Essays in Macroeconomics and Financial Economics

by

Wenting Song

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2021

Doctoral Committee:

Professor John Leahy, Co-Chair
Assistant Professor Pablo Ottonello, Co-Chair
Professor Linda Tesar
Professor Toni Whited

Wenting Song

wtsong@umich.edu

ORCID iD: 0000-0002-7019-0116

To Aixiang He and Mingxiang Song

Acknowledgments

I am grateful to the members of my committee, John Leahy, Pablo Ottonello, Linda Tesar, and Toni Whited, for their thoughtful feedback that has guided the dissertation. Pablo's patience and support from the very beginning allowed me to find my identity as a researcher. The time and attention he has given to my work since then far exceed any reasonable expectations. His contagious passion for economics, boundless sense of optimism, and unfounded confidence in me have inspired me many times over the past six years to become a better researcher. It is a privilege to work with such a brilliant economist and kind human being. John inspired my interest in information and guided the writing of the dissertation with his many detailed and insightful comments. His careful introduction to macro theory provided a foundation on which I can always rely. Jointly, John, Linda, Pablo, and Toni created an intellectual environment in which any ambitious student can succeed.

I would also like to thank Costas Azariadis, Bruce Petersen, John Nachbar, Lee Benham, Eva Russo, John Shareshian, Thomas Noe, Kevin Sheppard, Stefan Nagel, Christopher House, Dominick Bartelme, Andrei Levchenko, Adam Stevenson, Daniil Manaenkov, Hongbin Qu, Sonali Das, Helge Berger, and Michael McMahon, who provided support and inspiration at various stages of my studies.

Fellow Michigan students have been incredible sources of support and beyond generous in sharing with me their vast reservoir of knowledge. In particular, I would like to thank my co-authors Sam Stern and Andrew Usher, as well as Ruoyan Sun, Bhanu Gupta, Xinzhu Chen, Xing Guo, Frank Li, Alberto Arredondo Chavez, Luis Baldomero Quintana, Michelle Lam, Patrick Wu, Nafisa Lohawala, Mos Laoprapassorn, Guangye Cao, Mattan Alalouf, Ellen Stuart, and Matthew Wilson, for countless discussions and for motivating each other through the long journey. Being surrounded by such a group of talented, ambitious and supportive peers has made it impossible not to work hard and challenge myself. Thanks also to Geno for going out of his way to cheer up every graduate student he meets in Lorch and to Frank Turner for his music.

To my oldest and dearest friends, Xiaofei, Lidi, Chenchen, Ran, and Jiemai. When I think about how we have stayed together despite wildly diverging life paths, it gives me hope that my model will one day converge too. Thank you for being there all these years.

Most importantly, I would like to express my deepest gratitude to my parents, who have encouraged me to stay true to myself and believed in me unconditionally. They have worked to provide me with everything I would need to pursue a fulfilling life, and because of them I have a place in this world.

Table of Contents

List of Figures

Figure

List of Tables

Table

List of Appendices

Appendix

Abstract

This dissertation contains three essays in macroeconomics and financial economics that aim to understand the empirical importance of information and intermediaries on aggregate fluctuations, and consequently, the role of macroeconomic policy.

Chapter 1 provides direct evidence of the importance of firm attention to macroeconomic dynamics. This chapter develops a text-based measure of firm attention to macroeconomic news and documents firm attention that is polarized and countercyclical. Differences in attention lead to asymmetric responses to monetary policy: expansionary monetary shocks raise stock returns of attentive firms more than those of inattentive firms, and contractionary shocks lower returns of attentive firms by less. In a quantitative model of rationally inattentive firms with parameters for information frictions calibrated using the text-based measure, firms invest in attention endogenously and face heterogeneous information costs. Less attentive firms adjust prices slowly in response to monetary innovations, which yields non-neutrality. As average attention varies over the business cycle, so does the efficacy of monetary policy.

Chapter 2 provides empirical evidence of the causal effects of changes in financial intermediaries' net worth in the aggregate economy. The empirical strategy developed in this chapter identifies financial shocks as the high-frequency changes in the market value of intermediaries' net worth in a narrow window around their earnings announcements, based on tick-by-tick data. News of declines in U.S. intermediaries' net worth leads to significant declines in the market value of nonfinancial firms. These effects are more pronounced for small firms and when the aggregate net worth of financial intermediaries is low. In addition,

this chapter discusses channels through which intermediaries affect nonfinancial firms, which provides evidence of the effect of intermediaries on corporate borrowing costs.

Chapter 3 provides empirical evidence that economic narratives influence consumer sentiment. It develops a framework that captures news narratives of economic events using natural language processing and traces consumers' exposure to different narratives using retweeting activities. The framework is applied to study the narratives surrounding the yield curve inversion. Exposure to the negative narrative of an imminent recession causes consumers to display a more pessimistic sentiment, while exposure to the positive narrative that recession concerns are overblown leads to no change in consumer sentiment.

Chapter 1

Firm Inattention and the Transmission of Monetary Policy: A Text-Based Approach

Wenting Song and Samuel Stern

1.1 Introduction

Public information often goes unused because attention is scarce. Rational inattention models pioneered by Sims (2003) and a broader set of incomplete-information models (Mankiw and Reis, 2002; Woodford, 2009)[1] consider firm managers who gather information to maximize value while facing cognitive costs of processing information. Inattention provides an intuitive microfoundation for monetary policy non-neutrality in which firm managers misinterpret nominal monetary policy as shocks to real demand, yet empirically assessing the importance of attention is challenging because neither a firm's allocation of attention nor information-processing costs are readily observable.

This paper provides the first direct evidence of the importance of firm attention to macroeconomic dynamics using a novel text-based measure of firm attention. We document countercyclical firm attention and uncover substantial heterogeneity in attention across firms. Moreover, our measure is consistent with the asymmetric prediction of inattention models

[1] Additional work includes Lucas (1972); Angeletos and La'O (2013); Gabaix (2019); Farhi and Werning (2019).

that attentive firms exhibit higher profit semi-elasticities in response to expansionary monetary shocks and lower semi-elasticities following contractionary shocks. We then use this measure to calibrate information costs in a quantitative general equilibrium model with rationally inattentive firms and show that firm inattention generates monetary non-neutrality. Together with our empirical evidence on countercyclical firm attention, this result suggests that aggregate attention to macroeconomic conditions is an important dimension of state-dependence in monetary policy.

To construct our attention measure, we compile a corpus based on approximately 200,000 annual SEC filings of US publicly-traded firms and search each document for macroeconomic keywords. We define two measures of attention: "prevalence", whether firm managers discuss macro conditions at all, and "intensity", the frequency at which managers discuss macro conditions.

We document two stylized facts about firm attention. First, firm attention is polarized. The majority of firms in our sample either mention macroeconomic conditions in every filing or in none of their filings. Second, attention is countercyclical. Among the remaining firms with time-varying attention, the number of firms that mentioned macroeconomic news rose notably during recessions.

Our main empirical result validates that our text-based methodology effectively measures attention by testing for an asymmetry in firm performance that is predicted by inattention models: following a macroeconomic shock, firms with greater information-processing capacity should respond closer to the optimal response regardless of the shock's direction. Therefore, more attentive firms should exhibit higher profit elasticities in response to positive shocks and lower elasticities in response to negative shocks as they update prices more accurately than inattentive competitors. We test for this asymmetry using an event-study design that exploits high-frequency variation in firms' market values around FOMC announcements. This test requires combining our prevalence attention measure with daily CRSP stock prices, quarterly Compustat firm financials, and high-frequency monetary shocks constructed as in

Gürkaynak et al. (2005) and Nakamura and Steinsson (2018a).

Consistent with the theoretical prediction, expansionary monetary shocks raise stock returns of attentive firms by 2% more than those of their inattentive peers, whereas contractionary shocks lower returns of attentive firms by 6% less. The suboptimal responses to monetary shocks by inattentive firms are direct evidence of the cost of inattentive behavior. Moreover, the asymmetry invalidates some concerns about measuring firm attention with text analysis. Concern that filings contain macroeconomic buzzwords as a form of cheap talk to appease investors would imply a zero effect; concern that firms mention keywords solely as a function of exposure to monetary policy would imply symmetric responses to monetary shocks; and concern that stock returns vary with investor attention rather than firm attention would also fail to explain the asymmetric responses.

We then use our attention measure in a quantitative rational inattention model to study the aggregate implications of the heterogeneity in firm attention. Firms with heterogeneous information costs optimally trade off between the precision of their signals on aggregate demand and the cost of acquiring and processing information. Information-processing costs and the distribution of firm attention are calibrated using our text-based attention measure. Consistent with our empirical findings, attentive firms in the calibrated model have higher semi-elasticities to expansionary monetary shocks and lower semi-elasticities to contractionary shocks. We incorporate the empirical countercyclicality of firm attention to show that the efficacy of monetary policy declines as the fraction of attentive firms increases and more firms set prices closer to the optimum. This new interpretation of attention-dependent monetary policy implies that central banks should expect the effects of policy to be weaker when an aggregate shock has already drawn firm attention to macroeconomic policy.

Related Literature Our paper contributes to four strands of literature. First, we contribute to the empirical literature on macroeconomic expectations by developing an ongoing, broad-based measure of firm attention that extends back to the mid-1990s. Recent literature has highlighted the importance of expectations for macroeconomic policy[2] and consequently the need for empirical measures[3]. To study expectations in a macroeconomic context requires measurement beyond lab evidence (Reutskaja et al., 2011) or individual consumers (McCaulay, 2020). Our methodology complements survey-based evidence on firm expectations by Tanaka et al. (2019), Coibion et al. (2018), Afrouzi (2020), and Candia et al. (2021) and enables researchers to explore questions that lie outside the coverage of existing surveys.

Second, our findings on firm inattention lend empirical support to a broad body of theoretical work on incomplete information as a source of monetary non-neutrality (Sims, 2003; Mankiw and Reis, 2002; Woodford, 2009). Microfoundations proposed in rational inattention and sticky information models are successful in explaining firm pricing (Mackowiak and Wiederholt, 2009; Afrouzi and Yang, 2021), asset prices (Van Nieuwerburgh and Veldkamp, 2009), discrete choices (Matějka and McKay, 2015; Caplin et al., 2019), and reconciling micro and macro evidence (Auclert et al., 2020). However, the lack of measurement on firm attention makes it challenging to assess the empirical relevance of these microfoundations. Our results estimate a substantial cost of information frictions in the US data, providing direct support for these theories.

Our findings on the relationship between countercyclical attention and monetary policy efficacy relates to existing literature on state dependencies of monetary policy. Tenreyro and Thwaites (2016) estimate non-linear responses in monetary policy which are weaker in recessions than in expansions. Vavra (2014), McKay and Wieland (2019) and Ottonello and Winberry (2020) consider volatility, durable consumption, and default risk as other channels through which state dependency arises. This paper suggests that attention may be an important source of state dependency of monetary policy.

Finally, our paper relates to a broader and emerging literature that brings natural language processing techniques to economics. The seminal work of Loughran and McDonald

[2]See, for example, Coibion and Gorodnichenko (2015); Coibion et al. (2020); Malmendier and Nagel (2016)

[3]See Gabaix (2019) for a comprehensive survey of existing measure of attention.

(2011) applies the "bag of words" method to firm filings and develops word lists specific to economic and financial texts. Recent work has used textual analysis to measure financial constraints (Buehlmaier and Whited, 2018), central bank communication (Hansen et al., 2018), firm-level political risk (Hassan et al., 2016), and uncertainty (Handley and Li, 2020). We contribute to the literature by constructing a dictionary of macroeconomic keywords with detailed categories based on releases of macroeconomic series.

Road map The rest of the paper proceeds as follows: in Section 1.2 we describe our methodology for measuring attention and present evidence of the stylized facts listed above; in Section 1.3 we present a theoretical framework that incorporates attention and exposure to macro shocks and derive the predicted asymmetry; in Section 1.4 we outline an empirical strategy for testing the effects of attention on expected returns and present our results; in Section 1.5 we construct a quantitative model of rational inattention and conduct policy counterfactuals; Section 1.6 concludes.

1.2 Textual measure of attention

This section presents our measure of firm attention to macroeconomic news for the universe of US publicly-traded firms between 1994 and 2019. We then document several stylized facts about firm attention before conducting the main empirical analysis in Section 1.4.

1.2.1 SEC filings

To measure firm attention, we employ the universe of annual 10-K filings with the U.S. Securities and Exchange Commission (SEC) between 1994 and 2019. Under Regulation S-K, all public companies are required to disclose financial statements and business conditions in these filings. The annual filings (Form 10-K) requires a more extensive discussion of business conditions and audited financial statements, while the quarterly filings (Form 10-Q) is usually less descriptive and only requires unaudited financial statements. Our sample contains 201,751 unique annual 10-K filings by 35,655 firms. Table 1.1 shows the summary statistics on the 10-K filings. The average length of 10-Ks is 30,647 words with 2,433 unique words.

Table 1.1: Summary statistics on 10-K filings

	N	**Mean**	**Median**	**SD**	**Min**	**Max**
Total word count	201,751	30,647	26,133	23,031	152	199,520
excl. stopwords	201,751	18,912	16,128	14,232	98	164,734
Unique word count	201,751	2,433	2,496	1,039	74	7,937
excl. stopwords	201,751	2,337	2,395	1,026	68	7,822

Discussion of economic conditions in an SEC filing typically appears in two contexts: recent or future firm performance and the risk factors that shareholders face by investing in the company. The former context usually appears in Item 7 of 10-K and 10-Q filings, which requires managers to discuss and analyze the firm's financial conditions and results of operations. This section is written as a narrative and can vary in length across firms (for instance, Item 7 of Alphabet's 2020 10-K filing is 17 pages long). Economic conditions in the context of risk factors commonly appears in Items 1A and 7A, which detail general firm risks and near-term market risks, respectively.

1.2.2 Methodology

Textual measure of firm attention To construct our main measures of firm attention to macroeconomic news, we employ dictionary-based frequency counts in natural language processing. We identify instances in which firms discuss the following nine macroeconomic topics: general economic conditions, output, labor market, consumption, investment, monetary policy, housing, and oil. Each topic is matched with a keyword dictionary that consists of names of major macroeconomic releases from `Econoday` (the data provider behind `Bloomberg`'s economic calendar) as well as words and phrases that commonly appear in popular articles on each topic. Any words or phrases that might apply to both aggregate-and firm-specific conditions are removed to avoid misidentification. For example, the phrase "interest rates" is excluded from the monetary policy dictionary because firms may mention interest rates in the context of their own liabilities. The dictionary of topics and associated keywords appears in Table A.1.

We then construct two measures of attention based on these keywords. Attention *prevalence*, d_{it}^k, indicates whether a firm i mentioned any keyword related to a given topic k in period t:

$$d_{it}^k = \mathbb{1}(\text{Total topic-k words}_{it} > 0) \qquad \text{(prevalence)}$$

Attention *intensity*, s_{it}^k, records the rate at which keywords are mentioned as a share of total words in the filing. We interpret this measure as the average intensity with which firms pay attention to economic conditions:

$$s_{it}^k = \frac{\text{Total topic k words}_{it}}{\text{Total words}_{it}} \qquad \text{(intensity)}$$

Total word count is generated by following the parsing strategy in Loughran and McDonald (2011). First, a text is stripped of all numbers and "stop words" such as articles. The text is then mapped onto a dictionary of words constructed by extending 2of12inf, a commonly-used collection of English words, to include additional words in 10-K documents.

Sense check of the textual measure As a preliminary sense check of the textual measure, Table A.2 in the Appendix reports the summary statistics of firm characteristics by attention. Attentive firms, whose prevalence attention to the general topic is nonzero in any year in the sample period, tend to be larger, older, and slightly less levered than their inattentive counterparts.

We then investigate the cross-industry variation in attention. Figure 1.1 reports the share of firms that pays attention to each topic by industry. Industry is measured using 2-digit NAICS from `Compustat`. The quality of our attention measure varies by topic so

Figure 1.1: Firm attention by industry

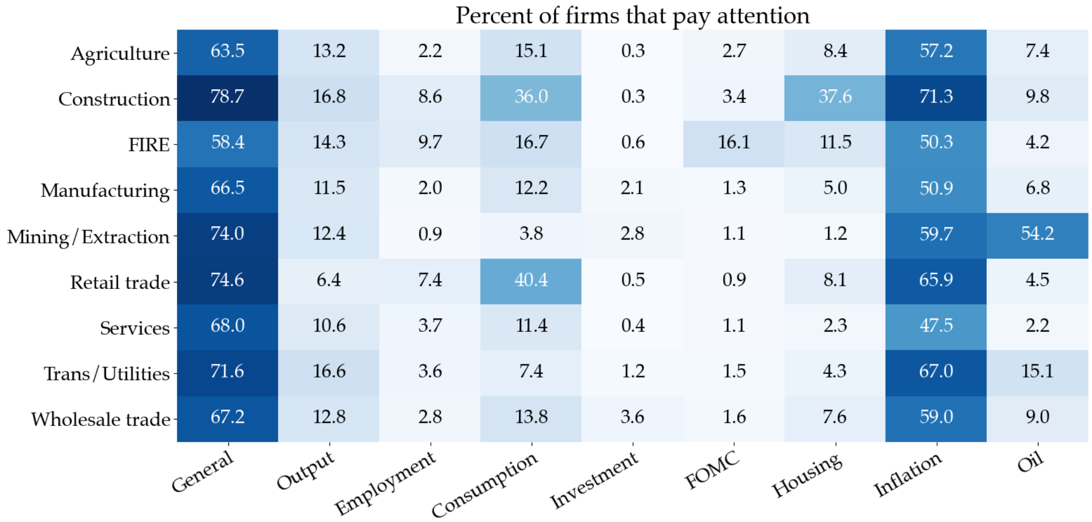

Notes: Heat map of the fraction of firms in an industry that pay attention to each macroeconomic topic. Industry is defined as 2-digit NAICS. Darker color represents a higher fraction of firms that pay attention.

these results should be interpreted across industry rather than across topic.

For each macro topic, attention is highest in industries for which profits are most sensitive to the topic a priori. For example, Mining, Oil, and Gas (NAICS 21) has the highest share of firms that pay attention to news about oil prices; Retail trade (NAICS 44-45) pays the greatest attention to news about consumption; and finance (NAICS 52) pays the greatest attention to news about FOMC meetings.

Furthermore, some industries appear to pay greater overall attention than others. Finance ranks among the most attentive industries to employment, FOMC, output, and interest rates, while agriculture (NAICS 11) and Professional, Scientific, and Technical Services (NAICS 54) appear least attentive overall.

The two features of cross-industry variation described above are fairly unsurprising and serve as sense checks of our attention measure. Put simply, industries whose profitability depends more on a certain macro topic have a higher share of firms that pay attention to that topic, and some industries appear to have greater overall interest in aggregate economic conditions.

