

DISCUSSION PAPER SERIES

DP13438

THE LONG-RUN INFORMATION EFFECT OF CENTRAL BANK COMMUNICATION

Stephen Hansen, Michael McMahon and Matthew
Tong

**MONETARY ECONOMICS AND
FLUCTUATIONS**

THE LONG-RUN INFORMATION EFFECT OF CENTRAL BANK COMMUNICATION

Stephen Hansen, Michael McMahon and Matthew Tong

Discussion Paper DP13438
Published 09 January 2019
Submitted 11 December 2018

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **MONETARY ECONOMICS AND FLUCTUATIONS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Stephen Hansen, Michael McMahon and Matthew Tong

THE LONG-RUN INFORMATION EFFECT OF CENTRAL BANK COMMUNICATION

Abstract

Why do long-run interest rates respond to central bank communication? Whereas existing explanations imply a common set of signals drives short and long-run yields, we show that news on economic uncertainty can have increasingly large effects along the yield curve. To evaluate this channel, we use the publication of the Bank of England's Inflation Report, from which we measure a set of high-dimensional signals. The signals that drive long-run interest rates do not affect short-run rates and operate primarily through the term premium. This suggests communication plays an important role in shaping perceptions of long-run uncertainty.

JEL Classification: E52, E58, C55

Keywords: monetary policy, communication, Machine Learning

Stephen Hansen - stephen.hansen@economics.ox.ac.uk
University of Oxford and CEPR

Michael McMahon - michael.mcmahon@economics.ox.ac.uk
University of Oxford and CEPR

Matthew Tong - matthew.tong@bankofengland.co.uk
Bank of England

Acknowledgements

Omiros Papaspiliopoulos and David Rossell provided numerous helpful suggestions on the statistical methodologies used in this paper. We would like to additionally thank, for comments, discussion and suggestions, Jan Marc Berk, Michael Ehrmann, Refet Gurkaynak, Emi Nakamura, Alessandro Riboni, Jon Steinsson, and Jonathan Talmi as well as seminar participants at the Banco do Brazil, Bank of Canada, Bank of England, Bilkent University, De Nederlandsche Bank, Durham, ECB, Empirical Macroeconomics Conference at Ghent University, Higher School of Economics, HKUST, New Economic School, Norges Bank, Oxford/NY Fed Conference, RCEA conference, University of Glasgow, and University of Pennsylvania. This paper has been prepared while McMahon was under the Lamfalussy Fellowship Program sponsored by the ECB. This views expressed here do not necessarily represent the views of the Bank of England, the Monetary Policy Committee, Financial Policy Committee, the ECB, the Eurosystem or anyone else other than the authors. Any errors remain ours alone.

The Long-Run Information Effect of Central Bank Communication*

Stephen Hansen[†] Michael McMahon[‡] Matthew Tong[§]

December 11, 2018

Abstract

Why do long-run interest rates respond to central bank communication? Whereas existing explanations imply a common set of signals drives short and long-run yields, we show that news on economic uncertainty can have increasingly large effects along the yield curve. To evaluate this channel, we use the publication of the Bank of England's Inflation Report, from which we measure a set of high-dimensional signals. The signals that drive long-run interest rates do not affect short-run rates and operate primarily through the term premium. This suggests communication plays an important role in shaping perceptions of long-run uncertainty.

Keywords: Monetary Policy, Communication, Machine Learning

JEL Codes: E52, E58, C55

*Omiros Papaspiliopoulos and David Rossell provided numerous helpful suggestions on the statistical methodologies used in this paper. We would like to additionally thank, for comments, discussion and suggestions, Jan Marc Berk, Michael Ehrmann, Refet Gürkaynak, Emi Nakamura, Alessandro Riboni, Jon Steinsson, and Jonathan Talmi as well as seminar participants at the Banco do Brazil, Bank of Canada, Bank of England, Bilkent University, De Nederlandsche Bank, Durham, ECB, Empirical Macroeconomics Conference at Ghent University, Higher School of Economics, HKUST, New Economic School, Norges Bank, Oxford/NY Fed Conference, RCEA conference, University of Glasgow, and University of Pennsylvania. This paper has been prepared while McMahon was under the Lamfalussy Fellowship Program sponsored by the ECB. This views expressed here do not necessarily represent the views of the Bank of England, the Monetary Policy Committee, Financial Policy Committee, the ECB, the Eurosystem or anyone else other than the authors. Any errors remain ours alone.

[†]University of Oxford, Alan Turing Institute and CEPR. Email: stephen.hansen@economics.ox.ac.uk

[‡]University of Oxford, CEPR, CAGE (Warwick), CfM (LSE), and CAMA (ANU). Email: mcmahon@economics.ox.ac.uk

[§]Bank of England. Email: matthew.tong@bankofengland.co.uk

1 Introduction

For the last two decades, central banks have increasingly used public communication to support their policy goals, and more specifically to manage the expectations that link the policy rates that central banks control to the market interest rates that determine economic decisions (Woodford 2001, Blinder 2008). As nominal policy rates were cut close to close to the zero lower bound, communication became a policy tool in its own right. Even as the major economies move towards more normal conditions, academics (Blinder 2018) and policymakers (Draghi 2017) expect communication to remain an important instrument. Understanding the effects of communication is therefore a key issue.

While the academic literature has established that central bank communication moves market interest rates at various maturities (Gürkaynak et al. 2005, Boukus and Rosenberg 2006, Blinder et al. 2008, Carvalho et al. 2016, for example), the channels through which this occurs are often unclear. In particular, there is an ongoing debate about why central bank communication moves *long-run* interest rates well outside the window within which central banks seek to obtain their policy goals.¹ Two recent and important contributions present different explanations. First, Nakamura and Steinsson (2017) argue that this is due to monetary policy shocks transmitting information about economic fundamentals that affects long-run market expectations of economic conditions, termed the *information effect* by Romer and Romer (2000). In contrast, Hanson and Stein (2015) argue that news about short-term policy expectations is propagated to longer-maturity bonds by the trading activity of yield-oriented investors. According to their model, decreases (increases) in short rates induce these investors to switch to (from) longer-maturity bonds, driving the yields on such bonds down (up) through changes in the term premium.

The first contribution of this paper is to draw attention to a third channel through which central bank communication can affect long-run interest rates, by providing news on risk and uncertainty around economic conditions, and thereby generating a change in the long-run term premium. This channel operates not by changing long-run expectations of economic conditions, but by changing the perceived variance of those conditions. We introduce a theoretical framework in which a central bank transmits information on both the level and uncertainty of short-run economic conditions. Both kinds of information can move long-run interest rates, but in different ways. Information on levels moves long-run rates when economic conditions are sufficiently persistent, in which case there is a similar move in short- and long-run market rates. By contrast, information on uncertainty has a much larger impact on long-run rates than short-run rates, as the impact of uncertainty accumulates over time.

¹This is part of a more general comovement of short- and long-term rates in response to other news as in, for example, Gürkaynak et al. (2005).

In the existing information effect literature, the focus is on news that changes market expectations of levels. These signals should move short-run interest rates by as least as much as long-run rates. The same is true for the channel in Hanson and Stein (2015), since the long-run term premium adjusts directly in response to the news that changes short-run expectations. In contrast, uncertainty signals can move long-run interest rates by more than short-run rates.

The second contribution of the paper is to construct a novel form of event study to identify the importance of these information channels. Implementing this empirical test requires finding an environment in which central bank communication provides information solely on economic conditions rather than policy stances, since this may confound observed market reactions. For this we use the publication of the Bank of England’s Inflation Report (IR) from February 1998 through May 2015. The IR contains information about the Bank of England’s economic forecasts, but does not provide explicit forward guidance on future policy. Moreover, during our sample period, the IR was published according to a fixed, quarterly schedule one week *after* the announcement of the contemporaneous policy decision. It therefore constitutes a policy-free central bank information shock. This allows us to directly assess the market impact of news about the Bank of England’s private views without having to decompose a policy change into separate information and policy shock components.²

A standard approach in the literature (Cook and Hahn 1989, Kuttner 2001, Bernanke and Kuttner 2005, among many others) is to use market rate reactions around central bank announcements to infer their information content. We instead directly measure a high-dimensional set of signals that the IR contains and use modern machine learning methods to analyze which of these signals are responsible for driving rate responses at different maturities. We also do the same for separate expectation and term premium components of these yields. The construction of a set of detailed, granular signals transmitted by the central bank allows us to open up the ‘black box’ of what precisely drives market responses, and to distinguish the uncertainty channel according to whether a set of signals drives long-run rate reactions distinct of short-run rates.

The IR contains numerical information in the form of the central bank’s inflation and GDP forecasts, as well as distributional information around those forecasts (variance and skew). It also contains narrative information in the form of extensive text, which we quantify using latent Dirichlet allocation (LDA), a popular probabilistic topic model (Blei et al. 2003) previously applied in monetary economics by Hansen and McMahon (2016) and Hansen et al. (2018). LDA represents each IR as a distribution over a finite set of topics that capture common themes in the data. For example, a particular IR might devote 20% of its space to inflation; 15% of its space to financial market conditions; and

²See Miranda-Agrippino and Ricco (2015) for an example of the latter.

so on. In all, we obtain 75 of dimensions of variation across every IR: 15 numeric signals, and 60 narrative signals.

Numeric information is explicitly designed to convey news on the state of the economy, but less clear is whether narrative information adds any news beyond this. To test this, we identify the number of significant narrative signals selected by an elastic net regression to explain the residual variance in bond yields (controlling for the effect associated with the numerical forecasts on both). We compare this to the number selected when we randomly permute the yield residuals and re-estimate the model. We overwhelmingly reject the hypothesis that narrative information is unrelated to the size of the interest rate moves not explained by numeric forecasts at both short- and long-term maturities. This allows us to exploit the high dimensionality of narrative to assess the different channels discussed above.

The third contribution of the paper is to provide robust evidence in favor of the uncertainty channel as the dominant factor in explaining why long-run rates react to IR publication. An initial finding is that news in the higher moments of the numerical forecasts explains an increasing proportion of interest rate variation at longer horizons. We then use the narrative signals in three ways to show that uncertainty signals drive the change in long-run interest rates:

1. We identify the narrative signals that are most robustly linked to yield movements at each maturity using a bootstrap procedure. There is little overlap in those signals that drive short-run interest rate changes and those that drive long-run changes. Moreover, the narrative signals that explain long-run rates feature words suggestive of uncertainty. We also repeat the analysis by yield component. The expectations (term premium) component drives the overall short-run (long-run) variation. Furthermore, the key signals that drive expectations components are highly correlated across the yield curve, whereas variation in term premiums is driven by independent signals. This suggests that the standard information effect on level expectations operates in the background, but does not account for most of the movement in long-run rates.
2. The IR contains two narrative parts, one describing current economic conditions and another describing the forecast and the risks around it. The signals that drive movements in short-run (long-run) rates come disproportionately from the former (latter) part.
3. We record the residual variance that the key narrative signals at each maturity explain and conduct a placebo test in which we replace them with the key signals from other maturities. The key signals for long-run (short-run) yields explain around a third of their residual variance, but little to none of the residual variance

for short-run (long-run) yields. The analysis by component again shows some evidence for the information effect on level expectations. An additional finding is that the key signals for short-run expectations do not explain changes in the long-run term premium.

In short, while an information effect appears to operate on long-run level expectations, it does not explain the overall long-run market rate reaction to the IR. Instead, the evidence is consistent with an uncertainty channel as the primary source of the reaction, via changes in the term premium. Unlike in Hanson and Stein (2015), this term premium effect is the direct result of news rather than an indirect result of trading activity.

Our findings are also related to a growing empirical literature on the effects of monetary policy events on term premiums. For example, Bundick et al. (2017) show that shocks to uncertainty about future interest rate decisions yields significant moves in long-term term premiums in the US. And Cieslak and Schrimpf (2018) show that monetary events in the US, Euro Area, UK and Japan are associated with risk preference shocks identified off the movement in term premiums across the yield curve. This literature does not emphasise the role of central bank communication about uncertainty as driving the effects. In fact, the findings could be unrelated to information effects as in the risk premium channel of monetary policy in Drechsler et al. (2018). Tang (2015) and Leombroni et al. (2018) link information effects from central bank communication to long-term interest rate movements with an explicit role for uncertainty. Unlike in our paper, the communication is about level expectations which then interact with given uncertainty that prevails at the time of the signal. In our paper, the central bank signals are themselves about uncertainty.

The main implication of our findings is that central bank communication can help shape market beliefs about long-run uncertainty. In our view, this channel has been under-appreciated in the literature on monetary policy, but is something that central banks should potentially consider as part of their overall communication strategies, especially as there is increasing evidence that uncertainty has macroeconomic effects (Bloom 2009, Fernandez-Villaverde et al. 2011, Jurado et al. 2015, Baker et al. 2016). This is also relevant for understanding the role that so-called ‘Delphic’ forward guidance can play, which Campbell et al. (2012) define as communicating a view on future conditions and at the same time describing a likely policy response to that view. Our paper shows that market reactions arise from communicating such views even in the absence of information on policy. Finally, a broader policy implication is the use of narrative as an instrument for managing expectations. Shiller (2017) recently introduced the notion of Narrative Economics, which emphasizes the role of narratives in spreading beliefs. In monetary policy, central banks have an important role in shaping public narrative (Haldane and McMahon 2018), and our work suggests this can generate different patterns of beliefs

among economic agents.

From a methodological perspective, we illustrate how to combine event study analysis with a high-dimensional set of regressors that measure signals from unstructured text data. Rather than simply describe how these features correlate with interest rates, we propose a framework that allows us to distinguish channels based on heterogeneity in the correlation patterns across yields. Given the popularity of event studies in the monetary policy literature, and the preponderance of text that accompanies many central bank communication events, the methods we propose have broad applicability. For example, Gürkaynak et al. (2018) find that the change in interest rates around central bank communication events in the US is only partially captured by headline numeric information. Our approach allows researchers to directly analyze the ‘missing’ information not accounted for in traditional analysis.

The paper is organized as follow. Section 2 describes the Inflation Report and the yield curve data; Section 3 introduces a framework that incorporates information effects on levels and uncertainty; and Section 4 explains how we measure the numeric and narrative information in the IR. Section 5 presents our core empirical findings, and Section 6 presents robustness results. Section 7 concludes.

2 IR Communication and the Yield Curve

In this section we motivate the focus on the Bank of England’s Inflation Report and the information it delivers. We then discuss the interest rate data we use, and show using an event study that the Report’s publication has an impact on market rates at a range of maturities.³

2.1 The Inflation Report

Following the adoption of inflation targeting in the UK in 1993, the Inflation Report (IR) has been published quarterly by the Bank of England. When the Bank of England was granted operational independence for monetary policy in May 1997, a nine-person Monetary Policy Committee (MPC) was established to set policy on a monthly basis in a way consistent with meeting its inflation target remit. Since independence, the IR has become the quarterly communication vehicle for the MPC and contained the Committee’s forecasts for GDP growth and inflation. In its own words, the IR “sets out the economic

³Reeves and Sawicki (2007) conducted an event study on the effect of IR publication on market rates, and find significant effects. Here we extend their analysis with a longer sample and different market rates.

analysis and inflation projections that the Monetary Policy Committee uses to make its interest rate decisions.”

Our sample comprises 70 IR publications. It starts in February 1998 when the MPC reached its fully operational size and began publishing forecasts on a consistent basis. During our sample, the IR was published one week after the MPC policy rate decision was publicly announced but before the publication of the minutes that explained the decision. Our sample ends in May 2015, after which the Bank moved to a new publication schedule where the IR is published at the same time as the policy rate and minutes, which makes isolating the impact of communication difficult.⁴

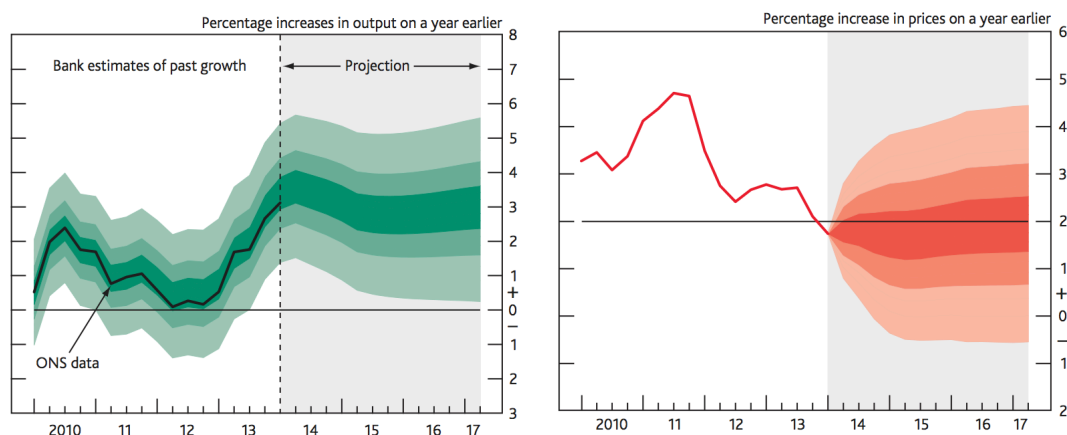


Figure 1: Numeric Information: Fan Charts

Notes: The left-hand figure shows the May 2014 Inflation Report fan chart for the GDP growth projection based on market interest rate expectations and other policy measures as announced. The right-hand figure shows the analogous fan chart associated with the CPI inflation forecast. Source: Bank of England.

The Inflation Report is a rich source of information. Its headline information is modal forecasts for GDP and inflation over the following two years (and from August 2004 onwards for the following three years), as well as distributional information around those modes in the form of a variance and skew. It presents these projections in the form of fan charts, as illustrated in Figure 1.⁵ The IR also contains extensive narrative information in the form of written text. In Section 4 we describe how we quantify both the numeric and narrative information in the IR.

The IR report is also important for what it does *not* contain, i.e. formal forward guidance. It is primarily a vehicle for delivering the MPC’s views on the development of economic conditions, and it does not provide explicit signals about how it will react

⁴During our sample, the MPC also engaged in other forms of communication like member speeches, but these do not systematically fall on IR publication days.

⁵The numeric values for forecast distributions are made publically available on the IR website alongside the graphical representation as fan charts.

to those developments. We therefore view it as an ideal setting to study how a central bank’s transmitting information solely about conditions can generate market news. Such an exercise would, for example, not be possible in the US since the Federal Reserve does not publish contemporaneous Greenbook forecasts.⁶

In summary, IR publication has several advantages. It is published according to an exogenously set schedule; there is no policy rate announcement on the same day as IR publication that could confound its market impact; it contains news just on the outlook; and it has a rich set of potential signals that allow one to dig deeply into the information it conveys.

2.2 Event study description and yield curve data

In order to study the impact of IR publication on market interest rates, we use an event study approach that has now become quite standard in the literature. Suppose the IR is published on day t in month $m(t)$. Figure 2 depicts the timeline of events. The month $m(t)$ policy rate $i_{m(t)}$ is announced seven days prior to IR publication. At the close of market trading on day $t - 1$, the market information set is I_{t-1}^{MK} and we observe some market interest rate. At the close of trade on day t , the market information set is I_t^{MK} and we observe a new market interest rate. The assumption in the event study literature is that $I_t^{\text{MK}} \setminus I_{t-1}^{\text{MK}}$ is generated entirely by IR publication, and that this additional information in turn generates the observed change in market interest rates on day t . One can then use the absolute observed change in market interest rates on day t to assess the news contained in IR publication. While the literature increasingly uses tight, intra-day windows around communication events,⁷ we use daily changes since it may take markets longer to incorporate the length and complexity of the IR.

For our analysis, we use daily data during the period 01/01/1998 to 31/07/2015 on four different maturity market rates derived from UK government bond prices: the one-year spot rate; the three-year forward rate; the five-year forward rate; and the five-year ahead, five year forward rate (equivalent to the average forward rate five to ten years ahead). We use nominal rather than real rates because obtaining reliable short-run real rates during our sample in the UK is difficult. In Section 6.2 we repeat the analysis on medium- and long-run real rates, and show results that are very similar to those with nominal rates.

