



BANK FOR INTERNATIONAL SETTLEMENTS

## **BIS Working Papers**


No 339

# Measuring disagreement in UK consumer and central bank inflation forecasts

by Richhild Moessner, Feng Zhu and Colin Ellis

Monetary and Economic Department

February 2011



JEL classification: C14,C82,E31,E58

Keywords: Adaptive kernel method, adaptive multimodality test, consumer survey, inflation forecasts, nonparametric density estimation.

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

Copies of publications are available from:

Bank for International Settlements  
Communications  
CH-4002 Basel, Switzerland

E-mail: [publications@bis.org](mailto:publications@bis.org)

Fax: +41 61 280 9100 and +41 61 280 8100

This publication is available on the BIS website ([www.bis.org](http://www.bis.org)).

© *Bank for International Settlements 2011. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)

ISBN 1682-7678 (online)

# Measuring disagreement in UK consumer and central bank inflation forecasts\*

Richhild Moessner

Feng Zhu

Bank for International Settlements

Bank for International Settlements

Colin Ellis

University of Birmingham

## Abstract

We provide a new perspective on disagreement in inflation expectations by examining the full probability distributions of UK consumer inflation forecasts based on an adaptive bootstrap multimodality test. Furthermore, we compare the inflation forecasts of the Bank of England's Monetary Policy Committee (MPC) with those of UK consumers, for which we use data from the 2001-2007 February GfK NOP consumer surveys. Our analysis indicates substantial disagreement among UK consumers, and between the MPC and consumers, concerning one-year-ahead inflation forecasts. Such disagreement persisted throughout the sample, with no signs of convergence, consistent with consumers' inflation expectations not being 'well-anchored' in the sense of matching the central bank's expectations. UK consumers had far more diverse views about future inflation than the MPC. It is possible that the MPC enjoyed certain information advantages which allowed it to have a narrower range of inflation forecasts.

**Keywords:** Adaptive kernel method, adaptive multimodality test, consumer survey, inflation forecasts, nonparametric density estimation.

**JEL Classification:** C14, C82, E31, E58

---

\*The views expressed in this paper belong to the authors and do not represent those of the Bank for International Settlements (BIS). We would like to thank seminar participants at the BIS for helpful comments.

# 1 Introduction

Inflation expectations play a key role in the monetary transmission mechanism, and central banks carefully monitor the public's inflation expectations for monetary policy purposes. Accurate measures of private sector inflation expectations matter for monetary policy for several reasons. First, as inflation expectations are an essential element in price and wage setting processes, they can have a significant impact on actual inflation outturns. For example, Leduc, Sill and Stark (2007) found that the US Federal Reserve accommodated temporary shocks to expected inflation before 1979, which then led to persistent increases in actual inflation. Second, changes in private sector inflation expectations are reflected in expected real rates of interest, which affect consumer spending and business investment decisions. Third, inflation expectations serve as a useful indicator of the effectiveness and credibility of monetary policy for central banks aiming to achieve price stability. Inflation expectations help guide monetary policy, and they are a key input for regular monetary policy deliberations.

Inflation forecasts from independent sources are valuable to central banks, and survey-based measures of inflation expectations provide rich information. Based on data extracted from surveys of households and professional forecasters, past empirical work has mostly focused on three areas, besides measurement and accuracy issues. First, following the seminal work of Lucas and Sargent on rational expectations, researchers have examined the expectations formation process by testing alternative hypotheses of adaptive, regressive and rational expectations. Expectations are rational if no systematic errors are made in agents' inflation forecasts, and they are efficient if agents make use of all existing relevant information when making their forecasts.

The evidence on rationality is mixed. Based on survey data, Figlewski and Wachtel (1981), Gramlich (1983), Pesaran (1985) and Zarnowitz (1985) all rejected the null hypotheses of unbiased and efficient inflation forecasts. Dotsey and DeVaro (1995) and DeLong (1997) reported a statistically significant bias in inflation expectations, the evolution of which lags behind actual inflation developments. In particular, expectations tended to underestimate inflation in periods of high and rising inflation, and overestimate it in periods of low and falling inflation. Ang, Bekaert and Wei (2006) suggested that inflation forecasts from the Livingston survey, University of Michigan survey and the Survey of Professional Forecasters were biased. Forsells and Kenny (2002) found evidence of lack of rationality in private agents' expec-

tations for the euro area. On the other hand, Rich (1989) found no evidence of departures from rationality for inflation forecasts based on the University of Michigan surveys, and Keane and Runkle (1990) also could not reject rationality. Using real-time data, Croushore (2006) found no evidence of bias in inflation forecasts based on the Livingston Survey and the Survey of Professional Forecasters.

A second line of research examines the extent of uncertainty and disagreement, exploring heterogeneity in existing survey data of inflation expectations. If inflation expectations are not uniform and disagreement among economic agents is substantial, this could have a significant impact on monetary policy transmission, for instance in terms of the future impact of relative price shocks, and the responses of agents to changes in both inflation and policy. Analysing fifty years of US data based on surveys of households and professional forecasters, Mankiw, Reis and Wolfers (2003) documented substantial disagreement in inflation expectations. Giordani and Söderlind (2003) found that disagreement was a better proxy of inflation uncertainty than indicated by earlier literature, while D'Amico and Orphanides (2006) suggested that disagreement about the mean inflation forecast might be a weak proxy for forecast uncertainty. Rich and Tracy (2006) examined matched point and density forecasts of inflation from the Survey of Professional Forecasters to analyse the relationship between expected inflation, disagreement, and uncertainty in the United States. Boero, Smith and Wallis (2008) have examined measures of disagreement and uncertainty using the Bank of England's Survey of external forecasters. They found that inflation uncertainty was reduced following the granting of operational independence to the Bank of England in 1997.

Third, a number of economists have evaluated policymakers' ability to forecast inflation using survey data. Measures of inflation expectations have been increasingly used to assess macroeconomic prospects, and have served as a useful input in determining the stance of monetary policy. A growing number of central banks have adopted inflation targeting, including the Bank of Canada, the Bank of England, the Sveriges Riksbank, the Reserve Bank of New Zealand, and the Central Bank of Norway. These central banks regularly publish their central projections of future inflation along with confidence bands of the forecasts, known as fan charts. The fan charts of the Bank of England's Monetary Policy Committee (MPC), which are published in the Inflation Report, have been evaluated in several studies. Elder, Kapetanios, Taylor and Yates (2005) find that at most forecast horizons, inflation out-

comes have been dispersed broadly in line with the MPC’s fan chart bands, suggesting that the fan charts gave a reasonably good guide to the probabilities and risks facing the Bank. By contrast, Dowd’s (2004, 2007) findings suggest that the Bank’s fan charts overestimated inflation risk. Wallis (2003) examined the performance of the inflation density forecasts over horizons of zero to four quarters ahead, and found that his model performed well over the current-quarter horizon but over-predicted four-quarter-ahead inflation risks; Clements (2004) obtained similar results. Wallis (2004) found that the central point forecasts of the fan charts were unbiased, but also that the MPC’s density forecasts substantially overstated forecast uncertainty.

Our paper aims to make contributions to literature in the following aspects. First, despite the recent attempts to use more distributional information from survey inflation expectations, little work has been done to analyse the evolution of the full-range distributions of such forecasts. An accurate description of these could provide rich information and sharpen debates on the true extent of uncertainty and disagreement about inflation expectations. In light of this, we propose a new perspective for studying disagreement in consumer inflation forecasts. In particular, the existence of multiple modes in consumer forecast densities is proposed as evidence for substantial disagreement. The extent of disagreement in consumer forecasts is one indicator of the effectiveness of monetary policy. Second, we compare policymakers’ best probabilistic inflation forecasts with the probability distribution of survey-based inflation forecasts. Disagreements between the two sets of forecasts could be a good measure of policymakers’ ability to anchor inflation expectations.

We use a relatively new survey data set on UK consumers’ inflation expectations, and characterise UK inflation expectations beyond the first moments. We find that first moments evolve in a similar way for UK consumer inflation expectations and for the MPC’s inflation forecasts, but measures based on higher moments suggest that the two distributions behave rather differently. This is important because summary statistics become less useful when probability distributions are not unimodal as usually assumed for inflation forecasts. We therefore use adaptive kernel density estimation and mode testing methods to obtain more accurate estimates of consumer inflation forecast densities as a whole. This analysis is free from any *a priori* parametric assumptions about the functional form of inflation forecast densities, instead relying entirely on the observed data set to generate the distribution. We identify the existence of multiple modes in UK consumer inflation forecasts,

around 0% inflation and at inflation rates higher than the central modes. There was no sign of convergence in consumer forecasts, and forecasts even diverged somewhat in 2006 and 2007. The increase in the probability mass surrounding the higher mode towards the end of the sample period could be associated with a sustained rise in energy prices.

Furthermore, we also examine policy effectiveness or the ability of the MPC to anchor private sector inflation expectations, by comparing the MPC's inflation forecast densities with density estimates of consumer inflation forecasts over time. We find substantial discrepancies between UK consumers' and the MPC's inflation forecasts throughout the sample period. Most obviously, we find that UK consumers' inflation expectations were far more dispersed than the MPC's inflation forecasts: the 90% confidence intervals of the MPC's inflation forecast fan charts cover less than 50% of the probability mass of consumers' forecasts. Consumers tend to have a far more diverse view about one-year-ahead inflation. The apparent discrepancies between the MPC's and consumers' views could be explained by an information advantage and analytical ability of the MPC, further enhanced by a policy advantage, as the MPC can change interest rates if inflation moves too far from its central forecast.

The paper is organised as follows. Section 2 describes the consumer survey data we use, and provides a detailed account of our methodology and proposed new measures of disagreement, Section 3 presents the main empirical results, and Section 4 concludes.

## 2 Data and methodology

The first section describes the data. We then outline the methodology for assessing disagreement in consumer inflation expectations and for evaluating policymakers' inflation forecasts. We examine whether UK consumers' inflation expectations were homogeneous, stable and converging over time, and whether these were consistent with the MPC's fan charts over our sample.

## 2.1 The GfK NOP consumer survey data<sup>1</sup>

In an effort to better gauge inflation expectations and attitudes towards monetary policy, the Bank of England decided to conduct quarterly surveys of the general public. After trials from November 1999 to November 2000, the Bank of England/GfK NOP Inflation Attitudes Survey was first published in February 2001. The survey is conducted on behalf of the Bank by GfK NOP, as part of their regular Omnibus survey. GfK NOP uses random-location samples designed to be representative of adults living in the United Kingdom, and interviews are carried out face-to-face in private homes. Individual data are then weighted to match the UK's demographic profile, and summaries of these weighted data are subsequently published on the Bank of England's website.

The quarterly surveys are carried out after the publication of the Bank of England Inflation Report in February, May, August and November. The full survey has a total of 14 questions,<sup>2</sup> which cover the relationship between interest rates and inflation, views on past and future interest and inflation rates, the impact of inflation and interest rates on the economy and individuals, and how satisfied people are with the way the Bank of England is doing its job of setting interest rates to meet the inflation target.

However, only nine questions are asked every quarter, as early trials showed that responses to the other five questions varied little from quarter to quarter, and the full 14 questions are only asked once a year, in February. Each February survey includes about 4000 respondents (Ellis 2006), double the size of surveys in other quarters. The relatively large number of survey respondents allows a more accurate estimation of consumer inflation forecast densities, and disagreement among respondents could provide a good proxy for uncertainty about inflation forecasts. We focus on the February surveys conducted between 2001 and 2007, excluding the initial trial period and the period when the recent global financial crisis erupted. Focusing on February surveys allows us to exploit larger sample sizes and to avoid possible seasonal variations in the response patterns. The endpoint in 2007 was chosen in order to analyse a period marked by low and stable inflation and when inflation expectations were widely regarded as well-anchored, avoiding a pos-

---

<sup>1</sup>This paper uses data on the GfK NOP consumer surveys published by the Bank of England on its website and in its Quarterly Bulletins (see <http://www.bankofengland.co.uk/statistics/nop/index.htm>).

