

Central Bank Density Forecasts: Do Higher-order Moments Matter?*

Jonathan Benchimol,[†] Ryan Rholes,[‡] and Tatevik Sekhposyan[§]

Preliminary Draft - Please do not cite or circulate.

November 16, 2023

Abstract

This paper considers the Bank of England's density forecasts and its revisions to quantify the effects of information flow on the financial markets and survey forecasters. Central banks increasingly rely on published forecasts to communicate their economic outlook to market participants. Point forecasts and their revisions have been shown to move financial markets. However, the effects of the higher-order moments have not been investigated thoroughly, primarily due to data limitations. The Bank of England, on the other hand, has been publishing information on its density forecasts since the late 1990s, making it useful for our analysis. Using daily information on the financial markets, we find that the updates of higher moments are more important in moving financial markets than the revisions to the first central moment of the density forecasts, making them relevant for monetary policy communication. Information about output matters more than information about inflation, and the effect of information is state-contingent. Finally, we see that the consensus forecast and disagreement among professional forecasters strongly correlate with updates in higher-order forecast moments.

Keywords: Bank of England, density forecasts, forecast revisions, information channel, asset prices, uncertainty, communication, monetary policy.

JEL codes: C13, C32, C53.

*We thank Luba Petersen, seminar participants at the Texas A&M University, Bank of England, Federal Reserve Bank of San Francisco, and participants of the 21st IWH-CIREQ-GW Macroeconomic Workshop on "Forecasting and Uncertainty" and Computational and Financial Econometrics Conference for valuable comments and discussions.

[†]Research Department, Bank of Israel, Jerusalem, Israel.

[‡]Department of Economics, University of Oxford, Oxford, UK. Corresponding author. Email: ryan.rholes@economics.ox.ac.uk

[§]Department of Economics, Texas A&M University, College Station, USA.

1 Introduction

Central banking has undergone a decades-long transparency revolution that brought about an increased reliance on communication to manage the expectations of economic agents to achieve policy objectives (Blinder et al., 2008). Communication is a particularly important tool whenever policy rates are constrained at the effective lower bound, which has been the case in the past decade across the world. Arguably, the first link in the context of monetary policy transmission is the effect of the communications on the market interest rates which are relevant for consumption and investment decisions in the economy — a topic we study in this paper. More specifically, we characterize the evolution of economic outlook across the monetary policy cycle and quantify the effects on the market interest rates and survey forecasts.

A primary component of a central bank’s communication strategy is its projections of the key macroeconomic variables, including inflation and output. These summarize the conditions to which a central bank is reacting through its policy instruments. Many central banks publish density forecasts alongside their point predictions, conveying the central bank’s uncertainty and outlook on the balance of risks. The Bank of England (BoE) was the first to publish ‘fan charts’ of its macroeconomic projections in 1998, which is now standard practice across the board.

There is a wide literature on monetary policy communication, mainly focusing on the publication of point forecasts. The main conclusion of this literature is that communication facilitates policy and acts as a coordination device for expectation formation. Swanson (2006), Blinder et al. (2008), Hubert (2014, 2015) support this finding with empirical analyses, while Kryvtsov and Petersen (2021) and Ahrens et al. (2023) provide experimental support. There is further evidence that central bank communication affects the market interest rates (Gürkaynak et al., 2005; Andrade and Ferroni, 2021), making it relevant for policymaking. The communication of economic outlook has also proven to be important under the information channel of monetary policy discussed in seminal works of Nakamura and Steinsson (2018) as well as Miranda-Agrippino and Ricco (2021), Hoesch et al. (2023), among others.

Though the literature on central bank communication is rather extensive, the use of density forecasts and the information they contain has been rather scant. By observing the various central bank behaviors, it also seems that there is no consensus in policy circles on what the communication of uncertainty delivers. As Petersen and Rholes (2022) highlight, with the COVID-19 crisis the central banks around the world have moved in two different directions — for example, both the Federal Reserve and the BoE abandoned predictive densities in exchange for scenario analysis at the onset of the pandemic. On the other hand, in the April 2020 Monetary Policy Report, the Bank of Canada rid itself of point predictions and

instead published only ranges. Some central banks, such as the European Central Bank, made no change to their communication strategy. The second interesting observation is that communication of uncertainty appears to be deemed important since, over time, the central banks are moving to a more timely release of that information. For instance, per Reuters article on November 5, 2020 “The Federal Reserve will publish new color around policymaker outlooks for interest rates and the economy, and release some details earlier, changes that could give fresh insight into rate-setting decisions ... [the changes should] provide a timely perspective on the risks or uncertainties that surround the modal or baseline projections,” Powell said, “thereby highlighting some of the risk management considerations that are relevant for monetary policy.” The heterogeneity in nature of central bank projections throughout the current global pandemic highlights the lack of consensus about when, how, or why to use predictive densities. Thus, it is not obvious how or even if information regarding risk and uncertainty contained in density forecasts matters from a policy perspective, even though density forecasts could conceivably influence financial markets and economic activity.

Despite this widespread publication of density forecasts, there is very limited evidence on how higher-order moments matter from a policy perspective. Rholes and Petersen (2021) show experimentally that communicating uncertainty alongside a point projection of inflation can increase individual-level forecast errors and forecast uncertainty relative to communicating only point projections. Hubert and Maule (2021) show empirically that private expectations respond strongly to the BoE’s signals about future economic activity as conveyed in its Quarterly Inflation Report (QIR), suggesting that density forecasts might operate through a signaling channel. Hansen et al. (2019) use textual analysis of the BoE’s QIR to show that the economic uncertainty conveyed by text (and which is orthogonal to information conveyed numerically at the 2-year forecast horizon) can have increasingly large effects along the yield curve. Further, the heterogeneity in the nature of central bank projections throughout the current global pandemic highlights the lack of consensus about when, how, or why to use predictive densities. Thus, it is not obvious how or even if information regarding risk and uncertainty contained in density forecasts matters from a policy perspective, even though density forecasts could conceivably influence financial markets and economic activity.

The goal of this paper is to characterize and understand whether and how the financial markets and private sector expectations respond to the revisions of higher-order moments of the economic outlook. We use the data from the BoE because the BoE provides the longest time series of relevant data for an empirical study. We rely on a high-frequency identification approach, taking the revisions from one publication to the next, announced on a specific calendar day to identify the effects of central bank economic outlook (revisions) on the UK term structure

of interest rates and the expectations of professional forecasters. The paper closest to our work is Hansen et al. (2019) who use textual analysis as well as density forecast updates to understand the long interest rate sensitivity to macroeconomic news (relative to the short end of the yield curve). The main difference between their work and ours is that we differentiate between uncertainty and skewness to better understand the various properties that are typically important in considering density forecasts.

We find that revisions to higher-order moments matter more than revisions to first-order moments for interest rate dynamics. In fact, there is not much action due to the first-moment movements. On the other hand, an increase in the higher moments of output growth densities inverts the yield curve, typically known to forecast a recession. On the other hand, an increase in inflation uncertainty does not seem to affect the financial markets much, while skewness appears to be important. Moreover, the effect of higher moments is state-dependent — uncertainty revisions play through in expansions, while skewness revisions provide informative signals in contractions. Historical decomposition suggests yields have responded to higher-order moments for many decades, and the response was strong during the financial crisis.

When looking at the Blue Chip Financial Forecasts, we find that the three-, six- and twelve-month ahead consensus forecasts (and forecast disagreement) of short- and long-term interest rates are strongly correlated with revisions to uncertainty, while revisions to the balance of risks, expressed through revisions to skewness, are not.

Our paper is organized as follows. Section 2 presents our key data sets and our data transformations and introduces the institutional details for the BoE communication through the monetary policy cycle. Section 3 lays out our identifying assumptions and estimation strategy. Section 4 presents our findings while Section 5 concludes.

2 Data

The BoE's nine-person Monetary Policy Committee (MPC) began publishing density forecasts, termed 'fan charts', of inflation, output growth, and unemployment in its QIR in 1997.¹ The MPC, which adheres to an inflation-targeting regime adopted in the U.K. 1993, publishes fan charts to convey the inherent uncertainty surrounding its economic outlook and to provide its collective outlook on the balance of risks. These fan charts present deciles of subjective estimates of the probability distribution of the Banks forecast of each of these key macroeconomic vari-

¹The BoE, in conjunction with the National Institute of Economic and Social Research, began publishing a measure of uncertainty surrounding inflation forecasts as early as February 1996.

ables² (Elder, 2005; Mitchell and Weale, 2023). We provide an example of these projections for GDP growth, inflation, and unemployment in Figure 1.

