



# Department of Economics Discussion Paper Series

## Narrative-Driven Fluctuations in Sentiment: Evidence Linking Traditional and Social Media

Alistair Macaulay and Wenting Song

Number 973  
June, 2022

# Narrative-Driven Fluctuations in Sentiment: Evidence Linking Traditional and Social Media\*

Alistair Macaulay  
University of Oxford

Wenting Song  
Bank of Canada

June 29, 2022

## Abstract

This paper studies the role of narratives for macroeconomic fluctuations. Micro-founding narratives as directed acyclic graphs, we show how exposure to different narratives can affect expectations in an otherwise-standard macroeconomic framework. We identify such competing narratives in news media reports on the US yield curve inversion in 2019, using techniques in natural language processing. Linking this to data from Twitter, we show that exposure to the narrative of an imminent recession causes consumers to display a more pessimistic sentiment, while exposure to a more neutral narrative implies no such change in sentiment. Applying the same technique to media narratives on inflation, we estimate that a shift to a viral narrative of inflation damaging the real economy in 2021 accounts for 42% of the fall in consumer sentiment in the second half of the year.

---

\*Macaulay ([alistair.macaulay@economics.ox.ac.uk](mailto:alistair.macaulay@economics.ox.ac.uk)): University of Oxford. Song ([wsong@bank-banque-canada.ca](mailto:wsong@bank-banque-canada.ca)): Bank of Canada. We thank John Leahy, Michael McMahon, Pablo Ottonello, Chris Roth, Leif Anders Thorsrud, and participants at various seminars and conferences for useful comments and suggestions. We also thank Andrew Usher for his help with the Twitter data. Alexa Kaminski provided excellent research assistance. Macaulay acknowledges financial support from the John Fell OUP Research Fund. The views expressed herein are those of the authors and not necessarily those of the Bank of Canada.

# 1. Introduction

Many decisions made by households and firms are influenced by expectations of future macroeconomic developments, but the factors determining these developments are often varied and nuanced. To help individuals make such decisions, popular narratives in the media provide simple causal stories of how a variable affects another. [Shiller \(2017\)](#) tracks the spread of certain “viral” narratives, which could cause or exacerbate existing macroeconomic fluctuations through their effect on expectations. To what extent do these narratives affect expectations once they have spread? And consequently, how substantial are the aggregate effects of shifting popular narratives about the macroeconomy?

In this paper, we answer those questions by linking traditional newspaper articles of macroeconomic events and engagement with that coverage on social media. Our theoretical framework specifies narratives as *directed acyclic graphs* (DAGs) (as in [Eliaz and Spiegler, 2020](#); [Andre, Haaland, Roth and Wohlfart, 2022b](#)), which drive fluctuations in expectations. Motivated by the theory, we capture competing narratives in traditional news media using natural language processing and trace the influence of those narratives by comparing the sentiment of Twitter users before and after engaging with a particular narrative. We study two importance macroeconomic events that have seen competing narratives: focusing on an episode of yield curve inversion in the US, we provide direct evidence that exposure to an imminent recession narrative causes consumers to display a more pessimistic sentiment; applying our framework to study recent inflation narratives, we discover a state-dependent effect of competing inflation narratives that is stronger when inflation is high.

Our paper begins by developing a theoretical framework of how narratives affect expectations. We start with a textbook consumption-and-saving problem faced by households and specify narratives as *directed acyclic graphs* (DAGs), or network representations of the underlying models, which have natural interpretation as “causal” stories. We consider two competing narratives: a baseline narrative in which expectations of future income depend on the current income and interest rate, and an extraneous narrative in which expectations of future income also depend on an extraneous variable, such as a popular recession indicator. The key contribution of this model is an equivalence result: we show that while the extraneous variable can enter into a household’s narrative about future incomes in a variety

of ways—as a *shock* affecting future income or as a *signal* of other variables—the resulting DAGs have observationally equivalent effects on expectations. This equivalence result implies that we do not need to distinguish between different narratives involving the extraneous variable. It is sufficient to identify only whether such a link exists between this variable and other variables in the baseline narrative.

Motivated by the theoretical framework, we measure narratives as the media’s competing interpretations of the same economic event. To do so, we use topic models from natural language processing on the news articles devoted to an economic event. We obtain empirical estimates of both the prevailing narratives and each article’s reliance on the narratives.

We then study the empirical importance of the identified narratives and their effects on consumer sentiment. The main empirical challenge for studying the relationship is that consumer sentiment can be influenced by the underlying economic event as well as the narratives around it. Our empirical strategy exploits an episode of a yield curve inversion episode in 2019—a popular recession indicator in the US with a nebulous theoretical foundation. The precise timing of the yield curve inversion was plausibly exogenous with respect to other macroeconomic news and monetary policy, which provides a quasi-experiment to isolate the effect of narratives on sentiment.

We uncover two competing narratives from major news outlets’ coverage: a “recession” narrative that links the inverted yield curve to an imminent recession and a “nonrecession” narrative that does not. We then identify the effects of narratives on the readers who are exposed. The most novel part of our data is the link from narratives in newspaper coverage to rich social network data from Twitter, which allows us to measure the spread of narratives. We use retweeting activities on Twitter to trace whether a consumer has engaged with news articles containing certain narratives. We find that exposure to negative narratives of an imminent recession causes consumers to display a more pessimistic sentiment, while exposure to the more neutral narrative without a strong link to recessions leads to no change in consumer sentiment. The drop in sentiment following engagement with a recessionary narrative is persistent, remaining significant 30 days after the retweet.

To assess the potential for viral narratives to drive aggregate sentiment, we then turn to narratives around inflation. These generally receive engagement from a wider audience, particularly since the second half of 2021 when inflation has become a prominent story across

news outlets. Applying the insights from our theory, we again use topic modeling to measure whether a given news article on inflation relates inflation to the real economy or not, and estimate whether exposure to a particular type of narrative is associated with a change in sentiment. In periods with low levels of realized inflation, there is no significant effect of the narratives on sentiment, consistent with low levels of attention when inflation is low and stable.<sup>1</sup> When realized inflation is high, however, engaging with a “New Keynesian” narrative in which inflation is linked to outcomes in the real economy is associated with a substantial decline in sentiment, while engaging with a “RBC”-type narrative, which mostly relates inflation to nominal variables, is associated with an improvement in sentiment. A back-of-the-envelope calculation suggests that 42% of the fall in aggregate sentiment in the Michigan Survey of Consumers in the second half of 2021 can be attributed to the shift towards narratives in which inflation can have damaging effects on the real economy. This suggests that changing narratives can be an important source of aggregate fluctuations.

**Related literature** Our paper relates to four strands of the literature. First, a growing literature pioneered by [Shiller \(2017\)](#) studies the role of narratives in economics.<sup>2</sup> Our contribution to this literature is twofold. Theoretically, we microfound narratives in a macroeconomic framework building on the Bayesian network literature ([Spiegler, 2016, 2020a; Eliaz and Spiegler, 2020](#)). Empirically, we develop a text-based measure of competing narratives that is directly connected to the theoretical framework, and link this to rich social media microdata for assessing the impacts on sentiments.

Our empirical methodology complements the semantics-based approach developed [Ash, Gauthier and Widmer \(2021\)](#) that is able to capture causal directions in narratives. We instead use topic models to capture narratives by leveraging the theoretical insight that DAGs with the same skeletons are observationally equivalent. Closely related to our methodology of topic models, [Larsen and Thorsrud \(2019\)](#) study the effects of narratives on business cycle fluctuations, defining narratives as significant economic events that are extracted using topic models on the corpus of newspaper articles. We instead capture narratives as news media’s competing interpretations of the *same* underlying economic event, which is motivated by our

---

<sup>1</sup>See, for example, [Pfäuti \(2022\)](#).

<sup>2</sup>Also see the body of work that highlights importance of political narratives, which includes, for example, [Gentzkow, Shapiro and Sinkinson \(2014\)](#), [Levy \(2021\)](#), and [Bianchi, Kung and Cram \(2021\)](#).

theoretical framework and models of competing narratives (Eliaz and Spiegler, 2020). We provide direct evidence on the importance of narratives by exploiting the natural experiment of the yield curve inversion in 2019, complementing survey-based evidence from Andre et al. (2022b) who conduct survey experiments to establish the causal effects of narratives on expectations, and Macaulay (2022) who presents evidence from UK household surveys on the importance of inflation narratives. Our empirical framework has the benefit of providing an ongoing measure of narratives outside of existing surveys, as we illustrate with the analysis on inflation narratives.

Second, narratives provide a way for individuals to interpret economic news and translate that into expectations, and therefore also relate to studies of differences of opinion (Harris and Raviv, 1993; Patton and Timmermann, 2010; Xiong and Yan, 2010; Atmaz and Basak, 2018) and subjective models (Dräger, Lamla and Pfajfar, 2016; Andrade, Gaballo, Mengus and Mojon, 2019; Molavi, 2019; Andre, Pizzinelli, Roth and Wohlfart, 2022a).

This paper also relates to the broader literature of belief formation. Empirical evidence documents the deviations by households and firms from full-information rational expectations (see Coibion, Gorodnichenko and Kamdar, 2018, for a comprehensive survey). Previous literature points to inattention (Sims, 2003; Mankiw and Reis, 2002), personal experiences (Malmendier and Nagel, 2016), salience (Cavallo, Cruces and Perez-Truglia, 2017), heuristics (Bordalo, Gennaioli and Shleifer, 2018), wishful thinking (Caplin and Leahy, 2019), among others, as important drivers of individuals' expectations. We provide direct evidence on the importance of narratives.

Third, we relate to the literature on sentiment and media. Our results highlight the role of economic narratives in shaping household sentiments, which are important sources of macroeconomic fluctuations (see, for example, Angeletos and La'O, 2013; Greenwood and Shleifer, 2014; Levchenko and Pandalai-Nayar, 2020; Maxted, 2019; Krishnamurthy and Li, 2020; Acharya, Benhabib and Huo, 2021). We contribute to the literature by showing that narratives constructed by the media provides a microfoundation for fluctuations in sentiment. We highlight, in particular, the role of media in curating news and constructing narratives, consistent with theories of news media as optimizing agents whose news reporting drives aggregate fluctuations (Nimark, 2014; Chahrour, Nimark and Pitschner, 2021).

Lastly, our unique data linking news coverage to its influence on social media allows us

to measure the impact of narratives constructed by the media on household beliefs, which relates to the growing literature that uses unstructured data sources to study the economic effects of news (see, for example, [Calomiris and Mamaysky, 2019](#); [Bybee, Kelly, Manela and Xiu, 2020](#); [Nyman, Kapadia and Tuckett, 2021](#)), and that exploits rich social network data to study the effects of policy (see, for example, [Bailey, Cao, Kuchler and Stroebel, 2018](#); [Gorodnichenko, Pham and Talavera, 2021](#); [Bianchi et al., 2021](#); [Matveev and Ruge-Murcia, 2021](#); [Haldane, Macaulay and McMahon, 2021](#); [Ehrmann and Wabitsch, 2022](#)).

**Outline** The rest of the paper proceeds as follows: in Section 2 we present our theoretical framework that connects narratives with expectations and derive conditions for observationally-equivalent narratives; in Section 3 we describe the episode of yield curve inversion in 2019; in Section 4 we describe our data and sample; in Section 5 we conduct our main empirical analysis on the narratives surrounding the yield curve inversion by linking news articles and social media; in Section 6 we further apply our empirical framework to study inflation narratives; Section 7 concludes.

## 2. Model

In this section we develop a framework to analyse the role of narratives in shaping household expectations and actions, in an otherwise standard consumption-saving problem. A narrative is defined as a causal ordering of variables, represented by a DAG. Importantly, we show that certain groups of narratives are observationally equivalent, in that they always produce the same household expectations. This will guide our approach to distinguishing between relevant narratives in the data in Sections 5 and 6.

### 2.1. Households

We start with a standard partial-equilibrium consumption problem faced by households. Each household chooses consumption to maximise their life-time utility subject to the budget constraint under the expectation  $E_{it}$ , taking the interest rate and income as given. Each

household  $i$  has preferences over consumption given by

$$\sum_{s=0}^{\infty} \beta^s \mathbb{E}_{it} u(C_{it+s})$$

where  $\beta$  is the discount factor;  $\mathbb{E}_{it}$  is the subjective expectation of household  $i$  given the time- $t$  information set; and the instantaneous utility function is CRRA, specified as

$$u(C_{it}) = \frac{C_{it}^{1-\frac{1}{\sigma}} - 1}{1 - \frac{1}{\sigma}}$$

Each period, the household receives real income  $Y_t$ , and can purchase one-period bonds  $B_{it}$  with a real interest rate of  $R_t$ . Their budget constraint is therefore given by

$$C_{it} + B_{it} = R_{t-1} B_{it-1} + Y_t$$

For simplicity, we take income  $Y_t$  to be exogenous to household  $i$ 's decisions.

The optimization leads to a standard consumption Euler equation. Log-linearizing the Euler equation and the budget constraint about a steady state in which  $\beta R = 1$  and  $B_i = 0$  gives the household's time- $t$  *consumption function* as

$$c_{it} = (1 - \beta) \sum_{s=0}^{\infty} \beta^s \mathbb{E}_{it} y_{t+s} - \sigma \beta \sum_{s=0}^{\infty} \beta^s \mathbb{E}_{it} r_{t+s} \quad (1)$$

where lower case  $c_{it}, y_t, r_t$  denote log-deviations of consumption, income, and real interest rates from their respective steady states.

