UNIVERSITY OF LEICESTER

Department of Economics

Extracting the Information Shocks from the Bank of England Inflation Density Forecasts

Carlos Diaz Vela, University of Leicester

Working Paper No. 16/13

Extracting the Information Shocks from the Bank of England Inflation Density Forecasts

Carlos Díaz*

University of Leicester

Abstract

This paper shows how to extract the density of the shocks of information perceived by the Bank of England between two consecutive releases of its inflation density forecasts. These densities are used to construct a new measure of *ex ante* inflation uncertainty, and a measure of news incorporation into subsequent forecasts. Also dynamic tests of point forecast optimality is constructed. It is shown that inflation uncertainty as perceived by the Bank was decreasing before the financial crisis, increasing sharply during the period 2008-2011. Since then, uncertainty seems to have stabilized, but it remains still above its pre-crisis levels. Finally, it is shown that forecast optimality is lost at some points during the financial crisis, and that there are more periods of optimal forecasts in long term than in short term forecasting. This could be also interpreted as that short term forecasts are subject to profound revisions.

Keywords: Inflation, density forecast, uncertainty, revisions, optimal forecasts.

JEL codes: C22, C53, C63, E31, E37, E58.

*Email: cdv7@leicester.ac.uk. This research used the ALICE High Performance Computing Facility at the University of Leicester. Thanks to Prof Wojciech Charemza and Prof José Luis Gallego for valuable comments on an earlier draft of this paper. Thanks also to the participants in the seminar at the University of Cantabria, the 36th International Symposium on Forecasting, and 3rd Annual Conference of the International Association for Applied Econometrics for their feedback.

1 Introduction

Density forecasting has become a useful tool for researchers and policymakers as a way to evaluate the uncertainty surrounding key economic and financial variables. In particular, inflation density forecasts are used by some central banks as informative tools in the process of monetary policy implementation. The Bank of England regularly produces inflation density forecasts presented to the public in the form of the so-called *fan charts*. These forecasts are updated each quarter to reflect how the information gathered by the Bank is expected to affect inflation in the following quarters. The main aim of this paper is to develop a method to extract the density of these shocks of information by comparing the revisions of density forecasts. These densities are informative of the effect that the Bank expects new information will have in future inflation. In this sense, the variance of these shocks of information can be interpreted as a measure of the expected (*ex-ante*) uncertainty can be constructed from them which could be used as an input in forward looking models. This exercise can be done for different forecast horizons, constructing something resembling to a *term structure* of uncertainty.

A second use of the density of the shocks of information is to evaluate the optimality of point forecasts along the lines of Alessi et al. (2014), who stress the importance of assessing the forecast accuracy of central banks' forecast ability in the light of the recent financial crisis. There's a vast literature on the evaluation of the Bank of England point and density forecasts. Wallis (2004), Clements (2004) or Dowd (2007) evaluate the forecasting ability of the fan charts showing that, although they are more accurate for short term forecasting, uncertainty is generally overestimated. Internal assessment of the fan charts within the Bank of England has been performed by Britton et al. (1998), Elder et al. (2005) or more recently Fawcett et al. (2015). More recently Gneiting and Ranjan (2011) or Galbraith and van Norden (2012) re evaluate these forecasts pointing again to the overestimation of risks and the lack of resolution. All these tests are static in the sense that they check performance over a time span. In this paper it is shown how the density of information shocks can be used to check dynamically a set of regularities of optimal point forecasts noted by Mankiw and Saphiro (1986), Nordhaus (1987) or Patton and Timmermann (2012), namely, the unbiasedness of optimal point forecasts and the decreasing variance of forecast errors as the forecasted outcome approaches in time. Quarters in which these conditions are not met within the sample period are identified. This could help identify periods in which optimality is lost and help improve forecasting methods. Although it will be left for future work, this method could be used to check density forecast efficiency, extending the idea of Mincer and Zarnowitz (1969) as explored by Mitchell (2008).

The main results of the paper, which focus on the period 1998-2015, show that inflation uncertainty was decreasing during the period before the financial crisis, with uncertainty being higher for long term forecasts than for short therm. Uncertainty increased rapidly during the period 2007-2011, having stabilized now at levels higher than its pre-crisis values. It is also shown that new information forms around 60% of the information contained in revisions of density forecasts for short and long term forecasting. Finally, it is shown that point forecasts can be considered optimal except during some quarters during the financial crisis. Also, short term forecasts are subject to deeper revisions more often than long term forecasts. Such periods of significant changes in point forecasts are identified.

The paper is structured as follows: section 2 describes the method used to extract the density of the socks of information via convolutions, defines a measure of news absorbtion and proposes dynamic tests of forecast optimality; section 3 applies this technique to the Bank of England density forecasts, and section 4 concludes.

