The Long-Run Information Effect of Central Bank Narrative[*]

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

March 31, 2018

Abstract

Central bank communication is an increasingly important policy tool, but its impact on market expectations arises through many potential channels. We use the publication of the Bank of England's Inflation Report to study the market impact of communication about expected future economic conditions can move short- and long-run market rates. We find compelling evidence for a long-run information effect driven more by narrative than by quantitative forecasts. The dominant channel of this effect operates largely via its impact on the term premium through narrative about risks and uncertainty. We conclude that central banks can impact long-run interest rates without making explicit future policy commitments.

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, Michael Ehrmann and Jonathan Talmi as well as seminar participants at the Banco do Brazil, Bank of England, ECB, Empirical Macroeconomics Conference at Ghent University, Higher School of Economics, New Economic School, Norges Bank, Oxford/NY Fed Conference, RCEA conference, and UPenn. 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: `michael.mcmahon@economics.ox.ac.uk`

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

1 Introduction

Central banks in the major economies communicate frequently nowadays, but this is a relatively recent phenomenon. Before January 1994, the Federal Open Market Committee (FOMC) didn't even immediately announce their monetary policy decisions. The main reason for the increase in communication appears to be a realization among central banks that expectations are a key link between their policy instruments and the interest rates that really matter for economic activity, and that communication could aid expectations, and hence economic, management (Woodford 2001, Blinder 2008). While initially the focus was on building credibility through accountability and the adoption of specific frameworks such as inflation targeting (a form of central bank communication), more recent interest has focused on greater transparency and more frequent communication.

There is well-established evidence that central bank communication does indeed drive interest rates of different maturities (Gürkaynak et al. 2005, Boukus and Rosenberg 2006, Blinder et al. 2008, Carvalho et al. 2016, for example), but the channels through which that occurs are often unclear. The same communication event can transmit information about the current policy stance, the central bank's views on economic conditions, and the central bank's future policy stance (Campbell et al. 2012). Moreover, central banks often release several pieces of information simultaneously – for example, an announcement of current policy and an accompanying statement. The extensive literature that measures monetary policy news based on high-frequency movements in short-term policy rates (Gertler and Karadi 2015, Nakamura and Steinsson 2017, Jarociński and Karadi 2017, for example) captures the aggregate impact of these separate, heterogeneous information sources. Each source can itself have multiple effects: an announcement of a policy rate decision both informs markets of the current policy stance and can signal the central bank's views of economic conditions that led to that stance (Romer and Romer 2000, Melosi 2017, Zhang 2017).

One specific example of the importance of understanding the channels through which communication operates is the ongoing debate about why long-run market interest rates react to short-run monetary news. Nakamura and Steinsson (2017) argue that this is due to monetary policy shocks transmitting information about economic fundamentals that affects long-run market expectations, 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 behavioral response 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. Given the current set of empirical methods in the literature, it is a serious challenge to precisely identify the transmission mechanism and duration of the information effect.

In this paper we study a controlled environment in which we isolate communication solely about economic fundamentals. We show that quantitative forecast information plays a relatively important role in moving interest rates at the short-end of the yield curve. However, we show that the narrative of the communication plays a relatively larger role at longer maturities. Moreover, this effect comes increasingly through term premiums for longer-maturity assets. Importantly, we find that different information affects the short- and long-end of the yield curve — there is, in addition to shorter-term news, a distinct long-term information effect of the central bank narrative. Our analysis suggests this new channel is the dominant one, at least in our sample.

Specifically, we exploit the unique environment of 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.[1]

The IR is also of interest due to its rich information structure. It contains detailed, quantitative inflation and output forecasts in the form of fan charts that place probabilities on different future realizations that are summarized with mode, variance, and skew. Much of the information effect literature focuses on markets' updating their beliefs based on such quantitative information. The IR also contains rich, narrative information in the form of text data that both describes recent economic developments and the near-term outlook, as well as the forecast itself. This structure allows us to assess the market impact of narrative information while controlling for the impact quantitative information.

To study the effect of the IR publication, we adopt an event-study methodology made popular by Cook and Hahn (1989), Kuttner (2001) and Bernanke and Kuttner (2005). We use daily data on four nominal rates at varying maturities 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. We then analyze the extent to which information in the IR can explain the absolute values of changes in these rates on the publication day, which is our measure of news. We find that the quantitative information has strong explanatory power for explaining the one-year spot rate, but much less for

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

longer rates. This is consistent with information on the quantitative forecasts' having an information effect at the short run, but casts doubt on whether it can substantially explain long-run rate movements.

We next ask whether narrative information can explain the residual moves in asset prices, conditional on the quantitative information published. We reduce the dimensionality of the text with Latent Dirichlet Allocation (LDA), a popular probabilistic topic model (Blei et al. 2003) previously applied in monetary policy by Hansen and McMahon (2016) and Hansen et al. (2017). 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 so on. Using these LDA topics, we create text shocks by regressing topic shares on the same quantitative forecast information we use to explain asset price movements, and back out the residual. This construction is the textual analogue of the method for computing monetary policy shocks of Romer and Romer (2004), and it yields text shocks that purge from language the endogenous variation due to the current economic outlook. We then face a high-dimensional statistical problem: there are many potentially relevant text characteristics and only seventy (quarterly) publication dates in our sample. To test for information content, we examine the number of significant characteristics identified by the Least Absolute Shrinkage and Selection Operator (LASSO) via cross validation and compare that to the number identified when we randomly permute the data a re-estimate the LASSO. We overwhelmingly reject the hypothesis that narrative information is unrelated to the size of the interest rate moves at both short and long-term maturities.

While a useful starting point, our baseline test is not sufficient to conclude that the narrative has a direct informational impact on long-run rates rather than simply being an information effect on short-term rates propagated to long-term rates by market mechanics. Instead, we need to explore *what* narrative information drives short- and long-term rates. In particular, the main hypotheses in the existing literature point to the information effects at all maturities being driven by the same information. In Hanson and Stein (2015), short-run monetary news is propagated to long-run term premiums while in Nakamura and Steinsson (2017), the central bank communicates expectations about a persistent variable.

To explore the content that drives the information effect, we use a bootstrapping procedure to rank topics separately for each asset maturity according to their ability to robustly predict price movements. The ranking of topics for the one-year spot rate is uncorrelated with the ranking for all longer-term rates. Moreover, the topics that explain longer-term rate movements appear to be more associated with narrative around the forecast rather than with narrative around current economic conditions. This evidence points to a long-run information effect mediated by the text narrative separate from any effect due to quantitative forecasts. Moreover, the effect seems to be driven by different content at the long and short ends. We suggest a new channel related to communication about uncertainty in the economy.

Finally, we decompose use four affine term structure models to decompose rate changes into expectation and term premium components and repeat our analysis by component. We show that the long-run information effect of central bank text is driven more by changes in the long-run term premium than long-run expectations of the modal rate. Our baseline test of information indicates that both components react to IR information at all maturities. But the narrative appears particularly important for explaining movements in the long-run term premium, with just a handful of key topics driving more variation than the more than one dozen quantitative variables. This is consistent with theories of macroeconomic conditions driving the term premium (Cochrane 2011, Bansal and Shaliastovich 2013, Martin 2013), and in particular with information about long-term risks around the modal Bank of England forecasts driving longer-maturity term premiums in the UK. This supports our proposed new channel which we find to be the dominant one.

These findings point to several important contributions. First, we further the academic debate on the long-run information effect by showing that communication about economic fundamentals alone can move long-term interest rates. We also show that this effect manifests itself particularly via the impact of narrative information on the long-run term premium, which is a novel channel. In other words, we would not necessarily expect a long-run information effect in the absence of narrative information, and the observed movement in long-run rates need not correspond to a change in modal expectations of future central bank policy.

Second, an important policy implication of our work is that communication solely about future economic conditions is sufficient to move long-run interest rates. Delphic forward guidance as formulated in Campbell et al. (2012) refers to transmitting a view on future conditions and at the same time describing a likely policy response to that view. We show that the latter component is not necessary for generating market effects. Another policy insight is that, while the channel in Hanson and Stein (2015) requires policymakers to generate a change in short-term interest rates in order to move long-term interest rates, our channel simply requires communicating economic conditions. We therefore expect it to be effective even in the absence of a change in short-term interest rates or in periods when they are constrained by a lower bound. 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.

Third, we make a methodological contribution by proposing a new test of long-run information in mixed data. The pairing of low-dimensional quantitative information with high-dimensional text data is not unique to the IR. As mentioned above, the Fed releases a textual statement along with its policy rate decision after each FOMC meetings. One could apply our empirical strategy directly to such events to see whether they transmit long-run information.[2] The economic interpretation of such information would be potentially different to that in the IR, but nevertheless our test of its presence could advance the broader empirical literature beyond this paper.[3]

The remainder of this paper is structured as follows. In Section 2, we provide more detail on the IR and evidence that IR publication days tend to give rise to market news. Section 3 then describes how we convert the IR into quantitative measures to conduct our empirical analysis, which we then use in section 4 to assess, separately, the market impact of the quantitative and narrative information in the Inflation Report. Section 5 uses a variety of tests to distinguish between potential driving channels of our long-maturity information effect. We then conclude.

2 Communication and the Yield Curve

In this section we start by presenting a simple framework to help think about the potential channels through which central bank communication can have an information effect on market interest rates, before going on to explain the interest rate data we use and establishing that the Bank of England's Inflation Report does provide significant market news. We 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, we mean by information effect any systematic market reaction to communication about economic fundamentals via any medium.

[2] Gürkaynak et al. (2005) show that the market reaction to Fed statements can be summarized with two factors, one associated with short-run rate movements and another with long-run movements. While they interpret the second factor as a response driven by text, they do not analyze the content of the statement directly. Lucca and Trebbi (2009) show that Fed rate announcements drive short-run rates and hawkish sentiment in Fed statements drives longer-run rates. The important part of our test is not the finding that text drives residual variance in asset prices, but that *different* components drive different maturities.

[3] A complementary literature has begun to extract policy intentions rather than private information on economic conditions from central bank communication using machine learning approaches (Tobback et al. 2017, for example).

2.1 A simple framework

The Bank of England was granted operational independence for monetary policy in May 1997. A nine-person committee, the Monetary Policy Committee (MPC), was established to set policy on a monthly basis in a way consistent with meeting its inflation target remit. The MPC publishes minutes of its meeting following each decision. The Bank had been publishing a quarterly 'Inflation Report' since 1993, following the adoption of inflation targeting in the UK. From the point of independence that Report became a 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 analysis and inflation projections that the Monetary Policy Committee uses to make its interest rate decisions." MPC members also give regular speeches setting out their own individual view of the outlook and appropriate policy stance.

We focus on the Bank of England's Inflation Report for several reasons. Since no policy decision occurs on its publication date during our sample period, and it does not contain formal forward guidance, it allows us to isolate communication solely about the Bank's views on economic conditions. This allows us to interpret the market reaction we observe on IR days as relating to the outlook component of *Delphic* communication, unlike in other environments where central banks transmit potentially multiple kinds of signals.[4] Second, since the IR is published on a fixed schedule, the timing is not endogenous to evolutions in traders' beliefs nor market conditions more broadly. Third, it contains heterogeneous information in the form of quantitative forecasts and textual narrative, which allows us to explore which type of information drives the different market responses.