Equation (1) shows that households' current consumption is driven by their expectations of future real income and real interest rates. Households observe the history of  $y, r$  up to the current period, but to form expectations of future realizations they must combine this with a belief about the evolution of both variables. We introduce narratives as the source of these beliefs relating observations to expectations.

## 2.2. Narratives

We follow [Eliaz and Spiegel \(2020\)](#) and [Andre et al. \(2022b\)](#) and define a narrative as a *directed acyclic graph* (DAG), that defines a series of causal relationships between variables.



For a thorough review of this approach to modeling expectations, see [Spiegler \(2020a\)](#).

**Definition 1** (narrative as a DAG). *A narrative for income and interest rates is defined as a DAG consisting of:*

1. *a set of nodes  $\mathcal{N}$ , where each element is a real-valued economic variable; and*
2. *a set of links  $\mathcal{L}$  which define the directed causal links between nodes.*

*The set of nodes  $\mathcal{N}$  contains current and future values of  $y$  and  $r$ , and potentially other additional variables. The links  $\mathcal{L}$  are acyclic: they are such that the graph contains no directed path from a node back to itself.*

The nodes of the DAG correspond to variables in the household environment. The links correspond to perceived causal relationships between those variables.

These links are crucial in the household’s decision problem. To choose consumption, the household forms expectations of future income and interest rates, *conditional* on observed current variables. To form that conditional expectation, they require a belief about the joint distribution of the variables involved. The narrative guides that belief, through the Bayesian factorization formula

$$\tilde{p}(x_{\mathcal{N}}) = \prod_{n \in \mathcal{N}} p(x_n | x_{\mathcal{L}(n)}) \quad (2)$$

where  $x_{\mathcal{N}}$  denotes the set of all variables in the narrative; and  $x_{\mathcal{L}(n)}$  denotes the subset of those variables which have a direct causal link to variable  $x_n$  in the narrative.

A narrative therefore specifies which conditional distributions should be involved in forming their beliefs about the joint distribution of all variables  $\tilde{p}(x_{\mathcal{N}})$ , which may or may not equal the true joint distribution. The perceived causal links between variables imply a series of conditional independence assumptions, that will affect how the households interpret data on the variables in their environment. We follow [Eliaz and Spiegler \(2020\)](#) and assume that the households observe a long time series of such data on each variable, and so are able to accurately recover the true conditional distributions  $p(x_n | x_{R(n)})$  involved in this factorization.


Expectations are then formed using the perceived joint distribution  $\tilde{p}(x_{\mathcal{N}})$ . This means that if the household’s narrative correctly accounts for the true causal links between variables,

Equation (2) yields the true joint distribution of the variables in their environment. Such a household will therefore have rational expectations.

However, if the narrative incorrectly specifies the true causal links between variables, the implied  $\tilde{p}(x_{\mathcal{N}})$  may not coincide with the true joint distribution. A household with such a narrative interprets data through the lens of a misperceived causal model, which may cause them to use incorrect assumptions about the conditional (in)dependence of certain variables. That, in turn, may generate incorrect beliefs about the joint distribution of variables in their environment. In that case, the expectations of these households will not coincide with rational expectations.

**Baseline Narrative** The first narrative we consider is displayed in Figure 1. In this narrative, real income is persistent, so  $y_t$  has a causal effect on  $y_{t+1}$ . In addition, real incomes affect contemporaneous interest rates, for instance because the central bank reacts to demand conditions through a standard Taylor rule. Changes in real interest rates then in turn affect real incomes with a lag, so  $r_t$  has a causal effect on  $y_{t+1}$ .<sup>3</sup> We do not specify here whether this narrative represents a correct understanding of causal relationships in the equilibrium of the economy, or whether it is a misperception of true economic relationships.

**Figure 1:** DAG representation of the baseline narrative



We refer to this narrative as the “baseline narrative”. The only variables that matter for expectations of  $y_t$  and  $r_t$  are lags of those variables themselves. Formally, the baseline narrative is defined as follows:

**Definition 2** (baseline narrative). *Let  $nRm$  denote a directed link from node  $n$  to node  $m$ . The baseline narrative is a DAG consisting of:*

1. *the set of nodes,  $\mathcal{N} = \{y_s, r_s\}_{s=t}^{\infty}$ ; and*
2. *the set of links,  $\mathcal{L} = \{y_s R y_{s+1}, y_s R r_s, r_s R y_{s+1}\}$ .*

---

<sup>3</sup>This lag is important to ensure that the graph remains acyclic, as required by the definition of a DAG.

**Extraneous Narratives** We now introduce a competing group of narratives, which introduce an extraneous variable,  $z$ , into the causal ordering of variables.

These “extraneous narratives” could reflect true causal relationships in the economy, or the extraneous variable could be entirely spurious. Politicians or news media may have incentives to create such spurious narratives to influence expectations or household behavior (Gentzkow and Shapiro, 2008; Eliaz and Spiegel, 2020).

In particular, we consider a class of narratives in which  $z_s$  is perceived to be related to real income in periods  $s$  and  $s + 1$ .

**Definition 3** (extraneous narratives). *The extraneous narratives are DAGs consisting of:*

1. the set of nodes,  $\mathcal{N} = \{y_s, r_s, z_s\}_{s=t}^\infty$ ; and
2. one of the sets of links  $\mathcal{L}_a$ ,  $\mathcal{L}_b$ , or  $\mathcal{L}_c$ , where:

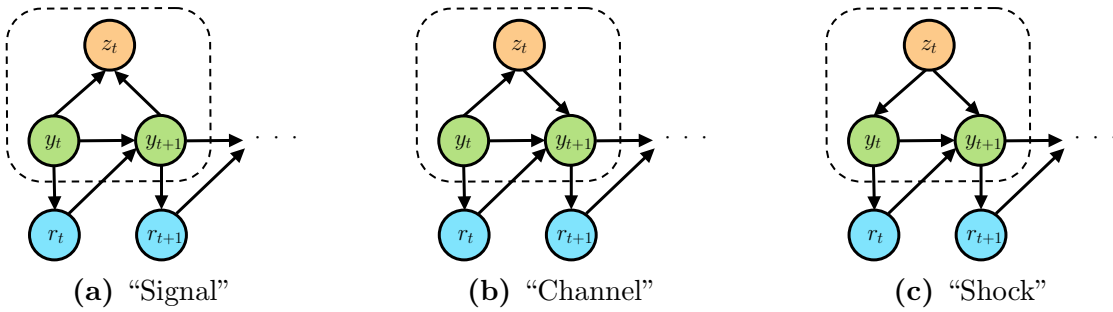
$$(a) \mathcal{L}_a = \mathcal{L} \cup \{y_s R z_s, y_{s+1} R z_s\};$$

$$(b) \mathcal{L}_b = \mathcal{L} \cup \{y_s R z_s, z_s R y_{s+1}\};$$

$$(c) \mathcal{L}_c = \mathcal{L} \cup \{z_s R y_s, z_s R y_{s+1}\};$$

Importantly, even restricting extraneous narratives to this class, there are still three possible ways for  $z$  to enter the household’s causal model. These are shown in Figure 2.

**Figure 2:** DAG representations of extraneous narratives



In the narrative in Panel (a),  $z$  is caused by the income process and is a symptom of the underlying economic fundamentals. It therefore signals changes in income without being a cause of that change. In the narrative in Panel (b),  $z$  is a channel through which the current income affects the future income. In the narrative in Panel (c),  $z$  is an exogenous shock that affects income.

We now go on to derive the processes for expectations implied by these narratives.

### 2.3. Expectations

To find how narratives affect expectations, we first find the Bayesian factorization formulae for each of the narratives described above.

For the baseline narrative, we have

$$\tilde{p}(r_s, r_{s+1}, y_s, y_{s+1}, z_s) = p(r_s|y_s)p(r_{s+1}|y_{s+1})p(y_s)p(y_{s+1}|r_s, y_s)p(z_s) \quad (3)$$

The equivalent factorization formulae for the three extraneous narratives are

$$\tilde{p}_a(r_s, r_{s+1}, y_s, y_{s+1}, z_s) = p(r_s|y_s)p(r_{s+1}|y_{s+1})p(y_s)p(y_{s+1}|r_s, y_s)p(z_s|y_s, y_{s+1}) \quad (4)$$

$$\tilde{p}_b(r_s, r_{s+1}, y_s, y_{s+1}, z_s) = p(r_s|y_s)p(r_{s+1}|y_{s+1})p(y_s)p(y_{s+1}|r_s, y_s, z_s)p(z_s|y_s) \quad (5)$$

$$\tilde{p}_c(r_s, r_{s+1}, y_s, y_{s+1}, z_s) = p(r_s|y_s)p(r_{s+1}|y_{s+1})p(y_s|z_s)p(y_{s+1}|r_s, y_s, z_s)p(z_s) \quad (6)$$

Systematically distinguishing between these different extraneous narratives in media would be challenging. However, despite the different interpretations of these three DAGs, Proposition 1 shows that their effects on households' beliefs are in fact observationally equivalent.

**Proposition 1** (observational equivalence of extraneous narratives). *The Bayesian factorization formulae for the three extraneous narratives are equivalent:*

$$\tilde{p}_a(r_s, r_{s+1}, y_s, y_{s+1}, z_s) = \tilde{p}_b(r_s, r_{s+1}, y_s, y_{s+1}, z_s) = \tilde{p}_c(r_s, r_{s+1}, y_s, y_{s+1}, z_s)$$

*Proof.* We begin by showing  $p_c(\cdot) = p_b(\cdot)$ . By the definitions of joint and conditional probabilities:

$$\begin{aligned} \tilde{p}_c(r_s, r_{s+1}, y_s, y_{s+1}, z_s) &= p(r_s|y_s)p(r_{s+1}|y_{s+1})\frac{p(y_s, z_s)}{p(z_s)}p(y_{s+1}|r_s, y_s, z_s)p(z_s) \\ &= p(r_s|y_s)p(r_{s+1}|y_{s+1})p(y_s)p(z_s|y_s)p(y_{s+1}|r_s, y_s, z_s) \\ &= \tilde{p}_b(r_s, r_{s+1}, y_s, y_{s+1}, z_s) \end{aligned}$$

Similarly, we can show  $p_b(\cdot) = p_a(\cdot)$ :

$$\begin{aligned}\tilde{p}_b(r_s, r_{s+1}, y_s, y_{s+1}, z_s) &= p(r_s|y_s)p(r_{s+1}|y_{s+1})p(y_s)\frac{p(y_{s+1}, z_s|r_s, y_s)}{p(z_s|r_s, y_s)}p(z_s|y_s) \\ &= p(r_s|y_s)p(r_{s+1}|y_{s+1})p(y_s)p(y_{s+1}|r_s, y_s)p(z_s|y_s, y_{s+1}) \\ &= \tilde{p}_a(r_s, r_{s+1}, y_s, y_{s+1}, z_s)\end{aligned}$$

where the penultimate equality uses that  $p(z_s|y_s, r_s) = p(z_s|y_s)$ , as  $r_s$  is not directly causally related to  $z_s$ .  $\square$

Intuitively, if one household believes that a rise in the variable  $z$  causes incomes to fall (“shock”), and another believes instead that falling incomes cause  $z$  to rise (“signal”), then both will revise their expected incomes down when they observe higher  $z$ . Formally, this property emerges because all of the extraneous narrative DAGs are “perfect”: the direct causes of any downstream variable are all themselves directly linked together (“all parents are married”). All perfect DAGs with the same skeleton necessarily share the same Bayesian factorization formula (Verma and Pearl, 1990).<sup>4</sup>

Proposition 1 implies that we do not need to consider the three extraneous narratives in Figure 2 separately. From here, we therefore refer to *the* extraneous narrative to mean any narrative satisfying Definition 3.

The next Proposition shows that, in general, the baseline narrative and the extraneous narrative generate different Bayesian factorization formulae.

**Proposition 2** (nonequivalence of baseline and extraneous narratives). *If  $z_s$  is correlated with  $y_s$  and/or  $y_{s+1}$ , then the following two Bayesian factorization formulae are nonequivalent:*

1.  $\tilde{p}(r_s, r_{s+1}, y_s, y_{s+1}, z_s) = p(r_s|y_s)p(r_{s+1}|y_{s+1})p(y_s)p(y_{s+1}|r_s, y_s)p(z_s)$
2.  $\tilde{p}_a(r_s, r_{s+1}, y_s, y_{s+1}, z_s) = p(r_s|y_s)p(r_{s+1}|y_{s+1})p(y_s)p(y_{s+1}|r_s, y_s)p(z_s|y_s, y_{s+1})$

*Proof.* Since  $z_s$  is correlated with  $y_s$  and/or  $y_{s+1}$ , they are not conditionally independent. As a result,  $p(z_s) \neq p(z_s|y_s, y_{s+1})$ , which implies  $\tilde{p}(r_s, r_{s+1}, y_s, y_{s+1}, z_s) \neq \tilde{p}_a(r_s, r_{s+1}, y_s, y_{s+1}, z_s)$ .  $\square$

---

<sup>4</sup>See Spiegel (2020b) for a detailed discussion of the implications of perfection in DAGs used to represent the causal mental models of decision-makers in a variety of economic contexts.