2 An evaluation of inflation density forecasts

2.1 Recovering the densities of information shocks via convolutions

Let z_{t+h} be the random variable being forecasted and let $f_{t+h}(\cdot)$ be its unknown density function. At moment t a forecast for this random variable is made, say $z_{t+1}(h-1)$, in the form of a density forecast, $\hat{f}_{t|h}(\cdot)$, constructed using all the information available to the forecaster up to time t -let I_t denote this information set. At time $t+1$ a new forecast for the same random variable is produced, say $z_{t+1}(h-1)$ with density $\hat{f}_{t+1|h-1}(\cdot)$. This new density forecast is made using all information available up to $t+1$ (i.e. using information set I_{t+1}) so, if forecasts are efficient, any difference between $\hat{f}_{t+1|h-1}(\cdot)$ and $\hat{f}_{t|h}(\cdot)$ should reflect the effect of the incorporation of information made available between t and $t+1$. We can, then, decompose the revision $z_{t+1}(h-1)$ using the well known forecast update equation as

$$z_{t+1}(h-1) = z_t(h) + \epsilon_{t+1}(h-1), \tag{1}$$

where $\epsilon_{t+1}(h-1)$ can be interpreted as the information or news shock perceived between periods t and $t+1$. Notice that both density forecasts can be misspecified, so $\epsilon_{t+1}(h-1)$ should be interpreted as a forecast revision, not as an econometric or a forecast error. The forecasts are subjective in the sense that they only incorporate information perceived by the forecaster. If the efficient forecast at time $t+1$ should be made using information set Ω_{t+1} and $I_{t+1} \subset \Omega_{t+1}$, then $\hat{f}_{t+1|h-1}(\cdot)$ would be a biased estimate of $f_{t+h}(\cdot)$ -the same could happen to $\hat{f}_{t|h}(\cdot)$. In this case, the density of $\epsilon_{t+1}(h-1)$ would only reflect changes in density forecasts due to perceived information which need not be the full information set. Nonetheless, information extracted from the density of these revisions or information shocks can be highly relevant, as they provide with *ex-ante* measures of perceived risks of the forecasted outcome: a high variance of this shock would imply that the forecaster expect the outcome to become more uncertain, and nonnegative skewness could be interpreted as an increase in the perceived risk of extreme event outcomes.

Although conceptually similar to the Functional Autoregressive method used by Chaudhuri et al. (2016) to obtain density forecasts of UK inflation rates, the method proposed here is based on the notion of convolution. Notice that $\epsilon_{t+1}(h-1) = z_{t+1}(h-1) - z_t(h)$, where $z_{t+1}(h-1)$ and $z_t(h)$ are two random variables with known densities. If they were uncorrelated, the density function of the shock $\epsilon_{t+1}(h-1)$ could be obtained as the convolution of a difference of random variables[1]. However, $z_{t+1}(h-1)$ will in general be correlated with $z_t(h)$, as the information sets on which those forecasts are constructed overlap (in general $I_t \subset I_{t+1}$). In that case, the density

[1]If X and Y are two independent random variables with densities $f_X(\cdot)$ and $f_Y(\cdot)$ respectively, then $f_X * f_Y(z) = \int_{-\infty}^{\infty} f_X(x) f_Y(x-z)dx$.

of the information shock should be obtained from the the joint density of $z_{t+1}(h-1)$ and $z_t(h)$, say $f_{z_{t+1}(h-1)z_t(h)}(\cdot,\cdot)$ as

$$\hat{f}_{\epsilon,h-1}(z) = \int_{-\infty}^{\infty} f_{z_{t+1}(h-1)z_t(h)}(x, x-z)dx. \tag{2}$$

As the forecasts are made at different times and with possibly different models, this joint density is not normally available. All that is available are the original density forecasts, that can be thought of as the marginal densities derived from this joint distribution. The proposed method consists of reconstructing the joint density of $z_{t+1}(h-1)$ and $z_t(h)$ from its known marginal densities, and then recovering the density of the information shock using (2). To this aim, copulas are helpful tools. A copula is a distribution function that models the dependency structure of two or more random variables with known marginal densities. More formally, let $\hat{F}_{t+1|h-1}(\cdot)$ and $\hat{F}_{t|h}(\cdot)$ be the marginal cumulative distribution functions (cdfs) of $z_{t+1}(h-1)$ and $z_t(h)$ respectively. Following Sklar (1959), if these distribution functions are continuous, there exists a unique function $C(\cdot,\cdot|\theta) : [0,1]^2 \to [0,1]$ such that, in this bivariate case, the copula function is reduced to a bivariate cumulative distribution function with uniform marginals:

$$\tilde{F}_{z_{t+1}(h-1)z_t(h)}(x_1, x_2) = C\left[\hat{F}_{t+1|h-1}(x_1), \hat{F}_{t|h}(x_2)|\theta\right], \qquad x_1, x_2 \in \Re, \tag{3}$$

where $\tilde{F}_{z_{t+1}(h-1)z_t(h)}(x_1, x_2)$ is an estimate of the joint cumulative distribution function of $z_{t+1}(h-1)$ and $z_t(h)$, and $\theta \in \Theta^m$ is an m-vector of parameters that controls the degree of dependence between both random variables. The joint probability density function can be easily obtained from (3) as

$$\tilde{f}_{z_{t+1}(h-1)z_t(h)}(x_1, x_2) = c\left[\hat{F}_{t+1|h-1}(x_1), \hat{F}_{t|h}(x_2)|\theta\right] \hat{f}_{t+1|h-1}(x_1)\hat{f}_{t|h}(x_2), \tag{4}$$