Denote the MPC's month m forecast variables as

$$\boldsymbol{\omega}_m = \left(\pi^{CB}_{m;h}, \tilde{y}^{CB}_{m;h}\right)^T \tag{1}$$

where $\pi^{CB}_{m;h} \equiv \mathbb{E}^{CB}_m \left[\, \pi_{m+h} \,\right] - \pi^*$ is the central bank's (in this case MPC's) h-month ahead forecast of inflation made in month m, here expressed as a deviation from the inflation target π^*; and $\tilde{y}^{CB}_{m;h} \equiv \mathbb{E}^{CB}_m [\, \tilde{y}_{m+h} \,]$ is the central bank's forecast of the output gap. The MPC's policy action in month m can be approximated with a standard monetary policy rule as

$$i_m = r^*_m + \pi^* + \boldsymbol{\phi}^T \boldsymbol{\omega}_m + \epsilon_m. \tag{2}$$

Here r^*_m is the MPC's view of the equilibrium real interest rate, which we allow to be time-varying, albeit slow moving, as in Laubach and Williams (2003); the $\boldsymbol{\phi}$ coefficients

[4] Such an exercise would, for example, not be possible in the US since the Fed does not publish contemporaneous Greenbook forecasts.

relate economic conditions to the nominal rate decision; and the residual ϵ_m is the part of the policy rate that cannot be explained by the central bank's typical reaction to the information set and is, therefore, the common definition of a monetary policy shock as in Romer and Romer (2004) or Cloyne and Hürtgen (2016). We take the $\boldsymbol{\phi}$ coefficients as fixed in the discussion below. In a model with time-varying coefficients, one can think of $\boldsymbol{\phi}$ as the average weights placed on forecasts, and the monetary policy shock as capturing deviations in meeting t from these averages.

Under standard asset pricing theory we can write longer-term 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.

The k-month ahead forward rate on day t can therefore be written as

$$f_{k,t} = \mathbb{E}\left[\, i_{m(t)+k} \;\middle|\; Z_t \,\right] + \mathrm{TP}(f_{k,t}). \tag{3}$$

Using (2), we can express that forward rate as

$$f_{k,t} = \mathbb{E}\left[\, r^*_{m(t)+k} \;\middle|\; Z_t \,\right] + \pi^* + \boldsymbol{\phi}^T \mathbb{E}\left[\, \boldsymbol{\omega}_{m(t)+k} \;\middle|\; Z_t \,\right] + \mathbb{E}\left[\, \epsilon_{m(t)+k} \;\middle|\; Z_t \,\right] + \mathrm{TP}(f_{k,t}). \tag{4}$$

The pricing of the forward interest rate therefore requires the market to predict the MPC's future view of the real interest rate $r^*_{m(t)+k}$; future preferences $\epsilon_{m(t)+k}$; and future *h-period* ahead forecasts $\boldsymbol{\omega}_{m(t)+k}$. As explored in Cochrane (2011), there are a number of competing theories for the existence of term premiums but Bansal and Shaliastovich (2013) and Martin (2013), in paritcular, discuss the role of risks and uncertainties around the outlook for real income growth and inflation in driving term premiums. If we observe a systematic reaction in market interest rates on the day of an IR announcement, then the market must receive information about one of these components. We now describe how quantitative and narrative information can deliver such information.

2.2 Effects of quantitative information

The quantitative information in the IR release on day t is the forecast $\boldsymbol{\omega}_{m(t)}$ and the distributions around this. At the time of release, the interest rate decision $i_{m(t)}$ is already known. The most direct impact of the revelation of $\boldsymbol{\omega}_{m(t)}$ on future expectations would be if it were also informative of future forecasts, since (2) makes clear these will affect future interest rates. To make the point more formal, define $\Delta\hat{\omega}_{m(t)+k,t} \equiv \mathbb{E}\big[\,\omega_{m(t)+k} \bigm| Z_t\,\big] - \mathbb{E}\big[\,\omega_{m(t)+k} \bigm| Z_{t-1}\,\big]$. A direct effect on the k-month ahead forward rate is present if $\Delta\hat{\omega}_{m(t),t}$ (the market's updated belief on the current forecast variables induced by the IR publication) is correlated with $\Delta\hat{\omega}_{m(t)+k,t}$. While persistence in macroeconomic conditions might generate this correlation at shorter horizons, it is less likely to generate significant correlations beyond one ($k = 12$) or two ($k = 24$) years, a point we return to in the next section. A similar point can be made with respect to the distributional information around the forecasts. While this in principle can affect future term premiums, whether it is persistence enough to do so is unclear.

Quantitative information can also have important indirect effects. The policy rule (2) presents the market with an identification problem. While $i_{m(t)}$ is known prior to day t, the inputs into the rule that led to the policy rate are unknown. A higher-than-expected rate might be due to the MPC's having higher beliefs on equilibrium real rates, higher inflation or growth forecasts, or having a more hawkish tilt in its policy stance. The publication of the forecasts eases this problem. Holding fixed a value for $r^*_{m(t)}$, knowledge of $\boldsymbol{\omega}_{m(t)}$ allows the market to back out $\epsilon_{m(t)}$ and vice versa. Furthermore, both equilibrium real rates and policy stances can be highly persistent. The real interest rate is often modeled as a unit-root process, which implies that any revisions to current views on its value would propagate one-for-one into future views. Also, for example, if a more-hawkish-than-expected tilt coincides with the appointment of new members to the MPC, the market might infer a more hawkish stance would persistent throughout their terms, which are three or more years.[5]

2.3 Effects of narrative information

The narrative information in the IR can also have important effects on future expectations 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 $\boldsymbol{\omega}_{m(t)}$. 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

[5]In our sample period, minutes containing individual MPC members' votes were published around two weeks after the IR, so any inference made on the policy stance due to forecast variable publication must be made for the MPC as a whole.

transitory price movements.

Second, along similar lines, the narrative can provide context about the variance and skew around the quantitative forecasts that can affect term premiums in forward rates. Moreover, as we discuss below, in practice there is very little change in the variance of forecasts in our sample, so these may not be accurate signals of the uncertainty facing the MPC about future economic conditions. The narrative can therefore guide markets about the risks the MPC considers when making its forecasts.

Third, 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, r_m^* 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. As discussed above, if markets understand that the MPC has revised its view on real rates, this can have effects on forward rates further into the future. The narrative can also communicate uncertainty about the level of the real rate, which would propagate into long-run term premiums.

A final point is that the information content of the quantitative and narrative signals is not in general separable. For example, the narrative can provide a signal on r_m^*, which, when combined with the published $\boldsymbol{\omega}_{m(t)}$, can fully overcome the identification problem that (2) presents markets. In our empirical strategy, we purge asset price movements *and* narrative information of their variation due in the quantitative forecasts in order to isolate the impact of narrative information.

2.4 Yield curve data

To capture the effects across the yield curve, in this paper we use data on four nominal interest rates derived from UK government bond prices: the one-year spot rate; three-year forward rate; five-year forward rate; and five-year ahead, five-year forward rate (equivalent to the average forward rate five to ten years ahead).[6] Our sample period runs from 1998, when the MPC started to publish consistent forecasts in the quarterly IR one week after the policy decision was announced, through to mid-2015, after which the Bank changed its communication approach and moved to publishing the IR alongside the policy announcement, a new policy statement and the minutes. In total, our sample includes 70 IR publications.

While the literature increasingly uses tight, intra-day windows around communication

[6] The rates we use are those published on the Bank of England's website and are zero-coupon rates calculated from end-day market bond prices using a variable roughness penalty, spline-based model.

events,[7] we use daily changes since it may take markets longer to incorporate the length and complexity of the IR. As shown in Figure 1, if market rates on day $t-1$ incorporate all information available prior to day t, denoted Z_{t-1}, then an IR publication on day t can then be thought of as transmitting two kinds of information: quantitative forecast information $\boldsymbol{\omega}_{m(t)}$ and qualitative, narrative information $\boldsymbol{\chi}_{m(t)}$. We denote here the IR information by month since its content relates directly to the current month's policy decision $i_{m(t)}$ made one week before day t.[8] Our identifying assumption is that $Z_t = Z_{t-1} \cup \boldsymbol{\omega}_{m(t)} \cup \boldsymbol{\chi}_{m(t)}$, i.e. all new information relevant for market expectations of future monetary policy on day t is contained in the IR. Given this assumption, we can attribute absolute changes in market rates on IR publication days to the news contained in the IR. Figure 2 shows the time series for the absolute values of the daily yield movements.

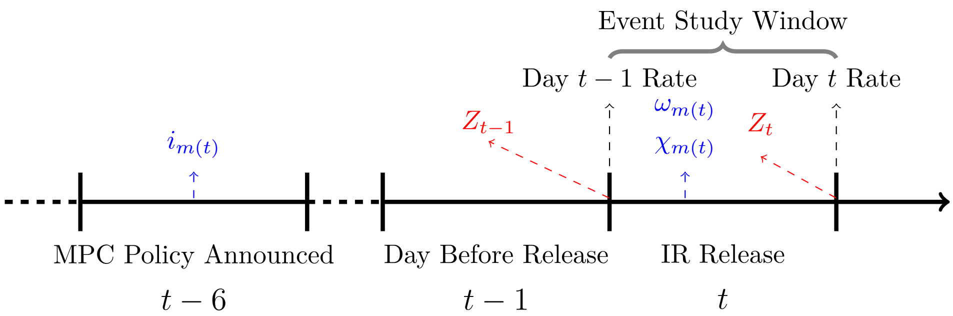

Figure 1: Event Study Time Line for IR Publication on Day t

As above, we can express our four interest rates on a given day t as[9]

1. 1-year spot rate:

$$i_{0:12,t} = \frac{i_{m(t)} + \mathbb{E}\left[\, i_{m(t)+1} \;\middle|\; Z_t \,\right] + \ldots + \mathbb{E}\left[\, i_{m(t)+11} \;\middle|\; Z_t \,\right]}{12} + \mathrm{TP}(i_{0:12,t}) \qquad (5)$$

2. 3-year forward rate:

$$f_{36,t} = \mathbb{E}\left[\, i_{m(t)+36} \;\middle|\; Z_t \,\right] + \mathrm{TP}(f_{36,t}) \qquad (6)$$

[7] Gürkaynak 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.

[8] In the week between the announcement of the MPC decision and the IR release, there is not typically monetary news.

[9] Here 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.

Table 1: Variance Decomposition of Changes in Market Interest Rates by Expectation and Term Premium

	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.

3. 5-year forward rates:

$$f_{60,t} = \mathbb{E}\left[\, i_{m(t)+60} \mid Z_t \,\right] + \mathrm{TP}(f_{60,t}) \tag{7}$$

4. 5-year, 5-year rates:

$$f_{60:120,t} = \frac{\mathbb{E}\left[\, i_{m(t)+60} \mid Z_t \,\right] + \ldots + \mathbb{E}\left[\, i_{m(t)+119} \mid Z_t \,\right]}{60} + \mathrm{TP}(f_{60:120,t}) \tag{8}$$

One common way to decompose changes in yields into expectations and term premiums is to use an affine term structure model. Some of these models use only the past behaviour of the market yield curve to estimate this 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.[10] 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] Specifically 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).