Note that the narratives only imply different perceived joint distributions if  $z$  is correlated with current or future real incomes. However, this does not imply that the extraneous narrative describes the true causal relationships between the variables, as this reduced-form correlation may be present even if there is no true causal relationship between  $z$  and  $y$ . Some other variable not in any household's narrative, for example, could simultaneously cause both variables. In that case households believing the extraneous narrative are basing their expectations on a spurious correlation.

This model therefore predicts that a household's expectations depend on which narrative they are exposed to. Using the extraneous narrative, expectations of real income one period ahead are

$$\mathbb{E}_{it}^e(y_{t+1}|\mathcal{I}_t) = \int y_{t+1}p(y_{t+1}|r_t, y_t, z_t)dy_{t+1} \quad (7)$$

In contrast, using the baseline narrative, the same expectation is

$$\begin{aligned} \mathbb{E}_{it}^b(y_{t+1}|\mathcal{I}_t) &= \int y_{t+1}p(y_{t+1}|r_t, y_t)dy_{t+1} \\ &= \int \int y_{t+1}p(y_{t+1}|r_t, y_t, z_t)p(z_t|r_t, y_t)dz_tdy_{t+1} \end{aligned} \quad (8)$$

From Proposition 2, these are different whenever the extraneous variable  $z$  is correlated with real income. In particular, expectations react to realized  $z_t$  under the extraneous narrative, but do not react under the baseline narrative.


$$\frac{\partial \mathbb{E}_{it}^e(y_{t+1}|\mathcal{I}_t)}{\partial z_t} \neq 0, \quad \frac{\partial \mathbb{E}_{it}^b(y_{t+1}|\mathcal{I}_t)}{\partial z_t} = 0 \quad (9)$$

This guides our empirical exercise. In the following sections, we consider narratives in which the extraneous variable  $z$  is whether the yield curve on US Treasuries is inverted or not. Following Proposition 1, we only attempt to distinguish whether a media narrative implies a causal link between yield curve inversion and future incomes or not. Matching these media accounts to data from Twitter, we test whether exposure to such an extraneous narrative implies a differential response of expectations to the yield curve inversion in 2019, as implied by Equation (9).


### 3. Yield Curve Inversion

Given our theoretical framework, we now study the impact of narratives on household sentiment. The key empirical challenge is that narratives are ingrained with the underlying economic events, making it difficult to isolate the effects of narratives. To address this, we exploit a unique episode of yield curve inversion in 2019 to provide direct evidence on the importance of narratives.

**Figure 3:** Timeline of the yield curve inversion episode



(a) Treasury spreads



(b) Media coverage and Google searches for “yield curve”

*Notes:* Panel (a) shows the spread between 10-year treasury yield and 3-month treasury yield (“10Y3M”) and the spread between 10-year treasury yield and 2-year treasury yield (“10Y2Y”) in 2019. Dates when the spreads first turn negative and revert back to positive are annotated. Panel (b) shows the number of news articles from Factiva containing the term “yield curve” and the Google search frequency in 2019. Google search frequency for the term “yield curve” has been scaled so the maximum value is 100.

Yield curve inversions have been a closely-watched indicator of upcoming recessions in the U.S. since [Harvey \(1988\)](#) documented their predictive power for major recessions from the 1960s to the 1980s. Figure [A.1](#) in the Appendix shows that the spread between the 10-year and 2-year Treasury bond yields has turned negative within 12 months before every recession in the US for the past 40 years.

When the yield curve inverted in 2019, it therefore received substantial attention from households and the media. Figure [3a](#) plots the timeline of the inversion, showing that the most widely-watched 10-year-over-2-year (10Y2Y) term spread inverted on August 28 and un-inverted on August 30. Figure [3b](#) shows that media coverage<sup>5</sup> and Google searches for the term “yield curve” spiked before and during the inversions of both the 10Y2Y term spread and the 10-year-over-3-month (10Y3M) term spread, with a peak of interest right before the inversion of the 10Y2Y spread.

Against the backdrop of a booming labor market and the longest expansion in US history, the inversion received several different interpretations in the media. The first interpretation is that a recession is looming. An example of such a recession narrative is Cristina Alesci’s article for CNN<sup>6</sup>:

Navarro is wrong on two fronts: The inversion did happen, and it’s not a good sign for the economy. Although the inversion was brief and small, major banks took note of it. [...] Yield curve inversions often signal recessions, which is why economic prognosticators pay so much attention to them.

which draws on the track record yield curve inversion to predict a recession and paints a negative picture on the economic outlook. Notably, the argument draws on both the “signal” narrative in Figure [2](#) (“inversions often signal recessions”) and the “shock” narrative (“major banks took note of it”). This highlights the intuition for Proposition [1](#): both of these narratives imply readers should update their expectations towards believing a recession is likely. It also underlines the importance of Proposition [1](#) for our empirical exercise, as it implies we do not need to disentangle these often-combined narratives to estimate the effects of the narrative on expectations.

---

<sup>5</sup>We measure media coverage using weekly data from Factiva. We obtain the number of nonduplicate news articles containing the term “yield curve” and restrict articles to be in English and specific to the US.

<sup>6</sup>“Fact-checking Peter Navarro’s claims that the yield curve is not inverted” by Cristina Alesci on August 19, 2019. [Link](#) to the article on CNN.



The second common interpretation is that the yield curve inversion is no longer an informative signal. Peter Coy illustrates such a narrative for Bloomberg<sup>7</sup>:

Well, guess what, folks? It’s still rainbows and pots of gold out there. Contrary to what seems to have become the overnight conventional wisdom in politics, a recession before Election Day 2020 remains a less than 50-50 proposition.

which goes on to explain that the long end of the yield curve has been trending down because of low and stable inflation and the strong fundamentals of the economy, suggesting that recession concerns are overblown. This corresponds to the “baseline narrative” in section 2.

The articles by Cristina Alesci and Peter Coy are strong examples of each of these narratives. Some other media reports on the yield curve inversion instead present a mix between the two narratives. For example, Brian Chappatta’s Bloomberg article<sup>8</sup> explains the nature of the yield curve and the historical significance of its inversion:

What’s a yield curve? [...] What are flat and inverted yield curves? [...] Why does it matter?

This defines an inverted yield curve, explains its history of proceeding recessions, but does not draw strong conclusions of what the inversion implies for the current economy.

Do these narratives influence the outlook of their readers? And if so, how much influence does each narrative have?

## 4. Data

### 4.1. Newspaper articles

We capture narratives as media’s different interpretations of the yield curve inversion (see Section 5.1 for details of our measurement approach). To form the media corpus for our analysis, we collect news articles covering the inversion of the 10Y2Y spread. Our data source is Factiva, a news database, and news outlets’ websites. To separate the effects of economic narratives from political narratives, we focus on news outlets classified as “centrist”

---

<sup>7</sup>“What a Yield-Curve Inversion Really Says About the U.S. Economy: A reliable recession indicator has lost some of its power to predict” by Peter Coy on August 22, 2019. [Link](#) to the article on Bloomberg.

<sup>8</sup>“The Yield Curve Is Inverted! Remind Me Why I Care” by Brian Chappatta. [Link](#) to the article on Bloomberg.

**Table 1:** Media outlets and coverage on the yield curve inversion

Outlet	Ideology placement	Twitter handle	# base tweets	# articles
MSNBC	Liberal/Center	msnbc	4	1
CNN	Liberal/Center	cnn	8	4
NBC News	Center	nbcnews	4	1
CBS News	Center	cbsnews	3	3
Bloomberg	Center	business	143	68
ABC News	Center	abc	1	1
USA Today	Center	usatoday	1	1
Yahoo News	Center	yahoonews	3	3
Wall Street Journal	Center	wsj	9	6
Fox News	Conservative/Center	foxbusiness	0	0
Total			176	88

*Notes:* Media outlets with centrist political leaning and their coverage of the yield curve inversion. Data on media outlets’ political placement is from (Jurkowitz et al., 2020), which determines the political ideology of an outlet by surveying the political leaning of its audience. The twitter handles of news outlets are hand searched. The tweets and articles on the yield curve are collected as described in Section 5.1.

by the Pew Research Center and exclude news aggregators such as Google News.<sup>9</sup> The 10 news outlets included in our sample in listed in Table 1.

During the event window of August 19 to September 13, 2019 (one week before the inversion and two weeks after the un-inversion, respectively)<sup>10</sup>, we search for tweets by news outlets which contains both “yield curve” and any of the stems from “invert”, “invers”, or “recession”. These “base tweets” by news outlets contain URLs to their webpages containing the full-length news articles, which form the corpus from which we extract narratives. Table 1 shows that the search criteria lead to 176 base tweets, linking to 88 unique articles.

## 4.2. Twitter

Our Twitter data consists of three parts. First, as described in the last subsection, we use outlet’s base tweets to identify news articles related to the yield curve inversion. We collect base tweets using Twitter’s Enterprise Search API, which contains the full archive of tweets since the start of Twitter in 2006.

<sup>9</sup>Jurkowitz, Mitchell, Shearer and Walker (2020) determine the political bias of a media outlets by surveying the political ideology of its audience.

<sup>10</sup>Although the yield curve was inverted from August 26 to August 30, media coverage and Google search trends in Figure 3b suggest that the interests in the yield curve rose before the actual inversion and stayed elevated after the un-inversion. Therefore, we expand the search window for news articles to one week before the inversion and two weeks after the un-inversion.

**Table 2:** Descriptive statistics on base tweets and retweeting users**(a)** Outlets’ base tweets on the yield curve

	Mean	SD	5th Pctl	Median	95th Pctl	Obs
Quote retweet count	8.5	39.1	0	3	28.2	178
Retweet count	45.4	89.9	0	23	162.6	178
Reply count	8.8	25.0	0	4	25.3	178
Favorite count	67.4	120.6	0	35	235.8	178

**(b)** Retweeting users

	Mean	SD	5th Pctl	Median	95th Pctl	Obs
# tweets	3,863	14,948	6	637	15,368	404
# outlets	3.5	2.5	1	3	8	404

*Notes:* Panel (a) reports descriptive statistics of media outlets’ tweets about the yield curve inversion between August 19 and September 13, 2019. The table reports descriptive statistics of the numbers of quote retweets, retweets, replies and favorites of media outlets’ tweets. Panel (b) reports descriptive statistics of users’ Twitter activity based on tweets one month before and one month after the quote retweets of the base tweets.

Second, we use the rich network data available from Twitter to measure a user’s exposure to narratives. Twitter provides four ways of interacting with posted tweets: quote retweet, retweet, reply and like. A “retweet” is when a user forwards a tweet without adding any comments, while a “quote retweet” requires that a user writes additional text when retweeting. The additional commentaries added by quote retweeters imply the absorption of new information contained in the articles linked in the base tweets. Therefore, we use quote retweets as the main measure of exposure to narratives. Through Twitter’s Standard API, we have information on the first 100 users who have quote retweeted each base tweet. Table 2a summarizes the retweeting activities of the base tweets on the yield curve. On average the base tweets in the sample have 9 quote retweets, and the 95 percentile has 28 quote retweets, far below the API constraint of 100 users.

Third, we measure changes in Twitter users’ sentiment after they are exposed to a narrative by measuring the sentiment of their tweets on all subjects. For all users who have quote retweeted any of the base tweets on the yield curve, we collect every tweet posted in a 1-month window around the quote retweet, again using the Enterprise API. Table 2b reports descriptive statistics of tweeting activity for the users in our sample, which shows that the median user is active and posts around 10 tweets per day. We measure the sentiment of a

tweet using a naïve Bayes classifier trained specifically to analyze the colloquial language on Twitter (for more details see Appendix C).<sup>11</sup> The sentiment score lies between 0 and 1, which is a uniform scale increasing with sentiment. A score greater than 0.5 corresponds to positive sentiment, and a score less than 0.5 corresponds to negative sentiment. To validate the sentiment measure, we present in Appendix Table A.1 the top 5 positive and negative tweets related to the yield curve, which demonstrates that the trained naïve Bayes classifier provides an accurate measure of tweet sentiment.

## 5. Narrative-Driven Fluctuations in Sentiment

### 5.1. Measuring narratives with topic models

As the theoretical framework in Section 2 illustrates, the distinguishing feature between narratives is their network structures. CNN’s “fact checking Navarro” presents a direct causal connection between the yield curve inversion and macroeconomic output, corresponding to an “extraneous narrative”. Bloomberg’s “rainbows and pots of gold,” on the other hand, dismisses the possibility of the inversion predicting an imminent recession. Under this “baseline narrative”, the yield curve inversion is disconnected from output and incomes. The coverage by Bloomberg’s Brian Chappata can be empirically interpreted as a mix between the two narratives.

We extract these economic narratives from news articles using latent Dirichlet allocation (LDA), as developed by Blei, Ng and Jordan (2003) for natural language processing. Appendix B provides details on the LDA model.<sup>12</sup> LDA is a Bayesian factor model that uncovers topics in the articles and represents each article in terms of these topics. It reduces the dimensionality of the text from the entire corpus of articles to just  $K$  “topics”, or groupings of words that tend to appear together.