where $c(\cdot,\cdot|\theta)$ is the copula density. A plethora of parametric copulas has been proposed in the literature to model the dependence between random variables (see e.g. Joe, 1997, for a review). Parametric copulas are convenient as they have closed form expressions for the distribution function $C(\cdot,\cdot)$ and the corresponding copula density $c(\cdot,\cdot)$, and different families allow to accommodate a wide range of dependence structures. In the bivariate case, closed form expressions of the Kendall and Spearman rank correlation coefficients can be obtained for each copula as a function of the copula parameters. However, parametric copulas depend on the vector of parameters θ that need to be estimated. A common procedure to estimate these parameters is to use the Inference Function for Marginal procedure described in Joe and Xu (1996), consisting on estimating the copula parameters by maximum likelihood using (4), having previously characterized the marginal distributions. This method would give a single estimate of the parameter vector θ for the entire sampling period, making necessary the strong assumption that the dependence structure between forecasts is constant over time. In order to overcome this restriction a minimum distance estimator of the copula parameters is proposed here. Notice that, given that $\epsilon_{t+1}(h-1)$ and $z_t(h)$ are independent by construction, an approximation to $\hat{f}_{t+1|h-1}(\cdot)$ can be obtained by the convolution of the density of the forecast at time t and the estimate of the density of the shock as

$$\hat{g}_{t+1|h-1}(z|\theta) = \int_{-\infty}^{\infty} \hat{f}_{t|h}(z-x)\hat{f}_{\epsilon,h-1}(x)dx. \tag{5}$$

As $\hat{f}_{\epsilon,h-1}(\cdot)$ is constructed using (4), $\hat{g}_{t+1|h-1}(z)$ depends on the copula parameter vector θ. The proposed estimation method for θ consists on minimizing the Kullback-Leibler information criterion distance measure (see e.g. Mitchell and Hall, 2005) that measures the discrepancy between the 'true' density $\hat{f}_{t+1|h-1}(\cdot)$ and the approximated one $\hat{g}_{t+1|h-1}(z)$. Thus, the proposed estimator for θ is

$$\hat{\theta}_{MDE} = \arg\min_{\theta} \int_{-\infty}^{\infty} \hat{f}_{t+1|h-1}(x) \ln\left[\frac{\hat{f}_{t+1|h-1}(x)}{\hat{g}_{t+1|h-1}(x|\theta)}\right] dx. \tag{6}$$

The copula parameters θ are chosen so that the expected difference of the log-scores of both densities are minimal. This estimator falls into the class of minimum distance estimators described in Basu et al. (2011) and, although the finite sample properties of this estimator are left for further research, the estimator is consistent. This proposed estimator has two main practical advantages with respect to the IFM method: first, it can be performed dynamically, so that every time a revision of a forecast is made a new estimate of the copula parameters is obtained. Also, the method does not need data to estimate the parameters. This is of great practical relevance. If, as commented earlier, we interpret the variance of the information shock as a measure forecast uncertainty of z_{t+h}, the proposed method provides a way to obtain measure of *ex-ante* uncertainty calculated at the moment of releasing an update of the density forecast. By obtaining this densities for different forecast horizons, h, we can study the 'term structure' of uncertainty for a fixed point event.

2.2 A measure of news absorbtion

The decomposition (5) evaluated under the optimal value of θ can also be used to construct a measure of information incorporation to updates of the density forecast. Assume that in (5) $\hat{g}_{t+1|h-1}(z) = \hat{f}_{t+1|h-1}(z)$, i.e., that the fit is perfect. The convolution theorem (see for example Bracewell, 1999) states that the Fourier transform of a convolution can be decomposed as the product of the Fourier transform of the densities involved in the convolution[2], so that

$$FT[\hat{f}_{t+1|h-1}(z)] \equiv s_{t+1|h-1}(\lambda) = s_{t|h}(\lambda)s_{\epsilon,h}(\lambda), \qquad \lambda \in [-\pi, \pi]. \tag{7}$$

The Fourier transform allows us to express the behaviour of the density at frequency λ as the product of both components at the same frequency. Making this relation additive by taking logarithms in (7), we can construct an approximate measure of the relative importance of ϵ on the updated density forecast at frequency λ:

$$d_\epsilon(\lambda) = \frac{\ln[s_{\epsilon,h}(\lambda)]}{\ln[s_{t+1|h-1}(\lambda)]}, \qquad 0 \leq d_\epsilon(\lambda) \leq 1.$$

Finally, a measure of the absorbtion of new information into the updated density forecast can be obtained aggregating $d_\epsilon(\lambda)$ at all frequencies as

$$Abs_{t+1|h-1} = \frac{1}{2\pi} \int_{-\pi}^{\pi} d_\epsilon(\lambda) d\lambda \tag{8}$$

[2]This can be easily proven as $FT[f_{t+1|h-1}(z)] = \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f_{t|h}(z-x)\hat{f}_{\epsilon,h}(x)dx \right] e^{-i\lambda z} dz = \int_{-\infty}^{\infty} \hat{f}_{\epsilon,h}(x)e^{-i\lambda z} \left[\int_{-\infty}^{\infty} f_{t|h}(z-x)e^{-i\lambda(z-x)}d(z-x) \right] dx = s_{t|h}(\lambda)s_{\epsilon,h}(\lambda)$.

This measure, which ranges between zero and one, gives an approximate estimate of the proportion of new information incorporated to the revision of the inflation forecast.

