaxed without changing the underlying theory, but the expressions would become less transparent.

We developed the theory and we applied this theory to inflation forecasts by the BoE and climate sensitivity estimates by the IPCC. Our applications assume that the BoE and the IPCC behave in a Bayesian way and think in terms of priors and posteriors. This may not be so. But it doesn’t matter because we are all Bayesians in the sense that we all have priors. Scientific Bayesians have both explicit and implicit priors; scientific frequentists only have implicit priors. But we all have priors and the purpose of this paper has been to make these priors explicit. Of course, the field of application of the ideas developed in this paper is not confined to inflation or climate issues. It can be used more generally to determine people’s biases, for example politicians or scientists.

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 is 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 (as in the case of the BoE and the NIESR) 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?

Appendix: The lognormal distribution

Recall the fundamental formula for “completing the square”,

$$\frac{(y - \beta)^2}{\sigma_0^2} + \frac{(\beta - a_1)^2}{\sigma_1^2} = \frac{(\beta - a_2)^2}{\sigma_2^2} + \frac{(y - a_1)^2}{\sigma_0^2 + \sigma_1^2}, \quad (21)$$

where

$$a_2 = \frac{\sigma_0^2 a_1 + \sigma_1^2 y}{\sigma_0^2 + \sigma_1^2}, \quad \sigma_2^2 = \frac{\sigma_0^2 \sigma_1^2}{\sigma_0^2 + \sigma_1^2}.$$

Multiplying both sides by $-1/2$ and taking exponentials gives

$$f(y; \beta, \sigma_0^2) f(\beta; a_1, \sigma_1^2) = f(\beta; a_2, \sigma_2^2) f(y; a_1, \sigma_0^2 + \sigma_1^2), \quad (22)$$

where

$$f(x; \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp -\frac{1}{2} \left(\frac{x - \mu}{\sigma} \right)^2$$

denotes the normal density. This, of course, is just Bayes’ formula

$$f(y|\beta) f(\beta) = f(\beta|y) f(y)$$

and shows that a normal likelihood plus a normal prior results in a normal posterior.

Now let g denote the lognormal density

$$g(x; \mu, \sigma^2) = \frac{h'(x)}{\sigma\sqrt{2\pi}} \exp -\frac{1}{2} \left(\frac{h(x) - \mu}{\sigma} \right)^2 \quad (x > 0).$$

where $h(x) = \log x$ and consequently $h'(x) = 1/x$. Then it follows immediately from (21) that, for $\beta > 0$,

$$f(y; h(\beta), \sigma_0^2) g(\beta; a_1, \sigma_1^2) = g(\beta; a_2, \sigma_2^2) f(y; a_1, \sigma_0^2 + \sigma_1^2). \quad (23)$$

Hence, a normal likelihood plus a lognormal prior results in a lognormal posterior, so that prior and posterior are conjugate distributions. The function h can be any monotonic transformation of x , not necessarily $\log x$.

Another way of arriving at this result is to realize that the lognormal distribution is not really a new distribution but rather a reparametrization. The parameter of interest remains β , but the analysis is performed on $\log \beta$ (more generally on $h(\beta)$). So we have a normal likelihood with mean $\log \beta$ and a normal prior on $\log \beta$, resulting in a normal posterior on $\log \beta$.

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 Number JP20H01477 and the program of the Joint Usage/Research Center for Behavioral Economics at ISER, Osaka University.

References

- Aldrin, M., M. Holden, P. Guttorp, R. B. Skeie, G. Myhre, and T. K. Berntsen (2012). Bayesian estimation of climate sensitivity based on a simple climate model fitted to observations of hemispheric temperatures and global ocean heat content. *Environmetrics*, 23, 253–271.
- Andrade, P. and H. Le Bihan (2013). Inattentive professional forecasters. *Journal of Monetary Economics*, 60, 967–982.
- Bank of England (2016a). *Inflation Report—May 2016*, London, UK.
- Bank of England (2016b). *Inflation Report—August 2016*, London, UK.
- Bank of England (2020a). *Monetary Policy Report—January 2020*, London, UK.
- Bank of England (2020b). *Monetary Policy Report—May 2020*, London, UK.
- Bank of England (2020c). *Monetary Policy Report—August 2020*, London, UK.
- Bender, F. A.-M., A. M. L. Ekman, and H. Rodhe (2010). Response to the eruption of Mount Pinatubo in relation to climate sensitivity in the CMIP3 models. *Climate Dynamics*, 35, 875–886.