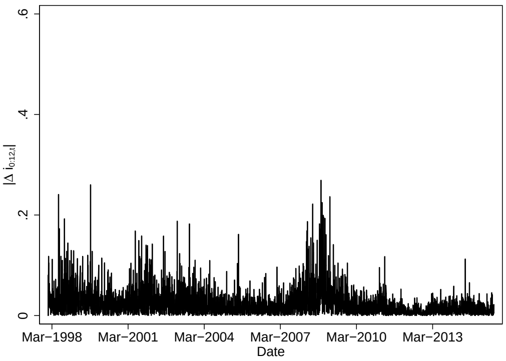

(a) Time Series of $|\Delta i_{0:12,t}|$

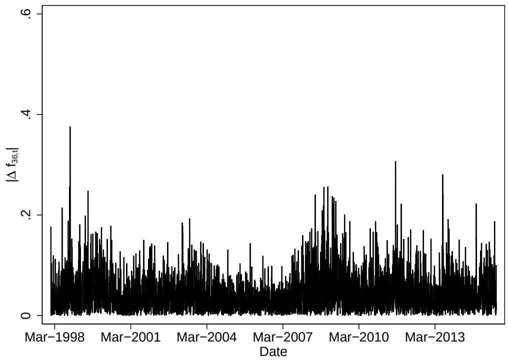

(b) Time Series of $|\Delta f_{36,t}|$

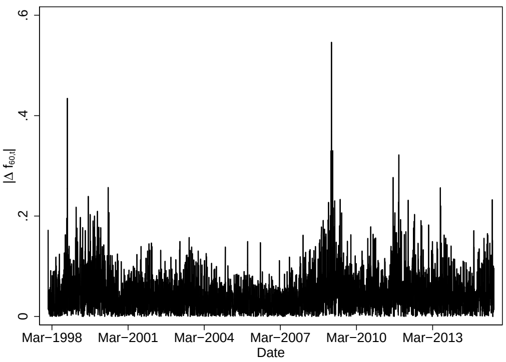

(c) Time Series of $|\Delta f_{60,t}|$

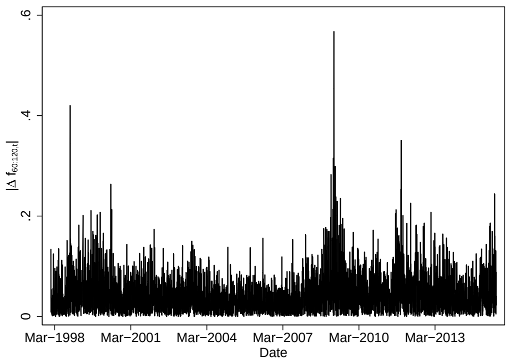

(d) Time Series of $|\Delta f_{60:120,t}|$

Figure 2: Daily Nominal Rate Moves

Notes: This figure plots the time series of the absolute change in the one-year spot rate; the three-year rate; the five-year rate; and the five-year ahead, five-year rate. These are nominal rates derived from daily Gilt prices from 01/01/1998-31/07/2015.

2.5 Impact of the Inflation Report

Before exploring the content and channels through which Inflation Report communication affects long-term interest rates we first seek to establish that the IR does indeed provide significant news in the context of usual market movements and relative to other forms of monetary policy communication. We classify each day in our sample 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. 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 consistently a 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.

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. Figure 4 superimposes the IR kernel densities from figures 3a-3d on the same plot for ease of comparison. If anything, the movements in the shortest rate (i.e. the 1-year spot) are the smallest of any yield: its mode is the highest, meaning there is a greater concentration of small moves, and its kernel denisty also has thinner tails than the others.

Another 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, 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.

3 Measuring the content of the Inflation Report

Having established that the Inflation Report does contain significant news for financial markets, we next seek to establish what the information content of the Inflation Report is and how to measure it.

3.1 Quantitative information

The main quantitative communication in the IR consists of the MPC's fan chart projections for GDP growth and inflation over the following two years (and from August 2004 onwards for the following three years). Those projections are summarized by a mode $g_{m;h}^{CB}$ and $\pi_{m;h}^{CB}$, variance $\text{Var}(g_{m;h}^{CB})$ and $\text{Var}(\pi_{m;h}^{CB})$, and skew $\text{Skew}(g_{m;h}^{CB})$ and $\text{Skew}(\pi_{m;h}^{CB})$ respectively. Since 1998 these forecasts have been consistently conditioned on the path for the policy rate (called 'Bank Rate') implied by market interest rates. Figure 5 provides an illustration of such fan charts. While projections are provided for each quarter over the forecast period, in our anlaysis 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.

One reason why this quantitative information might move long-term interest rates is if it is very persistent. To guage the extent to which this is the case, Table 3 shows the

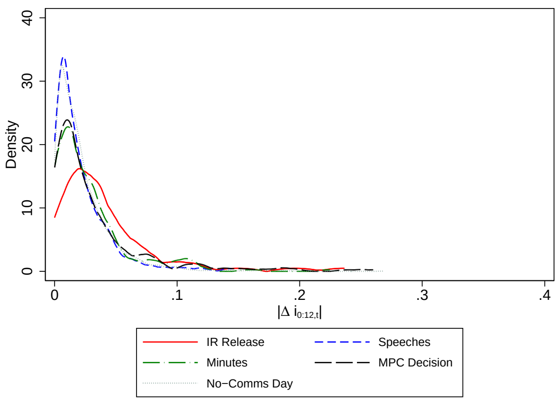

(a) Kernel Density of $|\Delta i_{0:12,t}|$.

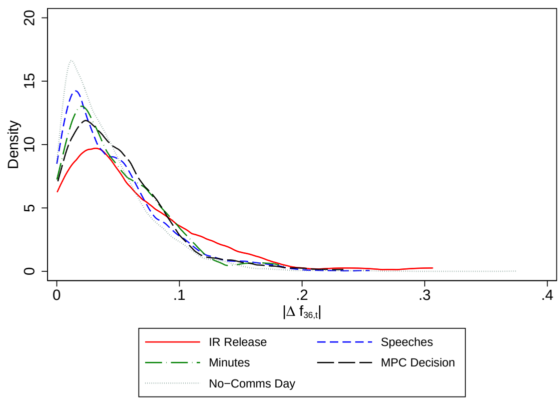

(b) Kernel Density of $|\Delta f_{36,t}|$.

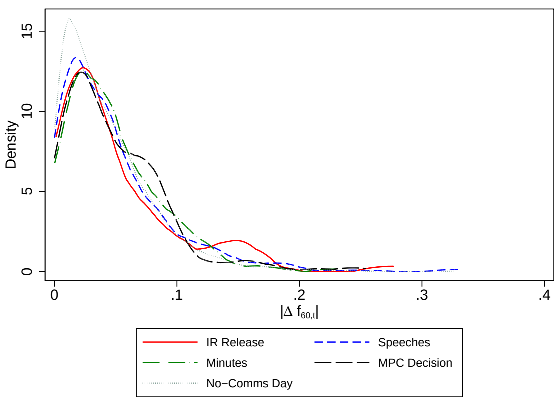

(c) Kernel Density of $|\Delta f_{60,t}|$.

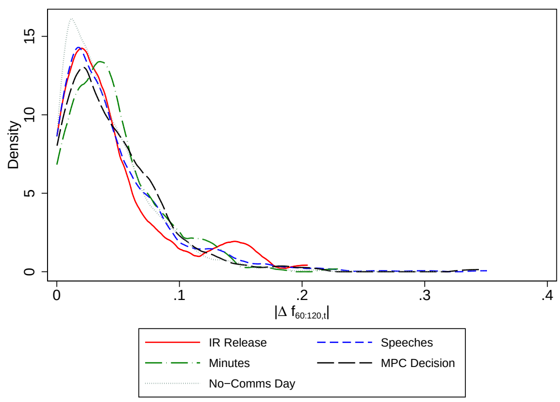

(d) Kernel Density of $|\Delta f_{60:120,t}|$.

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.

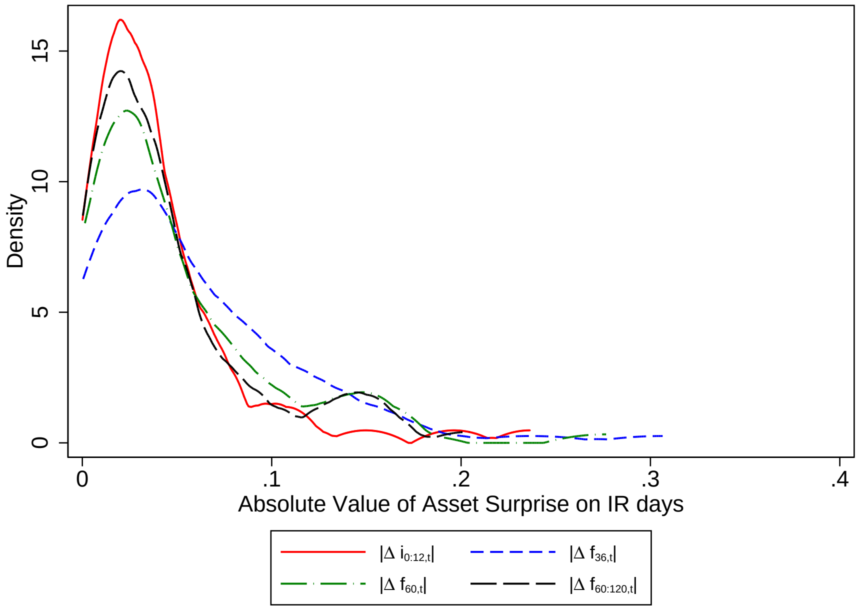

Figure 4: Comparison of Yield Movement Magnitudes on IR Publication Days

Notes: This figure shows the four kernel distributions for yield curve movements on IR publication days from figures 3a-3d on the same plot.

Table 2: Magnitude of Rate Moves on IR Days

Asset	$\geq$ 5 bps	$\geq$ 10 bps	$\geq$ 20 bps
$\lvert\Delta i_{0:12,t}\rvert$	0.24	0.09	0.01
$\lvert\Delta f_{36,t}\rvert$	0.46	0.20	0.03
$\lvert\Delta f_{60,t}\rvert$	0.36	0.13	0.01
$\lvert\Delta f_{60:120,t}\rvert$	0.27	0.11	0.01

Notes: This table tabulates the frequency of yield movements on IR publication days that exceed five, ten, and twenty basis points.