2.3 Dynamic evaluation of forecasts

In this section it will be shown how the densities of the information shocks, interpreted as forecast revisions, can be used to check two well known regularities of optimal forecasts: unbiasedness and the monotonic increase of the variance of forecast revision with the forecast horizon. Under a quadratic loss function the optimal forecast for z_{t+h} made at moment t is obtained as

$$\hat{z}_t(h) = \underset{\tilde{z} \in Z}{\arg\min}\, E[(z_{t+h} - \tilde{z})^2 | I_t], \qquad Z \subset \Re.$$

The point forecast $\hat{z}_t(h)$ is said to be efficient if it incorporates all information available to the forecaster at time t. Testing for point forecast efficiency is commonly performed using the popular Mincer and Zarnowitz (1969) regression-based test, or extensions to more general settings proposed more recently by, for example, Rossi and Sekhposyan (2015). Efficiency tests based on forecast revisions have also been proposed by Mankiw and Saphiro (1986) or Patton and Timmermann (2012). All these tests are static in the sense that the efficiency is checked for a given sample. However, it would be useful to check whether the properties of the forecasts are violated during the sampling period. It is know from Nordhaus (1987) that forecast revisions should be unpredictable, and therefore have conditional zero mean, and be independent of past revisions. Although the second part of the statement could prove hard to check dynamically, the first one (namely, the unbiasedness property of forecasts) can be checked using the density of the shocks of news derived earlier. Notice that from (1) $\epsilon_{t+1}(h-1)$ is a forecast revision, and that

$$E[\epsilon_{t+1}(h-1)] = E\{[z_{t+h} - \hat{z}_t(h)] - [z_{t+h} - \hat{z}_{t+1}(h-1)]\} = E[\hat{e}_t(h) - \hat{e}_{t+1}(h-1)], \qquad (9)$$

where $\hat{z}_t(h)$ is the optimal point forecast and the h-step forecast errors is defined as $\hat{e}_t(h) = z_{t+h} - \hat{z}_t(h)$. Under a square loss function, the forecasts should be unbiased, so that $E[\epsilon_{t+1}(h-1)|I_t] =$

$E[\epsilon_{t+1}(h-1)|I_{t+1}] = 0$, i.e., $\epsilon_{t+1}(h-1)$ must have a conditional and unconditional mean of zero. A nonzero mean would imply that a significant revision has been made, implying that the previous forecast did not incorporate all available information, or that new information is expected to affect substantially the forecasted variable. Therefore, unbiasedness can be though of as a necessary but not sufficient condition for efficiency. A dynamic test of unbiasedness can be defined as

$$H_0 : \mu_{t+1|h-1} = 0 \qquad \text{against} \qquad H_1 : \mu_{t+1|h-1} \neq 0$$

where $\hat{\mu}_{t+1|h-1} = E(\epsilon_{t+1}(h-1))$. This can be easily performed as for each revision the density of $\epsilon_{t+1}(h-1)$ is known. A rejection of the null can, then, be useful to identify periods in which efficiency is potentially breached.

A second well known regularity that can be tested dynamically using the densities of the shocks of information is the well known fact (see e.g. Mankiw and Saphiro, 1986) that the variance of the forecast revisions decreases as the forecast horizon decreases. This regularity can be intuitively explained by the fact that, as we approach the outcome date, the variance of the forecast error decreases indicating a decrease in uncertainty driven by the incorporation of information into the forecast, while the variance of the forecast itself increases (see e.g. Granger and Newbold, 1986). This trade-off between the variances of the forecast and the forecast error is well known, and its implications to forecast revisions is studied in Isiklar and Lahiri (2007) and is the base used by Patton and Timmermann (2012) to construct their rationality test. This regularity can be tested dynamically using the previous setting. We know that if forecasts are efficient $Var[\epsilon_t^2(h)] \geq Var[\epsilon_{t+j}^2(h-j)]$ for $j > 0$, and let $VR_{t,h} = Var[\epsilon_t^2(h)]/Var[\epsilon_{t+h-1}^2(1)]$ for $h > 1$, the testing problem is

$$H_0 : VR_{t,h} \geq 1 \qquad \text{against} \qquad H_1 : VR_{t,h} < 1,$$

where in d_h the variance of the revisions are compared to the revision of the forecast one period before the outcome date. As forecast errors at horizon h have an $MA(h-1)$ correlation structure, this should prevent problems derived from the correlation between numerator and denominator in d_j. Notice that the null distribution of d_h is unknown, so it must be simulated. Realizations of $\epsilon_t^2(h)$ are obtained from its distribution using the Inverse Transform Sampling method[3]. Again, a non-rejection of the null hypothesis does not necessarily imply a lack of efficient incorporation of information into subsequent revisions of forecasts. These two tests proposed in this section can act as warning signals about periods in which efficiency is lost, but a full dynamic efficiency test should be derived. The approach proposed in this paper appears a promising method of achieving this.