⁶We take a broader view of what Romer and Romer (2000) call the information effect. In their sample, the Fed engaged in very limited public communication, and so the observed policy rate change was one of the main means by which markets could infer Fed forecasts. In this paper, the “information effect” includes any systematic market reaction to communication about economic fundamentals via any medium.

⁷Gürkaynak et al. (2005), Nakamura and Steinsson (2013), and Gertler and Karadi (2015) all use high-frequency identification relying on news about monetary policy in a 30-minute window surrounding scheduled Federal Reserve announcements.




Figure 2: Event Study Time Line for IR Publication on Day t

Under standard asset pricing theory, we can write forward interest rates as a combination of an expectation and term premium. If investors were unconcerned about the risks around future interest rates, the term structure of interest rates—the ‘yield curve’—should equal the expected path for short-term interest rates. This is often called the ‘pure expectations hypothesis’ and arises from the ability of investors to choose between buying a long-term bond or investing in a series of short-term bonds. In practice, however, market interest rates deviate from the pure expectations hypothesis, with any additional return referred to as the ‘term premium’, which we denote by TP.

A general expression for the k -month ahead forward rate on day t is therefore

$$f_{k,t} = \mathbb{E}[i_{m(t)+k} \mid I_t^{MK}] + \text{TP}_k(I_t^{MK}), \quad (1)$$

and our four particular market rates can be expressed as⁸

1. 1-year spot rate:

$$i_{0:12,t} = \frac{i_{m(t)} + \sum_{i=1}^{11} \mathbb{E}[i_{m(t)+i} \mid I_t^{MK}]}{12} + \text{TP}_{0:12}(I_t^{MK}) \quad (2)$$

2. 3-year forward rate:

$$f_{36,t} = \mathbb{E}[i_{m(t)+36} \mid I_t^{MK}] + \text{TP}_{36}(I_t^{MK}) \quad (3)$$

3. 5-year forward rates:

$$f_{60,t} = \mathbb{E}[i_{m(t)+60} \mid I_t^{MK}] + \text{TP}_{60}(I_t^{MK}) \quad (4)$$

⁸Here we have expressed nominal rates in terms of expectations formed at a monthly frequency for notational convenience; in practice, the forward rates are computed using a notional instantaneous rate of interest, and the 1-year spot and 5-year, 5-year rates are integrals under the curve corresponding to these instantaneous rates.

4. 5-year, 5-year rates:

$$f_{60:120,t} = \frac{\sum_{i=60}^{119} \mathbb{E}[i_{m(t)+i} | I_t^{MK}]}{60} + \text{TP}_{60:120}(I_t^{MK}) \quad (5)$$

In our theoretical and empirical analysis below (although not in the simple event study of this section), we distinguish between the effect that IR publication has on expectations and term premiums separately. One common way to perform an empirical decomposition of the yield curve into these two components is to use an affine term structure model. Some of these models use only the past behavior of the market yield curve to estimate the decomposition, whereas others supplement that with survey or other additional data on expectations. The specification of the model can lead to quite large differences in the estimates. In our analysis, therefore, we use an average of four differently-specified models, two of which supplement the yield curve data with survey information.⁹

Table 1 shows the contribution of each component to explaining the overall variance in yields on IR publication days in our sample. The term premium plays an increasingly important role in accounting for movements in interest rates at longer horizons, and is the primary driver of changes in the five-year, five-year forward rate.¹⁰

Table 1: Variance Decomposition of Market Interest Rate Changes

	Total Var	Var(Exp)	Var(TP)	2 x Cov
1 Year Spot	0.0032	0.0024	0.0001	0.0007
	100	75	3	22
3 Year Forward	0.0066	0.0037	0.0009	0.0020
	100	56	14	30
5 Year Forward	0.0050	0.0026	0.0015	0.0009
	100	52	29	19
5 Year, 5 Year Forward	0.0039	0.0018	0.0023	-0.0002
	100	47	59	-6

Notes: This table reports the variance decomposition of different yields by expectation and term premium components on our 70 IR release days. Var(Exp) is the variance explained by expectations; Var(TP) is the variance explained by term premiums; and Cov is the covariance between the components.

2.3 Event study results

For the event study, we classify each day in our sample of market interest rates according to whether (1) an IR is released; (2) a policy decision from the MPC is announced; (3) an

⁹Specifically we use the benchmark and survey models in Malik and Meldrum (2016), the model in Vlieghe (2016), and the model in Andreasen and Meldrum (2015).

¹⁰One reason for the relatively low variance of the 1-year spot rate is that our sample includes a period in which short-maturity interest rates were at the effective lower bound.

MPC member makes a public speech; (4) minutes from MPC meetings are released; or (5) none of the above. We then plot kernel densities for each of these five categories in Figures 3a-3d. For one-year spot and three-year forward rates, the IR release dates appear to generate a consistently large amount of news relative to other forms of communication. For longer-horizon rates there is more similarity in the impact across communication events, but there is a mass of large tail moves in interest rates on IR publication dates not present on other communication event dates. In Appendix A, we conduct a more formal assessment of the relative market impact of the IR using regression analysis, and find a similar pattern as in the kernel densities.

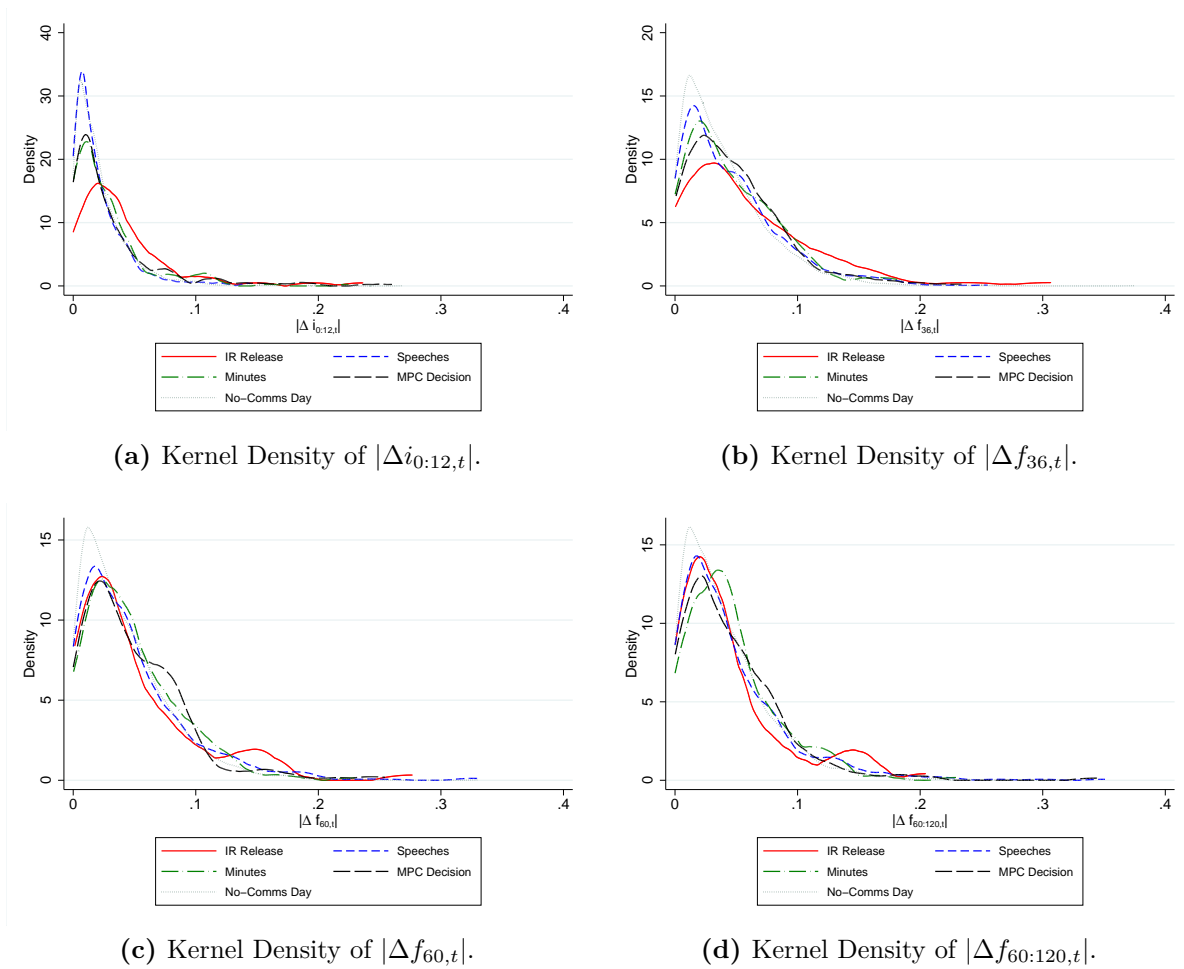


Figure 3: Kernel Densities of Yield Changes by Type of Communication

Notes: These figures show the kernel-density distribution of changes in expected interest rates at different maturities. We use the Epanechnikov kernel, and set the half-width to the value that would minimize the mean integrated squared error if the underlying distribution were Gaussian.

One concern might be that long-run rate movements on IR dates are too small to have policy relevance, so that explaining them is not of first-order importance. One way

of assessing the importance of long-run rate moves is by examining the fractions of IR publication dates on which there are large yield moves, as shown in Table 2. For all yields, and despite revealing no policy actions, a quarter or more of IR publication dates lead to at least a five-basis point change, with the proportion growing to nearly a half for three-year forward rates. Moreover, movements of ten basis points are also not uncommon, and there are even occasional twenty basis point moves. All of which suggests that there is indeed meaningful variation in longer maturity rates in our sample.

Table 2: Magnitude of Rate Moves on IR Days

Asset	≥ 5 bps	≥ 10 bps	≥ 20 bps
$ \Delta i_{0:12,t} $	0.11	0.02	0.00
$ \Delta f_{36,t} $	0.33	0.08	0.00
$ \Delta f_{60,t} $	0.34	0.08	0.01
$ \Delta f_{60:120,t} $	0.32	0.07	0.01

Notes: This table shows the fraction of IR publication dates in our sample on which there are yield moves of at least 5bps, 10bps and 20bps.

3 Theoretical Channels for the Information Effect

While the event study results above show that some information contained in the IR generates market news, an important question is *which* information is driving the response, especially at longer maturities. In this section we present a simple model of how news about the outlook for economic conditions from the central bank can affect market interest rates. The formal proofs are contained in Appendix B. In particular, we include a stochastic form of uncertainty and show that central bank signals about uncertainty shocks can have an increasing impact on rates at greater maturities under sufficiently high persistent of the underlying volatility.

This theoretical framework enables us to distinguish three potential channels through which central bank communication could affect interest rates. We model two information effect channels; one concerning level expectations and the other concerning signals about uncertainty. The third channel, unmodelled here, is an investor demand channel as in Hanson and Stein (2015). The uncertainty channel is distinguished by the correlation structure it induces between signals across yields and their components.

3.1 Model Environment

As in macroeconomic models with forward-looking monetary policy, the central bank is assumed to set nominal interest rates as a function of forecasts of future economic

conditions. We denote month- m economic conditions as $\omega_m \in \mathbb{R}$, $\mathbb{E}[\omega_{m+h} \mid I_m^{\text{CB}}]$ as the central bank's h -period-ahead forecast of economic conditions given its month m information set I_m^{CB} , and ϕ as the central bank's reaction coefficient. The short-term nominal interest rate in month m is therefore

$$i_m = \phi \mathbb{E}[\omega_{m+h} \mid I_m^{\text{CB}}] + \epsilon_m \quad (6)$$

where the monetary policy shock is assumed to be $\epsilon_m \sim \mathcal{N}(0, \sigma_\epsilon^2)$ and uncorrelated across months. This could be expanded to a vector of state variables $\boldsymbol{\omega}_m$ that included, for example, the expected output gap, expected inflation, and the equilibrium real interest rate, along with an associated vector of reaction coefficients $\boldsymbol{\phi}$. The analysis below would then apply to each component separately, with the overall effect of central bank communication then being the sum over the effect on each component. Moreover, we ignore the effective lower bound on interest rates but we shall return to this in section 6.

We assume economic conditions ω_m evolve according to an AR(1) process

$$\omega_m = \rho \omega_{m-1} + \underbrace{\mu_m + \varepsilon_m}_{= v_m} \text{ where } 0 < \rho \leq 1. \quad (7)$$

The shock to economic activity in month m , which we denote v_m , is comprised of two components. The first is μ_m which is drawn independently every month from $\mu_m \sim \mathcal{N}(0, s^2)$. We assume the central bank obtains information that allows it to forecast the level of μ_m (details below). We therefore view uncertainty in μ_m as reducible with improvements in forecasting ability or new information. If the central bank's forecasting ability is high enough, we can even treat μ_m as fully observable.

In contrast, the central bank cannot forecast the level of the second component of the shock, ε_m . This represents the fundamental, or irreducible, uncertainty in the economy. We assume it is drawn independently each month from $\varepsilon_m \sim \mathcal{N}(0, \sigma_m^2)$, where the amount of fundamental uncertainty in the economy σ_m^2 is stochastic.

We follow much of the finance literature and model σ_m^2 as

$$\log \sigma_m^2 = \rho_\sigma \log \sigma_{m-1}^2 + (1 - \rho_\sigma) \log \sigma_0^2 + u_m \text{ where } 0 < \rho_\sigma \leq 1. \quad (8)$$

Here σ_0^2 is some baseline level of uncertainty and we assume $u_m \sim \mathcal{N}(0, \sigma_u^2)$. This assumption generates a lognormal distribution for σ_m^2 . It is important to note that, while the level of ε_m is not forecastable, the level of uncertainty σ_m^2 is forecastable given information about shocks $u_{m'}$ for $m' < m$. This is an important mechanism that will lead to changes in long-run interest rates in the model.

Central Bank Information Set. In every month m , we assume that the central bank

observes ω_m perfectly. In addition, it observes some collection of signals correlated with the means of the forecastable shocks over the forecast horizon $\mu_{m+1}, \dots, \mu_{m+h}$. Moreover, these signals accumulate every month. So, for example, suppose in month m the central bank observes a first signal about μ_{m+h} . Then, in month $m+1$, it may observe a second signal about μ_{m+h} that it combines with its first signal to form a new, more precise, forecast, and so on through month $m+h$, when we assume that the shock v_{m+h} is fully revealed through observation of ω_{m+h} . This is consistent with the idea that the central bank revises both its mean forecast of future conditions, as well as the forecast uncertainty around that mean, as time proceeds. Rather than model the precise details of the signal structure, we summarize the signals the central bank observes in terms of the updated belief on μ_{m+i} that they induce. More specifically, we assume that

$$\mu_{m+i} | I_m^{\text{CB}} \sim \mathcal{N}\left(\widehat{\mu}_{m+i,m}^{\text{CB}}, (s_{m+i,m}^{\text{CB}})^2\right) \text{ for } i = 1, \dots, h. \quad (9)$$

where $\widehat{\mu}_{m+i,m}^{\text{CB}}$ is the central bank's point estimate of μ_{m+i} in month m , and $s_{m+i,m}^{\text{CB}}$ captures the forecast uncertainty around that mean.¹¹ We assume that the central bank forms more precise forecasts as it accumulates additional signals over time, so $s_{m+i,m+1}^{\text{CB}} \leq s_{m+i,m}^{\text{CB}}$. We also assume the central bank is Bayesian, so $\mathbb{E}[\widehat{\mu}_{m+i,m+k}^{\text{CB}} | \widehat{\mu}_{m+i,m}^{\text{CB}}] = \widehat{\mu}_{m+i,m}^{\text{CB}}$ for all $i > k > 1$.

While the central bank does not receive information that improves its forecast of the level of the second component of economic shocks ε_m , we assume that in every month m the central bank observes time-varying fundamental uncertainty σ_{m+h}^2 perfectly. This is clearly a strong assumption, as it implies that in every month m the central bank knows perfectly the sequence u_{m+1}, \dots, u_{m+h} . We could relax this assumption and instead allow the central bank to receive noisy signals of the future shocks to fundamental uncertainty. However, this increases notational clutter without leading to any fundamental new insights.

In summary, the central bank's information set I_m^{CB} consists of:

1. ω_m or, equivalently, the entire history of shocks v up to month m .
2. Signals for each μ over the forecast horizon that leads to beliefs as specified in (9).
3. σ_{m+h}^2 or, equivalently, the sequence of fundamental uncertainty shocks u up to month $m+h$.

Market Information Set. Suppose the inflation report (IR) is published on day t , and let $m = m(t)$ be the month in which day t falls. I_t^{MK} is defined as the market's information

¹¹Since μ_{m+i} is assumed to be drawn from a normal distribution, we can assume that the central bank observes normally distributed signals to arrive at normally distributed posterior beliefs.

set on day t , and we assume that $I_m^{\text{CB}} = I_t^{\text{MK}}$, so that the information contained in the IR is a sufficient statistic for whatever else the market knows about the economy on day t . On day $t - 1$, we assume the market has observed ω_m , as well as signals on the values of μ within the forecast horizon that lead to beliefs

$$\mu_{m+i} \mid I_{t-1}^{\text{MK}} \sim \mathcal{N}\left(\widehat{\mu}_{m+i,t-1}^{\text{MK}}, (s_{m+i,t-1}^{\text{MK}})^2\right) \text{ for } i = 1, \dots, h. \quad (10)$$

These signals may have come from the previous IR publication, or from independent market forecasts. In any case, we assume that the IR contains relevant additional information in the sense that $s_{m+i,t-1}^{\text{MK}} > s_{m+i,t}^{\text{MK}} = s_{m+i,m}^{\text{CB}}$: after IR publication, the market updates its beliefs from $\widehat{\mu}_{m+i,t-1}^{\text{MK}}$ to $\widehat{\mu}_{m+i,t}^{\text{MK}} = \widehat{\mu}_{m+i,m}^{\text{CB}}$ and has lower forecast uncertainty about the value of μ_{m+i} for all $i = 1, \dots, h$.

On day $t - 1$, we assume that the market has observed the sequence of shocks to fundamental uncertainty u up to month $m + h - 1$ only. Thus IR publication reveals u_{m+h} , which gives the market a new source of information for predicting σ_{m+i}^2 for all $i > h$. This is one particular way of modeling the idea that the IR contains news on fundamental uncertainty shocks. In a more complex model, the market would hold signals on day $t - 1$ about u_{m+1}, \dots, u_{m+h} that IR publication then added to, as we have assumed for the μ terms, but the basic idea would be the same as in this setup.¹²

In summary, the markets's day $t - 1$ information set I_{t-1}^{MK} consists of:

1. ω_m .
2. i_m , the month m policy rate published one week prior to the IR.
3. Signals for each μ in the forecast horizon that leads to beliefs as specified in (10).
4. σ_{m+h-1}^2 or, equivalently, the sequence of fundamental uncertainty shocks u up to month $m + h - 1$.

3.2 Expectations Channel

To assess the impact of information in the IR publication on interest rates, we distinguish between two separate channels. Recall that the k -month-ahead forward rate is given in (1) by

$$f_{k,t} = \mathbb{E}\left[i_{m+k} \mid I_t^{\text{MK}}\right] + \text{TP}_k(I_t^{\text{MK}}).$$

¹²Another channel we do not consider is that the IR could deliver indirect information about ϵ_m , which is itself persistent. This variable in the interest rate rule can be thought of as capturing the central bank's preferences, which indeed might feature autocorrelation. However this is unlikely to be true five or ten years ahead, when the membership of the MPC is likely to have changed completely.