LDA is our preferred tool for capturing narratives for two reasons. First, Proposition 1 shows that the direction of the links in the DAGs for the different extraneous narratives does not affect how a narrative influence a consumer’s expectations. In other words, what

---

<sup>11</sup>As recognized by Buehlmaier and Whited (2018), naïve Bayes is one of the oldest tools in natural language processing and has better out-of-sample performance in text-based tasks than alternative models (Friedman, Hastie, Tibshirani et al., 2001).

<sup>12</sup>Also see Hansen, McMahon and Prat (2018) for a discussion on LDA and its application in macroeconomics.

differentiates narratives empirically is whether phrases related to “yield curve” and “recession” are grouped together in one topic, and not the direction of causality between these words.<sup>13</sup> This is precisely what LDA is designed to capture. Second, one of the outputs of LDA is  $\theta(d, k) \in (0, 1)$ , the loading of article  $d$  on narrative  $k$ , which allows for the possibility that an article can contain multiple narratives and provides estimated loadings on each narrative. Therefore, LDA can capture polarizing articles containing a single narrative as well as balanced ones with multiple narratives.

We estimate the LDA with  $K = 5$  and symmetric Dirichlet priors<sup>14</sup>. Since LDA is a multi-membership model, the word “recession” can appear in multiple topics.  $K = 5$  is the smallest number of topics that ensures at least one topic does not contain “recession”.

## 5.2. Yield-curve-inversion narratives

The estimated topics from the LDA are shown in Figure 4. Two topics, in particular, consist of groupings of words that correspond to the theoretical definitions of the yield curve narratives in Section 2. The first topic in Panel (a) features the terms such as “recession,” “yield curve,” “economy” and “Trump,” mapping naturally to a “recession” narrative, corresponding to the extraneous narrative in our theoretical framework. It discusses the economic policy by the Trump administration in conjunction with the yield curve inversion and recession risks. The second topic in Panel (b) contains a broader discussion of other factors affecting the economy and investment opportunities in the bond and stock markets. Since it does not directly connect the slope of the yield curve to a coming recession, we interpret it as a “nonrecession” narrative, corresponding to the baseline narrative in our theoretical framework.

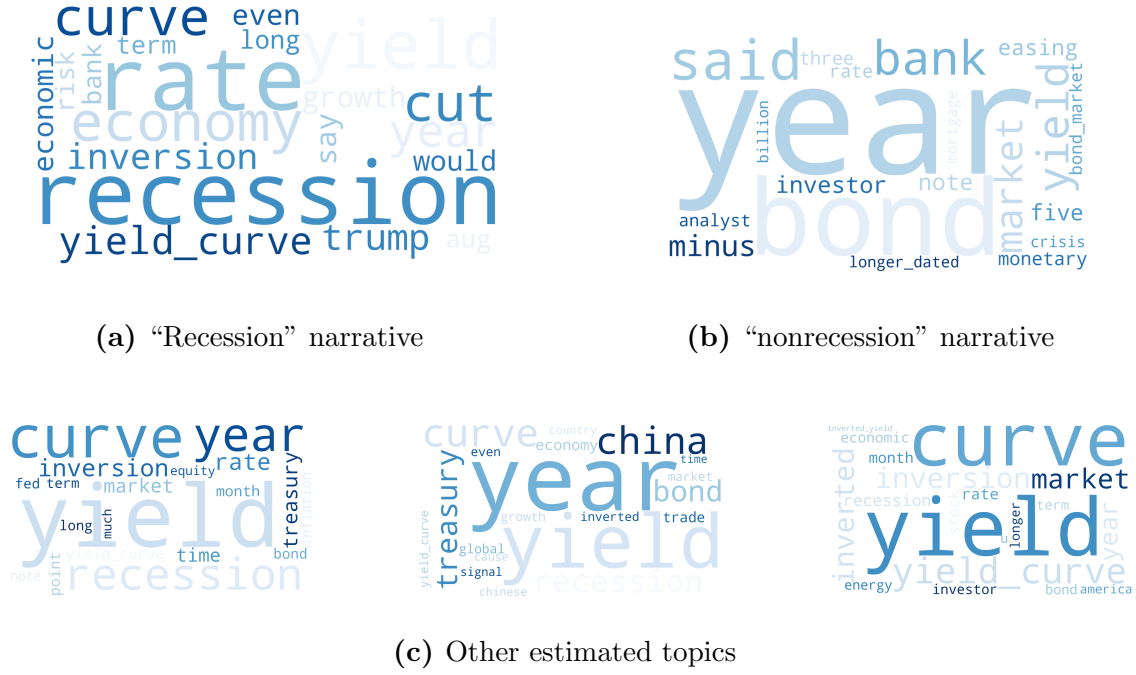
Since most news articles start with introducing the yield curve inversion as a recession predictor regardless of the narrative, the multi-membership feature of LDA allows for the word “recession” to appear in multiple topics, even when it is not the main thrust of the narrative. To match the theoretical definitions of narratives, we therefore consider the topic in Panel (a) as representing the recession narrative, because it has the highest probability of

---

<sup>13</sup>An alternative approach developed by Ash et al. (2021) uses semantic role labelling to capture the direction of causality in narratives.

<sup>14</sup>The pre-processing of texts includes removing stop words and numbers, lemmatizing, and representing the documents with a bigram model.

**Figure 4:** Economic narratives of the yield curve inversion: LDA outputs



*Notes:* This figure reports topics estimated with the LDA model on articles about the yield curve, with  $K = 5$  and symmetric Dirichlet priors. The size of a term represent the likelihood for it to appear in a topic. Raw values for this figure are reported in Appendix Table A.2.

the word “recession” appearing. The remaining three estimated topics are reported in Panel (c) for completeness.

The model performs well in capturing the narratives conveyed in news articles. To illustrate, for Peter Coy’s article discussed in Section 3 that argues the yield curve has lost its predictive power, the model estimates a loading of  $\theta(\text{nonrecession}) = 0.96$  on the nonrecession narrative and  $\theta(\text{recession}) = 0.01$  on the recession narrative. In contrast, for Cristina Alesci’s article emphasizing the recession risks, the model estimates  $\theta(\text{recession}) = 0.84$  and  $\theta(\text{nonrecession}) = 0.05$ . For the neutral coverage by Brian Chappata which introduces the yield curve, the model produces more balanced loadings of  $\theta(\text{recession}) = 0.67$  and  $\theta(\text{nonrecession}) = 0.11$ .

Based on these LDA outputs, we construct two measures of the narratives conveyed in an article. The first measure is  $\theta(d, k)$ , the estimated loading of article  $d$  on narrative  $k$ , where  $k$  is either the recession narrative or the nonrecession narrative. The second measure,  $\mathbb{1}(d, k)$ , is a binary measure to capture a narrative’s salience in an article relative to other

media coverage. We define  $\mathbb{1}(d, k) \equiv \mathbb{1}(\theta(d, k) > \frac{1}{D} \sum_{d \in D} \theta(d, k))$ , which takes the value 1 if the article loading exceeds the cross-sectional average loading of the narrative and 0 otherwise.

### 5.3. Empirical importance of narratives

We now use these measures to test whether different narratives of the yield curve inversion affect consumer sentiment. Our empirical model is a high-frequency event-time regression. For consumer  $i$  who has read news article  $d$ , the baseline model is:

$$\Delta s_{id} = \alpha + \beta_r \cdot \mathbb{1}(d, \text{recession}) + \beta_{nr} \cdot \mathbb{1}(d, \text{nonrecession}) + \varepsilon_{id}. \quad (10)$$

The dependent variable,  $\Delta s_{id}$ , is the change in a consumer's tweet sentiment 24 hours before and after the exposure to a narrative, where sentiment is measured with the naïve Bayes classifier described in Section 4. The exposure to a narrative is measured using quote retweeting activities on Twitter. We focus on the high-frequency changes in consumer sentiment 24 hours before and after the exposure to isolate the effect of the narrative. The timing is normalized so that the time when a consumer is exposed to a narrative is  $t = 0$ . Therefore, the time dimension of the baseline model in (10) is collapsed. The explanatory variables are narratives conveyed in an article. The binary variable  $\mathbb{1}(d, k)$  measures whether the loading of an article  $d$  on narrative  $k \in \{\text{recession}, \text{nonrecession}\}$  is above the cross-sectional mean. We also consider an alternative specification using the continuous measure of narratives  $\theta(d, k)$  (the loading of article  $d$  on narrative  $k$ ). The parameters of interest are  $\beta_r$  and  $\beta_{nr}$ , which estimate the effects of recession and nonrecession narratives on consumer sentiment, respectively.

Interpretation of the empirical model specified in (10) relies on four identifying assumptions. The first assumption is that the underlying event is exogenous. This is plausible because even though the Federal Reserve affects treasury yields through its open market operations, it does not control the exact timing of the yield curve inversion.

The second assumption is that sentiment changes are driven by narratives around the yield curve inversion and not by news about the inversion itself. Even though the yield curve has a good track record of predicting recessions, it has also given false-positive signals

(for example in 1966). The spread between long-term and short-term treasury yields is influenced by investors’ expectations of monetary policy and risk factors, along with other factors, and does not in theory predict a recession with certainty. The implication of the yield curve inversion—specifically whether it predicts an imminent recession—is therefore open for interpretation.

The third assumption is that a news subscription is uncorrelated with unobservable factors affecting the changes in sentiment within the high-frequency window. One obvious such unobservable factor that could potentially influence how consumers respond to the yield curve is political affiliation. To focus on consumers with similar political views, we restrict our analysis to centrist media outlets. Additionally, the high-frequency approach allows for isolating the effect of the exposure to a narrative. Note that because we study *changes* in sentiment, it is possible for Twitter users to have pre-existing differences in sentiment.

The last assumption concerns the direction of causality. The measure of exposure to a narrative that we use is retweeting. The implicit assumption is that retweeting implies the absorption of new information. However, consumers might selectively retweet articles that confirm their existing agenda. We conduct robustness checks below which impose a limit on the number of outlets that can appear in a user’s timeline to remove the most extreme users of this type from the sample.

Table 3 contains our main results from estimating variants of (10). Column 1 reports our baseline estimates of  $\beta_r$  and  $\beta_{nr}$ , displayed in basis points. Exposure to the recession narrative leads to a significantly more pessimistic outlook. A consumer who is exposed to an article emphasizing the recession narrative more strongly than the average article displays 1.3-basis-point more pessimistic sentiment in the 24 hours after absorbing the narrative. In contrast, the exposure to the nonrecession narrative leads to no significant changes in consumer sentiment. This is not surprising, since the nonrecession narrative downplays the scenario of a potential recession and conveys that there is no change in economic fundamentals. These results are robust to different measures of narratives, as reported in Column 2.

As a robustness check, we use univariate models to estimate the effect of each narrative individually. Columns 3 and 4 in Table 3 confirm the baseline results that the exposure to the recession narrative leads to a more pessimistic outlook. Both the economic and



**Table 3:** Effects of narratives on consumer sentiment

	(1)	(2)	(3)	(4)	(5)	(6)
	Consumer sentiment					
Recession narrative						
$\mathbb{1}(d, k)$	-1.29** (0.65)		-1.25** (0.62)			
$\theta(d, k)$		-1.74** (0.82)		-1.65** (0.80)		
Nonrecession narrative						
$\mathbb{1}(d, k)$	-0.11 (0.47)				0.15 (0.46)	
$\theta(d, k)$		-0.28 (0.64)				0.03 (0.63)
$R^2$	0.012	0.013	0.011	0.012	0.000	0.000
Observations	352	352	352	352	352	352

*Notes:* This table reports results from estimating variants of the baseline specification in (10). Column (1) reports  $\beta_r$  and  $\beta_{nr}$  from estimating the baseline specification

$$\Delta s_{id} = \alpha + \beta_r \cdot \mathbb{1}(d, \text{recession}) + \beta_{nr} \cdot \mathbb{1}(d, \text{nonrecession}) + \varepsilon_{id},$$

where  $\Delta s_{id}$  denotes changes in user  $i$ 's tweet sentiment 24 hours before and after reading article  $d$ ; and  $\mathbb{1}(d, k)$  for  $k \in \{\text{recession}, \text{nonrecession}\}$  denotes an indicator variable for whether the loading of article  $d$  on narrative  $k$  is above the cross-sectional mean. Tweet sentiment is measured with naïve Bayes classifier and an article's loading on a narrative is measured with the LDA model, as described in the main text. Column (2) reports  $\beta_r$  and  $\beta_{nr}$  from estimating  $\Delta s_{id} = \alpha + \beta_r \cdot \theta(d, \text{recession}) + \beta_{nr} \cdot \theta(d, \text{nonrecession}) + \varepsilon_{id}$ , where  $\theta(d, k)$  denotes the loading of article  $d$  on narrative  $k$ . Columns (3) through (6) report  $\beta$  from estimating univariate models  $\Delta s_{id} = \alpha + \beta \cdot x_{dk} + \varepsilon_{id}$ , where  $x_{dk}$  is  $\mathbb{1}(d, \text{recession})$ ,  $\theta(d, \text{recession})$ ,  $\mathbb{1}(d, \text{nonrecession})$ , or  $\theta(d, \text{nonrecession})$ . Standard errors are in parentheses. \* ( $p < 0.10$ ), \*\* ( $p < 0.05$ ), \*\*\* ( $p < 0.01$ ).

statistical significance are similar to those from the baseline estimates. Columns 5 and 6 also confirm that the nonrecession narrative is not associated with significant changes in consumer sentiment.