3 Evaluating the Bank of England fan charts

3.1 Extraction of shocks

Since 1993, and following the introduction of inflation targeting for UK monetary policy in 1992, the Bank of England produces its inflation forecasts[4] in the form of density forecasts. Each quarter the Bank releases its forecasts for the next 12 quarters (8 quarters until the second quarter of 2013) under two assumptions: constant official interest rates and the assumption that interest rates will follow market interest rates. The graphical representation of these forecasts is known as the *fan chart* and it depicts the Bank's balance of risks around the central tendency of inflation. The Bank reports the Monetary Policy Committee judgement on the future evolution of the central tendency (mode, μ), uncertainty (σ) and skewness[5] (η) of inflation. The value of these parameters and a report containing all the extra assumptions underlying the construction of the fan charts are publicly available on the Bank of England's website[6]. Once these parameters are decided, the

[3]The Inverse Transform Sampling method is suitable to obtain values of a random variable X when the only information available is its cumulative distribution function, $F(x)$. The method works as follows: obtain a realization u of a standard uniform distribution on the interval [0,1]. The draw from X is the value x such that $F(x) = u$.

[4]RPIX inflation until December 2003 and CPI inflation since.

[5]Skewness is measured as the difference between the mean and the mode.

[6]`http://www.bankofengland.co.uk/publications/Pages/inflationreport/irprobab.aspx`

Figure 1: Moments of errors

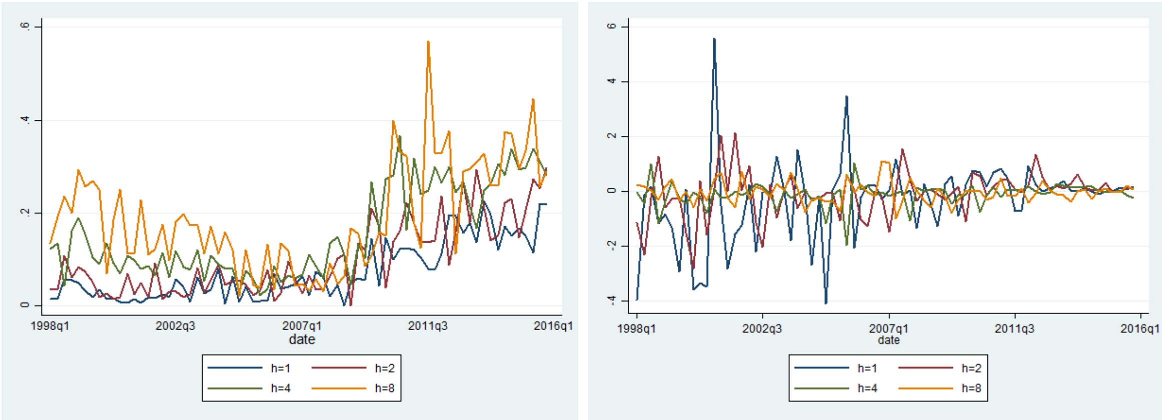

(a) variance

(b) skewness

fan charts are constructed using the Two Piece Normal (TPN) distribution

$$f(x) = \begin{cases} A \exp\left[-\frac{(x-\mu)^2}{2\sigma_1^2}\right] & \text{if } x < \mu \\ A \exp\left[-\frac{(x-\mu)^2}{2\sigma_2^2}\right] & \text{otherwise} \end{cases} \tag{10}$$

with $A = \sqrt{2/\pi}(\sigma_1 + \sigma_2)^{-1}$. The parameters of the distribution (μ, σ_1 and σ_2) are chosen so that it replicates the evolution of mode, uncertainty and skewness of inflation projected by the Bank using the following relations that can be derived from Britton et al. (1998) and Wallis (2004)

$$b = \frac{\pi\eta^2}{2\sigma^2}, \quad g = \text{sign}(\eta) \times \left[1 - \left(\frac{(\sqrt{1+2b}-1)^2}{b}\right)\right], \quad \sigma_1 = \frac{\sigma}{\sqrt{1+g}}, \quad \sigma_2 = \sqrt{\frac{(1+g)}{(1-g)}}\sigma_1$$

with μ equal to the mode projected by the Monetary Policy Committee (MPC).

Every quarter the Bank releases a new set of density forecasts which imply a revision of the previous ones incorporating the new information made available during the previous quarter. This new information may or may not alter the balance of risk assessment made by the Bank which could, then, affect the shape of the updated density forecasts of inflation. The analysis presented here starts in January 1998, the year in which the Bank of England was granted independence.

The method described in section 2.1 will be applied here to obtain the densities of these shocks of information. As noted earlier, a copula function is needed to model the joint density of two density forecast; the Frank copula (Frank, 1979; Genest, 1987) is used in (4). The Frank copula, with density function defined as

$$c(u_1, u_2) = \theta\eta e^{-\theta(u_1+u_2)}/[\eta - (1 - e^{-\theta u_1})(1 - e^{-\theta u_2})]^2, \qquad \eta = \theta - 1,$$

and parameter $\theta > 0$, allows for negative dependence between the marginals, and symmetric dependency and does not allow for tail dependence. It is, then, useful to model strong positive or negative dependence where this dependence is centered in the centre of the distribution. This is an appropriate choice for this problem. Dependence between revisions are likely to be strongly positive, and focused on central moments rather than in extreme events.