Figure 1.2: Time series of attention to “economic conditions”

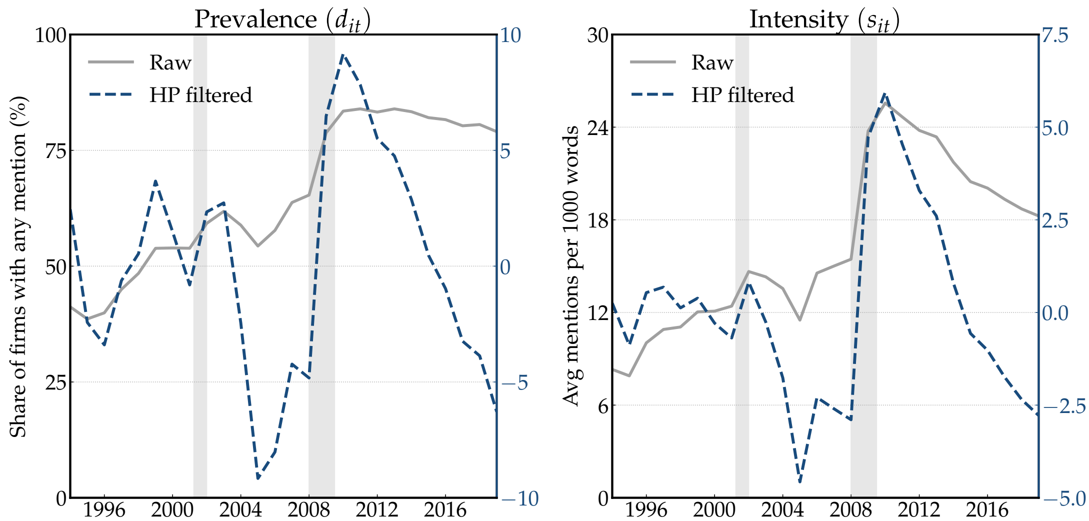

Notes: Time series of firm attention to the keyword “economic conditions”. Left panel plots the prevalence measure and reports the share of firms that mention the keyword. The right panel plots the intensity measure and reports the average mentions of the keyword per 1,000 words. “Raw” refers to the unfiltered series and “HP filtered” refers to the cyclical components of the HP-filtered series. Shares are reported in percent.

1.2.3 Stylized facts about firm attention

We now apply our prevalence and intensity measures to document two stylized facts about time and firm variation in attention: firm attention in the US is countercyclical and polarized.

Countercyclical attention to economic conditions Both the share of firms that mention macro keywords and the intensity with which firms mention macro keywords vary countercyclically over the business cycle. To illustrate this, we plot the time series related to the keyword “economic conditions”. Figure 1.2 plots the share of firms that mention the keyword. The left panel reports the prevalence measure, and the right panel reports the intensity measure. Both panels also show the cyclical components of the HP-filtered series.

The share of firms that mention “economic conditions” increased over the sample period, with faster growth during recessions. The share of firms jumped by about 15 percentage points during the Great Recession and has moderated to approximately 80% in subsequent years.

Figure 1.3: Share of filings that mention "economic conditions"

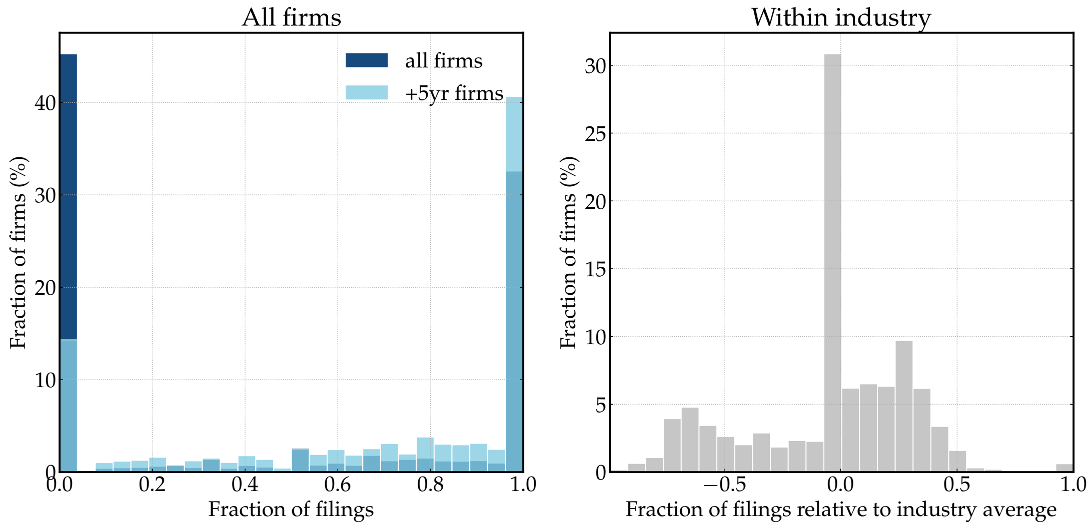

Notes: Histograms of the share of filings by a firm that mention "economic conditions". The left panel shows the histogram of the average fraction of filings that mention the keyword "economic conditions" over the sample period of 1994-2019. Dark blue bars correspond to the distribution of all firms, and light blue bars correspond to firms appearing for at least 5 years in the sample. The right panel shows the histogram of the time series averages of the residuals of firm attention to "economic conditions" after regressing on industry fixed effects. Shares of firms on the vertical axes are reported in percent.

The intensity related to the keyword "economic conditions" across all filings displays a stronger cyclical trend than the share of firms mentioning output. The share of words increases more during recessions and falls faster during recoveries compared to the share of firms mentioning output.

Countercyclical attention exhibited in Figure 1.2 is consistent with predictions by Mackowiak and Wiederholt (2009), who model firms that allocate attention between aggregate and idiosyncratic conditions. Their model predicts that firms will pay more attention to aggregate conditions in downturns if those conditions become more uncertain. This result is also consistent with Chiang (2021), who develops a generalized information structure in which agents pay greater attention to uncertain aggregate conditions when expecting a bad economic state, which subsequently generates countercyclical attention and uncertainty.

Polarization in firm attention Heterogeneous attention to publicly available news about U.S. output provides the clearest evidence that firms are limited in their capacity to process available information. The profitability of all publicly traded firms in our sample is arguably exposed to variation in U.S. economic conditions, and we should expect firms with unlimited information-processing bandwidth to incorporate this news into their decision making. Evidence of heterogeneity is to the contrary and provides new insights into how firms allocate attention differently.

The left panel of Figure 1.3 plots the histogram of firms by average attention over the sample period. The number of bins matches the number of annual observations in our sample and can be doubly interpreted as the number or fraction of filings in which firms pay attention. A firm with a value of 0 for the fraction of filings on the horizontal axis has never mentioned "economic conditions" over the sample period, whereas a firm with a value of 1 has mentioned that phrase in every filing. Most notably, firms are concentrated at each extreme: either never mentioning a macroeconomic keyword in their filings or mentioning a macroeconomic keyword in every filing. Despite the countercyclical variation found above, it appears that most variation in attention occurs across firms and that attention is largely invariant over time.

To test whether this polarization is driven by firms with few filings, we replicate the histogram using a restricted sample of firms with at least five years of filings. Although this restriction greatly reduces the number of firms that never pay attention to macroeconomic news in our sample, the polarization between always- and never-attentive firms remains.

We also test whether polarized attention is attributable to industry patterns in attention. The right panel of Figure 1.3 demeans firm attention by industry to isolate within-industry heterogeneity. This panel depicts a large degree of variation in attention even after accounting for industry averages. Aside from a high concentration of attention at the industry average, demeaned attention also appears bimodally dispersed.

The concentration at the industry average raises concern about the text-based measure:

Does the frequency of macroeconomic keywords in 10-K filings capture firm attention to macroeconomic news or firm exposure to aggregate conditions? It is entirely plausible that a firm does not discuss the macroeconomy because its profits are insensitive to aggregate fluctuations. Our main empirical analysis in Section 1.4 will focus on disentangling our hypothesized attention channel from this alternative exposure channel. We test our hypothesis by separately estimating the response of stock prices to positive and negative macro shocks. If firms discuss macro news more often because they are more exposed to aggregate fluctuations then "attentive" firms would profit more from a positive shock and lose more from a negative shock, generating symmetric relative responses to macro shocks. On the other hand, if the text-based measures indeed capture attention, then attentive firms would outperform inattentive competitors regardless of the direction of the shock, resulting in asymmetric relative responses. The theoretical framework in Section 1.3 discusses the mechanism in detail.

1.3 Illustrative framework

Motivated by the evidence that firms are heterogeneous in their attention to macroeconomic news, we set out to study how firm attention affects the transmission of macroeconomic policy. Before doing so, we address a key identification challenge: whether our text-based attention measures identify differences in firm attention to macroeconomic conditions, conditional on firm characteristics, rather than differences in exposure to those conditions. To confront the identification challenge, we lay out a stylized model in which firms are heterogeneous in both attention and exposure. For the two sources of heterogeneity, the model yields contrasting predictions for stock return responses to monetary shocks, which we then exploit to guide our regression specifications. The model environment is minimal to highlight the key mechanisms for attention and exposure. In Section 1.5, we expand the model environment to incorporate more realistic assumptions.

Environment Time is static. Consider a firm whose profits, $\pi(s,a)$, depend on an aggregate state variable, s, and a firm action, a. Assume that $\pi(s,a)$ is twice continuously differentiable, a single-peaked function of a, and maximized at $a^* = s$. For concreteness, we think of a as the price that a monopolistically competitive firm sets and s as the exogenous optimal price determined by factors outside of that firm's control, as in Woodford (2009).

Firm profits can be approximated under a second-order log approximation around the non-stochastic steady state as[4]:

$$\hat{\pi}(\hat{s},\hat{a}) = \pi_s(\bar{s},\bar{a})\bar{s}\hat{s} + \frac{1}{2}\left(\pi_{ss}(\bar{s},\bar{a})\bar{s}^2 - \pi_{aa}(\bar{s},\bar{a})\bar{a}^2\right)\hat{s}^2 + \frac{1}{2}\pi_{aa}(\bar{s},\bar{a})\bar{a}^2(\hat{a}-\hat{s})^2 \tag{1.1}$$

where $\bar{s}$ and $\bar{a}$ denote the steady-state values, $\hat{\pi}$, $\hat{s}$ and $\hat{a}$ denote the log deviations from the steady state, and $\pi_s \equiv \frac{\partial}{\partial s}\pi(s,a)$, $\pi_{aa} \equiv \frac{\partial^2}{\partial a^2}\pi(s,a)$ and $\pi_{ss} \equiv \frac{\partial^2}{\partial s^2}\pi(s,a)$.

Lastly, assume that firm profits are increasing in s, $\pi_s > 0$, and that the second-order condition for a stable equilibrium holds, $\pi_{aa} < 0$.

Attention and Exposure We can now define attention and exposure in the model. A firm is more exposed to aggregate conditions if its profits are more sensitive to aggregate shocks, while a firm is more attentive if its action are more sensitive more to shocks. Definitions 1 and 2 formalize these ideas.

Definition 1 (attention)**.** *Let a firm's action be a function of the state: $\hat{a} = f(\hat{s})$, with $f(0) = 0$ and $0 < f'(\hat{s}) \leq 1$. Firm i is attentive to macroeconomic conditions if $f_i'(\hat{s}) = 1$, and firm j is inattentive to macroeconomic conditions if* $0 < f_j'(\hat{s}) < 1$.

An attentive firm reacts one-for-one with innovations to the aggregate state, whereas an inattentive firm responds less than one-for-one. The simplified definition of inattention

[4] Under this approximation, $\pi_a(s,a)$ drops out because of the first-order condition and assumption that $a^* = s$ at the optimum. Appendix A.5.1 contains detailed derivations of the approximation.

is consistent with that in rational inattention models such as Sims (2003) which yields a steady-state Kalman gain between 0 and 1.

Definition 2 (exposure). *Firm i is more exposed to macroeconomic conditions than firm j if $\pi_s^i(s,a) > \pi_s^j(s,a)$.*

Differences in attention and exposure We now derive model predictions for heterogeneity in attention and exposure that guide the empirical analysis to come.

We first construct stock returns, which is the dependent variable in our empirical analysis. As in Gorodnichenko and Weber (2016), a firm's stock price is equal to its firm value, which in the simple static setting equals its profits:

$$v = \pi(s,a)$$

Realized equity returns, measuring the log change in a firm's value around an aggregate shock, are given by:

$$r = \hat{v} - \hat{v}_{-1} \tag{1.2}$$

where $\hat{v} \equiv \log V - \log \bar{V}$ denotes the log deviation of firm value from the steady state, and $\hat{v}_{-1} \equiv \log \mathbb{E}_{-1} V - \log \bar{V}$ denotes the log deviation of firm value before the shock is realized.

Proposition 1 highlights the asymmetric responses of stock returns to positive and negative aggregate shocks that result from the attention channel and the symmetric responses from the exposure channel.

Proposition 1. *The return elasticity with respect to aggregate shocks for the exposure and the attention channels can be characterized as below:*

1. ***Exposure**: If firm i is more exposed to macroeconomic conditions than firm j, then holding all else equal the return elasticity of firm i with respect to the aggregate shock*

is higher than the return elasticity of firm j for all shocks:

$$\frac{\partial r_i}{\partial \hat{s}} > \frac{\partial r_j}{\partial \hat{s}} \quad \forall \hat{s}$$

2. ***Attention**: Suppose firm i is attentive to macroeconomic conditions and firm j is inattentive. Then, holding all else equal, the return elasticity of a positive (expansionary) shock is higher for the attentive firm i than that of the inattentive firm j. For negative (contractionary) shocks, the return elasticity for the attentive firm i is lower than for the inattentive firm j. For zero shocks, the return elasticities for attentive and inattentive firms equal:*

$$\begin{cases} \frac{\partial r_i}{\partial \hat{s}} > \frac{\partial r_j}{\partial \hat{s}} & \text{if } \hat{s} > 0 \\ \frac{\partial r_i}{\partial \hat{s}} = \frac{\partial r_j}{\partial \hat{s}} & \text{if } \hat{s} = 0 \\ \frac{\partial r_i}{\partial \hat{s}} < \frac{\partial r_j}{\partial \hat{s}} & \text{if } \hat{s} < 0 \end{cases}$$

Proof. See Appendix A.5.2 □

Figure 1.4 illustrates the predictions from Proposition 1. In Panel (a), firms are heterogeneous in their exposures to aggregate shocks, and those with high exposure exhibit higher return elasticities to aggregate shocks regardless of the sign of the shock.Panel (b) illustrates the mechanism of attention. Attentive firms are better at tracking the state variable, so their stock returns outperform those of inattentive firms after any aggregate disturbance. In response to a positive shock, stock returns of both attentive and inattentive firms rise, but returns of attentive firms rise more. In response to a negative shock, returns of both types of firms decrease, but returns of attentive firms drop by less.

This asymmetry in return elasticities is a unique feature of the attention channel and allows us to distinguish between the effects of firm attention and exposure to macro news. In the next section, we use this predicted asymmetry to show that our text-based measure

Figure 1.4: Model predictions for exposure vs attention

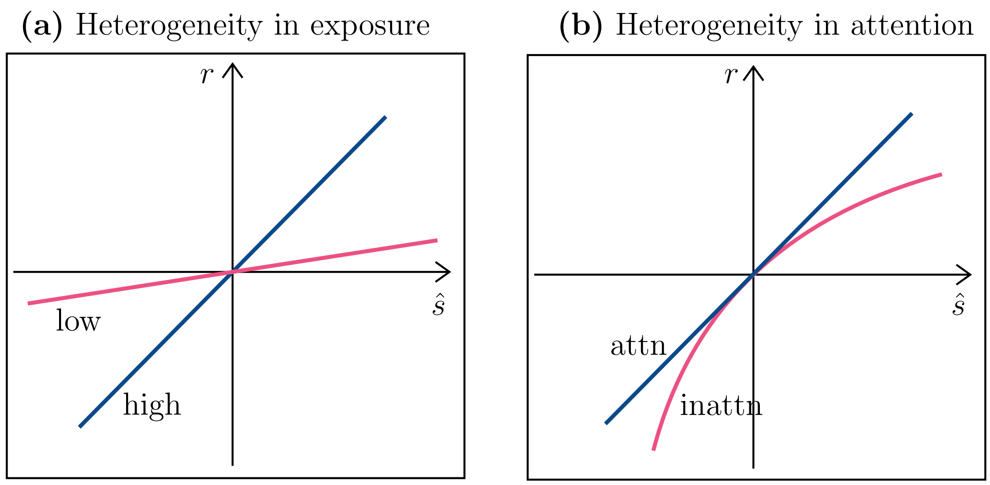

Notes: Illustration of model predictions of return elasticity with respect to aggregate shocks. Verticle axes represent conditional realized return, and horizontal axes represent the magnitude of shocks. Left panel shows return elasticity for firms that are highly exposed to macro conditions (*high*) and firms that are unexposed (*low*). Right panel shows return elasticity for attentive firms (*attn*) and inattentive firms (*inattn*). Exposure and attention are as defined in the main text.

correctly identifies firm attention and then estimate the cost of inattention based on the difference in return elasticities for positive and negative shocks.

1.4 Empirical analysis

Given our attention measures and theoretical predictions, we set out to test the hypothesis that attentive firms respond to macro shocks better than inattentive firms. We use a high-frequency identification strategy that isolates plausibly exogenous shocks to monetary policy from FOMC announcements and compares changes in stock prices of attentive and inattentive firms within a similarly narrow window around these announcements. We implement our empirical analysis with monetary policy shocks since they are familiar and well-identified[5], though the mechanism highlighted in our stylized inattention model is general to any aggregate shock.

Stock prices are a particularly informative outcome variable because they are forward-

[5] Ramey (2016) provides a comprehensive survey on the efforts on identifying monetary shocks.

looking and quickly reflect changes in expected future profits. By focusing on the high-frequency windows of stock price movements, we are able to separate effects of monetary surprises from other confounding factors. More direct measures of firm responses such as investment and hiring decisions are only observed over longer time horizons and are confounded by other factors that influence firms' choices.

To best isolate the effects of attention, our baseline specification controls for firm size, age, leverage, and industry measured by 4-digit NAICS. The underlying identifying assumption is that firms have similar exposure to monetary policy shocks within a narrowly defined industry after conditioning on firm characteristics and financial structure. Residual variation in stock prices can then be attributed to firm attention rather than cross-firm variation in the exposure to monetary policy.

1.4.1 Data

Monetary policy shocks are constructed using the high-frequency identification strategy developed by Cook and Hahn (1989) and citetgurkaynak2005actions, and used recently in Gorodnichenko and Weber (2016), Nakamura and Steinsson (2018a), and Ottonello and Winberry (2020). These shocks are measured as the change in the fed funds futures rate within a one-hour window surrounding FOMC announcements. Any changes within such a narrow window can be attributed to unanticipated changes to monetary policy as it is unlikely that other shocks occurred within the same window.

Monthly fed funds futures contracts clear at the average daily effective fed funds rate over the delivery month, so rate changes are weighted by the number of days in the month that are affected by the monetary policy shock. Following notation in Gorodnichenko and Weber (2016), the final shock series is defined as,

$$v_t = \frac{D}{D-\tau}(ff^0_{t+\Delta t^+} - ff^0_{t-\Delta t^-}), \tag{1.3}$$

where t is the time of the FOMC announcement, $ff^0_{t+\Delta t^+}$ and $ff^0_{t-\Delta t^-}$ are the fed funds futures rates 30 minutes before and after the announcement, D is the number of days in the month of the announcement, and τ is the date of the announcement. We use the series published by Gorodnichenko and Weber (2016) and Nakamura and Steinsson (2018a) for monetary shocks from 1994 to 2014.

Firm outcome and control variables are constructed using `CRSP` and `Compustat` data. Daily stock returns are measured as the open-to-close change in stock prices on the day of an FOMC announcement, matching the precise timing of the FOMC announcements which fall between 1pm to 3pm. Firm size, age, and industry controls are constructed as described in Section 1.2.3.

Firm attention is measured using the *prevalence* measure, d_{it}, described in Section 1.2. To better suit a high-frequency methodology, firm attention at the time of an FOMC announcement is identified using the firm's most recent annual filing rather than the filing in the same year as the FOMC announcement. This modification excludes the possibility that firms are identified as attentive to an FOMC announcement using a subsequent filing.