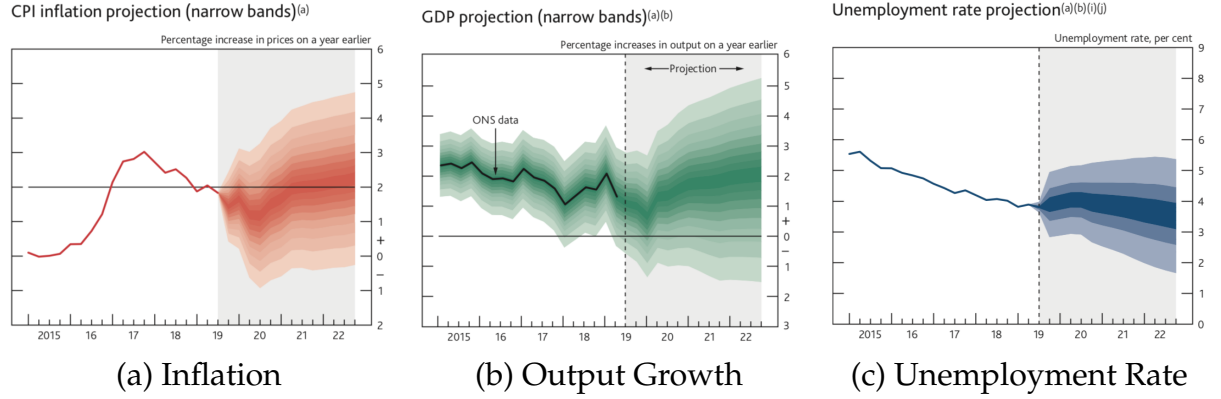


Figure 1: This figure shows the density forecasts of CPI inflation, GDP growth and unemployment rate from the November 2019 Monetary Policy Report published by the Bank of England.

The MPC constructs its fan charts using a split-normal distribution with a common mode but two different variances.³ The BoE's particular representation centers the distribution at the mode (μ) with an uncertainty (σ^2), while the skewness (ζ) controls the relative behavior of the two halves of the distribution. Thus, the probability density function (pdf), $f(x)$, is given by

$$f(x) = \begin{cases} Ae^{-\frac{[-(x-\mu)]^2}{2\sigma_1^2}}, & x \leq \mu \\ Ae^{-\frac{[-(x-\mu)]^2}{2\sigma_2^2}}, & x \geq \mu \end{cases} \quad (1)$$

where the following definitions hold

- $\sigma^2 = (1 + \gamma)\sigma_1^2 = (1 - \gamma)\sigma_2^2$,
- $E(X) = \mu + \sqrt{\frac{2}{\pi}}(\sigma_2 - \sigma_1)$
- $var(X) = (1 - \frac{2}{\pi})(\sigma_2 - \sigma_1)^2 + \sigma_1\sigma_2$
- $A = \left(\frac{\sqrt{2\pi}(\sigma_1 + \sigma_2)}{2}\right)^{-1}$
- $\gamma^2 = 1 - 4\left(\frac{\sqrt{1 + \pi\zeta^2/\sigma^2} - 1}{\pi\zeta^2/\sigma^2}\right)^2$

²The QIR itself includes graphical depictions of these fan charts. However, the BoE makes available the numerical information used to construct these fan charts at www.bankofengland.co.uk.

³These values are identical whenever the distribution is symmetric about the mode, which will equal the mean.

and parameter values are specific to forecast horizons and forecast date. That is, the MPC assigns potentially unique values of σ , γ for each forecasting horizon at each time t . For example, when forecasting period $t+k$ while in period t , the MPC sets values of $\sigma_{t,t+k}$, $\gamma_{t,t+k}$ for $k \in K \equiv \{1, 2, \dots, K\}$, where $\sigma_{t,t+k}, \sigma_{t,t+j}$, $j, k \in K$ can differ and $\sigma_{t,t+k}, \sigma_{(t-1),(t+k+1)}$ can differ.

Our focus in this paper is on how both the UK's yield curve and private expectations respond to revisions to the BoE's outlook on uncertainty (σ , the second central moment) and risk (γ , the third central moment).⁴ For comparative purposes, we also consider how yields and expectations respond to revisions of the point forecasts (μ , the first central moment).

An intuitive way to think of σ is as the uncertainty surrounding the MPCs point projections of some key economic variable. Visually, an increase in this parameter creates a more diffuse set of uncertainty bands surrounding the MPC's point projection⁵ One can think of γ as a measure of the MPC's subjective outlook on the balance of risks associated with some key economic variable. Whenever $\gamma = 0$, the split-normal density function reduces to the pdf of the normal distribution. Positive values of γ indicate that upside risk dominates downside risk. Visually, this would mean the MPC's density forecast is positively-skewed. The opposite is true for negative values of γ .⁶

To get a sense of how the MPC sets σ , γ , and how its approach might have changed over time, we consider the MPC's outlook on uncertainty and risk surrounding both the 2008 collapse of Lehman Brothers and the 2016 'Brexit' referendum.

When Lehman Brothers filed for bankruptcy on September 15, 2008, the firm held more than \$700 billion in liabilities, which obviates why the failure of Lehman's shocked global financial markets. The MPC addressed Lehman's collapse directly in its QIR in November 2008: "A number of institutional failures, and in particular, the collapse of Lehman Brothers on 15 September, led to rising anxieties about the survival of other financial institutions internationally." This sense of 'rising anxieties' is reflected by the MPC's upward revision of GDP growth and inflation uncertainty measures at all forecast horizons in its 2008Q4 inflation report, which we show in Figure 3. The MPC responded to this event by revising its balance of risks so that its density forecasts of GDP growth and inflation were exactly symmetric.

Less than a decade later, on June 23, 2016, the United Kingdom's electorate voted to leave the European Union. The MPC expressed concern that the vote to leave the EU generated considerable economic uncertainty and argued in the Au-

⁴ γ is the difference between the mean μ and the mode of the distribution.

⁵Note that σ is equal to the standard deviation whenever the distribution is symmetric (i.e. $\gamma = 0$)

⁶Wallis (2005) provides a more in-depth discussion of the BoE's density forecast.

gust 2016 QIR that short- and medium-term output growth would fall as a consequence. Additionally, the MPC focused on the devaluation of the pound sterling and the prospect of ensuing supply constraints. These concerns prompted the MPC to increase its level of forecast uncertainty about GDP growth, relative to its May 2016 report. Regarding inflation, the MPC made clear that the declining exchange rate was a secondary concern relative to impending supply constraints and dampened growth projections. The MPC, charged with maintaining full output and price stability, faced a trade-off between an inflationary policy and recessionary pressures. Ultimately, the MPC declared a willingness to endure temporary inflation above its 2% target to boost demand and supply to avoid larger declines in output growth. Thus, the MPC decided to cut the policy rate by 25 basis points following the electorates vote to leave the EU. These considerations and the MPC's policy action is reflected by the MPC's decision to leave unchanged its GDP growth skewness parameters while increasing skewness parameters for its short- and medium-term inflation forecast.

2.1 Timing

The MPC sets policy monthly and conveys its economic outlook quarterly (in February, May, August, and November) in its QIR. From November 1997 through May 2015, the BoE published the QIR one week after announcing information about that month's monetary policy decision. Since August of 2015, the BoE has released information about monetary policy decisions and its economic outlook simultaneously in the QIR.⁷ Thus, we can consider the impact of uncertainty and skewness shocks independent of information about the BOE's policy decisions through the second quarter of 2015 but not after. Figure 2 provides an overview of the BoE's information release schedule.

2.2 Density forecast data

Density forecast data for inflation from 2004 and for GDP growth from 2007 are publicly available as part of the BoE's quarterly inflation report.⁸ Density forecast data for inflation from 1997 through 2005 and for GDP growth from 1997 through 2007 are available through the U.K.'s national archive.⁹ Density forecasts of inflation from 2004 and onward use the consumer price index (CPI) to measure inflation while older inflation data uses retail purchases excluding mortgage payments (RPIX) to measure inflation.

⁷The BoE began referring to this joint release as its Monetary Policy Report in November 2019.

⁸This data is available for download here: <https://www.bankofengland.co.uk/inflation-report/inflation-reports>

⁹This data is available for download here: <https://webarchive.nationalarchives.gov.uk/20170831105150/http://www.bankofengland.co.uk/publications/Pages/inflationreport/irprobab.aspx>

	Current schedule	New schedule
Wednesday		"Pre-MPC" meeting with staff presentations to the MPC (joint MPC-FPC briefing meetings will take place four times a year).
Thursday		Stage 1: MPC deliberation meeting
Friday	"Pre-MPC" meeting with staff presentations to the MPC	
Monday		Stage 2: MPC policy discussion meeting
Tuesday		
Wednesday	Stage 1: MPC deliberation meeting	Stage 3: MPC decision meeting
Thursday	<ul style="list-style-type: none"> • Stage 2: MPC discussion and decision meeting • Announcement of monetary policy decision 	<ul style="list-style-type: none"> • Announcement of monetary policy decision and simultaneous publication of minutes • Inflation Report and press conference in Inflation Report months
Wednesday one week later	Inflation Report and press conference in Inflation Report months	
Wednesday two weeks later	Minutes published	

Figure 2: This figure, taken from the BoE's transparency and accountability report published December 11, 2014, summarizes the two different policy and communication schedules used by the BoE over our time sample.

The BoE publishes two different density forecasts for both GDP growth and inflation, each following a different assumption about the interest rate path. The bank forms one set of density forecasts by assuming that the prevailing nominal interest rate will continue throughout the forecast horizon. We call this the constant-rate assumption. The second set of density forecasts assumes the nominal interest rate matches the market's expected nominal rate throughout the forecast horizon. We call this the market-rate assumption. The BoE formed projections under the constant-rate assumption for 9 quarters, including a current-quarter forecast, from the beginning of our sample through May 2013 and for 13 quarters thereafter. The bank makes projections under market assumptions, when available, out to 13 quarters.

These two alternative interest rate assumptions typically yield different values of μ (point forecast) but have no impact on σ (uncertainty) or γ (skewness). Thus, we need not discern between these two interest rate assumptions when considering how rates and expectations respond to higher-order forecast moments.

