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 information shocks from the revisions of the Bank of England's inflation density forecasts. An information shock is defined in this paper as a random variable that contains the set of information made available between two consecutive forecasting exercises and that has been incorporated into a revised forecast for a fixed point event. Studying the moments of these information shocks can be useful to understand how the Bank has changed its assessment of risks surrounding inflation in the light of new information, and how it has modified its forecasts accordingly. The variance of the information shock is interpreted in this paper as a new measure of ex-ante inflation uncertainty, that measures the uncertainty that the Bank anticipates that information perceived in a particular quarter will pose on inflation. A measure of information absorption that indicates the approximate proportion of the information content in a revised forecast that

*University of Leicester School of Business. Astley Clarke Building. University Road. Leicester LE1 7RH. United Kingdom. Email: cdv7@leicester.ac.uk. This research used the ALICE High Performance Computing Facility at the University of Leicester. Thanks to Wojciech Charemza, Stephen Hall, José Luis Gallego and Riccardo Masolo 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 and two anonymous referees for their feedback.

is attributable to information made available since the last forecast release is also proposed.

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

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

1 Introduction

Density forecasting has become a useful tool for researchers and policymakers as a way to evaluate the uncertainty surrounding the central path of key economic and financial variables. In particular, inflation density forecasts are commonly used by some central banks as informative tools in the process of monetary policy implementation. The Bank of England periodically produces inflation and GDP growth density forecasts for the following quarters, presented to the public in the form of so-called *fan charts*. These forecasts are updated every quarter with new information made available since the previous release. Any difference between the density forecasts for a fixed point event produced in two consecutive forecasting exercises will be due to the effect of the incorporation of this new information into the revised forecast. The random variable that produces this change is what will be called in this paper an *information shock*. Knowing the density and the moments of this information shock can be helpful, among other things, to evaluate the process of information incorporation into subsequent forecasts. For example, if this density is non symmetric, this can be interpreted as sign that the Bank expects that new information will have a bigger impact on a particular side of the distribution than the one anticipated in the previous forecasting exercise. The skewness coefficient of the information shock can be seen, then, as a measure of how the Bank has changed its assessment of risks affecting inflation in the light of new information. Similarly, its variance can be interpreted as a measure of the uncertainty that the Bank anticipates that new information will have on future inflation. The method presented here to extract the density of this information shock just requires knowledge of the density forecasts produced in two consecutive releases for a fixed point event. Repeating this exercise for every forecasting round available, and looking at the change in the moments of the information shocks the researcher can study how the Bank has reacted to new information, changed its assessment of

risks about the future evolution of inflation accordingly, and incorporated this new information and assessment into the forecasts. This paper describes the tools to do so.

The paper contributes to the literature in two main areas. First, it complements other measures of macroeconomic uncertainty, such as the ones proposed by Jurado et al. (2015), Baker et al. (2016) or Clements (2014). The uncertainty measure proposed here differs from the previous ones in that this is an *ex-ante* measure based on density forecast revisions. It can be viewed not an overall measure of inflation uncertainty, but rather as a measure of the 'amount' of uncertainty that the Bank expects a particular set of news will cause on uncertainty for a fixed point event. This could help researchers trace back events that may have had a significant impact on an increase in overall uncertainty. A second area in which this paper contributes is in the evaluation of the forecasting practice of the Bank of England. As Alessi et al. (2014) stress, it is important to assess the forecast accuracy of Central Banks' forecasting ability in the light of the recent financial crisis. There is already a vast literature on the evaluation of the Bank of England's point and density forecasts. Wallis (2004), Clements (2004) or Dowd (2007), for example, evaluate the forecasting ability of the fan charts showing that, while they are 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). Gneiting and Ranjan (2011) or Galbraith and van Norden (2012) re-evaluate these forecasts pointing out again the overestimation of risks and the lack of resolution. This paper takes a different approach in the evaluation of these forecasts. It rather looks at the way in which the Bank has incorporated information into subsequent forecasts over the last two decades by analysing the change in the moments of the information shocks over time. Given the interpretation of the skewness coefficient above, it is found that during the Great Moderation the Bank seems to have been able to make an assessment on whether new information posed an upward or a downward risk to future inflation, as the skewness coefficients of the information shocks are non-zero -specially for short term forecasting. However, since the start of the global financial crisis, as the effect that the Bank expects that new information will have on inflation uncertainty is increasingly large, the densities of the information shocks become symmetric. This

can be interpreted as sign that, in a period of increased uncertainty, it has been more difficult for the Bank to assess properly the sign of the impact of new information on inflation. The analysis also shows that new information makes up to about 60% of the information content of a revised forecast. This figure is common to all forecast horizons, but the percentage seems to be slowly declining for long term forecasting since the financial crisis. This can be indicative of the Bank being increasingly cautious before modifying long term forecasts.

The paper is structured as follows: section 2 describes the method used to extract the density of the information shocks via convolutions, and defines a measure of information absorption. Section 3 applies this technique to recover the densities of the information shocks from the Bank of England density forecasts, and proposes and interprets a measure of inflation uncertainty. An evaluation of information absorption is also performed in this section. Section 4 concludes.