In Appendix Table A.3, we consider potential confounding economic factors by controlling for market conditions and macroeconomic uncertainty, measured by the S&P 500 Index and the VIX Index respectively. Our estimates are little changed, which suggests that the impact on sentiment is not driven by current economic conditions or uncertainty, but rather by the media's interpretations of the yield curve inversion.

**Focusing on susceptible consumers** In Shiller (2017)’s epidemiological model of narratives, the economy consists of three types of agents: susceptibles, infectives, and recovered. We now focus on measuring the effects of narratives on susceptible consumers, the households most likely to react to a new narrative. To do so, we limit our sample to users who retweet articles from a small number of news outlets only. The assumption here is that a Twitter user who is “infected” by a particular narrative will tend to retweet a large number of news outlets to promote their story. We rule out such users by restricting the maximum number of different news outlets to be 4, the mean number of outlets in the sample.


Appendix Table A.4 shows that, as in our main exercise, the recession narrative causes a decline in sentiment and the nonrecession narrative has no effect. However, the impact of recession narrative is about 50% stronger on susceptible users than on the general population. We can alternatively interpret the results in Table A.4 as a robustness check, ensuring that the effects are not driven by users who selectively retweet many articles with a particular narrative to promote that agenda, rather than processing the information contained in a narrative.

**Persistent effects of narratives** Finally, we study the persistence of the effects of narratives on sentiment. For each narrative  $k \in \{\text{recession}, \text{nonrecession}\}$  and horizon  $h$ , we estimate in the style of Jordà’s 2005 local projections

$$\Delta_h s_{id} = \alpha + \beta_{kh} \cdot \mathbb{1}(d, k) + \varepsilon_{idh}, \quad (11)$$

where  $\Delta_h s_{id}$  denotes the average change in consumer  $i$ ’s tweet sentiment between 1 day before and  $h$  days after the exposure to a narrative; and  $\mathbb{1}(d, k)$  denotes the binary measure of whether the loading of an article  $d$  on a narrative  $k$  is above the cross-sectional mean. As before, we collapse the time dimension by normalizing the time when a consumer is exposed to a narrative to be  $t = 0$ .

**Figure 5:** Dynamic effects of narratives



*Notes:* Panels (a) and (b) report  $\beta_{\text{recession},h}$  and  $\beta_{\text{nonrecession},h}$ , respectively, from estimating local projection in (11):  $\Delta_h s_{id} = \alpha + \beta_{kh} \cdot \mathbb{1}(d,k) + \varepsilon_{idh}$  for  $k \in \{\text{recession}, \text{nonrecession}\}$ , where  $\Delta_h s_{id}$  denotes the average change in consumer  $i$ 's tweet sentiment between 1 day before and  $h$  days after the exposure to a narrative; and  $\mathbb{1}(d,k)$  denotes an indicator variable of whether the loading of an article  $d$  on a narrative  $k$  is above the cross-sectional mean. We estimate (11) separately for each horizon  $h = 1, \dots, 30$ . Shaded areas represent 90% confidence intervals.

Figure 5 displays the results. Panel (a) shows that the negative effects of the recession narrative are persistent. In the month after reading the interpretation that the yield curve inversion signals an imminent recession, consumers become on average 15 basis points more pessimistic. Panel (b) shows that the exposure to the nonrecession narrative has no such effect.

## 6. Further Application: Inflation Narratives

The results above highlight the importance of narratives in shaping consumers' economic outlook. Indeed, while the yield curve inversion provides a laboratory to observe these effects, competing narratives are prevalent in the coverage of all economic news. Most prominent at present are the narratives around the current elevated levels of inflation (Andre et al., 2022b). In this section, we apply our empirical framework to study these inflation narratives. We document stylized facts about their evolution, and estimate that shifting narratives can have substantial effects on aggregate sentiment.

## 6.1. Data

Our textual sample consists of news articles on inflation, Twitter users who have quote retweeted such articles, and content of tweets from these users. We include all outlets from [Jurkowitz et al. \(2020\)](#) to study the broad implications of inflation narratives. For each outlet, we identify—from its official Twitter account—a list of base tweets that contain the keywords “PPI”, “CPI”, or “inflation”, and collect the corresponding news articles. We focus on US inflation and exclude news on non-US countries. These news articles form the corpus from which we capture inflation narratives. Table 4 lists the outlets included in our sample. In total, our sample consists of 28 news outlets, posting 5,128 base tweets on inflation, which links to 3,327 news articles. As a measure of a news outlet’s influence, we also obtain daily frequency counts of the mentioning of the outlet on Twitter (excluding self mentions). Our sample starts in 2014, when traffic on Twitter becomes active, and ends in 2021, which covers the onset of high inflation in the aftermath of the coronavirus pandemic.

For each base tweet on inflation, we identify Twitter users who have interacted with the base tweets through quote retweeting. We then collect tweets from these users’ timelines to measure changes in their sentiment. Appendix Tables A.5 and A.6 provide descriptive statistics for engagement activities with the inflation base tweets, and the Twitter activity of users who have engaged with those base tweets through quote retweeting. The median user is active, with 30 tweets in the 24-hour window around the exposure to inflation news.

Finally, we obtain macro series on Consumer Price Index (CPI) from FRED to study the state-dependent effects of inflation narratives. We also obtain survey data on the aggregate consumer sentiment from the University of Michigan Survey of Consumers to study the macroeconomic importance of narratives.

## 6.2. Inflation narratives

The theoretical framework in Section 2 can be adapted to the study of inflation narratives. In this case, the “extraneous variable”  $z$  is inflation, rather than a yield curve inversion. From Proposition 1, we therefore need to consider two competing narratives on inflation: one that suggests inflation is disconnected from household income, and another that suggests inflation is connect to real variables that affect households. These correspond to the “baseline” and “extraneous” narratives in Section 2. Motivated by this, we use the LDA model to capture

**Table 4:** Media outlets and coverage on inflation

Outlet	Twitter handle	# Base tweets	# Articles
Bloomberg	business	2647	1705
The Economist	TheEconomist	613	198
The Guardian	guardian	551	477
Wall Street Journal	WSJ	534	402
ABC News	abc	111	103
Washington Post	washingtonpost	84	69
CNN	CNN	73	47
Slate	slate	73	10
New York Times	nytimes	59	48
Breitbart	BreitbartNews	55	43
CBS News	CBSNews	53	34
USA Today	USAToday	40	26
Politico	politico	40	36
NBC News	NBCNews	34	23
NPR	NPR	31	24
Sean Hannity Show	seanhannity	24	13
New Yorker	NewYorker	18	7
Yahoo News	YahooNews	17	15
MSNBC	MSNBC	17	6
Al Jazeera America	AJEnglish	16	11
The Blaze	theblaze	8	7
Fox News	FoxNews	8	8
PBS	PBS	6	6
Huffington Post	HuffPost	5	4
Glenn Beck Program	glennbeck	4	0
Buzzfeed	BuzzFeed	4	3
BBC	BBCWorld	2	2
The Daily Show	TheDailyShow	1	0
All outlets		5128	3327


*Notes:* This table reports descriptive statistics of media outlets in our sample that have posted base tweets containing “PPI”, “CPI”, or “inflation” from 2014 to 2021.

these narratives from newspaper coverage, as we did for yield curve narratives in Section 5.<sup>15</sup>

Figure 6 presents the two narratives (topics) captured by the LDA as word clouds. The first topic in Panel (a) groups inflation with words such as “fed”, “central bank”, and “policy”, relating inflation to monetary policy. It also groups inflation with words such as “bond” and “interest rate”, consistent with the relationship under the Fisher equation.

<sup>15</sup>We specify the number of topics to be  $K = 2$  and Dirichlet priors to be symmetric in our estimation of the LDA.

**Figure 6:** Inflation narratives: LDA outputs



*Notes:* This table reports results from estimating the LDA model on articles about inflation, with  $K = 2$  and symmetric Dirichlet priors. The size of a term represents the likelihood for it to appear in a topic. Raw values for this figure are reported in Appendix Table A.7.

What is absent in this topic are variables that relate to household income. Therefore, we map it to the “baseline” narrative that suggests inflation is disconnected from real variables that affect household income. Since this narrative does not link inflation to real variables, we label this the “real-business-cycle” (RBC) narrative. In contrast, the second topic in Panel (b) groups inflation with “wage”, “cost”, and “consumer”, which suggest links to household income and the real economy. Since this connects inflation with real variables that affect household, we label this the “New Keynesian” (NK) narrative.


Having extracted these narratives, we first present the prevalence of these inflation narratives over time. We measure the prevalence of a narrative as the weighted sum of articles’ loadings on the narrative, weighted by news outlets’ influence on Twitter. Denoting news outlets with  $j \in \mathcal{J}$ , and articles posted by each outlet with  $d \in \mathcal{D}_j$ , we define the prevalence of a narrative  $k$  as

$$v_t(k) = \sum_{j \in \mathcal{J}} \sum_{d \in \mathcal{D}_j} \omega_{jt} \theta(d, k), \quad \text{where } \omega_{jt} = \frac{N_{jt}}{\sum_{j \in \mathcal{J}} N_{jt}}. \quad (12)$$

Here,  $\theta(d, k)$  is the LDA loading of an article  $d$  on a narrative  $k$ .  $N_{jt}$  is the number of times a news outlet  $j$  is mentioned on Twitter in day  $t$  (excluding the times an outlet mentions itself). Therefore,  $\omega_{jt}$  measures the influence of a news outlet.

Figure 7 reports the prevalence of each inflation narrative over our sample period, normalized so that the maximum value corresponds to 100. The coverage on both narratives is

**Figure 7:** Prevalence of inflation narratives



*Notes:* This figure reports the prevalence of inflation narratives defined in (12). For news outlets  $j \in \mathcal{J}$  and articles posted by each outlet  $d \in \mathcal{D}_j$ , the prevalence of a narrative  $k$  is defined as  $v_t(k) = \sum_{j \in \mathcal{J}} \sum_{d \in \mathcal{D}_j} \omega_{jt} \theta(d, k)$ , where  $\theta(d, k)$  is the LDA loading of an article  $d$  on a narrative  $k$ , and  $\omega_{jt} = \frac{N_{jt}}{\sum_{j \in \mathcal{J}} N_{jt}}$  is the frequency count of outlet  $j$  on Twitter (excluding self mentions) as a fraction of the frequency counts of all sample outlets. Prevalence has been scaled so the highest value is 100.

similar and minimal for most of the sample period, when inflation is low and stable. The coverage on both inflation narratives spikes during 2021, when realized inflation rises.

The two narratives do not, however, spread in equal amounts. The prevalence of the NK narrative increases dramatically relative to that of the RBC narrative in 2021, showing a sign of becoming “viral”. As inflation rises to a historically high level, media outlets shift their attention to cover inflation news. In doing so, they also shift the narrative that they use to discuss inflation, towards one that disproportionately emphasizes that inflation is an economic phenomenon with real consequences.

### 6.3. State-dependent effects of inflation narratives

Using these empirical measures of inflation narratives, we now estimate their effects on consumer sentiment. Importantly, unlike with the yield curve inversion studied above, inflation is not a single discrete event. A narrative linking inflation with household income may induce very different responses of expectations when inflation is high than when it is low. Indeed, recent empirical evidence suggests households perceive substantially more negative consequences of inflation for real variables when they perceive that realized inflation is high

(Drager, Lamla and Pfajfar, 2020; Macaulay, 2022). Formally, the inflation narratives we study generate conditional expectations *functions* (i.e.  $E_{it}^k(y_{t+1}|\mathcal{I}_t)$  in Equations (7) and (8)). These functions take information on inflation as an input to expectations, and so the effects of a change in narrative will vary depending on that information.<sup>16</sup>

To account for this, we allow our estimates of the effect of narratives on sentiment to vary with the realized level of inflation at the time the user engages with the narrative. For user  $i$  who quote retweets article  $d$  containing narrative  $k$  at time  $t$ , we estimate the state-dependent effect of a narrative  $k \in \{\text{RBC}, \text{NK}\}$ , depending on the realized level of inflation.

Our empirical model takes the form

$$\Delta s_{itd} = \alpha_k + \beta_{kc} \cdot \mathbb{1}(\pi \geq c) + \gamma_{kc} \cdot \underbrace{\mathbb{1}(\pi \geq c)}_{\text{inflation}} \times \underbrace{\mathbb{1}(d, k)}_{\text{narrative}} + \varepsilon_{itd}, \quad (13)$$

where  $\Delta s_{itd}$  is the change in a Twitter user’s textual sentiment 24 hours before and after quote retweeting the inflation base tweet;  $\alpha_k$  is a constant;  $\mathbb{1}(\pi \geq c)$  is an indicator variable which takes the value 1 if the annualized CPI inflation is greater or equal than  $c\%$ ;  $\mathbb{1}(d, k)$  is our binary measure of narratives, which takes the value 1 if the loading of an article on the narrative is above the cross-sectional mean; and  $\varepsilon_{itd}$  is a random error. The coefficient of interest is  $\gamma_{kc}$ , which measures the impact of narrative  $k$  on sentiment changes for a given level of inflation. We estimate (13) separately for each integer level of annualized CPI inflation  $\pi \geq -2\%, \dots, 9\%$ .