Any observed change in $f_{k,t}$ must arise from a change in either expected future nominal rates or the term premium. We refer to the *expectations channel* as the effect of IR publication on expected future interest rates, and define

$$EXP_m(k) \equiv \mathbb{E}[i_{m+k} \mid I_t^{MK}] - \mathbb{E}[i_{m+k} \mid I_{t-1}^{MK}].$$

Using the policy rule as specified in (6), we can express the k -month-ahead policy rate as

$$i_{m+k}(I_{m+k}^{CB}) = \phi \mathbb{E}[\omega_{m+k+h} \mid I_{m+k}^{CB}] + \epsilon_{m+k}. \quad (11)$$

We can expand the ω_{m+k+h} term as

$$\omega_{m+k+h} = \rho^{k+h} \omega_m + \sum_{i=1}^{k+h} \rho^{k+h-i} v_{m+i} \quad (12)$$

and so

$$\mathbb{E}[\omega_{m+k+h} \mid I_{m+k}^{CB}] = \underbrace{\rho^{k+h} \omega_m + \sum_{i=1}^k \rho^{k+h-i} v_{m+i}}_{= \omega_{m+k}} + \sum_{i=k+1}^{k+h} \rho^{k+h-i} \widehat{\mu}_{m+i, m+k}^{CB}. \quad (13)$$

In month $m+k$, the central bank observes ω_{m+k} by assumption. The final term in (13) is the central bank's forecasts for the shocks that will hit the economy within the month $m+k$ forecasting horizon.

While the market in month m can observe ω_m , IR publication provides news in two senses. First, it provides signals on the v terms within the month- m forecasting horizon, which will feed into the market expectations for ω_{m+k} for all k due to the autoregressive process specified in (7). Second, it provides indirect signals on $\widehat{\mu}_{m+i, m+k}^{CB}$ for $i = 1, \dots, h$. Because the central bank is Bayesian, the market's best guess for $\widehat{\mu}_{m+i, m+k}^{CB}$ after observing the IR is $\widehat{\mu}_{m+i, m}^{CB}$. The overall effect is described in the following result.

Proposition 1 *The impact of IR publication in month m on interest rates through the expectations channel is*

$$EXP_m(k) = \alpha \rho^k \sum_{i=1}^h \rho^{h-i} (\widehat{\mu}_{m+i, m}^{CB} - \mathbb{E}[\widehat{\mu}_{m+i, m}^{CB} \mid \widehat{\mu}_{m+i, t-1}^{MK}, i_m])$$

Moreover, if $\rho < 1$ then $|EXP_m(k)|$ is strictly decreasing in k and $\lim_{k \rightarrow \infty} |EXP_m(k)| = 0$, while if $\rho = 1$ then $|EXP_m(k)|$ is independent of k .

On day $t-1$ the market has its own forecasts of future conditions $\widehat{\mu}_{m+i, t-1}^{MK}$ and the current

policy rate i_m as relevant indicators of $\hat{\mu}_{m+i,m}^{\text{CB}}$ for $i = 1, \dots, h$,¹³ which are in turn the best predictors of (13) following IR publication. The size of the expectations channel effect on market interest rates depends on the degree to which the market's views on $\hat{\mu}_{m+i,m}^{\text{CB}}$ change after observing the IR.

For our purposes, more relevant is how the expectations channel operates at different points on the forward yield curve. Proposition 1 distinguishes between two cases. The first is when $\rho < 1$ and (7) is a stationary process, as is plausible for macro variables like inflation and GDP growth. Here the size of the expectations channel effect is declining in maturity of the forward rate. The reason is simply that the influence of current shocks on ω_{m+k} is declining in k due to mean reversion in the autoregressive process. Put another way, additional signals on v_{m+1}, \dots, v_{m+h} are less and less useful for predicting ω_{m+k} as k grows. In the very long run, as k grows large, the additional signals contain vanishingly little information and the expectations channel is zero.

The other case is when $\rho = 1$ and (7) is a unit-root process. For example, there is evidence that shocks to the natural rate of interest are highly persistent (Laubach and Williams 2003). In this case, long-run expectations of ω_{m+k} would react to updated beliefs of those shocks. Such information in the IR publication would, therefore, induce a one-for-one movement at all maturities in the forward yield curve. This observation in turn leads to the result

Corollary 1 *Any move in long-run forward rates due to the expectations channel must generate at least an equivalent move in short-run forward rates. The impact of the expectations channel on long-run rates is maximal when $\rho = 1$, in which case the impact on short-run rates is identical.*

This implication of the model is important for empirically evaluating theories of long-run yield curve movements that rely on central bank communication's shifting market expectations of some economic fundamental. This mechanism is plausible only when the fundamental is highly persistent. Moreover, in this case the same signals that move long-run expectations must necessarily also move short-run expectations, and by essentially the same magnitude.¹⁴ We assess the evidence for such effects in the empirical analysis.

3.3 Uncertainty Channel

We now analyze the effect on forward nominal interest rates that information on uncertainty contained in the IR publication —the *uncertainty channel*—could have and

¹³ i_m provides information since it is a function of $\hat{\mu}_{m+i,m}^{\text{CB}}$ for $i = 1, \dots, h$. It does not perfectly reveal the central bank's information set since it also depends on the stochastic shock ϵ_m .

¹⁴Another possibility is that the central bank could send separate short- and long-maturity-specific information that led market participants to update their views on each end of the yield curve independently, but this is not the case in the Inflation Report.

define

$$\text{UNC}_m(k) \equiv \text{Var}[i_{m+k} \mid I_t^{\text{MK}}] - \text{Var}[i_{m+k} \mid I_{t-1}^{\text{MK}}].$$

Any news contained in the IR that only affects $\text{UNC}_m(k)$ and not $\text{EXP}_m(k)$ must by definition affect only the term premium. Moreover, a large macrofinance literature provides theoretical foundations for why uncertainty about economic conditions indeed affects the term premium (e.g. Bansal and Shaliastovich 2013, Martin 2013). The purpose of our model is to disentangle the various sources of news contained in central bank communication, and their effects at different points on the yield curve. As such, we do not explicitly model the link between changes in $\text{UNC}_m(k)$ and changes in the term premium, but maintain the reasonable assumption that a change in $\text{UNC}_m(k)$ generates a corresponding change in the term premium of the k -period-ahead forward rate.

In our model there is an important distinction between short- and long-run effects of the uncertainty channel, so we analyse each in turn.

3.3.1 Short-run effects

By short-run effects, we mean the effect on forward rates within the central bank's forecast horizon, i.e. $k < h$. The relevant quantity that IR publication could affect here is the market's perceived variance of the central bank's month $m+k$ expectation of economic conditions in month $m+k+h$, as in equation (13) above.¹⁵ There are two relevant sources of news in the IR. First, the market receives additional signals about the shocks that will hit economic conditions, which reduces variance in the forecasts of v_{m+i} for $i \leq k$. Second, since the market learns $\hat{\mu}_{m+i,m}^{\text{CB}}$ from the IR, it can also better predict $\hat{\mu}_{m+i,m+k}^{\text{CB}}$ for $k < i \leq h$ since the signals that go into forming $\hat{\mu}_{m+i,m}^{\text{CB}}$ also enter $\hat{\mu}_{m+i,m+k}^{\text{CB}}$. The uncertainty channel for $k < h$ is therefore quite straightforward to derive, and we state the following without proof.

Proposition 2 *Suppose that $k < h$. Then the uncertainty channel is*

$$\text{UNC}_m(k) = \alpha^2 \left\{ \begin{array}{l} \sum_{i=1}^k \rho^{2(k+h-i)} \left((s_{m+i,t}^{\text{MK}})^2 - (s_{m+i,t-1}^{\text{MK}})^2 \right) + \\ \sum_{i=k+1}^h \rho^{2(k+h-i)} \left(\text{Var}[\hat{\mu}_{m+i,m+k}^{\text{CB}} \mid I_t^{\text{MK}}] - \text{Var}[\hat{\mu}_{m+i,m+k}^{\text{CB}} \mid I_{t-1}^{\text{MK}}] \right) \end{array} \right\}.$$

Both sources of news serve to reduce the variance of future interest rates, and so $\text{UNC}_m(k) < 0$. However, the dependence of $\text{UNC}_m(k)$ on k is difficult to pin down at this level of abstraction. Characterizing this requires determining conditions under which the reduction in variance in future cyclical shocks can be compared to the reduction in variance in the central bank's future beliefs about cyclical shocks. Rather than provide such results, we interpret Proposition 2 as saying that the publication of the IR impacts

¹⁵We ignore the variance in future rates arising from ϵ_m since IR publication does not affect this.

the short-run term premium due to a reduction in uncertainty about short-run nominal interest rates, but that the effect need not have a clear relationship with the maturity of forward rates within the central bank's forecasting horizon.

3.3.2 Long-run effects

We next consider the impact of uncertainty news on forward rates outside of the forecast horizon, i.e. $k \geq h$. The crucial difference in this case is that the variance of future nominal rates now depends on the variance of ε_{m+i} for $i \in \{h, \dots, k\}$, which the market does not know on day $t - 1$. Instead, it forms a forecast of these terms with variance σ_{m+h-1}^2 , which is in I_{t-1}^{MK} . IR publication then reveals σ_{m+h}^2 (or, equivalently, u_{m+h}), which leads the market to update its forecast on future fundamental uncertainty and thus its view on future nominal rate volatility. This effect is absent in the short run ($k < h$) since the short-run fundamental uncertainty in the economy is known prior to IR publication.

The extent to which learning σ_{m+h}^2 affects market forecasts of long-run fundamental uncertainty depends on ρ_σ , the persistence of shocks to fundamental uncertainty in the model defined in (8). With low persistence, the effect of u_{m+h} dies away quickly and forecasts of the variance of ε_{m+i} are relatively unaffected as i grows past h . With high persistence, the opposite is true. In our next result, we characterize an upper bound on the uncertainty channel in the long run by considering the limiting behavior as ρ_σ approaches 1 and (8) becomes a unit-root process.

Proposition 3 *Suppose that $k \geq h$. Then the uncertainty channel satisfies*

$$\lim_{\rho_\sigma \rightarrow 1} \text{UNC}_m(k) = \alpha^2 \left\{ \begin{array}{l} \sum_{i=1}^h \rho^{2(k+h-i)} \left((s_{m+i,t}^{\text{MK}})^2 - (s_{m+i,t-1}^{\text{MK}})^2 \right) + \\ \sum_{i=h}^k \rho^{2(k+h-i)} \exp\left(\frac{(i-h)\sigma_u^2}{2}\right) \left[\sigma_{m+h}^2 - \sigma_{m+h-1}^2 \exp\left(\frac{\sigma_u^2}{2}\right) \right] \end{array} \right\}.$$

Moreover,

$$\lim_{k \rightarrow \infty} \lim_{\rho_\sigma \rightarrow 1} |\text{UNC}_m(k)| = \infty.$$

whenever $\sigma_{m+h}^2 \neq \sigma_{m+h-1}^2 \exp\left(\frac{\sigma_u^2}{2}\right)$.

The effect in the first line of the expression for $\text{UNC}_m(k)$ is also present in the short run, and represents a reduction in uncertainty due to additional signals on the cyclical shocks that will hit economic conditions over the next h months. Its value is declining in k whenever $\rho < 1$ because, as with the expectations channel, the impact of short-run shocks fades away in the long run in a stationary autoregressive process.¹⁶

¹⁶The short-run uncertainty channel also depended on the change in the variance of the central bank's future beliefs on economic conditions. This effect is absent in the long run because IR publication provides no news on $\hat{\mu}_{m+i,m+k}^{\text{CB}}$ for $k > h$.

The effect in the second line is specific to the long run, and reflects the impact of revised forecasts of future fundamental uncertainty. It can be positive or negative depending on the sign of $\sigma_{m+h}^2 - \sigma_{m+h-1}^2 \exp\left(\frac{\sigma_u^2}{2}\right)$, which captures whether IR publication increases or decreases the expected value of σ_{m+i}^2 for $i > h$.¹⁷ As k grows, the absolute value of this effect also grows. This is because the number of shocks to fundamental uncertainty between months $m+h$ and $m+k$ increases in k . So fundamental uncertainty, and therefore the impact of forecast revisions, accumulates as one moves further out in the yield curve. In the limit as k grows very large, IR publication induces an unboundedly large absolute change in the expected variance of the policy rate outside of the measure zero event $\sigma_{m+h}^2 \neq \sigma_{m+h-1}^2 \exp\left(\frac{\sigma_u^2}{2}\right)$. While we have not modeled the precise mapping between the uncertainty channel and changes in the term premium, our model strongly suggests that news contained in central bank communication relevant for forecasting fundamental uncertainty can have a large impact on long-run term premiums.¹⁸

The case of high ρ_σ is an empirically plausible assumption; Bansal and Shaliastovich (2013) estimate a stochastic volatility model similar to ours, and find the persistence in uncertainty shocks to be well above 0.9. We expect a similar mechanism to operate at the long run, and similar results to arise, in situations with a low value of ρ_σ if the central bank alternatively provides signals on the baseline uncertainty (σ_0^2) rather than the innovations to uncertainty.

3.4 Distinguishing theories of long-run rate movements

A central motivation of this paper is to ascertain why central bank communication moves long-run interest rates, and so before the empirical analysis we conclude this section by summarizing competing explanations and how they can be distinguished by data.

Expectations channel. Central bank communication changes modal expectations of long-run economic conditions. This is the channel emphasized in Nakamura and Steinsson (2017), and is plausible only when the central bank transmits information about shocks to highly persistent variables like the equilibrium real interest rate. Proposition 1 shows that such information should change short-run expectations at least as much as it does long-run expectations.

Uncertainty channel. Central bank communication changes the perceived variance of interest rates by transmitting information on persistent uncertainty. As in Proposition

¹⁷The change in future expected value does not depend simply on $\sigma_{m+h}^2 - \sigma_{m+h-1}^2$ because σ_{m+i}^2 is lognormally distributed.

¹⁸Martin and Ross (2018) present a non-Gaussian bond-pricing framework in which a similar effect could arise if signals altered transition probabilities between persistent states of the world. Cieslak and Schrimpf (2018) discuss a similar channel (without formally modeling it), and the idea is also consistent with Bansal and Shaliastovich (2013) and Ellison and Tischbirek (2018).

3, this information should have its most prominent effect on long-run rates and we expect this channel to operate through the term premium.

Investor Demand channel. This channel is not present in our model, but is modelled by Hanson and Stein (2015). In that model, the effect on long-run rates comes from a change in demand from yield-oriented investors who react to monetary news that affects short-run expectations by trading longer-term debt to maximize the yield in their portfolios. The main impact on long rates comes via the term premium but is driven by identical information to that driving the change in short-run rates.

The first two channels rely on the central bank providing direct information relevant for long-run beliefs. While the first is present in the literature, the uncertainty channel is novel. The demand channel instead relies on information relevant for short-run beliefs that then propagates to the long run through trading activity.

While not explicitly included in the model, each Inflation Report event can be thought of as the central bank sending a vector of signals to the market, and in the next section we will explicitly construct an empirical proxy for this vector. In fact the model is silent about which signals are responsible for generating which channels, and on whether the same signal could simultaneously convey information on the level and variance of future economic conditions. For example, one signal contained in each IR is the change in the inflation forecast at the forecast horizon relative to the previous IR. From this signal, market participants may update their expectations for the inflationary state of the economy, but they may also, as a result of a large change (or non-change), update their views of uncertainty going forward. More generally, we view the total effect of any given signal the central bank sends as potentially coming from all three channels, and our empirical exercise does not attempt to argue that some channels are present while others are not. Instead, the goal is to identify which channel appears most responsible for the long-run rate reactions we observe after IR publication.

The main distinguishing feature of the uncertainty channel is that, insofar as a signal conveys information about persistent shocks to fundamental uncertainty, its effect on interest rates should be higher for long-run maturities than for short-run maturities as per the arguments of Proposition 3. In contrast, the signals that generate the other two channels must necessarily move both short- and long-run interest rates. This is argued explicitly for the expectations channel in corollary 1. In the demand channel, the long-run term premium necessarily moves in tandem with short-run expectations affected by monetary news. In contrast, in our model the long-run term premium can move independently of short-run expectations if the signals that generate the expectations channel differ from those that generate the uncertainty channel. Therefore, in our empirical analysis we propose various tests of whether the signals that move long-run interest rates

correspond to those that move short-run rates.

4 Measuring Inflation Report Signals

Using our theoretical framework to interpret interest rates moves on IR publication days requires us to measure the vector of signals that the IR contains. Since we do not model these signals explicitly, our approach is to be as flexible as possible and build a high-dimensional set of measures from both the numeric and narrative data in the IR, each of which in principle can convey news to markets. Here the richness of the information in the Report is crucial, as it allows us to study with a great deal of granularity the information that drives different maturities in the market data and the different components (expectation and term premium) of the asset price response.

4.1 Numeric information

As described in Section 2.1, the numeric forecast information in the IR on day t are the forecasts for GDP growth and inflation, which form part of $\mathbb{E}[\omega_{m+h} \mid I_m^{\text{CB}}]$, and distributional information around them. In our analysis we use a set of 15 core numeric signals contained in each IR published in month m .

We take the modal growth and inflation projections, denoted g_m^{CB} and π_m^{CB} respectively. We also include the variances, $\text{Var}(g_m^{\text{CB}})$ and $\text{Var}(\pi_m^{\text{CB}})$, and skews, $\text{Skew}(g_m^{\text{CB}})$ and $\text{Skew}(\pi_m^{\text{CB}})$. Since 1998, these forecasts have been consistently conditioned on the path for the policy rate (called ‘Bank Rate’) implied by market interest rates. While projections are provided for each quarter over the forecast period, in our analysis we focus our attention on the projections at the two-year horizon as that is the horizon that has tended to be focused on in the Bank’s monetary policy communication as the one most relevant for the current stance of policy.

Rather than the rate of GDP growth g_m^{CB} , the potentially more relevant variable for interest rate expectations is the MPC’s view of the h -month ahead output gap $\tilde{y}_{m,h}^{\text{CB}}$. It may be that investors infer the MPC’s view of the future output gap from its GDP growth forecast. To measure this, we construct an implied modal output gap using the MPC’s growth forecasts together with private-sector estimates of long-run potential growth.¹⁹

It is also important to capture not just the IR forecast levels but also how much these deviate from market expectations of the IR forecast, since the surprise component in the forecast should be the driver of any change in market interest rates. Ideally we

¹⁹Specifically, we grow the real GDP series implied by the forecast at the rate of long-run growth from Consensus Economics. We then pass the resulting series through a Baxter-King Bandpass filter to isolate movements between 2 and 36 quarters. The output gap estimate for the IR release in month m , $\tilde{y}_{m,h}^{\text{CB}}$, is the percentage deviation of the forecast level of real GDP from the BK-filtered trend series.

would compare each of the MPC’s forecast measures to equivalent expectations from the private sector. The Bank of England collects equivalent market forecasts ahead of each IR publication, and we denote the absolute difference between g_m^{CB} and its expected value as $\text{Supr}(g_m^{CB})$, and similarly $\text{Supr}(\pi_m^{CB})$. For the variance and skew variables, we do not have measures of private-sector expectations, and we define surprise as the difference between the current and previous value of the variable.

For the output gap surprises, we include the absolute deviation of the two-year-ahead implied output gap from the two-year-ahead implied output gap in the previous IR forecast, i.e. $|\tilde{y}_{m,24}^{CB} - \tilde{y}_{m-1,24}^{CB}|$. We also control for a like-for-like comparison by comparing the two-year-ahead (8 quarters) output gap forecast from the last IR with the 7-quarter-ahead output gap in the current forecast, i.e. $|\tilde{y}_{m,21}^{CB} - \tilde{y}_{m-1,24}^{CB}|$.

In total we obtain 15 variables associated with the numeric information in the IR communication published on day t , which we label \mathbf{q}_t . We divide these into two groups. The first contains variables that directly represent the level or news on the expectational component of the forecast: the modal growth and inflation forecasts; the surprise in these forecasts relative to market expectations; the output gap; and the two measures of the evolution of the output gap from the previous IR. We denote these variables $\mathbf{q}_t^{\text{EXP}}$. The remaining eight variables represent information on the distributions around these forecasts, and broadly measure uncertainty. We label these $\mathbf{q}_t^{\text{UNC}}$. Of course, the same set of signals can operate through multiple channels if they convey, either directly or indirectly, information on both the level and uncertainty around future economic conditions.