1.4.2 Methodology

We separately estimate the slope of the interaction between monetary shocks and firm attention for positive and negative shocks, and then test whether these two coefficients are statistically different.

For a firm i in industry j on day t, our baseline model takes the form,

$$\begin{aligned} r_{it} = \delta_j + \delta_j v_t + \beta_d d_{it} + \beta_1 \mathbb{1}_{v_t>0} + \beta_{v+} v_t \mathbb{1}_{v_t>0} + \beta_{v-} v_t \mathbb{1}_{v_t<0} \\ + \beta_{dv+}(d_{it} \cdot v_t \cdot \mathbb{1}_{v_t>0}) + \beta_{dv-}(d_{it} \cdot v_t \cdot \mathbb{1}_{v_t<0}) + \mathbf{\Gamma'}(\mathbf{X}_t + \mathbf{X}_t v_t) + \varepsilon_{it}, \end{aligned} \tag{1.4}$$

where d_{it} is the attention prevalence, v_t is the monetary policy shock, $\mathbb{1}_{v>0}$ indicates positive monetary policy shocks, $\mathbb{1}_{v<0}$ indicates negative monetary policy shocks, and $\mathbf{X}_t$ is a set of controls including the indicator variable for positive shocks and quarterly firm controls for size, age, and leverage. We also control for the interaction of monetary shocks with industry fixed effects, $\delta_j v_t$, and with firm controls, $\mathbf{X}_t v_t$, to capture the effects of firm characteristics on differential responses to monetary shocks. Standard errors are clustered by FOMC announcement to allow for correlated errors across firms at each FOMC announcement.

The coefficients of interest are β_{dv+} and β_{dv-}. The theoretical framework in Section 1.3 hypothesizes β_{dv+} to be positive and β_{dv-} to be negative, implying attentive firms should outperform inattentive firms in response to both expansionary and contractionary monetary shocks. To formally test the hypothesis, we conduct a Wald Test with the null hypothesis $H_0 : \beta_{dv+} = \beta_{dv-}$.

1.4.3 Empirical results

Our baseline results are reported in Table 1.2. In the first column, we estimate the effect of high-frequency monetary shocks without our attention measures and find that a 25 basis point unanticipated increase in the fed funds rate is associated with about a one percent increase in stock prices. This result is consistent with existing estimates from Gorodnichenko and Weber (2016) and Nakamura and Steinsson (2018a). The second column introduces the unconditional interaction between monetary shocks and firm attention. We find that attentive firms experience slightly higher stock returns than their inattentive counterparts but our estimate is not statistically distinguishable from zero. This result is consistent with the framework outlined in Section 1.3, which remains agnostic as to the average interaction over the entire range of monetary shocks.

The main results from Equation (1.4) are presented in the third column. We test whether attention leads to differential responses to positive and negative monetary shocks. Consistent with predictions from rational inattention models, attentive firms appear to experience larger increases in stock returns following expansionary monetary shocks and smaller decreases in stock returns following contractionary monetary shocks. The coefficients are statistically different from zero, and the Wald Test of whether these coefficients are equivalent is rejected

Table 1.2: Baseline results

	(1) Average	(2) Exposure	(3) Attention	(4) excl. ZLB
Shock	4.55*	4.55*		
	(2.53)	(2.65)		
Attention		-0.01	-0.07	-0.03
		(0.05)	(0.06)	(0.06)
Shock × Attn		1.07		
		(0.64)		
Shock × $\mathbb{1}_{v_t>0}$			4.93*	6.54**
			(2.74)	(2.75)
Shock × $\mathbb{1}_{v_t<0}$			-3.57	-0.95
			(3.72)	(3.69)
Shock × Attn ×$\mathbb{1}_{v_t>0}$			2.02***	1.55**
			(0.72)	(0.72)
Shock × Attn ×$\mathbb{1}_{v_t<0}$			-5.87*	-5.77*
			(3.18)	(3.30)
Observations	575667	575667	575667	432458
R^2	0.022	0.022	0.026	0.027
Clustered SE	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes
excl. ZLB	no	no	no	yes
Wald Test p-value			0.026	0.050

Notes: Results from variants of estimating the baseline specification (1.4):

$$r_{it} = \delta_j + \delta_j v_t + \beta_d d_{it} + \beta_1 \mathbb{1}_{v_t>0} + \beta_{v+} v_t \mathbb{1}_{v_t>0} + \beta_{v-} v_t \mathbb{1}_{v_t<0} \\ + \beta_{dv+}(d_{it} \cdot v_t \cdot \mathbb{1}_{v_t>0}) + \beta_{dv-}(d_{it} \cdot v_t \cdot \mathbb{1}_{v_t<0}) + \mathbf{\Gamma}'(\mathbf{X}_t + \mathbf{X}_t v_t) + \varepsilon_{it},$$

where v_t is the monetary shock, d_{it} is the prevalence attention measure, δ_j and $\delta_j v_t$ are an industry fixed effect and its interaction with the shock, and $\mathbf{X}_t$ contains firm-level controls of size, age and leverage, and $\mathbf{X}_t v_t$ contains the interactions between firm controls and the shock. We have normalized the sign of the monetary shock ν_t so that a positive shock is expansionary (corresponding to a decrease in interest rates). Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

at 5% significance.

Finally, the fourth column ends the sample in 2007 to exclude the zero lower bound period following the Great Recession. This excludes periods of forward guidance and unconventional monetary policy and allows us to focus on conventional monetary transmission. Results are both qualitatively and quantitatively similar as in the full sample, suggesting our findings are not driven by anomalies from the financial crisis, the zero lower bound periods, or the information effects of monetary policy.

The *asymmetric* responses to positive and negative shocks are consistent with heterogeneous responses predicted by a model of inattention and rules out alternative interpretations of the textual measure that predict symmetric responses. The first alternative interpretation, discussed in detail in Section 1.3, is that the textual measure misidentifies firms' profit exposure to macroeconomic conditions as attention. In this case, symmetric responses to positive and negative monetary shocks would yield a positive and significant effect from the interaction term between shock and attention (β_{dv}) in the second column, which is inconsistent with our empirical findings. A second alternative hypothesis is that firms attribute poor performance to broader economic forces and are more likely to mention FOMC meetings when underperforming. We would then expect attentive firms to underperform in response to negative monetary shocks, corresponding to a positive coefficient for β_{dv-} in the third column, which is also at odds with our empirical findings. Another concern is that investor attention is more important to stock price than firm attention. Inattentive investors would then systematically under-react to both positive and negative shocks, which fails to explain the observed asymmetry. A final concern is that firms may differ in price stickiness even within a narrow sector beyond controlled characteristics. For price stickiness to explain our empirical results, it must be correlated with macro keyword counts in SEC filings, which seems unlikely.

For additional robustness, the next subsection shows that our results are robust to controlling for firm management quality, past exposure to monetary policy, information effects of monetary policy, and macro variables.

The suboptimal responses to monetary policy by inattentive firms reported in Table 1.2 together with the large fraction of inattentive firms documented in Figure 1.3 provide some of the first direct evidence on the empirical consequences of firm inattention in the US. We estimate that inattentive firm returns rise by 2% less following positive shocks and drop by 6% more following negative shocks compared to those of their attentive peers. These differences are substantial given the average stock return response of 5%.

1.4.4 Additional empirical results

The Appendix contains three sets of additional empirical results.

Drivers of firm attention The first set of additional results in Appendix A.2 investigates firm characteristics that may drive attention to macroeconomic news. We now treat our prevalence attention measure as the dependent variable and test its association with firm size, age, leverage, management quality, and exposure to monetary policy announcements. For each set of results, we separately estimate cross- and within-firm effects by including sector and time fixed effects for the former and firm fixed effects for the latter.

In the cross-section, higher attention is associated with younger firms, consistent with Cloyne et al. (2018)'s findings of stronger monetary responses by young firms, yet appears to increase over time within-firm. Attention is also higher among larger firms and is weakly associated with lower leverage, which is consistent with Ottonello and Winberry (2020)'s findings of stronger monetary responses from low-leverage firms. Firm attention is associated with higher management quality when measured as the share of board members who hold a graduate degree, and with greater exposure to monetary policy when measured as a rolling five-year average stock price response to high frequency monetary shocks.

Robustness checks Appendix A.3 checks whether our baseline results in Section 1.4 are sensitive to additional drivers of firm attention or potentially confounding effects to high frequency monetary shocks. We first control for the aforementioned drivers of firm attention. Our baseline specification already includes controls for industry and firm covariates (age, size and leverage). Tables A.6 and A.7 additionally control for management quality and exposure to monetary shocks, respectively, and show that our baseline asymmetric semi-elasticities are robust in each case.

Two concerns that have been raised about high frequency monetary shocks are that i) an "information effect" confounds the direct effects of a change to interest rates (Nakamura and Steinsson, 2018a), and ii) monetary shocks may be correlated business cycle fluctuations. Following Miranda-Agrippino and Ricco (2021), we control for each FOMC announcement's information effect using Greenbook forecast revisions between FOMC announcements and show that our main results are little changed in Table A.8. We then incorporate macro controls in Table A.9 including lagged unemployment, real output growth, and inflation. Again, our main results are robust to these controls.

Limitations and promise of textual measures Recycled or *boilerplate* language is a key concern of using regulatory filings to measure firm attention. 10-K filings are often written collaboratively between managers and legal departments, and evidence suggests that firms include certain statements within 10-K filings to appease investors or lower liability (**?**). Moreover, firms likely save time and resources by revising their filing from the prior year rather than starting from scratch. Boilerplate language is a concerning source of measurement error when it includes keywords that identify firm attention. In Appendix A.4.3, we test whether boilerplate language contaminates our main results by measuring the diversity in filing language with a Jaccard score of lexical similarity and restricting our analysis to the most linguistically diverse 10-K sections. Results in Table A.10 are qualitatively and quantitatively similar as the baseline results.

Even greater measurement error may come from misidentifying attentive firms as inattentive (Type II error), which raises concerns about *underestimating* overall firm attention. False negatives may occur if our text analysis fails to identify discussion of economic topics due to limited sophistication, or if attention is not uniformly publicized in 10-K filings across firms. For the purposes of this paper, underestimated attention would attenuate our results and imply that our current estimate for the cost of information frictions serves as a lower bound.

Text analysis methods also hold tremendous promise for uncovering a more refined depiction of firm attention and expectations formation. We illustrate these capabilities with two approaches for identifying the context in which firms discuss economic conditions. The first approach (Appendix A.4.2) uses a Latent Dirichlet Allocation (LDA) unsupervised model to categorize words that neighbor a given keyword. The second approach (Appendix A.4.1) uses the itemized structure of 10-K filings to identify which sections of the filing contained the most keywords.

1.5 Quantitative model

Motivated by the empirical heterogeneity in firm attention, we now construct a general-equilibrium model with rationally-inattentive firms to understand the aggregate implications of heterogeneous firm attention. Key parameters of the model are calibrated using the attention measure and empirical moments from the sections above. Using the quantitative model, we explore the effects of inattention on monetary policy.

1.5.1 Model environment

The model mechanism is an extension of the stylized model outlined in Section 1.3. Time is discrete and infinite. The economy consists of households, firms and the central bank. Households and the central bank have full information about the economy, while firms face information frictions. We start with a standard general equilibrium model with rationally inattentive firms as in Mackowiak and Wiederholt (2009) and Afrouzi and Yang (2021). Attention is modeled with the Shannon mutual information following Sims (2003) and is an endogenous choice by the firm (Luo et al., 2017). Then we incorporate heterogeneous costs of information and connect model objects to the data to calibrate parameters for information frictions.

Household A representative household maximizes its life-time utility,

$$\max_{C_{it},N_t} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (\log C_t - \psi N_t), \tag{1.5}$$

where N_t denotes the labor supply and ψ represents the disutility of labor. Consumption C_t is aggregated over each good type i with a CES aggregator,

$$C_t = \left(\int_0^1 C_{it}^{\frac{\varepsilon_p - 1}{\varepsilon_p}} dj \right)^{\frac{\varepsilon_p}{\varepsilon_p - 1}}, \tag{1.6}$$

where ε_p is the elasticity of substitution. In addition to the wage income, households have access to a one-period bond B_t with the interest rate ι_t and receive a lump-sum transfer D_t from the government. The household budget constraint is given by:

$$\int_0^1 P_{it} C_{it} di + B_t \leq W_t N_t + (1 + \iota_t) B_{t-1} + D_t \tag{1.7}$$

Central Bank The central bank targets aggregate money supply similar to Caplin and Spulber (1987) and Gertler and Leahy (2008). As a result, the nominal aggregate demand follows an autoregressive process:

$$\Delta \log Q_t = \rho \Delta \log Q_{t-1} + \nu_t, \quad \nu_t \sim N(0, \sigma_\nu^2). \tag{1.8}$$

Firms Firms are owned by a risk-neutral agent and have production technology that is linear in labor:

$$Y_{it} = N_{it}.$$

Figure 1.5: Firm's timeline

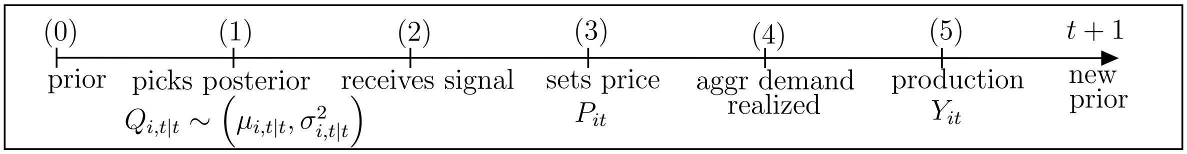

The functional form of a firm's information flow is specified with Shannon's mutual information:

$$\mathcal{I}(Q_{i,t|t-1}, Q_{i,t|t}) = \frac{1}{2} \log \frac{\sigma^2_{i,t|t-1}}{\sigma^2_{i,t|t}}, \tag{1.9}$$

which captures the expected reduction in entropy from prior $Q_{i,t|t-1}$ to posterior $Q_{i,t|t}$. The Shannon mutual information is decreasing in the posterior variance, so that more precise posteriors are more expensive. The marginal cost of information per nat, $2\omega_i$ is heterogeneous across firms and can be either high or low:

$$\omega_i \in \{\omega_H, \omega_L\}.$$

This heterogeneity is motivated by our empirical finding of polarized firm attention.

Figure 1.5 shows a firm's timeline. It enters a period with a prior on the aggregate demand. Then it chooses the posterior distribution. Since the Shannon mutual information in (1.9) does not depend on the posterior mean, it is optimal for a firm to center the posterior distribution around the true mean. So the firm's information choice is only of the posterior variance $\sigma^2_{i,t|t}$. Based on the chosen posterior distribution, the firm receives a signal on the aggregate demand and sets its price P_{it} based on the posterior belief. Then, the aggregate demand is realized, the firm produces and enters the next period with a new prior.

A firm's value function is given by

$$V(\sigma^2_{i,t|t-1}) = \max_{P_{it},\sigma^2_{i,t|t}} \mathbb{E}_t \left[\underbrace{\frac{Y_{it}}{P_t}(P_{it} - MC_t)}_{\text{flow op. profits}} - \underbrace{2\omega_i \mathcal{I}(Q_{i,t|t-1}, Q_{i,t|t})}_{\text{info costs}} + \beta \underbrace{V(\sigma^2_{i,t+1|t})}_{\text{cont. value}} \middle| \sigma^2_{i,t|t} \right], \quad (1.10)$$

which consists of flow operational profits that are maximized when firms successfully track the aggregate demand, information costs that depend on firms' information acquisition choices, and a continuation value. The expectation operator of a firm is based on its time-t information set. The problem of a firm's manager in each period is to maximize the firm value by jointly setting prices and investing in attention.

Firms optimize subject to the following constraints:

$$Y_{it} = (P_{it}/P_t)^{-\varepsilon_p} C_t \qquad \text{(demand)}$$

$$\sigma^2_{i,t+1|t} = \rho^2 \sigma^2_{i,t|t} + \sigma^2_\nu \qquad \text{(law of motion for prior)}$$

$$0 \leq \sigma^2_{i,t|t} \leq \sigma^2_{i,t|t-1} \qquad \text{(no forgetting)}$$

The demand function comes from the household's problem, and the law of motion for a firm's prior belief is derived from the central bank's monetary rule. The no-forgetting constraint prohibits firms from discarding previously-acquired information to make room for new information, ensuring the Shannon information costs are non-negative.

Equilibrium The equilibrium consists of the household allocation, $\{C_t, \{C_{it}\}_{i\in[0,1]}, N_t\}_t$, firms allocations, $\{\sigma^2_{i,t|t}, P_{it}, Y_{it}\}_t$, and a set of prices $\{P_t, W_t\}_t$ such that:

1. Given prices and the firms' choices, the household optimizes (1.5);
2. Given an initial prior $\sigma^2_{i,0|-1}$, prices and the households' choices, firms optimize (1.10);
3. Monetary policy follows (1.8);
4. All markets clear.

Model Solution Following Mackowiak and Wiederholt (2009) and Afrouzi and Yang (2021), we approximate firm's flow profits with second order log approximations around the full-information steady state.[6] This approximation yields an imperfect-information firm value, $\tilde{v}$. We decompose a firm's total value under log approximation, v, into a full-information value, v^*, representing the firm's value under optimal pricing with full information, and the imperfect information value, $\tilde{v}$, representing the loss in firm value from imperfect information.

The firm's imperfect information problem is solved numerically using the algorithm for dynamic rational inattention problems (`DRIPs`) developed by Afrouzi and Yang (2021).

1.5.2 Calibration

Calibration features two sets of parameters: standard parameters unrelated to information frictions and parameters related to information frictions. Importantly, we calibrate parameters related to information frictions to match the stylized facts on attention and the empirical elasticities estimated in the empirical analysis.

Standard parameters The top panel of Table 1.3 shows the calibration for predetermined parameters. The model period is a quarter, so the discount rate is set as $\beta = 0.95^{1/4}$. The monetary shock process is calibrated using quarterly US nominal output between 1994 and 2019. To match our empirical specification, which compares firms within a sector, we restrict our attention to nominal output in the manufacturing sector. The persistence of the shock is calibrated to $\rho = 0.89$ and the standard deviation is calibrated to $\sigma_\nu = 0.063$. The elasticity of substitution is set to $\varepsilon_p = 11$, implying a steady-state markup of 10%, and the disutility of labor is set to $\psi = 0.91$ to offset the steady-state distortions from monopolistic competition.

[6] Details of the approximation can be found in Appendix A.6.1. Log-quadratic approximation is a common simplifying assumption in rational inattention models to address the curse of dimensionality that arises from firms having the joint distribution of prices and nominal aggregate demand as the state variable. Sims (2003) shows that the optimal distribution under Gaussian priors and quadratic payoffs is also Gaussian, so log-quadratic approximation of the profit function greatly reduces the dimensionality of the problem.

Table 1.3: Calibration

Parameter	Description	Value
Standard parameters		
β	discount rate	$0.95^{1/4}$
ρ	shock persistence	0.89
σ_ν	shock std. dev.	0.063
ε_p	elasticity of substitution	11
ψ	disutility of labor	0.91
Information-friction parameters		
θ	fraction of attentive firms	65%
ω_L	cost of information	30
ω_H	cost of information	47

Information-friction parameters The bottom panel of Table 1.3 contains calibrations for parameters $(\theta, \omega_L, \omega_H)$. To calibrate these important parameters governing the degree of information frictions in the model, we use our text-based measure of attention and the empirical moments from Section 1.4.

The fraction of attentive firms is set to $\theta = 65\%$ to match the average fraction of firms that have paid attention to the keyword "economic conditions" over the sample period. Attention to economic conditions conveys firm attention to aggregate demand, which is a direct counterpart of the model state variable that firms track.