Figure 8 reports the results estimating the effects of inflation narratives. We plot with blue bars the estimated effects for the RBC narrative and red bars those for the NK narratives, with whiskers representing standard errors. When annualized inflation is below the Fed’s targeted 2%, neither narrative has a significant effect on consumer sentiment. This is consistent with evidence on the cyclical nature of macroeconomic attention, and in particular low levels of attention when inflation is low and stable (Pfäuti, 2022; Song and Stern, 2022). However, when inflation rises above 2%, the two narratives have significant and diverging effects on consumer sentiment. The RBC narrative, which informs consumers that nominal

---

<sup>16</sup>In the language of Macaulay (2022), our narratives specify only the subjective model component of expectations, so the effects of changing subjective model will depend on the information component of expectations at which the change occurs.



**Figure 8:** Effects of inflation narratives



Notes: This figure reports  $\gamma_k$  for  $k \in \{\text{RBC}, \text{NK}\}$  from estimating the specification in (13)

$$\Delta s_{itd} = \alpha_k + \beta_{kx} \cdot \mathbb{1}(\pi \geq c) + \gamma_{kx} \cdot \mathbb{1}(\pi \geq c) \times \mathbb{1}(d, k) + \varepsilon_{itd},$$


where  $\Delta s_{itd}$  is the change in a Twitter user  $i$ 's textual sentiment 24 hours before and after reading an article  $d$ ;  $\mathbb{1}(\pi \geq c)$  is an indicator variable of whether the annualized CPI inflation is greater or equal than  $c\%$ ; and  $\mathbb{1}(d, k)$  is an indicator variable if the LDA loading of an article  $d$  on the narrative  $k$  is above the cross-sectional mean. Points estimates for the “real-business-cycle” and “New Keynesian” narratives are represented by blue and red bars, respectively. Whiskers represent standard errors.

inflation is disconnected from their income, raises the sentiment of Twitter users who are exposed to it.

In contrast, after being exposed to the NK narrative that inflation affects their income, Twitter users display a more pessimistic outlook. The negative effects of the NK narrative are increasing with the realized levels of inflation. When the annualized inflation is greater than or equal to 7%, exposure to a NK narrative lowers consumer sentiment by 40 basis points.<sup>17</sup>

<sup>17</sup>Note that the effects of the two narratives are symmetric because the loadings of topics in each article add up to 1. The yield curve results are not symmetric because we extracted more than two topics for each article.

**Figure 9:** Changes in aggregate sentiment in the Michigan Survey of Consumers



*Notes:* This figure displays changes in aggregate sentiment in the University of Michigan Survey of Consumers. Panel (a) displays the Consumer Sentiment Index for our sample period from 2014 to 2021. The 1966 index value is normalized to be 100. Panel (b) displays the fraction of respondents that report higher prices as reasons for worse personal finances. Raw monthly data is reported in solid blue lines, and 3-month moving averages are reported in dashed grey lines.

#### 6.4. Macroeconomic effects of inflation narratives

While these results are not to be interpreted as causal<sup>18</sup>, they provide suggestive evidence on the effects of inflation narratives. To gauge the potential macroeconomic importance of these narratives, we perform a simple back-of-the-envelope calculation using our empirical results.

Figure 9a shows that as inflation rose sharply in late 2021, the Index of Consumer Sentiment in the University of Michigan Surveys of Consumers declined rapidly, despite many other indicators suggesting the US was not close to recession (Sahm, 2022). At the same time, the balance of narratives around inflation in US media outlets shifted towards “New Keynesian” narratives, in which inflation can have damaging effects on the real economy. This shift in narratives coincided with a steep rise in the percent of survey respondent listing higher prices as reasons for worse personal finances, as shown in Figure 9b.

Table 5 summarizes results of this back-of-the-envelope calculation. During the second half of 2021, the Consumer Sentiment Index declined by 19%. Taking the estimated state-dependent effects of each narrative in Figure 8, and weighting by the daily prevalence of NK

<sup>18</sup>Other factors related to inflation may simultaneously influence consumer sentiment. See, for example, the factors surveyed in D’Acunto, Malmendier and Weber (2022).

**Table 5:** Macroeconomic effects of narratives

	% decline Jun vs. Dec 2021	As % of CSI decline
Consumer Sentiment Index (CSI)	19%	-
Effects of narratives	13%	68%
from state-dependent effects of narratives	5%	26%
from shift towards NK narratives	8%	42%

*Notes:* This table reports the percentage change in the Consumer Sentiment Index (CSI) in the University of Michigan Survey of Consumers from June 2021 to December 2021. We compute the effects of narratives on sentiment as  $\sum_t \sum_k \tilde{v}_{tk} \beta_{kc}$  where  $\tilde{v}_{tk} = v_{tk} / \sum_k v_{tk}$  is the relative prevalence of narrative  $k \in \text{RBC, NK}$  based on the daily prevalence measures  $v_{tk}$  in (12); and  $\beta_{kc}$  is the estimated state-dependent effects of narrative  $k$  from (13), with  $c$  denoting the integer floor of CPI inflation in month  $t$ . We calculate the effects from the state-dependent effects of narratives as  $\sum_t \sum_k \tilde{v}_{t-12,k} \beta_{kc}$ , that is by replacing the actual relative prevalence of each narrative in 2021H2 with the proportions from the same periods in 2020. The remaining effects,  $\sum_t \sum_k \tilde{v}_{tk} \beta_{kc} - \tilde{v}_{t-12,k} \beta_{kc}$ , are attributed to the shift towards NK narratives.

and RBC narratives in Figure 7, we estimate that the effects of narratives led to 13% decline in consumer sentiment, accounting for 68% of the decline in aggregate sentiment.<sup>19</sup>

The shift in the media towards New-Keynesian narratives played an important role in the decline of consumer sentiment. To see this, we conduct a counterfactual analysis in which we replace the true relative prevalence of each narrative with the proportions from the same periods in 2020, before the New-Keynesian narratives went viral. In this counterfactual, consumer sentiment only falls by 5%, substantially less than what was observed in the data. This indicates that the worsening of sentiment arises mostly from the rising prevalence of NK narrative. The shift in narratives in news media towards one that emphasizes the real damage of inflation therefore explains 42% of the observed decline in the Consumer Sentiment Index.

<sup>19</sup>Our calculation assumes that the effects of narratives on Twitter users extend to the general population. Perrin and Anderson (2019) show that 22% of US adults use Twitter. Of these, Wojcik and Hughes (2019) document that while Twitter users are representative of US adults in terms of gender and ethnicity, they are younger, more likely to identify as Democrats, more highly educated, and have higher income than US adults overall.

## 7. Conclusion

Narratives are increasingly seen as an important factor in how economic agents form their expectations, by both academics (Shiller, 2017, 2020) and policymakers (Schnabel, 2020).<sup>20</sup> We provide evidence that exposure to particular narratives in the media does indeed have significant effects on consumer sentiment.

Formalizing narratives as directed acyclic graphs, we show that certain groups of narratives will in fact have exactly the same effect on expectations. In the context of the inversion of the U.S. yield curve in 2019, the distinguishing feature between a “recession” narrative and a “nonrecession narrative” is, therefore, whether there is a link connecting the inverted yield curve with an upcoming recession.

Standard tools from topic modeling in natural language processing are well suited to making this distinction. We do this in a large corpus of articles from traditional news media, which is a key source of macroeconomic narratives (Andre et al., 2022b). Linking these articles with rich data on Twitter activity, we find that engaging with an article advancing a “recession” narrative causes a significant and persistent decline in the sentiment of that Twitter user, as embodied in their other activity on the social media site at the time. In contrast, engaging with a “nonrecession” narrative has no such effect on sentiment. This is precisely what would be predicted by models in which viral narratives affect aggregate behaviour by shifting expectations. It also suggests a powerful role for the media in influencing aggregate sentiment, as in e.g. Nimark (2014).

We confirm this aggregate implication by further applying our methodology to narratives around inflation. As inflation took off in the U.S. in 2021, media discussion of inflation also grew rapidly. The narrative that we find went viral in this period was one that explicitly linked inflation to the real economy, emphasizing its effects beyond nominal financial variables. Engaging with this narrative on Twitter during periods of high inflation is associated with substantial declines in sentiment, consistent with the sharp declines in other measures of consumer sentiment at the time. Indeed, a simple calculation suggests that the shift towards this “New Keynesian” narrative of inflation in the media accounts for 42% of

---

<sup>20</sup>See, for example, the speech by Isabel Schnabel, Member of the Executive Board of the ECB, at the Karlsruhe Law Studies Society entitled “*Narratives about the ECB’s monetary policy – reality or fiction?*” (Schnabel, 2020).

the decline in sentiment in the second half of 2021.

Our approach using tools from natural language processing to extract relevant groups of narratives from text can be used in other settings. For example, while news media is an important source of narratives, similar techniques can be used to study economic narratives created by policymakers in monetary and fiscal policy statements and by firm managers in earnings reports. These data sources are naturally occurring, which means that our method can be deployed to track the evolution of narratives and their ongoing effects—potentially providing a useful input to discussions of macroeconomic policy.

## References

- Acharya, Sushant, Jess Benhabib, and Zhen Huo, “The anatomy of sentiment-driven fluctuations,” *Journal of Economic Theory*, 2021, *195*, 105280.
- Andrade, Philippe, Gaetano Gaballo, Eric Mengus, and Benoît Mojon, “Forward guidance and heterogeneous beliefs,” *American Economic Journal: Macroeconomics*, 2019, *11* (3), 1–29.
- Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart, “Subjective models of the macroeconomy: Evidence from experts and a representative sample,” *Review of Economic Studies*, 2022.
- , Ingar Haaland, Christopher Roth, and Johannes Wohlfart, “Narratives about the macroeconomy,” Technical Report, University of Bonn and University of Cologne, Germany 2022.
- Angeletos, George-Marios and Jennifer La’O, “Sentiments,” *Econometrica*, 2013, *81* (2), 739–779.
- Ash, Elliott, Germain Gauthier, and Philine Widmer, “Text semantics capture political and economic narratives,” *arXiv preprint arXiv:2108.01720*, 2021.
- Atmaz, Adem and Suleyman Basak, “Belief dispersion in the stock market,” *Journal of Finance*, 2018, *73* (3), 1225–1279.
- Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel, “The economic effects of social networks: Evidence from the housing market,” *Journal of Political Economy*, 2018, *126* (6), 2224–2276.
- Bianchi, Francesco, Howard Kung, and Roberto Gomez Cram, “Using social media to identify the effects of congressional partisanship on asset prices,” Technical Report, National Bureau of Economic Research 2021.
- Blei, David M, Andrew Y Ng, and Michael I Jordan, “Latent Dirichlet allocation,” *Journal of Machine Learning Research*, 2003, *3*, 993–1022.

- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Diagnostic expectations and credit cycles,” *Journal of Finance*, 2018, *73* (1), 199–227.
- Buehlmaier, Matthias MM and Toni M Whited**, “Are financial constraints priced? Evidence from textual analysis,” *Review of Financial Studies*, 2018, *31* (7), 2693–2728.
- Bybee, Leland, Bryan T Kelly, Asaf Manela, and Dacheng Xiu**, “The structure of economic news,” Technical Report, National Bureau of Economic Research 2020.
- Calomiris, Charles W and Harry Mamaysky**, “How news and its context drive risk and returns around the world,” *Journal of Financial Economics*, 2019, *133* (2), 299–336.
- Caplin, Andrew and John V Leahy**, “Wishful thinking,” Technical Report, National Bureau of Economic Research 2019.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia**, “Inflation expectations, learning, and supermarket prices: Evidence from survey experiments,” *American Economic Journal: Macroeconomics*, 2017, *9* (3), 1–35.
- Chahrour, Ryan, Kristoffer Nimark, and Stefan Pitschner**, “Sectoral media focus and aggregate fluctuations,” *American Economic Review*, 2021, *111* (12), 3872–3922.
- Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar**, “The formation of expectations, inflation, and the phillips curve,” *Journal of Economic Literature*, 2018, *56* (4), 1447–91.
- Dräger, Lena, Michael J. Lamla, and Damjan Pfajfar**, “Are survey expectations theory-consistent? The role of central bank communication and news,” *European Economic Review*, 2016, *85*, 84–111.
- Drager, Lena, Michael J Lamla, and Damjan Pfajfar**, “The hidden heterogeneity of inflation expectations and its implications,” 2020.
- D’Acunto, Francesco, Ulrike Malmendier, and Michael Weber**, “What do the data tell us about inflation expectations?,” Technical Report, National Bureau of Economic Research 2022.