We show the first moment of the BoE's density forecasts of GDP growth and

inflation formed using both interest rate assumption in Figure 3, and σ and γ for both GDP growth and inflation in Figure 4.

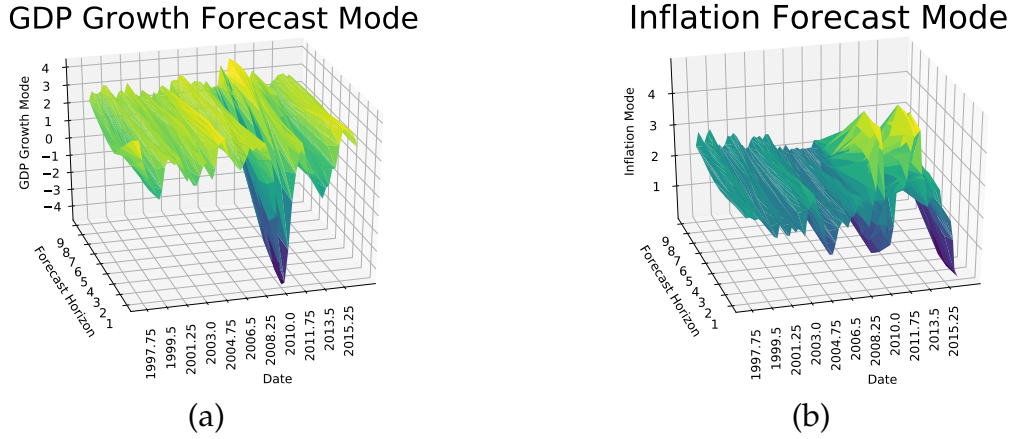


Figure 3: This figure shows the first moments of the BoE’s density forecast of GDP growth (panel a) and inflation (panels b) under the assumption of market-expected interest rates (z-axis) for a nine-quarter forecast horizon (y-axis) starting in the last quarter of 1997 and ending in the second quarter of 2015 (x-axis).

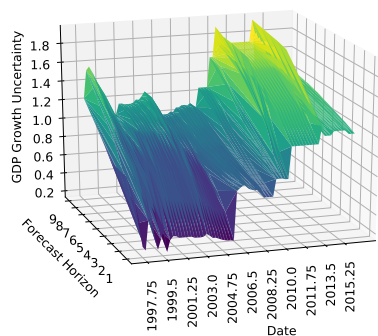
Note that movements in σ and γ are highly correlated over time across their respective forecast horizons. Also worth noting is that σ_π , σ_γ , for all forecast horizons, both exhibit considerably more variability before and through the early parts of the Great Recession than in subsequent forecasts, suggesting that MPC may have changed its approach to parameter selection. This change seems to correspond to the peak of a drastic increase in uncertainty surrounding density forecasts of both GDP growth and inflation that occurred following the global financial crisis.

Notice also that γ_π rarely takes on positive values at any forecast horizon. This means the BoE rarely expects downside risk to dominate across our time sample. Also interesting to note is that γ_γ doesn’t behave similarly across the two most recent recessions, which aligns with the idea that the MPC has begun taking a much different approach to how it selects values of γ and σ when forming its density forecasts. Finally, it is also noteworthy that neither of γ_π nor γ_γ appears to be strongly counter-cyclical, which is true of most other measures of economic uncertainty (Bloom, 2014).

2.3 Daily data for the UK’s yield curve

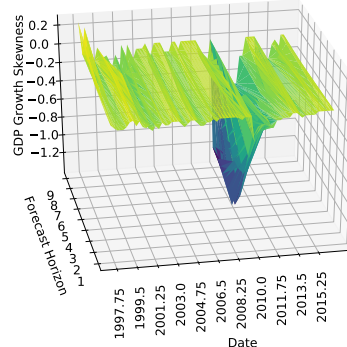
The BoE forms two different daily estimates of maturities along the UK’s yield curve. The first estimate is for maturities ranging from one to 60 sixty months in one-month intervals. Following the BoE’s nomenclature, we refer to this as the ‘short-end’ estimate. The second estimate is for maturities ranging from six

GDP Growth Forecast Uncertainty



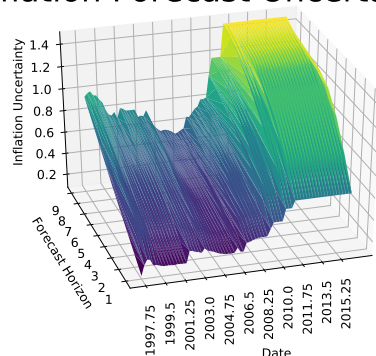
(a)

GDP Growth Forecast Skewness



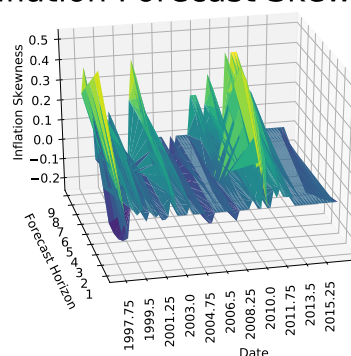
(b)

Inflation Forecast Uncertainty



(c)

Inflation Forecast Skewness



(d)

Figure 4: This figure shows the uncertainty and skewness (z-axis) of GDP growth (panels a and b) and inflation (panels c and d) for a nine-quarter forecast horizon (y-axis) starting in the last quarter of 1997 and ending in the second quarter of 2015 (x-axis).

months to 25 years in six-month intervals. The BoE has more recently begun estimating maturities out to 40 years. However, we only consider yields out to 25 years since these estimates are available for our entire time sample. We refer to this as the ‘long-run’ estimates. Data for both estimates are publicly available through the BoE’s website.¹⁰

The BoE bases its daily estimates of maturities on UK government bonds (gilts) and yields in the general collateral (GC) repurchase agreement (repo) market. The bank generates synthetic zero-coupon bonds from GC repo rates to improve its estimates of shorter maturities, where gilts tend to be less liquid. The BoE uses the variable roughness penalty (VRP) method, based on the spline-based technique proposed by Waggoner (1997), to estimate the yield curve.

¹⁰This data is available for download at <https://www.bankofengland.co.uk/statistics/yield-curves>

	2010Q1	2010Q2	2010Q3	2010Q4	2011Q1	2011Q2	2011Q3	2011Q4	2012Q1	2012Q2
2010Q1	H0	H1	H2	H3	H4	H5	H6	H7	H8	
2010Q2		Z0	Z1	Z2	Z3	Z4	Z5	Z6	Z7	Z8

Figure 5

2.4 Data Transformations

Revisions

We obtain our main results using a high-frequency identification strategy. We project changes in maturities across the UK's yield curve that occur in a small window surrounding QIR releases onto shocks to the first three central moments of the BoE's density forecasts of GDP growth and inflation. Thus, we transform our data to obtain two primary components: a left-hand-side variable capturing these maturity changes and right-hand-side variables capturing shocks to our moments of interest for inflation and GDP growth (i.e. information shocks).

To obtain our information shocks, we exploit the fact that the BoE often revises the value of our parameters of interest used to forecast economic values for some fixed point in time in sequential forecasts. For example, consider the BoE's forecasts of some key economic variable made in the first and second quarters of 2010, which we depict in Section 2.4.

Given these two forecasts, we first compute the change in parameter values μ , σ , γ associated with the BoE's density forecasts of inflation and output growth for each possible forecast horizon. Using values depicted in Section 2.4, we compute these revisions for the second quarter of 2010 as $\Delta Z_0 = Z_0 - H_1$, ΔZ_1 , ΔZ_2 , ..., ΔZ_7 . Next, we use principle component analysis (PCA) to extract the first principle component from ΔZ_0 , ΔZ_1 , ΔZ_2 , ..., ΔZ_7 to obtain $PC_{x,j,t}$ where $x = \{\mu, \sigma^2, \zeta\}$, $j = \{\pi, Y\}$. This approach allows us to summarize the majority of the variation contained in our information shocks while abstracting away from certain horizon-specific idiosyncrasies. We graph these information shocks for GDP growth and inflation in Figure 6.¹¹ Section 2.4 gives the level of correlation between these factors.

We follow Diebold and Li (2006) to extract time-varying level, slope, and curvature factors from our daily yield estimate that models the UK's yield curve. Because we have daily data, we obtain daily estimates of these factors and can use the change in these estimated factors that occurs on inflation report release dates as outcome variables in the high-frequency identification scheme described above. To extract these factors, we use all available yield curve data ranging from the last quarter of 1997 through the second quarter of 2015. For maturities ranging from one to 60 months, we use data from the BoE's short-end yield curve estimates.

¹¹We provide similar graphs for levels in Figure 12 our appendix.