<sup>2</sup>See Ellis (2006) for details of the survey questions.



sible large shift in consumer sentiment and inflation forecasts caused by the recent financial crisis and the subsequent recession.

There are concerns that consumer forecasts might not accurately reflect the true extent of consumer disagreement about future inflation trends. This could be due to differences in the composition of consumption baskets of each individual survey respondent. Furthermore, each surveyed individual could also interpret survey questions in a different way. However, the majority of consumers tend to consume similar goods and services in similar proportions, and consumers can observe, sometimes with acute awareness, the price evolution of certain goods and services they do not purchase directly or frequently. Consumers also tend to be heavily influenced by opinion leaders and views of people close to them. Happily, the relatively large number of survey respondents helps to average out individual oddities or idiosyncracies. Most importantly, central banks still need to know whether and to what extent consumers perceive and forecast inflation differently, knowing the inherent differences between individuals. Measures of disagreement would inevitably reflect differences in individual circumstances and abilities to gather and process information, not solely differences in opinion.

The consumer survey data provided by GfK NOP are in the form of histograms of 8 bins. Except for the two bins at the left and right ends and the point mass at 0% inflation level, the five other bins are centred around 0.5%, 1.5%, 2.5%, 3.5% and 4.5%, with a  $\pm 0.5\%$  band around each bin centre. The bins located at the two extremes are open-ended, covering all values below 0% and above 5%, respectively.

To infer the underlying probability distributions for the GfK NOP survey data for one-year-ahead UK consumer inflation expectations, we need to make some assumptions on how they relate to the histograms. We use three different data resampling schemes. We first count the number of respondents in each of the eight bins in the original histogram. Taking these numbers as given, we randomise the consumer inflation forecasts data “within” each bin, except for the point mass at 0% inflation (*i.e.*, responses of “no price changes” where no resampling is needed).

In the first scheme, we use a uniform distribution within each bin, assuming that a survey participant has no particular preferred inflation value within a bin, say  $[2\%, 3\%]$ , when the participant picks a response that prices have gone “up by 2% but less than 3%”. We assume that the left- and right-end bins are of interval  $[-1\%, 0\%)$  and  $[5\%, 6\%]$ , respectively, so that the whole distribution has bounded support within the  $[-1\%, 6\%]$ -range. The

motivation for doing so is that since the survey offers the choice of a range of bins, we implicitly assume that expected inflation would be within a relatively tight interval surrounding those bins; otherwise, the number or size of bins in the survey should be increased for the survey to remain informative. Another interpretation would be that values further away from the end point of the specified bins should be considered irrational or misinformed estimates.<sup>3</sup> The second scheme assumes that the consumer forecasts cluster around each bin centre. For each bin, we resample the data located within that bin with a non-truncated normal distribution with its mean equal to the central bin value and a standard deviation of 0.5% percentage points. The resulting data set will have a different number of respondents in each bin due to the resampled tail responses, but the difference is very small. The new data set is of unbounded support as it is in the case for GfK NOP Surveys. The third scheme is a combination of the first two schemes: data in the range (0%, 5%] are resampled as in the first scheme, and data in the bins [-1%, 0%) and [5%, 6%) are resampled as in the second scheme. The resampled data are again open-ended. Data produced with the resampling schemes I and III would yield exactly the same histograms as in the original surveys, but the second scheme produces slightly different histograms.

## 2.2 A test of disagreement in inflation expectations: adaptive mode testing

Economists and market practitioners often focus on the mean, median or standard deviation of survey expectations.<sup>4</sup> The first and second moments are informative, but their information is so limited that they could be similar even with substantially different underlying distributions, and may consequently fail to provide a faithful picture of the evolution of inflation expectations. Measures beyond the first and second moments help to quantify the shape of the distribution further. In this paper, we focus on ways to extract more detailed information about the distribution of inflation expectations from existing data. Specifically, we examine the evolution of the complete probability distribution of consumers' inflation forecasts in the United King-

---

<sup>3</sup>However, identification of an additional mode at the highest bin of [5%, 6%) in some years could be sensitive to this assumption.

<sup>4</sup>One exception was Diebold, Tay and Wallis (1998), who analysed the probability distribution of inflation expectations based on surveys of professional forecasters.

dom, going beyond measures of central tendency and uncertainty.

The conceptual framework we rely on to assess consumer disagreement in inflation expectations may be expressed as a finite mixture model for the GfK NOP survey data of consumer inflation forecasts  $\{\pi_i\}_{i=1}^I$ , where  $i$  indexes each individual forecast, which is assumed to be drawn from a finite mixture of  $M$  distributions, denoted by  $p(\pi)$ ,

$$p(\pi) = \sum_{m=1}^M \theta_m f_m(\pi) \quad \pi \in \mathfrak{R} \quad (1)$$

where  $\theta_m > 0$ ,  $\sum_{m=1}^M \theta_m = 1$ ,  $f_m(\pi) \geq 0$ , and  $\int_{\mathfrak{R}} f_m(\pi) d\pi = 1$ , for  $m = 1, \dots, M$ . The survey data  $\{\pi_i\}_{i=1}^I$  is said to be generated from a finite mixture of  $M$  distributions,  $M$  is the number of existing modes in the density of inflation forecasts, and  $p(\cdot)$  is a finite mixture density function. The parameters  $\theta_1, \dots, \theta_M$  are known as the mixing proportions or mixing weights, and  $f_1, \dots, f_M$  are the component densities. In some cases, it is possible to specify functional forms for the component densities. We have little or no *a priori* information about the functional forms for each component density  $f_M$ , as economic theory provides little guidance about the specific composition of the distribution of consumer inflation forecasts. Instead we adopt a nonparametric approach of bump-hunting or mode testing as advocated in Zhu (2005).

The most striking feature of a mixture density is often the presence of multiple bumps and modes, which can be identified if the component densities in a population are unimodal and are well-separated. Multimodality or the existence of multiple bumps is indicative of the existence of distinct component densities, hence population heterogeneity. The presence of multiple modes in the probability distribution of consumer inflation forecasts indicates the existence of distinct groups, each having its own probability distribution of inflation forecasts reflecting, for example, the group's distinct socio-economic knowledge and experience. Essentially, each group has a distinct economic "model" generating a unimodal distribution of forecasts. Our definition of disagreement is based on a formal statistical test of multimodality.

**Definition 1** A *mode* in a density  $f$  is a local maximum, and a *bump* with no flat part is the portion of the density  $f$  lying between two points of inflection  $a$  and  $b$ , such that  $f$  is concave in the interval  $[a, b]$  but not outside. A density  $f$  is *multimodal* if it has more than one mode.

**Definition 2** Given a sample of  $I$  individual inflation forecasts  $\{\pi_i, i = 1, \dots, I\}$  and the estimated empirical density function  $\hat{f}^I$ , consumer disagreement is said to exist in the inflation forecasts if a multimodality test concludes that  $M > 1$ , *i.e.*, there exist more than one modes or bumps in  $\hat{f}^I$ . The greater the number of modes, the deeper is the disagreement among consumers. The extent of disagreement between any pair of uncovered consumer groups can be measured as the distance between the corresponding modes.

Our definition of disagreement is rather different from the definitions which appeared in the previous literature, which focussed mostly on distances between individual point estimates, or between point estimates and the average inflation forecast. Conventional measures of disagreement are quantile-based, including range, inter-quartile range (IQR), and “quasi-standard-deviation”, the half distance between the 16th and 84th percentiles of inflation forecasts (see Giordani and Söderlind, 2003). In cases of surveys of professional forecasters where each surveyed individual often provides a density forecast, disagreement measures can be computed in terms of second moments (*e.g.*, standard deviation) instead of point forecasts. This is known as disagreement in uncertainty. Disagreement based on the concept of multimodality has the same applicability as measures such as interquartile range, if they are computed on the same type of data such as consumer forecasts. But the multimodality-based concept could be more useful and robust by providing far more information than summary statistics.

In this paper, we define disagreement in inflation expectations in terms of the complete probability distribution of consumer inflation forecasts. This has advantages over the partial snap shot provided by conventional measures. For instance, the distribution of inflation expectations may change dramatically over time but the range and IQR estimates could remain the same. In addition, by focusing on the entire distribution, disagreement measures based on mode testing allow us to classify consumers into different groups, each having its model or data generating process. Inflation expectations of consumers in different groups could therefore be generated on the basis of socio-economic knowledge and experience common to consumers within one group but rather different from group to group. However, we do not pursue the factors underlying the grouping of consumer inflation forecasts, given data limitations and a lack of guidance from economic theory.

We use an adaptive version of kernel density estimation and bootstrap

multimodality test procedures as proposed in Zhu (2005).<sup>5</sup> The original fixed-bandwidth Silverman (1981, 1983) bootstrap multimodality test is intuitive and easy to implement, but it can be sensitive to spurious noise in the tails. Without properly controlling the sensitivity to local data density, multimodality tests are error-prone. Adaptive kernel density estimation makes it possible to purge spurious noise in the tails while preserving important characteristics in the fat part of the distribution. Bandwidth tests based on the adaptive method are more accurate. Zhu's (2005) adaptive and simultaneous version of the test is designed to cope with the relative sparseness of data in the tails and possible fluctuations in the estimated  $p$ -values.

### 2.3 The MPC's inflation forecasts

Each quarter, the Bank of England's Monetary Policy Committee (MPC) publishes its latest forecasts in the Inflation Report. The forecasts are presented in probabilistic terms, acknowledging explicitly uncertainties surrounding such projections. The so-called "fan charts", a graphical representation of such uncertainties, are presented in sequenced 10%-probability bands.<sup>6</sup> The central band contains the single most likely outcome (central projection, or the mode). The further away from the central band, the greater the degree of uncertainty, in terms of moving into the tails of the forecast distribution. Since 1997, the MPC has published the summary parameters for its fan charts on the Bank of England's internet site, which can be used to reconstruct the fan charts, and indeed the entire probability distribution for the MPC's forecasts. The MPC's probability distributions (reflecting uncertainty of the MPC) can be compared with the empirical distributions of consumer inflation expectations (reflecting disagreement among consumers, which may be taken as a proxy for uncertainty).<sup>7</sup>

There are two caveats with this exercise. First, in August 2004 the MPC shifted emphasis to focus more on forecasts and fan charts based on market rate expectations, as opposed to constant rates (this emphasis has arguably reversed recently). This change could, in principle, represent improved MPC

---

<sup>5</sup>See Appendix III for details.

<sup>6</sup>Britton *et al* (1998) and Bank of England (2002) provide more information on the fan charts.

<sup>7</sup>Uncertainty and disagreement are not identical. But disagreement is expected to be generally correlated with uncertainty. Appendix II provides details for the two-piece normal distribution which is used to reconstruct the MPC's probability distributions.

forecasts, but was unlikely to have had a significant impact on UK consumers' inflation expectations.