3.2 Uncertainty measure

Figure 3.1 contains the plots of the variance (panel a) and skewness coefficient (panel b) of the information shocks extracted from the Bank of England density forecasts for the one quarter ahead (blue line), two quarters ahead (red line), one year ahead (green line) and two years ahead (yellow line). The picture shows a decrease in the variance of the information shocks from the end of the 90s to the end of 2007. Also for a given forecasted quarter in this period the variances of the shocks uniformly decrease with the forecast horizon suggesting that the second regularity of optimal forecasts commented on earlier (uniform decrease of the variance of revisions as we approach the outcome date) is likely to hold. This suggests also that it is likely to assume that the Bank was efficiently incorporating information to its forecasts during that period. As mentioned earlier, the shock expected for inflation in two years' time should have a high variance reflecting the uncertainty surrounding the outcome variable. The incorporation of information in the subsequent revision of this forecast in the following quarters should decrease uncertainty as measured by the variance of the errors extracted from the fan charts. During the period 2008-2011 this mechanism of information incorporation seems to break. The variances of the errors increase sharply at all forecast horizons reflecting the increase of economic uncertainty, and there is no clear evidence of the decrease of variances with the forecast horizon. This could be due to the inability of the Bank

Table 1: Correlation of $Var[\epsilon_{t+h}(h)]$ with uncertainty measures

	Correlation			p-value		
h	EPU	JLN-F	JLN-M	EPU	JLN-F	JLN-M
1	0.6634	-0.2778	-0.3664	0.00	0.01	0.00
2	0.6764			0.00		
4	0.6162	0.0079	-0.3526	0.00	0.47	0.00
8	0.5215			0.00		

to incorporate or even process new information in the density forecasts. After 2011 the variances of the shocks, although still quite high with respect to its pre-crisis values, seem to have stabilized. Also, a pattern of efficient incorporation of information seems to have been re-established.

We can interpret the variance of the shocks as a measure of perceived uncertainty at different horizons. For example, the blue line in Figure 1a represents the expected uncertainty surrounding inflation in the next quarter, while the yellow line can be interpreted as the expected uncertainty of inflation in two years' time. In order to compare the relationship of this measure with other popular measures of uncertainty, the rank correlation coefficients of the proposed measure with the UK Economic Policy Uncertainty (EPU) index of Baker et al. (2013) and the macroeconomic and financial uncertainty indices of Jurado et al. (2015) (JLN-M and JLN-F respectively) are presented in table 1. While there is a strong positive correlation between EPU and the variance of the information shocks for the selected horizons, the correlation seems to be negative for the JNL measures. While this last result is striking at first glance, it should be noted that this latter measure is specific of the US. The proposed measure of uncertainty has an advantage as compared to the previous ones: it can be calculated every time a revision of the forecast is performed, so it is in this sense an *ex-ante* measure of uncertainty that could be incorporated into forward looking models. Also, as it is a measure of inflation uncertainty expected by the Bank of England, it could be used to check whether uncertainty plays a role in monetary policy actions.

With respect to the skewness coefficient, the distribution of information shocks for one and two years ahead seem to be quite symmetrical, meaning that the Bank does not have a strong belief in the likelihood of extreme event outcomes. However, the MPC seems to pay attention to this possibility when producing short term forecasts. The skewness coefficient of the information

Figure 2: News absorbtion

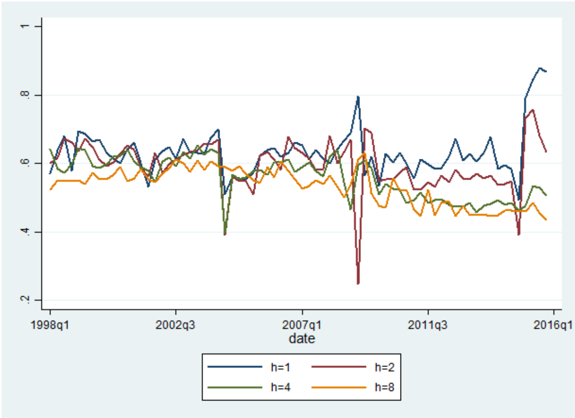

shocks fluctuates widely until the beginning of the financial crisis for short term forecast revisions. However, since then the distribution of the shocks perceived by the Bank seem to have a more symmetrical distribution for all forecast horizons. This result, together with the fact that the variance of the information shocks increased during the financial crisis, can be interpreted to mean that the Bank is perceiving a more uncertain environment in which large positive shocks are as likely as large negative ones.

Finally, figure 2 shows the measure of news absorbtion, which measures the degree of incorporation of new information in subsequent forecasts. Notice that for short term forecasts this measure is more volatile than for long term forecasts, although in all cases it fluctuates around 0.6. This means that approximately 60% of the information contained in a forecast comes from information made available after the previous release. Regarding long term forecasting (one and two years ahead) the proportion is similar, although it seems to be declining since the financial crisis. This result seems to indicate that in a period of greater uncertainty new information is taken with caution before incorporating it into long term forecasts. An open question is to decide how much of this new information is 'news' or 'noise' as noted in Lahiri and Liu (2006).