- Ehrmann, Michael and Alena Wabitsch**, “Central bank communication with non-experts—A road to nowhere?,” *Journal of Monetary Economics*, 2022, *127*, 69–85.
- Eliaz, Kfir and Ran Spiegler**, “A model of competing narratives,” *American Economic Review*, 2020, *110* (12), 3786–3816.
- Friedman, Jerome, Trevor Hastie, Robert Tibshirani et al.**, *The elements of statistical learning*, Vol. 1, Springer series in statistics New York, 2001.
- Gentzkow, Matthew and Jesse M Shapiro**, “Competition and truth in the market for news,” *Journal of Economic perspectives*, 2008, *22* (2), 133–154.
- , —, and **Michael Sinkinson**, “Competition and ideological diversity: Historical evidence from us newspapers,” *American Economic Review*, 2014, *104* (10), 3073–3114.
- Go, Alec, Richa Bhayani, and Lei Huang**, “Twitter sentiment classification using distant supervision,” *CS224N project report, Stanford*, 2009, *1* (12), 2009.
- Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera**, “Social media, sentiment and public opinions: Evidence from# Brexit and# USElection,” *European Economic Review*, 2021, p. 103772.
- Greenwood, Robin and Andrei Shleifer**, “Expectations of returns and expected returns,” *Review of Financial Studies*, 2014, *27* (3), 714–746.
- Haldane, Andrew, Alistair Macaulay, and Michael McMahon**, “The three E’s of central-bank communication with the public,” 2021.
- Hansen, Stephen, Michael McMahon, and Andrea Prat**, “Transparency and deliberation within the FOMC: a computational linguistics approach,” *Quarterly Journal of Economics*, 2018, *133* (2), 801–870.
- Harris, Milton and Artur Raviv**, “Differences of opinion make a horse race,” *Review of Financial Studies*, 1993, *6* (3), 473–506.
- Harvey, Campbell R.**, “The real term structure and consumption growth,” *Journal of Financial Economics*, 1988, *22* (2), 305–333.



- Hoffman, Matthew, Francis Bach, and David Blei**, “Online learning for latent Dirichlet allocation,” *Advances in Neural Information Processing Systems*, 2010, *23*, 856–864.
- Jordà, Òscar**, “Estimation and inference of impulse responses by local projections,” *American Economic Review*, 2005, *95* (1), 161–182.
- Jurkowitz, Mark, Amy Mitchell, Elisa Shearer, and Mason Walker**, “US media polarization and the 2020 election: A nation divided,” *Pew Research Center*, 2020, *24*.
- Krishnamurthy, Arvind and Wenhao Li**, “Dissecting mechanisms of financial crises: Intermediation and sentiment,” Technical Report, National Bureau of Economic Research 2020.
- Larsen, Vegard and Leif Anders Thorsrud**, “Business cycle narratives,” 2019.
- Levchenko, Andrei A and Nitya Pandalai-Nayar**, “TFP, news, and “sentiments”: The international transmission of business cycles,” *Journal of the European Economic Association*, 2020, *18* (1), 302–341.
- Levy, Ro’ee**, “Social media, news consumption, and polarization: Evidence from a field experiment,” *American Economic Review*, 2021, *111* (3), 831–70.
- Macaulay, Alistair**, “Heterogeneous information, subjective model beliefs, and the time-varying transmission of shocks,” 2022.
- Malmendier, Ulrike and Stefan Nagel**, “Learning from inflation experiences,” *Quarterly Journal of Economics*, 2016, *131* (1), 53–87.
- Mankiw, N Gregory and Ricardo Reis**, “Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve,” *Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.
- Matveev, Dmitry and Francisco Javier Ruge-Murcia**, “Tariffs and the exchange rate: Evidence from Twitter,” Technical Report, Bank of Canada Staff Working Paper 2021.
- Maxted, Peter**, “A macro-finance model with sentiment,” *Unpublished working paper*, 2019.


- Molavi, Pooya**, “Macroeconomics with learning and misspecification : A general theory and applications,” 2019.
- Nimark, Kristoffer P**, “Man-bites-dog business cycles,” *American Economic Review*, 2014, 104 (8), 2320–67.
- Nyman, Rickard, Sujit Kapadia, and David Tuckett**, “News and narratives in financial systems: Exploiting big data for systemic risk assessment,” *Journal of Economic Dynamics and Control*, 2021, 127, 104119.
- Patton, Andrew J. and Allan Timmermann**, “Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion,” *Journal of Monetary Economics*, 2010, 57 (7), 803–820.
- Perrin, Andrew and Monica Anderson**, “Share of US adults using social media, including Facebook, is mostly unchanged since 2018,” 2019.
- Pfäuti, Oliver**, “Inflation - who cares? Monetary policy in times of low attention,” Technical Report 2022.
- Ramos, Juan et al.**, “Using tf-idf to determine word relevance in document queries,” in “Proceedings of the first instructional conference on machine learning,” Vol. 242 Citeseer 2003, pp. 29–48.
- Rish, Irina et al.**, “An empirical study of the naive Bayes classifier,” in “IJCAI 2001 workshop on empirical methods in artificial intelligence,” Vol. 3 2001, pp. 41–46.
- Sahm, Claudia**, “We are not in a recession, nor is one inevitable,” *Stay-At-Home Macro Blog*, June 2022.
- Schnabel, Isabel**, “Narratives about the ECB’s monetary policy—reality or fiction?,” *Juristische Studiengesellschaft*, 2020.
- Shiller, Robert J**, “Narrative economics,” *American Economic Review*, 2017, 107 (42), 967–1004.
- , *Narrative economics: How stories go viral and drive major economic events*, Princeton University Press, 2020.

- Sims, Christopher A**, “Implications of rational inattention,” *Journal of Monetary Economics*, 2003, *50* (3), 665–690.
- Song, Wenting and Samuel Stern**, “Firm inattention and the efficacy of monetary policy: A text-based approach,” Technical Report, Bank of Canada 2022.
- Spiegler, Ran**, “Bayesian networks and boundedly rational expectations,” *Quarterly Journal of Economics*, 2016, *131* (3), 1243–1290.
- , “Behavioral implications of causal misperceptions,” *Annual Review of Economics*, 2020, *12*, 81–106.
- , “Can agents with causal misperceptions be systematically fooled?,” *Journal of the European Economic Association*, 2020, *18*, 583–617.
- Verma, Thomas and Judea Pearl**, “Causal networks: Semantics and expressiveness,” in “Machine Intelligence and Pattern Recognition,” Vol. 9, Elsevier, 1990, pp. 69–76.
- Wojcik, Stefan and Adam Hughes**, “Sizing up Twitter users,” *PEW Research Center*, 2019, *24*.
- Xiong, Wei and Hongjun Yan**, “Heterogeneous expectations and bond markets,” *Review of Financial Studies*, 2010, *23* (4), 1433–1466.

## APPENDICES

### A. Additional Tables and Figures

**Figure A.1:** Yield curve inversion and recessions in the US



*Notes:* Yield curve and recessions in the US for 1976–2019. The blue solid line displays the spread between 10-year treasury yield and 2-year treasury yield (“10Y2Y”). Recession dates as classified by NBER are shaded in grey.

**Table A.1:** Top positive and negative scores: tweets on yield curve**Panel (a):** Top negative tweets (most negative first)

	<b>Tweet</b>	<b>Score</b>	<b>Sentiment</b>
1	@USER @USER @USER Real recessions have real inverted yield curves. That really invert and stay there. Then the real Recession starts. Probably July, 2020 just in time for the election. Isn't that what the Deep State wants? But they'll blame it on "don't cry for me Argentina!"	0.211	negative
2	@USER: IT DIDN'T WORK: Despite the Fed, the yield curve is stuck in 'recession' mode, stocks are a mess, and manufacturing is ...	0.218	negative
3	@USER: Global mkts in bad mood after hawkish Fed cut. Stocks fell, yield curve flattened worryingly & dollar strengthened as ...	0.218	negative
4	@USER: It doesn't always mean a recession's coming, but you don't get a recession without an inverted yield curve. Therein lies the worr ...	0.225	negative
5	@USER: Economics can't be spun. An inverted yield curve is the sign of a sick economy. Period... Trump had tried to spin the ...	0.233	negative

**Panel (b):** Top positive tweets (most positive first)

	<b>Tweet</b>	<b>Score</b>	<b>Sentiment</b>
1	@USER: Nice article and agree 100%... the market is treating the "yield curve" inversion like the Ebola virus for stocks... REAL M...	0.677	positive
2	Japanese yen stands tall as US yield curve inversion stokes economic worries HTTPURL via @USER HTTPURL	0.668	positive
3	@USER: A simple graph does a better job of predicting recessions than the experts. @USER remind us why the yield curve matters ...	0.655	positive
4	@USER: U.S. yield curve flattens on supply, trade worries HTTPURL HTTPURL	0.651	positive
5	White House trade advisor Navarro: 'Technically we did not have a yield curve inversion' HTTPURL via @USER HTTPURL	0.634	positive

*Notes:* This table reports the top 5 positive and negative tweets about the yield curve classified by the naïve Bayes model described in Appendix Section C. User names and URLs have been anonymized to tokens "@USER" and "HTTPURL", respectively.

**Table A.2:** Topics estimated with LDA: yield curve inversion

Topic 1 “Recession”		Topic 2 “Nonrecession”	
Term	Probability	Term	Probability
<i>recession</i>	0.016	year	0.052
rate	0.016	bond	0.048
yield	0.011	said	0.036
economy	0.011	bank	0.025
cut	0.010	yield	0.021
curve	0.010	market	0.016
year	0.009	minus	0.015
yield curve	0.009	investor	0.015
trump	0.008	note	0.014
inversion	0.008	five	0.013
growth	0.008	easing	0.013
say	0.008	monetary	0.012
economic	0.008	three	0.011
even	0.008	rate	0.011
would	0.008	bond market	0.010
bank	0.006	analyst	0.010
risk	0.006	longer dated	0.010
long	0.006	mortgage	0.010
aug	0.006	crisis	0.009
term	0.006	billion	0.009

Topic 3		Topic 4		Topic 5	
Term	Probability	Term	Probability	Term	Probability
yield	0.040	yield	0.024	year	0.025
curve	0.036	curve	0.021	yield	0.023
yield curve	0.026	year	0.016	curve	0.016
inversion	0.016	<i>recession</i>	0.014	china	0.015
inverted	0.016	inversion	0.013	<i>recession</i>	0.014
market	0.015	rate	0.013	treasury	0.012
year	0.013	treasury	0.009	bond	0.012
<i>recession</i>	0.012	market	0.008	economy	0.011
rate	0.010	time	0.008	trade	0.010
stock	0.010	yield curve	0.008	global	0.008
month	0.010	point	0.008	growth	0.008
economic	0.009	month	0.008	market	0.008
term	0.008	bond	0.007	even	0.008
investor	0.008	fed	0.007	inverted	0.007
bond	0.008	long	0.007	signal	0.007
energy	0.008	term	0.007	yield curve	0.007
u	0.007	inflation	0.006	time	0.007
longer	0.007	note	0.006	country	0.006
america	0.007	much	0.006	chinese	0.006
inverted yield	0.007	equity	0.006	cause	0.006

*Notes:* This table reports topics estimated with the LDA on articles of the yield curve with  $K = 5$  and symmetric Dirichlet priors. For each topic, we report the distribution over vocabulary terms estimated with the LDA model.

**Table A.3:** Controlling for macroeconomic conditions

	(1)	(2)	(3)	(4)	(5)	(6)
	Consumer sentiment					
Recession narrative						
$\mathbb{1}(d, k)$	-1.13*		-1.26**			
	(0.65)		(0.63)			
$\theta(d, k)$		-1.63*		-1.62**		
		(0.83)		(0.80)		
Nonrecession narrative						
$\mathbb{1}(d, k)$	0.47				0.74	
	(0.60)				(0.58)	
$\theta(d, k)$		-0.01				0.32
		(0.67)				(0.65)
$R^2$	0.020	0.019	0.019	0.019	0.012	0.008
Observations	352	352	352	352	352	352
Macro controls	yes	yes	yes	yes	yes	yes

*Notes:* This table reports results from estimating variants of the baseline specification in (10) while controlling for macroeconomic and financial fluctuations. Column (1) reports  $\beta_r$  and  $\beta_{nr}$  from estimating the baseline specification

$$\Delta s_{id} = \alpha + \beta_r \cdot \mathbb{1}(d, \text{recession}) + \beta_{nr} \cdot \mathbb{1}(d, \text{nonrecession}) + \Gamma' Z_t + \varepsilon_{id},$$

where  $\Delta s_{id}$  denotes changes in user  $i$ 's tweet sentiment 24 hours before and after reading article  $d$ ; and  $\mathbb{1}(d, k)$  for  $k \in \{\text{recession}, \text{nonrecession}\}$  denotes an indicator variable for whether the loading of article  $d$  on narrative  $k$  is above the cross-sectional mean;  $Z_t$  is a vector of macro and financial controls including the S&P 500 and VIX indices. Tweet sentiment is measured with naïve Bayes classifier and an article's loading on a narrative is measured with the LDA model, as described in the main text. Column (2) reports  $\beta_r$  and  $\beta_{nr}$  from estimating  $\Delta s_{id} = \alpha + \beta_r \cdot \theta(d, \text{recession}) + \beta_{nr} \cdot \theta(d, \text{nonrecession}) + \Gamma' Z_t + \varepsilon_{id}$ , where  $\theta(d, k)$  denotes the loading of article  $d$  on narrative  $k$ . Columns (3) through (6) report  $\beta$  from estimating univariate models  $\Delta s_{id} = \alpha + \beta \cdot x_{dk} + \Gamma' Z_t + \varepsilon_{id}$ , where  $x_{dk}$  is  $\mathbb{1}(d, \text{recession})$ ,  $\theta(d, \text{recession})$ ,  $\mathbb{1}(d, \text{nonrecession})$ , or  $\theta(d, \text{nonrecession})$ . Standard errors are in parentheses. \* ( $p < 0.10$ ), \*\* ( $p < 0.05$ ), \*\*\* ( $p < 0.01$ ).