To calibrate the costs of attention, ω_L and ω_H, we target regression coefficients in Table 1.2 by running the same regressions with simulated model data. We first define model objects that match those observed in the data. Stock returns in the model are defined as the log change in a firm's value function in Equation (1.10), $r_{it} = \log V_{it} - \log \mathbb{E}_{t-1}(V_{it})$. We define attention in the model to be the Shannon mutual information. Since our main empirical specification uses the prevalence attention measure, we define a corresponding attention indicator, d_{it}, to equal 1 when a firm's attention is above the cross-sectional mean in a given period and 0 otherwise. Finally, we use ν_t as the monetary shocks. We simulate the model for a panel of 100 firms and for 1000 quarters, discarding the first 100 quarters as burn-in.

The cost of information for inattentive firms, ω_H, is calibrated to target $\hat{\beta}_v$ in Column

Figure 1.6: Sensitivity of simulated moments to costs of information

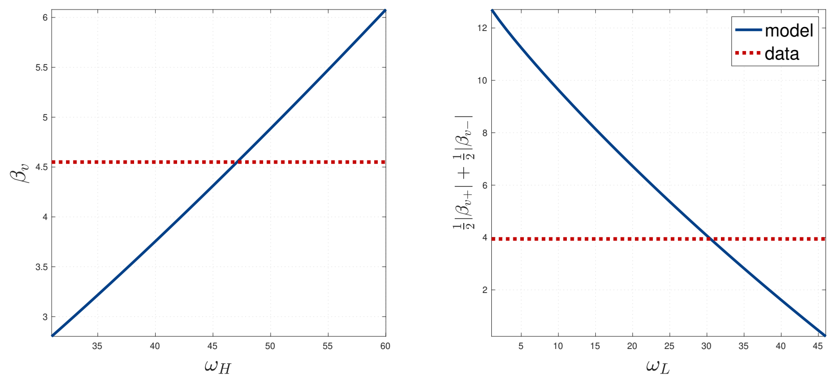

Notes: Simulated moments for a range of costs of information parameters. We simulate models for a panel of 100 firms and for 1000 periods with 100 periods burn-ins. Simulated moments are generated with regressions discussed in the text.

(2) of Table 1.2, which measures the average response of stock returns to monetary policy. With simulated data, we run the following regression:

$$r_{it} = c + \beta_v \nu_t + \beta_d d_{it} + \beta_{dv} d_{it} \nu_{it} + \varepsilon_{it},$$

and set ω_H so that the simulated β_v matches the empirical moment $\hat{\beta}_v$. The left panel of Figure 1.6 shows how ω_H is identified. We simulate the model for a range of values of ω_H. As the costs of information for attentive firms ω_H increases, the average response to monetary policy β_v increases monotonically.

For a given ω_H, we then set the cost of information for attentive firms, ω_L, to match $\hat{\beta}_{dv+}$ and $\hat{\beta}_{dv-}$ in Column (3) of Table 1.2, which measure the heterogeneous return semi-elasticity to monetary policy. The distance between ω_H and ω_L reflects the relative cost of information for inattentive firms compared to attentive firms. We run the regression with simulated data:

$$r_{it} = c + \beta_1 \mathbb{1}_{v>0} + \beta_{v+} v_t \mathbb{1}_{v>0} + \beta_{v-} v_t \mathbb{1}_{v<0} + \beta_d d_{it} + \beta_{dv+} d_{it} \nu_{it} \mathbb{1}_{v>0} + \beta_{dv-} d_{it} \nu_{it} \mathbb{1}_{v<0} + \varepsilon_{it}$$

In particular, the elasticity from Column (3) we target is $\frac{1}{2}|\hat{\beta}_{dv+}| + \frac{1}{2}|\hat{\beta}_{dv-}|$, which measures the relative stock return losses of firms that do not pay attention. The right panel of Figure 1.6

shows how ω_L is identified. Given a value of ω_H, we simulate the model for a range of ω_L. As ω_L increases and the gap between ω_H and ω_L narrows, the simulated elasticity monotonically decreases, implying lowering heterogeneity between attentive and inattentive firms. Figure A.5 in the appendix shows how simulated estimates of β_{dv+} and β_{dv-} change as a function of ω_L. β_{dv+} is positive and β_{dv-} is negative, suggesting that the stock returns of attentive firms outperform those of their inattentive peers for both positive and negative monetary shocks. As ω_L increases and the gap between the information costs for attentive and inattentive firms narrows, β_{dv+} decreases and β_{dv-} increases, implying a smaller difference in attention between attentive and inattentive firms.

The information cost parameters are calibrated to $\omega_L = 30$ and $\omega_H = 47$.[7] The calibration implies significant information costs for firms, which might seem surprising considering macroeconomic series are freely available. However, as plant-level evidence by Zbaracki et al. (2004) suggests, information costs involve not only information gathering costs but also information processing costs and communication costs. More recently, Abis and Veldkamp (2020) estimate the data production function which takes labor and capital inputs to process unstructured data into structure data and analyze data to produce knowledge. It requires significant manpower and expertise to process, summarize and forecast macroeconomic series into sufficient statistics that aids a firm's investment, production and pricing decisions, as highlighted in Reis (2006). The parameters of information costs in our model represent the costs of both acquiring and processing information.

1.5.3 Model dynamics

We now study how firm inattention results in monetary non-neutrality. Figure 1.7 shows the impulse responses to expansionary and contractionary monetary shocks of one standard

[7] The only preceding calibration for firm cost of attention is Afrouzi (2020), which studies the rational inattention problem of New Zealand firms under strategic complementarity and calibrates $\omega = 0.3$ using firm beliefs reported in New Zealand surveys. Our calibration differs by the sample of US firms and the abstraction from strategic complementarity. Flynn and Sastry (2021) build upon our calibration strategy of matching conditional regression moments to study countercyclical attention.

Figure 1.7: Firm impulse responses to monetary shocks

(a) Firm prices and operating profits

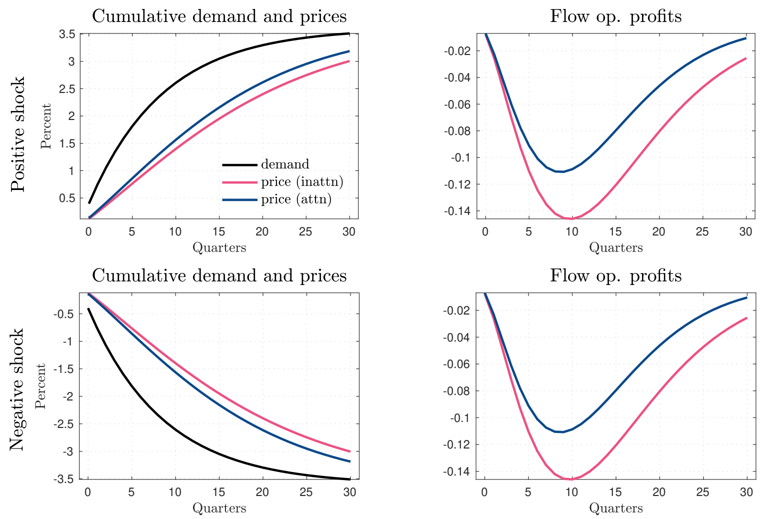

(b) Conditional realized returns

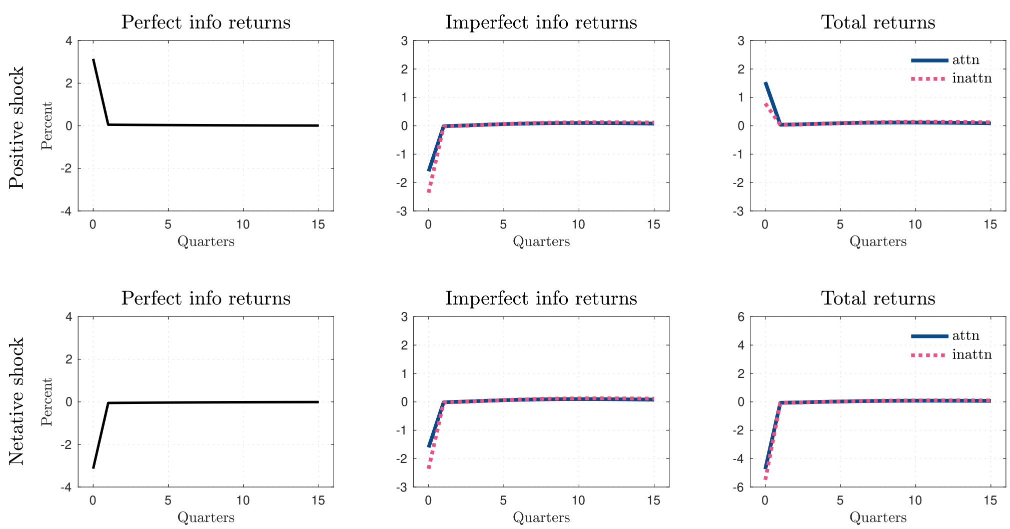

Notes: Firm impulse responses to a one standard deviation positive (expansionary) monetary shock and negative (contractionary) shock. Impulse responses are in percent deviations from the perfect-information steady state. "demand" refers the nominal aggregate demand. "attn" refers to the impulse responses of attentive firms, "inattn" refers to the impulse responses of inattentive firms.

deviation. Inattentive firms are shown in red, and attentive firms are shown in blue. Panel (a) shows the responses of firms' prices and flow operating profits. As nominal aggregate demand rises, firms' prices respond sluggishly, reflecting partial incorporation of noisy signals about demand. Attentive firms track aggregate demand better than inattentive firms and exhibit more responsive prices. Since we approximate firm profits around the full-information steady state, any deviation from the full-information benchmark results in a loss. Inattentive firms experience greater operational losses because they have less precise information about the aggregate demand. Inattentive firms also pay higher information costs despite acquiring less information because they face a higher marginal cost of information. With a constant marginal cost of information, firms' equilibrium choice of attention is not time-varying and therefore does not result in a change in returns.

Panel (b) shows the responses of stock returns. Following an expansionary monetary shock, full-information equity returns of both attentive and inattentive firms increase, since firms are monopolistically competitive. Returns of attentive firms increase by more than those of inattentive firms because attentive firms track the optimal price more closely. Returns of both imperfect-information firms are lower than those of a full-information firm that sets the optimal price. Following a contractionary shock, returns of attentive firms drop by less than those of inattentive firms.

In Figure 1.8, we study the aggregate responses of output and inflation to a one standard deviation expansionary monetary shock by aggregating attentive and inattentive firms.

In response to a nominal aggregate demand shock equivalent to a 25 basis point interest rate cut[8], annualized inflation and output increase by 0.04% and 0.07% at their peak, respectively. As a benchmark, Christiano et al. (2005) estimates the annualized peak effect of monetary policy shocks as 0.2% for inflation and 0.5% for output. With information as the only source of friction, our model generates about one seventh of their output response.

[8]Our model considers monetary policy shock to the nominal aggregate demand and Christiano et al. (2005) consider shocks to the interest rate. In Appendix A.6.3 we estimate the passthrough of interest rate on the nominal aggregate demand with manufacturing output data.

Figure 1.8: Aggregate responses to expansionary monetary shock

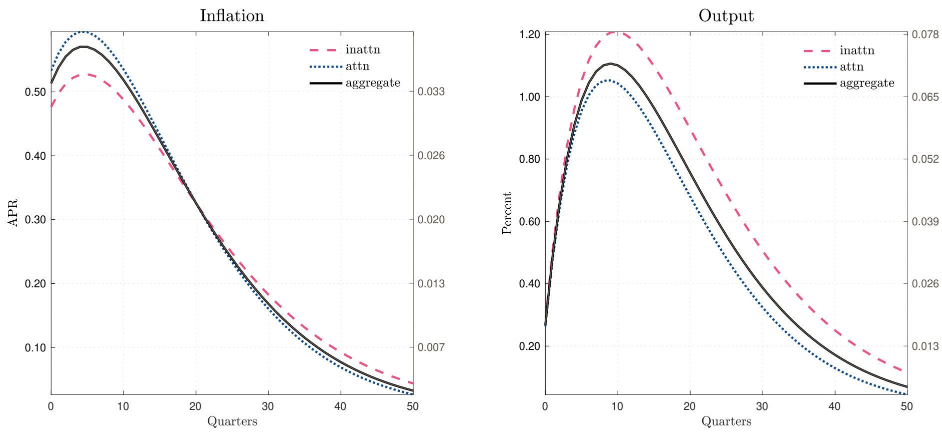

Notes: Impulse responses of inflation and output. The right scales show the impulse responses to a one standard deviation expansionary monetary shock, and the right scales show the impulse responses to an equivalent of 25 basis point expansionary monetary policy shock. Impulse responses are in percent deviations from the perfect-information steady state. "attn" refers to the impulse responses of attentive firms, "inattn" refers to the impulse responses of inattentive firms, and "aggregate" refers to the aggregate impulse responses.

1.5.4 Inattention and the efficacy of monetary policy

In our rational inattention model, monetary non-neutrality increases with both the fraction of inattentive firms and cost of information acquisition. Section 1.2 documents that firm attention evolves countercyclically over the business cycle, with attention rising during both the 2001 recession and the Great Recession.

The countercyclicality of aggregate attention suggests an important insight about the efficacy of monetary policy: when the Federal Reserve cuts rates during an ongoing recession, monetary policy is less powerful because firms are likely paying more attention to central bank decision-making. With a higher fraction of attentive firms, information frictions are less severe, monetary policy is closer to neutral, and monetary stimulus has a smaller effect on output. In contrast, preemptive monetary policy measures aimed at averting a potential recession are more powerful because a smaller fraction of firms are likely to respond attentively.

Table 1.4: Attention and monetary non-neutrality

	Least attentive	Baseline	Most attentive
Fraction of attentive firms (θ)	56%	65%	73%
Average output response (%)	0.1016	0.0992	0.0971

Notes: Dependence of output responses on the fraction of attentive firms in the economy. Average output responses are calculated over 50 periods. Calibration for the least and most attentive economy is described in the main text.

To illustrate the quantitative scope of the effect, we exogenously vary the fraction of attentive firms in the model and measure the average responses to a one standard deviation expansionary monetary shock. We start with the baseline calibration for the fraction of attentive firms, $\theta^{\texttt{baseline}} = 65\%$, which is the time series average of the prevalence measure of firm attention to aggregate demand between 1994 and 2019. Then, we decompose the time series of attention into the trend and cyclical components with the HP filter:

$$d_t = \tau_t + \zeta_t + \xi_t$$

where τ_t, ζ_t and ξ_t denote the trend, cyclical and error components of the attention measure d_t, respectively. The series frequency is annual and the smoothing parameter for the HP filter is set to 400. We then add the minimum (maximum) of the cyclical component to the baseline calibration to form the most (least) attentive calibration of the model:

$$\theta^{\texttt{least attn}} = \theta^{\texttt{baseline}} + \min(\zeta_t)$$
$$\theta^{\texttt{most attn}} = \theta^{\texttt{baseline}} + \max(\zeta_t)$$

where $\min(\zeta_t)$ and $\max(\zeta_t)$ correspond to the minimum and maximum of the HP-filtered prevalence measure in the left panel of Figure 1.2. Therefore, $\theta^{\texttt{least attn}} = 56\%$ and $\theta^{\texttt{most attn}} = 73\%$.

Then we study how aggregate responses to monetary policies change as we vary the fraction of attentive firms in the economy. Table 1.4 shows the average responses of output relative to the steady state over 50 periods. The average output response to monetary policy is 5% weaker in the most attentive calibration compared to the least attentive calibration. This suggests if the Federal Reserve cuts rates in the depth of a crisis period such as the COVID-19 pandemic when all firms are paying attention to macroeconomic policies, its monetary stimulus will be 5% weaker than if it cuts rates in a preemptive fashion to lean against the wind. This result is consistent with studies on the state dependency of monetary policy that find US monetary policy to be weaker in recessions than in expansions (Tenreyro and Thwaites, 2016).

1.6 Conclusion

The empirical evidence of information frictions that we document in this paper, along with growing evidence in the literature (Candia et al., 2021), highlights firms' deviation from full-information rational expectations (FIRE) in the US. To discipline models without FIRE, researchers require an understanding of firms' information sets and expectations formation processes.

In that direction, this paper presents a new text-based measure of firm attention to macroeconomic news, which will be made available publicly and updated on an ongoing basis. We validate that the measure indeed measures firm attention by testing for an asymmetric prediction of rational inattention on monetary policy transmission. We show that firms that pay attention to the FOMC have larger increases in stock returns after positive monetary shocks and smaller decreases in stock returns after negative monetary shocks, providing direct empirical evidence for the consequences of firm inattention.

The empirical measure can be used in combination with imperfect-information models to ground those theories with data. We demonstrate the value of this measure in a quantitative rational inattention model by showing that time variation in firm attention has important implications for the state dependency of monetary policy. In the model, average inattention drives the degree of monetary non-neutrality. The countercyclical nature of firm attention to macroeconomic news implies that the efficacy of monetary policy is weaker during recessions and should be considered in policy design.

Chapter 2

Financial Intermediaries and the Macroeconomy: Evidence from a High-frequency Identification

Pablo Ottonello and Wenting Song

2.1 Introduction

What effect do financial intermediaries have in the macroeconomy? This question, which has been central to macroeconomics at least since the Great Depression, received significant attention from researchers over the last decade (see, for example Gertler and Gilchrist, 2018). The main empirical challenge in addressing this question is that changes in macroeconomic conditions originating outside the financial system affect the balance sheets of intermediaries, making it challenging to isolate their aggregate effects on the economy.

In this paper, we propose a high-frequency (HF) identification strategy to study the causal effects of financial shocks in the aggregate economy. Our empirical strategy focuses on changes in individual financial intermediaries' net worth in a narrow window around their earnings announcements. In the spirit of the HF event-study approach to identify monetary-policy shocks (surveyed in Nakamura and Steinsson, 2018b), our empirical strategy exploits the fact that earnings announcements are lumpy, which leads to a discontinuity in the content of financial news released around these events. Using these shocks, we document that declines in the market value of U.S. intermediaries' net worth leads to substantial effects in the market value of nonfinancial firms. These effects are more pronounced for firms with small market capitalization and when the aggregate net worth of the financial system is low. We examine potential channels through which intermediaries' net worth affect nonfinancial firms and document intermediaries' persistent effects on borrowing costs of high-default-risk firms.

Our paper begins by constructing a HF measure of financial shocks in the U.S. economy. Our measure of financial shocks uses tick-by-tick data of intermediaries' stock prices in 60-minute windows around their earnings releases. We exploit the fact that publicly traded financial intermediaries have considerable market size, so idiosyncratic news about these intermediaries can have an effect in the aggregate economy, as in the recently proposed "granular" identification strategy (Gabaix and Koijen, 2020). We focus on commercial and investment banks because of their direct involvement in the external finance of nonfinancial firms, the main focus of our analysis. The combined net worth of our sample of intermediaries represents 64% of the total equity of U.S. depositary institutions in the period from 1998 to 2014. We provide three pieces of supportive evidence that these shocks reflect primarily information about intermediaries from their earnings releases and not about the rest of the economy. First, the stock price movements of intermediaries around their earnings announcements tend to be positively associated with the component of their earnings not anticipated by market participants. Second, intermediaries' stock prices exhibit larger volatility around earnings releases with respect to nonevent dates, which is less observed for nonfinancial firms, suggesting that intermediaries' information content around their announcements is larger than that of other firms. Finally, using a state-of-the-art machine-learning model, we show that financial shocks are not predictable by macroeconomic and financial variables available prior to the shock, suggesting that financial shocks are not driven by information in the rest of the economy available before intermediaries' earnings have been released.