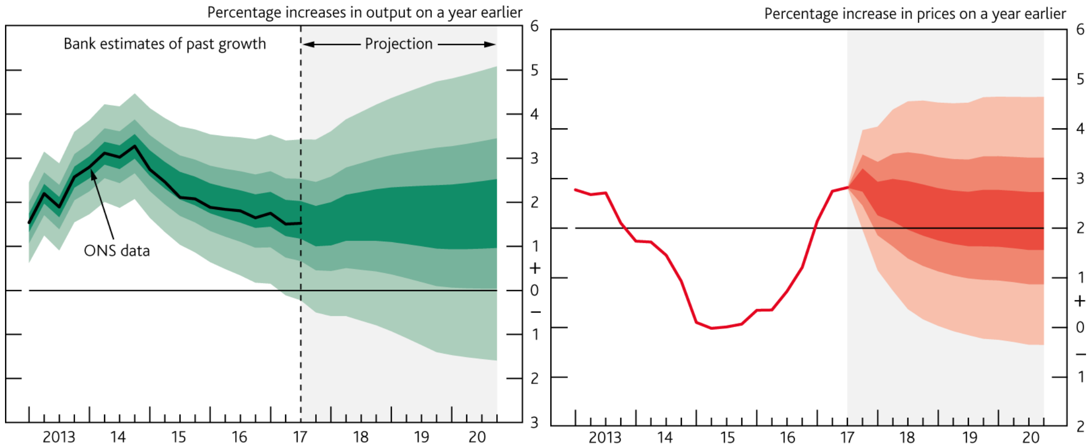

Figure 5: Quantitative Information: Fan Charts

Notes: The left-hand figure shows the November 2017 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 analagous fan chart associated with the CPI inflation forecast.

coefficients from estimating a simple AR(1) process on the forecast variables over our sample period as well as the implied change in future forecasts (one-, three- and five-years) following a unit shock to each variable. We also include a dummy variable for the period after the collapse of Lehman Brothers in September 2009. The MPC increased the variance of the forecast distributions for CPI and GDP growth at the onset of the crisis and has largely maintained these higher variances since. None of the forecast variables are highly persistent, meaning this effect cannot be the driver of the reaction of long-maturity yields.

Table 3: Persistence of Key Forecast Variables

Main Regressors	(1) $\pi^{CB}_{m;h}$	(2) $g^{CB}_{m;h}$	(3) Var($\pi^{CB}_{m;h}$)	(4) Var($g^{CB}_{m;h}$)	(5) Skew($\pi^{CB}_{m;h}$)	(6) Skew($g^{CB}_{m;h}$)
AR1 Coefficient	0.57***	0.58***	0.81***	0.75***	0.51***	0.63***
D(Post-Lehman)	-0.16**	0.092	0.33**	0.43**	0.017	-0.092
Constant	-0.0030	1.10***	0.11**	0.32**	0.016	-0.039
R-squared	0.544	0.374	0.976	0.870	0.285	0.511
After 1 year	0.10	0.11	0.42	0.32	0.07	0.15
After 3 years	0.00	0.00	0.08	0.03	0.00	0.00
After 5 years	0.00	0.00	0.01	0.00	0.00	0.00

Notes: D(Post-Lehman) is a dummy variable which takes on the value 1 after the collapse of Lehman Brothers in September 2009.

As explained in Section 2, rather than the rate of GDP growth $g^{CB}_{m;h}$ the potentially more relevant variable for interest rate expectations is the MPC's view of the output gap ($\tilde{y}^{CB}_{m;h}$). It may be that investors infer the MPC's view of the future output gap from its GDP growth forecast. To proxy this we construct an implied modal output gap using the MPC's growth forecasts together with private-sector estimates of long-run potential growth.[11]

Since we are looking at the absolute value of market news, it is also important to control for the news in the quantitative information. Ideally we would compare each of the MPC's forecast measures to equivalent expectations from the private sector. We include this for the modal forecast variables where we have comparable private sector forecast data on CPI inflation and GDP growth; we call these surprise variables $|\pi^{CB}_{m;h} - \pi^{P}_{m;h}| \equiv |\mathbb{E}^{CB}_{m}\left[\,\pi_{m+h}\,\right] - \mathbb{E}^{P}_{m}\left[\,\pi_{m+h}\,\right]|$ and $|g^{CB}_{m;h}) - (g^{P}_{m;h})| \equiv |\mathbb{E}^{CB}_{m}\left[\,g_{m+h}\,\right] - \mathbb{E}^{P}_{m}\left[\,g_{m+h}\,\right]|$ where $\mathbb{E}^{P}_{m}$ indicates a private sector forecast.

For the variance and skew variables, where equivalent private-sector measures are not readily available over our sample, we use the change from the previous IR release denoted as $\Delta\mathrm{Var}(X) \equiv \mathrm{Var}\,(X_m) - \mathrm{Var}\,(X_{m-1})$. For the output gap, 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 — $|\tilde{y}^{CB}_{m,h} - \tilde{y}^{CB}_{m-1,h}|$). 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 — $|\tilde{y}^{CB}_{m,h-1} - \tilde{y}^{CB}_{m-1,h}|$.

In total this produces fifteen variables associated with the quantitative information in the IR communication.

3.2 Narrative information

In addition to quantitative information, the IR also contains extensive narrative data in the form 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. As explained in Section 2, however, this does not mean that it does not potentially contain information relevant for long-run rate expectations.

[11]Specifically, we take the implied real GDP series from the the forecast and grow it 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, $\tilde{y}^{CB}_{m;h}$, is the percentage deviation of the forecast level of real GDP from the BK-filtered trend series.

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. (2017). 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. (2017) and introduced by Griffiths and Steyvers (2004); we refer interested readers to those papers for full details.[12]

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 *topics* (i.e. factors), each of which is a probability vector $\boldsymbol{\beta}_k \in \Delta^{V-1}$ 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 $\boldsymbol{\theta}_d$ (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 vth 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 $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ 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

[12] A 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.

interpretable set of topics.[13]

Figure 6 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. (2017) to obtain the distribution over topics in the IR published on day t,e given the estimated LDA model, denoted by $\boldsymbol{\theta}_t$. Since changes in topic coverage can also have potentially important market effects, we also include $\boldsymbol{\delta}_t \equiv |\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1}|$ to obtain a 60-dimensional representation of the text information in each IR —the 30 topic levels and the 30 absolute changes.

To illustrate this, Figure 7 plots the share over time of the content the IRs in our sample devote to two topics, for which we provide an alternative representation 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.

4 News in the Inflation Report content

Using the measures described above, in this section we assess the market impact of the quantitative and narrative information in the Inflation Report. The approach we take is to first control for the impact of the quantitative forecast publication, and then to test whether the narrative contains additional information beyond this. We thus determine whether the narrative moves market rates independently of the forecasts.

[13] There 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.

													Pro-cyclicality
Topic0	growth	wage	cost	earn	pay	pressur	labour	sector	real	privat	higher	rise	0.013
Topic1	measur	chang	use	differ	can	weight	indic	rel	provid	movement	base	tend	0.013
Topic2	rate	interest	market	offici	expect	impli	point	chart	rise	path	unchang	forward	0.013
Topic3	growth	quarter	gdp	slow	output	weaker	slowdown	strong	stronger	trend	grew	pick	0.01
Topic4	survey	sector	output	servic	suggest	manufactur	indic	report	evid	busi	agent	region	0.01
Topic5	invest	capit	busi	money	financi	deposit	compani	sector	corpor	broad	increas	privat	0.009
Topic6	price	oil	energi	rise	increas	cost	commod	higher	import	food	futur	ga	0.009
Topic7	month	year	rose	annual	fell	point	percentag	earlier	chart	around	increas	januari	0.007
Topic8	import	uk	export	trade	good	net	demand	domest	servic	sterl	account	world	0.006
Topic9	report	sterl	februari	exchang	august	rate	novemb	start	lower	time	point	averag	0.006
Topic10	estim	data	revis	gdp	suggest	account	recent	output	previous	past	inform	current	0.005
Topic11	household	spend	hous	consumpt	incom	consum	real	increas	save	rise	wealth	ratio	0.005
Topic12	may	recent	retail	level	year	sale	stock	reflect	time	underli	trend	inventori	0.004
Topic13	remain	year	past	recent	level	although	low	continu	littl	broadli	appear	risen	0.003
Topic14	project	period	forecast	central	mpc	assum	assumpt	path	around	next	report	profil	0.002
Topic15	employ	labour	unemploy	work	number	hour	market	job	peopl	increas	time	rise	0.0
Topic16	may	could	reduc	might	exampl	increas	possibl	adjust	lead	greater	addit	becom	-0.001
Topic17	chart	year	averag	tabl	probabl	show	around	forecast	distribut	rang	outturn	time	-0.001
Topic18	risk	committe	uncertainti	outlook	balanc	judg	view	central	downsid	prospect	overal	consider	-0.004
Topic19	effect	will	govern	impact	increas	public	plan	fiscal	year	nomin	chang	budget	-0.004
Topic20	inflat	target	cpi	pressur	price	rpix	higher	vat	inflationari	rise	contribut	downward	-0.004
Topic21	expect	term	near	littl	month	medium	short	outlook	rise	next	longer	fall	-0.006
Topic22	will	section	factor	depend	discuss	develop	extent	suppli	box	affect	impact	demand	-0.007
Topic23	economi	unit	euro	area	global	world	kingdom	activ	state	countri	prospect	econom	-0.008
Topic24	market	equiti	financi	asset	bond	yield	corpor	share	may	increas	return	reflect	-0.008
Topic25	fall	reflect	declin	chart	part	fallen	sharpli	larg	past	lower	fell	much	-0.008
Topic26	demand	continu	growth	recoveri	remain	domest	weak	activ	support	pace	gradual	eas	-0.011
Topic27	bank	polici	committe	purchas	mpc	meet	asset	monetari	rate	maintain	billion	econom	-0.015
Topic28	product	compani	capac	margin	suppli	recess	spare	demand	weak	economi	output	slack	-0.018
Topic29	bank	credit	condit	lend	fund	borrow	mortgag	loan	avail	compani	secur	cost	-0.02

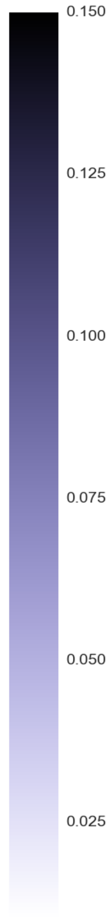

Figure 6: Topics Ranked by Pro-Cyclicality; Terms within Topics Ranked by Probability

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.

(a) Topic 15 'Labor Market' **(b)** Topic 26 'Demand'

Figure 7: 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 proprtional to its probability.

4.1 Impact of quantitative forecast information

The first step in our analysis is to estimate the regression model

$$|Y_t| = \beta_{0Y} + \boldsymbol{\beta}_{1Y}^T \mathbf{q}_t + \beta_{2Y}\text{VIX}_t + \varepsilon_t^Y, \tag{9}$$

where $\mathbf{q}_t$ is the set of 15 quantitative variables derived from the published forecasts described in Section 3.1, Y_t refers to a yield (1-year spot, 3-year forward, etc.), and t refers to an IR publication date. We also include the VIX uncertainty 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 4 displays the results. We are not particularly interested in the magnitudes nor significance of any individual quantitative measure,[14] but rather their joint explanatory

[14] In particular, some of the derived measures are highly correlated.