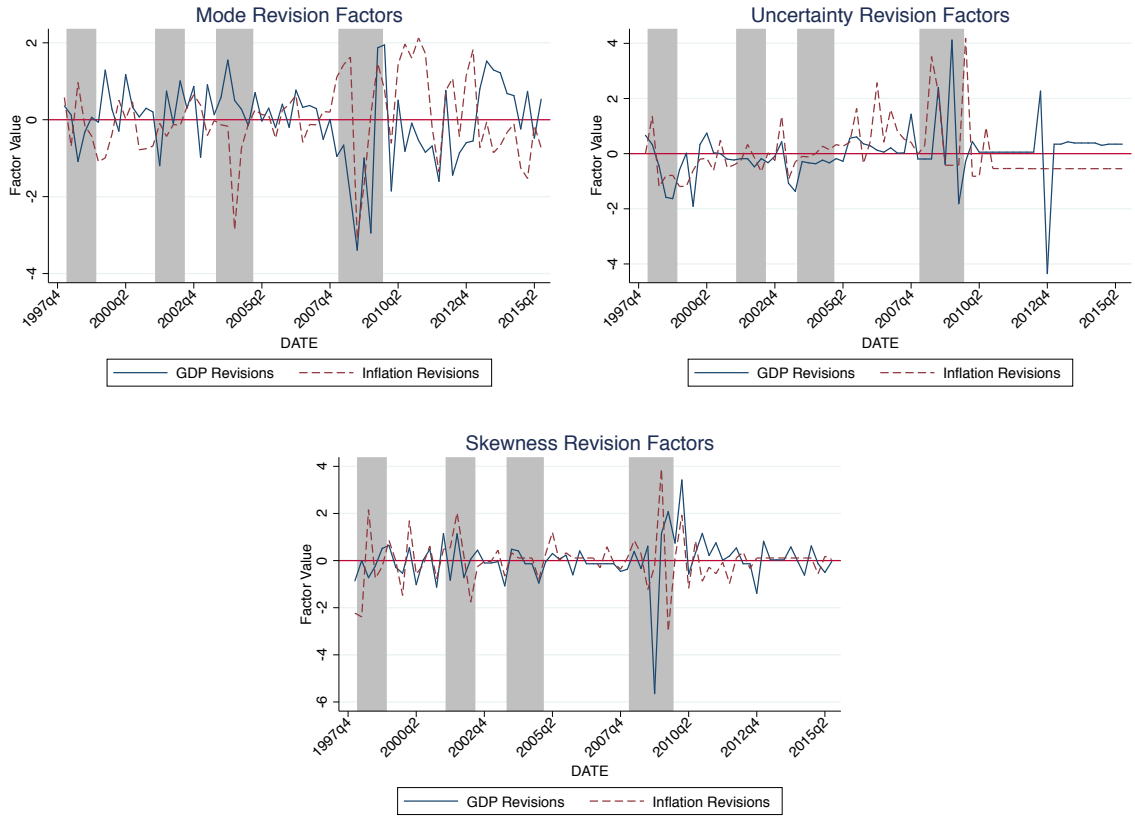


Figure 6: This figure shows the density forecasts of CPI inflation, GDP growth and unemployment rate from the November 2019 Monetary Policy Report published by the Bank of England.

These estimates are available for available for maturities at monthly intervals. For maturities further into the term structure, we use the BoE’s estimates of the full yield curve. These estimates are available for maturities ranging from six months to 25 years in six-month intervals. Using this data, we estimate Equation (2)

$$yield_t(m) = \beta_{1,t}X_1 + \beta_{2,t} \left(\frac{1 - e^{\lambda_t m}}{\lambda_t m} \right) + \beta_{3,t} \left(\frac{1 - e^{\lambda_t m}}{\lambda_t m} - e^{\lambda_t m} \right). \quad (2)$$

Interpretations of each $\beta_{i,t}$ that follows directly from the behavior of the corresponding loadings $X_{i,m}$. We show both the factors and the loadings in Figure 7.

The loading on $\beta_{1,t}$, X_1 , is a constant so that we can interpret $\beta_{1,t}$ as our long-term (i.e. level) factor. $X_{2,m}$ decays monotonically from one to zero and so we can interpret $\beta_{2,t}$ as our short-term (i.e. slope) factor. $X_{3,m}$ begins at zero, peaks at around 30 months, and then converges monotonically toward zero. Thus, we treat $\beta_{3,t}$ as the medium-term (i.e. curvature) factor.

Table of Correlations for Revision Factors						
	Y_μ	Y_{σ^2}	Y_ξ	π_μ	π_{σ^2}	π_ξ
Y_μ	1.00					
Y_{σ^2}	-.3170	1.00				
Y_ξ	.0088	.1773	1.00			
π_μ	-.0038	-.1795	.1413	1.00		
π_{σ^2}	-.0885	.1552	-.0399	.0727	1.00	
π_ξ	-.3030	.2649	.1937	.0311	-.0354	1.00
Prop.	.6	.67	.86	.69	.66	.83

Table 1: This table summarizes the correlations between each of our information shocks. The final row provides the proportion of variation captured by the first principle component if individual parameter revisions.

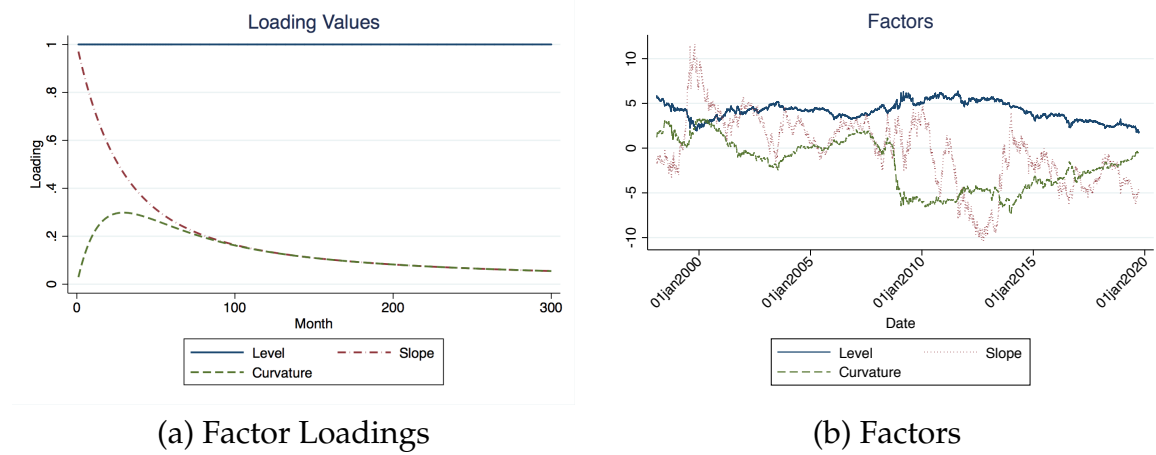


Figure 7: This figure shows factor loadings (panel a) used to extract our short-, medium-, and long-term yield curve factors (panel b).

3 Results

This section presents our main results, wherein we use a high-frequency identification approach, developed by Cook and Hahn (1989), Kuttner (2001), and Cochrane and Piazzesi (2002). For examples of this identification approach, see Nakamura and Steinsson (2018) or Gürkaynak et al. (2005).

In our context, this identification scheme involves projecting changes in yield curve factors that occur in the 24-hour window surrounding QIR releases onto our information. Our identifying assumption is that variation in our outcomes of interest that occur in our tight window surrounding the release of the inflation report is driven by the information shocks we extract from sequential QIR releases.

Using $\Delta\beta_{i,t}$ and $\Delta PC_{x,j,t}$, where $i = \{l, c, s\}$ we estimate:

$$\Delta\beta_{i,t} = \alpha_{i,t} + \sum_{x,j} \psi_{x,j} PC_{x,j,t} + \kappa_{i,t} FTSE_{t-1} + \epsilon_{i,t} \quad (3)$$

which we estimate using a Newey-West estimator. Here, $\psi_{x,j}$ captures the causal relationship between our information shocks and yield maturities, and $FTSE_{t-1}$ is a daily, market-based measure of economic uncertainty. Note that we standardize both the revision factors and the yield curve factors. This is helpful for two reasons. First, we can compare the magnitude of estimated effects for a given yield curve factor for each of our six revision factors. Second, for a given revision factor, we can make relative comparisons of its effect on the three different yield curve factors.¹²

We graph $\psi_{x,j}$ for each yield factor estimated using our full data sample, during expansions, and during contractions in Figure 8.¹³ Graphs in the left column show estimates of how the U.K.'s yield curve responds to revisions of the BoE's GDP growth forecast. Graphs in the right column do the same for inflation. Within each panel, there are three clusters of three coefficients each surrounded by confidence intervals that fade from 60% (darkest) and to 90% (lightest). Each cluster of coefficients corresponds to either the mean, uncertainty, or skewness revision factor. Within each cluster, we provide estimates from a model that uses revisions of the level, slope, or curvature factor as the dependent variable.

Estimates of Equation (3) using our full data indicate that revisions to the higher-order moments of both GDP growth and inflation matter at least as much as revisions of the respective first-order moments. We see that an increase in inflation forecast uncertainty puts upward (downward) pressure on long-term (short-term) maturities and that medium-term yields exhibit a negative response to an increase in outlook on the balance of risks. Also, we see that an increase in the MPC's outlook on the balance of risks for GDP growth can put upward (downward) pressure on short- and medium-term (long-term) rates. This suggests that information contained in the third-moment of MPC's predictive density of GDP growth might contribute to yield curve inversion, which is commonly interpreted as a sign of an impending recession.

We gain additional insight by considering the state-contingent effects of information shocks. To do this, we extract the cyclical component of U.K. GDP using the Hamilton filter Hamilton (2018), and partition our data into expansionary and contractionary periods whenever the cyclical component of GDP is positive or negative, respectively. We can then consider separately the effects of equivalent information shocks during expansions and contractions. We graph the results of this exercise in panels (c) through (f) of Figure 8.

¹²We also show the response of individual yields across the U.K.'s yield curve to information shocks in our appendix.