Second, the MPC switched to forecasting consumer prices index (CPI) inflation in February 2004, after the Government changed the Bank of England's inflation target to 2% CPI inflation in December 2003 from 2.5% RPIX inflation.<sup>8</sup> The two price indices evolved in similar patterns but differed in several ways, with CPI inflation being lower than RPIX inflation throughout our sample period. In order to make CPI and RPIX fan charts comparable, we need to make adjustments in the two sets of fan charts. Bank of England (2004) noted that the long-run gap between CPI and RPIX inflation was about 0.75 percentage points, a gap resulting from methodology and coverage.<sup>9</sup> We use this fact to adjust the fan charts. Our calculations confirmed that the average gap between the two series over the sample period has been of a similar size (Figure 1).

We focus on RPIX-based fan charts (ie we adjust the CPI fans) for two reasons. First, the Bank initially focused on the RPIX when the GfK NOP survey was launched, and the published fan charts were based on this.<sup>10</sup> Second, Goldman (2007) noted that the vast majority of UK wage settlements which made reference to a price index still used the RPI, despite a recent drive to move public sector wage negotiations towards a CPI benchmark. As RPI was probably more widely recognised during our sample, it is likely that the general public paid more attention to it than to other indices. In addition, median survey forecasts appeared to be closer to RPIX inflation than CPI inflation during our sample.

In the first line of analysis, we compare consumers' views of inflation outcomes one year ahead, as well as measures of uncertainty derived from these, with those based on Inflation Report fan charts since 2001. This allows us to quantify the disagreement between UK consumers' and the MPC's forecasts, and to evaluate the "well-anchoredness" of UK inflation expectations.

We also examine more than just the standard measures of disagreement. In particular, we consider robust moments as measures of disagreement among UK consumers. In the presence of outliers and thick tails, as in the case of the consumer surveys, it is useful to consider higher moments based

---

<sup>8</sup>RPIX is the retail prices index excluding mortgage interest payments.

<sup>9</sup>Nickell (2003) discusses the difference between the RPIX and CPI (previously known as HICP) in greater detail.

<sup>10</sup>Ellis (2005) noted that the change in inflation rate target from RPIX to CPI inflation seemed not to have affected respondents' interpretation of the GfK NOP survey questions.

on the percentiles of the distribution. Robust moments are less sensitive to outliers and observations in the far tails of the distribution. Moreover, we go beyond moments as a measure of disagreement, by directly estimating the consumers’ probability distribution using non-parametric methods, and performing statistical tests for multiple modes as a measure of disagreement. We consider the presence of two or more modes in the distribution as evidence of greater disagreement than the presence of only one mode.

### 3 Empirical results

Our sample period from 2000 Q2 to 2007 Q2 was characterised by relatively low and stable inflation rates. Four-quarter RPIX and CPI inflation rates averaged 2.4% and 1.5%, respectively (see Figure 1). Inflation volatility was also very low, compared with the UK’s experience in previous decades (see Bank of England, 2007). During much of the period, prices of consumer goods were falling in the United Kingdom, while those of consumer services rose (see Figure 2). This reflected falling import prices for finished manufactured goods, in part driven by the emergence of low-cost producers such as China and India. However, CPI and RPIX goods prices were affected by a sharp increase in energy prices, particularly petrol, electricity and gas, since 2004 Q3 (see Figure 3). A pick-up in prices of imported finished goods late in the sample period appeared to have also contributed to a rise in inflation.

In this section, we first examine the evolution of summary statistics, *i.e.* the standard and “robust” moments of the one-year-ahead consumer inflation forecasts and forecasts published in Bank of England Inflation Reports. We then discuss results on disagreement based on the estimation and mode testing of the entire distributions of consumer inflation forecasts. We then evaluate the “well-anchoredness” of UK consumers’ inflation expectations by comparing the empirical distributions of consumers’ inflation forecasts with our computed MPC inflation density forecasts.

#### 3.1 Robust moments of inflation forecasts

We compare the first and higher moments for one-year-ahead inflation forecasts based on UK consumer survey data with those based on the MPC’s fan charts. We consider the possibility of fat tails or a large number of outliers in the distribution of consumer inflation forecasts, and when little informa-

tion is available about the tail distribution, quantile-based “robust” moments could be useful.<sup>11</sup> The robust measures are defined on the basis of different quantiles of a distribution, they are less sensitive to outliers and help us focus more on the different parts of interests of the distribution.

The mean and median of the distribution for expected inflation derived from the consumer surveys and fan charts are shown in Figure 4.<sup>12</sup> These first moments moved in a similar fashion for much of the sample period. However, while mean and median consumer survey forecasts were little changed in 2007, the mean and median of the MPC’s forecast fell significantly. This could reflect the MPC’s view that the pickup in inflation up to 2006 was in large part due to temporary factors, and inflation would eventually fall back. Consumers, in contrast, appear either to be more backward-looking, or might have greater difficulty in determining the temporary nature of some shocks on inflation (inflation could lead to rises in consumer inflation expectations. In addition, consumer inflation expectations could be affected to a greater extent by the prices of non-discretionary expenditure (*e.g.* gas, electricity and food items), or consumers might assign a greater weight to a number of more frequently purchased items.

Figure 5 shows second moments for one-year-ahead inflation forecasts. The differences between consumers’ and the MPC’s forecasts are striking. First, as expected both the standard deviation and inter-quartile range are much larger for the consumers’ forecasts than for the MPC’s forecasts. For the consumer surveys the second moments can be considered as measures of disagreement, rather than uncertainty, while for the MPC’s forecasts they are explicitly measures of uncertainty. However, disagreement may be a good proxy for uncertainty, as suggested by Giordani and Söderlind (2003). Disagreement is generally thought to be correlated with both the level and uncertainty of inflation. The evidence in Figure 5 therefore suggests greater uncertainty surrounding consumer forecasts. On top of this, the measures of second moments for the MPC’s forecasts fell until 2006 but increased in 2007, while they increased in 2006 but fell in 2007 in the case of consumer forecasts.

Measures of skew from the inflation forecasts are presented in Figure 6. Skewness measures the degree and direction of asymmetry of a proba-

---

<sup>11</sup>See Appendix I for technical details.

<sup>12</sup>We adjust the values for 2004-2007 CPI fan-chart forecasts up by 0.75pp for compatibility with RPIX-based forecasts, as discussed above.



bility distribution. The two robust measures of skewness we use are based on quartiles and octiles, emphasising parts of the distribution closer to the centre or towards the tails, respectively. For symmetric distributions all skewness measures equal zero. The quartile-based skewness measure equals 1 for large skewness to the right,  $-1$  for large skewness to the left. The standard skewness measure appeared to have been far more volatile than its robust counterparts, for both sets of forecasts. For the MPC’s inflation forecasts, changes in all three skewness measures behaved in a similar fashion. This was not the case for consumer forecasts, where there was a substantial difference between the movements of the octile- and quartile-based skewness measures before 2003 and after 2005. This suggests major shifts in different parts of the probability distribution of consumer forecasts, which could only be revealed by examining changes in the entire probability distribution. Different skewness measures can therefore give different results depending on which part of the distribution is the focus of the specific measure.

Kurtosis, defined as the fourth central moment, measures the dispersion and shape of a distribution. Excess kurtosis measures (see Appendix I) are useful to determine the fatness of the tails and the peakedness of a distribution relative to that of a normal distribution. The fatter the tails are and the more peaked a distribution is, the higher is its kurtosis. Figure 7 shows the evolution of the kurtosis measures for inflation forecasts. However, the meaning of the standard kurtosis measure could be ambiguous, since distributions which are less-peaked but have fatter tails than a normal distribution are common. Consumer forecasts are a good example of this. The excess kurtosis for consumer forecasts ranged between  $-1$  and  $-0.5$  throughout the sample period, while the corresponding robust measures were above 0, suggesting a distribution that is less peaked or more dispersed at the centre (see Figure 7). On the other hand, since the MPC’s inflation forecasts are essentially derived from a two-piece normal distribution, both the standard and robust excess kurtosis of MPC forecasts were close to 0.

This analysis of distributional moments is revealing. While the first moments (mean and median) evolved largely in a similar way for both sets of inflation forecasts, this was not the case for higher moments. As expected, the dispersion was much greater for consumer inflation forecasts than for the MPC’s forecasts. Moreover, much of the time the consumers’ and the MPC’s forecasts appeared to have been skewed in opposite directions, reflecting different views on the likely risks around the central projections or “best guesses”. Excess kurtosis measures indicate that consumer forecasts were less

peaked and more dispersed than MPC forecasts. The quantile-based robust measures were particularly interesting. For instance, octile- and quartile-based skewness measures for consumer forecasts evolved in rather different ways, suggesting more complex intra-distributional dynamics which could not be revealed by simpler summary statistics. In the next section we go on to study probability density estimates of consumer forecasts in an attempt to uncover such dynamics.

### 3.2 Characteristics of consumer forecast distributions

We first examine three distributional statistics to obtain information about the distribution of one-year-ahead consumer inflation forecasts; we use histograms, Tukey boxplots and Q-Q plots for resampled GfK NOP survey data. Figure 8 shows the histogram estimates, and three aspects stand out. First, consumer forecasts have a wide support and dispersion is high. The tails of the distributions are fat throughout the sample period. Second, the uniform and normal resampling schemes (and hence also the mixed scheme) produced similar data properties for the 2001-2007 period. Third, besides the large spikes observed around 0% inflation, there seem to be multiple bumps or modes in the distributions of forecasts. The large probability mass at 0% UK RPIX inflation was close to or above 2%, and since the Bank of England’s inflation target was first 2.5% for the RPIX measure and subsequently 2% for CPI, it is difficult to understand why such a large proportion of survey participants expected no price changes (on average) over the coming 12 months. The apparent anomaly could be due to a misunderstanding of the survey question: survey participants might have thought that the question referred to changes in inflation rather than to changes in the price level.<sup>13</sup>

In Figure 9, we draw the Tukey boxplots (or box-whisker charts). The boxes have lines at the lower quartile, median and upper quartile values. Dashed lines extend from each end of the box to the most extreme data value within 1.5 times the inter-quartile range, beyond which lie the outliers indicated by “+”. Comparing the two boxplots, we find that the uniformly resampled data typically has a wider inter-quartile range than the normally

---

<sup>13</sup>One possible solution might be to change the GfK NOP Survey Question Q2 from “How much would you expect prices in the shops generally to change over the next twelve months?” to “How high would you expect price inflation in the shops generally to be over the next 12 months?” The same issue arises for Question Q1. Such a question would clearly require different bins for responses.

resampled data, but the latter has more outliers in both the left and right tail ends, most likely a result of open-ended sampling. After an initial rise to 2.6% in 2003, the median inflation forecasts of the uniformly and normally resampled data both fell to 2.2% in 2005, probably a result of lower import prices and successful policy actions. However, median inflation forecasts jumped to 2.8% in 2006 and stayed roughly the same in 2007.

The Q-Q plots in Figures 10 and 11 compare the uniformly and normally resampled data for each sample period with data generated from a normal distribution with mean and standard deviation equal to those estimated from the resampled data. The plots allow us to visually examine the hypothesis that the resampled survey data come from a normal distribution. Inspecting the graphs, we find that while normal distributions appear to fit the central part of distributions in the interval [1%, 4%] well, they fail to account for the fat tails at both ends. This result points to a platykurtic distribution for UK consumer inflation forecasts, contrary to Lahiri and Teigland's (1987) and Rich and Tracy's (2006) finding of leptokurtic distributions for professional forecasts. The data appear to be non-normal.<sup>14</sup>

### 3.3 Consumer disagreement and forecast convergence

We now investigate the hypotheses of heterogeneity and convergence in consumer inflation expectations, *i.e.*, the degree of disagreement in consumer inflation forecasts, and whether such disagreement tended to narrow over the sample period. We rely on the adaptive nonparametric density estimation and mode testing approach advocated by Zhu (2005). The empirical analysis is based on the normally resampled data, for two reasons: data under different resampling schemes are broadly similar; the normally resampled data have open-ended distributions in line with the original survey questions. The choice is conservative, as the normally sampled data typically have less dispersion and are more concentrated around bin centres.