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 4 for a separate specification in which we do not include VIX but just the IR controls $\mathbf{q}_t$. In either case, the key result is the same. The forecast 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 results on persistence in Table 3 already suggested that forecast publication alone was unlikely to have a direct impact on market expectations of future interest rates, but as explained in Section refsec:framework it could have long-run indirect effects by transmitting information about, for example, the real equilibrium interest rates or policy intentions. In the case of the Bank of England, these type of effects appear to have sharply declining effects on long-term interest rates.

4.2 Testing for additional news in the narrative

Given the declining importance of quantitative information in explaining the variance in market interest rates across the horizon of the rate, we next attempt to identify to what extent that residual variance can be explained by the narrative information in the IR.

One issue in assessing the information content of the narrative is that it is likely in part to be endogeneous 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. To address this we regress the time series across IRs for each topic $\boldsymbol{\theta}_t$ on the set of numerical forecast information above $\mathbf{q}_t$ and VIX_t. 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. Instead we construct 'narrative shocks' by extracting the exogenous component of the Inflation Report text.

Our challenge is then to determine whether these text shocks can explain the variance in the interest rate residuals from equation 9. One clearly cannot regress $\widehat{\varepsilon}_t^Y$ on the full set of text shock variables using ordinary least squares since there are are almost as many variables as there are observations. Instead, we adopt an elastic net regression approach as in Zou and Hastie (2005) and solve the problem

$$\min_{\gamma_Y} \sum_t \left(\widehat{\varepsilon}_t^Y - \boldsymbol{\gamma}_Y^T \widehat{\boldsymbol{\nu}}_t\right)^2 + \lambda \left[\alpha \sum_v |\gamma_{vY}| + (1-\alpha) \sum_v \gamma_{vY}^2\right]. \tag{10}$$

The first term is simply the objective function of an OLS regression of the yield Y

Table 4: Effect of Forecast Variable on Market Yields

Main Regressors	(1) $\lvert\Delta i_{0:12;t}\rvert$	(2) $\lvert\Delta f_{36;t}\rvert$	(3) $\lvert\Delta f_{60;t}\rvert$	(4) $\lvert\Delta f_{60:120;t}\rvert$
$\pi^{CB}_{m;h}$	-0.016	-0.014	0.016	0.025
	[0.535]	[0.698]	[0.645]	[0.432]
$\lvert\pi^{CB}_{m;h} - \pi^{P}_{m;h}\rvert$	0.034	-0.035	-0.042	-0.051
	[0.292]	[0.442]	[0.314]	[0.152]
$\text{Var}(\pi^{CB}_{m;h})$	-0.011	0.025	0.021	0.020
	[0.324]	[0.161]	[0.257]	[0.218]
$\Delta\text{Var}(\pi^{CB}_{m;h})$	0.038	-0.057	-0.12**	-0.085*
	[0.455]	[0.240]	[0.011]	[0.084]
$\text{Skew}(\pi^{CB}_{m;h})$	0.0045	0.0100	-0.0040	-0.0036
	[0.891]	[0.834]	[0.924]	[0.923]
$\Delta\text{Skew}(\pi^{CB}_{m;h})$	0.025	0.065	0.068	0.052
	[0.500]	[0.150]	[0.108]	[0.154]
$g^{CB}_{m;h}$	0.031**	0.00010	-0.012	-0.015
	[0.038]	[0.995]	[0.476]	[0.381]
$\lvert g^{CB}_{m;h} - g^{P}_{m;h}\rvert$	-0.048*	0.050	0.053*	0.024
	[0.063]	[0.140]	[0.084]	[0.373]
$\text{Var}(g^{CB}_{m;h})$	0.0010	-0.0052	-0.0021	-0.012
	[0.934]	[0.765]	[0.903]	[0.446]
$\Delta\text{Var}(g^{CB}_{m;h})$	0.0098	0.019	0.024	0.034*
	[0.539]	[0.236]	[0.153]	[0.058]
$\text{Skew}(g^{CB}_{m;h})$	-0.014	-0.071	-0.060	-0.045
	[0.688]	[0.171]	[0.202]	[0.260]
$\Delta\text{Skew}(g^{CB}_{m;h})$	-0.063**	-0.00055	0.0031	-0.022
	[0.049]	[0.990]	[0.938]	[0.533]
$\tilde{y}^{CB}_{m;h}$	0.41	3.35*	1.68	-0.50
	[0.722]	[0.097]	[0.369]	[0.747]
$\lvert\tilde{y}^{CB}_{m;h} - \tilde{y}^{CB}_{m-1;h}\rvert$	4.96**	1.41	-1.00	-1.43
	[0.015]	[0.685]	[0.767]	[0.595]
$\lvert\tilde{y}^{CB}_{m;h-1} - \tilde{y}^{CB}_{m-1;h}\rvert$	1.64	-0.35	-0.25	1.31
	[0.430]	[0.925]	[0.944]	[0.652]
VIX_t	0.0013**	0.0023	0.0018	0.0017
	[0.023]	[0.151]	[0.229]	[0.150]
Constant	-0.070*	-0.026	0.014	0.043
	[0.075]	[0.693]	[0.833]	[0.456]
R-squared	0.563	0.368	0.280	0.274
R-squared No Vix	0.526	0.303	0.229	0.215

Notes: This table reports estimates from regressing absolute changes in market yields on the quantitative forecast variables defined in Section 3.1.

residuals on the text shock variables. The second term is a penalty on non-zero values of the regression coefficients $\boldsymbol{\gamma}_Y$. The parameter α can range from 0, equivalent to a ridge regression, to 1 for 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 potential covariates are 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 a LASSO-like model but that is robust to not excluding any very high correlations in the text shock variables (Friedman et al. 2010).

Before estimating (10), one must also choose a value for the penalty parameter λ. A common approach is to use cross validation to select λ based on out-of-sample predictive performance. Given our small sample size, leave-one-out cross validation (LOOCV) 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 69 data points:
      i. Remove the point from the sample.
      ii. Fit (10) on the remaining 68 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 69 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, LASSO selects fewer covariates. This increases our confidence that any selected text shock variables have a robust relationship with market interest rates. Table 5 reports the value of the penalty weight chosen by LOOCV (λ_{CV}) and the number of text shock variables selected under this model.

Table 5: Number of Selected Narrative Shocks from LOOCV

	$\lvert\Delta i_{0:12,t}\rvert$	$\lvert\Delta f_{36,t}\rvert$	$\lvert\Delta f_{60,t}\rvert$	$\lvert\Delta f_{60:120,t}\rvert$
λ_{CV}	0.0006	0.001	0.002	0.001
# selected text variables	60	57	54	57

Notes: This table summarizes the estimation of (10) for four market rates by leave-one-out cross validation. At each maturity, a large number of text shock variables are chosen to accurately predict held-out yield residuals.

That a large number of variables are selected for each maturity suggests that they are indeed important in explaining the yield residuals. To test that this relationship is significant we need to compare to the distribution of the number of selected variables would be under the null hypothesis of their containing no information for market interest rates. To do that we use a simple permutation test. In each of 500 simulations for each interest rate, we randomly permute the residuals, re-estimate (10) by LOOCV, and observe the number of selected text shock variables. We then compare the values in Table 5 against these simulated distributions. The results are shown in Figure 8, where the red-shaded histogram is the distribution of the number of selected text variables generated by the simulations, and the blue dashed line is the number of text variables we select in our actual data. At all maturities we strongly reject the null hypothesis that the correlations we find are spurious. In the vast majority of permuted draws LOOCV selects no text shock variables at all, and in no draw is the number of selected variables is greater than or equal to the number we select with the original order. We conclude that indeed there is genuine explanatory power contained in the IR that is orthogonal to the numerical forecast variables.

One hypothesis might be that the IR narrative contains information in our sample only due to the period of quantitative easing that began in 2009. The communication of such an unconventional monetary policy may have been more reliant on or accompanied by a greater degree of narrative than for more conventional changes in the policy rate. We have repeated the exact analysis described in this section for the subsample of IRs ending in November 2008 to exclude the period of quantitiative easing, and continue to find strong supportive evidence for an information effect of the IR narrative.[15]

4.3 The content of the narrative news

The challenge is then to determine which topics are associated with yield movements. Although LDA reduces the dimensionality of text substantially, we have sixty topic variables in total and so OLS is still not feasible—when combined with the quantitative variables, there are more covariates than observations. Moreover, there is no guarantee that the set of selected topics under the elastic net procedure corresponds to the true set of predictors. Meinshausen and Bühlmann (2006) have shown that the number of selected features from an elastic net regression estimated via cross validation may be a superset of the relevant variables. Instead, we adopt a bootstrap procedure suggested by Hastie et al. (2015). First, we estimate the elastic net regression. Then, in each of 500 simulations we draw a bootstrap sample with replacement from our original data, compute coefficient estimates using the LOOCV procedure described in the previous section, and record whether each topic variable is selected. Across all the bootstrap draws, we

[15]Results available on request.

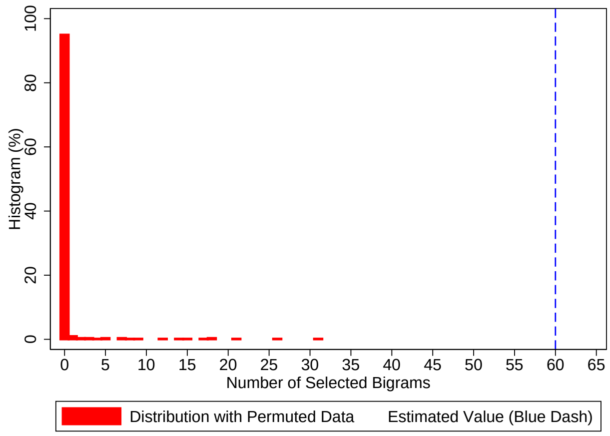

(a) 1-year Spot

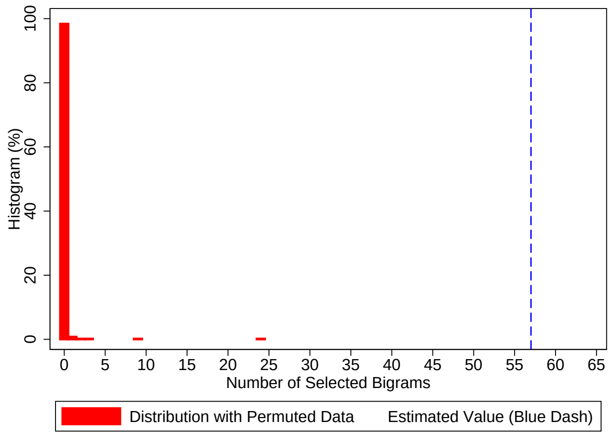

(b) 3-year Forward

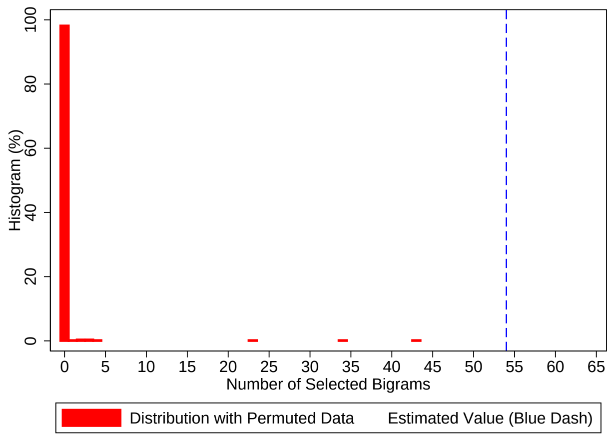

(c) 5-year Forward

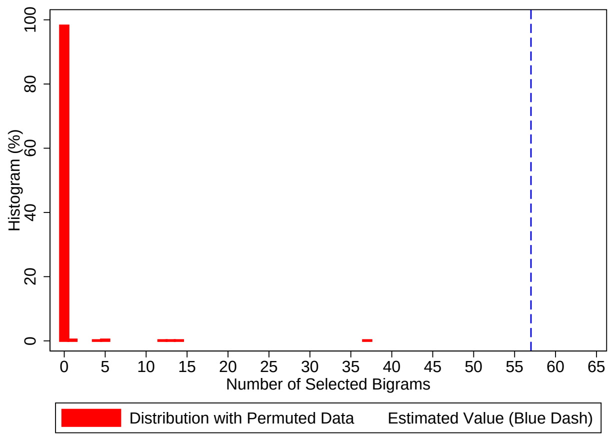

(d) 5-year, 5-year Forward

Figure 8: 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.

can compute the fraction of times that each topic variable is selected, and use this as an indicator of which are the key variables driving the market response to the IR.[16]

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.