3.3 Optimal forecasts regularity tests

Figure 3 shows the mean of the shocks of information for selected forecast horizons (right axis). The grey bars indicate a rejection of the null hypothesis of zero mean at the 10% significance level

Figure 3: Efficiency tests: Test for zero mean of the shocks

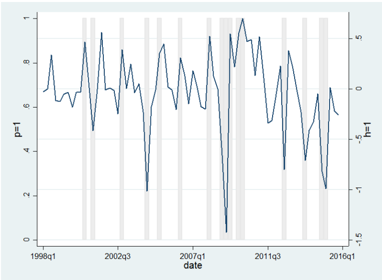

(a) h=1

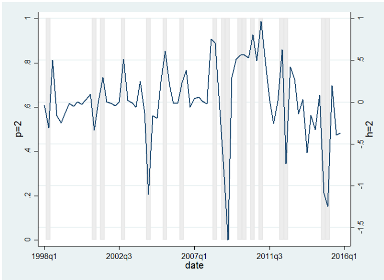

(b) h=2

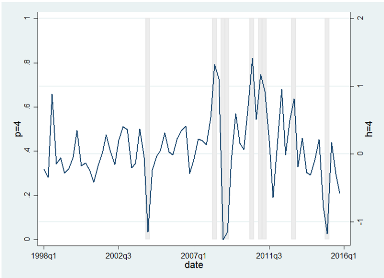

(c) h=4

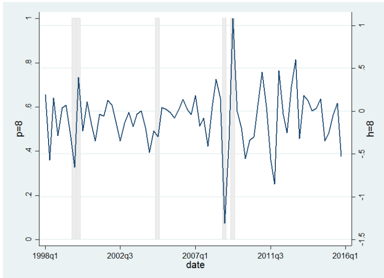

(d) h=8

and, a therefore, it identifies periods of biased forecasts or significant forecast revisions. Notice that the number of periods for which the null hypothesis is rejected is decreasing with the horizon and that for h=2, 75% of the forecasts can be considered optimal. However, that potential loss of optimality is focused consistently across horizons in the periods immediately after 2007 and during the financial crisis. This coincides with the period in which uncertainty, as measured by the variance of the shocks of information, starts to increase. This identifies a period in which economic and financial conditions worsened very quickly, and the Bank seem to have difficulty in processing the information properly. For the rest of forecast horizons, the proportion of non-optimal forecasts is below the 75%. Another period of lack of optimality that the test is consistently identifying is the period around January 2004, when the Bank changed its inflation target from RPIX to CPI, which could be considered spurious. The picture painted does not contradict previous results that prove the efficiency of the Bank of England forecasts (see e.g. Bank of England, 2015; Elder et al., 2005), but instead complements it, identifying certain periods in which efficiency is lost. It can also be noted that forecasts are more significantly revised for short term forecasts than for long term forecasting, which is consistent with the higher uncertainty of inflation for horizons larger

Figure 4: Efficiency tests: Test for decreasing variance of the shocks

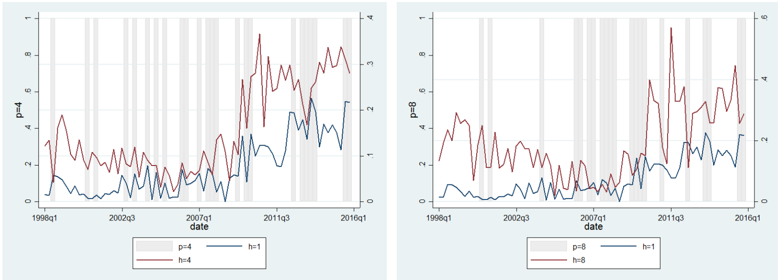

(a) h=4 (b) h=8

than two.

Finally, figure 4 plots the variance of the shocks for one and two years ahead forecasts. Grey bars indicate rejection at 10% significance level of the null hypothesis that these variances are significantly larger than the variance of the shocks one month ahead. Although the number of periods for which the null is rejected is higher than in the previous test, the results are similar: there seems to be a loss of optimality in the aftermath of the financial crisis. It should be noted that the rejection of the null during the years of the so-called Great Moderation could be due to the effect of the Friedman-Ball hypothesis -low inflation expected in periods of low inflation uncertainty. What is observed here is the fact that inflation was low and the economic conditions were stable, which resulted in a period of low uncertainty. However, the results indicate that this hypothesis may not hold after the financial crisis, where inflation uncertainty is still large but the level of inflation remains at low levels. Further investigation is needed in this aspect.

4 Conclusions

This paper proposes a new method to extract the density of information shocks from the revision of density forecasts. Studying this density gives information about the change in the balance of risks and it could help understand the mechanism of information incorporation to subsequent forecasts. The paper focuses on extracting this information from the revision of inflation density forecasts produced by the Bank of England. The variance of these shocks at different forecast horizons is proposed as a new measure of *ex-ante* inflation uncertainty. The results show that until the beginning of the financial crisis the Bank was perceiving a period of decreasing uncertainty. Since then, inflation uncertainty perceived by the Bank rose sharply and, although it has stabilized, it is still at levels above its pre-crisis values. New information released between two vintages, which represent about 60% of the information contained in revisions, seems to have been incorporated optimally. Periods of unbiasedness of forecasts are identified which may correspond with periods in which forecast efficiency does not hold. These moment focus on the 2008-2011 period when there seems to be no clear learning mechanism that incorporates information into revisions of forecasts. This results and the methodology proposed open interesting areas of research: it should help improve density forecast evaluation tests, as it can be used to eliminate the dependence present in multi-step forecasts. Also, the uncertainty shocks obtained can be used to study if they have affected monetary policy outcomes and in what ways.

References

Alessi, L., Ghysels, E., Onorante, L., Peach, R. and Potter, S. (2014), "Central Bank Macroeconomic Forecasting During the Global Financial Crisis: The European Central Bank and Federal Reserve Bank of New York Experiences", *Journal of Business and Economic Statistics*, 32(4), pp. 483-500.