¹³We also provide coefficient estimates in tabular form in the appendix.

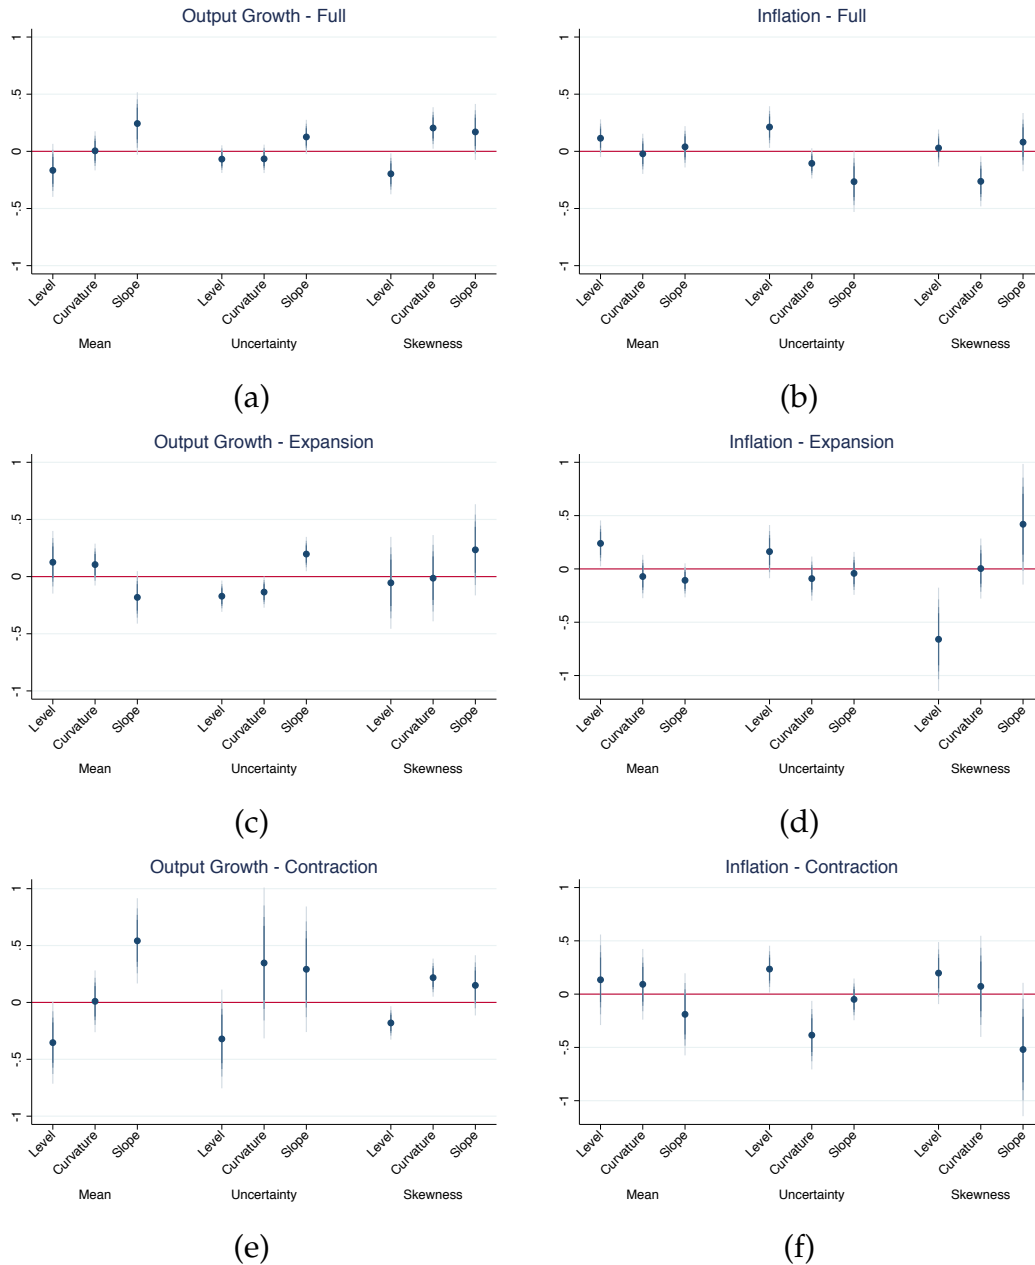


Figure 8: This figure plots the relationship between changes to the level, slope, and curvature yield factors and shocks to the first three moments of the BoE’s forecast of inflation (right column) and GDP growth (left column). We provide estimates using our full data sample (top row), during expansions (middle row), and during expansions (bottom row).

Panels (c) and (d) of Figure 8 depict the results of estimating Equation (3) over expansionary periods. For GDP growth, the response of yield factors to first-moment revisions is reversed relative to full-sample estimates. Further, we see that yields exhibit a muted response to skewness shocks during expansions relative to full-sample estimates.

We also find qualitative differences in how yields respond to information shocks surrounding inflation during expansions relative to full-sample estimates.

R^2	Full	Contraction	Expansion
<i>Level</i>	0.21	0.43	0.31
<i>Curvature</i>	0.26	0.47	0.09
<i>Slope</i>	0.22	0.45	0.29

The response of long- and short-term yields to skewness shocks is exacerbated while the response of the yield curve to uncertainty shocks is muted.

First-order moments of the MPC’s output growth forecast are much more important during contractions than during expansions, with positive shock putting significant upward (downward) pressure on short-term (long-term) yields. Further, we now see that shocks to both higher-order moments of the GDP growth forecast put upward pressure on short- and medium-term yields, and downward pressure on long-term yields.

Finally, we consider how yields respond to inflation information shocks during a contraction. Most interesting to note is that the yield curve exhibits a qualitatively different response to inflation skewness shocks during contractions relative to expansions.

We ensure the validity of our results by performing a placebo check wherein we re-estimate Equation (3) using changes to yield curve factors that occur on days when the BoE did not release a QIR. If our main results reflect a true causal relationship between financial markets and information contained in the higher-order moments of the BoE’s forecast, we would expect placebo estimates to be relatively precise zeros, with some tolerance for type one errors. We show the results of this placebo test in Figure 9, which we interpret as compelling evidence that our main results are not driven by spurious correlation.¹⁴

We also consider how well our information shocks explain changes in short-, medium-, and long-term yields over our full sample, during contractions, and during expansions. To do this, we compare the coefficient of determination obtained from estimating Equation (3) over each possible data sample. We report these results in Section 3. Overall, the information shocks better explain yield curve movements during contractions than expansions.

3.1 Expectations

This section of the results deals with the relationship between higher-order moments of central bank forecasts and private expectations. Much empirical and experimental work supports the idea that central banks can use point forecasts to coordinate expectations and nudge boundedly-rational agents toward ex-ante

¹⁴We provide corresponding numerical coefficient estimates in tabular form in our appendix.