We use HF financial shocks to study the effect of changes of intermediaries' net worth on nonfinancial firms. We provide evidence from two empirical strategies. One is an event-study approach, whose identifying assumption is that, in a 60-minutes window around intermediaries' earnings announcement, changes the stock price of intermediaries releasing earnings are driven by information contained in these announcements. The other is a heteroskedasticity-based identification strategy (Rigobon, 2003; Rigobon and Sack, 2004; Hébert and Schreger, 2017), whose identifying assumption is that the variance of intermediaries' stock price during earnings-announcement events is larger than in nonevents, while those of nonfinancial firms are the same during event and nonevent dates. Using these two strategies, we document that a one-percent change in intermediaries' net worth leads to a 0.4-to-0.7-percent change in the market value of nonfinancial firms in the S&P 500. These effects are larger for small firms, as measured by the returns of the S&P SmallCap 600 and Russell 2000 indices. Aggregate conditions in the financial system play an important role in determining these effects; the effects we identify are governed by periods in which the aggregate net worth for the financial system is low.

We then study potential channels through which intermediaries affect nonfinancial firms. Using bond-level data, we document that financial shocks have persistent effects on corporate bond spreads. These effects are concentrated on spreads of high default-risk firms: For corporate bonds rated CCC or lower, a one-percent decline in the market value of intermediaries' net worth leads to 0.1–0.3-percentage-point increases in spreads in the month following an event. These effects are observed even for individual bonds for which intermediaries releasing earnings have no holdings, suggesting an important role for aggregate net worth channels (as stressed, for instance, in Gertler and Kiyotaki, 2010).

Related literature Our paper is related to three strands of the literature. First, an important body of theories argues that financial intermediaries play an important role for macroeconomic dynamics and asset prices Gertler and Kiyotaki (2010); He and Krishnamurthy (2011, 2013); Brunnermeier and Sannikov (2014). Empirical work documenting the role of these aggregate effects have used time series methods (see, for example, Bernanke, 2018; Gertler and Gilchrist, 2018), the combination of cross-sectional and regional data (Gertler and Gilchrist, 2019), and model-based inference (see, for example, Christiano et al., 2015; Herreño, 2020). Our empirical analysis provides evidence of intermediaries in the aggregate economy, as well as for the role of aggregate intermediaries' net worth shaping these effects, based on a HF identification strategy. We see our method as complementary to existing empirical work, with the advantage of requiring milder assumptions for identification (as discussed by Nakamura and Steinsson, 2018a, in the context of monetary policy shocks).

Second, our paper is related to a large body of empirical work providing evidence that the net worth of financial intermediaries affects firms (e.g., Khwaja and Mian, 2008; Chodorow-Reich, 2014; Huber, 2018) and asset prices (see, for example, Coval and Stafford, 2007; Adrian et al., 2014; He et al., 2017; Siriwardane, 2019; Morelli et al., 2019). See also He and Krishnamurthy (2018) for a recent survey. An important element in the identification strategy developed in this body of work is the cross-sectional exposure of firms or assets to intermediaries. Our paper complements this literature by documenting intermediaries' aggregate effects.

Third, our paper is related to the macroeconomic literature that seeks to use direct empirical evidence to inform macroeconomic models Nakamura and Steinsson (2018b). Methodologically, our work is close to the HF approach to study the effect of monetary policy shocks in the economy (see, for example, Cook and Hahn, 1989; Kuttner, 2001; Cochrane and Piazzesi, 2002; Gürkaynak et al., 2004; Bernanke and Kuttner, 2005; Gorodnichenko and Weber, 2016; Nakamura and Steinsson, 2018a) and of the effects of sovereign default risk in the economy (Hébert and Schreger, 2017).

Outline The rest of the paper is organized as follows. Section 2.2 describes the data used in the empirical analysis. Section 2.3 describes the construction of the HF financial shocks and discusses their properties. Section 2.4 studies the effect of financial shocks on the economy.

Section 2.5 concludes.

2.2 Data

Our measure of financial shocks uses tick-by-tick data of intermediaries' stock prices in a window around their earning releases. We obtain the stock-price data from the NYSE's *Trade and Quote* (TAQ) and the earnings announcements' specific dates and times from *IBES Academic*. Our baseline sample focuses on commercial banks, investment banks, and securities dealers listed on the S&P 500 during the period 1998 to 2014.[1] We focus on these types of intermediaries because their direct involvement in lending activities in the economy makes them more likely to be linked to the macroeconomy, our main focus of analysis. In additional analysis, we create broader measures of financial shocks to include other types of financial intermediaries. Table 2.1 details the set of 18 financial intermediaries selected with our main criteria, together with the period in which they are included in our analysis. In our period of analysis, we obtain 870 announcements of earnings, roughly four per institution–year.

We study the effect on stock markets using daily indices data from *FRED*, including the S&P 500, S&P Small Cap 600, and Russell 2000. Table B.1 presents descriptive statistics of daily stock returns in our period of analysis, showing that days with financial shocks exhibit descriptive statistics similar to those of the whole period of analysis. We complement this analysis with tick-by-tick data for nonfinancial firms from TAQ.

We also study the effect of financial shocks in the corporate bond market by using daily individual bond-level data from the constituents of the Intercontinental Exchange Bank of America's (ICE BofA) AAA and CCC and Lower U.S. Corporate indices. For each of these bonds, we have information on bond option-adjusted spreads, together with characteristics including index weightings, ratings, residual maturities, average trailing 30-day spreads, and month-to-date changes in spreads. Following Anderson and Cesa-Bianchi (2020), we use the

[1]The financial intermediaries we use in the analysis correspond to NAICS 522110 and 523110.

Table 2.1: Financial intermediaries included in the sample

Financial Intermediary	**Ticker**	**Start**	**End**	**Avg Equity ($ billion)**	**Share of Sample**	**Share of Aggr Equity**
Citicorp	CCI, C	1998Q1	2014Q4	148.8	26.8%	13.2%
Bank of America	BAC	1998Q1	2014Q4	136.4	24.6%	12.1%
Wells Fargo	WFC	1998Q1	2014Q4	73.6	13.3%	6.5%
Goldman Sachs	GS	2002Q3	2014Q4	51.7	6.8%	3.9%
Morgan Stanley	MWD, MS	1998Q1	2014Q4	37.3	6.7%	3.3%
J.P. Morgan Chase	CMB, JPM	1998Q1	2000Q4	36.0	1.1%	6.3%
Wachovia	WB	1998Q1	2008Q4[a]	35.8	4.2%	4.0%
Merrill Lynch	MER	1998Q1	2008Q4	25.4	3.0%	2.8%
U.S. Bankcorp	USB	1998Q1	2014Q4	22.1	4.0%	2.0%
Bank One	ONE	1998Q1	2004Q2[b]	19.8	1.3%	3.0%
Bank of New York Mellon	BK	1998Q1	2014Q4	18.7	3.4%	1.7%
Fleet Boston Financial	FBF	1998Q1	2004Q1[c]	14.9	0.9%	2.3%
Lehman Brothers	LEH	1998Q1	2008Q3	12.6	1.4%	1.4%
Ameriprise Financial	AMP	2005Q4	2014Q4	8.6	0.8%	0.6%
First Chicago	FCN	1998Q1	1998Q4[d]	8.2	0.0%	1.5%
MNBA Corp	KRB	1998Q1	2005Q4	7.6	0.6%	1.0%
Bankboston	BKB	1998Q1	1999Q3[e]	4.9	0.1%	0.9%
Northern Trust	NTRS	1998Q1	2014Q4	4.6	0.8%	0.4%
Mean				37.1	5.56%	3.71%
SD				42.4	8.04%	3.68%
Min				4.6	0.04%	0.41%
Max				148.8	26.82%	13.16%

Notes: This table lists the financial intermediaries included in the sample and their tickers in TAQ. "Avg Equity" is the time series average of total shareholder equity (SEQ in Compustat) of the financial intermediary. "Share of Sample" measures a financial intermediary's equity as a share of equity of all financial intermediaries in the sample. "Share of Aggr Equity" represents a financial intermediary's equity as a share of aggregate equity of U.S. depository institutions. [a]Acquired by Wells Fargo. [b]Merged with J.P. Morgan Chase. [c]Acquired by Bank of America. [d]Merged with Banc One to form Bank One. [e]Merged with Fleet to form Fleet Boston.

option-adjusted spread, defined as the amount by which the government spot curve is shifted to match the present value of discounted cash flows to the corporate bond's price, as the main measure of corporate bond spread because it incorporates both a maturity adjustment (Gilchrist et al., 2009), by computing the spread relative to a risk-free security of matching maturity, and an option adjustment (Duffee, 1998), by removing the price of the embedded option. The AAA index consists of 293 bonds and the CCC and Lower index consists of 3,308 bonds. Table B.2 presents descriptive statistics of the individual bond spread together with those for days with and without earning releases of financial intermediaries, which exhibit similar descriptive statistics.

We analyze the heterogeneous impacts of financial shocks on corporate bond markets using additional data from Bloomberg. We study whether the effects vary depending on the share held by the intermediary that releases the information. For each individual bond in our sample,[2] we obtain from Bloomberg (at the CUSIP level) the share held by each reporting financial institution. This information is available at the quarterly frequency, so we collect data for each financial shock and for each outstanding bond on the quarter before the shock. Table B.3 reports descriptive statistics of bond holdings. On average the financial intermediaries in our sample represent 1,760 of the reported holdings of these bonds. These holdings exhibit heterogeneity, with a standard deviation of 11,337, which can be used to study the differential effects of bonds that are more or less strongly held by institutions releasing earnings reports.

2.3 High-frequency financial shocks

2.3.1 Construction and descriptive statistics

We define the HF financial shocks as the change in the stock price of the intermediaries reporting earnings in a narrow window around their earnings announcements:

$$\varepsilon_t^{\mathrm{F}} = \sum_{i \in \mathcal{I}_t} \theta_{i,q(t)} (\log P_{i,\tau(i,t)+\Delta^+} - \log P_{i,\tau(i,t)-\Delta^-}), \tag{2.1}$$

where $\mathcal{I}_t$ denotes the set of intermediaries reporting their announcement in day t, $\tau(i,t)$ is the time of an announcement for institution i in day t (expressed in minutes within a day), $P_{i,\tau}$ is the stock price of institution i in period τ, Δ^+ and Δ^- control the size of the window around the announcement, and $\theta_{i,q(t)}$ is the market capitalization of institution i as a share of the total market capitalization of institutions in our sample in the quarter before announcement day $q(t)$. For announcements made within trading hours, we select

[2] To study the heterogeneous impacts of financial shocks, we restrict the sample of bonds to those rated CCC or lower issued by firms with at least 10 bonds outstanding.

$\Delta(t)^-$ to be 20 minutes before the announcement and $\Delta(t)^+$ to be 40 minutes after the announcement, following Nakamura and Steinsson (2018a) for monetary-policy shocks. For announcements that occur after trading hours, we compute the financial shock as change between the closing and opening log prices. Given that our measure is more precise for announcements made within trading hours, we create two measures of financial shocks: a "narrow" measure that includes only this type of announcements, and a "broad" one that includes both type of shocks. Appendix Figure B.1 illustrates our HF identified shocks with four graphical examples. Panels (a) and (b) show two shocks that occur inside trading hours, with their magnitudes corresponding to median positive and negative shocks inside trading hours; Panels (c) and (d) illustrate shocks that occur outside of trading hours.

Table 2.2 reports the shocks' descriptive statistics. The first two columns show the HF change in log prices of reporting institutions around their earnings announcements. On average price changes are close to zero, with a standard deviation of 2.4%. The median positive and negative shocks are close to 1%. Shocks that occur inside trading hours exhibit roughly similar descriptive statistics to those occurring outside market hours. The third and fourth columns show descriptive statistics of the HF financial shocks—which, as shown in (2.1), weights each change in log price of reporting institutions by their market share. Weighting overall reduces the magnitude of the shocks, resulting in a standard deviation of 0.23% and median positive and negative shocks of 0.05% and −0.05%, respectively.

Summary statistics on shocks

2.3.2 Characterization of the HF financial shocks

The empirical analysis proposed in this paper relies on two properties of HF financial shocks.

Granularity Gabaix and Koijen (2020) discuss how the idiosyncratic shocks of large players in the economy that affect aggregates constitute powerful instruments. The first property of the HF financial shocks is that they are based on the change in the net worth of large

Table 2.2: Financial shocks

	Changes in Stock Prices		**HF Financial Shocks**	
	Broad	**Narrow**	**Broad**	**Narrow**
Mean	0.06	−0.16	0.00	−0.03
Median +	1.07	1.22	0.05	0.06
Median −	−1.08	−1.22	−0.05	−0.08
Std Deviation	2.37	2.68	0.23	0.30
5th Percentile	−3.53	−4.59	−0.34	−0.56
95th Percentile	3.65	3.81	0.30	0.31
Observations	870	343	870	343

Notes: This table shows descriptive statistics of financial shocks. "Unweighted" shocks are constructed with changes in stock prices of financial intermediaries as described in the main text. "Weighted" shocks are weighted by the market net worth of the financial intermediary as a fraction of the total market net worth of the sample in the quarter. "Narrow" refers to shocks with earning releases inside of market trading hours, including pre-market and extended trading, if available. "Broad" refers to shocks with earning releases outside of market trading hours. "All" includes all financial shocks. "Median +" and "Median −" refer to median positive and median negative shocks.

financial institutions, whose individual changes in net worth represent a significant change in the net worth of financial intermediaries of the aggregate economy. In particular, Table 2.1 shows that financial intermediaries in our sample represent 64% of the total equity of U.S. depositary institutions, measured by the Federal Reserve's Flow of Funds. In addition, Appendix Figure B.2 shows that shocks to the market value of an intermediary releasing earnings leads to a 0.2% increase in the market value of other financial intermediaries in the sample that do not disclose their earnings, suggesting that the change in the market value of intermediaries disclosing earnings do not lead to opposing forces in other intermediaries not disclosing their earning that could offset the effects of the shocks.

Financial content The second important property of financial shocks for our empirical analysis is that they primarily reflect information of financial intermediaries and not about the rest of the economy where we want to study their effects. Three pieces of evidence support this view. First, Appendix B.2.1 uses data on unexpected earnings in announcements to show that the stock price movement from financial institutions tend to be positively associated with their surprise earnings, suggesting that financial shocks encode information of earnings released in the announcements.

Second, financial shocks are not systematically linked to information available at the moment of earnings releases. Appendix B.2.2 uses a state-of-the-art machine-learning model and shows that the HF financial shocks are not predictable with existing macroeconomic or financial data available before the shocks, suggesting that financial shocks are not driven by information in the rest of the economy available before intermediaries' earnings have been released.

Third, Table B.4 reports the volatility of the stock price of financial intermediaries and nonfinancial firms in events of intermediaries earnings announcements and comparing with those of nonevent days. These moments show that the volatility of financial intermediaries' stock prices during their earnings announcements increases by substantially more than those of nonfinancial firms during these events, which is consistent with intermediaries' earnings announcements reflecting more information about financial intermediaries than about nonfinancial firms. Based on this, in our empirical analysis of the next section, we conduct a heteroskedasticity-based identification, which can be conducted even if factors of nonfinancial firms affect the market value of intermediaries during their earnings announcements, as long as the variance of intermediaries' stock price during earnings-announcement events is larger than in nonevent dates, while those of nonfinancial firms are the same during events of earnings releases of financial intermediaries and nonevent dates.

2.4 The effects of financial shocks on the aggregate economy

We now document the effect of financial shocks in the aggregate economy. Section 2.4.1 shows their effects on the market value of firms and Section 2.4.2 their effects on corporate borrowing costs.

2.4.1 Effects of financial shocks on the market value of nonfinancial firms

Event-study approach Our first empirical strategy to study the effect of financial shocks on the market value of nonfinancial firms is an event-time study, which consists of estimating the regression

$$\Delta \log y_t = \alpha + \beta \varepsilon_t^{\mathrm{F}} + \varepsilon_t, \tag{2.2}$$

where y_t denotes a stock-market index for nonfinancial firms at closing date t of financial intermediaries' earnings announcements, and ε_t is a random error term. The coefficient of interest, β, measures the elasticity of the market value of nonfinancial firms to financial shocks. The identifying assumption to interpret these effects as causal is that, in the 60-minute window around intermediaries' earnings announcements, changes the stock price of intermediaries releasing earnings are driven by information contained in these announcements and not by other factors affecting stock prices of nonfinancial firms in an announcement-date event, contained in ε_t.

Panel (A) of Table 2.3 shows the results of estimating the baseline event-study regression (2.2). The first column shows the effects on nonfinancial firms included in the SP500, with an estimated elasticity of 0.7. The second and third columns show the effects on the SmallCap 600 and Russell 2000, showing an elasticity of 1.1, indicating that the effects of financial shocks are larger for small firms. Table 2.3 shows that these effects are robust if we include in the regression only dates in which the announcements were made within trading hours and if we include a set of macroeconomic controls in (2.2).

We conduct two "placebo" exercises to provide further evidence for our interpretation of these event-time results. The first exercise shows that the effects we identify for financial shocks are not found if we follow a similar procedure to identify shocks originating in nonfinancial firms. To conduct this exercise, we follow a similar procedure to that developed in

Table 2.3: Effects of financial shocks on the market value of nonfinancial firms

		SP500 Ex-Fin	**SmallCap**	**Russell**	**Obs**
A. Event-Study	Baseline	0.720***	1.085***	1.124***	486
		(0.179)	(0.213)	(0.229)	
	Narrow Windows only	0.924***	1.348***	1.453***	272
		(0.241)	(0.296)	(0.313)	
	Macro Controls	0.870***	1.271***	1.373***	272
		(0.242)	(0.297)	(0.314)	
B. Identifcation through Heteroskedasticity	Baseline	0.403***	0.477***	0.498***	3,532
		(0.026)	(0.028)	(0.029)	
	Macro Controls	0.399***	0.472***	0.492***	3,532
		(0.027)	(0.029)	(0.029)	
C. By Aggr Net Worth	Well-Capitalized	0.259	0.447	0.398	135
		(0.483)	(0.541)	(0.599)	
	Under-Capitalized	1.059***	1.535***	1.673***	137
		(0.306)	(0.390)	(0.405)	
D. By Institution	Commercial Banks	0.918***	1.399***	1.456***	336
		(0.227)	(0.267)	(0.287)	
	Invm Banks & Security Dealers	0.329	0.463	0.462	252
		(0.328)	(0.382)	(0.415)	
E. By Dealer Status	Primary Dealers	0.241	0.492	0.547	308
		(0.281)	(0.324)	(0.349)	
	Nonprimary Dealers	1.183***	1.684***	1.719***	313
		(0.264)	(0.310)	(0.333)	

Notes: Panel (A) shows the results from estimating $\Delta \log y_t = \beta_0 + \beta_1 \varepsilon_t^{\mathrm{F}} + \Gamma_t' X_t + \varepsilon_t$, where y_t denotes the price of a stock index (Nonfinancial firms in the SP500, SmallCap 600, and the Russell 2000), "Baseline" shows the results without controls X_t, "Narrow window only" shows the results for regressions that include only the announcements made during trading hours, "With controls" shows the results including a set of monthly macroeconomic controls in the vector X_t (industrial production and employment). Panel (B) shows the results for identification by heteroskedasticity (following Rigobon, 2003; Rigobon and Sack, 2004), described in Section 2.4.1.Confidence intervals were obtained following the procedure in Hébert and Schreger (2017). First-state F-statistics are 70.26, 66.88, and 24.76 for baseline, including controls (S&P500, VIX, industrial production, and payrolls), and using S&P 500 nonfinancial firm earnings releases as nonevents, respectively. Panels (C) shows the results from estimating the event-study model for dates in which the financial system exhibits different degrees of capitalization, with well- (under-)capitalized referring to when the HP-filtered series of depository institutions' equity capital is above (below) mean. Panel (D) shows the effects when we decompose the HF financial shocks into those from two types of institutions, commercial banks (NAICS code 522110) and investment banks and securities dealers (NAICS code 523110). Panel (E) shows the effects when we decompose the HF financial shocks into those from institutions with different dealership status: primary dealers and nonprimary dealers (Data source: Federal Reserve Bank of New York).