4.2 Narrative information

The narrative information in the IR consists of text broadly organized into two parts. A set of economics sections assess the current state of the economy, covering recent developments in and the near-term outlook for financial conditions, demand, supply, costs and prices. A forecast section describes the MPC’s forecasts, the risks around those forecasts, and the potential trade-offs for policy. The IR does not contain explicit forward guidance, understood as an explicit commitment to a future policy rule, or even a suggested response as to how future interest rates may evolve. However, this does not mean that it does not potentially contain information relevant for long-run rate expectations.

The narrative information can have important effects on future expectations and uncertainty for several reasons. First, there are many hundreds of hard and soft indicators of economic activity that the MPC regularly monitors, including surveys, disaggregate activity and inflation series, and information from regional agents. These indicators are all (potentially) endogenously related to each other and to the inflation and output forecasts

contained in the fan charts. The narrative in the IR provides the Bank of England’s views about the nature of these endogenous relationships, as well as what are the key drivers of the current forecasts. This can influence market views of likely future MPC forecasts. For example, the IR can reveal whether the inflation forecast is driven by persistent or transitory price movements.

Additionally, monetary policymakers in general, and the MPC specifically, do not publish quantitative views on the value of latent macroeconomic variables such as the equilibrium real interest rate. While an important driver of the policy action, the equilibrium real rate is an inherently elusive variable that depends on quantities such as the unobserved productive capacity of the economy about which there may be significant disagreement. In this context, the narrative may be the *only* way the MPC can signal its view of, and uncertainty about, the level of the real rate.

Another advantage of using narrative information is that it is inherently much richer than the numeric forecasts. This allows us to capture much more precisely the different signals that central banks send to markets. However this advantage also presents several statistical challenges that we must address. These include the issue of how to quantitatively represent the text for statistical analysis, as discussed in the rest of this section. It also includes the issues of the endogeneity of text to economic conditions and forecasts, and how to determine which topics are driving interest rate changes when the topic space is high dimensional, both of which are discussed in Section 5.

In the 70 Reports in our sample, there are 15,023 paragraphs. We first pre-process the text by removing all non-alphabetic terms, as well as extremely common words that are uninformative about the content such as ‘the’, ‘and’, and so on —so-called *stopwords*. We then stem each remaining term into its linguistic root using the Porter stemmer. Stems need not be an English word: for example, the stem of ‘inflation’ is ‘inflat’. Following these steps gives us 754,884 total terms in the dataset and 4,382 unique terms.

In order to reduce the dimensionality of the dataset we represent the text using a probabilistic topic model called Latent Dirichlet Allocation (LDA), first used in the economics literature by Hansen et al. (2018). Here we provide a high-level overview of the concept. Our estimation follows the same Markov Chain Monte Carlo procedure described in Hansen et al. (2018) and introduced by Griffiths and Steyvers (2004); we refer interested readers to those papers for full details.²⁰

LDA is a Bayesian factor model for discrete data. Suppose there are D documents (we treat each paragraph as a document, so $D = 15,023$) that comprise a corpus of texts with V unique terms (so here $V = 4,382$). The first important objects in LDA are K

²⁰A precursor to LDA is Latent Semantic Analysis (LSA), a non-probabilistic model that applies a singular value decomposition to the matrix of term counts in a corpus. Boukus and Rosenberg (2006) and Hendry and Madeley (2010) use LSA to assess the market impact of Fed and Bank of Canada communications, respectively, but do not propose tests for the information effect.

topics (i.e. factors), each of which is a probability vector $\beta_k \in \Delta^{V-1}$ over the V unique terms in the data. The choice of probability distributions is important since it allows the same term to appear in different topics with potentially different weights. Informally, one can think of a topic as a weighted word list that groups together those words that express the same underlying theme.

Each document can belong to multiple topics. Formally, each document d has its own distribution over topics given by θ_d (i.e. factor loadings). Informally, θ_d^k is topic k 's "share" of document d . The probability that any given word in document d is equal to the v th term is therefore $p_{dv} \equiv \sum_k \beta_k^v \theta_d^k$ and the overall likelihood is $\prod_d \prod_v p_{d,v}^{n_{d,v}}$ where $n_{d,v}$ is the number of times terms v appears in document d .

Importantly, LDA reduces the dimensionality of each document substantially. In the document-term matrix, documents live in a V -dimensional space. After estimating LDA, one obtains a representation of each document in terms of the (estimated) θ_d , which lives in the $K - 1$ simplex, and in general $K \ll V$. Importantly, though, LDA does not ignore any dimensions of variation in the raw term counts since the underlying topics are free to lie anywhere in the $V - 1$ simplex.

LDA places Dirichlet priors over the β and θ probability vectors, and the inference problem is to approximate their posterior distributions. The main model selection choice is the number of topics K . We use a model with $K = 30$, which provides a generally interpretable set of topics.²¹

Figure 4 represents the 30 topics that LDA estimates in our data and demonstrates that they are indeed interpretable. Topic 6, for example, appears to capture discussion of commodity prices; Topic 14 of the forecast; Topic 24 of financial markets; and so on. Since topics have no natural ordering, we define our own based on whether an IR is published during a cycle of rate increases (i.e. the previous rate change was an increase) or rate decreases (i.e. the previous rate change was a decrease). For each topic, we compute its average share of time in the IR during both cycle, and order topics based on the difference. Topic 0, about the pace of wage and labour cost growth, is most associated with an increasing rate cycle. While Topic 29, financial market conditions, which were of primary concern during the crisis, is most associated with a decreasing cycle.

While we estimate LDA at the paragraph level to exploit variation across thousands of examples of text, we are ultimately interested in the content of each IR in its entirety. We follow the procedure detailed in Hansen et al. (2018) to obtain the distribution over topics in the IR published on day t , which we denote θ_t . Since changes in topic coverage can also have potentially important market effects, we also include $\delta_t \equiv |\theta_t - \theta_{t-1}|$ to

²¹There is a well-known trade-off between interpretability and goodness-of-fit in the machine learning literature (Chang et al. 2009). While objective measures of goodness-of-fit can be used to determine a choice for K , our goal is to obtain an interpretable description of IR content, for which defining objective criteria is challenging.



Figure 4: Topics Ranked by Pro-Cyclicality; Terms within Topics Ranked by Probability

Notes: This table summarizes the 30 separate distributions over vocabulary terms that LDA estimates to represent topics. We order these distributions from 0 to 29 based on a pro-cyclicality index that computes the difference in average time the IR spends discussing the corresponding topic when interest rates are in tightening and loosening cycles, respectively. Within each row, terms are ordered left to right by the probability they appear in each topic, with differential shading indicating approximate probability values.

obtain a 60-dimensional representation of the text information in each IR —the 30 topic levels (later denoted ‘L’) and the 30 absolute changes (denoted ‘D’).



(a) Topic 15 ‘Labor Market’

(b) Topic 26 ‘Demand’

Figure 5: Illustrative Topic Variation across Inflation Reports

Notes: These figures plot the prevalence of two illustrative topics in the Inflation Reports in our sample. Recession periods are shaded in gray. The distributions over terms that each topic induces are represented as word clouds, where the size of term is approximately proportional to its probability.

To illustrate LDA output, Figure 5 plots the share over time of two representative topics within the Inflation Report. We provide an alternative representation of these topics in terms of word clouds. Topic 15 reflects discussion of labor markets. This was fairly stable until 2014, when the Bank started increasing its analysis of the labor market in response to the puzzle that domestic inflationary pressure had remained subdued even as unemployment fell. Topic 26 about demand had a marked increase at the onset of the financial crisis, and has remained high reflecting the MPC’s concerns about the pace of the recovery.

A final point is that, since LDA is an unsupervised learning algorithm, the topics have no objective labels. While the most frequent words in topics are certainly strongly suggestive of a concrete meaning, one should proceed with caution when using these in any evaluation of economic mechanisms. Although we comment on which topics drive

which market interest rates below, we also provide several tests that do not rely on any specific interpretation.

5 Market Impact of Inflation Report Signals

This section presents the main empirical results of the paper. Whereas the event study in Section 2 established that IR publication leads to a greater-than-average absolute change in market interest rates, this section tries to identify which of the IR signals extracted in Section 4 are most responsible for generating the observed rate movements and hence shed light on which channels described in Section 3 may be responsible. A key prediction of that theoretical framework is that the signals that generate an effect through the uncertainty channel can have independent effects on long-run interest rates that do not move short-run interest rates, and we organize much of what follows around the evidence for this.

From a statistical viewpoint, ours is a high-dimensional regression problem since there are more signals in the IR (the 15 numeric signals and the 60 narrative signals) than there are IR publication days (70 in total, but only 69 used for analysis due to differencing across IRs to generate some signals). We therefore use regularization methods below for estimation, but we do not treat numeric and narrative information symmetrically. The numeric variables are, almost by definition, viewed by the Bank of England as informative about future economic conditions and so, while one may not be sure of the exact channels they enter, there is a strong argument for including them in the set of signals that generate any overall information effect. The narrative signals, by contrast, may or may not provide market news. For this reason, we include the numeric variables as regressors throughout the analysis, and only regularize the narrative signals.

5.1 Numeric news signals

We begin the analysis by estimating a regression model that explains yield curve movements using just the numeric information in the IR as captured by the variables in \mathbf{q}_t described in Section 4.1. The model is

$$|Y_t| = \beta_{0Y} + \boldsymbol{\beta}_{1Y}^T \mathbf{q}_t + \beta_{2Y} \text{VIX}_t + \varepsilon_t^Y, \quad (14)$$

where Y_t is the change in the market interest rates (1-year spot, 3-year forward, etc.) observed on IR publication date t . We also include the VIX volatility index as a control, since bond prices may tend to be more volatile on days with increased levels of general market volatility regardless of the level of news in central bank communication.

Table 3 displays the results of regressions in which we use the pre-defined split between central expectations and uncertainty forecast information to separate \mathbf{q}_t into $\mathbf{q}_t^{\text{EXP}}$ and $\mathbf{q}_t^{\text{UNC}}$. We are not particularly interested in the magnitudes nor significance of any individual quantitative measure,²² but rather the overall (joint) explanatory power as measured by the R^2 (R-Squared) statistic for each of the four interest rate horizons. We also report the R^2 statistic in each column of Table 3 for a separate specification in which we do not include VIX but just the IR controls. In either case, the key results are the same.

The first key result is that the forecast variables, and especially the central expectation variables, are an important driver of variance at the short-end of the yield curve, as measured by changes in the 1-year spot rate, with an R^2 statistic of over 0.5. But the impact declines monotonically in maturity, with quite weak effects on market rates at five-years ahead and beyond. The split between expectation and uncertainty signals allows us to show that the declining explanatory power of the forecast variables is driven by a declining role of $\mathbf{q}_t^{\text{EXP}}$. When economic conditions follow a stationary AR1 process (as is natural for GDP growth and inflation) and signals operate through the expectations channel, we showed in proposition 1 that central bank communication would have its largest effect on short-run rates with a monotonically declining effect in horizon. The empirical results on the market impact of $\mathbf{q}_t^{\text{EXP}}$ is consistent with this view.

The second key result is that the distribution signals are increasingly important as we move out the yield curve, which is consistent with an effect through the uncertainty channel in our model. We measure the relative explanatory power of the uncertainty variables by the partial R^2 , i.e. the proportion of the remaining unexplained variance that is captured by adding those additional variables. The partial R^2 grows with maturity of the interest rate.

We next replace the model in (14), in which the absolute change in the overall yield is the dependent variable, with two separate regressions in which the dependent variables are the absolute change in yield expectations and the absolute change in the term premium from the decomposition described in Section 2.2. Table 4 reports these results.²³ One can observe that the relative impact (as measured by partial R^2) of the $\mathbf{q}_t^{\text{EXP}}$ variables on expectations is not greater at the long than the short end of the yield curve, whereas the relative impact on long-run term premiums is greater than that on the short run. This is another result that is consistent with the uncertainty channel, which should have its largest impact via the long-run term premium. Although, one should be careful in over-interpreting the R^2 from these models as the amount of variation in term premiums

²²In particular, because some of the measures are highly correlated.

²³Table D.3 in Appendix D reports the equivalent of table 4 with standard errors of coefficients reported.

Table 3: Effect of Forecast Variables on Market Yields

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$ \Delta i_{0:12;t} $	$ \Delta i_{0:12;t} $	$ \Delta f_{36;t} $	$ \Delta f_{36;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60:120;t} $	$ \Delta f_{60:120;t} $
VIX_t	0.0016***	0.0013**	0.0015	0.0023	0.0011	0.0018	0.0012	0.0017
π_m^{CB}	-0.037	-0.016	-0.060**	-0.014	-0.014	0.016	-0.0014	0.025
$\text{Supr}(\pi_m^{CB})$	0.0018	0.034	-0.018	-0.035	-0.0026	-0.042	-0.022	-0.051
g_m^{CB}	0.018	0.031**	-0.0032	0.00010	-0.015	-0.012	-0.024	-0.015
$\text{Supr}(g_m^{CB})$	-0.0097	-0.048*	0.042*	0.050	0.042	0.053*	0.037	0.024
$\tilde{y}_{m;24}^{CB}$	1.00	0.41	1.92*	3.35*	0.76	1.68	0.067	-0.50
$ \tilde{y}_{m;24}^{CB} - \tilde{y}_{m-1;24}^{CB} $	4.01	4.96**	-0.28	1.41	-2.36	-1.00	-2.05	-1.43
$ \tilde{y}_{m;21}^{CB} - \tilde{y}_{m-1;24}^{CB} $	0.25	1.64	1.47	-0.35	2.28	-0.25	1.88	1.31
$\text{Var}(\pi_m^{CB})$		-0.011		0.025		0.021		0.020
$\Delta \text{Var}(\pi_m^{CB})$		0.038		-0.057		-0.12**		-0.085*
$\text{Skew}(\pi_m^{CB})$		0.0045		0.0100		-0.0040		-0.0036
$\Delta \text{Skew}(\pi_m^{CB})$		0.025		0.065		0.068		0.052
$\text{Var}(g_m^{CB})$		0.0010		-0.0052		-0.0021		-0.012
$\Delta \text{Var}(g_m^{CB})$		0.0098		0.019		0.024		0.034*
$\text{Skew}(g_m^{CB})$		-0.014		-0.071		-0.060		-0.045
$\Delta \text{Skew}(g_m^{CB})$		-0.063**		-0.00055		0.0031		-0.022
Constant	-0.053*	-0.070*	0.019	-0.026	0.052	0.014	0.074	0.043
R-squared	0.403	0.563	0.250	0.368	0.122	0.280	0.114	0.274
Partial R-squared	-	0.285	-	0.320	-	0.564	-	0.584
R-squared No Vix	0.336	0.526	0.216	0.303	0.099	0.229	0.082	0.215
Component	Total	Total	Total	Total	Total	Total	Total	Total

Notes: This table reports estimates from regressing absolute changes in market yields on the numeric forecast variables defined in Section 4.1.

Table 4: Effect of Forecast Variables on Market Yields: Split by Component

(a) Expectations								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$ \Delta i_{0:12;t} $	$ \Delta i_{0:12;t} $	$ \Delta f_{36;t} $	$ \Delta f_{36;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60:120;t} $	$ \Delta f_{60:120;t} $
VIX _t	0.0013**	0.0010*	0.0010	0.0014	0.00072	0.0011	0.00074	0.00096
π_m^{CB}	-0.031	-0.016	-0.054***	-0.027	-0.043***	-0.017	-0.032**	-0.021
Supr(π_m^{CB})	0.014	0.040	-0.016	-0.014	-0.011	-0.015	-0.013	-0.012
g_m^{CB}	0.017	0.028*	0.0072	0.012	-0.00041	0.0017	-0.00055	0.0014
Supr(g_m^{CB})	-0.013	-0.049**	0.028	0.018	0.028*	0.027	0.023	0.016
$\tilde{y}_{m:24}^{CB}$	0.54	-0.034	1.65***	2.44*	1.43**	2.60**	0.89	1.64
$ \tilde{y}_{m:24}^{CB} - \tilde{y}_{m-1:24}^{CB} $	2.68	3.41*	2.19	2.88	1.26	2.13	1.34	1.70
$ \tilde{y}_{m:21}^{CB} - \tilde{y}_{m-1:24}^{CB} $	-0.20	1.09	-0.29	-0.29	0.64	-0.17	0.33	0.24
Var(π_m^{CB})		-0.0100		0.013		0.012		0.0063
Δ Var(π_m^{CB})		0.036		0.0078		-0.014		0.0070
Skew(π_m^{CB})		-0.00092		0.021		0.012		0.0051
Δ Skew(π_m^{CB})		0.017		0.035		0.041		0.025
Var(g_m^{CB})		0.0019		-0.0021		-0.00050		-0.00093
Δ Var(g_m^{CB})		0.0073		0.0080		0.0093		0.0057
Skew(g_m^{CB})		-0.0080		-0.039		-0.049		-0.034
Δ Skew(g_m^{CB})		-0.056*		-0.011		0.00064		0.00020
Constant	-0.045	-0.060	-0.010	-0.042	0.0091	-0.019	0.0058	-0.010
R-squared	0.390	0.544	0.354	0.443	0.297	0.429	0.331	0.408
Partial R-squared	-	0.283	-	0.201	-	0.308	-	0.189
R-squared No Vix	0.331	0.513	0.325	0.396	0.277	0.387	0.300	0.366
Component	Exp	Exp	Exp	Exp	Exp	Exp	Exp	Exp

(b) Term Premiums								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$ \Delta i_{0:12;t} $	$ \Delta i_{0:12;t} $	$ \Delta f_{36;t} $	$ \Delta f_{36;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60:120;t} $	$ \Delta f_{60:120;t} $
VIX _t	0.00034***	0.00030**	0.00056	0.00078	0.00069	0.00087	0.00054	0.00075
π_m^{CB}	-0.00022	0.0018	0.0030	0.0045	0.022	0.018	0.033	0.031
Supr(π_m^{CB})	-0.0055	-0.00034	-0.0027	-0.019	-0.015	-0.020	-0.020	-0.019
g_m^{CB}	0.00017	0.0012	-0.0070	-0.0079	-0.012	-0.0094	-0.015	-0.0075
Supr(g_m^{CB})	0.00085	-0.0025	0.015	0.019	0.017	0.0031	0.014	-0.017
$\tilde{y}_{m:24}^{CB}$	0.26*	0.19	0.047	0.32	-0.76	-0.94	-1.22	-1.93**
$ \tilde{y}_{m:24}^{CB} - \tilde{y}_{m-1:24}^{CB} $	0.88**	0.99**	-0.95	-0.79	0.57	0.0058	1.81	0.97
$ \tilde{y}_{m:21}^{CB} - \tilde{y}_{m-1:24}^{CB} $	0.46	0.64	0.73	0.15	-0.28	0.98	-0.58	1.85
Var(π_m^{CB})		-0.0019		0.0084		0.0091		0.012
Δ Var(π_m^{CB})		0.0027		-0.035		0.020		0.046
Skew(π_m^{CB})		0.0011		0.0016		0.013		0.018
Δ Skew(π_m^{CB})		0.0077		0.0076		-0.0014		-0.00030
Var(g_m^{CB})		-0.00081		-0.0051		-0.0082		-0.013
Δ Var(g_m^{CB})		0.0021		0.0067		0.0029		0.0083
Skew(g_m^{CB})		-0.0033		-0.027		-0.015		-0.015
Δ Skew(g_m^{CB})		-0.0060		0.011		-0.0016		-0.019
Constant	-0.0025	-0.0019	0.024	0.021	0.042	0.035	0.059	0.044
R-squared	0.325	0.457	0.091	0.168	0.142	0.227	0.141	0.280
Partial R-squared	-	0.289	-	0.459	-	0.374	-	0.497
R-squared No Vix	0.215	0.390	0.058	0.115	0.107	0.172	0.128	0.252
Component	TP	TP	TP	TP	TP	TP	TP	TP

Notes: This table reports estimates from regressing absolute changes in market yields on the numeric forecast variables defined in Section 4.1.

at the short end of the yield curve is very small (as shown in Table 1); the R^2 in column 2 of panel (b) tells us that the overall regression is explaining about half of very little variation.