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

$\lvert\Delta i_{0:12,t}\rvert$		$\lvert\Delta f_{36,t}\rvert$		$\lvert\Delta f_{60,t}\rvert$		$\lvert\Delta f_{60:120,t}\rvert$	
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 (i.e. an element of $\boldsymbol{\theta}_t^R$), while a D indicates a residuals in the absolute change in the topic level (i.e. an element of $\boldsymbol{\delta}_t^R$).

Figure 9a shows a histogram of all the unique bilateral correlations between the 60 text shocks – levels (L) and change (D). It shows that the majority are zero and only 377 out of a possible 1770 are statistically significant. Figure 9b shows that the correlations amongst our 'top' topics is even more centred on zero and has only 7 of 55 bilateral correlations that are statistically-significant at 5% level.

4.4 The relative importance of narrative news across maturities

We know from Table 4 that the effect of the forecast variables declines as the maturity of the interest rate increases. Here we examine the extent to which the narrative information can explain the variation in interest rates across maturity and relative to the forecast variables.

To begin, Figure 10 plots the evolution of the R^2 as we add topic variables to the baseline regressions presented in Table 4 in the order of the frequency with which they are selected in the bootstrap draws. The vertical intercepts of the curves in the figure correspond to those in the baseline. As we add more topic variables, the R^2 is of course guaranteed to increase, but the rates of increase across yields are informative. Starting from a much lower baseline R^2, topic variables initially have a high marginal impact on

[16]The elastic net regression 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.

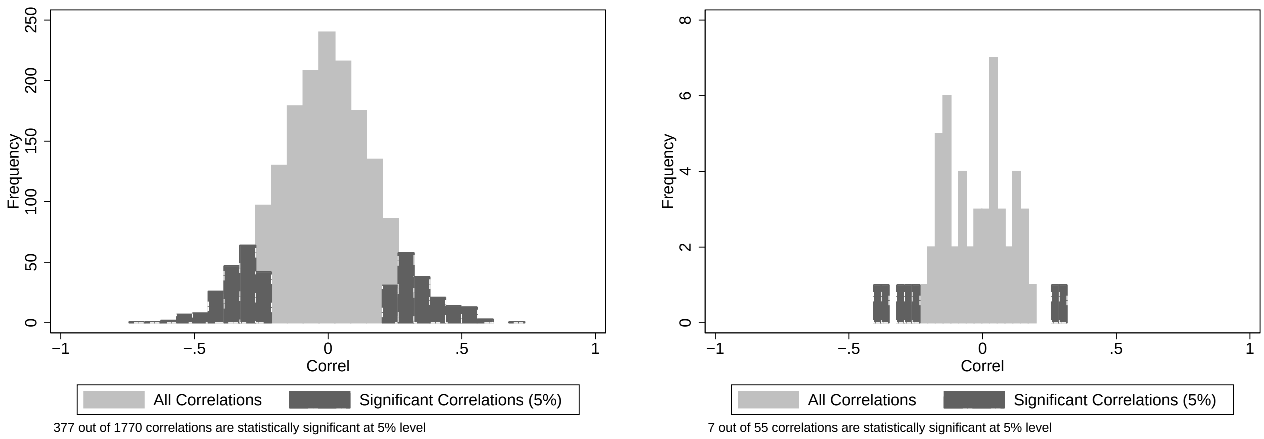

(a) All text shocks

(b) Top text shocks

Figure 9: Correlation between text shocks

Notes: This figure .

goodness-of-fit for long rates before tapering off after the inclusion of around eight. On the other hand, the marginal impact of topics for short rates is much steadier. These findings are also in line with the selection percentages reported in Table 6. The dropoff in frequency for topic variables for long rates is much sharper than for the 1-year spot rate. Altogether, we find that a potentially broader range of narrative signals is important for driving short-run market expectations but that each has a moderate impact on R^2, while a narrower range of signals matter for long-run expectations but these have a large impact.

The baseline R^2 statistics displayed in Figure 10 are the result of including both quantitative forecast controls and VIX in asset price regressions, which makes determining the relative impact of forecast variables versus topic variables hard to directly determine. Instead, in Table 7 we begin in column (1) by only regressing the change in 1-year spot rates on VIX. In column (2) we add the forecast controls and compute the partial R^2 (the fraction of the remaining variance that the forecasts explain after controlling for VIX) as 0.39. In column (3) we add the top four topic variables corresponding to those in Table 6, and again compute the partial R^2 (the fraction of remaining variance that topics explain after controlling for VIX and forecasts) as 0.29. In this sense, we find that quantitative forecasts have a larger impact on short-run rates than narrative. Columns (4)-(6) of Table 7 repeat the exercise for three-year forward rates, and Table 8 for the longer forward rates. Along the yield curve, we find an increasing relative importance of narrative. For the longest rate (five year, five year forward), only four topic variables explain substantially more movement than fifteen forecast variables.

Table 7: Contribution of Quantitative vs. Narrative Variables to R^2

Main Regressors	(1) $\lvert\Delta i_{0:12;t}\rvert$	(2) $\lvert\Delta i_{0:12;t}\rvert$	(3) $\lvert\Delta i_{0:12;t}\rvert$	(4) $\lvert\Delta f_{36;t}\rvert$	(5) $\lvert\Delta f_{36;t}\rvert$	(6) $\lvert\Delta f_{36;t}\rvert$
VIX_t	0.0024***	0.0013**	0.00050	0.0023**	0.0023	0.0037**
	[0.00]	[0.02]	[0.34]	[0.03]	[0.15]	[0.02]
T29			1.80***			
			[0.01]			
d7			-0.66			-1.39
			[0.13]			[0.16]
T18			-1.32**			
			[0.02]			
T8			-1.34**			
			[0.03]			
d29						-3.18
						[0.10]
T9						1.74**
						[0.03]
d16						1.17
						[0.21]
Constant	-0.011	-0.070*	0.014	0.010	-0.026	-0.065
	[0.41]	[0.08]	[0.81]	[0.61]	[0.69]	[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

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 quantitative in (2) and then, in (3), quantitative and narrative information captured by the top 4 topics.

Table 8: Contribution of Quantitative vs. Narrative Variables to R^2

Main Regressors	(1) $\lvert\Delta f_{60;t}\rvert$	(2) $\lvert\Delta f_{60;t}\rvert$	(3) $\lvert\Delta f_{60;t}\rvert$	(4) $\lvert\Delta f_{60:120;t}\rvert$	(5) $\lvert\Delta f_{60:120;t}\rvert$	(6) $\lvert\Delta f_{60:120;t}\rvert$
VIX_t	0.0014	0.0018	0.0035***	0.0011	0.0017	0.0030***
	[0.20]	[0.23]	[0.01]	[0.19]	[0.15]	[0.00]
T9			1.48**			
			[0.02]			
d4			1.41*			1.77**
			[0.08]			[0.01]
d20			-2.39			-3.34*
			[0.23]			[0.06]
T6			1.39*			2.07***
			[0.08]			[0.00]
d12						-3.26***
						[0.01]
Constant	0.020	0.014	-0.053	0.020	0.043	-0.0070
	[0.33]	[0.83]	[0.44]	[0.26]	[0.46]	[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

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 quantitative in (2) and then, in (3), quantitative and narrative information captured by the top 4 topics.

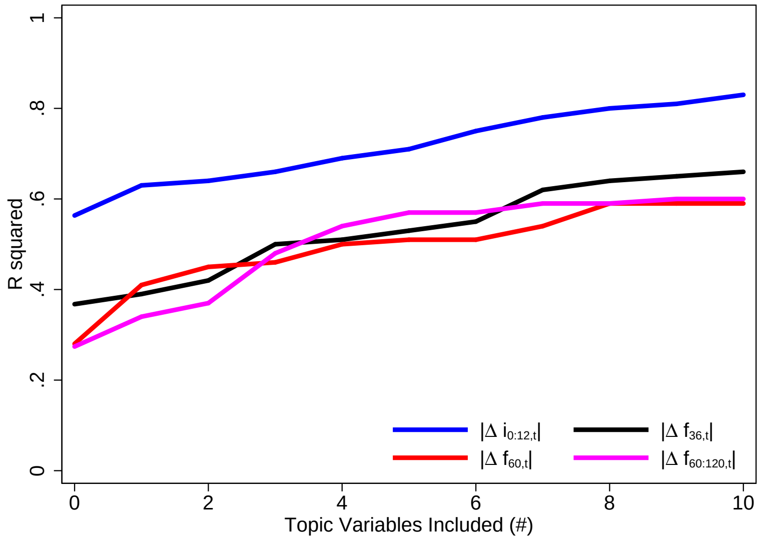

Figure 10: Change in R^2 with Topic Variables

Notes: This figure plots the evolution of the R^2 statistic as we move from including no topic variables to up to ten topic variables. Topic variables are included starting with the topic that has the highest number of times it is included in the boostrap draws. There is a separate R^2 evolution for each of the four maturities of yield considered.

5 What Drives The Long-Run Information Effect?

While the results so far strongly suggest that narrative helps shape market expectations independently of quantitative forecasts, this interesting finding does not address the more fundamental question concerning the channel of a long-maturity information effect. The main goal of this section is to distinguish these alternatives.

5.1 Distinguishing Information Channels

The multidimensionality of text is key for our test of long-run information. We distinguish four channels:

HS Channel In Hanson and Stein (2015), the long-maturity effect comes via short-run monetary news that affects short-term expectations and these expectations are then propagated to longer rates via yield-oriented investors. The effects arise largely because of changes in longer-rate term premiums.

NS1 Channel In Nakamura and Steinsson (2017), the central bank provides expectational information on a persistent variable such as r^*.this leads market participants to update their view of the persistent variable into the future and the effect is mostly on the expectational component of the asset price.