Section 2.3 for financial shocks, but focus on earnings announcements of nonfinancial firms included in the Dow Jones index. Table B.5 shows the results from estimating (2.2) but using the shock to nonfinancial firms instead of the financial shock, $\varepsilon_t^{\mathrm{F}}$. Results indicate a baseline estimate that is not statistically significant and is unstable across specifications (e.g., has a negative point estimate when we use the narrow version of the shocks, only announcements during trading hours). Furthermore, the point estimates are smaller when we study their impact in smaller firms, suggesting that the effects that we identify in our empirical model are specific for financial intermediaries.

The second placebo exercise, shown in Appendix Figure B.3, shows that the HF shocks do not have an effect on the market value of nonfinancial firms during the days before the shock, suggesting that the effects are not driven by pretrends. This figure also shows that the HF shocks do not have an impact on the days after the shocks, suggesting that the information in financial shocks are incorporated into the value of nonfinancial firms in the day of the shock and that there are not offsetting forces in consecutive days that revert the impact effects of these shocks.

Heteroskedasticity-based identification One potential concern about the event-time study approach is that factors unrelated to the release of earnings of intermediaries may ultimately be related to the stock prices of nonfinancial firms even within a narrow window around earnings announcements. We address this concern by conducting an alternative estimation based on a heteroskedasticity-based identification strategy (Rigobon, 2003; Rigobon and Sack, 2004). This strategy can be conducted even if other factors of nonfinancial firms do affect the market value of intermediaries during their earnings announcements, as long as the variance of intermediaries' stock price during earnings-announcement events dates is larger than in nonevents dates, while those of nonfinancial firms are the same during earning releases of financial intermediaries and nonevent dates.

To conduct the estimation based on identification through heteroskedasticity, we consider the bivariate model,

$$\Delta\nu_t = \beta_F \Delta x_t + \Gamma'_F X_t + \eta_{Ft}$$
$$\Delta y_t = \beta \Delta\nu_t + \Gamma' X_t + \eta_t,$$

where $\Delta\nu_t$ is the log change in a value-weighted index of intermediaries' stock prices in period t and X_t is a vector of control variables. Unlike the event-time analysis estimating (2.2), the heteroskedasticity-based approach uses data from both dates in which intermediaries release their announcements and those in which they do not. We estimate the coefficient of interest and its confidence interval following the procedure described in Hébert and Schreger (2017). We use 1,000 repetitions of a stratified bootstrap, resampling with replacement from events and nonevents.

Panel (B) of Table 2.3 shows the results from estimating the effects of financial shocks on nonfinancial firms using the heteroskedasticity-based approach. The estimated elasticity is 0.4, which, although smaller than the point estimate from the event-study approach, is statistically significant and economically relevant. Using this approach, we also find that small firms are more impacted, as measured by those in the SmallCap 600 and the Russell 2000.

Additional results Table 2.3 provides three additional results, decomposing the effects of financial shocks on nonfinancial firms. First, Panel (C) shows that the effect of financial shocks on nonfinancial firms are driven by their effects in dates in which the financial system is undercapitalized (i.e., when the market value of intermediaries' net worth is below its HP trend). When the financial system is well-capitalized, the effects of financial shocks on nonfinancial firms are economically small and statistically insignificant. This result indicates that a key component driving the general equilibrium effects of intermediaries in the economy are the overall conditions of the financial system.

Panels (D) and (E) decompose financial shocks by institution type. Panel (D) shows that the effects are stronger for commercial banks (NAICS code 522110) than for investment banks and security dealers (NAICS code 523110). Panel (E) shows that the effects are driven by institutions that are nonprimary dealers. Examples of nonprimary commercial banks in our sample include Ameriprise, Bank of America, Bank One, BankBoston, Fleet Boston, Keycorp, MBNA, Northern Trust, Bank of New York Mellon, U.S. Bancorp, Wachovia, and Wells Fargo. As these intermediaries play an important role in the lending to nonfinancial firms, these results are consistent with the lending channel playing an important role in the effect of intermediaries on firms. The next section further analyzes this channel by studying the effect of financial shocks on the corporate bond market.

2.4.2 Effects of financial shocks on corporate borrowing costs

We study the effect of financial shocks on corporate bond spreads by estimating the Jordà (2005) local projections,

$$\Delta_h z_{ijt} = \alpha_{hj} + \beta_{hj} \varepsilon_t^{\mathrm{F}} + \epsilon_{ijth}, \tag{2.3}$$

where z_{ijt} is the spread of bond i with credit rating j in period t of earnings announcements, and $h \geq 0$ indexes days after these announcements. As described in Section 2.2, for this analysis, we use daily individual bond-level data from the constituents of Intercontinental Exchange Bank of America (ICE BofA). The coefficient of interest, β_{hj}, measures the semielasticity of corporate bonds of rating j to financial shocks at horizon h.

Figure 2.1 reports the effects of estimating (2.3) for different credit ratings and time horizons. Panel (a) shows that financial shocks have no significant effects on the option-adjusted spreads of AAA-rated bonds. Panel (b) shows that financial shocks have a persistent effect on the spreads of bonds with a rating CCC or lower rating. The magnitude of these effects indicate that a one-percent decline in the market value of intermediaries' net worth leads to 0.1–0.3-percentage-point increases in spreads in the month following an event.

Figure 2.1: Effects of financial shocks on corporate bond spreads

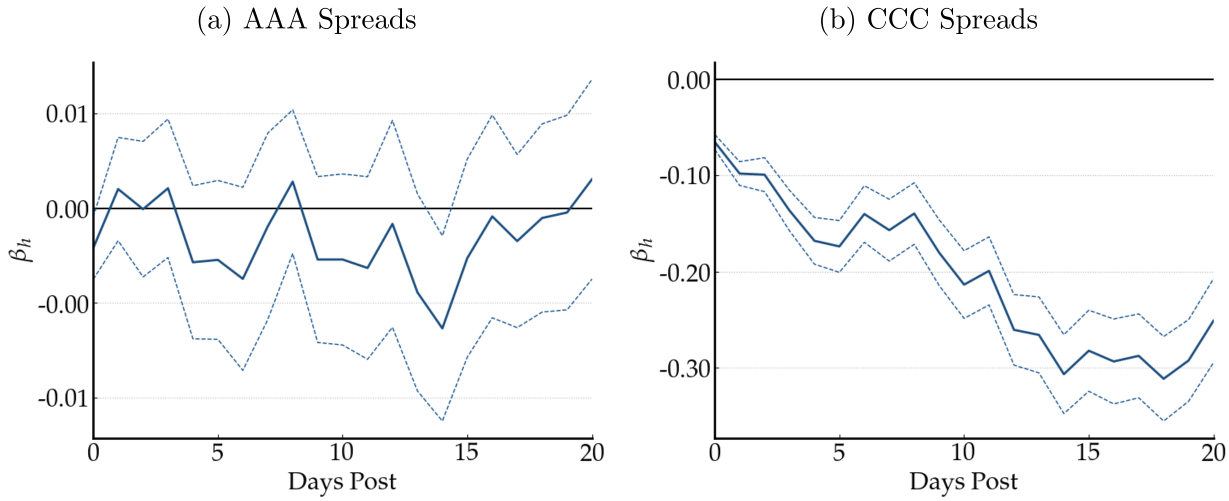

Notes: The figures show the estimated cumulative responses, β_h, of bond spreads to HF shocks at horizon h from estimating local projections

$$\Delta_h z_{ijt} = \alpha_{hj} + \beta_{hj}\varepsilon_t^{\mathrm{F}}.$$

The left panel reports the responses of nonfinancial constituent bonds in the AAA index, and the right panel reports the responses of nonfinancial constituent bonds in the CCC or Lower index. Solid lines represent point estimates of the local projection at each horizon, and dotted lines represent 90% confidence intervals.

To further understand the mechanisms through which financial shocks affect the spreads of corporate bonds Table 2.4 uses data on bond holdings of individual financial intermediaries. The first column of the table reports the impact effect, which is the same as that reported in Panel (B) of Figure 2.1 for $h = 0$. The second column of the table reports the effects for bonds that are held by financial intermediaries reporting earnings announcements and those that are not held by these intermediaries. Results are larger for bonds for which intermediaries releasing earnings have positive holdings, which is consistent with the view that financial institutions acquire specialized knowledge about certain bonds for trading purposes and that, in the presence of trading frictions, financial shocks have larger impact in the prices of these bonds (see, for example, Morelli et al., 2019). However, results also indicate substantial effects for bonds that are not held by financial intermediaries releasing information, with an impact semielasticity of -0.07, suggesting the importance of aggregate net worth of financial intermediaries affecting corporate borrowing costs (as stressed, for instance, in Gertler and

Table 2.4: Effects of financial shocks on corporate bond spreads
by intermediaries' individual bond holdings

	(1)	(2)
Fin shock	-0.069***	
	(0.010)	
Fin shock $\times \mathbb{1}(e_{k,it-1} > 0)$		-0.123***
		(0.034)
Fin shock $\times \mathbb{1}(e_{k,it-1} = 0)$		-0.064***
		(0.011)
Observations	15231	15231
R^2	0.113	0.113
Firm fixed effect	✓	✓

Notes: This table shows the effects of a bond's exposure to an intermediary on how spreads respond to the shocks generated by the intermediary. The specification in Column (2) takes the form:

$$\Delta cs_{kj(i)t} = \delta_i + \beta_1 \cdot \nu_t \cdot \mathbb{1}(e_{kj(i)t-1} > 0) + \beta_2 \cdot \nu_t \cdot \mathbb{1}(e_{kj(i)t-1} = 0) + \varepsilon_{kj(i)t},$$

where $\Delta cs_{kj(i)t}$ is the change in spreads of bond j issued by firm i around the earning release of bank k in day t; δ_i is a firm fixed effect; and $e_{kj(i)t-1}$ is the share of bond j issued by firm i that is held by bank k, measured at the previous quarter of the earning release. Control variables include industrial production, nonfarm payrolls, recession indicator, bond remaining maturity, average spreads in the previous 30 days, and month-to-date changes in spreads.

Kiyotaki, 2010).

2.5 Conclusion

In this paper we proposed a new measure of financial shocks, based on the HF change in the market value around intermediaries' earnings announcements. We then exploited the "granularity" of financial shocks, stemming from the fact that U.S. publicly traded financial intermediaries have considerable size, to study the effects of financial shocks in the aggregate economy. We document intermediaries' substantial effects on the market value and borrowing costs of nonfinancial firms. The effects are stronger for small firms and when the financial system is undercapitalized.

The HF financial shocks developed in the paper can be directly used by researchers conducting empirical research in macroeconomics, similarly to the large body of evidence developed using HF monetary-policy shocks. Our empirical findings about the effect of intermediaries in the aggregate economy can also be useful to be combined with macrofinance models aimed at understanding role of financial frictions determining the aggregate transmission of shocks. We leave the combination of models with these empirical estimates for future research.

Chapter 3

Economic Narratives and Consumer Sentiment: Evidence from Twitter

Wenting Song

3.1 Introduction

Political narratives shape ideological diversity (Gentzkow et al., 2014), political polarization (Levy, 2021), and asset prices (Bianchi et al., 2021). Less studied, however, is the role of economic narratives. As Shiller (2020) points out:

> We need to incorporate the contagion of narratives into economic theory. Otherwise, we remain blind to a very real, very palpable, very important mechanism for economic change, as well as a crucial element for economic forecasting.

This paper develops an empirical framework to study the role of economic narratives. The empirical analysis is motivated by Shiller (2017), who models the contagion of economic narratives with an epidemic SIR model developed by Kermack and McKendrick (1927):

$$\frac{\mathrm{d}S}{\mathrm{d}t} = -cSI \tag{3.1}$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = cSI - rI \tag{3.2}$$

$$\frac{\mathrm{d}R}{\mathrm{d}t} = RI \tag{3.3}$$

The total population, N, is divided into three compartments: S, the number of "susceptibles" who are vulnerable to a narrative, I, the number of "infected" who believe in a narrative and can further convince their friends of it, and R, the number of "recovereds" who have moved on from the narrative. Whether a narrative becomes "viral" depends on two parameters: the contagion rate $c > 0$ and the recovery rate $r > 0$.

The paper provides evidence on the empirical importance of economic narratives and also estimates the values of the contagion rates c.

I start by developing a procedure to capture economic narratives. In Eliaz and Spiegler (2020), narratives are defined as causal models (directed acyclic graphs) that map actions into outcomes, while weaving other random variables into the story. In the context of media-created macroeconomic narratives, I define "narratives" as the media's interpretations of how policies affect macroeconomic outcomes. Applying topic models from natural language processing on the news articles devoted to an economic event, I obtain empirical estimates of both the prevailing narratives and each article's reliance on the narratives.

I then study the empirical importance of the identified narratives and estimate the extent of contagion. Social media provides an ideal platform to monitor a narrative's popularity and influence. The microblogging service Twitter, in particular, is widely used for news dissemination and commentary. I use quote retweeting activities on Twitter to measure users' exposure to a given narrative and trace the sentiment changes displayed in users' tweets after their exposure.

Applying the framework to study the 2019 yield curve inversion, a recession indicator in the US (Harvey, 1988), I show that economic narratives significantly affect consumer sentiment: consumers exposed to the negative narrative of an imminent recession display a more pessimistic outlook, while consumers exposed to the positive narrative of recession concerns being overblown experience no change in their sentiment.

Related literature The paper is related to three strands of the literature. First, it is most closely related to the literature on economic narratives, which have received several different formalizations. Eliaz and Spiegler (2020) define narratives as causal models which map actions into outcomes. Shiller (2017) defines narratives as catchy stories, measures the popularity of narratives using Google search frequencies, and models narratives using an epidemiology model. Bénabou et al. (2018) study moral decision making and define narratives as signals that affect agents' beliefs on the externality of their actions. This paper focuses on empirical testing for the importance of narratives. The definition of narratives follows Shiller (2017) and Eliaz and Spiegler (2020), both of whom represent narratives as models that piece together actions, outcomes, and a selected subset of data. Diverging from the empirical approach in Shiller (2017), the framework proposed in this paper does not require researchers to identify the narratives ex ante and, therefore, allows for real-time detection and tracking of narratives.

Second, the paper relates to a growing literature which uses natural language processing to study the economic effects of news, such as Bybee et al. (2020), Nyman et al. (2021) and Calomiris and Mamaysky (2019). In particular, Larsen and Thorsrud (2019) study the effects of narratives on business cycle fluctuations, defining narratives as significant economic events that are extracted empirically using topic models on the corpus of Wall Street Journal articles. The differences between this paper and previous ones stem from the different theoretical definitions of narratives. Narratives in this paper are news media's different interpretations of the *same* underlying economic event, extracted using topic models on the articles on the given event. In addition, the data of retweeting activities provides estimates of the contagion rate of a narrative.

Third, the paper relates to the empirical literature of belief formation, which points to personal experiences (Malmendier and Nagel, 2016), salience (Cavallo et al., 2017) and heuristics (Bordalo et al., 2018), among others, as important drivers of individuals' expectations. This paper contributes to the literature by highlighting the role of economic narratives in shaping sentiment.

Outline The rest of the paper lays out the empirical framework using the yield curve inversion as an example. Section 3.2 describes the yield curve inversion event. Section 3.3 formally defines narratives in the context of the inversion. Section 3.4 conducts the empirical analysis. The empirical section starts with the specification of the empirical model in Section 3.4.1. Then it describes the data sources and empirical approaches used to construct each component of the empirical model: Section 3.4.2 introduces the topic model and data sources used to capture narratives; Section 3.4.3 describes the measure of narrative exposure based on retweeting activities; Section 3.4.4 describes the measure of consumer sentiment obtained using a naïve Bayes classifier. Section 3.4.5 assembles the components and reports the empirical results. Section 3.5 concludes.

3.2 Yield curve inversion

Harvey (1988) connects expected consumption growth to the term spread of yield curve and documents the predictive power of the yield curve inversion for major recessions from the 1960s to the 1980s. Since then, the yield curve inversion has been a closely-watched and reliable recession indicator in the US. Figure C.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.

As expected, the yield curve inversion in 2019 has attracted a lot of attention. Figure 3.1b shows that Google searches for the term "yield curve" spiked before and during the inversions

Figure 3.1: Timeline of the yield curve inversion episode

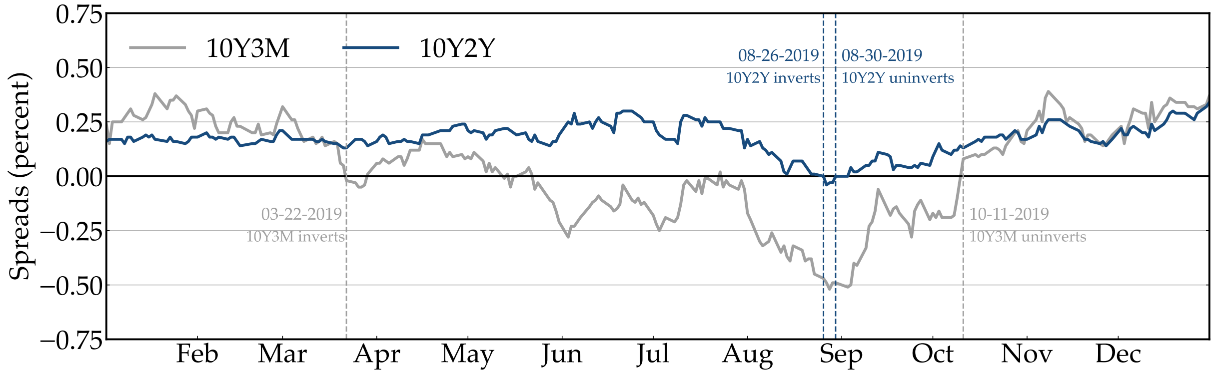

(a) Treasury spreads

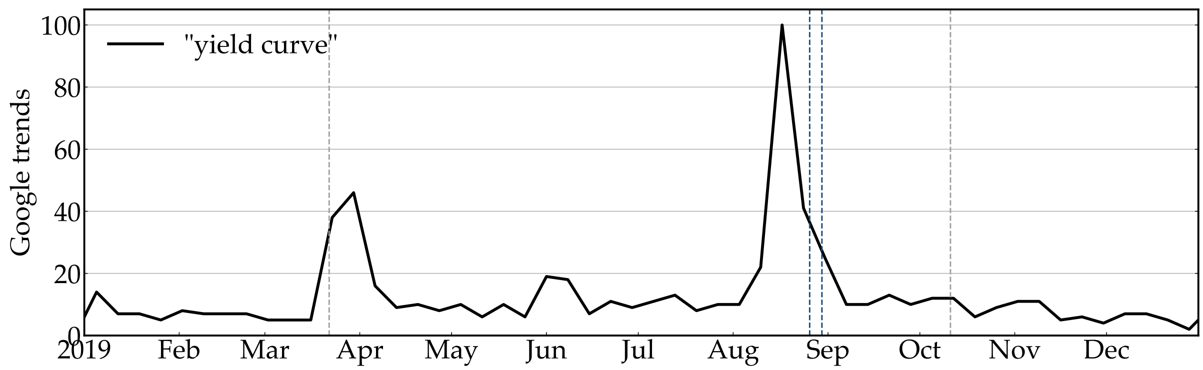

(b) Google searches

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 Google search frequency for the term "yield curve" in 2019. The highest search frequency has been scaled to be 100.

of both the 10-year-over-2-year (10Y2Y) term spread and the 10-year-over-3-month (10Y3M) term spread, with a peak of the searches right before the inversion of the most wildly-watched 10Y2Y spread.