These results provide basic, preliminary evidence that central bank communication operates through multiple channels, and that movements in long-run interest rates can arise from signals on uncertainty that are distinct from the signals that move short-run interest rates.

5.2 Narrative news signals

The analysis of the numeric forecast information is consistent with the uncertainty channel being a primary driver of longer-maturity yields. We now use the narrative signals to explore this idea more specifically, as its dimensionality provides a means to explore the different communication channels that can drive the long-run information effect (expectations, uncertainty, investor demand) in more detail.

5.2.1 Is there news in narrative?

Before analyzing the market impact of narrative signals, one needs to establish that they contain news at all over and above the information in the numeric forecasts. The first step is to purge variation in the narrative variables that is endogenous to the numerical forecast information. For example, an increase in the narrative content about inflation may be associated with a deviation of inflation from target in the modal forecast.

We fit the models

$$\theta_{t,k} = \alpha_{0k} + \boldsymbol{\alpha}_{1k}^T \mathbf{q}_t + \alpha_{2k} \text{VIX}_t + v_{Lt}^k \quad (15)$$

$$\delta_{t,k} = \beta_{0k} + \boldsymbol{\beta}_{1k}^T \mathbf{q}_t + \beta_{2k} \text{VIX}_t + v_{Dt}^k \quad (16)$$

The estimated residuals \widehat{v}_{Lt}^k and \widehat{v}_{Dt}^k represent the part of the level of topic k , and its change from the previous IR, not explained by the forecasts.

Our construction is similar to that in Romer and Romer (2004) and Cloyne and Hürtgen (2016), who construct monetary policy shocks by regressing interest rate decisions on numerical forecast variables for the Federal Reserve and the Bank of England, respectively. Here we construct ‘narrative shocks’ by extracting the exogenous component of the Inflation Report text. This yields 60 narrative shocks that we denote $\widehat{\mathbf{v}}_t = (\widehat{\mathbf{v}}_{Lt}, \widehat{\mathbf{v}}_{Dt})$.

One concern in high-dimensional regression is that correlation among regressors can impede the ability to identify variables with a ‘true’ relationship with the outcome. In fact, our narrative shocks display rather little correlation on average. Figure 6 shows a histogram of all the unique bilateral correlations between the 60 shocks. The modal

correlation is zero and only 377 out of 1770 are statistically significant (at 5% level) indicating that the narrative shocks can largely be thought of as independent signals.




Figure 6: Correlation between Narrative Shocks

Notes: Histogram of the 1770 pairwise correlations between narrative shocks. The darker shaded bars indicate correlations that are significantly different from zero at the 5% significance level.

Statistically, establishing whether narrative shocks contain news is equivalent to determining whether the narrative shocks explain the residuals from equation (14), that part of the overall market news not explained by numerical forecasts. As explained above, ordinary least squares is not feasible in this setting since there are more narrative shock variables than degrees of freedom in the model. Instead, we use an elastic net regression as in Zou and Hastie (2005) and solve

$$\min_{\gamma_Y} \sum_t (\hat{\varepsilon}_t^Y - \gamma_Y^T \hat{\mathbf{v}}_t)^2 + \lambda [\alpha \|\gamma_Y\|_1 + (1 - \alpha) \|\gamma_Y\|_2^2]. \quad (17)$$

The first term is the objective function of an OLS regression of the yield- Y residuals on the narrative shock variables. The second term is a penalty on non-zero values of the regression coefficients γ_Y . The parameter α can range from 0, equivalent to a ridge regression, to 1, equivalent to the least absolute shrinkage and selection operator (LASSO). The $\alpha = 1$ LASSO specification is useful because it induces sparse solutions but when two or more covariates are significantly correlated it typically only generates a non-zero coefficient for one of them. We set $\alpha = 0.99$ to induce a degree of sparsity akin to the LASSO, while maintaining robustness to the (relatively few) high correlations in the narrative shocks (Friedman et al. 2010).

Before estimating (17), one must choose a value for the penalty parameter λ . We use the common approach of selecting λ using cross validation based on out-of-sample predictive performance. We describe this procedure in detail in Appendix C and here just focus on the results. We find that a large number of narrative shocks are selected at each maturity, as displayed in Table 5.

That a large number of narrative shocks is selected for each maturity suggests that

Table 5: Number of Selected Narrative Shock in Baseline

	$ \Delta i_{0:12,t} $	$ \Delta f_{36,t} $	$ \Delta f_{60,t} $	$ \Delta f_{60:120,t} $
# selected narrative shocks	60	57	54	57

Notes: This table displays the number of selected covariates arising from the estimation of (17) for four market rates by leave-one-out cross validation.

they are indeed important in explaining the yield residuals. To test that this result is not spurious, we need to compare it to the distribution of the number of selected variables under the null hypothesis that narrative shocks are independent of yield residuals. To approximate this hypothesis test, we use a simple permutation test. In each of 500 simulations for each interest rate, we randomly permute the residuals, re-estimate (17), and record the number of selected narrative shock variables. We then compare the values in Table 5 against these simulated distributions. The results are shown in Figure 7, where the red-shaded histogram is the distribution of the number of selected narrative shocks generated by the simulations, and the blue dashed line is the number of narrative shocks we select in our actual data. At all interest rate maturities we strongly reject the null hypothesis that the correlations we find are spurious. In the vast majority of permuted draws, the elastic net selects no shocks, and in no draw is the number of selected shocks greater than the number we select with the non-permuted data. We conclude that there indeed appears to be genuine explanatory power contained in the IR narrative that is orthogonal to that in the numerical forecast variables.²⁴

The finding that there is market news in the narrative of central bank communication is of independent interest and another contribution of the paper. Most studies of central bank communication that analyze their content focus on numerical information, but there is evidence that some other factor is needed to explain the full market reaction around communication events (Gürkaynak et al. 2018). We find it natural to view narrative as an important aspect of this ‘missing’ information in event studies, and the approaches we develop are more broadly relevant to the literature.

5.2.2 Testing distinct long-run information I: key narrative signals

While we have established that narrative shocks are a source of news, we do not know which are most important for explaining market rates, nor whether these differ by maturity. The baseline elastic net regression selects nearly every shock at every maturity,²⁵

²⁴We also repeat this methodology in order to test for narrative information in the IR by component (expectations and term premiums). We do not report here in the interest of space, but we continue to find strong evidence for such information at all maturities for each component.

²⁵This is not necessarily surprising. As Meinshausen and Bühlmann (2006) have shown, the number of selected features from an elastic net regression estimated via cross validation may be a superset of the relevant variables.

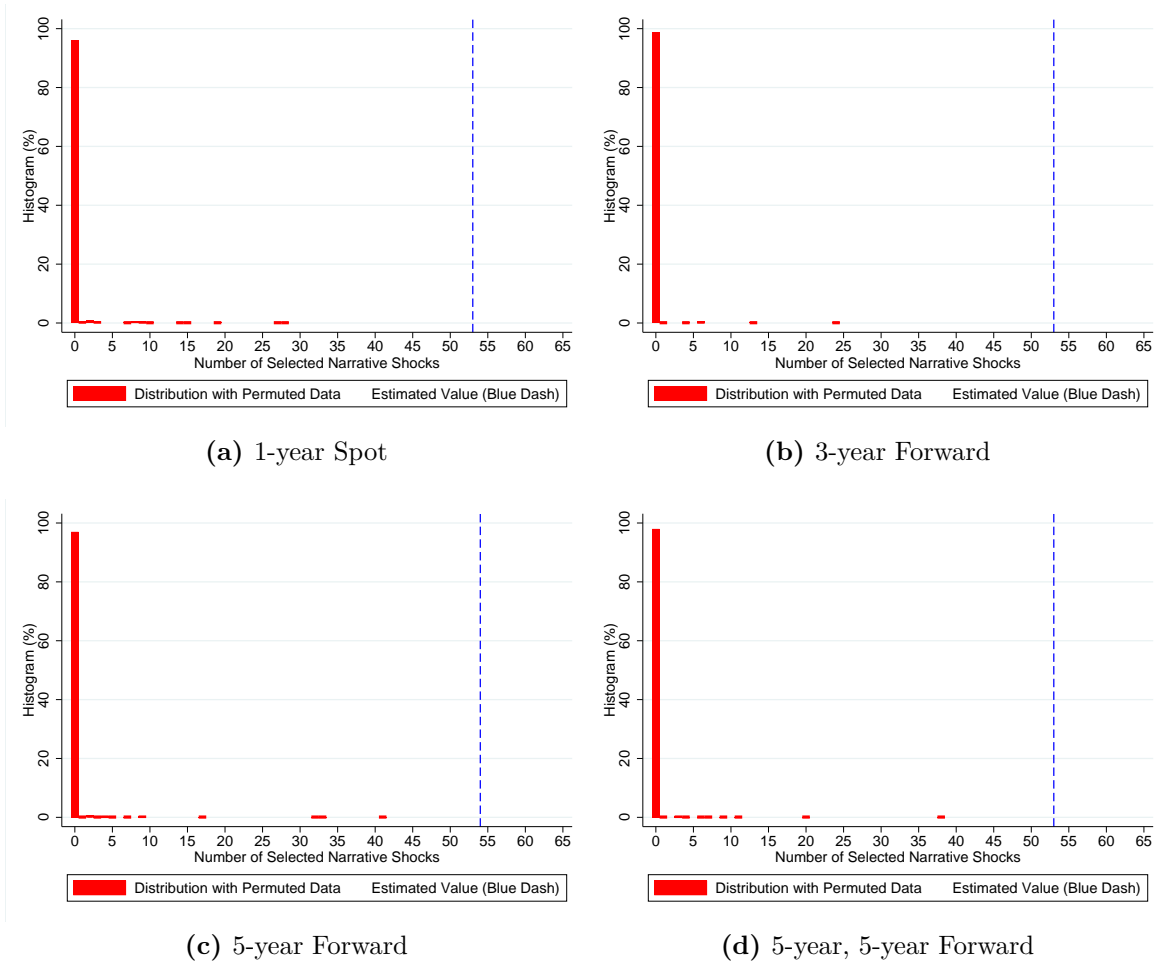


Figure 7: Permutation Test for Narrative News

Notes: These figures describe a permutation test for narrative news. The blue, dashed vertical lines for each yield plot the number of selected text variables from Table 5. The red histograms describe the distribution of selected features in 500 different random permutations of yield residuals for which we used the same cross validation procedure as on the original data. In no permutation do we select as many features as with the true order.

so we need some way of discriminating among signals according to their information content. To do this, we adopt a bootstrap procedure suggested by Hastie et al. (2015). In each of 500 simulations we draw a bootstrap sample with replacement from our original data, compute coefficient estimates using the same cross validation procedure as in the baseline, and record whether each topic variable is selected. Across all the bootstrap draws, we can compute the fraction of times that each topic variable is selected, and use this as an indicator of which variables are key in driving the market response to the IR.²⁶

Table 6 lists the top four topic variables for each yield based on the bootstrap draws, and reports the fraction of draws in which they appear. Simple observation reveals that the topics that are most likely to drive short and long rates diverge considerably: the top topics for 1-year spot rates and 5-year ahead, 5-year forward rates contain no overlap. We formalize this below.

Table 6: Top Topics for Different Yields (L=Level; D=Change)

$ \Delta i_{0:12,t} $		$ \Delta f_{36,t} $		$ \Delta f_{60,t} $		$ \Delta f_{60:120,t} $	
Var	Selection %	Var	Selection %	Var	Selection %	Var	Selection %
L25	0.958	D24	0.858	L28	0.876	D17	0.91
D24	0.954	D25	0.844	D17	0.784	D18	0.896
L5	0.932	L28	0.826	D18	0.772	L20	0.836
L26	0.91	D14	0.76	L20	0.722	D13	0.808

Notes: This table lists the top four topics for each yield according to fraction of times they are selected across 500 bootstrap draws. An L indicates the topic variable corresponds to a residual in levels, while a D indicates a residuals in the absolute change in the topic level.

The most likely words within the key topics for the shortest- and longest-maturity assets we examine, listed above in Table 6, are presented graphically as word clouds in Figures 8 and 9. As noted in the discussion of LDA, topics do not come with labels and interpreting them is a subjective exercise. Still, the key topics that drive the different yields are suggestive of the channels from the theoretical model. Those that drive the 1-year spot rates appear to relate to current economic conditions and include those that vary most with the interest rate cycle (see Figure 5). On the other hand, the key topics for the five-year ahead, five-year forward rate appear to relate to the forecasts and their uncertainty and are less cyclical. This is consistent with a model in which the central bank sends signals about the levels of mean-reverting economic conditions that affect short-run rates, and signals about economic uncertainty that affects long-run rates.

Table 7 formalizes the finding that different narrative signals drive different rate ma-

²⁶The elastic net regression we estimate is very similar to the LASSO, which has a probabilistic formulation as a Bayesian regression model with Laplace priors on the coefficients. The bootstrap procedure can be thought of as a shortcut for doing a full posterior simulation exercise that would allow one to compute marginal inclusion probabilities.

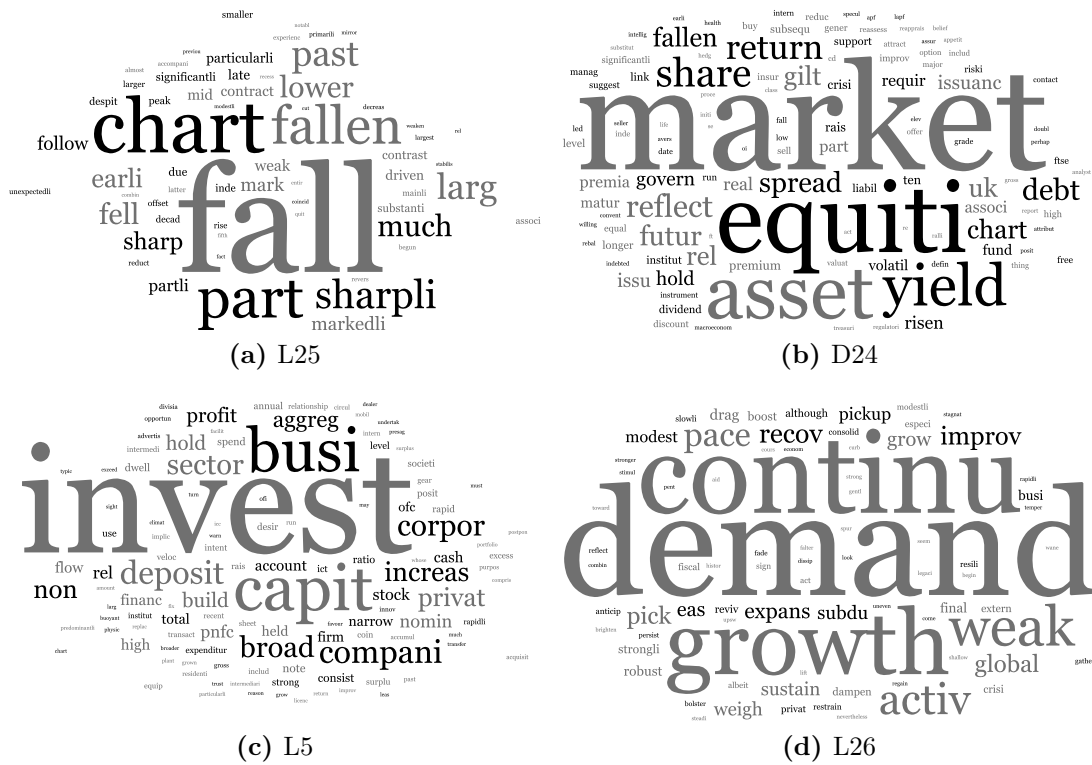


Figure 8: Key Topics for Market Reaction to Narrative: 1-Year Spot Rate

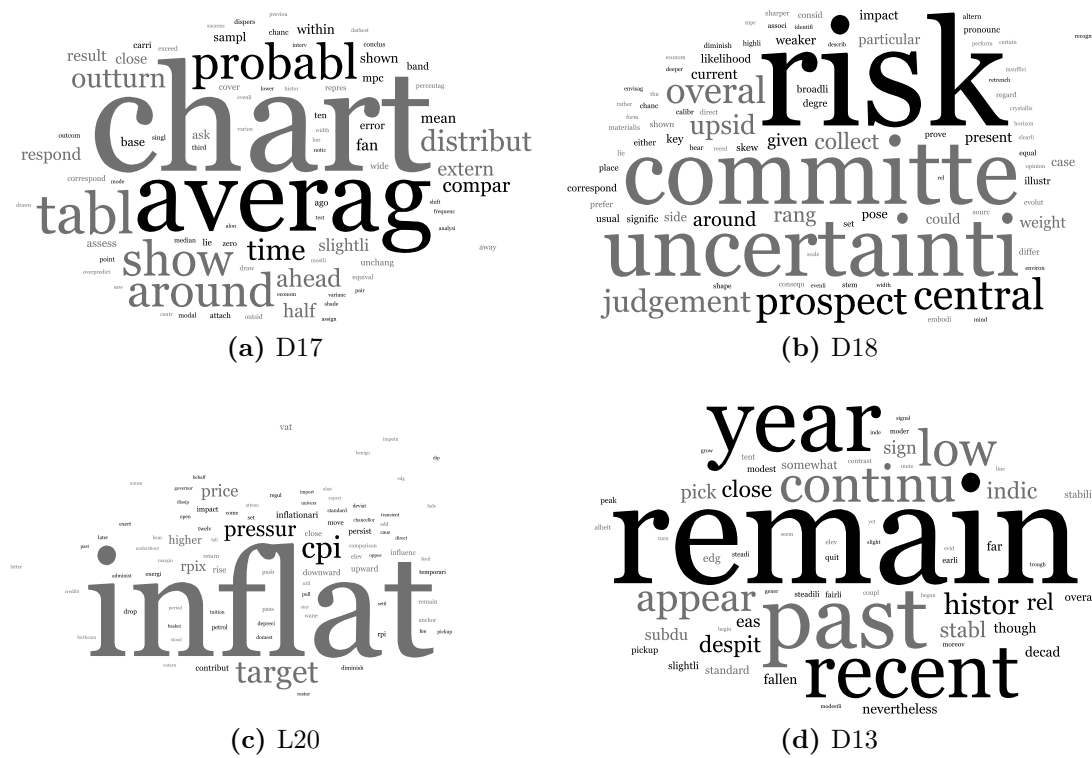


Figure 9: Key Topics for Market Reaction to Narrative: 5-Year, 5-Year Forward Rate

turities. It reports the Pearson correlation coefficients between narrative shocks based on the fraction of bootstrap draws in which they are selected. The selection percentages for 1-year spot rates are in fact uncorrelated with those for longer rates, while the selection percentages associated with all other rates are significantly correlated. Importantly, this finding casts doubt on either the expectations channel or the investor demand channel being the primary driver of long-run rate movements in response to the IR. Both channels require the signals that move short-run rates to also move long-run rates.

Table 7: Pearson Correlations of Narrative Signals' Selection Percentage Across Yields

	$ \Delta i_{0:12,t} $	$ \Delta f_{36,t} $	$ \Delta f_{60,t} $	$ \Delta f_{60:120,t} $
$ \Delta i_{0:12,t} $	1	.	.	.
$ \Delta f_{36,t} $	0.21	1	.	.
$ \Delta f_{60,t} $	0.04	0.68***	1	.
$ \Delta f_{60:120,t} $	0.06	0.45***	0.84***	1

Notes: This table reports the Pearson correlation coefficient of the topics' selection percentages across 500 bootstrap draws for different yields. *** denotes significance at the 1% significance level.