2 An evaluation of density forecasts

2.1 Recovering the densities of information shocks via convolutions

The process of elaboration of the Bank of England’s inflation density forecasts, commonly known as ‘fan charts’, is not a purely econometric one. As described in Bank of England (2015), a typical forecasting round¹ is structured as a series of meetings between the technical staff of the Bank and the Monetary Policy Committee (MPC). At the beginning of the process the technical staff collects all information available about the state of the economy that has been released since the last forecasting round, and produces a series of preliminary forecasts of the evolution of inflation and other macroeconomic variables for the following twelve quarters. In the light of this new information the MPC reassesses its balance of risks about future inflation, and modifies accordingly the forecasts presented to it by the technical staff. Therefore, the final density forecasts released to the public with the Quarterly Inflation Report represent the ‘best collective judgment’ of the MPC (see Tay and Wallis, 2000), contain a considerable degree of judgment. Since these new forecasts are revisions of the ones published the previous quarter, any difference

¹The expressions ‘forecasting round’, ‘forecasting exercise’ and ‘forecasting release’ will be used interchangeably.

between them for a fixed point event will be due to the effect that the MPC anticipates that new information made available during the last quarter will have on inflation at that given point in the future. This change is produced by what will be called an *information shock*. It is important to note that the same information set could be incorporated differently at different forecast horizons (e.g., an oil shock that is perceived as transitory is likely to make the MPC change the short term forecasts but not the long term ones), and that the MPC can understand that not all new information will impact all forecast horizons. Hence, in every forecasting round there will be a different information shock for every forecast horizon. Since these information shocks have a high component of judgment, knowing their densities will help us not only understand how the forecasts have been revised, but also to evaluate how new information has changed (if at all) the Bank's assessment of risks for future inflation at each forecast horizon. In particular, if the variance of the information shock at a given forecast horizon is large, it would indicate that the MPC expects that information obtained during the last quarter will increase the risks associated to the outcome of inflation at that particular point in the future. This section presents a method to extract the densities of these information shocks, so that after every forecasting round a full set of densities (one for each forecast horizon, except the last one²) will be obtained. By repeating this exercise for every forecasting exercise for which we have information, we can see the evolution over time of the densities of the information shocks and their moments at each forecast horizon. This will help study, for example, the change in the term structure of uncertainty.

To illustrate the method, let z_{t+h} be the variable to be forecast (e.g. inflation in period $t+h$). At moment t a density forecast for this variable is produced using all information available to the forecaster up to that point. Let I_t denote this information set. Let also $z_{t+h|t}$ be a random variable ('the forecast') with density function equal to the density forecast, say $f_{t+h|t}^z(\cdot)$. At time $t+1$ a new density forecast for z_{t+h} is produced, $f_{t+h|t+1}^z(\cdot)$. This new density forecast is constructed using all information available up to time $t+1$ (i.e., using information set I_{t+1}), so any difference between $f_{t+h|t}^z(\cdot)$ and $f_{t+h|t+1}^z(\cdot)$ will reflect the effect of the incorporation of information made

²The last forecast produced in a given forecasting exercises will be first time that variable is forecasted, so there will be no previous density forecast to compare it with.

available between periods t and $t + 1$ to the updated forecast, $z_{t+h|t+1}$. The well known updating equation of forecasts decomposes the revised forecast as

$$z_{t+h|t+1} = z_{t+h|t} + \epsilon_{t+h|t+1}. \quad (1)$$

where $\epsilon_{t+h|t+1}$ is the *information shock* which density is to be recovered. It is important to note that (1) establishes a relationship between random variables and not point forecasts. This equation also defines the mechanism of incorporation of information made available between t and $t + 1$ into the revised forecast $z_{t+h|t+1}$ that will be used to define the measure of information absorption later on. It is worth pointing out that this information shock contains all the information perceived by the forecaster and incorporated into the revision, but does not contain information not observed by the forecaster but relevant for the outcome, or information observed but not considered relevant to be incorporated. In this sense, $\epsilon_{t+h|t+1}$ is informative of how the forecaster has adapted the revised forecast, and it contains an important degree of judgment. Its density function will help us analyzed this subjective assessment of risks that the forecaster identifies that new information poses on the objective variable. The aim now is to obtain the density of the information shock $\epsilon_{t+h|t+1}$. The method proposed here for this is based on the notion of convolution. Notice from (1) that we can define $\epsilon_{t+h|t+1} = z_{t+h|t+1} - z_{t+h|t}$, where $z_{t+h|t+1}$ and $z_{t+h|t}$ are two random variables with known densities. If these two random variables were independent, the density function of $\epsilon_{t+h|t+1}$ could be obtained by convolution of their densities³. However, $z_{t+h|t+1}$ will normally be dependent on $z_{t+h|t}$, as the information sets on which these forecasts are constructed overlap (in general $I_t \subset I_{t+1}$). In that case, the density of the information shock can be obtained as the density of the difference between two dependent random variables by convolution from the joint density of $z_{t+h|t+1}$ and $z_{t+h|t}$:

$$f_{t+h|t+1}^\epsilon(z) = \int_{-\infty}^{\infty} f_{z_{t+h|t+1}z_{t+h|t}}(x, x - z)dx. \quad (2)$$

³If X and Y are two independent random variables with densities $f_X(\cdot)$ and $f_Y(\cdot)$ respectively, then the density of $Z = X - Y$ can be obtained as $f_Z(z) = \int_{-\infty}^{\infty} f_X(z - x)f_Y(x)dx$.