Against the backdrop of a booming labor market and the longest expansion in US history, the yield curve inversion in 2019 has 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[1]:

[1] "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.

> 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. The second interpretation is that the yield curve inversion is no longer an informative signal. Peter Coy illustrates such a narrative for Bloomberg[2]:

> 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. The third is a neutral coverage, providing both sides of the previous narratives. An example of such coverage is Brian Chappatta's Bloomberg article[3], explaining 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?

which defines an inverted yield curve, explains its history of proceeding recessions, but does not draw any 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? Before answering the questions empirically, I first formally define "narratives" within the context of the yield curve inversion.

[2] "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.

[3] "The Yield Curve Is Inverted! Remind Me Why I Care" by Brian Chappatta. Link to the article on Bloomberg.

3.3 Theoretical definition of narratives

As in Eliaz and Spiegler (2020), a narrative is represented by a *directed acyclic graph* (DAG). Acyclicity means that the graph contains no directed path from a node to itself and therefore can be interpreted as a causal model. A news reporter forms the narrative of how the economic policy, a, influences the economic outcome, y. Both a and y take binary values: $a = 0$ (1) denotes the current (counterfactual) policy, and $y = 0$ (1) denotes a recession (an expansion). The reporter has the option to weave the yield curve inversion into the narrative. Let s denote the slope of the yield curve. $s = 0$ represents an inverted yield curve, while $s = 1$ represents the common upward-sloping yield curve.

An objective narrative, G^1, points out that the slope of the yield curve is influenced by a variety of factors, θ, such as economic fundamentals, investor expectations, and market liquidity. These economic fundamentals are what ultimately determine the outcome y. Such a narrative can be represented by the DAG:

$$G^1 : \quad a \rightarrow s \leftarrow \theta \rightarrow y$$

Under the narrative G^1, there is no direct causal link between the inverted yield curve and a recession. On the other hand, a sensationalized narrative connects the inverted yield curve to an imminent recession. Such a narrative can be represented by the DAG:

$$G^2 : \quad a \rightarrow s \rightarrow y$$

Under the narrative G^2, the current economic policy is directly responsible for the inversion of the yield curve, which will lead to a recession.

Consumers who are exposed to these narratives would then form different subjective beliefs. Let $x = (a, s, \theta, y)$ denote the collection of the aforementioned variables. Readers of G^1 recognize that the inverted yield curve is merely a symptom of economic fundamentals and policies. To them, the inverted yield curve does not necessarily predict a recession:

$$p_{G^1}(x) = p(a)p(\theta)p(s|a,\theta)p(y|\theta)$$

In contrast, readers of G^2 perceive a direct causal link between the yield curve inversion and a recession:

$$p_{G^2}(x) = p(a)p(s|a)p(y|a,s)$$

Appendix C.2 makes functional form assumptions on payoffs and long-run probabilities to illustrate how G^1 and G^2 can co-exist in the steady state. Given the historical correlation between the yield curve inversions and recessions, the sensationalized narrative is not inconsistent with the data. Its causal structure predicts a better future if consumer switch to the counterfactual policy, promising higher expected utility under the subjective belief. As more consumers switch to the counterfactual policy, the gap between the current and counterfactual policies shrink, diminishing the sensationalized narrative's ability to promise a higher payoff, which leads to an equilibrium policy under which both narratives co-exist.

The general definitions follow the Bayesian network literature (Spiegler, 2016; Eliaz and Spiegler, 2020). First, I define policy and outcomes, which are components of a narrative.

Definition 3 (policy and outcome)**.** *Let $X = X_1 \times \cdots \times X_m$ be a finite set of states with $X_i = \{0, 1\}$ for each $i = 1, \cdots, m$ with $m > 2$. For every $N \subseteq \{1, \cdots, m\}$, denote $X_N = \prod_{i \in N} X_i$. For any $x \in X$, x_1 and x_m, denoted a and y, are the economic policy and the economic outcome.*

In the example of the yield curve, $m = 4$ and $X = X_1 \times X_2 \times X_3 \times X_4 = A \times S \times \Theta \times Y$. The full set of states X_{N_1} where $N_1 = \{1, 2, 3, 4\}$ contains the element $x_1 = (a, s, \theta, y)$, and a partial set of states X_{N_2} where $N_2 = \{1, 2, 4\}$ contains the element $x_2 = (a, s, y)$. Both x_1 and x_2 contain the policy, a, and the outcome, y.

Definition 4 (narrative)**.** *A directed acyclic graph (DAG) is a pair of nodes and directed links $G = (N, R)$, where $N \subseteq \{1, \cdots, m\}$ is a set of nodes and $R \subseteq N \times N$ is a set of directed links. Acyclicity excludes directed path from a node to itself. Let $\mathcal{G}$ be a collection of DAGs. Then a narrative is an element $G \in \mathcal{G}$ satisfying:*

1. *$\{1, m\} \subseteq N$, that is, all feasible narratives contain policies and outcomes;*
2. *$|N| \leq n$ where $n \in \{2, \cdots, m\}$, that is, narratives have a maximum complexity;*
3. *1 is an ancestral node, that is, the policy has no prior causes.*

Definition 5 (subjective belief)**.** *Let iRj denote a directed link from node i to the node j, and let $R(i) = \{j \in N \mid jRi\}$ denote the set of prarent nodes of i. Suppose p is the objective distribution. Then, the subjective belief over X_N induced by the narrative $G = (N, R)$ is given by:*

$$p_G(x_N) = \prod_{i \in N} p(x_i | x_{R(i)}). \tag{3.4}$$

Both aforementioned narratives of the yield curve inversion satisfy Definition 4. The narrative $a \rightarrow s \leftarrow \theta \rightarrow y$ can be represented with the DAG: $G^1 = (N^1, R^1)$, where the nodes $N^1 = \{1, 2, 3, 4\}$ and the links $R^1 = (1R2, 3R2, 3R4)$. The sensationalized narrative $a \rightarrow s \rightarrow y$ can be represented by the DAG: $G^2 = (N^2, R^2)$, where $N^2 = \{1, 2, 4\}$ and $R^2 = (1R2, 2R4)$.

G^1 treats the inverted yield curve as an exogenous factor which is not a consequence of the economic policy and only one of the factors determining the economic outcome. Under this narrative, the recession risk is overblown. Therefore, I refer to it as the "overblown" narrative:

$$G^{\mathsf{ovb}}: \quad \text{current economic policy} \ \rightarrow \ \text{yield curve inversion} \leftarrow \text{other factors} \rightarrow \text{recession}$$

The coverage by Bloomberg's Peter Coy is an example of such narrative. It states that even though the yield curve has inverted, the economic fundamentals are solid, which implies a low recession risk. Readers who are exposed to the overblown narrative form no causal link between the inversion and the recession. Their subjective beliefs of the event follow $p_{G^{\mathtt{ovb}}}(x) = p(a)p(\theta)p(s|a,\theta)p(y|\theta)$.

In contrast, G^2 weaves in an inverted yield curve as being causal of a recession. Under this narrative, the inversion implies an imminent recession. Therefore, I refer to it as the "recession" narrative:

$$G^{\mathtt{rec}}: \quad \text{current economic policy} \rightarrow \text{yield curve inversion} \rightarrow \text{recession}.$$

The coverage by CNN's Cristina Alesci posits the yield curve inversion as a signal of recession, and therefore establishes a direct causal link from the inverted yield curve to heightened recession risks. Readers exposed to the recession narrative develop the subjective beliefs $p_{G^{\mathtt{rec}}}(x) = p(a)p(s|a)p(y|a,s)$. They are expected to develop a more pessimistic economic outlook since they connect yield curve inversions to recessions.

3.4 Empirical analysis

Given the definitions of narratives, I now study the importance of narratives for consumer sentiment and measure the rate of their contagion. The empirical model takes the form of a high-frequency event study. Data needed to conduct the empirical analysis draws on unstructured sources. I describe data construction after specifying the empirical model and then discuss the results.

3.4.1 Empirical model

The main hypothesis that I test for is whether a narrative of the yield curve inversion affects consumer sentiment. For consumer i who has read news article d with narrative

$k \in \{\text{recession}, \text{overblown}\}$, the baseline model is:

$$\Delta\texttt{sentiment}_{id} = \alpha + \beta \cdot \texttt{Narrative}_d^k + \varepsilon_{id}. \tag{3.5}$$

The dependent variable is the change in consumer sentiment after exposure to a narrative. To isolate the effect of the narrative, I focus on the high-frequency changes in consumer sentiment 24 hours before and after the exposure. 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 is collapsed. The next few subsections discuss the data sources and data construction of the inputs into the baseline model, including measures of a narrative, an article's loading on a narrative, and the consumer sentiment. The parameter of interest is β which estimates the effect of a narrative on consumer sentiment. Additionally, the fraction of consumers with exposure to each narrative provides an estimate of its contagion rate.

Identifying assumptions The empirical model specified in (3.5) contains three underlying 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 a news subscription is uncorrelated with unobservable factors affecting the changes in sentiment within the high-frequency window. The assumption allows for pre-existing differences in sentiment. An obvious unobservable factor that can potentially influence how consumers respond to the yield curve is political affiliation. I study news media with center political placement to focus on consumers with similar political views. Additionally, the high-frequency approach allows for isolating the effect of the exposure to a narrative.

The third assumption concerns the direction of causality. The measure of exposure to a narrative that I use is retweeting. The implicit assumption is that retweeting implies the absorption of new information. However, consumers might selectively retweet articles that

Table 3.1: Media outlets and coverage on the yield curve inversion

Outlet	Pew ideology placement	Twitter handle	# 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

Notes: Media outlets with centerist political leaning and their coverage of the yield curve inversion. Data source for 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 handle of news outlets are hand searched. The tweets and articles on the yield curve are collected as described in the main text in Section 3.4.2.

confirm their existing agenda. I conduct robustness checks which impose a limit on the number of outlets that can appear in a user's timeline to remove such users from the sample.

3.4.2 Measuring economic narratives

The first input we need is an empirical measure of narratives. To match the theoretical definition in Section 3.3, I empirically measure narratives as media outlets' different interpretations of the yield curve inversion. Bloomberg's "rainbows and pots of gold," for example, interprets the inverted yield curve as one of many factors in determining the economic outlook, while CNN's "fact checking Navarro" interprets a direct causal link from the current policy to the yield curve inversion to an imminent recession. The third neutral coverage by Bloomberg's Brian Chappata can be empirically interpreted as a mix between the two narratives, giving a balanced representation of both sides of the story.

Data source To measure the narratives created by the news media on the yield curve inversion, I first collect news coverage of the inversion of the 10Y2Y spread. To separate the effects of economic narratives from political narratives, I focus on the news outlet rated as “center” by the Pew Research Center (Jurkowitz et al., 2020)[4] and exclude news aggregators such as Google News. Even though the event window of the inverted yield curve is from August 26 to August 30, Google search trends in Figure 3.1b suggest that the interests in the yield curve rose before the actual inversion and stayed elevated after the un-inversion. Therefore, I expand the search window for news articles to be from August 19 to September 13, one week before the inversion and two weeks after the un-inversion, respectively.

Table 3.1 contains the list of media outlets included in the sample and their coverage of the yield curve inversion. As an illustration, Bloomberg tweeted on August 24:

> Here’s what the yield curve says about when the next recession could happen `https://t.co/DF4OlRpwIX?amp=1`

which contained a link to an in-depth analysis by Lauren Leatherby and Katherine Greifeld posted on `Bloomberg.com` on August 15. The number of tweets reported in Table 3.1 is the number of “base tweets” related to the yield curve[5], illustrated by the August 24 Bloomberg base tweet. The number of articles reported in Table 3.1 is based on the articles linked from Twitter to the outlets’ websites, such as the one written by Leatherby and Greifeld. Since outlets typically tweet to direct traffic to their websites, I focus on the linked articles rather than the tweets to study economic narratives.

Methodology The main tool I use to extract economic narratives from the news articles is Latent Dirichlet Allocation (LDA) developed by Blei et al. (2003) for natural language processing.

LDA is a Bayesian factor model[6] designed to uncover topics in the articles and represent 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.

[4]The methodology used in Jurkowitz et al. (2020) to determine the political bias of a media outlets is through surveying the political ideology of its audience.

[5]Specifically, the criteria is that a tweet needs to contain both “yield curve” and any of the stems from “invert”, “invers” or “recession”. The search window is from August 19 to September 13.

[6]Detailed description of LDA and its application to economics can be found in Hansen et al. (2018).

LDA is a suitable tool for capturing narratives for two reasons. First, as an unsupervised learning algorithm, it does not require researchers to assign labels to the observations. The few inputs required by the algorithm makes it possible to detect economic narratives in real time. 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 provide estimated loadings on each narrative. Therefore, LDA can capture polarizing articles containing a single narrative as well as balanced ones with multiple narratives.

I estimate the LDA with $K = 5$ and symmetric Dirichlet priors[7]. An important feature of the LDA model is that it is a multi-membership model. For example, the word "recession" can appear in multiple topics. Most news articles start with introducing the yield curve inversion as a recession predictor regardless of the narrative. Given the prominent role of the word, it has a high probability of appearing in multiple topics. $K = 5$ is the smallest number of topics that ensures at least one topic does not contain the word recession and allows for capturing different narratives of the yield curve inversion.

The estimated topics from the LDA are shown in Figure 3.2. Two topics, in particular, contain the groupings of words that correspond to the theoretical definitions of the yield curve narratives. The first topic in Panel (3.2a) features the terms such as "recession," "yield curve," "economy" and "Trump," mapping naturally to the "recession" narrative $G^{\texttt{rec}}$. It discuss the economic policy by the Trump administration in conjunction with the yield curve inversion and recession risks. The second topic in Panel (3.2b) contains a broader discussion of other factors affecting the economy and investment opportunities in the bond and stock markets, mapping to the "overblown" narrative $G^{\texttt{ovb}}$.

Unsurprisingly, the word "recession" appears in all but one topics. To match the theoretical definitions of narratives, I focus on the topic in Panel (3.2a) to represent the "recession" narrative as it has the highest probability of the word appearing. The remaining three estimated topics are reported in Figure 3.2 for completeness.

[7]The pre-processing of texts includes removing stop words and numbers, lemmatizing, and representing the documents with a bigram model.

Figure 3.2: Economic narratives of the yield curve inversion: LDA outputs

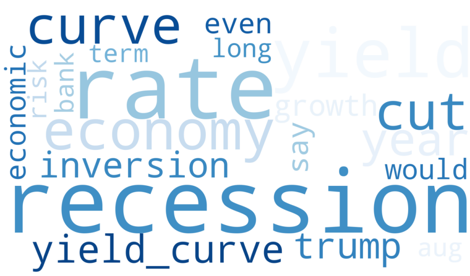

(a) “Recession” narrative

(b) “Overblown” narrative

(c) Other estimated topics

Notes: Results from estimating the LDA on articles of 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.

The model performs well in capturing the narratives conveyed in news articles. To illustrate, for Peter Coy’s article discussed in Section 3.2 which states that the yield curve has lost its predictive power and that “it’s still rainbows and pots of gold,” the model estimates a loading of $\theta^{\texttt{ovb}} = 0.96$ on the overblown narrative and $\theta^{\texttt{rec}} = 0.01$ on the recession narrative. In contrast, for Cristina Alesci’s article that fact-checks Peter Navarro and emphasizes the recession risks, the model estimates a loading of $\theta^{\texttt{rec}} = 0.84$ on the recession narrative and $\theta^{\texttt{ovb}} = 0.05$ on the overblown narrative. For the neutral coverage by Brian Chappata which introduces the yield curve, the model produces more balanced loadings of $\theta^{\texttt{rec}} = 0.67$ and $\theta^{\texttt{ovb}} = 0.11$.

Based on the LDA outputs, I construct two measures of narratives conveyed in an article. The first measure is θ_d^k, the loading of article d on narrative k estimated with the LDA, where k is either the recession narrative or the overblown narrative. The second measure is a binary measure $\mathbb{1}(\theta_d^k > \frac{1}{D}\sum_{d\in D}\theta_d^K)$, which takes the value 1 if the article loading exceeds the cross-sectional average loading of the narrative and 0 otherwise.

3.4.3 Measuring consumer exposure to narratives

The second input needed to estimate the empirical model is a measure of consumer exposure to narratives. The microblogging service Twitter provides rich data on the network among users and on the dissemination of information.

Twitter provides four ways of interacting with posted tweets: quote retweet, retweet, reply and favorite. 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, I use quote retweets as the main measure of exposure to narratives.

I collect information on the users who have quote retweeted the base tweets using Twitter's API[8], which provides the list of first 100 (quote) retweeters of any tweet. Table 3.2 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 100 constraint. Therefore, the information collected through Twitter's API provides a representative sample of the users who have read the base tweets.

For everyone who has quote retweeted any of the base tweets on the yield curve, I then collect tweets from their timelines. The goal is to study the changes in user sentiment in response to the exposure to economic narratives. The event window in the baseline empirical model is of the high frequency of 24 hours around the time of quote retweet. In addition to the tweets posted in the narrow windows, I also collect tweets up to one month around the time of quote retweet to conduct a set of robustness checks. Specifically, to ensure that retweeting implies the absorption of new information, I exclude users who selectively retweet

[8] `GET statuses/retweets/:id`. Details of the API can be found on Twitter's documentation page.

Table 3.2: Descriptive statistics on 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

Notes: Descriptive statistics of media outlets' tweets about the yield curve inversion between August 19 and September 13, 2019. The table reports the descriptive statistics of the numbers of quote retweets, retweets, replies and favorites of media outlets' tweets.

Table 3.3: Descriptive statistics on 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: Descriptive statistics of users' timelines based on tweets one month before and one month after the quote retweets of the base tweets.

articles to promote their existing agendas by limiting the number of news outlets that can appear in a user's timeline. The summary statistics of users' timelines are reported in Table 3.3.

3.4.4 Measuring consumer sentiment

Based on the tweets from users' timelines collected as described in the previous subsection, I 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 $(t_1, \cdots, t_n)$ as:

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

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. I 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, I 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:

$$\text{tf-idf}_{t,d} = \underbrace{\frac{w_{t,d}}{\sum_{\tau \in d} w_{\tau,d}}}_{\text{term frequency}} \cdot \underbrace{\log \frac{D}{|\{d \in D : t \in d\}|}}_{\text{inverse document frequency}} \tag{3.6}$$

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 I use the naïve Bayes algorithm to classify the sentiment of tweets. Specifically, I represent the probability that a tweet j conveys positive sentiment as a function of the tf-idf-weighted terms $t_1, \cdots, t_n$ of in the tweet:

$$\tilde{p}_j(\text{positive}) = f(t_1, \cdots, t_n) \tag{3.7}$$

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

I pre-train the naïve Bayes classifier using $100,000$ pre-classified tweets in Go et al. (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, I construct two measures of consumer sentiment of consumer i in day t who have quote retweeted the coverage on the yield curve inversion:

$$s_{it} = \frac{1}{J} \sum_j \tilde{p}_j(\text{positive}) \tag{3.8}$$

$$d_{it} = \frac{1}{J} \sum_j \mathbb{1}(\tilde{p}_j > 0.5) \quad \text{for } j \text{ posted in day } t \tag{3.9}$$

where s_{it} and d_{it} measure the average sentiment of tweets posted by the consumer in a day. The difference is that the sentiment measure used in s_{it} is based on the predicted probability of a tweet having positive sentiment ($\tilde{p} \in (0, 1)$), and that the sentiment measure used in d_{it} is based on the predicted sentiment, with the value 1 representing a positive tweet and 0 representing a negative tweet. The higher the values of s_{it} and d_{it}, the more optimistic a consumer is of the outlook.