NS2 Channel Of course, while not the channel emphasised in Nakamura and Steinsson (2017), the central bank could send short-maturity- and long-maturity-specific information that separately provide expectational information. These signals would lead market participants to update their views of each end of the yield curve independently, although as in NS1, the effect would come through the expectational component of the asset price.

TP Channel An alternative explanation, which builds off the analysis in Martin (2013), Bansal and Shaliastovich (2013) and Ellison and Tischbirek (2018), relies on the central bank sending signals that cause changes in the longer term uncertainty surrounding the evolution of key persistent variables. Such signal would *only* affect longer-maturity rates but would come through changes in the term premiums.

Through our analysis of the key content, we will now distinguish between these alternative channels. That is, knowing that the relative importance of the narrative grows with asset maturity, we now wish to examine *what* the information drives this narrative information effects.

5.2 The heterogeneity of information across different maturities

Under the first two explanations above, HS and NS1 channels, short and long rates must be moved by the same information. Therefore, if we can show that the same topics drive both short and long rates, then one can plausibly argue there is no long-run information transmission: the topics can move rates at the short run due to information, but then correlate with long-run rates due to the mechanical reaction of these to short-run rate movements. If, however, different topics drive short and long rates, such an explanation is no longer adequate since under pure market mechanics, any topic that affects long rates should also affect short rates.

The content of the key topics for the shortest- and longest-maturity assets we examine, listed above in table 6, is presented graphically as word clouds in Figures 11 and 12. The key topics for 1-year spot rates appear to relate to conjunctural economic news, and include those that vary most with the interest rate cycle. On the other hand, the key topics for the five-year ahead, five-year forward rate appear to relate to the forecasts and the risks around them and are less cyclical. This casual inspect suggests an interesting distinction in the information that drives short and long rates.

Table 9 formalizes the idea that different dimensions drive different rate maturities. It reports the Pearson correlation coefficients between topic variables at different yields 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

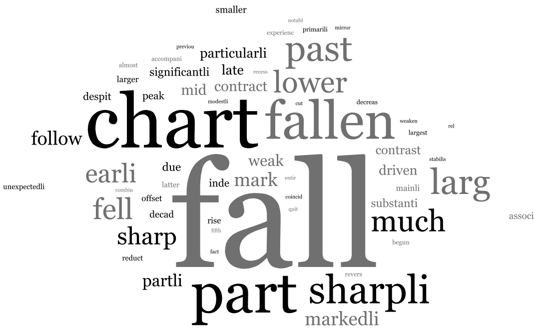

(a) L25

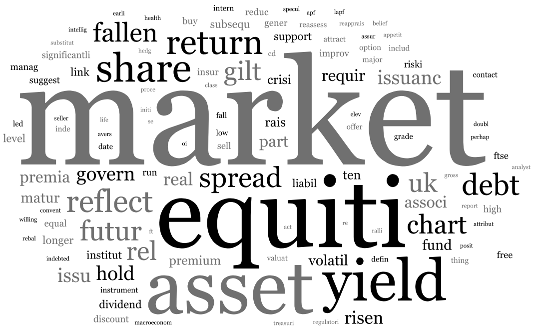

(b) D24

(c) L5

(d) L26

Figure 11: Key Topics for Market Reaction to Narrative: 1-Year Spot Rate

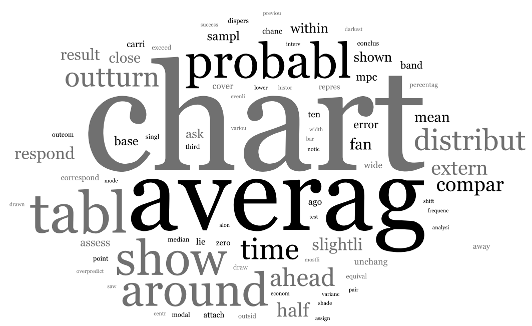

(a) D17

(b) D18

(c) L20

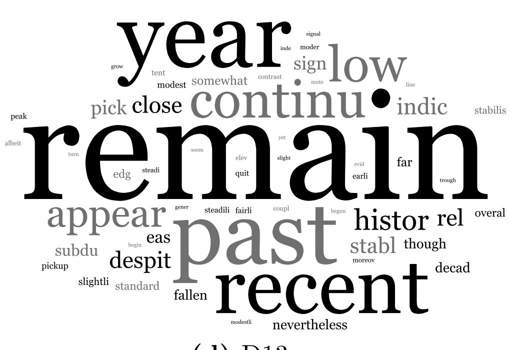

(d) D13

Figure 12: Key Topics for Market Reaction to Narrative: 5-Year, 5-Year Forward Rate

the selection percentages associated with all other rates are significantly correlated. This provides strong evidence that the narrative produces a direct long-run information effect and begins to cast doubt over the HS and NS channels.

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

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

To further evidence this point, we compute the fraction of time the different sections of the IR, those related to economic analysis and those related to the forecast, spend discussing the top topics associated with each rate, which table 10 reports. There is a clear pattern; as one moves out into the yield curve, the top topics are increasingly discussed in the forecast section and discussion in the economics section decreases. This provides yet more evidence that different narrative signals drive short and long rate movements.

Table 10: Top Topic Coverage by Inflation Report Section

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

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

As a final exercise, we conduct a Placebo regression analysis. This involves rotating the top topics from each of the asset classes as explanatory variables for each maturity of asset price news. Table 11 shows the Partial R^2 from this placebo test. For example, the first row shows the Partial R^2 explained by the use of each of the sets of top topics on the 1-year spot rate asset news; the top topics from the asset itself add most information while the top topics from the 5y5y forward rates are essentially uninformative. A mirrored result is found in the final row for the effect of short-maturity topics explaining long-maturity asset news.

In summary, we have found that different aspects of the IR narrative drive short- and long-run rates, which points towards there being a long-run information effect of

Table 11: Matrix of Partial R^2 from top narrative signals across assets

Asset News	Topics			
	Spot 1y	Fwd 3y	Fwd 5y	Fwd 5y5y
Spot 1y	0.29	0.08	0.06	0.01
Fwd 3y	0.15	0.24	0.18	0.11
Fwd 5y	0.14	0.21	0.31	0.24
Fwd 5y5y	0.12	0.13	0.34	0.37

Notes:

central bank narrative. Moreover, the narrative that matters for the long run appears to reflect information about forecasts rather than economic conditions, and also uncertainty around the forecasts.

5.3 Term premium information effects

In the existing information effect literature, emphasis is placed on outlook shocks the central bank's private information changing the mean of private sector expectations on future outcomes. But our earlier results suggest that neither HS nor NS1 are consistent with the long-maturity narrative information effect we find. In order to explore whether there is a key role for information on uncertainty driving the narrative effects, we examine whether the effects on the market news in the expectations and term premium components separately.

The first set of results, which we do not report here in the interest of space, repeat the methodology of section 4 to test for narrative information in the IR by component. We continue to find strong evidence for such information at all maturities and for each asset class.

We next repeat the bootstrap procedure to identify key topics for driving each component. Table 12 shows the results. One feature of interest is that there is substantial overlap between the topic variables that drive short-run movements in the overall yield change and the topic variables that drive the short-run expectations component. In contrast, there is an overlap between the top long-run overall and term premium topics. This indicates that the long-run information effect we have identified is driven largely by movements in long-run term premiums. We return to this point below.

Relatedly, the top topics for explaining long-run term premia are not the top topics for explaining the overall change in short-run yields. This is further evidence that long-run rate movements are not solely driven by mechanical market responses but by news. Table 13 is the analogue of table 9 by component. Here we find an interesting distinction between expectation and term premium topics: the topic selection percentages for expectations are much more correlated across maturities than that of term premiums. This is consistent with a common signal driving short- and long-run modal expectations, such as news about a persistent variable that enters into the nominal rate rule as in our framework. The information that drives the term premium at different maturities appears more heterogeneous, as would be natural if short-run and long-run risks depended on different aspects of the economy. The first results suggests there is some evidence for NS1, even though it didn't appear in the overall analysis, and reduces the likelihood of NS2. The second results provides strong evidence of our new TP channel.

We can also explore the contribution of narrative to the R^2 as above. Figure 13 replicates figure 10 from above showing, for the two elements of the decomposition, the R^2 as more topic variables are added. One first can observe that the quantitative forecast variables impact expectations more strongly than term premiums. This may also lend support to NS1 as occuring through the quantitative channel rather than the narrative.

Second, the topic variables are especially important for explaining longer-maturity term premiums. To more directly compare expectation versus term premium topics, Tables B.1 and B.2 in Appendix B include the top topics from each component as regressors for explaining the overall change in yields. As table 1 shows, for all maturities up to the five-year forward rate, the expectations component contributes more to the overall variation. For this reason, we find that for these maturities the expectations topics explain more than the term premium topics. For the five year, five year rate, by contrast, the expectations topics explain essentially none of the overall variation. The entire effect of narrative at this long horizon comes via changes in the term premium.

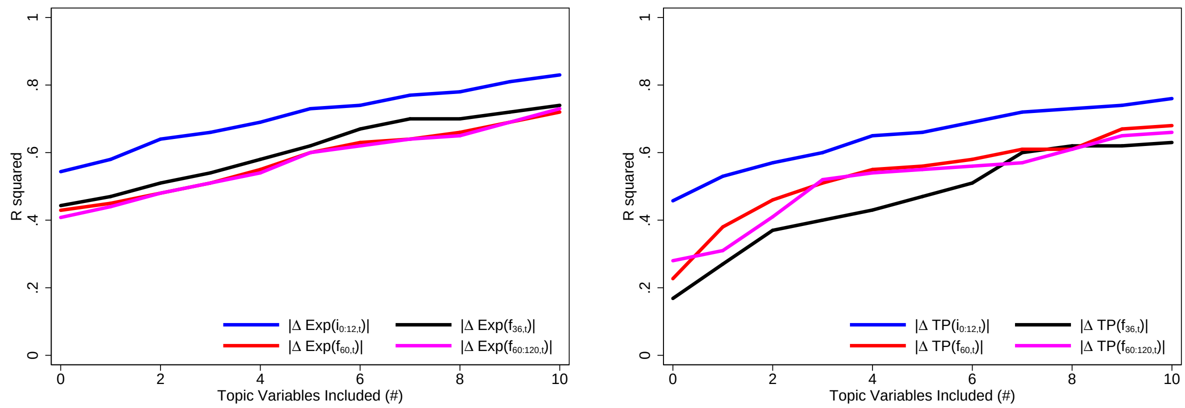

(a) Nominal Yields: Expectations **(b)** Nominal Yields: Term Premiums

Figure 13: Effect of adding more controls: R^2 - Nominal Yields Decomposition

Finally, table 14 repeats the placebo regression analysis for all possible combinations of asset decompositions and the top topics associated with each. Short-run expectations topics explain little of variation in long-rate term premiums; this clearly rules out the HS

Table 12: Top Topics for Different Yields (L=Level; D=Change)

(a) Expectations

$\|\Delta i_{0:12,t}\|$		$\|\Delta f_{36,t}\|$		$\|\Delta f_{60,t}\|$		$\|\Delta f_{60:120,t}\|$	
Var	Selection %	Var	Selection %	Var	Selection %	Var	Selection %
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

$\|\Delta i_{0:12,t}\|$		$\|\Delta f_{36,t}\|$		$\|\Delta f_{60,t}\|$		$\|\Delta f_{60:120,t}\|$	
Var	Selection %	Var	Selection %	Var	Selection %	Var	Selection %
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 (i.e. an element of $\boldsymbol{\theta}_t^R$), while a D indicates a residuals in the absolute change in the topic level (i.e. an element of $\boldsymbol{\delta}_t^R$).