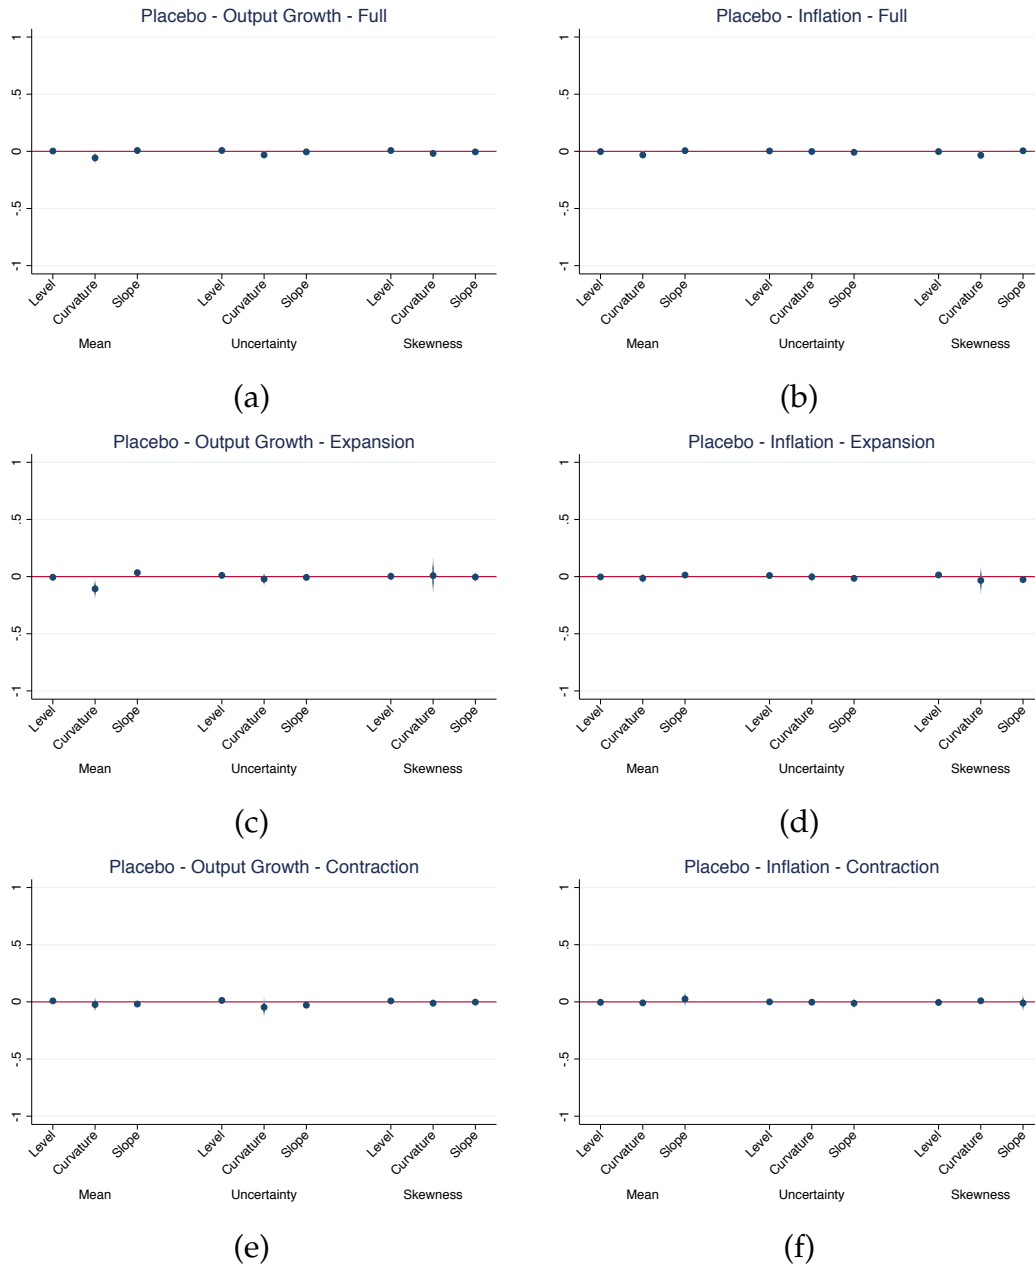


Figure 9: This figure shows placebo estimates of how Yield curve factors respond to changes in revision factors estimated using our full data sample, during recessions, and during expansions prior to 2015Q3.

rationality, thereby reducing economic volatility. However, there is very limited evidence regarding how higher-order moments influence expectation formation, if at all. We provide some suggestive evidence that professional forecasters do respond to higher-order forecast moments.

To do this, we estimate Equation (4) using as our outcome variables the three-, six-, and 12-month-ahead consensus forecast and forecast disagreement¹⁵ for the

¹⁵Measured as within-period forecast variation.

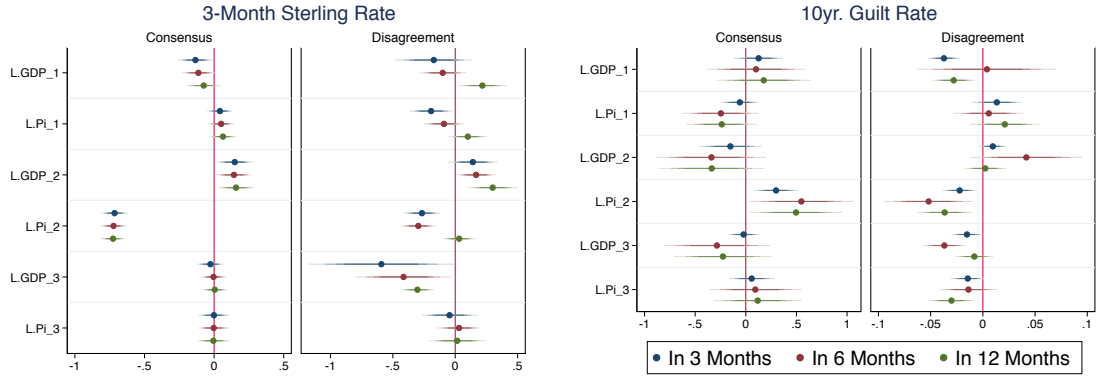


Figure 10: This graph shows how the three-month-, six-month-, and twelve-month-ahead consensus forecast and forecast disagreement of the three-month sterling and the 10-year Guilt rate respond to a one standard deviation change in each revision factor.

three-month Sterling rate and the 10-year Guilt rate, available at a monthly frequency from Blue Chip Financial Forecast (BCFF). The BCFF is a monthly survey administered at the beginning of each month. Because the BCFF is administered early in each month, survey respondents have already provided responses before the QIR release in QIR release months. This means that expectations cannot respond contemporaneously to information contained in the QIR. We account for this by projecting each of our three forecast outcomes onto lagged values of our information shocks. For example, the BoE releases its first-quarter inflation report about halfway through February each year. To understand how private expectations respond to information in this report, we project our variables of interest from the BCFF survey collected in March onto information shocks extracted from the February report. We graph these results in Figure 10.

$$\Delta Measure_{i,t} = \alpha_{i,t} + \sum_{x,j} \psi_{x,j} PC_{x,j,t-1} + \kappa_{i,t} FTSE_{t-1} + \epsilon_{i,t} \quad (4)$$

We first consider the relationship between short-term interest rate forecasts and our information shocks. Note that consensus forecasts at all horizons are at least as strongly correlated to forecast uncertainty as they are to first-moment revisions. This correlation is particularly strong for inflation uncertainty, which corresponds to our full-sample estimate of how short-term maturities respond to inflation uncertainty shocks. This logical coherence between expectations and yield movements also holds for GDP growth uncertainty, where both expectations and actual maturities exhibit a small, positive response to output growth uncertainty shocks. However, private expectations of short-term rates are seemingly uncorrelated with the third-moment of either predictive density.

Additionally, we see that inflation uncertainty and forecast disagreement are negatively correlated but GDP growth uncertainty and disagreement are posi-

tively correlated. This suggests market participants in the UK may better understand how the MPC, which officially follows an inflation-targeting regime, will respond to inflation uncertainty than to output growth uncertainty. Finally, we also see that expectations of short-term rates are negatively correlated with the third-moment of the MPC's forecast of GDP growth but are seemingly uncorrelated with the third-moment of the inflation forecast. This suggests that market participants better understand how the MPC will respond to upside inflation risk than to upside output risk.

Second, we consider the relationship between long-term interest rate forecasts and information shocks. Similar to short-term rates, we see that private expectations are at least as strongly correlated with higher-order information as they are with information about the central moment of both inflation and GDP growth density forecasts. We also see the same logical coherence between how expectations and maturities respond to information shocks.

Overall, the coherence we observe between how private expectations and maturities respond to second-order information shocks suggests that market participants better understand how maturities respond to uncertainty than to information surrounding risk.

3.2 Historical Decomposition

This section considers the historical importance of higher-order information for explaining yield curve changes by decomposing changes in level, slope, and curvature factors over time. We plot these historical decompositions in Figure 11.

This decomposition exercise shows that our results are not driven by anomalous events (i.e., the Great Recession). Rather, we see that higher-order information shocks have played a consistently-important role in explaining yield curve movements. Further, we see that higher-order information shocks have historically been at least as important for explaining yield changes as have first-moment information shocks.

Historical Decomposition by Yield Factor

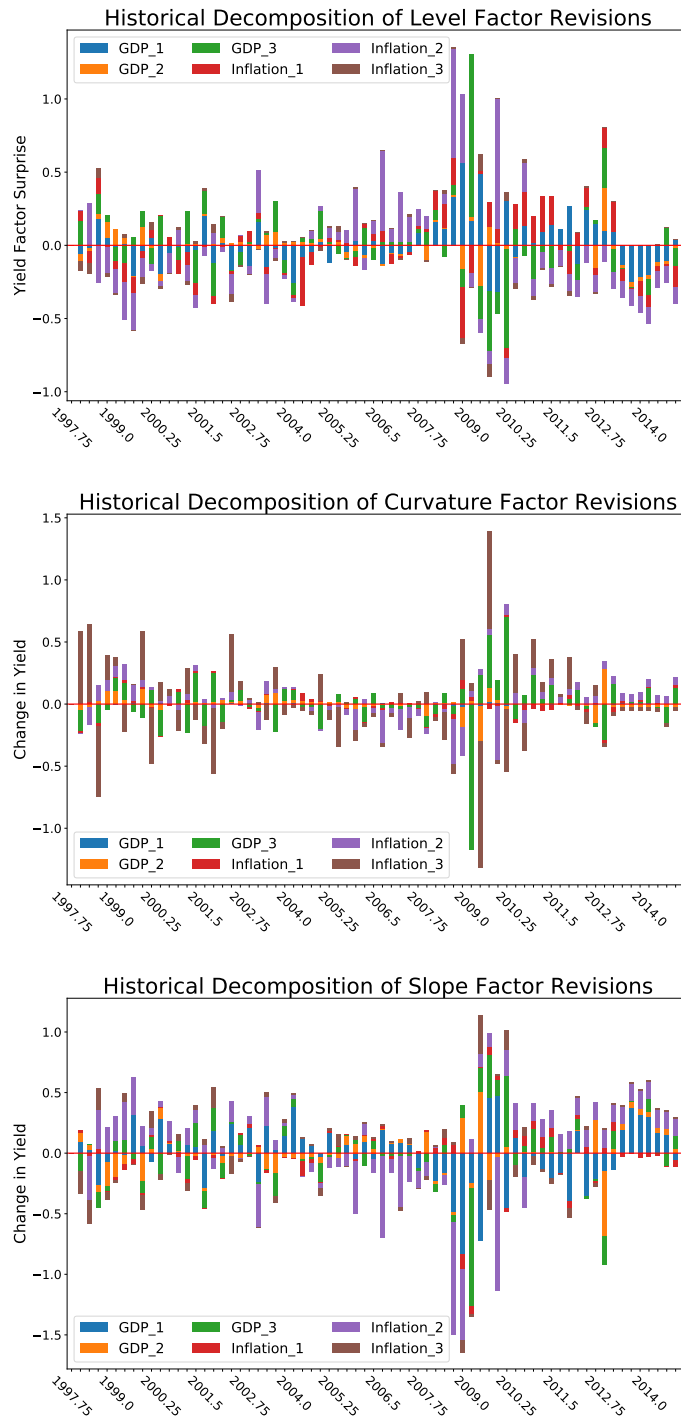


Figure 11: This figure shows the individual effect of each of our six information shocks on changes in our three yield curve factors over time. The top panel depicts the decomposition of the level factor, the middle the curvature factor, and the bottom the slope factor. GDP_i and $Inflation_i$ correspond to the i th moment of output growth and inflation density forecasts, respectively.