If the uncertainty channel drives long-run interest rates, we expect it to do so through the term premium rather than expectations, while the reverse is true of the expectations channel. We repeat the bootstrap procedure on each separate yield component, Table 8 shows the results. There is substantial overlap between the narrative signals that drive short-run movements in the overall yield change and the signals that drive the short-run expectations component. In contrast, there is an overlap between the top long-run overall and term premium signals. Relatedly, the key signals for explaining long-run term premiums are not the key signals for explaining the overall change in short-run yields. This is further evidence that long-run rate movements do not arise solely from trading activity in response to movements at the short end, as would be the case in the investor demand channel.

Table 9 is the analogue of Table 7 by yield component. Here we find an interesting distinction between expectation and term premium signals: the topic selection percentages for expectations are much more correlated across maturities than that of term premiums. This suggests that the results on the overall change potentially mask an underlying expectations component that is in fact rather persistent since a set of common signals drive short- and long-run expectations. However, this is not the dominant source of variation in long-run interest rates since these signals are not the key drivers of the overall long-run rate. This accords with our view that the uncertainty channel is not the only factor driving the long-run information effect, but is the primary one.

Table 8: Top Narrative Signals for Different Yields (L=Level; D=Change)

(a) Expectations							
Var	$ \Delta i_{0:12,t} $ Selection %	Var	$ \Delta f_{36,t} $ Selection %	Var	$ \Delta f_{60,t} $ Selection %	Var	$ \Delta f_{60:120,t} $ Selection %
L26	0.95	D24	0.962	D24	0.936	D24	0.984
L25	0.936	D25	0.92	D1	0.834	D25	0.94
D28	0.908	D1	0.876	D25	0.818	L1	0.876
L5	0.892	L25	0.856	L25	0.784	D8	0.856

(b) Term Premiums							
Var	$ \Delta i_{0:12,t} $ Selection %	Var	$ \Delta f_{36,t} $ Selection %	Var	$ \Delta f_{60,t} $ Selection %	Var	$ \Delta f_{60:120,t} $ Selection %
L9	0.97	D17	0.844	D17	0.96	D13	0.962
D24	0.932	L28	0.818	D18	0.958	D18	0.928
D3	0.894	L7	0.804	D13	0.874	D17	0.914
D7	0.854	D18	0.8	D20	0.842	D8	0.85

Notes: This table lists the top four topics for each yield and component according to fraction of times they are selected across 500 bootstrap draws. An L indicates the topic variable corresponds to a residual in levels, while a D indicates a residuals in the absolute change in the topic level.

Table 9: Pearson Correlations of Topic Variables' Selection Percentage Across Yields

(a) Expectations				
	$ \Delta i_{0:12,t} $	$ \Delta f_{36,t} $	$ \Delta f_{60,t} $	$ \Delta f_{60:120,t} $
$ \Delta i_{0:12,t} $	1	.	.	.
$ \Delta f_{36,t} $	0.43***	1	.	.
$ \Delta f_{60,t} $	0.25*	0.92***	1	.
$ \Delta f_{60:120,t} $	0.21	0.91***	0.92***	1

(b) Term Premiums				
	$ \Delta i_{0:12,t} $	$ \Delta f_{36,t} $	$ \Delta f_{60,t} $	$ \Delta f_{60:120,t} $
$ \Delta i_{0:12,t} $	1	.	.	.
$ \Delta f_{36,t} $	0.02	1	.	.
$ \Delta f_{60,t} $	0.12	0.54***	1	.
$ \Delta f_{60:120,t} $	0.21	0.14	0.70***	1

Notes: This table reports the Pearson correlation coefficient of the topics' selection percentages across 500 bootstrap draws for different yields and for different components of the yield curve. * and *** denote significance at the 10% and 1% significance levels, respectively.

5.2.3 Testing distinct long-run information II: source of key narrative signals

As we emphasize above, any attribution of meaning to topics in the absence of external validation is subjective. So, although the idea that the key narrative signals that drive short-run rates displayed in Figure 8 reflect short-run economic news and that those in Figure 9 reflect long-run uncertainty is plausible, other interpretations are possible. To aid our interpretation, we exploit the structure of the Inflation Report, which is divided into two broad parts, one that reports on current economic conditions and another that explains the forecast and risks around it (see Section 4.2 for more details). Instead of forming topic distributions at the level of the whole IR, we form them within each of these two sections, and compute the average fraction of content in them devoted to discussing the key topic signals at different maturities reported in Table 6. Table 10 reports the results.

Table 10: Key Signal Coverage by Inflation Report Section

	Economics	Forecast
$ \Delta i_{0:12,t} $	0.136	0.105
$ \Delta f_{36,t} $	0.120	0.171
$ \Delta f_{60,t} $	0.099	0.215
$ \Delta f_{60:120,t} $	0.100	0.204

Notes: This table reports the fraction of time that each section of the Inflation Report spends discussing the key signals reported in table 6 for each yield.

We find that the news that drives different maturities tends to arise from different parts of the IR, with the long-run news coming more from the forecast part and less from the current economic conditions part. This is important because the sections provide a means of objectively interpreting the key signals. The fact that the section devoted to providing information about short-run economic conditions does not devote much attention to the topic signals most associated with long-run rate movements highlights that signals on the levels of conditions are not their main driver. More broadly, this is further evidence that fundamentally different information drives long-run rates, which is the distinguishing feature of the uncertainty channel.

5.2.4 Testing distinct long-run information III: placebo regressions

Although we have found distinct narrative signals in the bootstrap test, one might still question whether these actually explain different amounts of movement in yield residuals. Our final test of distinct long-run information assesses to what extent the key narrative signals associated with each maturity actually explain more of the yield residual at that maturity than the key narrative signals associated with other maturities.

To begin, we regress the absolute yield changes on just the VIX, and add in the numeric forecast variable \mathbf{q}_t defined above. We then compute the partial R^2 associated with also including \mathbf{q}_t , which is recorded in Table 11. These regressions replicate those in Table 3, and again one can note the declining impact of the numeric forecasts along the yield curve. Finally, we add in the top four narrative signals from Table 6, and record the updated overall R^2 and partial R^2 in Table 11. One finding of interest here is that, while the relative impact of the numeric forecasts as measured by partial R^2 is decreasing in yield maturity, the relative impact of narrative signals is increasing. That is, narrative information has its largest relative effect for longer-run rates. (Appendix Tables D.1 and D.2 display more details of the regressions on which Table 11 is based.)

Table 11: Summary of R^2 Statistics from Yield Regressions

		Vix Only	Vix & \mathbf{q}_t	Vix, \mathbf{q}_t & Key Narrative Signals
$ \Delta i_{0:12;t} $	R^2	0.286	0.563	0.693
	Partial R^2	.	0.39	0.29
$ \Delta f_{36;t} $	R^2	0.150	0.368	0.518
	Partial R^2	.	0.26	0.24
$ \Delta f_{60;t} $	R^2	0.065	0.280	0.502
	Partial R^2	.	0.23	0.31
$ \Delta f_{60:120;t} $	R^2	0.058	0.274	0.542
	Partial R^2	.	0.23	0.37

Notes: This table reports the R^2 (and Partial- R^2) statistics from regressing absolute changes in market yields on Vix (column 3), Vix and the numeric forecast variables defined in Section 4.1 (column 4), and Vix, the numeric forecast variables and also the top 4 topic shocks for each yield (column 5).

To test whether the key narrative signals at each maturity explain more yield residual variance at that maturity, we replace them with key narrative signals at other maturities, and record the partial R^2 in exactly the same way as in Table 11. Table 12 displays the results and indicates the significance level of an F-test on the joint significance of the narrative shocks given the inclusion of VIX and \mathbf{q}_t . Here we have not accounted for the post-selection inference problem of how to conduct hypothesis tests on the significance of selected regressors from an elastic net regression. Our question of interest is not whether the selected signals at each maturity are significant in a statistical sense, though, but how their explanatory power compares to other narrative signals selected at other maturities. This is another way of asking whether we have selected independent signals to explain long-run rate movements.

Table 12: Matrix of Partial R^2 from Placebo Regressions

Asset News	Narrative Shocks			
	$i_{0:12;t}$	$f_{36;t}$	$f_{60;t}$	$f_{60:120;t}$
$i_{0:12;t}$	0.29***	0.08*	0.06	0.01
$f_{36;t}$	0.15	0.23*	0.18	0.11
$f_{60;t}$	0.14	0.21	0.31**	0.24*
$f_{60:120;t}$	0.12	0.13	0.34***	0.37***

Notes: This table reports the Partial- R^2 statistics from a regression of absolute changes in market yields on Vix and the numeric forecast variables (\mathbf{q}_t) to which four narrative shocks are added. Each row reports the results for a different yield. Each column indicates the yields from which the top narrative shocks are estimated. For example, the cell in the bottom right indicates the partial- R^2 from adding the top four narrative shocks estimated from the 5y-5y forward rates on 5y-5y forward rate news; it is the same result as in table 11. Stars (***, **, and *) indicate the significance level (1%, 5% and 10%) of an F-test on the joint significance of the narrative shocks given the inclusion of VIX and q_t .

As one can see from Table 12, the answer is clear. The top narrative signals for the short run explain only 12% of residual variance for five-year, five-year forward rates. Even more striking, the top narrative signals for the long run explain essentially none (1%) of the residual variance for the one-year spot rate. The effects of the top narrative signals for the five-year forward rate and the five-year, five-year rate are similar for both maturities, which is not surprising given there is overlap in the selected topics. The bottom line is that the selected narrative signals for the long-run are not some indirect proxy for the information that moves short-run rates, but distinct information. This reinforces the plausibility of the uncertainty channel.

As a further exercise, we repeat the same analysis also including the separate components of the overall yield curve, with results in Table 13. This allows us to explore in more detail the impact of the expectations and demand channels. The results in Table 13 show that the narrative signals that explain expectations in the three-year forward rate also explain expectations in all the rates to a similar degree. More generally, most of the top narrative signals for expectations at any maturity produce a significant F-test when regressed on the expectations residuals at any other maturity. While this is consistent with an expectations channel that provides information on short-run conditions and also long-run conditions via persistence, it is not the main driver of the overall movement in long-run rates. The top four narrative signals for explaining five-year, five-year expectations explain only 4% of the overall residual variance. On the other hand, the top four signals for explaining the five-year, five-year term premium explain 23% of the overall residual variance.

The investor demand channel requires the information that moves short-run expectations to also move the long-run term premium. Instead, the top signals for explaining expectations in the one-year spot and three-year forward rates explain only 12% (10%)

Table 13: Full Matrix of Partial R^2 from Placebo Regressions including Components of the Asset Response

Asset News	Narrative Shocks											
	$i_{0:12;t}$			$f_{36;t}$			$f_{60;t}$			$f_{60:120;t}$		
	Overall	EXP	TP	Overall	EXP	TP	Overall	EXP	TP	Overall	EXP	TP
$i_{0:12;t}$	0.29***	0.31***	0.11	0.08	0.24***	0.06	0.06	0.08*	0.06	0.01	0.06	0.04
EXP($i_{0:12;t}$)	0.30***	0.32***	0.08	0.10*	0.23***	0.06	0.06	0.10*	0.06	0.01	0.06	0.04
TP($i_{0:12;t}$)	0.17*	0.10	0.36***	0.13	0.19	0.02	0.06	0.13	0.02	0.06	0.13	0.04
$f_{36;t}$	0.15	0.11	0.15	0.23*	0.15***	0.15	0.18	0.24***	0.07	0.11	0.16**	0.05
EXP($f_{36;t}$)	0.19***	0.14*	0.14*	0.19*	0.25***	0.07	0.08	0.21***	0.05	0.07	0.21***	0.03
TP($f_{36;t}$)	0.11	0.17	0.13	0.15	0.03	0.31**	0.35**	0.15	0.18	0.28*	0.04	0.18
$f_{60;t}$	0.14	0.17	0.12	0.21	0.04	0.26**	0.31**	0.19*	0.14	0.24*	0.04	0.15
EXP($f_{60;t}$)	0.18**	0.12	0.18	0.21*	0.18***	0.11	0.16	0.21***	0.05	0.11	0.18**	0.05
TP($f_{60;t}$)	0.12	0.12	0.08	0.06	0.08	0.38***	0.38***	0.08	0.42***	0.44***	0.09*	0.39***
$f_{60:120;t}$	0.12	0.17	0.08	0.13	0.04	0.28**	0.34***	0.13	0.24*	0.37***	0.04	0.23*
EXP($f_{60:120;t}$)	0.22***	0.14*	0.16	0.24***	0.27***	0.10	0.12	0.24***	0.05	0.09	0.22***	0.05
TP($f_{60:120;t}$)	0.14	0.12	0.04	0.07	0.10	0.29**	0.32***	0.08	0.39***	0.40***	0.11*	0.36***

Notes: This table replicates table 12 reporting the Partial- R^2 statistics from adding four narrative shocks to explain different yields but expands the analysis to include the breakdown by yield components. The top four topics added are indicated by the column and each row reports the results for a different yield or its component. Stars (**, **, and *) indicate the significance level (1%, 5% and 10%) of an F-test on the joint significance of the narrative shocks given the inclusion of VIX and q_t .

and 10% (8%) of the variation in the five-year, five-year (five-year forward) term premium. Moreover, the top signals for explaining the long-run term premium explain almost nothing about short-run expectations (4% for the one-year spot and 3% for the three-year forward).

Together, the results show that the expectations channel may indeed operate, but there is little evidence of the investor demand channel. In any case, the main driver of long-run interest rates is narrative signals that explain the long run independently of any impact on short- or medium-term expectations. This long-run impact comes largely via the term premium. All of these facts point to the uncertainty channel as the primary mechanism through which the Inflation Report operates on long-run interest rates.

6 Robustness

6.1 No zero lower bound

Our model does not consider the impact of the Zero Lower Bound (ZLB). One concern is that in a situation where the ZLB is binding, the short-end of the yield is downwardly restricted. This means a signal that economic conditions will be persistently weaker, which would be expected to shift the short- and long-end of the yield curve equally, can only move the yield curve further out. This would be an expectations shock that appeared to only move the long-end. Moreover, to the extent that the persistent weakness of the economy might signal that the economy may have switched into a new regime in which deflation risk is higher, those effects may come through a change in term premiums.

To assess the role that this plays in driving our results, we analyze the effect of removing the ZLB period from our analysis. In our sample, the ZLB was binding from March 2009 and so the pre-ZLB sample is reduced by 25 IR events.

In spite of the small sample, the key results from the whole sample are present in this period. First, Table 14 shows that the basic pattern of the effect of the numeric forecast information holds; expectational signals have the greater effect at the short end while distributional information is relatively more important further out. In fact, the numeric information explains more of the market reaction at all points on the yield curve to the IR before the ZLB.

Second, we repeat the information test of Section 5.2.1 using just the 45 IRs and their corresponding narrative signals. Notwithstanding their being a smaller residual left to explain, the narrative continues to contain important information that helps to explain this residual reaction. Table 15 presents the results of the information test.

Table 16 shows the correlation across the full sample and the pre-ZLB sample of the top topics measured using the topic variables' selection percentage from the bootstrap

Table 14: Effect of Forecast Variables on Market Yields - Pre-ZLB Sample

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$ \Delta i_{0:12;t} $	$ \Delta i_{0:12;t} $	$ \Delta f_{36;t} $	$ \Delta f_{36;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60:120;t} $	$ \Delta f_{60:120;t} $
VIX _t	0.00093	-0.00037	0.00046	-0.00037	0.00020	-6.0e-06	0.00015	-2.4e-08
π_m^{CB}	-0.065***	-0.056*	-0.038	-0.041	0.011	0.0013	0.0091	0.0012
Supr(π_m^{CB})	0.0072	-0.074*	-0.028	-0.12**	-0.016	-0.083*	-0.030	-0.078
g_m^{CB}	0.032**	0.0099	-0.018	-0.039*	-0.017	-0.029	-0.018	-0.023
Supr(g_m^{CB})	0.0051	0.0040	0.078***	0.036	0.087***	0.036	0.075***	0.019
$\tilde{y}_{m;24}^{CB}$	0.58	0.74	0.077	-0.34	-0.60	-1.54	-1.38	-2.63
$ \tilde{y}_{m;24}^{CB} - \tilde{y}_{m-1;24}^{CB} $	6.33**	13.6***	2.44	10.1***	-1.01	3.65	-0.92	1.58
$ \tilde{y}_{m;21}^{CB} - \tilde{y}_{m-1;24}^{CB} $	0.26	-6.19***	0.67	-3.99	-2.04	-2.68	-3.14	-1.94
Var(π_m^{CB})		-0.12**		-0.14		-0.073		-0.055
Δ Var(π_m^{CB})		0.070		0.041		-0.028		0.038
Skew(π_m^{CB})		0.045		-0.057		-0.079*		-0.043
Δ Skew(π_m^{CB})		0.042		0.072**		0.072**		0.037
Var(g_m^{CB})		0.11***		0.12***		0.067*		0.045
Δ Var(g_m^{CB})		-0.067**		-0.072*		-0.034		-0.033
Skew(g_m^{CB})		-0.25***		-0.16*		-0.052		-0.063
Δ Skew(g_m^{CB})		0.14***		0.049		-0.017		-0.020
Constant	-0.082**	-0.086**	0.055	0.053	0.063	0.063	0.075	0.071
R-squared	0.655	0.896	0.413	0.649	0.296	0.481	0.248	0.420
Partial R-squared	-	0.269	-	0.363	-	0.385	-	0.409
R-squared No Vix	0.640	0.895	0.409	0.648	0.295	0.481	0.248	0.420
Component	Total	Total	Total	Total	Total	Total	Total	Total

Notes: This table reports estimates from regressing absolute changes in market yields on the numeric forecast variables defined in Section 4.1.

Table 15: Results of the Information Tests on the Pre-ZLB Sample

	$ \Delta i_{0:12,t} $	$ \Delta f_{36,t} $	$ \Delta f_{60,t} $	$ \Delta f_{60:120,t} $
# of Selected Narrative Shocks	34	33	31	35
Permuted draws \geq Selected Narrative Shocks	0	0	1	0

Notes: The first row of this table reports the number of narrative shocks selected by leave-one-out cross validation using only the pre-ZLB sample and is analogous to the information in table 5. The second row shows the results of a permutation test for narrative news (as in figure 7); it reports the number of draws out of 500 that exceed the selected number of narrative shocks using the correct order.

draws. This table reassures us that it is unlikely to be the ZLB period that drives our results given that the top topics are very similar in the pre-ZLB period. Even more reassuringly, the information that explains the residual asset price news at the short-end of the yield curve is different to the signals that explain the residual movements at the longer maturities. This is shown in Table 17.

Table 16: Correlation between narrative signal inclusion probabilities in Pre-ZLB and Full Sample

Asset	$ \Delta i_{0:12,t} $	$ \Delta f_{36,t} $	$ \Delta f_{60,t} $	$ \Delta f_{60:120,t} $
Correlation	0.68***	0.48***	0.42***	0.62***

Notes: This table shows the correlation across the full sample and the pre-ZLB sample of the topic rankings measured using the topic variables' selection percentage from the bootstrap draws.

Table 17: Pearson Correlations of Narrative Signals' Selection Percentage Across Yields - Pre-ZLB period

	$ \Delta i_{0:12,t} $	$ \Delta f_{36,t} $	$ \Delta f_{60,t} $	$ \Delta f_{60:120,t} $
$ \Delta i_{0:12,t} $	1	.	.	.
$ \Delta f_{36,t} $	-0.02	1	.	.
$ \Delta f_{60,t} $	0.01	0.66***	1	.
$ \Delta f_{60:120,t} $	0.14	0.59***	0.88***	1

Notes: This table reports the Pearson correlation coefficient of the topics' selection percentages across 500 bootstrap draws for different yields and for different components of the yield curve. *** denotes significance at the 1% significance level.