Table 13: 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.

channel as the main driver of our information effect. Instead, the top topics explaining term premium moves for longer maturity rates explain an even greater proportion of the asset news than the top topics for the total asset news.

The overall message is that the long-run information effect appears to operate via both expectations and term premiums up to five years ahead, but for longer horizons it operates almost exclusively via term premiums. We find this a compelling counterpoint to arguments against the existence of a long-run information effect due to the implausibility of market point estimates of economic conditions changing many years in the future. Instead we show that the TP channel of the information effect need not directly change modal expectations to change long-run market rates. By conveying narrative about risk and uncertainty, central banks can affect second-order and higher moments of market beliefs, and thereby term premiums.

6 Conclusion

Conventional monetary policy typically moves a short-term interest rate in order to influence the interest rates banks, firms and households face across a variety of different maturities. Communication has offered an additional tool to do this and the relative importance of this channel likely increases when confronted with an effective lower bound on short-term interest rates.[17] While there is a growing literature showing the effects of central bank communication, there remain many important questions as to the channels through which it works.

We have used the Bank of England Inflation Report to show that there is an information effect of communication solely about the central bank's view on the outlook for economic fundamentals. Quantitative forecast information plays a relatively important role in moving interest rates at the short-end of the yield curve. However, we show that the narrative of the communication plays a relatively larger role at longer maturities. Moreover, this effect comes increasingly through term premiums for longer-maturity assets. Importantly, we find that different information affects the short- and long-end of the yield curve — there is, in addition to shorter-term news, a distinct long-term information effect of the central bank narrative. Our analysis suggests this new channel is the dominant one, at least in our sample.

Of course, while more remains to be done to understand fully the channels of central bank communication (and we hope that our methodological contributions in this paper

[17]For 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.

Table 14: Placebo Regressions: Partial R^2 when using Top Topics from Different Asset News

News	Topics											
	Spot1y	Exp-1y	TP-1y	3y	Exp-3y	TP-3y	5y	Exp-5y	TP-5y	5y5y	Exp-5y5y	TP-5y5y
Spot1y	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-1y	0.30	0.32	0.08	0.08	0.23	0.06	0.06	0.10	0.06	0.01	0.06	0.04
TP-1y	0.17	0.10	0.36	0.12	0.19	0.02	0.06	0.13	0.02	0.06	0.13	0.04
3y	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-3y	0.19	0.14	0.14	0.17	0.25	0.07	0.08	0.21	0.05	0.07	0.21	0.03
TP-3y	0.11	0.17	0.13	0.15	0.03	0.31	0.35	0.15	0.18	0.28	0.04	0.18
5y	0.14	0.17	0.12	0.19	0.04	0.26	0.31	0.19	0.14	0.24	0.04	0.15
Exp-5y	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-5y	0.12	0.12	0.08	0.08	0.08	0.38	0.38	0.08	0.42	0.44	0.09	0.39
5y5y	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-5y5y	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-5y5y	0.14	0.12	0.04	0.08	0.10	0.29	0.32	0.08	0.39	0.40	0.11	0.36

can be helpful in answering these), our main findings have an important policy implication. 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 there developments. Our results suggest that communication solely about future economic conditions, and the distribution of risks and uncertainties around these, may be sufficient to move long-run interest rates. That is, open-mouth operations may allow policymakers to talk down longer-term interest rates without the need to confirm how policy will react – the possibility of policy-free forward guidance.

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.

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., 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.

Boukus, E. and Rosenberg, J. (2006). The information content of FOMC minutes. Technical report, Federal Reserve Bank of New York.

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*.

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.

Cochrane, J. H. (2011). Presidential address: Discount rates. *The Journal of Finance*, 66(4):1047–1108.

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.

Ellison, M. and Tischbirek, A. (2018). Beauty Contests and the Term Structure. mimeo, University of Oxford.

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., 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).

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. (2017). Transparency and Deliberation within the FOMC: A Computational Linguistics Approach. *The Quarterly Journal of Economics*. Forthcoming.

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.

Jarociński and Karadi, P. (2017). Central Bank Information Shocks. Technical report, mimeo.

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.

Lucca, D. O. and Trebbi, F. (2009). Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements. NBER Working Papers 15367, National Bureau of Economic Research, Inc.

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. 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.

Melosi, L. (2017). Signalling Effects of Monetary Policy. *Review of Economic Studies*, 84(2):853–884.

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.

Tobback, E., Nardelli, S., and Martens, D. (2017). Between hawks and doves: measuring central bank communication. Working Paper Series 2085, European Central Bank.

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.

Zhang, D. (2017). Term Structure, Forecast Revision and the Information Channel of Monetary Policy. Technical report, mimeo.

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 **??** 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

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

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 **??**, 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) $\lvert\Delta i_{0:12;t}\rvert$	(2) $\lvert\Delta i_{0:12;t}\rvert$	(3) $\lvert\Delta i_{0:12;t}\rvert$	(4) $\lvert\Delta i_{0:12;t}\rvert$	(5) $\lvert\Delta i_{0:12;t}\rvert$
IR	0.016***	0.0073	0.014***	0.019**	0.031
	[0.000]	[0.317]	[0.002]	[0.013]	[0.151]
Announcement	0.0084***	0.00076	0.0023	0.0033	0.038
	[0.002]	[0.522]	[0.236]	[0.394]	[0.149]
Speech	-0.0022**	-0.0013**	-0.0019**	-0.0024	-0.0034
	[0.033]	[0.039]	[0.030]	[0.108]	[0.396]
Minutes	0.0046**	-0.0011	0.0019	0.0065**	0.024***
	[0.029]	[0.302]	[0.344]	[0.035]	[0.002]
VIX_t	0.00098***	0.00027***	0.00063***	0.0013***	0.0030***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Constant	0.0024*	0.0022***	0.0039***	0.0040**	0.0022
	[0.093]	[0.001]	[0.000]	[0.015]	[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) $\lvert\Delta f_{36;t}\rvert$	(2) $\lvert\Delta f_{36;t}\rvert$	(3) $\lvert\Delta f_{36;t}\rvert$	(4) $\lvert\Delta f_{36;t}\rvert$	(5) $\lvert\Delta f_{36;t}\rvert$
IR	0.018***	0.0041	0.011	0.027**	0.061***
	[0.005]	[0.327]	[0.137]	[0.050]	[0.000]
Announcement	0.0071***	0.0027	0.011***	0.0097***	0.0088
	[0.007]	[0.274]	[0.000]	[0.004]	[0.339]
Speech	0.0039**	-0.0016	0.0025	0.0068***	0.022***
	[0.030]	[0.143]	[0.312]	[0.008]	[0.005]
Minutes	0.0042	0.0016	0.0028	0.012**	-0.0027
	[0.109]	[0.404]	[0.464]	[0.031]	[0.591]
VIX_t	0.00092***	0.00025***	0.00091***	0.0013***	0.0030***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Constant	0.022***	0.0098***	0.014***	0.030***	0.045***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
R-squared	0.053				
Quantile	OLS	.25	.5	.75	.95
Sample	All	All	All	All	All

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

(a) 5-year forward rate

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

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

B Contribution of Expectations vs Term Premiums

Table B.1: Effect of Top Expectations and Term Premium Topics on R^2

	(1)	(2)	(3)	(4)	(5)	(6)
Main Regressors	$\lvert\Delta i_{0:12;t}\rvert$	$\lvert\Delta i_{0:12;t}\rvert$	$\lvert\Delta i_{0:12;t}\rvert$	$\lvert\Delta f_{36;t}\rvert$	$\lvert\Delta f_{36;t}\rvert$	$\lvert\Delta f_{36;t}\rvert$
VIX_t	0.0011**	0.0011**	0.0011*	0.0030*	0.0020	0.0029
	[0.05]	[0.05]	[0.06]	[0.09]	[0.25]	[0.11]
Topic L29	0.29	0.29			0.15	
	[0.41]	[0.40]			[0.79]	
d7	1.42		1.46	1.02	1.06	
	[0.19]		[0.24]	[0.60]	[0.57]	
Topic L18	0.81	0.97				
	[0.24]	[0.21]				
Topic L8	0.20	0.16				
	[0.66]	[0.74]				
d9		0.73				
		[0.36]				
Topic L2			0.31			
			[0.63]			
d13			0.46			
			[0.68]			
d22			0.76			
			[0.58]			
d29				-0.19	-0.28	
				[0.83]	[0.83]	
Topic L9				-2.02*		-2.11*
				[0.08]		[0.06]
d16				-1.25		
				[0.42]		
d27					-2.26	
					[0.34]	
d4						-0.65
						[0.67]
Topic L22						-0.35
						[0.70]
d20						1.34
						[0.29]
Constant	-0.11**	-0.12**	-0.082*	0.025	-0.019	0.0019
	[0.02]	[0.03]	[0.07]	[0.68]	[0.79]	[0.97]
R-squared	0.610	0.600	0.591	0.443	0.383	0.448
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	0.11	0.08	0.06	0.11	0.02	0.13

Table B.2: Effect of Top Expectations and Term Premium Topics on R^2

	(1)	(2)	(3)	(4)	(5)	(6)
Main Regressors	$\lvert\Delta f_{60;t}\rvert$	$\lvert\Delta f_{60;t}\rvert$	$\lvert\Delta f_{60;t}\rvert$	$\lvert\Delta f_{60:120;t}\rvert$	$\lvert\Delta f_{60:120;t}\rvert$	$\lvert\Delta f_{60:120;t}\rvert$
VIX_t	0.0026	0.0025	0.0018	0.0018	0.0018	0.0018
	[0.11]	[0.13]	[0.27]	[0.17]	[0.13]	[0.15]
Topic L9	-2.09*	-2.33**				
	[0.06]	[0.03]				
d4	0.12		-0.72	-0.26		-0.63
	[0.93]		[0.65]	[0.85]		[0.66]
d20	0.45		0.63	1.02		0.92
	[0.73]		[0.63]	[0.40]		[0.43]
Topic L6	0.23			0.85		
	[0.73]			[0.14]		
d7		0.85			-0.34	
		[0.64]			[0.82]	
d27		-1.64			-1.34	
		[0.35]			[0.44]	
d29		-0.60			0.15	
		[0.48]			[0.87]	
d12			0.26	0.49		0.34
			[0.81]	[0.63]		[0.70]
d6			-0.070			
			[0.93]			
d10					1.49**	1.37**
					[0.02]	[0.04]
Constant	0.030	0.057	0.012	-0.0021	0.029	0.020
	[0.61]	[0.32]	[0.85]	[0.97]	[0.62]	[0.71]
R-squared	0.374	0.383	0.287	0.335	0.352	0.358
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	0.12	0.14	0.01	0.08	0.10	0.12