This joint density $f_{z_{t+h|t+1}z_{t+h|t}}(\cdot, \cdot)$ is usually not available, as all that is known to the forecaster are the density forecasts themselves (i.e., the densities of $z_{t+h|t+1}$ and $z_{t+h|t}$), which can be thought of as the marginal densities that arise from this joint density. Knowledge of the marginal densities can be used to reconstruct the joint density via copulas, and then the density of the information shock can be obtained using (2). A copula is a multivariate distribution with uniform marginals that can be used to model the dependency structure of two or more random variables with known marginal densities. More formally, let $F_{t+h|t+1}^z(\cdot)$ and $F_{t+h|t}^z(\cdot)$ be the marginal cumulative distribution functions (cdfs) of $z_{t+h|t+1}$ and $z_{t+h|t}$ respectively. Following Sklar (1959), if these distribution functions are continuous, there exists a unique function $C(\cdot, \cdot|\theta) : [0, 1]^2 \rightarrow [0, 1]$ such that, in this bivariate case,

$$\tilde{F}_{z_{t+h|t+1}z_{t+h|t}}(x_1, x_2) = C[F_{t+h|t+1}^z(x_1), F_{t+h|t}^z(x_2)|\theta], \quad x_1, x_2 \in \mathfrak{R}, \quad (3)$$

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

$$\tilde{f}_{z_{t+h|t+1}z_{t+h|t}}(x_1, x_2) = c[F_{t+h|t+1}^z(x_1), F_{t+h|t}^z(x_2)|\theta] f_{t+h|t+1}^z(x_1) f_{t+h|t}^z(x_2), \quad (4)$$

where $c(\cdot, \cdot|\theta)$ is the copula density. A plethora of parametric copulas have been proposed in the literature to model different types of dependence structures between random variables (see e.g. Joe, 1997, for a review). Parametric copulas are convenient since they have closed form expressions for the distribution function $C(\cdot, \cdot|\theta)$ and the corresponding copula density $c(\cdot, \cdot|\theta)$. Also, closed form expressions of the Kendall and Spearman rank correlation coefficients can be obtained for each copula in the bivariate case. However, parametric copulas depend on the vector of parameters θ that needs to be estimated. The method proposed here allows to estimate a different copula parameter for every release, t . Traditional methods of copula estimation provide a single estimate of the copula parameters, θ , for the whole sample, which would imply that the

dependence between revisions is constant over time. Furthermore, since the objective of this paper is to obtain the density of the information shocks before the realisation of the forecasted variable is observed, a constant copula parameter would imply assuming that this degree of dependence between revisions will carry on to the future, which may be quite a restrictive assumption in a changing environment. It would be expected that in periods of relative tranquility the forecaster will make just small adjustments to the forecasts, being the dependence between them higher than in periods of rapid change, in which new events will make the forecaster modify the forecasts more substantially. Therefore, the copula parameter should be allowed to change over time.

To describe the proposed estimation method of the copula parameters first of all notice that, given that $\epsilon_{t+h|t+1}$ and $z_{t+h|t}$ are constructed on non overlapping information sets, both random variables are independent of each other by construction, and an approximation to $f_{t+h|t+1}^z(\cdot)$ can be obtained by the convolution of their respective densities from (1) as

$$g_{t+h|t+1}^z(z|\theta) = \int_{-\infty}^{\infty} f_{t+h|t}^z(z-x) \tilde{f}_{t+h|t+1}^{\epsilon}(x|\theta) dx, \quad (5)$$

where $\tilde{f}_{t+h|t+1}^{\epsilon}(\cdot|\theta)$ is an estimate of the density of the information shock obtained substituting (4) into (2). The density function $g_{t+h|t+1}^z(z|\theta)$ is, then, an approximation to the known density of $z_{t+h|t+1}$ that depends on the copula parameter vector θ that needs to be estimated. The proposed estimation method for θ consists of minimizing the Kullback-Leibler (KL) measure (see e.g. Mitchell and Hall, 2005; Bao et al., 2007). The KL is a measure of the expected divergence between the 'true' density $f_{t+h|t+1}^z(\cdot)$ and the approximated one $g_{t+h|t+1}^z(\cdot)$. Thus, the proposed estimator for θ can be defined as

$$\hat{\theta}_{t+1} = \arg \min_{\theta} \int_{-\infty}^{\infty} f_{t+h|t+1}^z(x) \ln \left[\frac{f_{t+h|t+1}^z(x)}{g_{t+h|t+1}^z(x|\theta)} \right] dx. \quad (6)$$

This estimator falls into the class of minimum distance estimators (MDE) described in Basu et al. (2011). Study of the properties of some MDEs of copula parameters can be found in Biau and Wegkamp (2005). The Kullback-Leibler divergence is used here because of its intuitive

interpretation. This divergence measure is commonly known in the Information Theory literature as 'relative entropy', and it is interpreted as a measure of the information lost when a function g is used to approximate another known one f . If both functions are the same, then KL will be zero, and it will be positive otherwise. Its value can be directly used to compare different approximations to the same function f , and is the base of the well known Akaike's Information Criterion (see Burnham and Anderson, 2002). Therefore, the idea behind (6) is that the copula parameters θ should be chosen so that the information lost when $g_{t+h|t+1}^z(x|\theta)$ is used to approximate $f_{t+h|t+1}^z(x)$ is minimal. This principle of minimum cross-entropy has been used previously by Miller and Liu (2002), Chu (2011) or Gronneberg and Hjort (2014) to estimate copula parameters. The crucial difference between these papers and the proposed method is that here the 'true' density $f_{t+h|t+1}^z(x)$ is known, so it doesn't need to be recovered using historical information. Notice also that the only information needed in (6) is the density forecast of a fixed point event in two consecutive forecasting exercises. Hence, it is possible to estimate a different copula parameter vector for every forecasting round and forecast horizon, allowing the dependence between revisions to change over time and across horizons, as initially desired. Akaike (1973) shows that, when the true density in (6) is known⁴, the minimum cross entropy estimator is equivalent to the maximum likelihood estimator, so the proposed method provides consistent estimators of the copula parameters (see Soofi, 2000, and the appendix for a summary of this result). This method can be seen as a way to extract from the revised forecast $z_{t+h|t+1}$ all the information content up time t , so that this ensures that the information set on which the density of the information shock is formed by $I_{t+1} - I_t$.