**Table A.4:** Limiting the number of outlets in user timelines

	(1)	(2)	(3)	(4)	(5)	(6)
	Consumer sentiment					
Recession narrative						
$\mathbb{1}(d, k)$	-1.74*		-1.74*			
	(0.99)		(0.96)			
$\theta(d, k)$		-2.34*		-2.23*		
		(1.26)		(1.23)		
Nonrecession narrative						
$\mathbb{1}(d, k)$	-0.01				0.29	
	(0.69)				(0.67)	
$\theta(d, k)$		-0.34				0.04
		(0.91)				(0.89)
$R^2$	0.014	0.015	0.014	0.014	0.001	0.000
Observations	227	227	227	227	227	227

*Notes:* This table reports results from estimating variants of the baseline specification in (10), restricting the sample to users whose Twitter timelines contain no more than 4 different news outlets in the 2-month window around their quote retweets. Column (1) reports  $\beta_r$  and  $\beta_{nr}$  from estimating the baseline specification

$$\Delta s_{id} = \alpha + \beta_r \cdot \mathbb{1}(d, \text{recession}) + \beta_{nr} \cdot \mathbb{1}(d, \text{nonrecession}) + \varepsilon_{id},$$

where  $\Delta s_{id}$  denotes changes in user  $i$ 's tweet sentiment 24 hours before and after reading article  $d$ ; and  $\mathbb{1}(d, k)$  for  $k \in \{\text{recession}, \text{nonrecession}\}$  denotes an indicator variable for whether the loading of article  $d$  on narrative  $k$  is above the cross-sectional mean. Tweet sentiment is measured with naïve Bayes classifier and an article's loading on a narrative is measured with the LDA model, as described in the main text. Column (2) reports  $\beta_r$  and  $\beta_{nr}$  from estimating  $\Delta s_{id} = \alpha + \beta_r \cdot \theta(d, \text{recession}) + \beta_{nr} \cdot \theta(d, \text{nonrecession}) + \varepsilon_{id}$ , where  $\theta(d, k)$  denotes the loading of article  $d$  on narrative  $k$ . Columns (3) through (6) report  $\beta$  from estimating univariate models  $\Delta s_{id} = \alpha + \beta \cdot x_{dk} + \varepsilon_{id}$ , where  $x_{dk}$  is  $\mathbb{1}(d, \text{recession})$ ,  $\theta(d, \text{recession})$ ,  $\mathbb{1}(d, \text{nonrecession})$ , or  $\theta(d, \text{nonrecession})$ . Standard errors are in parentheses. \* ( $p < 0.10$ ), \*\* ( $p < 0.05$ ), \*\*\* ( $p < 0.01$ ).



**Table A.5:** Descriptive statistics on inflation base tweets

(a) All containing “CPI”, “PPI”, “inflation”

	Mean	SD	5th Pctl	Median	95th Pctl	Obs
Quote retweet count	4.2	11.6	0	2	14	7741
Retweet count	28.4	56.3	4	16	85	7741
Reply count	7.5	26.1	0	3	24	7741
Favorite count	44.1	123.7	6	21	127	7741

(b) Excluding non-US

	Mean	SD	5th Pctl	Median	95th Pctl	Obs
Quote retweet count	4.4	12.4	0	2	15.0	5128
Retweet count	28.5	58.6	3	16	84.0	5128
Reply count	9.3	31.3	0	3	30.6	5128
Favorite count	50.8	148.0	5	22	149.6	5128

*Notes:* This table reports descriptive statistics of media outlets’ tweets about inflation from 2014 to 2021. Reported descriptive statistics include the numbers of quote retweets, retweets, replies and favorites of media outlets’ tweets. Panel (a) presents descriptive statistics for all base tweets containing “CPI”, “PPI”, or “inflation”. Panel (b) presents descriptive statistics that exclude non-US news.

**Table A.6:** Tweets in the timelines of quote retweeters of inflation base tweets

	Mean	SD	5th Pctl	Median	95th Pctl	Obs
# tweets	433.7	9583.5	1	30	867.3	14935

*Notes:* This table reports descriptive statistics of users’ timelines based on tweets one day before and one day after the quote retweets of the base tweets on inflation.

**Table A.7:** Topics estimated with LDA: Inflation

Topic 1 "RBC"		Topic 2 "New Keynesian"	
Term	Probability	Term	Probability
inflation	0.023	price	0.015
rate	0.016	inflation	0.012
fed	0.012	year	0.010
said	0.011	said	0.007
year	0.011	cost	0.005
bank	0.011	consumer	0.004
percent	0.010	month	0.004
policy	0.008	economy	0.004
market	0.008	would	0.004
central	0.007	higher	0.004
price	0.007	also	0.004
economy	0.006	increase	0.004
central bank	0.005	wage	0.003
interest	0.005	time	0.003
growth	0.005	last	0.003
month	0.004	rise	0.003
target	0.004	government	0.003
would	0.004	people	0.003
bond	0.004	one	0.003
interest rate	0.004	since	0.003

*Notes:* This table reports topics estimated with the LDA model on articles about the yield curve, with  $K = 5$  and symmetric Dirichlet priors. For each topic, we report the distribution over vocabulary terms estimated with the LDA model.

## B. Latent Dirichlet allocation

Latent Dirichlet allocation (LDA) developed by [Blei et al. \(2003\)](#) is a generative probabilistic model that is aimed at reducing the dimensionality of text corpus. This section presents details of the model.

We represent each *word* from our vocabulary as a basis vector of length  $V$  with a single component equal to 1 and all other components equal to zero. For example, the  $v$ th word is denoted as  $w = (0, \dots, 0, 1, 0, \dots, 0)$  where  $w_v = 1$  and  $w_u = 0$  if  $u \neq v$ . Then, an *article* is a vector consisting of  $N$  words, i.e.,  $w = (w_1, \dots, w_N)$  where  $w_n$  is the  $n$ th word. Finally, A *corpus* is a collection of  $M$  articles, i.e.,  $D = \{w_1, \dots, w_M\}$ .

Consider a  $k$ -dimensional Dirichlet random variable  $\theta$  with a parameter vector  $\alpha = (\alpha_1, \dots, \alpha_K)$ , whose probability density over a  $(k - 1)$ -simplex is given by

$$p(\theta|\alpha) = \frac{\Gamma(\sum_{i=1}^k \alpha_i)}{\prod_{i=1}^k \Gamma(\alpha_i)} \theta_1^{\alpha_1-1} \dots \theta_k^{\alpha_k-1} \quad (14)$$

where  $\Gamma(x)$  is the Gamma function. Then, LDA assumes the following data generating process for each article  $d$  in our corpus  $D$ :

1. Draw  $N \sim \text{Poisson}(\xi)$ ;
2. Draw  $\theta \sim \text{Dirichlet}(\alpha)$ ;
3. Each word  $w_n$  is generated from a two-step process:
  - (a) Draw a topic  $z_n \sim \text{Multinomial}(\theta)$ ;
  - (b) Draw a word  $w_n$  from  $p(w_n|z_n, \beta)$ , the multinomial probability conditioned on the topic;

where  $\beta$  denotes a  $k$ -by- $V$  matrix with  $\beta_{ji} = p(w_j = 1|z_i = 1)$  that represent word probabilities.

Given the parameters  $\alpha, \beta$ , the distribution over a topic  $\theta$ , a set of topics  $z$ , and a set of  $N$  words, the joint likelihood is given by

$$p(\theta, z, w|\alpha, \beta) = p(\theta|\alpha) \prod_{n=1}^N p(z_n|\theta) p(w_n|z_n, \beta). \quad (15)$$

We can integrate over  $\theta$  and sum over  $z$  to obtain the marginal distribution of an article as

$$p(w|\alpha, \beta) = \int p(\theta|\alpha) \left( \prod_{n=1}^N \sum_{z_n} p(z_n|\theta) p(w_n|z_n, \beta) \right), \quad (16)$$

and we can obtain the probability of a corpus by taking the product of all marginal probabilities of single documents

$$p(D|\alpha, \beta) = \prod_{d=1}^M \int p(\theta_d|\alpha) \left( \prod_{n=1}^{N_d} \sum_{z_{dn}} p(z_{dn}|\theta_d) p(w_{dn}|z_{dn}, \beta) \right) \quad (17)$$

The inference problem that we solve with the LDA is to compute the posterior distribution of the unobserved variables given a document:

$$p(\theta, z|w, \alpha, \beta) = \frac{p(\theta, z, w|\alpha, \beta)}{p(w|\alpha, \beta)} \quad (18)$$

where

$$p(w|\alpha, \beta) = \frac{\Gamma(\sum_i \alpha_i)}{\prod_i \gamma(\alpha_i)} \int \left( \prod_{i=1}^k \theta_i^{\alpha_i-1} \right) \left( \prod_{n=1}^N \prod_{i=1}^k \prod_{j=1}^V (\theta_i \beta_{ij})_{w_n^j} \right) d\theta, \quad (19)$$

which we approximate using the online variational Bayes algorithm developed by [Hoffman, Bach and Blei \(2010\)](#).

Our text preprocessing is standard. We remove stop words such as “a” and “the”, numbers, words with a single character, and capitalization. We reduce the dimensionality of the corpus by lemmatizing, grouping together words with different forms that express the same meaning into a single token (for example, “curve” and “curves” are both lemmatized to “curve”).

## C. Measuring tweet sentiment

Based on the tweets from users’ timelines collected as described in the previous subsection, we estimate consumer sentiment using the naïve Bayes classifier developed by [Rish et al. \(2001\)](#). Using the Bayes law, the classifier represents the probability of the sentiment  $y = \{0, 1\}$  of a tweet consisting of terms  $(t_1, \dots, t_n)$  as:

$$p(y|(t_1, \dots, t_n) \propto p(y) \prod_{i=1}^n p(t_i|y) \quad (20)$$

As recognized by [Buehlmaier and Whited \(2018\)](#), naïve Bayes is one of the oldest tools in natural language processing and has better out-of-sample performance in text-based tasks than alternative models ([Friedman et al., 2001](#)). The special features in tweets require additional preprocessing. We convert all user mentions and links into single tokens (@USER and HTTPURL), remove special characters (RT and FAV), and fix common typos. For example, a raw tweet:

RT @UMich @UMichFootball: Victors valiant, champion of the west! <https://umich.edu/>

will be transformed to:

@USER @USER: victors valiant, champion of the west! HTTPURL

After pre-processing, we vectorize tweets using term-frequency inverse-document-frequency (tf-idf), which weighs a token by its importance to a document relative to the corpus ([Ramos et al., 2003](#)). The weighting is specified as:

$$\text{tf-idf}_{t,d} = \underbrace{\frac{w_{t,d}}{\sum_{\tau \in d} w_{\tau,d}}}_{\text{term frequency}} \cdot \underbrace{\log \frac{D}{|\{d \in D : t \in d\}|}}_{\text{inverse document frequency}} \quad (21)$$

where  $w_{t,d}$  represent the frequency count of term  $t$  in document  $d$ ,  $D$  represents the total number of documents, and  $|\{d \in D : t \in d\}|$  is the number of documents term  $t$  appears. Tf-idf reduces the importance of words that appear with high frequency, such as “the” or “we.”

Then we use the naïve Bayes algorithm to classify the sentiment of tweets. Specifically, we represent the probability that a tweet  $j$  conveys positive sentiment as a function of the

tf-idf-weighted terms  $t_1, \dots, t_n$  of in the tweet:

$$\tilde{p}_j(\text{positive}) = f(t_1, \dots, t_n) \quad (22)$$

where tildes indicate that the probability  $\tilde{p}$  is predicted by the naïve Bayes classifier.

We pre-train the naïve Bayes classifier using 100,000 pre-classified tweets in [Go, Bhayani and Huang \(2009\)](#), who use emoticons to automatically classify the sentiment of tweets as positive and negative. For example, smiley faces :) indicate positive tweets, and sad faces :( indicate negative tweets.

Based on the predicted sentiment from the naïve Bayes classifier, we define the sentiment of consumer  $i$  in day  $t$  as:

$$s_{it} = \frac{1}{J} \sum_j \tilde{p}_j(\text{positive}) \quad \text{for } j \text{ posted in day } t \quad (23)$$

where  $s_{it}$  measures the average sentiment of tweets posted by the consumer in a day. Values of  $s_{it}$  lie between 0 and 1, with values greater than 0.5 corresponding to positive sentiment. The higher the values of  $s_{it}$ , the more optimistic a consumer is of the outlook.