Figure 12 presents the density estimates for the one-year-ahead consumer inflation forecasts for years 2001 to 2007, computed using critical bandwidths determined by Zhu's (2005) adaptive version of Silverman's (1981,1986) bootstrap multimodality test. The exercise is effectively a statistical test of the

---

<sup>14</sup>Surveys of professional forecasters are expected to have smaller dispersion than consumer surveys. Carlson (1975) used Livingston survey data to test for normality of price expectations. He concluded that the sample distributions were skewed to the right and more peaked than normal distributions.

degree of heterogeneity or disagreement among consumers with respect to their expectations of future inflation. Comparison of density estimates over time would provide information on convergence in inflation forecasts. Examining the “critically” smoothed density estimates allows us to draw the following conclusions.

First, forecast densities were clearly multimodal according to our tests, possibly indicating pronounced expectational heterogeneity. This is in strong contrast with the unimodal distributional assumption underlying the MPC’s inflation fan charts. In fact, consumer forecast densities were trimodal throughout the sample period except for 2005, the year with the lowest median inflation forecast. If one assumes that the spikes around 0% inflation were a survey artifact due to a misunderstanding of the survey questions by some respondents, density estimates would still be bimodal except in 2005. Taking the number of modes in consumer forecast densities as a barometer of monetary policy effectiveness, 2005 would be a year of greater effectiveness. The existence of multiple modes is a clear indication of strong heterogeneity and disagreement among UK consumers with respect to their inflation forecasts, suggesting a significant degree of uncertainty in household expectations. In particular, over the vast majority of our sample a substantial proportion of consumers held views about future inflation which departed significantly from consumers’ central tendencies (*e.g.* modes) and from the MPC’s forecasts.

Second, consumer forecast densities were relatively stable and there was no convergence over time. Forecast densities changed little between 2006 and 2007, and these were little different from the 2003 density estimate (see Figure 12). The three years were characterised by a relatively high median inflation forecasts. Annual CPI inflation was only 1.4% in 2003, taking the year as a whole, but higher during 2006 and 2007 at 2.1%. Compared to other years, forecast densities in 2003, 2006 and 2007 were characterised by a larger probability mass surrounding the mode at the higher inflation rate above 4%. Consequently higher actual inflation was associated with a sizeable shift of forecasts from modes at lower inflation towards the mode at higher inflation. Consumers could have become better informed of the UK inflation outlook, or perceived increased inflation pressures since 2004, or both. There was no tendency in our sample for expectations to converge towards a unique mode. If anything, expectations probably diverged in the last two years of the sample period.

On the other hand, density estimates for 2001, 2002, 2004 and 2005 had similar shapes, with a considerably smaller probability mass surrounding the

third mode above 4%. Twelve-month CPI inflation in the Februarys of 2001, 2002, 2004 and 2005 averaged 1.3%, compared with 2.1% on average in the three other years of our sample (2003, 2006 and 2007). Consumers seemed to understand that inflation pressures increased, and many of them increased their inflation forecasts. There was a clear tendency for the spike around 0%-inflation to diminish over time (except for 2005), with much of the probability mass moving towards the upper part of the distribution. The 2007 forecast density had the smallest spike around 0%-inflation.

### 3.4 Comparison with the MPC's forecasts

The Bank of England sets interest rates to achieve the inflation target set by the UK Government, namely 2% CPI inflation since December 2003 and 2.5% RPIX inflation from May 1997 until that date. The sample period under consideration was characterised by low and stable CPI inflation rates averaging 1.7%, below the target. CPI inflation only rose above 2% since mid-2005. In this section we first reconstruct the probability distributions of the MPC's inflation forecasts. We then assess the MPC's ability to anchor consumers' inflation expectations, by comparing the MPC's forecast densities with density estimates of consumer inflation forecasts over time.

We construct the MPC's fan-chart inflation forecast densities using two-piece normal distributions parameterised using values published by the Bank of England.<sup>15</sup> The dispersion of the MPC's forecast densities measures uncertainty around the central inflation forecasts, rather than disagreement among policymakers. The MPC's forecast densities were unimodal, very peaked and tight, and most of them were almost symmetric (see Figure 13). From 2004 to 2006, the MPC's forecast densities were tighter and more peaked than in other years. Moreover, mean and median forecasts rose over the sample period, but fell significantly from 2006 to 2007.

In Figures 14 and 15, we compare the MPC's inflation forecasts with the critically-smoothed consumer-based estimates of forecast density derived in the previous section. We draw the following conclusions. First, UK consumer forecast densities were far more dispersed and far less peaked than the MPC's density forecasts in all years. Taking the dispersion in consumer forecasts as an approximate measure of forecast uncertainty, then UK consumers were far less sure about one-year-ahead inflation than the MPC, as one would expect.

---

<sup>15</sup>See Appendix II for details.

Second, while the MPC's forecast densities all had a unique mode by assumption, consumer forecast densities were clearly multimodal. In addition, distances between different modes were large, typically well over 2 percentage points, reflecting a high degree of heterogeneity in consumers' expectations formation. Moreover, such differences showed no tendency to diminish over time, i.e. no tendency towards expectational convergence. Although unimodality of the MPC's density forecasts is a reasonable assumption, representing policymakers' consensus views and the extent of their perceived uncertainty surrounding the agreed central inflation projections, it does not appear to be a good assumption for consumer inflation forecasts. High heterogeneity in the form of three modes suggests the possible existence of (at least) three distinct processes by which consumers form expectations, perhaps reflecting differences in age, education, social and professional experience, or other characteristics. Persistent multimodality in consumer forecast densities, during a period of low and stable inflation, demonstrates the difficulties policymakers face in reducing the disagreement and uncertainty that are inherent in inflation expectations. Moreover, persistent differences between consumers' forecast densities and those of the MPC imply that the MPC's forecasts were not fully credible to consumers.

Third, in most years during the sample period, the 90%-confidence bands for the MPC's inflation projections covered less than one third of the probability mass of consumer forecast densities, implying that less than one third of consumers had their one-year-ahead inflation forecasts falling within the MPC's 90%-confidence bands. A wide dispersion of consumer forecasts could also suggest a possible lack of credibility regarding the Bank's inflation target. The extent of the differences between consumers' expectations and the MPC's forecasts was striking. This happened even though summary measures of the inflation forecasts by the MPC and UK consumers - means, medians and modes - were quite close to each other. As a result, simply focusing on these moments ignores substantial differences between the MPC's and consumers' assessments of future inflation risks, and an analysis of the entire probability distribution can be useful. Furthermore, density estimation under unimodal parametric assumptions could not provide information on the type of heterogeneity we have uncovered. In particular, multimodality tests are a useful tool for ascertaining the degree of heterogeneity in density estimates, and consequently in inflation expectations formation.

The evident discrepancies between the MPC's and consumers' expectations could have a number of different interpretations. One possibility could

be that the MPC has superior and more timely information than individual consumers, which is further enhanced by in-depth economic analysis by skilled Bank staff. This could allow the MPC to greatly reduce its uncertainty about future inflation developments, leading to tight and peaked forecast densities. The Bank of England regularly monitors business and household activities and sentiments, but it takes time for this information to be transmitted to and absorbed by the private sector. Having limited resources and less access or ability to process information, consumers might resort to rules of thumb and place greater weight on past inflation developments. Consumers might also not fully grasp the temporary nature of certain shocks to inflation. Disparities in the ability to collect and analyse relevant information could therefore drive a wedge between the MPC's and consumers' forecasts.

A second possibility is that consumers themselves are a heterogeneous collection of individuals with distinct abilities, resources and willingness to obtain and process information concerning economic and price developments. The opportunity to communicate with each other, in order to exchange information and opinions with a view to improving their forecasts, is small compared with the tight structure of the MPC and its forecasting process. Access to economic and financial news and analyses in the public domain helps, but the extent and depth of such coverage may well be limited. By contrast, the MPC consists of a small group of experts supported by highly trained and experienced staff. MPC members regularly meet to discuss economic issues and exchange information, and this process is likely to help bridge or narrow differences in views. The MPC clearly enjoys a significant advantage here when formulating its forecasts.

A third possibility is that, unlike UK consumers, the MPC has monetary policy instruments at its disposal including Bank Rate, and it can take preemptive policy actions if inflation is expected to drift away from target, with a direct impact on the projected path of inflation. The MPC has a better knowledge of its actions than any other agents, which affords the MPC a policy advantage over consumers. Enjoying such information, internal communication and policy advantages, the MPC may be expected to be more certain about future UK inflation developments.

## 4 Conclusions

Inflation expectations play a crucial role in actual inflation dynamics, and central banks carefully monitor the public’s inflation expectations. This paper studies inflation expectations of UK consumers from the 2001-2007 February consumer surveys conducted by the GfK NOP on behalf of the Bank of England. We first study the evolution of and characterise the probability distribution of UK consumer inflation expectations based on the GfK NOP survey data via standard and robust moments, histograms, Tukey boxplots and Q-Q plots. This analysis provides valuable but limited and sometimes contradictory information. We therefore improve on this analysis by using nonparametric kernel methods to obtain density estimates and to identify the number of modes in the distribution of UK consumer forecasts. The number of modes and the distance between modes are proposed as new measures of disagreement in consumer forecasts. Finally, we compare the MPC’s inflation forecast densities with UK consumer inflation forecast densities derived from the Inflation Report fan charts to assess the “well-anchoredness” of UK inflation expectations.

We find it informative to characterise UK consumers’ inflation forecasts beyond the first moments, which describe central tendencies and serve as policy and communication focal points. Means and medians evolved in a similar fashion for both consumers’ and the MPC’s forecasts, masking the fact that the two forecast distributions behaved very differently in many other aspects. Disagreement in consumer inflation forecasts, a proxy for consumer forecast uncertainty that we measure using standard deviation and interquartile range, is far greater than uncertainty in the MPC’s forecasts. Estimates of standard and robust skewness suggest that consumers’ and the MPC’s forecasts were skewed in opposite directions in most years in the sample. Under the assumption of unimodal forecast densities, this implies that consumers (as a whole) and the Bank had opposite views on whether inflation risks were more on the upside or downside. In particular, the MPC’s forecasts were skewed to lower inflation rates while consumer forecasts were skewed to higher rates.

Moreover, quartile- and octile-based robust skewness measures for UK consumers’ expectations diverged in some years, suggesting more complex intra-distributional dynamics. Measures of excess kurtosis showed that the MPC’s forecasts were close to standard normal, while consumer forecasts were platykurtic, *i.e.*, less peaked in the centre of the distribution and with



thinner tails. However, measures of skewness and kurtosis, whether robust or standard measures, are not valid in cases of multimodality. The meaning of such distributional statistics is lost once the conventional unimodality assumption is proved to be invalid.

We propose a new perspective on disagreement and uncertainty in inflation expectations by examining the complete probability distributions of consumer inflation forecasts. Specifically, we test for possible disagreement in consumer inflation forecasts using Zhu's (2005) adaptive bootstrap multimodality test. Our tests suggest that consumer forecast densities were multimodal in all years considered, with the central modes staying fairly close to the MPC's central forecasts. Using nonparametric methods, we identified the existence of an additional mode in the consumer forecast distribution at relatively high rates of inflation in some years. This could reflect persistently high inflation expectations by a substantial proportion of survey respondents. The increase in the probability mass surrounding the higher mode towards the end of the sample period may be associated with a rise in energy prices. We also identified one other additional mode at zero inflation throughout the sample period. This could result from a possible misunderstanding of the original survey question and available options for replies, with some consumers confusing "no change in prices" with "no change in inflation", *i.e.* confusing a statement about the level with that about the first difference. Rephrasing this question might eliminate or diminish the zero-inflation mode.