- Boero, G., J. Smith, and K. F. Wallis (2015). The measurement and characteristics of professional forecasters' uncertainty. *Journal of Applied Econometrics*, 30, 1029–1046.
- Candia, B., O. Coibion, and Y. Gorodnichenko (2021). The inflation expectations of U.S. firms: Evidence from a new survey. *NBER Working Paper*, 28836.
- Delignette-Muller, M.L, and C. Dutang (2015). Fitdistrplus: An R package for fitting distributions. *Journal of Statistical Software*, 64, 1–34.
- Galvão, A.B., A. Garratt, and J. Mitchell (2021). Does judgment improve macroeconomic density forecasts? *International Journal of Forecasting*, 37, 1247–1260.
- Hantzsche, A., and G. Young (2020). Prospects for the UK Economy. *National Institute Economic Review*, 251, F4–F34.
- Hargreaves, J. C., J. D. Annan, M. Yoshimori, and A. Abe-Ouchi (2012). Can the last glacial maximum constrain climate sensitivity? *Geophysical Research Letters*, 39, L24702.
- Holden, P. B., N. R. Edwards, K. I. C. Oliver, T. M. Lenton, and R. D. Wilkinson (2010). A probabilistic calibration of climate sensitivity and terrestrial carbon change in GENIE-1. *Climate Dynamics*, 35, 785–806.
- Huber, M., I. Mahlstein, M. Wild, J. Fasullo, and R. Knutti (2011). Constraints on climate sensitivity from radiation patterns in climate models. *Journal of Climate*, 24, 1034–1052.
- Hwang, C., F. Reynès, and R. S. J. Tol (2013). Climate policy under fat-tailed risk: An application of DICE. *Environmental and Resource Economics*, 56, 415–436.
- IPCC (2007). *Climate Change 2007: The Physical Science Basis*. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK and New York, NY, USA.
- IPCC (2013). *Climate Change 2013: The Physical Science Basis*. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK and New York, NY, USA.

- IPCC (2021). *Climate Change 2021: The Physical Science Basis*. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, Cambridge, UK and New York, NY, USA, in press.
- Kirby, S., O. Carreras, J. Meaning, and R. Piggott (2016a). Prospects for the UK Economy. *National Institute Economic Review*, 236, 82–107.
- Kirby, S., O. Carreras, J. Meaning, R. Piggott, and J. Warren (2016b). Prospects for the UK Economy. *National Institute Economic Review*, 237, F42–F71.
- Knutti, R., M. A. A. Rugenstein, and G. C. Hegerl (2017). Beyond equilibrium climate sensitivity. *Nature Geoscience*, 10, 727–736.
- Köhler, P., R. Bintanja, H. Fischer, F. Joos, R. Knutti, G. Lohmann, and V. Masson-Delmotte (2010). What caused Earth’s temperature variations during the last 800,000 years? Data-based evidence on radiative forcing and constraints on climate sensitivity. *Quaternary Science Reviews*, 29, 129–145.
- Lenoël, C., and G. Young (2020). Prospects for the UK Economy. *National Institute Economic Review*, 252, F10–F43.
- Lenoël, C., R. Macqueen, and G. Young (2020). Prospects for the UK Economy. *National Institute Economic Review*, 253, F4–F34.
- Lewis, N. (2013). An objective Bayesian, improved approach for applying optimal fingerprint techniques to estimate climate sensitivity. *Journal of Climate*, 26, 7414–7429. doi:10.1175/JCLID-12-00473.1.
- Libardoni, A. G., and C. E. Forest (2011). Sensitivity of distributions of climate system properties to the surface temperature dataset. *Geophysical Research Letters*, 38, L22705. Correction in: *Geophysical Research Letters* (2013), 40, 2309–2311. doi:10.1002/grl.50480.
- Lin, B., L. Chambers, P. Stackhouse Jr., B. Wielicki, Y. Hu, P. Minnis, N. Loeb, W. Sun, G. Potter, Q. Min, G. Schuster, and T.-F. Fan (2010). Estimations of climate sensitivity based on top-of-atmosphere radiation imbalance. *Atmospheric Chemistry and Physics*, 10, 1923–1930.
- Lindzen, R. S., and Y. S. Choi (2011). On the observational determination of climate sensitivity and its implications. *Asia-Pacific Journal of Atmospheric Sciences*, 47, 377–390.

- McNees, S. K. (1990). The role of judgement in macroeconomic forecasting accuracy. *International Journal of Forecasting*, 6, 287–299.
- Olson, R., R. Sriver, M. Goes, N. M. Urban, H. D. Matthews, M. Haran, and K. Keller (2012). A climate sensitivity estimate using Bayesian fusion of instrumental observations and an Earth system model. *Journal of Geophysical Research: Atmospheres*, 117, D04103.
- Otto, A., F. E. L. Otto, O. Boucher, J. Church, G. Hegerl, P. M. Forster, N. P. Gillett, J. Gregory, G. C. Johnson, R. Knutti, N. Lewis, U. Lohmann, J. Marotzke, G. Myhre, D. Shindell, B. Stevens, and M. R. Allen (2013). Energy budget constraints on climate response. *Nature Geoscience*, 6, 415–416.
- Palaeosens Project Members (2012). Making sense of palaeoclimate sensitivity. *Nature*, 491, 683–691.
- Schmittner, A., N. M. Urban, J. D. Shakun, N. M. Mahowald, P. U. Clark, P. J. Bartlein, A. C. Mix, and A. Rosell-Melé (2011). Climate sensitivity estimated from temperature reconstructions of the last glacial maximum. *Science*, 334, 1385–1388. Response to comment in: *Science* (2012), 337, 1294.
- Schwartz, S. E. (2012). Determination of Earth’s transient and equilibrium climate sensitivities from observations over the twentieth century: Strong dependence on assumed forcing. *Surveys in Geophysics*, 33, 745–777.
- Turner, D. S. (1990). The role of judgement in macroeconomic forecasting. *Journal of Forecasting*, 9, 315–345.
- Wallis, K. F. (2014). The two-piece normal, binormal, or double Gaussian distribution: Its origin and rediscoveries. *Statistical Science*, 29, 106–112.
- Winkler, R. (1968). The consensus of subjective probability distributions. *Management Science* 15, B61–B75.