So far the method to extract the density of the information shocks has been explained considering just one forecast horizon in two consecutive releases (h for forecasts released at time t , or $h - 1$ for forecasts released at time $t + 1$, both referring to the same fixed point event z_{t+h}). In practice the forecaster will produce forecasts for several horizons in a given forecasting exercise. Applying this procedure to different forecast horizons of the same forecasting exercise will provide the densities of a set of information shocks, one for each forecast horizon. Repeating this exercise

⁴In our application, this 'known' or 'true' density is the revision of the density forecast produced by the Bank in moment $t + h$.

for all the forecasting releases available will produce a picture of the evolution over time of the densities of the information shocks at different forecast horizons. This will be helpful to study the change in the term structure of uncertainty.

2.2 A measure of information absorption

Equation (1) establishes a relationship between two successive forecasts for the same fixed point event. The revised forecast is made up of a part of 'inertia', coming from the forecast made in the previous period for the same variable, and a term $\epsilon_{t+h|t+1}$ that takes into account all the information that has been being perceived by the forecaster between periods t and $t+1$ and that has been incorporated into the revised forecast $z_{t+h|t+1}$. In order to evaluate how new information is used to update forecasts, it would be useful to know what is the proportion of new information embedded into a revised forecast. This section describes how the decomposition (5), evaluated under the optimal value of θ , can be used to construct a measure of information incorporation to updates of density forecasts.

Assume that $g_{t+h|t+1}^z(z) = f_{t+h|t+1}^z(z)$, i.e., that the fit is perfect. The convolution theorem (see e.g. 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 it⁵, so that

$$FT[f_{t+h|t+1}^z(z)] \equiv s_{t+h|t+1}^z(\lambda) = s_{t+h|t}^z(\lambda) s_{t+h|t+1}^\epsilon(\lambda), \quad \lambda \in [-\pi, \pi]. \quad (7)$$

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 relationship 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 :

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

⁵This can be easily proven as $FT[f_{t+h|t+1}^z(z)] = \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f_{t+h|t}^z(z-x) f_{t+h|t+1}^\epsilon(x) dx \right] e^{-i\lambda z} dz = \int_{-\infty}^{\infty} f_{t+h|t+1}^\epsilon(x) e^{-i\lambda x} \left[\int_{-\infty}^{\infty} f_{t+h|t}^z(z-x) e^{-i\lambda(z-x)} d(z-x) \right] dx = s_{t+h|t}^z(\lambda) s_{t+h|t+1}^\epsilon(\lambda)$.

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

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

This measure, which lies between zero and one, tells us which is the proportion of the information content in $z_{t+h|t+1}$ that is attributable to information made available and incorporated into this forecast between t and $t+1$. As with the estimation of the copula parameters in the previous section, this measure can be constructed for every forecast horizon and forecasting release available. This allows to study how the incorporation of information may change through time and across forecast horizons.

3 Evaluating the Bank of England fan charts

3.1 Extraction of the densities of information shocks

Since 1993, and following the introduction of inflation targeting for UK monetary policy in 1992, the Bank of England has been producing its inflation forecasts⁶ in the form of density forecasts. Each quarter the Bank releases its views about the evolution of inflation for the next 12 quarters (8 quarters until the second quarter of 2013) under two assumptions: constant official interest rates, and interest rates following market rates. These forecasts are presented in what has become popularly known as *fan charts*, which can be interpreted as a graphical representation of the Bank's balance of risks surrounding the expected central tendency of inflation for the next two years. With the release of projections every quarter, the Bank reports as well the three moments that characterise their distributions, the mode (μ), the uncertainty (σ), and the skewness⁷ (η). The value of this set of parameters and a report containing all other assumptions underlying the forecasting process are publicly available on the Bank of England's website⁸. These moments

⁶for RPIX inflation until December 2003 and for CPI inflation since.

⁷Skewness is measured as the difference between the mean and the mode.

⁸<http://www.bankofengland.co.uk/publications/Pages/inflationreport/irprobab.aspx>

are decided by the MPC as the final step of the procedure outlined in section 2. These three moments, that can be set to different values for each forecast horizon, represent the best collective judgement of the MPC about the future evolution of inflation in a given forecasting round. Once the parameters have been decided, the fan charts are constructed assuming that the density of inflation at every future quarter follows a Two Piece Normal (TPN) distribution, with moments matching the ones already agreed upon. The density of the TPN is defined as

$$f(x) = \begin{cases} A \exp \left[-\frac{(x-\mu)^2}{2\sigma_1^2} \right] & \text{if } x < \mu \\ A \exp \left[-\frac{(x-\mu)^2}{2\sigma_2^2} \right] & \text{otherwise} \end{cases} \quad (9)$$

with $A = \sqrt{2/\pi}(\sigma_1 + \sigma_2)^{-1}$. The parameters of this distribution (μ , σ_1 and σ_2) can be chosen so that they replicate the mode, uncertainty and skewness coefficients provided by the MPC for each horizon, using the following relationships (see Britton et al., 1998; Wallis, 2004):

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

This section describes and analyses the results of using the method proposed in section 2.1 to extract the densities of the information shocks from these density forecasts. The analysis starts in January 1998, which is the moment in which the Bank of England started to work with operational independence.