4 Conclusion

This paper considers the BoE's density forecasts and its revisions to quantify the effects of information flow on the financial markets and expectations. The paper contributes to the broader literature on the effects of news on financial markets and to the emerging literature on how information in the higher-order moments of central bank forecasts matters from a policy perspective.

We find that financial markets respond at least as strongly to the information contained in the higher-order moments of the BoE's density forecasts of output growth and inflation as to the information contained in the corresponding first moments. Further, we find that both the magnitude and direction of responses are state-contingent.

Additionally, we use Blue Chip Financial Forecast data to study how professional forecasters respond to information contained in higher-order forecast moments. We find that the consensus forecast and level of forecast disagreement of both short- and long-term interest rates are strongly correlated with higher-order forecast moments. Further, we observe a logical coherence between private expectations and realized yield changes, suggesting that market participants understand how rates will respond to higher-order forecast moments.

Overall, our results suggest that communicating high-order forecasts moments to market participants does affect their subsequent behavior. Important from a policy perspective is that higher-order moments can move markets, which suggests that density forecasting is a viable policy option.

References

- Ahrens, S., J. Lustenhouwer, and M. Tettamanzi (2023). The stabilizing effects of publishing strategic central bank projections. *Macroeconomic Dynamics* 27(3), 826–868.
- Andrade, P. and F. Ferroni (2021). Delphic and Odyssean monetary policy shocks: Evidence from the Euro Area. *Journal of Monetary Economics* 117, 816–832.
- Blinder, A. S., M. Ehrmann, M. Fratzscher, J. D. Haan, and D.-J. Jansen (2008). Central bank communication and monetary policy: a survey of theory and evidence. *Journal of Economic Literature* 46(4), 910–945.
- Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives* 28(2), 153–176.
- Cochrane, J. H. and M. Piazzesi (2002). The Fed and interest rates - a high-frequency identification. *American Economic Review* 92(2), 90–95.

- Cook, T. and T. Hahn (1989). The effect of changes in the federal funds rate target on market interest rates in the 1970s. *Journal of Monetary Economics* 24(3), 331–351.
- Diebold, F. X. and C. Li (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics* 130(2), 337–364.
- Elder, R. (2005). Assessing the mpc’s fan charts. Bank of England quarterly bulletin, Bank of England.
- Gürkaynak, R. S., B. Sack, and E. Swanson (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking* 1(1), 55–93.
- Hamilton, J. D. (2018). Why you should never use the Hodrick-Prescott filter. *The Review of Economics and Statistics* 100(5), 831–843.
- Hansen, S., M. McMahon, and M. Tong (2019). The long-run information effect of central bank communication. *Journal of Monetary Economics* 108(C), 185–202.
- Hoesch, L., B. Rossi, and T. Sekhposyan (2023). Has the information channel of monetary policy disappeared? Revisiting the empirical evidence. *American Economic Journal: Macroeconomics* 15(3), 355–387.
- Hubert, P. (2014). FOMC forecasts as a focal point for private expectations. *Journal of Money, Credit and Banking* 46(7), 1381–1420.
- Hubert, P. (2015). The influence and policy signalling role of FOMC forecasts. *Oxford Bulletin of Economics and Statistics* 77(5), 655–680.
- Hubert, P. and B. Maule (2021). Policy and macro signals from central bank announcements. *International Journal of Central Banking* 17(2), 255–296.
- Kryvtsov, O. and L. Petersen (2021). Central bank communication that works: Lessons from lab experiments. *Journal of Monetary Economics* 117(C), 760–780.
- Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics* 47(3), 523–544.
- Miranda-Agrippino, S. and G. Ricco (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics* 13(3), 74–107.
- Mitchell, J. and M. Weale (2023). Censored density forecasts: Production and evaluation. *Journal of Applied Econometrics* 38(5), 714–734.

- Nakamura, E. and J. Steinsson (2018). High-frequency identification of monetary non-neutrality: the information effect. *Quarterly Journal of Economics* 133(3), 1283–1330.
- Petersen, L. and R. Rholes (2022). Macroeconomic expectations, central bank communication, and background uncertainty: A COVID-19 laboratory experiment. *Journal of Economic Dynamics and Control* 143(C).
- Rholes, R. and L. Petersen (2021). Should central banks communicate uncertainty in their projections? *Journal of Economic Behavior & Organization* 183(C), 320–341.
- Swanson, E. T. (2006). Have increases in Federal Reserve transparency improved private sector interest rate forecasts? *Journal of Money, Credit and Banking* 38(3), 791–819.
- Waggoner, D. F. (1997). Spline methods for extracting interest rate curves from coupon bond prices. FRB Atlanta Working Paper 97-10, Federal Reserve Bank of Atlanta.
- Wallis, K. F. (2005). Combining density and interval forecasts: a modest proposal. *Oxford Bulletin of Economics and Statistics* 67(s1), 983–994.

5 Appendix

5.1 Tables and Figures

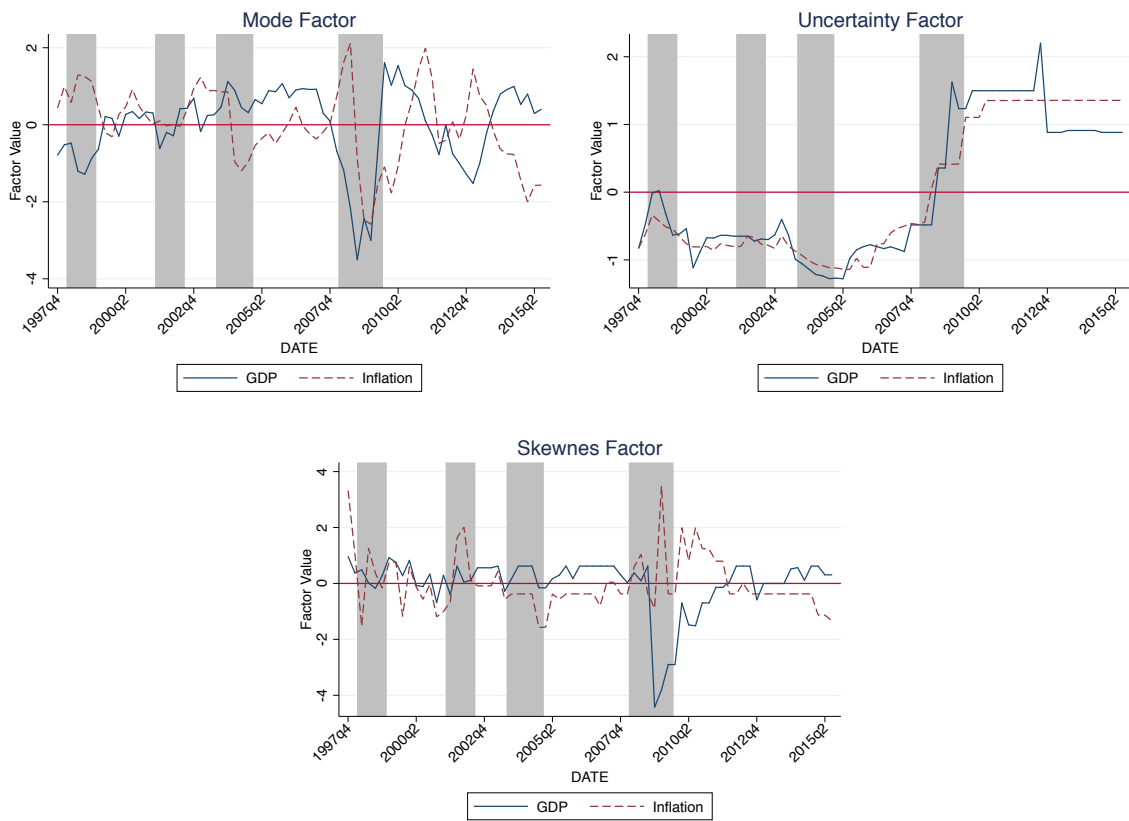


Figure 12: This figure shows the first principle component in levels for the first three forecast moments of inflation and GDP growth.