We examine the ability of the MPC to anchor private sector inflation expectations, by comparing the MPC's inflation forecast densities with density estimates of consumer inflation forecasts over time. Our analysis indicates that there is substantial disagreement not only among the consumers but also between the MPC and UK consumers. The latter group has far more diverse views about one-year-ahead inflation. The apparent discrepancies in the MPC's and consumers' views might be explained by the MPC's information advantage and analytical ability, further enhanced by a policy advantage.

## 5 Appendices

### 5.1 Appendix I: Definition of robust moments

Robust measures of higher moments are particularly useful in cases of fat-tailed distributions with a large number of outliers. The inter-quartile range (IQR) of a probability distribution is a robust measure of disagreement and uncertainty. It is defined as

$$IQR = Q_{0.75} - Q_{0.25} \quad (2)$$

where  $Q_{0.75}$  and  $Q_{0.25}$  are the 75th and 25th percentiles, respectively. The standard skewness measure is defined as the normalised third central moment

$$skew = E[(x - \mu)^3 / \sigma^3] \quad (3)$$

where  $\mu$  and  $\sigma$  are the mean and standard deviation of  $x$ .  $E(\bullet)$  is the expectations operator. Hinkley (1975) suggested a class of robust skewness measures of the following form

$$skew_R^p = \frac{(Q_{1-p} - Q_{0.5}) - (Q_{0.5} - Q_p)}{Q_{1-p} - Q_p} \quad (4)$$

where  $Q_p$  is the  $p$ -th quantile,  $p \in (0, 1)$ , and  $Q_{0.5}$  is the median. For  $p = 0.25$ , the *quartile*-based skewness is simply<sup>16</sup>

$$skew_R^{0.25} = \frac{(Q_{0.75} - Q_{0.5}) - (Q_{0.5} - Q_{0.25})}{Q_{0.75} - Q_{0.25}} \quad (5)$$

For  $p = 1/8$ , the *octile*-based skewness is defined likewise.

The standard excess kurtosis measure is defined as the normalised fourth central moment minus 3, the value of kurtosis for the standard normal distribution:

$$kurt = E[(x - \mu)^4 / \sigma^4] - 3 \quad (6)$$

Moors (1988) proposed a robust, *octile*-based measure of excess kurtosis<sup>17</sup>

$$kurt_R^{1/8} = \frac{(Q_{7/8} - Q_{5/8}) - (Q_{3/8} - Q_{1/8})}{Q_{6/8} - Q_{2/8}} - 1.23 \quad (7)$$

The value of 1.23 again corresponds to the Moors coefficient of kurtosis for the standard normal distribution.

<sup>16</sup>See also Bowley (1920) and Kim and White (2003).

<sup>17</sup>See also Kim and White (2003).

## 5.2 Appendix II: Deriving probability distributions from fan charts

The Bank of England publishes the mean, median and mode underlying the MPC's inflation forecast fan charts provided in the quarterly *Inflation Reports*, together with measures of uncertainty and skewness. According to Britton, Fisher and Whitley (1998), the fan charts are based on a two-piece normal distribution<sup>18</sup>

$$f(x) = \begin{cases} \frac{A}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-m)^2(1+\gamma)}{2\sigma^2}\right) & x \leq m \\ \frac{A}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-m)^2(1-\gamma)}{2\sigma^2}\right) & x > m \end{cases} \quad (8)$$

where  $m$  is the mode, and

$$A = 2 \left[ (1 + \gamma)^{-1/2} + (1 - \gamma)^{-1/2} \right]^{-1} \quad (9)$$

By definition, the distribution is unimodal. The uncertainty parameter  $\sigma$  equals the standard deviation only if the distribution is symmetric. The skewness parameter is defined as the mean minus the mode. The two-piece normal distribution can be reparametrised as<sup>19</sup>

$$f(x) = \begin{cases} \sqrt{\frac{2}{\pi}} \left( \frac{1}{\sigma_1 + \sigma_2} \right) \exp\left(-\frac{(x-m)^2}{2\sigma_1^2}\right) & x \leq m \\ \sqrt{\frac{2}{\pi}} \left( \frac{1}{\sigma_1 + \sigma_2} \right) \exp\left(-\frac{(x-m)^2}{2\sigma_2^2}\right) & x > m \end{cases} \quad (10)$$

where  $\mu = m + \sqrt{2/\pi}(\sigma_2 - \sigma_1)$ , and  $\sigma_{1,2} = (1 \pm \gamma)^{-1/2} \sigma$ . The parameter  $\gamma$  in equations (8) and (9) can be written in terms of the parameters  $(\mu, m, \sigma)$

$$\gamma^2 = 1 - 4 \left( \frac{\sqrt{1 + \pi s^2} - 1}{\pi s^2} \right)^2 \quad (11)$$

where  $s = (\mu - m)/\sigma$ , and  $\gamma$  takes the sign of  $s$ . We can therefore reconstruct the fan-chart forecast distributions using the values for the mean, median, mode and uncertainty published by the Bank of England. Note that when skewness is zero, *i.e.*, when the risks to the MPC's central projections are

<sup>18</sup>Two-piece normal distributions are discussed in detail in John (1982) and Johnson, Kotz and Balakrishnan (1994).

<sup>19</sup>See for example, Garvin and McClean (1997) and Wallis (2004).

balanced,  $m = \mu$  and  $\sigma_1 = \sigma_2$ , and the two-piece distribution becomes standard normal.

The second, third and fourth central moments around the mean  $\mu$  are, respectively

$$E(X - \mu)^2 = \left(1 - \frac{2}{\pi}\right) (\sigma_2 - \sigma_1)^2 + \sigma_2 \sigma_1 \quad (12)$$

$$E(X - \mu)^3 = \sqrt{\frac{2}{\pi}} (\sigma_2 - \sigma_1) \left\{ \left(\frac{4}{\pi} - 1\right) (\sigma_2 - \sigma_1)^2 + \sigma_2 \sigma_1 \right\} \quad (13)$$

$$\begin{aligned} E(X - \mu)^4 &= E(X - m)^4 + 4E(X - m)^3(m - \mu) + 6E(X - m)^2(m - \mu)^2 \\ &\quad + 4E(X - m)(m - \mu)^3 + (m - \mu)^4 \end{aligned} \quad (14)$$

The  $r$ -th central moment around the mode  $m$  is

$$E(X - m)^r = \frac{(\sigma_2 \sqrt{2})^{r+1} - (-\sigma_1 \sqrt{2})^{r+1}}{\sigma_1 + \sigma_2} \frac{\Gamma(r/2 + 1/2)}{\sqrt{2\pi}} \quad (15)$$

where  $\Gamma(\bullet)$  is the gamma function.

To obtain the robust moments for the MPC's fan-chart inflation forecast distributions, we calculate the octiles of the two-piece normal distribution by inverting the expression for the cumulative probability distribution function of the two-piece normal distribution<sup>20</sup>

$$F(x) = \begin{cases} \frac{2\sigma_1}{\sigma_1 + \sigma_2} \Phi\left(\frac{x-m}{\sigma_2}\right) & x \leq m \\ \sigma_1 - \sigma_2 + \frac{2\sigma_2}{\sigma_1 + \sigma_2} \Phi\left(\frac{x-m}{\sigma_2}\right) & x > m \end{cases} \quad (16)$$

where  $\Phi(\bullet)$  is the standard normal cumulative distribution function.

### 5.3 Appendix III: Adaptive kernel density estimation and mode testing

We use an adaptive version of kernel density estimation and bootstrap multimodality test procedures as proposed in Zhu (2005). The original fixed-bandwidth Silverman (1981, 1983) bootstrap multimodality test can be sensitive to spurious noise in the tails. Without properly controlling the sensitivity to local data density, multimodality tests are error-prone. It is known

---

<sup>20</sup>See also Kimber (1985).

that fixed-bandwidth kernel methods tend to produce spurious modes in the tails while masking essential details in the central part of size distributions. In the variable bandwidth case, a local bandwidth parameter  $\gamma_{ni}$  adapts the degree of smoothing to local data density. Adaptive kernel density estimation makes it possible to purge spurious noise in the tails while preserving important characteristics in the fat part of the distribution. Bandwidth tests based on the adaptive method are more accurate.

Suppose that in each period  $t$  we have a set of univariate observations  $Y_1, \dots, Y_n$ , the adaptive kernel density estimate can be computed in the following two-step procedure, adapted from Breiman *et al* (1977):

1. Obtain a pilot kernel estimate  $\tilde{f}(y)$  of the inflation forecast density  $f(y)$  in such a way that  $\tilde{f}(Y_i) > 0$ , for all  $i$ :

$$\tilde{f}(y) = (nh_0)^{-1} \sum_{i=1}^n K\left(\frac{y - Y_i}{h_0}\right) \quad (17)$$

where  $h_0$  is the initial bandwidth and  $K(\cdot)$  is the Gaussian kernel.

2. Compute the adaptive kernel estimate  $\hat{f}(y)$  as follows:

$$\hat{f}(y) = n^{-1} \sum_{i=1}^n h_i^{-1} K\left(\frac{y - Y_i}{h_i}\right) \quad (18)$$

where  $h_i = h\gamma_{ni}$ .  $h$  is the overall bandwidth or smoothing parameter and  $\gamma_{ni}$  is the local bandwidth parameter which varies according to the local data density:

$$\gamma_{ni} = \left(\frac{\tilde{f}(Y_i)}{\bar{f}}\right)^{-\alpha} \quad (19)$$

where  $\bar{f}$  is the geometric mean of the pilot estimates  $\tilde{f}(Y_i)$  and  $\alpha \in [0, 1]$  is a sensitivity parameter.

Define the number of modes of density  $f$  as:

$$M(f) = \#\{y \in \mathfrak{R}_+ : f'(y) = 0 \text{ and } f''(y) < 0\} \quad (20)$$

For  $m = 1, \dots, M$ , we test the null hypothesis that the underlying density  $f$  has  $m$  modes ( $H_0 : M(f) \leq m$ ), against the alternative that  $f$  has more

than  $m$  modes ( $H_1 : M(f) > m$ ). The bootstrap multimodality test is based on the notions of critical smoothing and critical bandwidth. Define the  $m$ -th critical bandwidth  $\hat{h}_{m,crit}$  by

$$\hat{h}_{m,crit} = \inf \left\{ h : M(\hat{f}_h) \leq m \right\} \quad (21)$$

where, for any density estimate  $\hat{f}_h$ ,  $\hat{h}_{m,crit}$  is the smallest overall bandwidth such that  $\hat{f}_h$  has at most  $m$  modes. Consequently,  $\hat{f}_{\hat{h}_{m,crit}}$  is known as the  $m$ -th critical density. Following Silverman (1986) and Efron and Tibshirani (1993), we assess the significance of  $\hat{h}_{m,crit}$  estimated from data against its bootstrap distribution.<sup>21</sup>

## References

- [1] **Ang, A, Bekaert, G and Wei, M (2006)**, ‘Do macro variables, asset markets or surveys forecast inflation better’, *Finance and Economics Discussion Series 2006-15*, Federal Reserve Board.
- [2] **Bank of England (2002)**, ‘The MPC’s fan charts’, Box in May *Inflation Report*, p. 48-49.
- [3] **Bank of England (2004)**, ‘The new inflation target’, Box in February *Inflation Report*, p. 36.
- [4] **Bank of England (2007)**, *Inflation Report*.
- [5] **Boero, G, Smith, J and Wallis, K (2008)**, ‘Uncertainty and disagreement in economic prediction: the Bank of England Survey of external forecasters’, *Economic Journal*, Vol. 118(530), p. 1107-27.
- [6] **Bowley, A (1920)**, *Elements of statistics*, New York, Charles Scribner’s Sons.
- [7] **Britton, E, Fisher, P and Whitley, J (1998)**, ‘The *Inflation Report* projections: understanding the fan chart’, *Bank of England Quarterly Bulletin*, February, p. 30-37.