In order to implement the method a copula function must be chosen to model the joint distributions of the forecasts defined in (4). Different copulas are able to model different types of dependence structures (see e.g. Joe, 1997). In this paper the copula used is the Frank copula (Frank, 1979; Genest, 1987), which density is

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

This copula has a single parameter, $\theta \geq 0$, that controls the degree of dependence with $\theta = 0$ being the independent case. The Frank copula, unlike other copulas, allows to model negative

Figure 1: Kullback-Leibler divergence measures

dependence between the marginals, but it does not allow for tail dependence (i.e. dependence between extreme events). Dependence modelled by this copula is stronger in the middle of the distribution and, as other popular copulas such as the Gaussian, dependence is symmetric in both tails. Unlike other magnitudes modelled via copulas, such as financial variables, the forecaster of macroeconomic magnitudes normally pays less attention to extreme events, and is more worried about forecasting central tendencies. As well, the revisions of forecasts, which are the variables that are modelled here, are likely to have a strong either positive or negative dependence. In periods of tranquility it is expected that a forecast will rely heavily in the forecast made in the previous moment for the same variable. In periods of uncertainty, that dependence might turn negative if new information suggests that the previous forecast could be over or underestimating future inflation. Hence, the properties of the Frank copula makes it a good and intuitive candidate for this analysis.

Once the copula has been chosen, the process of extracting the information shocks can be performed as explained in the previous section. For every forecasting round the densities of the information shocks affecting the following 8 quarters⁹ are recovered. A different copula parameter

⁹Although there are up to 12-quarters-ahead forecasts available in each forecasting round, the analysis focuses on the first 8 quarters to make the exercise extensible to forecasting rounds not presented here in which only 8 forecast horizons were available.

is estimated for every release and forecast horizon using (6), and the value of the copula parameter and its standard error, the Kullback-Leibler divergence measure and the first three moments of each information shock are calculated. This exercise is repeated for every forecasting round and forecast horizon available in the sample, which allows to analyze how the previous magnitudes have evolved through time. Figure 1 plots the value of the minimum Kullback-Leibler measure resulting from the copula estimation process for all the forecasting rounds, and selected forecast horizons. Remember that this measure can be interpreted as the expected difference between the 'true' density and the approximated one. Therefore, although its value can be difficult to interpret at first, since it does not have an upper boundary, the closer it is to zero, the better the copula model fits the joint density. From this plot we can see that the fit is better for long term forecasts than for short term forecasts, and that the fit for $h = 1$ and $h = 2$ improves in the second half of the sample. Only in one case is the value of this divergence larger than 0.5. In order to check how these results compare when other copulas are used, a robustness exercise has been performed. Focusing on the last release available in the sample (2015 Q4), the density of the information shocks has been obtained for all forecast horizons up to two years using 6 popular one-parameter copulas: Gaussian, Frank, Farlie-Gumbel-Morgenstern, Clayton and Gumbel. For each copula, the estimate of the copula parameter for each forecast horizon and its standard error (calculated from the Hessian) are reported, along with the value of the Kullback-Leibler measure at the minimum and the first three moments of the information shocks. Table 1 summarises these results. Notice that, since the Kullback-Leibler is directly comparable across models, it can be used to decide which copula fits better the data. The Frank copula has the lowest divergence measure for all forecast horizons, confirming its suitability for the analysis. Similar to the pattern identified previously, the fit seems to be better for long term forecasts than for short term forecasts, and this is true for all copulas analysed. Notice that, although the fit of the Gaussian copula is good, the estimator provides a corner solution (as the Farlie-Gumbel-Morgenstern does). The Gaussian copula is known to be not well behaved in the $[0,0]$ and $[1,1]$ regions, so it seems to be not flexible enough to model some of the cases considered here. Copulas such as Farlie-Gumbel-Morgenstern, Clayton or Gumbel produce very imprecise estimates of the corresponding copula

parameter, with extremely large values for their standard errors. These copulas are more suitable to model dependence in extreme values, and hence their apparent unsuitability in this case. Notice also that, while the means of the information shocks remain quite stable across copulas, higher moments seem to depend on the choice of the copula. The lack of precision of some copulas might be the reason for this. Given the good fit of the Frank copula as compared to competitors, and the properties that make it an intuitively good candidate for this analysis, results presented later in the section are based on this copula.

Table 1: Estimation results of the copula parameters and moments of the information shock for several copulas. Forecasts release date: 2015 Q4.