6.2 Analysis of real interest rates

We have thus far focused on nominal rates due to lack of data availability for real rates but it is real rates that should have the biggest impact on real economic decisions such

as investment. In particular, real rates data is not available over much of our sample for the 1-year spot interest rate. We can, nonetheless, carry out the basic analysis on the available real yields data.

Figure 10 plots the news in equivalent real and nominal yields on IR publication dates in our sample for the three real rates. The most striking observation, consistent with the findings in Nakamura and Steinsson (2017), is that real and nominal yields move closely together. The correlation between the real and nominal market reactions are 0.81, 0.85 and 0.8 for the 3y, 5y and 5y5y-forward rate respectively. All are significant at the 1% significance level. The implication is that inflation expectations, at least at longer maturities, do not react too much in response to the IR.

We also repeat the information test and can show that the narrative shocks have explanatory power for the real market news residual. And the key narrative signals which drive the nominal yield curve are significantly correlated with the key narrative signals driving the real market reaction, as shown in Table 18. In particular, the signals that move the medium-run rates are uncorrelated with those that move long-run rates which is inconsistent with the standard expectations channel on real rates.

Table 18: Correlation Between Narrative Signal Inclusion Probabilities for Real and Nominal Yields

Asset	$ \Delta f_{36,t} $	$ \Delta f_{60,t} $	$ \Delta f_{60:120,t} $
Correlation	0.71***	0.6***	0.45***


Notes: This table shows the correlation across the real and nominal topic rankings measured using the topic variables' selection percentage from the bootstrap draws.

Finally, although we cannot test the difference of the signals driving the short end and the long end, the 3-year forward, 5-year forward and 5-year, 5-year forward rates correlation structure looks similar to that in the nominals, as shown in Table 19


Table 19: Pearson Correlations of Narrative Signals' Selection Percentage Across Yields - Real Yields

	$ \Delta r_{36,t} $	$ \Delta r_{60,t} $	$ \Delta r_{60:120,t} $
$ \Delta r_{36,t} $	1	.	.
$ \Delta r_{60,t} $	0.34***	1	.
$ \Delta r_{60:120,t} $	-0.01	0.71***	1


Notes: This table reports the Pearson correlation coefficient of the topics' selection percentages across 500 bootstrap draws for different yields and for different components of the yield curve. *** denotes significance at the 1% significance level.



(a) 3-year Forward



(b) 5-year Forward



(c) 5-year, 5-year Forward

Figure 10: Market News in Real and Nominal Yields on IR Days

7 Conclusion

Communication has offered an additional tool to central banks to move interest rates faced by banks, firms and households face across a variety of different maturities. One mechanism for this is through the central bank conveying information about economic conditions. So far the literature on this information effect has focused on signals about the expectations of the level of economic activity. Our results show that, in addition to this conventional expectations channel, signals about the expected uncertainty around economic conditions can give rise to important effects, especially at the long-run. Using a novel combination of theory, unstructured data and event studies, we find that this uncertainty channel plays a dominant role in moving interest rates at the long-end of the yield curve in response to publication of the Bank of England Inflation Report.

These results suggest that central banks wishing to influence long-term interest rates should take seriously the communication of the distribution of risks and uncertainties around economic conditions. Of course, more remains to be done to understand fully the policy implications of this channel of central bank communication and, in particular, how to design communication strategies with this in mind. For instance, earlier work on Delphic forward guidance, an approach adopted by many central banks in the last decade, has stressed the need to combine views on the future evolution of the economy together with a description of how monetary policy will react to these developments whereas our results suggest the possibility of policy-free forward guidance. This may be particularly helpful in periods when the central bank is confronted with an effective lower bound on short-term interest rates.²⁷

²⁷For example, Carvalho et al. (2016) find that once US interest rates reached their zero-lower bound, communication continued to have effects on longer-maturity bonds even when shorter-maturity bonds stopped responding.

References

- Andreasen, M. and Meldrum, A. (2015). Market beliefs about the UK monetary policy lift-off horizon: a no-arbitrage shadow rate term structure model approach. *Bank of England working papers*, 541.
- Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring Economic Policy Uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636.
- Bansal, R. and Shaliastovich, I. (2013). A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets. *The Review of Financial Studies*, 26(1):1–33.
- Bernanke, B. S. and Kuttner, K. N. (2005). What explains the stock market’s reaction to federal reserve policy? *Journal of Finance*, 60(3):1221–1257.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
- Blinder, A. S. (2008). Talking about Monetary Policy: The Virtues (and Vices?) of Central Bank Communication. Working Papers, Princeton University, Department of Economics, Center for Economic Policy Studies. 1048, Princeton University, Department of Economics, Center for Economic Policy Studies.
- Blinder, A. S. (2018). Through a crystal ball darkly: The future of monetary policy communication. *AEA Papers and Proceedings*, 108:567–71.
- Blinder, A. S., Ehrmann, M., Fratzscher, M., Haan, J. D., and Jansen, D.-J. (2008). Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence. *Journal of Economic Literature, American Economic Association*, 46(4):910–45.
- Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685.
- Boukous, E. and Rosenberg, J. (2006). The information content of FOMC minutes. Technical report, Federal Reserve Bank of New York.
- Bundick, B., Herriford, T., and Smith, A. L. (2017). Forward Guidance, Monetary Policy Uncertainty, and the Term Premium. Research Working Paper RWP 17-7, Federal Reserve Bank of Kansas City.
- Campbell, J., Evans, C., Fisher, J., and Justiniano, A. (2012). Macroeconomic Effects of Federal Reserve Forward Guidance. *The Brookings Papers on Economic Activity*, Spring:1–54.
- Carvalho, C., Hsu, E., and Nechio, F. (2016). Measuring the effect of the zero lower bound on monetary policy. Working Paper Series 2016-6, Federal Reserve Bank of San Francisco.
- Chang, J., Gerrish, S., Boyd-Graber, J. L., and Blei, D. M. (2009). Reading tea leaves: How humans interpret topic models. In *Advances in Neural Information Processing Systems*.
- Cieslak, A. and Schrimpf, A. (2018). Non-monetary news in central bank communication.

- Cloyne, J. and Hürtgen, P. (2016). The Macroeconomic Effects of Monetary Policy: A New Measure for the United Kingdom. *American Economic Journal: Macroeconomics*, 8(4):75–102.
- Cook, T. and Hahn, T. (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.
- Draghi, M. (2017). Remarks at Central Bank Communications Conference. ECB.
- Drechsler, I., Savov, A., and Schnabl, P. (2018). A model of monetary policy and risk premia. *The Journal of Finance*, 73(1):317–373.
- Ellison, M. and Tischbirek, A. (2018). Beauty Contests and the Term Structure. mimeo, University of Oxford.
- Fernandez-Villaverde, J., Guerron-Quintana, P., Rubio-Ramirez, J. F., and Uribe, M. (2011). Risk Matters: The Real Effects of Volatility Shocks. *American Economic Review*, 101(6):2530–2561.
- Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software*, 33(1).
- Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.
- Griffiths, T. L. and Steyvers, M. (2004). Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(Suppl. 1):5228–5235.
- Gürkaynak, R. S., Kisacikoglu, B., and Wright, J. H. (2018). Missing Events in Event Studies: Identifying the Effects of Partially-Measured News Surprises. Centre for Economic Policy Research (CEPR) DP13153, Centre for Economic Policy Research (CEPR).
- Gürkaynak, R. S., Sack, B., and Swanson, E. (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).
- Gürkaynak, R. S., Sack, B., and Swanson, E. (2005). The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American Economic Review*, 95(1):425–436.
- Haldane, A. and McMahon, M. (2018). Central bank communication and the general public. *AEA Papers and Proceedings*, 1(1):Forthcoming.
- Hansen, S. and McMahon, M. (2016). Shocking Language: Understanding the Macroeconomic Effects of Central Bank Communication. In *NBER International Seminar on Macroeconomics 2015*, NBER Chapters. National Bureau of Economic Research, Inc.
- Hansen, S., McMahon, M., and Prat, A. (2018). Transparency and deliberation within the fomc: A computational linguistics approach*. *The Quarterly Journal of Economics*, 133(2):801–870.

- Hanson, S. G. and Stein, J. C. (2015). Monetary policy and long-term real rates. *Journal of Financial Economics*, 115(3):429–448.
- Hastie, T., Tibshirani, R., and Wainwright, M. (2015). *Statistical Learning with Sparsity: The Lasso and Generalizations*. Number 143 in Monographs on Statistics and Applied Probability. CRC Press.
- Hendry, S. and Madeley, A. (2010). Text Mining and the Information Content of Bank of Canada Communications. Working Papers 10-31, Bank of Canada.
- Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring Uncertainty. *American Economic Review*, 105(3):1177–1216.
- 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.
- Laubach, T. and Williams, J. C. (2003). Measuring the Natural Rate of Interest. *The Review of Economics and Statistics*, 85(4):1063–1070.
- Leombroni, M., Vedolin, A., Venter, G., and Whelan, P. (2018). Central Bank Communication and the Yield Curve. CEPR Discussion Papers 12970, C.E.P.R. Discussion Papers.
- Malik, S. and Meldrum, A. (2016). Evaluating the robustness of UK term structure decompositions using linear regression methods. *Journal of Banking and Finance*, 67:85–102.
- Martin, I. and Ross, S. (2018). Notes on the yield curve. Working Papers 17-12, London School of Economics.
- Martin, I. W. R. (2013). Consumption-Based Asset Pricing with Higher Cumulants. *The Review of Economic Studies*, 80(2):745–773.
- Meinshausen, N. and Bühlmann, P. (2006). High-Dimensional Graphs and Variable Selection with the LASSO. *The Annals of Statistics*, 34(3):1436–1462.
- Miranda-Agrippino, S. and Ricco, G. (2015). The Transmission of Monetary Policy Shocks. Discussion Papers 1711, Centre for Macroeconomics (CFM).
- Nakamura, E. and Steinsson, J. (2013). High Frequency Identification of Monetary Non-Neutrality. NBER Working Papers 19260, National Bureau of Economic Research, Inc.
- Nakamura, E. and Steinsson, J. (2017). High Frequency Identification of Monetary Non-Neutrality: The Information Effect. *The Quarterly Journal of Economics*. Forthcoming.
- Reeves, R. and Sawicki, M. (2007). Do Financial Markets React to Bank of England Communication? *European Journal of Political Economy*, 23(1):207–227.
- Romer, C. D. and Romer, D. H. (2000). Federal reserve information and the behavior of interest rates. *American Economic Review*, 90(3):429–457.

- Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084.
- Shiller, R. J. (2017). Narrative Economics. *American Economic Review*, 107(4):967–1004.
- Tang, J. (2015). Uncertainty and the signaling channel of monetary policy. Working Papers 15-8, Federal Reserve Bank of Boston.
- Vlieghe, G. (2016). Monetary policy expectations and long term interest rates. Speech given by Gertjan Vlieghe, External MPC member, Bank of England at London Business School on 19 May 2016.
- Woodford, M. (2001). Monetary policy in the information economy. *Proceedings - Economic Policy Symposium - Jackson Hole*, pages 297–370.
- Zou, H. and Hastie, T. (2005). Regularization and Variable Selection via the Elastic Net. *Journal of the Royal Statistical Society Series B*, 67(2):301–320.

A Inflation Report Event Study

In this section, we conduct an event study to assess the average market impact of IR publication and other Bank of England communications. This extends the work of Reeves and Sawicki (2007), who conduct a similar analysis on a shorter sample. See Section A in the main text for related discussion. We define the following events within our sample period: (1) IR publication; (2) policy rate announcement; (3) speech by MPC member; (4) release of minutes of MPC meeting. We define a dummy variable for each event, and estimate the model

$$|\Delta\text{Yield}|_t = \alpha + \beta_1 D(\text{IR})_t + \beta_2 D(\text{Rate})_t + \beta_3 D(\text{Speech})_t + \beta_4 D(\text{Min})_t + \varepsilon_t \quad (\text{A.1})$$

for each yield. The estimated coefficients from ordinary least squares (OLS) estimates are in column (1) of Tables A.1a-A.2b. In columns (2)-(5) of these tables we estimate quantile regressions at various points in the distribution.

Confirming the visual evidence from the kernel densities in Section 2.3, at shorter maturities the IR is a dominant mover of market interest rates. The OLS coefficients for the one-year spot and three-year forward rates are both highly significant and approximately twice as large as the coefficient for policy announcements. There is a drop in significance for the five-year forward rate, but the magnitude of the IR coefficient is equivalent to that for policy announcements. This suggests a lack of power given there are three times as many announcements as IR dates over the sample period. However, there is a significant effect of IR releases in the right tail, as seen in column (5). For the five-year ahead, five year forward rate there is a marginally significant coefficient in column (5), and its magnitude is again the largest of any type of communication.

Table A.1: Estimated Coefficients of Event-Study Regression

(a) 1-year spot rate					
Main Regressors	(1) $ \Delta i_{0:12;t} $	(2) $ \Delta i_{0:12;t} $	(3) $ \Delta i_{0:12;t} $	(4) $ \Delta i_{0:12;t} $	(5) $ \Delta i_{0:12;t} $
IR	0.016*** [0.000]	0.0073 [0.317]	0.014*** [0.002]	0.019** [0.013]	0.031 [0.151]
Announcement	0.0084*** [0.002]	0.00076 [0.522]	0.0023 [0.236]	0.0033 [0.394]	0.038 [0.149]
Speech	-0.0022** [0.033]	-0.0013** [0.039]	-0.0019** [0.030]	-0.0024 [0.108]	-0.0034 [0.396]
Minutes	0.0046** [0.029]	-0.0011 [0.302]	0.0019 [0.344]	0.0065** [0.035]	0.024*** [0.002]
VIX _t	0.00098*** [0.000]	0.00027*** [0.000]	0.00063*** [0.000]	0.0013*** [0.000]	0.0030*** [0.000]
Constant	0.0024* [0.093]	0.0022*** [0.001]	0.0039*** [0.000]	0.0040** [0.015]	0.0022 [0.664]
R-squared	0.121				
Quantile	OLS	.25	.5	.75	.95
Sample	All	All	All	All	All

(b) 3-year forward rate					
Main Regressors	(1) $ \Delta f_{36;t} $	(2) $ \Delta f_{36;t} $	(3) $ \Delta f_{36;t} $	(4) $ \Delta f_{36;t} $	(5) $ \Delta f_{36;t} $
IR	0.018*** [0.005]	0.0041 [0.327]	0.011 [0.137]	0.027** [0.050]	0.061*** [0.000]
Announcement	0.0071*** [0.007]	0.0027 [0.274]	0.011*** [0.000]	0.0097*** [0.004]	0.0088 [0.339]
Speech	0.0039** [0.030]	-0.0016 [0.143]	0.0025 [0.312]	0.0068*** [0.008]	0.022*** [0.005]
Minutes	0.0042 [0.109]	0.0016 [0.404]	0.0028 [0.464]	0.012** [0.031]	-0.0027 [0.591]
VIX _t	0.00092*** [0.000]	0.00025*** [0.000]	0.00091*** [0.000]	0.0013*** [0.000]	0.0030*** [0.000]
Constant	0.022*** [0.000]	0.0098*** [0.000]	0.014*** [0.000]	0.030*** [0.000]	0.045*** [0.000]
R-squared	0.053				
Quantile	OLS	.25	.5	.75	.95
Sample	All	All	All	All	All

Notes: These tables report quantile regressions to examine the effect on market interest rates according to whether (1) an IR is released, (2) a policy decision from the MPC is announced, (3) an MPC member makes a public speech, (4) minutes from MPC meetings are released, or (5) none of the above. These tables complement the kernel densities in figure 3.

Table A.2: Estimated Coefficients of Event-Study Regression

(a) 5-year forward rate

	(1)	(2)	(3)	(4)	(5)
Main Regressors	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $
IR	0.0065 [0.271]	0.0019 [0.578]	-0.0039 [0.449]	0.0069 [0.557]	0.043*** [0.002]
Announcement	0.0063* [0.055]	0.0032 [0.233]	0.0059 [0.106]	0.010*** [0.004]	0.0032 [0.857]
Speech	0.0044** [0.023]	-0.00041 [0.753]	0.0030 [0.181]	0.0057** [0.042]	0.025*** [0.000]
Minutes	0.0037 [0.159]	0.0046 [0.112]	0.0050* [0.096]	0.0040 [0.283]	-0.0055* [0.075]
VIX _t	0.0010*** [0.000]	0.00026*** [0.000]	0.00071*** [0.000]	0.0014*** [0.000]	0.0028*** [0.000]
Constant	0.021*** [0.000]	0.0097*** [0.000]	0.019*** [0.000]	0.031*** [0.000]	0.052*** [0.000]
R-squared	0.052				
Quantile	OLS	.25	.5	.75	.95
Sample	All	All	All	All	All

(b) 5-year ahead, 5-year forward rate

	(1)	(2)	(3)	(4)	(5)
Main Regressors	$ \Delta f_{60;120;t} $	$ \Delta f_{60;120;t} $	$ \Delta f_{60;120;t} $	$ \Delta f_{60;120;t} $	$ \Delta f_{60;120;t} $
IR	0.0021 [0.688]	-0.0015 [0.592]	-0.0014 [0.798]	-0.0014 [0.889]	0.037* [0.074]
Announcement	0.0049 [0.154]	0.000084 [0.978]	0.0014 [0.641]	0.0078 [0.179]	0.0046 [0.669]
Speech	0.0035* [0.070]	-0.00050 [0.685]	3.7e-06 [0.998]	0.0026 [0.437]	0.021*** [0.000]
Minutes	0.0036 [0.158]	0.0065* [0.081]	0.0047* [0.051]	0.0043 [0.491]	-0.0012 [0.838]
VIX _t	0.00097*** [0.000]	0.00027*** [0.000]	0.00061*** [0.000]	0.0013*** [0.000]	0.0027*** [0.000]
Constant	0.021*** [0.000]	0.0094*** [0.000]	0.020*** [0.000]	0.032*** [0.000]	0.051*** [0.000]
R-squared	0.049				
Quantile	OLS	.25	.5	.75	.95
Sample	All	All	All	All	All

Notes: These tables report quantile regressions to examine the effect on market interest rates according to whether (1) an IR is released, (2) a policy decision from the MPC is announced, (3) an MPC member makes a public speech, (4) minutes from MPC meetings are released, or (5) none of the above. These tables complement the kernel densities in figure 3.

B Omitted Proofs

B.1 Proof of Proposition 1

Proof. The first step is to take the expectation of (13) conditional on I_t^{MK} . Observe that

$$\mathbb{E}[v_{m+i} \mid I_t^{\text{MK}}] = \widehat{\mu}_{m+i,m}^{\text{CB}} \text{ for } i = 1, \dots, h.$$

Also, because the central bank is Bayesian, we have

$$\mathbb{E}[\widehat{\mu}_{m+i,m+k}^{\text{CB}} \mid I_t^{\text{MK}}] = \widehat{\mu}_{m+i,m}^{\text{CB}} \text{ for } i = 1, \dots, h; k < i.$$

Outside the current forecast horizon ($i > h$), we have $\mathbb{E}[v_{m+i} \mid I_t^{\text{MK}}] = \mathbb{E}[\widehat{\mu}_{m+i,m+k}^{\text{CB}} \mid I_t^{\text{MK}}] = 0$ since the IR in month m contains no relevant news about cyclical shocks more than h months ahead. Expectations over them are therefore computed using the zero-mean prior distribution from which μ is drawn.