Bank of England (2015), 'Evaluating Forecast Performance', Independent Evaluation Office, November 2015.

Baker, S. R., N. Bloom and S.J. Davis (2013), 'Measuring economic policy uncertainty', Stanford University and University of Chicago Booth School of Business.

Basu, A., H. Shioya and C. Park (2011), Statistical Inference: The Minimum Distance Approach, CRC Press.

Bracewell, R. (1999), *The Fourier Transform and its Applications*, 3rd ed. New York: McGraw-Hill.

Britton, E., Fisher, P. and Whitley, J. (1998), 'The Inflation Report Projections: Understanding the Fan Chart', Bank of England Quarterly Bulletin, 38, pp. 30-37.

Clements, M.P. (2004), 'Evaluating the Bank of England Forecasts of Inflation', *The Economic Journal*, 114, pp. 844-866.

Chaudhuri, K., Kim, M. and Shin, Y. (2016), 'Forecasting Distributions of Inflation Rates: The Functional Auto-Regressive Approach', *Journal of the Royal Statistical Society*, 179(1), pp. 65-102.

Dowd, K. (2007), 'Too Good to be True? The (In)credibility of the UK Inflation Fan Charts', *Journal of Macroeconomics*, 29, pp. 91-102.

Elder, R., Kapetanios, G., Taylor, T. and Yates, T. (2005), 'Assessing the MPC's Fan Charts', Bank of England Quarterly Bulletin, 45, pp. 326–348.

Fawcett, N., Korber, L., Masolo, R.M. and Waldron, M. (2015), 'Evaluating UK Point and Density Forecast from an Estimated DSGE Model: The Role of Off-Model Information over the Financial Crisis', Bank of England Staff Working Paper No. 538.

Frank, M.J. (1979), 'On the Simultaneous Associativity of $F(x, y)$' and $x + y \sim F(x, y)$, *Aequationes Mat.*, 19, pp. 194-226.

Galbraith, J.W. and van Norden, S. (2012), 'Assessing Gross Domestic Product and Inflation Probability Forecasts Derived from Bank of England Fan Charts', *Journal of the Royal Statistical Association, Series A*, 175(3), pp. 713-727.

Genest, C. (1987), 'Frank's Family of Bivariate Distributions', *Biometrika*, 74(3), pp. 549-555.

Gneiting, T. and Ranjan, R.(2011), 'Comparing Density Forecasts Using Threshold- and Quantile-Weighted Scoring Rules', *Journal of Business and Economic Statistics*, 29(3), pp. 411-422.

Granger, C.W.J. and Newbold, P. (1986), *Forecasting Economic Time Series* (2nd Edition). Academic Press: New York.

Isiklar,G. and Lahiri, K. (2007), 'How Far Ahead Can We Foecast? Evidence From Cross-Country Surveys' *International Journal of Forecasting*, 23, pp.167-187.

Joe, H. (1997), *Multivariate Models and Dependence Concepts*. Chapman and Hall. London.

Joe, H. and Xu, J.J. (1996), 'The Estimation Method of Inference Functions for Margins for Multivariate Models', Department of Statistics, University of British Columbia, Technical Report, 166.

Jurado, K., Ludvigson, S.C. and Ng, S. (2015), 'Measuring Uncertainty', *American Economic Review*, 105(3), pp. 1177-1216.

Lahiri, K. and Liu, F. (2006), 'Modelling Multi-Period Inflation Uncertainty Using a Panel of Density Forecasts'. *Journal of Applied Econometrics*, 21, pp. 1199-1219.

Mankiw, N.G. and Saphiro, M.D. (1986), 'News or Noise? An Analysis of GNP Revisions'. NBER Working Paper No. 1939.

Mincer, J. and Zarnowitz, W. (1969), 'The Evaluation of Economic Forecasts', in *Economic Forecasts and Expectations*, ed. J.Mincer, New York: National Bureau of Economic Research.

Mitchell, J. (2008), 'Density Forecast Revisions and Forecast Efficiency', National Institute of Economic and Social Research.

Mitchell, J. and Hall, S.G. (2005), 'Evaluating, Comparing and Combining Density Forecasts Using the KLIC with an Application to the Bank of England and NIESR 'Fan' Charts of Inflation'. *Oxford Bulletin of Economics and Statistics*, 67, pp. 995-1033.

Nordhaus, W.D. (1987), 'Forecast Efficiency: Concepts and Applications', *Review of Economics and Statistics*, 69(4), pp. 667-674.

Patton, A.J. and Timmermann, A. (2012), 'Forecast Rationality Tests Based on Multi-Horizon Bounds'. *Journal of Business and Economic Statistics*, 30(1), pp. 1-17.

Rossi, B. and Sekhposyan, T. (2015), 'Forecast Rationality Tests in the Presence of Instabilities, with Applications to Federal Reserve and Survey Forecsts'. *Journal of Applied Econometrics*, DOI: 10.1002/jae.2440.

Sklar, A. (1959), 'Fonctions de repatition á n dimensions et leurs marges'. Publ Inst Univ Paris 8, pp-229-231.

Wallis, K.F. (2004), 'An Assessment of Bank of England and National Institute Inflation Forecast Uncertainties', *National Institute Economic Review*, 189, pp. 64-71.