---

<sup>21</sup>See Zhu (2005) for details of the bootstrap resampling procedure.

- [8] **Breiman, L, Meisel, W and Purcell, E (1977)**, ‘Variable kernel estimates of multivariate densities’, *Technometrics*, Vol. 19, p. 135-44.
- [9] **Carlson, J (1975)**, ‘Are Price Expectations Normally Distributed?’, *Journal of the American Statistical Association*, Vol. 70, p. 749-54.
- [10] **Clements, M (2004)**, ‘Evaluating the Bank of England density forecasts of inflation’, *Economic Journal*, Vol. 114, p. 844-66.
- [11] **Croushore, D (2006)**, ‘An evaluation of inflation forecasts from surveys using real-time data’, *Federal Reserve Bank of Philadelphia Working Paper*, No. 06-19.
- [12] **D’Amico, S and Orphanides, A (2006)**, ‘Uncertainty and disagreement in economic forecasting’, *mimeo*.
- [13] **DeLong, J (1997)**, ‘America’s Peacetime Inflation: The 1970s’, in *Reducing Inflation: Motivation and Strategy*, edited by Romer, C and Romer, D, The University of Chicago Press, Chicago.
- [14] **Diebold, F, Tay, A and Wallis, K (1998)**, ‘Evaluating density forecasts of inflation: The survey of professional forecasters’, *mimeo*.
- [15] **Dotsey, M and DeVaro, L (1995)**, ‘Was the Disinflation of the Early 1980s Anticipated?’, *Federal Reserve Bank of Richmond Economic Quarterly*, Vol. 81(4), p. 41-59.
- [16] **Dowd, K (2004)**, ‘The inflation fan charts: an evaluation’, *Greek Economic Review*, Vol. 23, No. 1, p. 99-111.
- [17] **Dowd, K (2007)**, ‘Too good to be true? The (In)credibility of the UK inflation fan charts’, *Journal of Macroeconomics*, Vol. 29, p. 91-102.
- [18] **Elder, R, Kapetanios, G, Taylor, T and Yates, T (2005)**, ‘Assessing the MPC’s fan charts’, *Bank of England Quarterly Bulletin*, Autumn, p. 326-48.
- [19] **Ellis, C (2005)**, ‘Public attitudes to inflation’, *Bank of England Quarterly Bulletin*, Summer, p. 194-201.
- [20] **Ellis, C (2006)**, ‘Public attitudes to inflation’, *Bank of England Quarterly Bulletin*, Summer, p. 181-89.

- [21] **Efron, B and Tibshirani (1993)**, *An introduction to the bootstrap*, Chapman.
- [22] **Figlewski, S and Wachtel, P (1981)**, ‘The formation of inflationary expectations’, *The Review of Economics and Statistics*, Vol. 63, p. 1-10.
- [23] **Forsells, and Kenny (2002)**, ‘The rationality of consumers’ inflation expectations: survey-based evidence for the euro area’, ECB Working Paper No. 163.
- [24] **Garvin, J and McClean, S (1997)**, ‘Convolution and sampling theory of the binormal distribution as a prerequisite to its application in statistical process control’, *The Statistician*, Vol. 46, No. 1, pages 33-47.
- [25] **Geldman, A (2007)**, ‘Headline inflation still holds the key for pay setters’, *IRS Employment Review*, Issue 864, p. 16-20.
- [26] **Giordani, P and Söderlind, P (2003)**, ‘Inflation forecast uncertainty’, *European Economic Review*, Vol. 47, p. 1037-1059.
- [27] **Gramlich, E (1983)**, ‘Models of Inflation Expectations Formation: A Comparison of Household and Economist Forecasts’, *Journal of Money, Credit and Banking*, Vol. 15(2), p. 155-73.
- [28] **Hinkley, D (1975)**, ‘On power transformations to symmetry’, *Biometrika*, Vol. 62, p. 101-11.
- [29] **John, S. (1982)**, ‘The three-parameter two-piece normal family of distributions and its fitting’, *Communications in Statistics – Theory and Methods*, 11, p. 879-85.
- [30] **Johnson, N.L., Kotz, S. and Balakrishnan, N. (1994)**, *Continuous Univariate Distributions*, 2nd ed., vol 1, New York, Wiley.
- [31] **Keane, M and Runkle, D (1990)**, ‘Testing the Rationality of Price Forecasts: New Evidence from Panel Data’, *American Economic Review*, Vol. 80(4), p. 714-35.
- [32] **Kim, T-H, and White, H (2003)**, ‘Estimation, inference, and specification testing for possibly misspecified quantile regression’, in *Maximum Likelihood Estimation of Misspecified Models: Twenty Years Later*, edited by Fomby, T and Hill, C, Elsevier, New York.



- [33] **Kimber, A (1985)**, ‘Methods for the two-piece normal distribution’, *Communications in Statistics – Theory and Methods*, 14(1), p. 235-45.
- [34] **Lahiri, K, and Teigland, C (1987)**, ‘On the Normality of Probability Distributions of Inflation and GNP Forecasts’, *International Journal of Forecasting*, Vol. 3, p. 269-79.
- [35] **Leduc, S, Sill, K and Stark, T (2007)**, ‘Self-fulfilling expectations and the inflation of the 1970s: Evidence from the Livingston survey’, *Journal of Monetary Economics*, Vol. 54, p. 433-59.
- [36] **Mankiw, G, Reis, R and Wolfers, J (2003)**, ‘Disagreement about inflation expectations’, *NBER Working Paper 9796*.
- [37] **Moors, J (1988)**, ‘A quantile alternative for kurtosis’, *The statistician*, Vol. 37, p. 25-32.
- [38] **Nickell, S (2003)**, ‘Two current monetary policy issues’, *Bank of England Quarterly Bulletin*, Winter, p. 504-17.
- [39] **Pesaran, H (1985)**, ‘Formation of Inflation Expectations in British Manufacturing Industries’, *Economic Journal*, Vol. 95(380), p. 948-75.
- [40] **Rich, R (1989)**, ‘Testing the Rationality of Inflation Forecasts from Survey Data: Another Look at the SRC Expected Price Change Data’, *The Review of Economics and Statistics*, Vol. 71, No. 4, p. 682-86.
- [41] **Rich, R and Tracy, J (2006)**, ‘The Relationship between Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts’, *Federal Reserve Bank of New York Staff Reports*, no. 253.
- [42] **Silverman, B (1981)**, ‘Using Kernel Density Estimates to Investigate Multimodality’, *Journal of the Royal Statistical Society, Series B.*, Vol. 43:1, p. 97-99.
- [43] **Silverman, B (1983)**, ‘Some Properties of a Test for Multimodality Based on Kernel Density Estimates’, in J. F. C. Kingman and G. E. H. Reuter, eds., *Probability, Statistics and Analysis*, Cambridge: Cambridge University Press, p. 248-259.

- [44] **Silverman, B (1986)**, *Density Estimation for Statistics and Data Analysis*, Chapman.
- [45] **Wallis, K (2003)**, ‘Chi-squared tests of interval and density forecasts, and the Bank of England’s fan charts’, *International Journal of Forecasting*, Vol. 19, p. 165-75.
- [46] **Wallis, K (2004)**, ‘An assessment of Bank of England and National Institute inflation forecast uncertainties’, *National Institute Economic Review*, Vol. 189, p. 64-71.
- [47] **Zarnowitz, V (1985)**, ‘Rational Expectations and Macroeconomic Forecasts’, *Journal of Business and Economic Statistics*, Vol. 3, p. 293-311.
- [48] **Zhu, F (2005)**, ‘A nonparametric analysis of the shape dynamics of the US personal income distribution: 1962-2000’, *Bank for International Settlements Working Paper* No. 184.

## 6 Figures




Figure 1: UK RPIX and CPI inflation.<sup>22</sup>

---

<sup>22</sup>For CPI inflation, data before 1998 are ONS estimates




Figure 2: CPI goods and services inflation.




Figure 3: Energy prices and imported finished manufactures prices.<sup>23</sup>

---

<sup>23</sup>Energy prices include water and housing; imported finished manufactures have not been adjusted for missing trader intra-community (MTIC) fraud.




Figure 4: Mean and median for consumers' and the MPC's inflation forecasts.

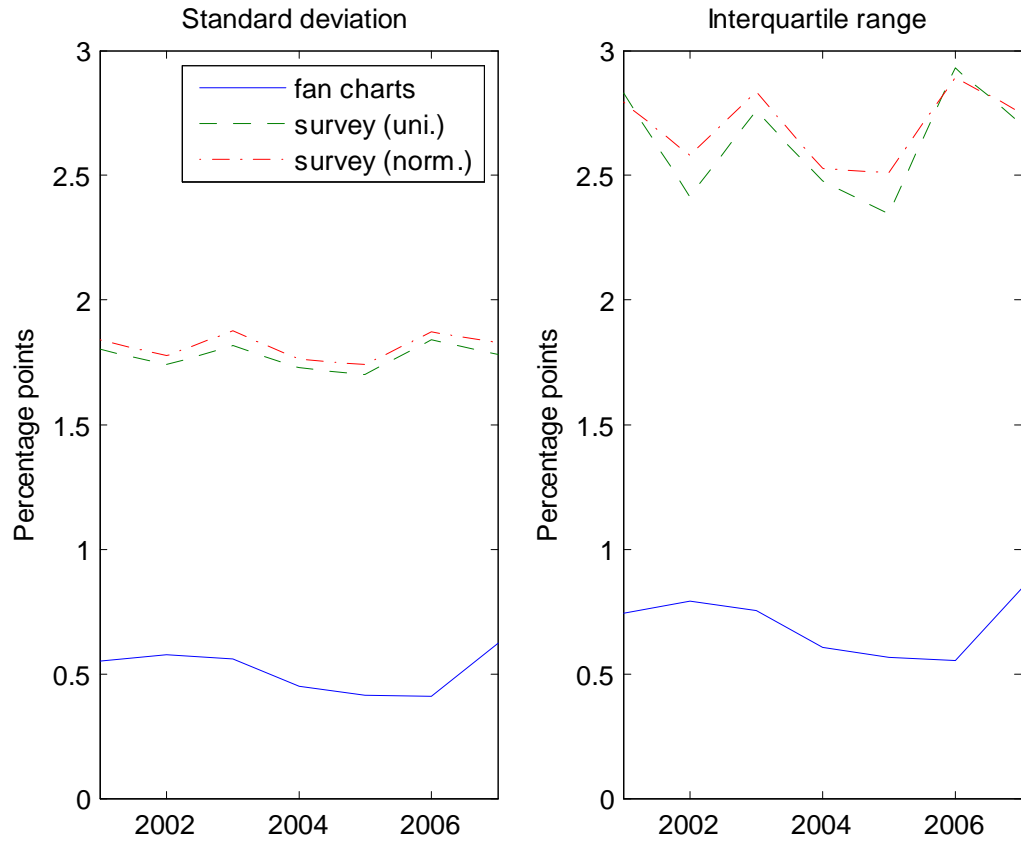


Figure 5: Dispersion measures for consumers' and the MPC's inflation forecasts.