Combining these observations gives

$$\mathbb{E}[\mathbb{E}[\omega_{m+k+h} \mid I_{m+k}^{\text{CB}}] \mid I_t^{\text{MK}}] = \rho^{k+h}\omega_m + \sum_{i=1}^h \rho^{k+h-i}\widehat{\mu}_{m+i,m}^{\text{CB}}.$$

By the law of iterated expectations we also obtain

$$\mathbb{E}[\mathbb{E}[\omega_{m+k+h} \mid I_{m+k}^{\text{CB}}] \mid I_{t-1}^{\text{MK}}] = \rho^{k+h}\omega_m + \sum_{i=1}^h \rho^{k+h-i}\mathbb{E}[\widehat{\mu}_{m+i,m}^{\text{CB}} \mid \widehat{\mu}_{m+i,t-1}^{\text{MK}}, i_m].$$

The proposition follows immediately. ■

B.2 Proof of Proposition 3

Proof. We begin with basic properties of the lognormal distribution. By expanding the stochastic process for volatility (8), we obtain

$$\log \sigma_{m+h+k}^2 = \rho_\sigma^k \log \sigma_{m+h}^2 + (1 - \rho_\sigma^k) \log \sigma_0^2 + \sum_{i=1}^k \rho_\sigma^{k-i} u_{m+h+i}$$

So $\log \sigma_{m+h+k}^2 \mid \sigma_{m+h}^2$ is normally distributed with mean $\rho_\sigma^k \log \sigma_{m+h}^2 + (1 - \rho_\sigma^k) \log \sigma_0^2$ and variance $\sum_{i=1}^k \rho_\sigma^{2(k-i)} \sigma_u^2$. Therefore we obtain

$$\mathbb{E}[\sigma_{m+h+k}^2 \mid \sigma_{m+h}^2] = \exp \left[\rho_\sigma^k \log \sigma_{m+h}^2 + (1 - \rho_\sigma^k) \log \sigma_0^2 + \frac{\sigma_u^2}{2} \sum_{i=1}^k \rho_\sigma^{2(k-i)} \right]$$

from which we obtain

$$\lim_{\rho_\sigma \rightarrow 1} \mathbb{E}[\sigma_{m+h+k}^2 \mid \sigma_{m+h}^2] = \exp \left[\log \sigma_{m+h}^2 + \frac{k\sigma_u^2}{2} \right] = \sigma_{m+h}^2 \exp \left[\frac{k\sigma_u^2}{2} \right].$$

By similar arguments

$$\lim_{\rho_\sigma \rightarrow 1} \mathbb{E} \left[\sigma_{m+h+k}^2 \mid \sigma_{m+h-1}^2 \right] = \sigma_{m+h-1}^2 \exp \left[\frac{(k+1)\sigma_u^2}{2} \right].$$

We begin by characterizing the variance of (13) conditional on I_t^{MK} . First note that the variance of $\hat{\mu}_{m+i, m+k}^{\text{CB}}$ does not depend on IR publication whenever $i > k \geq h$ because the IR in month m contains no information on the central bank's forecast more than h months ahead. It is therefore sufficient to compute

$$\text{Var} \left[\sum_{i=1}^k \rho^{k+h-i} v_{m+i} \mid I_t^{\text{MK}} \right] = \sum_{i=1}^k \rho^{2(k+h-i)} \text{Var} [v_{m+i} \mid I_t^{\text{MK}}]$$

where we can further expand

$$\text{Var} [v_{m+i} \mid I_t^{\text{MK}}] = \text{Var} [\mu_{m+i} \mid I_t^{\text{MK}}] + \text{Var} [\varepsilon_{m+i} \mid I_t^{\text{MK}}].$$

When $i \leq h$ these conditional variances are by assumption $\text{Var} [\mu_{m+i} \mid I_t^{\text{MK}}] = (s_{m+i,t}^{\text{MK}})^2$ and $\text{Var} [\varepsilon_{m+i} \mid I_t^{\text{MK}}] = \sigma_{m+i}^2$.

When $i > h$, we have $\text{Var} [\mu_{m+i} \mid I_t^{\text{MK}}] = s^2$ since no information is available on μ_{m+i} beyond the prior distribution, whose variance is s^2 . The conditional variance of ε_{m+i} can be decomposed by the law of total variance as

$$\begin{aligned} \text{Var} [\varepsilon_{m+i} \mid I_t^{\text{MK}}] &= \mathbb{E} [\text{Var} [\varepsilon_{m+i} \mid I_t^{\text{MK}}, \sigma_{m+i}^2] \mid I_t^{\text{MK}}] + \\ &\quad \text{Var} [\mathbb{E} [\varepsilon_{m+i} \mid I_t^{\text{MK}}, \sigma_{m+i}^2] \mid I_t^{\text{MK}}]. \end{aligned}$$

The second term is zero since this is the mean of the fundamental uncertainty shock in all periods. For the first term, we can use the results above to obtain

$$\lim_{\rho_\sigma \rightarrow 1} \text{Var} [\varepsilon_{m+i} \mid I_t^{\text{MK}}] = \lim_{\rho_\sigma \rightarrow 1} \mathbb{E} [\sigma_{m+i}^2 \mid \sigma_{m+h}^2] = \sigma_{m+h}^2 \exp \left[\frac{(i-h)\sigma_u^2}{2} \right].$$

Combining these results together yields²⁸

$$\begin{aligned} &\text{Var} \left[\sum_{i=1}^k \rho^{k+h-i} v_{m+i} \mid I_t^{\text{MK}} \right] = \\ &\sum_{i=1}^h \rho^{2(k+h-i)} \left((s_{m+i,t}^{\text{MK}})^2 + \sigma_{m+i}^2 \right) + \sum_{i=h+1}^k \rho^{2(k+h-i)} \left(s^2 + \sigma_{m+h}^2 \exp \left[\frac{(i-h)\sigma_u^2}{2} \right] \right). \end{aligned}$$

²⁸This expression is valid for $k > h$. For $k = h$ the correct expression is simply the first term in the sum.

The equivalent expression for the variance conditional on I_{t-1}^{MK} is

$$\begin{aligned} \text{Var} \left[\sum_{i=1}^k \rho^{k+h-i} v_{m+i} \mid I_{t-1}^{\text{MK}} \right] &= \sum_{i=1}^{h-1} \rho^{2(k+h-i)} \left((s_{m+i,t-1}^{\text{MK}})^2 + \sigma_{m+i}^2 \right) + \\ &\quad \rho^{2k} \left((s_{m+h,t-1}^{\text{MK}})^2 + \sigma_{m+h-1}^2 \exp \left[\frac{\sigma_u^2}{2} \right] \right) + \\ &\quad \sum_{i=h+1}^k \rho^{2(k+h-i)} \left(s^2 + \sigma_{m+h-1}^2 \exp \left[\frac{(i-h)\sigma_u^2}{2} \right] \exp \left[\frac{\sigma_u^2}{2} \right] \right). \end{aligned}$$

The statement of the proposition then follows directly. ■

C Details of Cross Validation Procedure

Here we detail the cross-validation procedure we use to select λ in our estimation of the elastic net regressions in the paper. Given our small sample size, leave-one-out cross validation is computationally feasible and we adopt it. The specific algorithm is:

1. For each of a sequence of possible λ penalty coefficients:
 - (a) For each of the N data points:
 - i. Remove the point from the sample.
 - ii. Fit (17) on the remaining $N - 1$ points.
 - iii. Calculate the forecasted value for the held-out point from the fitted model, and compute the squared error.
2. Select the highest value of λ that has a mean squared error (MSE) within one standard deviation of the MSE-minimizing λ across N out-of-sample forecasts.

The model selection rule at stage 2 is sparser than the model with the most accurate out-of-sample predictive power because, as λ increases, the elastic net selects fewer covariates. This increases our confidence that any selected text shock variables have a robust relationship with market interest rates.

D Supplementary Tables

Table D.1: Contribution of Numeric vs. Narrative Variables to R^2

	(1)	(2)	(3)	(4)	(5)	(6)
	$ \Delta i_{0:12;t} $	$ \Delta i_{0:12;t} $	$ \Delta i_{0:12;t} $	$ \Delta f_{36;t} $	$ \Delta f_{36;t} $	$ \Delta f_{36;t} $
VIX _t	0.0024*** [0.00]	0.0013** [0.02]	0.00050 [0.34]	0.0023** [0.03]	0.0023 [0.15]	0.0037** [0.02]
Topic L25			1.80*** [0.01]			
Topic D24			-0.66 [0.13]			-1.39 [0.16]
Topic L5			-1.32** [0.02]			
Topic L26			-1.34** [0.03]			
Topic D25						-3.18 [0.10]
Topic L28						1.74** [0.03]
Topic D23						1.17 [0.21]
Constant	-0.011 [0.41]	-0.070* [0.08]	0.014 [0.81]	0.010 [0.61]	-0.026 [0.69]	-0.065 [0.32]
R-squared	0.286	0.563	0.693	0.150	0.368	0.513
Include Vix	Yes	Yes	Yes	Yes	Yes	Yes
Include \mathbf{q}_t	No	Yes	Yes	No	Yes	Yes
Partial R^2	.	0.39	0.29	.	0.26	0.23
F-test p-value	.	0.045	0.001	.	0.020	0.070

Notes: Columns (1)-(3) ((4)-(6)) show how much market news for $|\Delta i_{0:12;t}|$ ($|\Delta f_{36;t}|$) can be explained by adding numeric in (2) and then, in (3), numeric and narrative information captured by the top 4 topics. This information is reflect in table 11 in the main text.

Table D.2: Contribution of Numeric vs. Narrative Variables to R^2

	(1)	(2)	(3)	(4)	(5)	(6)
	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60:120;t} $	$ \Delta f_{60:120;t} $	$ \Delta f_{60:120;t} $
VIX _t	0.0014 [0.20]	0.0018 [0.23]	0.0035*** [0.01]	0.0011 [0.19]	0.0017 [0.15]	0.0030*** [0.00]
Topic L28			1.48** [0.02]			
Topic D17			1.41* [0.08]			1.77** [0.01]
Topic D18			-2.39 [0.23]			-3.34* [0.06]
Topic L20			1.39* [0.08]			2.07*** [0.00]
Topic D13						-3.26*** [0.01]
Constant	0.020 [0.33]	0.014 [0.83]	-0.053 [0.44]	0.020 [0.26]	0.043 [0.46]	-0.0070 [0.89]
R-squared	0.065	0.280	0.502	0.058	0.274	0.542
Include Vix	Yes	Yes	Yes	Yes	Yes	Yes
Include \mathbf{q}_t	No	Yes	Yes	No	Yes	Yes
Partial R^2	.	0.23	0.31	.	0.23	0.37
F-test p-value	.	0.000	0.023	.	0.004	0.004

Notes: Columns (1)-(3) ((4)-(6)) show how much market news for $|\Delta f_{60;t}|$ ($|\Delta f_{60:120;t}|$) can be explained by adding numeric in (2) and then, in (3), numeric and narrative information captured by the top 4 topics. This information is reflect in table 11 in the main text.

Table D.3: Effect of Forecast Variables on Market Yields by Expectations and Term Premium

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$ Exp1y $	$ TP1y $	$ Exp3y $	$ TP3y $	$ Exp5y $	$ TP5y $	$ Exp5y $	$ TP5y $
π_m^{CB}	-0.016 [0.515]	0.0018 [0.544]	-0.027 [0.338]	0.0045 [0.773]	-0.017 [0.437]	0.018 [0.304]	-0.021 [0.294]	0.031 [0.161]
$Supr(\pi_m^{CB})$	0.040 [0.162]	-0.00034 [0.950]	-0.014 [0.669]	-0.019 [0.280]	-0.015 [0.568]	-0.020 [0.256]	-0.012 [0.617]	-0.019 [0.348]
$Var(\pi_m^{CB})$	-0.0100 [0.313]	-0.0019 [0.203]	0.013 [0.305]	0.0084 [0.241]	0.012 [0.242]	0.0091 [0.262]	0.0063 [0.458]	0.012 [0.222]
$\Delta Var(\pi_m^{CB})$	0.036 [0.400]	0.0027 [0.780]	0.0078 [0.851]	-0.035 [0.153]	-0.014 [0.676]	0.020 [0.624]	0.0070 [0.791]	0.046 [0.397]
$Skew(\pi_m^{CB})$	-0.00092 [0.975]	0.0011 [0.880]	0.021 [0.526]	0.0016 [0.935]	0.012 [0.646]	0.013 [0.564]	0.0051 [0.844]	0.018 [0.522]
$\Delta Skew(\pi_m^{CB})$	0.017 [0.609]	0.0077 [0.256]	0.035 [0.287]	0.0076 [0.708]	0.041 [0.128]	-0.0014 [0.945]	0.025 [0.291]	-0.00030 [0.989]
g_m^{CB}	0.028* [0.052]	0.0012 [0.573]	0.012 [0.308]	-0.0079 [0.342]	0.0017 [0.867]	-0.0094 [0.415]	0.0014 [0.882]	-0.0075 [0.599]
$Supr(g_m^{CB})$	-0.049** [0.040]	-0.0025 [0.618]	0.018 [0.455]	0.019 [0.177]	0.027 [0.206]	0.0031 [0.870]	0.016 [0.384]	-0.017 [0.469]
$Var(g_m^{CB})$	0.0019 [0.860]	-0.00081 [0.690]	-0.0021 [0.867]	-0.0051 [0.482]	-0.00050 [0.962]	-0.0082 [0.325]	-0.00093 [0.915]	-0.013 [0.186]
$\Delta Var(g_m^{CB})$	0.0073 [0.608]	0.0021 [0.267]	0.0080 [0.589]	0.0067 [0.339]	0.0093 [0.335]	0.0029 [0.715]	0.0057 [0.548]	0.0083 [0.439]
$Skew(g_m^{CB})$	-0.0080 [0.792]	-0.0033 [0.582]	-0.039 [0.264]	-0.027 [0.157]	-0.049 [0.111]	-0.015 [0.463]	-0.034 [0.183]	-0.015 [0.514]
$\Delta Skew(g_m^{CB})$	-0.056* [0.053]	-0.0060 [0.344]	-0.011 [0.696]	0.011 [0.526]	0.00064 [0.979]	-0.0016 [0.928]	0.00020 [0.992]	-0.019 [0.393]
$\tilde{y}_{m;24}^{CB}$	-0.034 [0.975]	0.19 [0.335]	2.44* [0.080]	0.32 [0.696]	2.60** [0.025]	-0.94 [0.219]	1.64 [0.105]	-1.93** [0.040]
$ \tilde{y}_{m;24}^{CB} - \tilde{y}_{m-1;24}^{CB} $	3.41* [0.062]	0.99** [0.016]	2.88 [0.204]	-0.79 [0.607]	2.13 [0.276]	0.0058 [0.997]	1.70 [0.294]	0.97 [0.507]
$ \tilde{y}_{m;21}^{CB} - \tilde{y}_{m-1;24}^{CB} $	1.09 [0.555]	0.64 [0.105]	-0.29 [0.903]	0.15 [0.924]	-0.17 [0.933]	0.98 [0.548]	0.24 [0.890]	1.85 [0.327]
VIX _t	0.0010* [0.059]	0.00030** [0.012]	0.0014 [0.147]	0.00078 [0.253]	0.0011 [0.201]	0.00087 [0.168]	0.00096 [0.180]	0.00075 [0.112]
Constant	-0.060 [0.113]	-0.0019 [0.746]	-0.042 [0.311]	0.021 [0.478]	-0.019 [0.612]	0.035 [0.323]	-0.010 [0.752]	0.044 [0.266]
R-squared	0.544	0.457	0.443	0.168	0.429	0.227	0.408	0.280
R-squared No Vix	0.513	0.390	0.396	0.115	0.387	0.172	0.366	0.252

Notes: This table reports estimates from regressing absolute changes in market yields on the numeric forecast variables defined in Section 4.1.

Table D.4: Effect of Top Expectations and Term Premium Narrative Shocks on R^2

	(1)	(2)	(3)	(4)	(5)	(6)
	$ \Delta i_{0:12;t} $	$ \Delta i_{0:12;t} $	$ \Delta i_{0:12;t} $	$ \Delta f_{36;t} $	$ \Delta f_{36;t} $	$ \Delta f_{36;t} $
VIX _t	0.00050 [0.34]	0.00038 [0.44]	0.0012** [0.03]	0.0037** [0.02]	0.0029 [0.11]	0.0036** [0.04]
Topic L25	1.80*** [0.01]	1.61** [0.01]			1.01 [0.42]	
Topic D24	-0.66 [0.13]		-0.49 [0.37]	-1.42 [0.11]	-1.21 [0.22]	
Topic L5	-1.32** [0.02]	-1.23** [0.03]				
Topic L26	-1.34** [0.03]	-1.86*** [0.01]				
Topic D28		1.25 [0.18]				
Topic L9			-0.21 [0.68]			
Topic D3			2.17 [0.19]			
Topic D7			1.58 [0.17]			
Topic D25				-2.91 [0.14]	-3.65** [0.05]	
Topic L28				1.35* [0.05]		1.23 [0.14]
Topic D14				2.00 [0.28]		
Topic D1					2.57** [0.04]	
Topic D17						0.55 [0.56]
Topic L7						-0.46 [0.63]
Topic D18						-1.32 [0.56]
Constant	0.014 [0.81]	0.046 [0.49]	-0.082* [0.06]	-0.058 [0.35]	-0.075 [0.25]	-0.039 [0.58]
R-squared	0.693	0.699	0.611	0.510	0.461	0.457
Topics	Overall	Expectation	Term Premium	Overall	Expectation	Term Premium
Include Vix	Yes	Yes	Yes	Yes	Yes	Yes
Include \mathbf{q}_t	Yes	Yes	Yes	Yes	Yes	Yes
Partial R^2	0.29	0.31	0.11	0.23	0.15	0.15
F-test p-value	0.001	0.002	0.234	0.120	0.006	0.353

Notes: Columns (1)-(3) [(4)-(6)] show how much market news for $|\Delta i_{0:12;t}|$ ($|\Delta f_{36;t}|$) can be explained by the overall yield top narrative shocks in (1) [(4)], the top narrative shocks for the yield expectations component in (2) [(5)], and the top narrative shocks for the term premium component in (3) [(6)].

Table D.5: Effect of Top Expectations and Term Premium Narrative Shocks on R^2

	(1)	(2)	(3)	(4)	(5)	(6)
	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60;t} $	$ \Delta f_{60:120;t} $	$ \Delta f_{60:120;t} $	$ \Delta f_{60:120;t} $
VIX _t	0.0035*** [0.01]	0.0033** [0.04]	0.0022 [0.14]	0.0030*** [0.00]	0.0021* [0.10]	0.0024** [0.04]
Topic L28	1.48** [0.02]	1.66** [0.02]				
Topic D17	1.41* [0.08]		1.53 [0.11]	1.77** [0.01]		1.54** [0.05]
Topic D18	-2.39 [0.23]		-2.71 [0.20]	-3.34* [0.06]		-3.30* [0.09]
Topic L20	1.39* [0.08]			2.07*** [0.00]		
Topic D24		-0.49 [0.62]			0.036 [0.97]	
Topic D1		0.47 [0.71]			0.20 [0.87]	
Topic D25		-1.93 [0.27]			-2.09 [0.19]	
Topic D13			-2.92* [0.07]	-3.26*** [0.01]		-3.37** [0.02]
Topic D20			0.75 [0.51]			
Topic D8					0.67 [0.33]	0.88 [0.14]
Constant	-0.053 [0.44]	-0.032 [0.63]	0.023 [0.72]	-0.0070 [0.89]	0.017 [0.79]	0.039 [0.49]
R-squared	0.502	0.425	0.385	0.542	0.303	0.439
Topics	Overall	Expectation	Term Premium	Overall	Expectation	Term Premium
Include Vix	Yes	Yes	Yes	Yes	Yes	Yes
Include \mathbf{q}_t	Yes	Yes	Yes	Yes	Yes	Yes
Partial R^2	0.31	0.19	0.14	0.37	0.04	0.23
F-test p-value	0.023	0.086	0.242	0.004	0.672	0.066

Notes: Columns (1)-(3) ((4)-(6)) show how much market news for $|\Delta f_{60;t}|$ ($|\Delta f_{60:120;t}|$) can be explained by the overall yield top narrative shocks in (1) [(4)], the top narrative shocks for the yield expectations component in (2) [(5)], and the top narrative shocks for the term premium component in (3) [(6)].