Revision Means and Factors

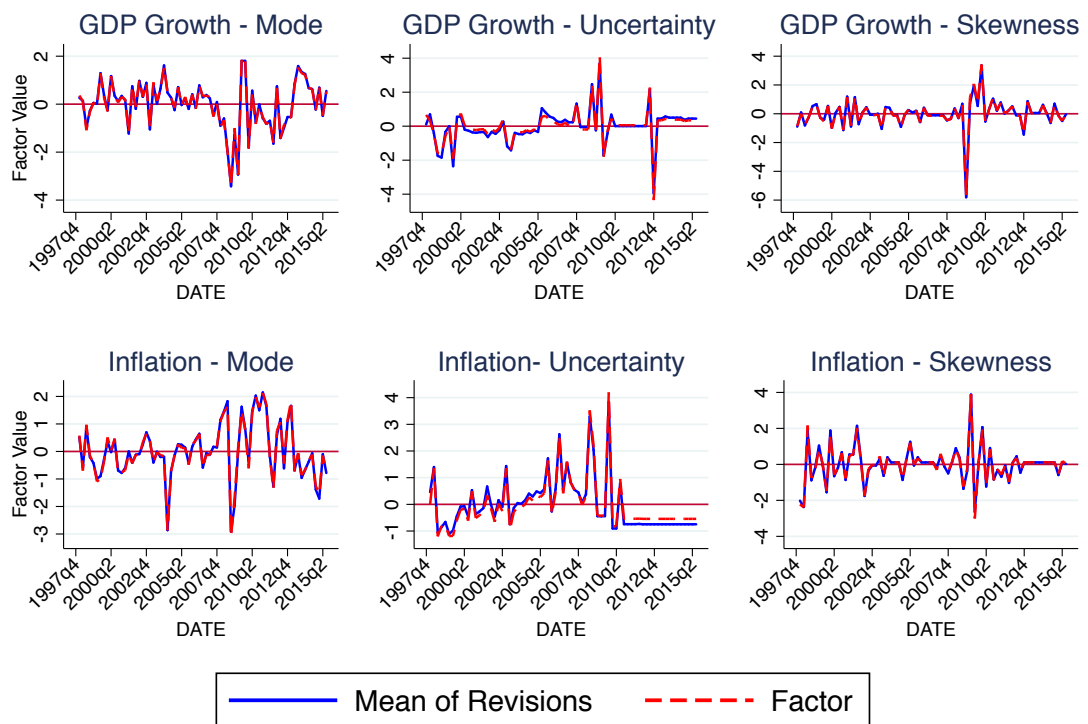


Figure 13: This plots the first principal component extracted from the revision of each of the first three central moments of GDP growth and inflation against their corresponding means.

Yield Responses to Information Shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	L _F	L _C	L _E	C _F	C _C	C _E	S _F	S _C	S _E
GDP_1	-0.166 (0.139)	-0.353 (0.204)	0.126 (0.162)	0.00443 (0.102)	0.00977 (0.153)	0.105 (0.108)	0.243 (0.163)	0.542 (0.211)	-0.18 (0.136)
GDP_2	-0.0680 (0.0725)	-0.321 (0.245)	-0.172 (0.0809)	-0.0657 (0.0743)	0.347 (0.374)	-0.135 (0.0807)	0.126 (0.0896)	0.292 (0.312)	0.197 (0.088)
GDP_3	-0.197 (0.107)	-0.180 (0.0827)	-0.0544 (0.238)	0.204 (0.109)	0.218 (0.0946)	-0.0138 (0.223)	0.170 (0.146)	0.150 (0.149)	0.23 (0.236)
Pi_1	0.115 (0.0985)	0.134 (0.240)	0.239 (0.127)	-0.0218 (0.105)	-0.190 (0.217)	-0.0714 (0.120)	0.0397 (0.108)	0.0921 (0.187)	-0.10 (0.094)
Pi_2	0.213 (0.109)	0.235 (0.124)	0.162 (0.148)	-0.105 (0.0791)	-0.0492 (0.111)	-0.0909 (0.122)	-0.265 (0.158)	-0.385 (0.182)	-0.041 (0.119)
Pi_3	0.0293 (0.0966)	0.197 (0.164)	-0.660 (0.286)	-0.263 (0.131)	-0.520 (0.353)	0.00321 (0.166)	0.0807 (0.152)	0.0729 (0.268)	0.419 (0.335)
Constant	-0.0272 (0.114)	0.0749 (0.209)	-0.162 (0.0941)	-0.0194 (0.107)	0.0885 (0.291)	0.159 (0.121)	0.0121 (0.141)	-0.219 (0.300)	0.25 (0.080)
R ²	0.21	0.43	0.31	0.26	0.47	0.09	0.22	0.45	0.29

Standard errors in parentheses

+ $p < 0.10$, * $p < 0.05$

Table 2: This table reports numerical estimates (standard errors in parentheses) of how the level, curvature, and slope factors respond to density moment revisions over our full sample, during expansions, and during contractions. Each column includes a two-letter label. Columns whose first letter are an L, C, or S use the level, curvature, or slope factor as the dependent variable, respectively. The subscript in each column label is either F, E, or C for full sample, expansionary, or contractionary. For example, column (5) reports estimates of how the curvature factor responds to moment revisions during contractions. Rows denote the first thru third moments of GDP growth and inflation.

Yield Responses to Information Shocks - Placebo

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	L _F	L _C	L _E	C _F	C _C	C _E	S _F	S _C	S _E
GDP_1	0.003 (0.007)	0.009 (0.012)	-0.006 (0.013)	-0.057 (0.029)	-0.023 (0.035)	-0.107 (0.048)	0.007 (0.008)	-0.019 (0.011)	0.035 (0.012)
GDP_2	0.008 (0.009)	0.013 (0.013)	0.011 (0.012)	-0.031 (0.022)	-0.047 (0.047)	-0.022 (0.033)	-0.005 (0.009)	-0.029 (0.012)	-0.007 (0.012)
GDP_3	0.007 (0.005)	0.008 (0.004)	0.003 (0.022)	-0.018 (0.017)	-0.012 (0.021)	0.008 (0.087)	-0.005 (0.006)	-0.002 (0.003)	-0.004 (0.026)
Pi_1	-0.002 (0.005)	-0.004 (0.013)	-0.003 (0.010)	-0.032 (0.020)	0.025 (0.034)	-0.015 (0.027)	0.006 (0.010)	-0.009 (0.015)	0.013 (0.016)
Pi_2	0.003 (0.005)	-0.000 (0.008)	0.009 (0.007)	-0.001 (0.018)	-0.012 (0.026)	-0.002 (0.026)	-0.008 (0.007)	-0.003 (0.009)	-0.015 (0.013)
Pi_3	-0.002 (0.009)	-0.005 (0.014)	0.014 (0.016)	-0.034 (0.023)	-0.011 (0.039)	-0.033 (0.067)	0.005 (0.009)	0.009 (0.011)	-0.026 (0.022)
Constant	-0.001 (0.006)	-0.007 (0.015)	0.004 (0.011)	-0.016 (0.025)	-0.151 (0.060)	-0.001 (0.036)	0.000 (0.008)	0.026 (0.024)	-0.004 (0.010)

Standard errors in parentheses

+ $p < 0.10$, * $p < 0.05$

Table 3: This table reports numerical estimates (standard errors in parentheses) from our placebo exercise, wherein we project changes in the level, curvature, and slope factors on non-QIR release days onto our information shocks. Each column includes a two letter label. Columns whose first letter are an L, C, or S use the level, curvature, or slope factor as the dependent variable, respectively. The subscript in each column label is either F, E, or C for full sample, expansionary, or contractionary. For example, column (5) reports estimates of how the curvature factor responds to moment revisions during contractions. Rows denote the first thru third moments of GDP growth and inflation.

5.2 Alternative Estimation Strategy

Here we consider how individual yields respond to information shocks using the same identification strategy used to produce our main results. Rather than using changes in yield curve factors as our outcome of interest, we instead consider the change in yields for each maturity along the UK's yield curve:

$$\Delta yield_{m,t} = yield_{m,t} - yield_{m,t-1}. \quad (5)$$

This gives the change in maturity m on day t , where t corresponds to the release date of one of the BoE's inflation reports.

Thus, we estimate

$$\Delta yield_{m,t} = \lambda_{i,t} + \sum_{x,j} \theta_{x,j} PC_{x,j,t} + \kappa_{i,t} FTSE_{t-1} + \epsilon_{i,t} \quad (6)$$

for $m \in \{6, 12, \dots, 300\}$.

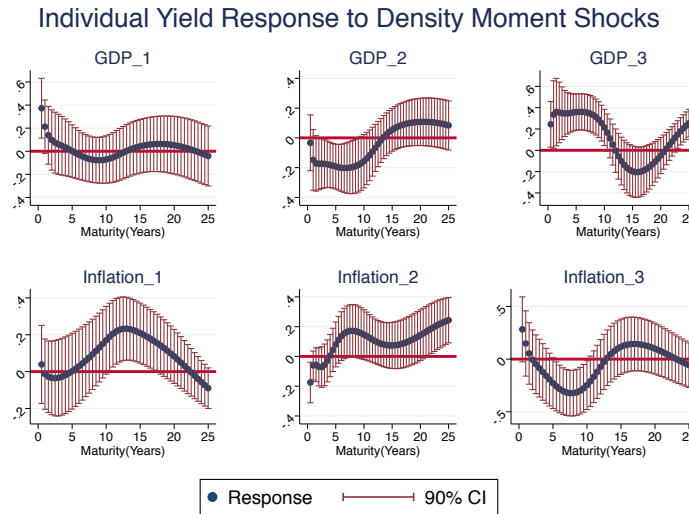


Figure 14: This graph shows how yields from the BoE's estimates of the full yield curve (left) and the short-end of the yield curve (right) respond to revisions of the first three central moments of GDP growth and inflation.