	Forecast horizon							
	h=1	h=2	h=3	h=4	h=5	h=6	h=7	h=8
Gaussian								
$\hat{\theta}$	0.999	0.999	0.999	0.999	0.999	0.990	0.999	0.999
$se(\hat{\theta})$	0.1136	0.0484	0.1066	0.1141	0.0135	0.3801	0.0692	0.0421
KL divergence	0.3643	0.2480	0.1321	0.0550	0.0427	0.0226	0.0417	0.0618
Mean	-0.2338	-0.3610	-0.4771	-0.5787	-0.5916	-0.6172	-0.5852	-0.5415
Variance	0.0861	0.0892	0.0275	0.0127	0.0243	0.0389	0.0310	0.0177
Skewness	0.5813	1.3120	5.4745	3.0835	2.6221	-0.5951	1.8756	2.6688
Frank								
$\hat{\theta}$	13.1095	21.2762	27.4918	21.5263	18.0450	18.6482	20.6416	17.1974
$se(\hat{\theta})$	24.8428	1.4434	1.0392	1.2168	7.2840	6.0764	1.3978	6.3639
KL divergence	0.1891	0.0770	0.0288	0.0135	0.0074	0.0085	0.0105	0.0143
Mean	-0.2599	-0.3704	-0.4725	-0.5721	-0.5824	-0.6171	-0.5798	-0.5300
Variance	0.2176	0.1897	0.1537	0.2210	0.2997	0.2952	0.2493	0.3550
Skewness	0.1800	0.2636	0.3246	-0.0220	-0.0773	-0.1208	-0.1392	-0.0861
Farlie-Gumbel-Morgenstern								
$\hat{\theta}$	1	1	1	1	1	1	1	1
$se(\hat{\theta})$	12.4509	11.7445	11.2422	10.7888	10.6099	10.5946	10.6499	10.7898
KL divergence	0.3426	0.2593	0.2070	0.1658	0.1507	0.1471	0.1500	0.1590
Mean	-0.2600	-0.3703	-0.4713	-0.5729	-0.5861	-0.6283	-0.5917	-0.5476
Variance	0.8116	1.3960	1.9625	2.3160	2.4741	2.5512	2.6261	2.7686
Skewness	0.0560	0.0420	0.0319	-0.0036	-0.0268	-0.0592	-0.0792	-0.0739
Clayton								
$\hat{\theta}$	37.8167	50.0000	34.8285	14.5037	31.8995	11.8634	17.1206	28.1105
$se(\hat{\theta})$	110.7537	618.1363	240.2038	55.0693	96.7548	29.5588	97.9643	111.1135
KL divergence	0.3462	0.2062	0.1356	0.0882	0.0516	0.0627	0.0529	0.0695
Mean	-0.2216	-0.3263	-0.4714	-0.5822	-0.5880	-0.6349	-0.6070	-0.5420
Variance	0.1461	0.1728	0.0887	0.1497	0.0765	0.1788	0.1143	0.0686
Skewness	0.3803	1.1834	0.5676	-0.5534	-0.0271	-0.9822	-0.7806	-0.0752
Gumbel								
$\hat{\theta}$	9.5690	50.1587	13.5784	10.1426	9.4449	9.1797	10.0401	11.7256
$se(\hat{\theta})$	39.0679	618.1363	88.6320	23.7282	22.8291	26.1719	31.4769	46.6867
KL divergence	0.3332	0.1808	0.1349	0.0636	0.0378	0.0255	0.0360	0.0530
Mean	-0.2893	-0.4678	-0.4888	-0.5832	-0.5851	-0.6265	-0.5888	-0.5318
Variance	0.1136	0.1263	0.0511	0.0596	0.0692	0.0797	0.0657	0.0495
Skewness	0.3821	0.6107	2.5777	1.1882	0.4425	0.3454	0.2911	0.8635

3.2 Measures of Inflation Uncertainty and Information Absorption

Figure 2: Moments of the information shocks

(a) variance

(b) skewness

Figure 2 plots the variances (panel a) and skewness coefficients (panel b) of the information shocks for selected forecast horizons and all forecasting rounds as described in the previous section. The skewness coefficient at a given round and forecast horizon can be interpreted as a measure of the upward or downward risk that the MPC understands that information made available during the last quarter will pose on future inflation. A positive value of the skewness coefficient is

interpreted as a sign that the MPC expects that new information will increase the probability of an inflationary process at a given point in the future with respect to the assessment made the previous quarter. From panel (b) in figure (2) we can see that the MPC seems to judge information to have a higher effect on upward or downward inflationary pressures for short term forecasts than for long term ones. However, after the financial crisis, this difference attenuates as the skewness coefficient at all horizons converge to zero. This is complemented with the picture painted by the variance of the information shocks presented in panel (a). This variance can be interpreted as a measure of the effect that the MPC assesses that new information will have on inflation uncertainty. The graph shows a decrease in the variance of the information shocks at all horizons during the *Great Moderation*. Along the line of common knowledge, the variance of the shocks seems to be generally higher for long term rather than short term forecast horizons. According to the interpretation given for this variance, it can be said that the Bank assumed that new information introduced increasingly less uncertainty into the forecasts. The Bank seemed to have perceived the macroeconomic increasingly more certain with respect to the evolution on inflation, as new information contributed marginally less to uncertainty. As pointed out by one referee, this decrease in uncertainty could also well be due to the Bank improving its density forecasting practice. It may also be due to a progressive adjustment of the variances reported in the fan charts in order to address the criticism about the possible overestimation of uncertainty at the start of the sample (see e.g. Dowd, 2007). In any case, the general reduction of uncertainty and the convergence of this measure for all forecasting horizons in the months previous to the start of the financial crisis speak of a period in which information did not change substantially the Bank's assessment of uncertainty. However, there is a sharp increase in this measure at all forecast horizons during the financial crisis of the late 2000s. Although it has stabilized since the beginning of 2012, it remains at levels higher than previous to the crisis. Combining the results derived from the second and third moments of the information shocks, we can distinguish two separate periods in the forecasting practice of the Bank: before the financial crisis the Bank was relatively more confident about its forecasting practice, judging whether new information posed an upward or downward risk on inflation with increased certainty. However, after the financial

crisis, the skewness coefficients of the information shocks become zero at all horizons, and their variances increase sharply. This can be interpreted as a sign that the Bank understands that new information will have an important impact on inflation, but it might not be able to clearly discern on which side of the distribution it will impact.