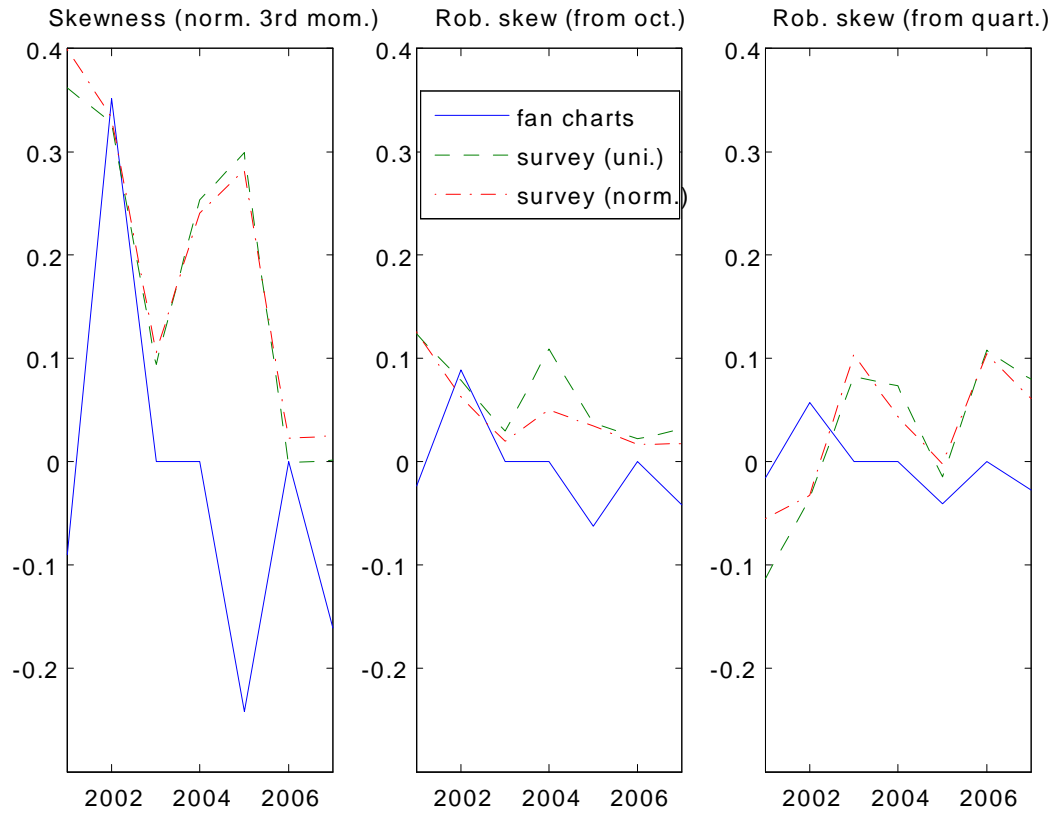


Figure 6: Skewness measures for consumers' and the MPC's inflation forecasts.



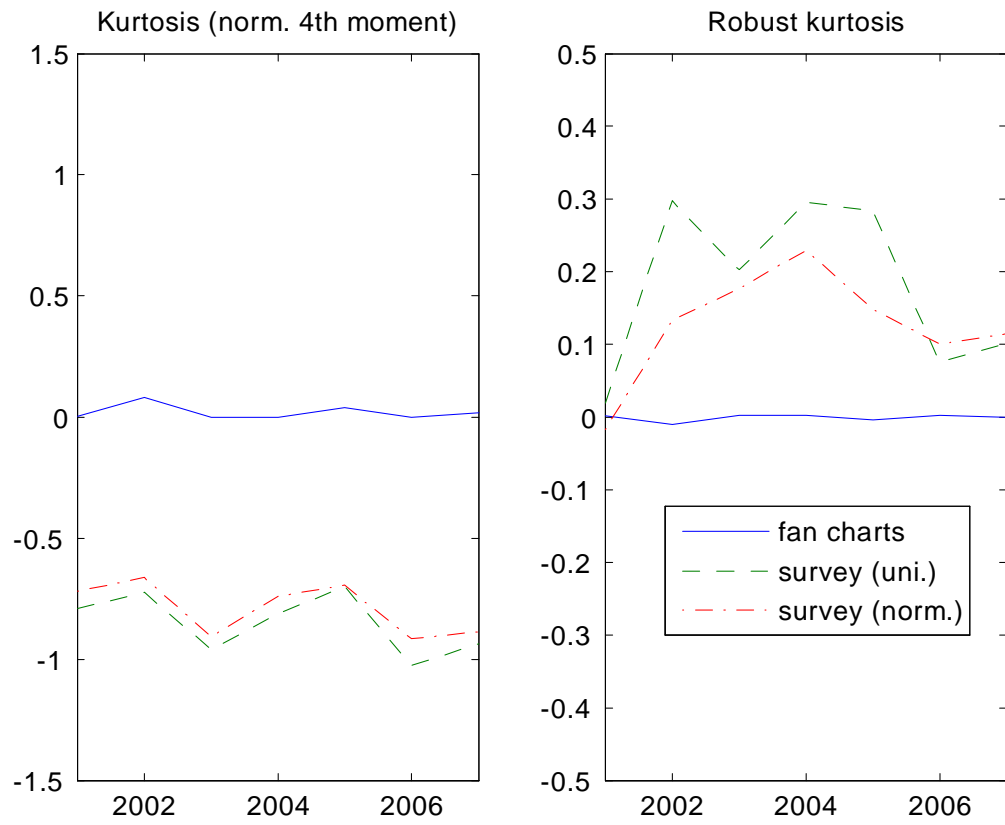


Figure 7: Excess kurtosis measures for consumers' and the MPC's inflation forecasts.

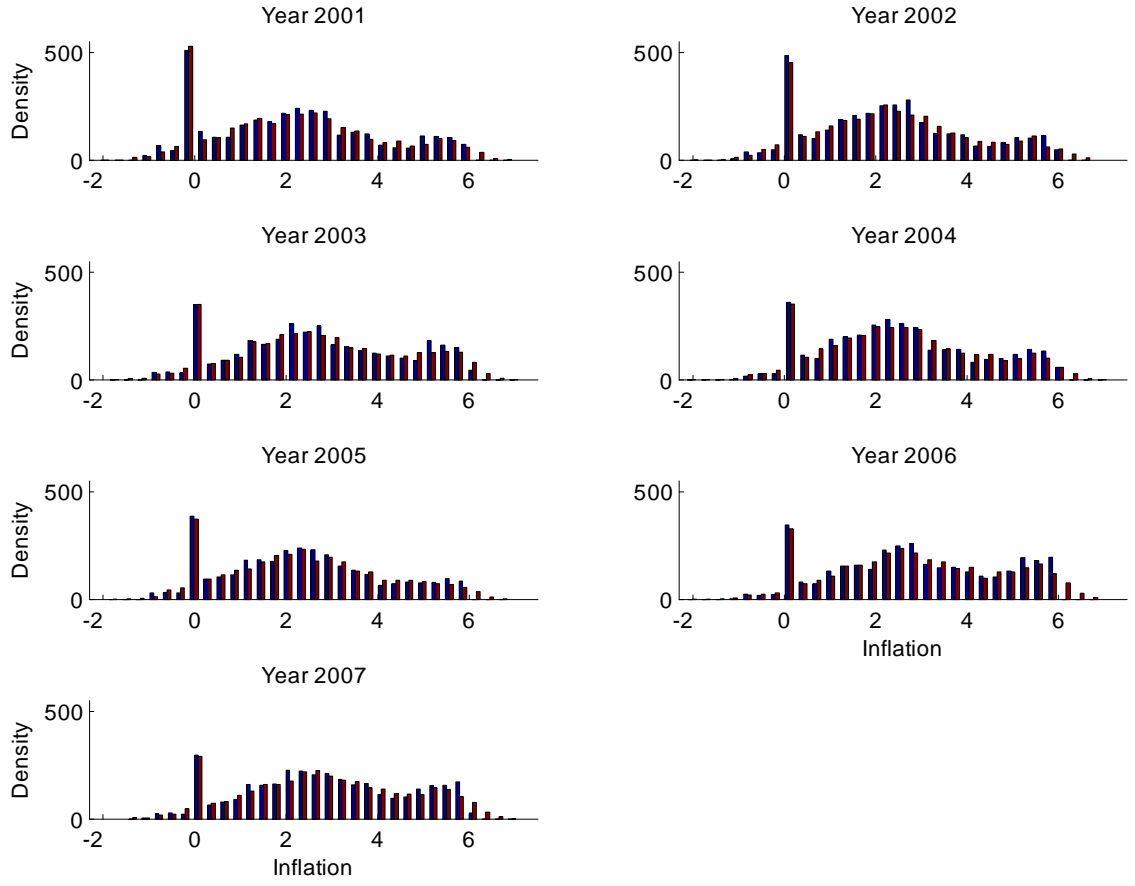


Figure 8: Histograms for consumer inflation forecasts 2001-2007.

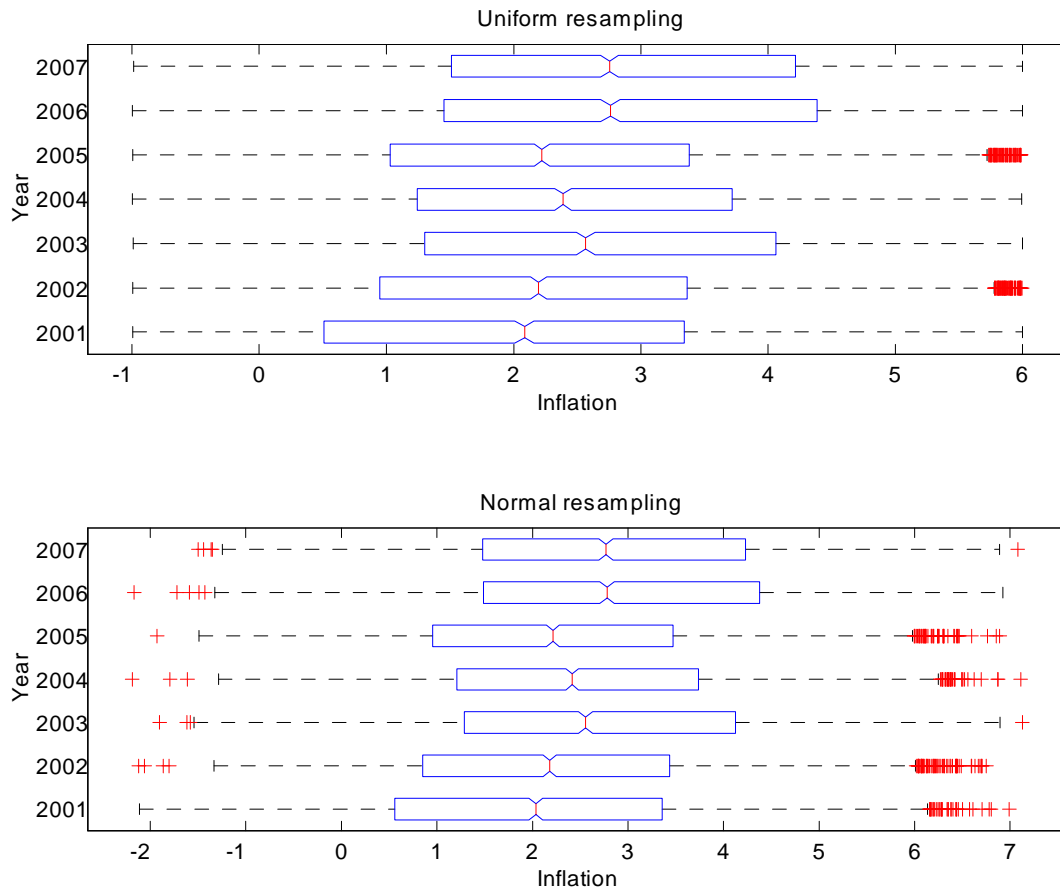


Figure 9: Tukey boxplots for resampled consumer inflation forecasts 2001-2007.

