Revealing priors from posteriors: Supplementary material

Masako Ikefuji
OSIPP, Osaka University, Japan

Jan R. Magnus
Department of Econometrics and Data Science,
Vrije Universiteit Amsterdam and Tinbergen Institute, The Netherlands

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

June 1, 2022

This file contains supplementary material to the paper *Revealing priors from posteriors* by the same three authors.

1 Introduction

This supplementary file contains a (much) fuller treatment of the illustration contained in Section 3 of the associated paper. This illustration concerns inflation forecasts by the Bank of England, especially under the uncertainty about Brexit and the Covid crisis (lockdown). We are interested to recover the prior beliefs of the decision maker, in this case the Bank of England.

In Section 2 we analyze how to recover the prior from the data and the posterior within the framework of the normal distribution, in the special case when there is only one parameter of interest. Equations (1) and (2) are taken from the paper, but the introduction of α and κ is new. Sections 3 and 4 contain our detailed analysis of the Bank of England's forecast of the interest rate.

2 One parameter of interest

In the special 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: $b_0 \sim \mathrm{N}(\beta, \sigma_0^2)$. If we add a prior $\beta \sim \mathrm{N}(b_1, \sigma_1^2)$, then we obtain the posterior $\beta \sim \mathrm{N}(b_2, \sigma_2^2)$, where

$$b_2 = \frac{\sigma_0^2 b_1 + \sigma_1^2 b_0}{\sigma_0^2 + \sigma_1^2}, \qquad \sigma_2^2 = \frac{\sigma_0^2 \sigma_1^2}{\sigma_0^2 + \sigma_1^2}. \tag{1}$$

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

$$b_1 = \frac{\sigma_0^2 b_2 - \sigma_2^2 b_0}{\sigma_0^2 - \sigma_2^2}, \qquad \sigma_1^2 = \frac{\sigma_0^2 \sigma_2^2}{\sigma_0^2 - \sigma_2^2}, \tag{2}$$

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

Defining α_m and α_v implicitly by

$$b_2 = \alpha_m \, b_0, \qquad \sigma_2^2 = \alpha_v \, \sigma_0^2, \tag{3}$$

we can rewrite (2) as

$$b_1 = \kappa_m \, b_0, \qquad \sigma_1^2 = \kappa_v \, \sigma_0^2, \tag{4}$$

where

$$\kappa_m = \frac{\alpha_m - \alpha_v}{1 - \alpha_v}, \qquad \kappa_v = \frac{\alpha_v}{1 - \alpha_v} \tag{5}$$

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

$$1 - \kappa_m = \frac{1 - \alpha_m}{1 - \alpha_v}, \tag{6}$$

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 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 $b_1 \approx b_2$.

3 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 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". In their study they need to separately identify the adjustment by the forecaster and the unadjusted mechanical density forecasts.

Our purpose is to investigate 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 the equations in Section 2 we can 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.

3.1 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 $\mathrm{N}(b_2, \sigma_2^2)$ 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.

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.

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

[1] 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.

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.

3.2 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 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 likely 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 $\mathrm{N}(b_0, \sigma_*^2)$. 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 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. An attempt to resolve this issue, based on the presence of positive correlation, was recently provided by Magnus and Vasnev (2022).

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.

4 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

[2]See https://www.niesr.ac.uk/business-conditions-forum.

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

	Data		Priors			
			BoE		NIESR	
Year/Quarter	b_0	σ_0	b_1	σ_1	b_1	σ_1
Referendum						
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
Lockdown						
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

4.1 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.

4.2 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.

4.3 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.

In Table 4 we present the key parameter ratios, as discussed in Section 2. We focus on the parameters related to the prior: κ_m and κ_v. Recall that

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

	Data		BoE				NIESR			
Year/Quarter	b_0	σ_0	α_m	α_v	κ_m	κ_v	α_m	α_v	κ_m	κ_v
Referendum										
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
Lockdown										
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

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.

References

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.

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.

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.

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.

Magnus, J.R. and A.L. Vasnev (2022). On the uncertainty of a combined forecast: The critical role of correlation. *International Journal of Forecasting*, revision submitted.

McNees, S. K. (1990). The role of judgement in macroeconomic forecasting accuracy. *International Journal of Forecasting*, 6, 287–299.

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.