To test for the validity of this interpretation of the measure of uncertainty, it has been compared with another popular measures of uncertainty available in the literature: the UK Economic Policy Uncertainty (EPU) index of Baker et al. (2016). The rank correlation coefficient between both measures is positive and significant, and higher than 0.5 for all forecast horizons. It is reasonable to expect these results, as both measures are based on news events. However, there are three main differences between the uncertainty measure presented here and the ones already in the literature: first, this measure is an *ex-ante* measure of uncertainty. It is observed at the moment of making the forecast, and no realisation of the forecasted variable is needed to construct it. Second, the measure is based on the process of acquisition and incorporation of information into the forecasts, so it has a substantial degree of judgment that could be useful as an input to some forward looking models. And third, since the forecaster in this application is the institution in control of the monetary policy, it could well be that this measure of expected uncertainty acts as one of the variables that the MPC has into account (albeit unconsciously) when taking some kind of policy decisions. Furthermore, if this is so and economic agents anticipate it, it would be expected that participants in the economy will take this measure of uncertainty into account when trying to anticipate monetary policy reactions. Hence, this measure could have some impact on other macroeconomic and financial variables. A full study of this mechanism is left for further study.

Finally, figure 3 shows the measure of information absorption, which indicates the degree of incorporation of new information into subsequent forecasts. For short term forecasts this measure is quite volatile, although for all horizons it fluctuates around 0.6. This result suggests that approximately 60% of the information contained in a revised forecast comes from information made available since the last release. The remaining proportion would be information already contained in previous forecasts. Regarding long term forecasts (one and two years ahead) the proportion of new information content is similar, although after the financial crisis this proportion

Figure 3: News absorption

is slowly declining. This seems to indicate that, while new events still represents a high proportion of the information content of a revised forecast, in a period of greater uncertainty these news are taken with more caution before incorporating them into long term forecasts. An open question is to decide, as Lahiri and Liu (2006) point out, how much of the change in a forecast for a fixed point event is 'news' and how much is 'noise'.

4 Conclusions

This paper proposes a method to extract the density of the information shocks from the revisions of the inflation density forecasts produced by the Bank of England. An information shock contains the set of information that has been perceived by the forecaster, and that has been incorporated into the revised forecast. An examination of the moments of the information shocks can help also understand to some extent how the Bank has processed new information. The variance of the information shock at a given forecast horizon and release is interpreted here as a measure of the uncertainty that the Bank of England expects that information perceived between that forecasting exercise and the previous one will have on inflation at a given point in the future. This could be understood as a measure of the uncertainty caused by a particular set of information. The skewness

coefficient can be interpreted as a measure of how new information has change in the Bank's assessment of risks regarding the level of inflation. The method presented here allows to recover the information shocks for every forecasting exercise and every forecast horizon. This allows to analyse, for example, the evolution of uncertainty over time and its term structure. In particular, it is found that before the financial crisis, the Bank was able to assess whether new information posed an upward or downward risk on future inflation as the skewness coefficients of the information shocks where nonzero. However, after the financial crisis, as new information introduce more uncertainty, the Bank has not been fully able to decide which side of the distribution this new information will impact on. This paper also proposes a measure of information absorption into subsequent forecasts. It is shown that new information represents around 60% of the information content of a revised forecast, being the remaining part information already contained in previous forecasts. This seems to be the for all forecast horizons, although this proportion seems to be slowly decreasing for long term forecasts after the financial crisis.

Since the measure of inflation uncertainty proposed here is an ex-ante measure, and that it is based on the subjective incorporation of information by the Bank of England, that is in charge of the monetary policy in the UK, this measure can be used to check whether market participants take this measure into account when anticipating changes in monetary policy. The method proposed in this paper to extract the densities of the information shocks could also be used to extend the tests of forecast efficiency based on revisions *à la* Mincer and Zarnowitz (1969) for point forecasts or Mitchell (2008) for density forecasts, to a dynamic setting. As described by Mankiw and Shapiro (1986), Nordhaus (1987) or Patton and Timmermann (2012), forecasts are said to be efficient if the differences between two consecutive forecasts for a fixed point event have zero mean and their variances decreases as the forecasted event becomes closer in time. Since the information shock can also be interpreted as the difference between two forecasts for the same event, and given that the procedure allows to obtain the density of this difference for every forecasting exercise and forecast horizon, a test could be constructed to check whether these regularities hold for every forecasting round and not on average over a sampling period as it is traditionally done. This kind of testing procedure could help the researcher identify specific

periods in which efficiency is lost, which in turn will help the forecaster improve the forecasting practice. This is left for future research.

Appendix

Let $f(x)$ and $g(x|\theta)$ two density functions. Assume that $f(x)$ represents the density of the 'true' data generating process, and $g(x|\theta)$ is a model used to approximate it that depends on the unknown parameter θ . The Kullback-Leibler divergence is defined as

$$KL = \int f(x) \log \left(\frac{f(x)}{g(x|\theta)} \right) dx. \quad (10)$$

If $f(x) = g(x|\theta)$ everywhere, $KL = 0$, and $KL > 0$ otherwise. Notice that we can rewrite (10) as

$$KL = \int f(x) \log(f(x)) dx - \int f(x) \log(g(x|\theta)), \quad (11)$$

$$= E_f[\log(f(x))] - E_f[\log(g(x|\theta))], \quad (12)$$

so that the Kullback-Leibler divergence can be interpreted as the expected difference between the log scores of both models, where the expectation is taken with respect to the true distribution $f(x)$.