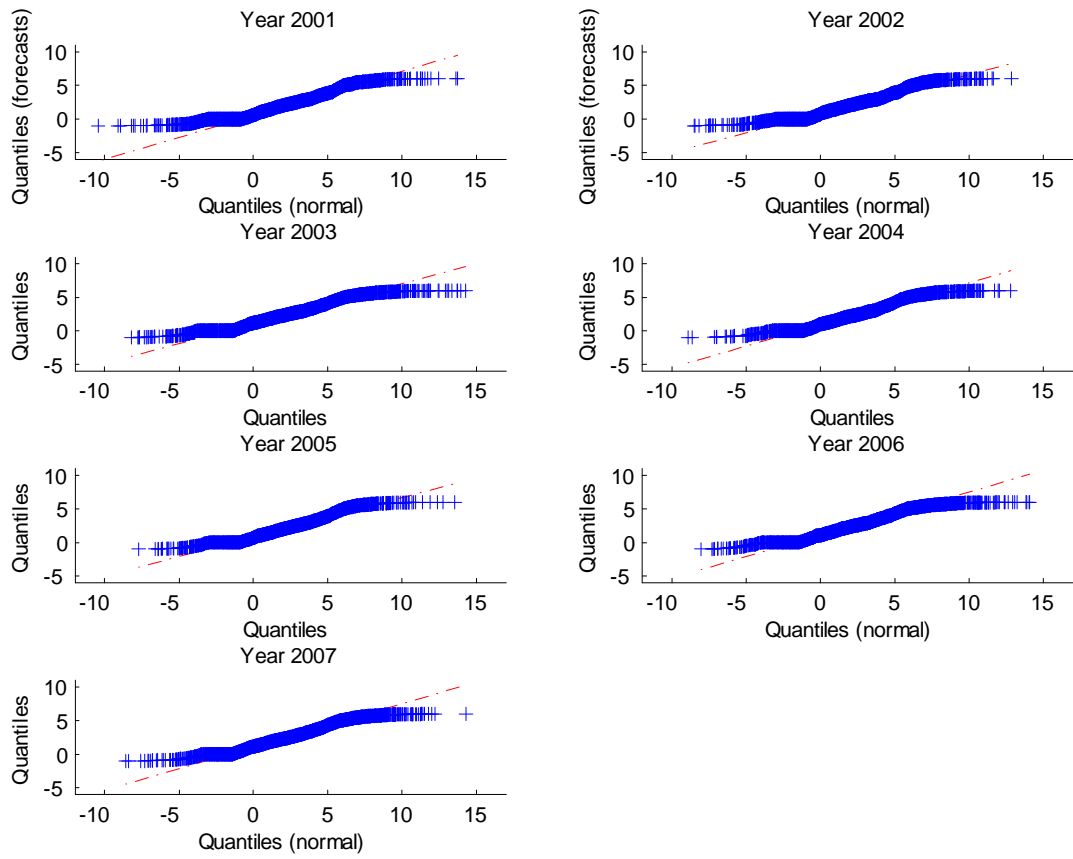


Figure 10: Q-Q plots for uniformly resampled consumer inflation forecasts 2001-2007.




Figure 11: Q-Q plots for normally resampled consumer inflation forecasts 2001-2007.

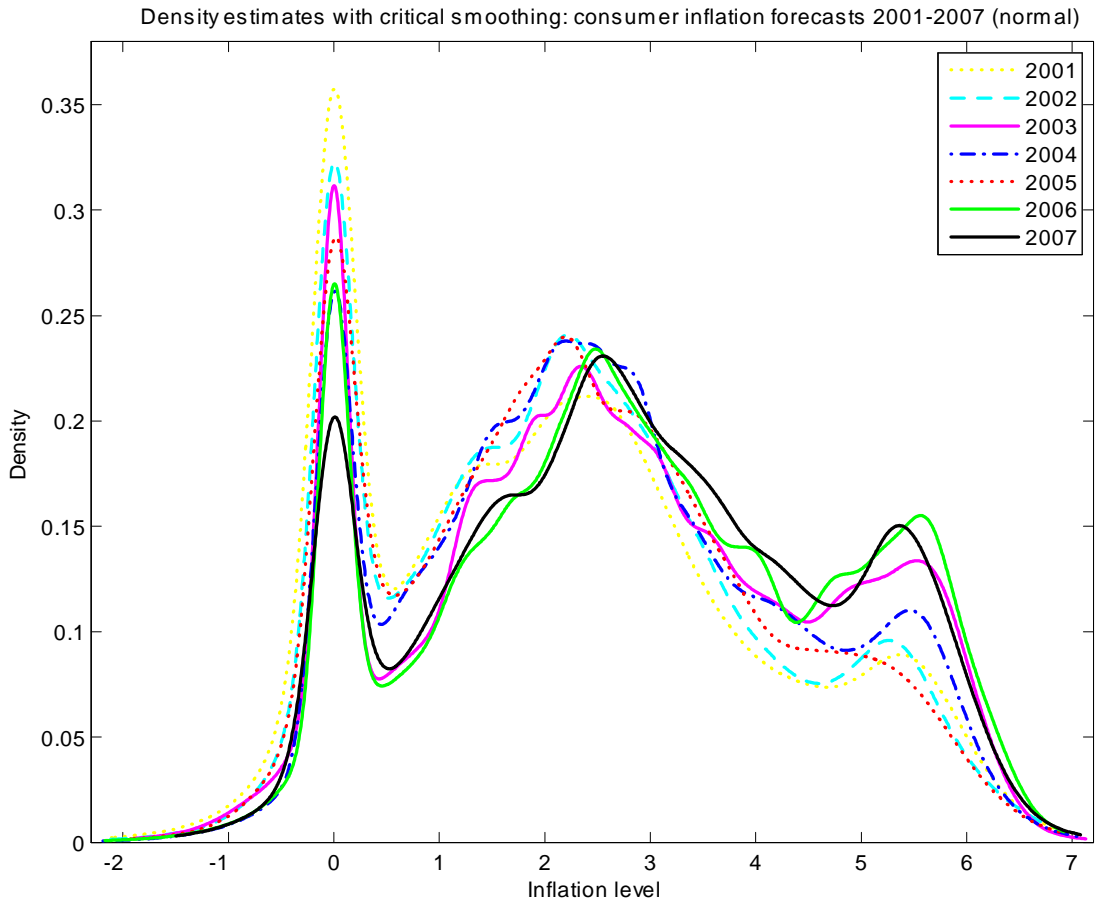


Figure 12: Density estimates for normally resampled consumer inflation forecasts 2001-2007.




Figure 13: Density estimates based on the MPC's inflation forecasts 2001-2007.

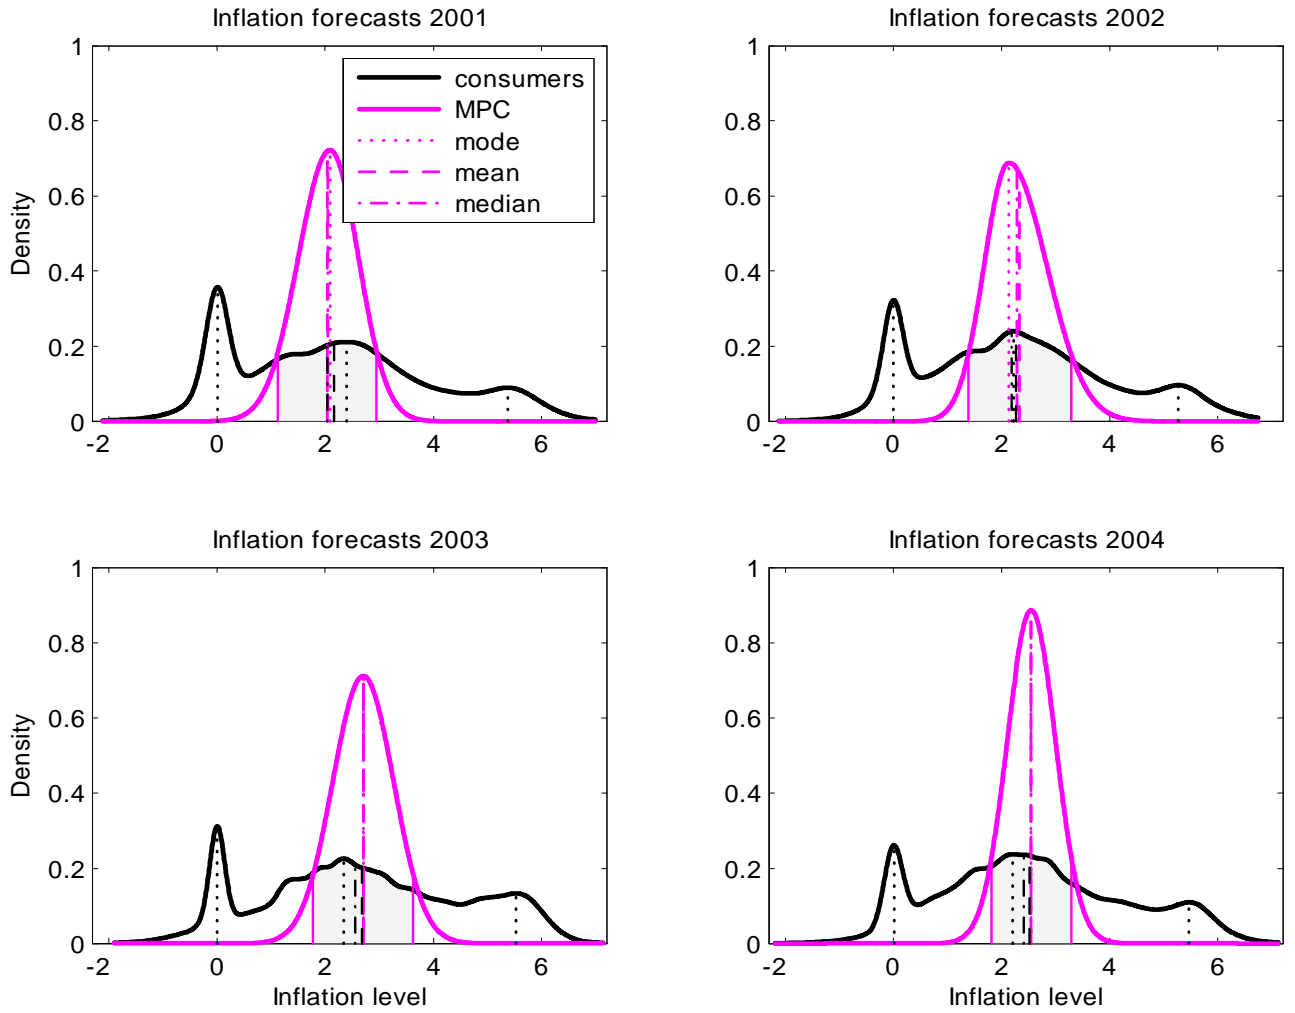


Figure 14: Densities estimates for UK consumers' and the MPC's inflation forecasts, 2001-2004.






Figure 15: Densities estimates for UK consumers' and the MPC's inflation forecasts, 2005-2007.