The minimum relative entropy principle implies choosing θ so that KL is minimal. The first order condition of this minimization problem can be expressed, from (12), as

$$E_f \left[\frac{d \log(g(x|\tilde{\theta}))}{d\theta} \right] = 0. \quad (13)$$

To show the equivalence of this method with the method of maximum likelihood, consider an *iid* random sample of size n , $\{x_i\}_{i=1}^n$, drawn from the true model $f(x)$. Equation (13) can be rewritten as

$$\frac{1}{n} \left[\sum_{i=1}^n \frac{d \log(g(x_i|\hat{\theta}))}{d\theta} \right] = 0,$$

which is the expression of the first order condition of the maximum likelihood estimation problem. Under some regularity conditions (see e.g. White, 1994) the MLE estimator $\hat{\theta}$ converges almost surely to $\tilde{\theta}$ as $n \rightarrow \infty$, the minimum relative entropy estimator.

References

- Akaike, H. (1973). Information Theory and an Extension of the Maximum Likelihood Principle. 2nd International Symposium on Information Theory, 267-281.
- 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), 483-500.
- Bank of England (2015). Evaluating Forecast Performance. Independent Evaluation Office. November 2015.
- Baker, S. R., N. Bloom and S.J. Davis (2016). Measuring economic policy uncertainty. Forthcoming in *The Quarterly Journal of Economics*. doi: 10.1093/qje/qjw024
- Bao, Y., T.-H. Lee and B. Saltoglu (2007). Comparing Density Forecast Models. *Journal of Forecasting*, 26, 203-225.
- Basu, A., H. Shioya and C. Park (2011). *Statistical Inference: The Minimum Distance Approach*. CRC Press.
- Biau, G. and M. Wegkamp (2005). A Note on Minimum Distance Estimation of Copula Densities. *Statistics and Probability Letters*, 73, 105-114.
- 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, 30-37.
- Burnham, K.P. and Anderson, D.R. (2002). Model Selection and Multimodel Inference. A Practical Information-Theoretic Approach. Second Edition. Springer: New York.
- Clements, M.P. (2004). Evaluating the Bank of England Forecasts of Inflation. The Economic Journal, 114, 844-866.
- Clements, M.P. (2014). Forecast Uncertainty - Ex Ante and Ex Post: US Inflation and Output Growth. Journal of Business and Economic Statistics, 32(2), 206-216.
- Clements, M.P. and Smith, J. (2000). Evaluating the Forecast Densities of Linear and Non-linear Models: Applications to Output Growth and Unemployment. Journal of Forecasting, 19, 255-276.
- 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), 65-102.
- Chu, B. (2011). Recovering Copulas from Limited Information and an Application to Asset Allocation. Journal of Banking and Finance, 35, 1824-1842.
- Diebold, F.X. and Lopez, J.A. (1996). Forecast Evaluation and Combination. In The Handbook of Statistics, vol. 14. G.S. Maddala and C.R. Rao eds.
- Dowd, K. (2007). Too Good to be True? The (In)credibility of the UK Inflation Fan Charts. Journal of Macroeconomics, 29, 91-102.
- Elder, R., Kapetanios, G., Taylor, T. and Yates, T. (2005). Assessing the MPC's Fan Charts. Bank of England Quarterly Bulletin, 45, 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, 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), 713-727.
- Genest, C. (1987). Frank's Family of Bivariate Distributions. *Biometrika*, 74(3), 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), 411-422.
- Granger, C.W.J. and Newbold, P. (1986). *Forecasting Economic Time Series* (2nd ed.). Academic Press: New York.
- Gronnerberg, S. and Hjort, N. L. (2014). The Copula Information Criteria. *Scandinavian Journal of Statistics. Theory and Applications*, 41, 436-459.
- Isiklar, G. and Lahiri, K. (2007). How Far Ahead Can We Forecast? 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.
- Jurado, K., Ludvigson, S.C. and Ng, S. (2015). Measuring Uncertainty. *American Economic Review*, 105(3), 1177-1216.
- Lahiri, K. and Liu, F. (2006). Modelling Multi-Period Inflation Uncertainty Using a Panel of Density Forecasts. *Journal of Applied Econometrics*, 21, 1199-1219.
- Mankiw, N.G. and Shapiro, M.D. (1986). News or Noise? An Analysis of GNP Revisions. NBER Working Paper No. 1939.
- Miller, D.J. and Liu, W. (2002). The Recovery of Joint Distributions from Limited Information. *Journal of Econometrics*, 107, 259-274.

- 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, 995-1033.
- Nordhaus, W.D. (1987). Forecast Efficiency: Concepts and Applications. *Review of Economics and Statistics*, 69(4), 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), 1-17.
- Rossi, B. and Sekhposyan, T. (2015). Forecast Rationality Tests in the Presence of Instabilities, with Applications to Federal Reserve and Survey Forecasts. *Journal of Applied Econometrics*, 31(3), 507-532.
- Sklar, A. (1959). Fonctions de repartition á n dimensions et leurs marges. *Publ Inst Univ Paris*, 8, 229-231.
- Soofi, E.S. (2000). Principal Information Theoretic Approaches. *Journal of the American Statistical Association*, 95(452), 1349-1353.
- Tay, A.S. and Wallis, K.F. (2000). Density Forecasting: A Survey. *Journal of Forecasting*, 19, 235-254.
- White, H. (1994). *Estimation, Inference and Specification Analysis*. Cambridge University Press. Cambridge, UK.
- Wallis, K.F. (2004). An Assessment of Bank of England and National Institute Inflation Forecast Uncertainties. *National Institute Economic Review*, 189, 64-71.