3.4.5 Results

To gather the results, I combine the empirical elements and estimate the effects of the recession and overblown narratives:

$$\Delta\texttt{sentiment}_{id} = \alpha_{\texttt{rec}} + \beta_{\texttt{rec}} \cdot \texttt{Narrative}_d^{\texttt{rec}} + e_{id} \tag{3.10}$$

$$\Delta\texttt{sentiment}_{id} = \alpha_{\texttt{ovb}} + \beta_{\texttt{ovb}} \cdot \texttt{Narrative}_d^{\texttt{ovb}} + u_{id}. \tag{3.11}$$

Narratives are captured using LDA on news articles of the yield curve inversion, as described in Section 3.4.2. To measure the narrative conveyed in an article, I use both the continuous measure, θ_d^k, the loading of article d on narrative k, and the binary meausre, $\mathbb{1}(\theta_d^k > \bar{\theta}^k)$, whether the loading of an article on the narrative is above the cross-sectional mean. The dependent variable is the change in consumer sentiment 24 hours before and after being exposed to the narrative. Measures of consumer sentiment, s_{it} and d_{it}, are predicted with the naïve Bayes classifier and described in Section 3.4.4. Whether a consumer is exposed to a narrative is measured using quote retweeting activities on Twitter, as described in Section 3.4.3.

The parameters of interest are $\beta_{\texttt{rec}}$ and $\beta_{\texttt{ovb}}$. They indicate whether a narrative leads a significant change in the sentiment of consumers who are exposed. The coefficients also estimate the contagion rates of a narrative within the 24-hour window.

Table 3.4 reports the results from estimating Equations (3.10) and (3.11). Panel 3.4a indicates that the exposure to the recession narrative leads to a more pessimistic outlook. Column 1 contains the baseline specification with article narrative measured by $\mathbb{1}(\theta_d^k > \bar{\theta}^k)$ and the consumer sentiment measured by Δs_{it}. A consumer who is exposed to an article that emphasizes the recession narratives more than the average articles displays 0.01% more pessimistic sentiment in the 24 hours after absorbing the narrative. The coefficient can be doubly interpreted as the contagion rate of the narrative, which determines how likely it will become "viral." The results of the recession narrative leading to more pessimistic sentiment is robust to the measures of sentiment or narratives, as reported in Columns (2) – (4).

In contrast, Panel 3.4b suggests that exposure to the overblown narrative leads to no significant changes in consumer sentiment. This is unsurprising since the overblown narrative downplays the scenario of a potential recession and conveys that there is no change in the economic fundamentals.

In Table C.1 in the Appendix, I report the robustness check with limit the maximum number of news outlet that can appear in a user's timeline to be 4, the mean number of outlets in the sample. Results are qualitatively and quantitatively similar, with the recession narrative causing a more pessimistic sentiment and the overblown narrative having no effects. Users who selectively retweet to further their agenda are likely fervent supporters of a narrative. By excluding them, I focus on the susceptible population to better estimate

Table 3.4: Effects of economic narratives on consumer sentiment

(a) Narrative: "recession"

Dependent variable: Consumer sentiment	(1) Δs_{it}	(2) Δd_{it}	(3) Δs_{it}	(4) Δd_{it}
$\mathbb{1}(\theta_d^{\texttt{rec}} > \bar{\theta}^{\texttt{rec}})$	-0.01**	-0.06*		
	(0.01)	(0.03)		
$\theta_d^{\texttt{rec}}$			-0.02**	-0.07
			(0.01)	(0.04)
Constant	0.00	-0.00	0.00	0.00
	(0.00)	(0.01)	(0.00)	(0.01)
Observations	352	352	352	352
R^2	0.011	0.008	0.012	0.007

(b) Narrative: "overblown"

Dependent variable: Consumer sentiment	(1) Δs_{it}	(2) Δd_{it}	(3) Δs_{it}	(4) Δd_{it}
$\mathbb{1}(\theta_d^{\texttt{ovb}} > \bar{\theta}^{\texttt{ovb}})$	0.00	0.01		
	(0.00)	(0.03)		
$\theta_d^{\texttt{ovb}}$			0.00	-0.02
			(0.01)	(0.03)
Constant	-0.00	-0.01	0.00	-0.00
	(0.00)	(0.01)	(0.00)	(0.01)
Observations	352	352	352	352
R^2	0.000	0.000	0.000	0.001

Notes: Panel (a) reports the baseline results from estimating the specification in Equation (3.10): $\Delta\texttt{sentiment}_{id} = \alpha_{\texttt{rec}} + \beta_{\texttt{rec}} \cdot \texttt{Narrative}_d^{\texttt{rec}} + e_{id}$. Panel (b) reports the baseline results from estimating the specification in Equation (15) (3.11): $\Delta\texttt{sentiment}_{id} = \alpha_{\texttt{rec}} + \beta_{\texttt{rec}} \cdot \texttt{Narrative}_d^{\texttt{ovb}} + e_{id}$. As described in the main text in Section 3.4.5, $\Delta\texttt{sentiment}_{id}$ is measured as Δs_{it} or Δd_{it} and $\texttt{Narrative}_d^{\texttt{rec}}$ is measured as $\mathbb{1}(\theta_d^k > \bar{\theta}^k)$ or θ_d^k. Standard errors are in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

the contagion rate of a narrative.

3.5 Conclusion

This paper provides empirical evidence that economic narratives influence consumer sentiment. Through illustrating a yield curve inversion event, I define narratives as causal models which relate policy to outcomes. The empirical strategy that I develop is able to capture the creation of narratives and trace their influence. The findings suggest that exposure to the negative narrative of an imminent recession causes consumers to display a more pessimistic sentiment, while exposure to the positive narrative that recession concerns are overblown leads to no change in consumer sentiment.

Appendix for Chapter 1

A.1 Additional tables and figures in Chapter 1

Table A.1 contains the list of keywords used in frequency search under each topic. The keywords are based on `Econoday`, which provides notifications for major economic news and is the service behind `Bloomberg` economic calendar.

Table A.1: Macroeconomic topics and keywords

Topic	**Keywords**
General	economic conditions
Output	GDP, economic growth, macroeconomic condition, construction spending, national activity, recession
Employment	unemployment, JOLTS, labor market, jobless claims, jobs report, nonfarm payroll, ADP employment report, empoyment cost index
Consumption	consumer confidence, consumer credit, consumer sentiment, durable goods, personal income, retail sales
Investment	business inventories, manufacturing survey, factory orders, business outlook survey, manufacturing index, industrial production, business optimism, wholesale trade
FOMC	FOMC, monetary policy, quantitative easing
Housing	home sales, home prices, housing starts, housing market
Inflation	price index, price level, consumer price index, CPI, PMI, PPI, inflation, inflationary, disinflation, disinflationary, hyperinflation, hyperinflationary
Oil	oil prices, oil supply, oil demand

Notes: Dictionary of keywords used in constructed text-based attention measures. Keywords are based on names of macroeconomic releases from `Econoday`, complemented with macroeconomic words and phrases from popular press.

Table A.2 contains the summary statistics of firm characteristics by attention. In this table, a firm is attentive if its prevalence attention to the general topic is nonzero in any year in the sample period.

Table A.2: Summary statistics of firm characteristics by attention

	N	**Mean**	**Median**	**SD**
Inattentive				
Total assets (Millions)	33,277	2,873.36	104.02	35,004.36
Age	33,796	7.78	7.00	4.98
Leverage	32,955	0.35	0.17	0.69
Attentive				
Total assets (Millions)	102,493	7,311.57	538.12	65,274.94
Age	103,312	11.57	10.00	7.37
Leverage	101,981	0.30	0.20	0.46
Total				
Total assets (Millions)	135,770	6,223.78	370.50	59,333.37
Age	137,108	10.64	9.00	7.05
Leverage	134,936	0.31	0.19	0.53

Notes: In this table, a firm is attentive if its prevalence attention to the general topic is nonzero in any year in the sample period. Firm size is measured by the log of total assets, age is measured as the number of years since the firm first appeared in our sample, and leverage is defined as the ratio of total debt to market equity.

A.2 Drivers of firm attention

This appendix investigates firm characteristics that drive firm attention to monetary policy. We study the cross-sectional and time-series variation in firm i's attention by estimating at the annual frequency:

$$\text{Cross-sectional variation:} \quad d_{it} = \delta_t + \delta_j + \beta \cdot x_{it} + \varepsilon_{it} \tag{A.1}$$

$$\text{Time-series variation:} \quad d_{it} = \delta_i + \beta \cdot x_{it} + \varepsilon_{it} \tag{A.2}$$

where x_{it} is one of the firm characteristics discussed in this section. The attention measure throughout this section is d_{it}, the prevalence measure of attention to FOMC news used in our baseline specification. Equation (A.1) includes a time fixed effect, δ_t, and a sector fixed effect, δ_j, at 4-digit NAICS level, to study the drivers of firm attention in the cross section. Equation (A.2) includes a firm fixed effect, δ_i, to study the drivers of firm attention over a firm's life cycle.

A.2.1 Management quality

Table A.3: Attention and firm management

	(1)	(2)
Management	0.0110***	0.0585***
	(0.0039)	(0.0054)
Observations	65392	65393
R^2	0.422	0.756
Sector × Time FE	yes	no
Firm FE	no	yes

Notes: Column (1) reports the estimated coefficient β from $d_{it} = \delta_t + \delta_j + \beta \cdot m_{it} + \varepsilon_{it}$, and Column (2) reports the estimated coefficient β from $d_{it} = \delta_i + \beta \cdot m_{it} + \varepsilon_{it}$, described as in the main text. d_{it} is the prevalence attention to FOMC news, m_{it} is the fraction of board members who have a master degree or above, δ_i is a firm fixed effect, δ_j is a sector fixed effect (4-digit NAICS), and δ_t is a time fixed effect.

Management quality is a part of a firm's infrastructure which determines its information-processing capacity. We obtain data on publicly-traded firms' board members and their education levels from `BoardEx`. Management quality, m_{it}, is measured as the fraction of firm i's board members in year t who have a master degree or above[1].

Table A.3 shows that firms with a highly-educated board are more attentive to monetary policy than their peers. Column (2) shows that a firm is more likely to be attentive when it has a highly-educated board.

A.2.2 Exposure to monetary policy

Prior exposure to monetary policy is a potential driver of firm attention to FOMC news. To measure a firm's exposure to the monetary surprises associated with the FOMC announcement at date τ , we first estimate the sensitivity of its stock prices to prior announcements over a 5-year rolling window using $t \in [\tau - 5\text{yr}, \tau)$:

$$\begin{aligned}
&\text{Baseline model:} && r_{it} = \alpha_{i\tau} + \beta_{i\tau}^{\texttt{baseline}} v_t + \varepsilon_{it} \\
&\text{CAPM model:} && r_{it} - r_t^f = \alpha_{i\tau} + \beta_{i\tau}^{\texttt{capm}} v_t + \beta_{i\tau}^{M}(r_t^M - r_t^f) + \varepsilon_{it} \\
&\text{FF3 model:} && r_{it} - r_t^f = \alpha_{i\tau} + \beta_{i\tau}^{\texttt{ff3}} v_t + \beta_{i\tau}^{1}(r_t^M - r_t^f) + \beta_{i\tau}^{2} SMB_t + \beta_{i\tau}^{3} HML_t + \varepsilon_{it}
\end{aligned}$$

where v_t is the high-frequency monetary shock, and r_{it} is the close-to-close returns of firm i at date t. In addition to the baseline model, we also estimate a stock's sensitivity controlling for the market factor (r^M) and Fama-French 3 factors (r^M, SML and HML), to isolate the sensitivity to monetary policy. We obtain the daily data on factors from Kenneth French's website.

Based on the estimated sensitivity, we then measure a firm's exposure to monetary policy

[1] Degrees counted as master-level or above include: MBA, MS, MSC, MA, JD, MD, MPA, MSE, PHD, and degree names which include "master" or "doctor."

as the absolute values of the beta's:

$$\theta_{i\tau}^{\lambda} = |\beta_{i\tau}^{\lambda}| \quad \text{for } \lambda \in \{\text{baseline, CAPM, FF3}\}$$

Table A.4 shows that firm attention is increasing in the exposure to monetary policy, both in the cross section and over the time series. The relationship is robust to the measures of monetary exposure. In Appendix A.3, we incorporate additional controls for monetary exposure in the baseline specification to show that even though exposure drives a firm's attention, the baseline results of differential monetary transmission by attention is not driven by firms' exposure to monetary policy.

A.2.3 Firm characteristics

Table A.5 reports the relationship between attention and firm characteristics documented in existing literature to be important for monetary transmission (Gertler and Gilchrist, 1994; Cloyne et al., 2018; Ottonello and Winberry, 2020).

We obtain data on firm balance-sheet items from `Compustat`. In this section, we follow the data construction procedures in Ottonello and Winberry (2020) to make our results comparable to existing literature. We define size as the log of total assets, age as the years since first occurrence in the Compustat sample, and leverage as the debt-to-asset ratio. All firm covariates are standardized, and leverage is in addition demeaned to capture the permanent differences across firms. Sample excludes the financial and utility sectors.

Panel A in Table A.5 shows the in the cross section in a given year, firm attention to FOMC news increases with size and decreases with age and leverage. Panel B in Table A.5 shows as a firm grows older and bigger, it is more likely to become attentive to monetary news. The lack of significance of the leverage measure over a firm's life cycle is unsurprising, because the demeaned measure leverage captures the permanent differences across firms.

Table A.4: Attention and exposure to monetary policy

Panel A: Time-sector level

	(1)	(2)	(3)
Exposure (baseline model)	0.0016***		
	(0.0006)		
Exposure (CAPM model)		0.0036***	
		(0.0006)	
Exposure (FF3 model)			0.0036***
			(0.0008)
Observations	74272	73649	72509
R^2	0.034	0.035	0.035
Time × Sector FE	yes	yes	yes
Firm FE	no	no	no
Firm controls	yes	yes	yes

Panel B: Firm level

	(1)	(2)	(3)
Exposure(baseline model)	0.0016***		
	(0.0005)		
Exposure (CAPM model)		0.0012**	
		(0.0006)	
Exposure (FF3 model)			0.0019***
			(0.0006)
Observations	74280	73657	72520
R^2	0.567	0.567	0.560
Time × Sector FE	no	no	no
Firm FE	yes	yes	yes
Firm controls	yes	yes	yes

Notes: Panel A reports the estimated coefficient β from $d_{it} = \delta_t + \delta_j + \beta \cdot \theta_{it}^{\lambda} + \varepsilon_{it}$, and Panel B reports the estimated coefficient β from $d_{it} = \delta_i + \beta \cdot \theta_{it}^{\lambda} + \varepsilon_{it}$, with θ_{it}^{λ} denoting the exposure to monetary policy with $\lambda \in \{\text{baseline, CAPM, FF3}\}$ and constructed as described in the main text. d_{it} is the prevalence attention to FOMC news, δ_i is a firm fixed effect, δ_j is a sector fixed effect (4-digit NAICS), and δ_t is a time fixed effect. Firm controls include age, size and leverage.

Table A.5: Attention and firm covariates

Panel A: Time-sector level

	(1)	(2)	(3)
Size	0.0037***		
	(0.0005)		
Age		-0.0080***	
		(0.0005)	
Leverage			-0.0006*
			(0.0004)
Observations	104786	104599	104395
R^2	0.031	0.033	0.030
Time × Sector FE	yes	yes	yes
Firm FE	no	no	no

Panel B: Firm level

	(1)	(2)	(3)
Size	0.0136***		
	(0.0010)		
Age		0.0070***	
		(0.0005)	
Leverage			-0.0001
			(0.0003)
Observations	104797	104609	104406
R^2	0.574	0.579	0.576
Time × Sector FE	no	no	no
Firm FE	yes	yes	yes

Notes: Panel A reports the estimated coefficient β from $d_{it} = \delta_t + \delta_j + \beta \cdot x_{it} + \varepsilon_{it}$, and Panel B reports the estimated coefficient β from $d_{it} = \delta_i + \beta \cdot x_{it} + \varepsilon_{it}$, where x_{it} is the firm size, age or leverage, constructed as described in the main text, d_{it} is the prevalence attention to FOMC news, δ_i is a firm fixed effect, δ_j is a sector fixed effect (4-digit NAICS), and δ_t is a time fixed effect. Sample excludes the financial and utility sectors.

A.3 Additional robustness

This appendix provides additional robustness checks of our baseline results reported in Table (1.2) from estimating the specification specification in (1.4).

Results robust to controlling for management quality Table A.6 shows that our baseline results are robust to controlling for management quality. We measure management quality as the fraction of board members with education level of the master degree or above, as described in Appendix A.2.1. Since good management can capitalize on expansionary shocks and mitigate contrationary shocks, we allow the controls for management to interact with the monetary shocks asymmetrically. The specification estimated in Table A.6 is:

$$\begin{aligned} r_{it} = \delta_j + \delta_j v_t + \beta_{v+} v_t \mathbb{1}_{v_t>0} + \beta_{v-} v_t \mathbb{1}_{v_t<0} + \beta_d d_{it} + \beta_m m_{it} \\ + \beta_{dv+} d_{it} v_t \mathbb{1}_{v_t>0} + \beta_{dv-} d_{it} v_t \mathbb{1}_{v_t<0} + \beta_{mv+} m_{it} v_t \mathbb{1}_{v_t>0} + \beta_{mv-} m_{it} v_t \mathbb{1}_{v_t<0} \\ + \boldsymbol{\beta}'_{\boldsymbol{X}} (\mathbf{X}_t + \mathbf{X}_t v_t) + \varepsilon_{it}, \end{aligned} \tag{A.3}$$

where m_{it} denotes management quality and d_{it} denotes our baseline prevalence attention measure. Column (1) of Table A.6 reports the results from our baseline specification using only the sample that overlaps with `BoardEx` data. Column (2) reports no significant effects of managment quality on responses to monetary policy. Column (3) reports the effects of attention controlling for management quality. The Wald test for the null hypothesis that β_{dv+} and β_{dv-} are equal is rejected at 1%, suggesting that the finding of asymmetric responses to monetary policy by attention is robust to controlling for management quality.

Results robust to controlling for monetary exposure The theoretical prediction of asymmetry from Section 1.3 confirms the baseline effects to be driven by firm attention rather than firm exposure to monetary policy. Nevertheless, we conduct additional robustness in Table A.7 to control for firms' exposure to monetary policy, constructed as described in

Appendix A.2.2 as rolling historical return sensitivity to monetary shocks. For all three measures of monetary exposure, θ_{it}^{λ} for $\lambda \in$ {baseline, CAPM, FF3}, the Wald tests for the null hypothesis that $\beta_{dv+} = \beta_{dv-}$ are rejected at 5%, showing that our results are not driven by firms' exposure to monetary policy.

Results not driven by the information effect of monetary policy Nakamura and Steinsson (2018a) documents that FOMC announcementa release information about the economic fundamentals, in addition to monetary policy. Following Miranda-Agrippino and Ricco (2021), we control for the information effects of monetary policy by including as controls the Greenbook forecast revisions between FOMC meetings. We obtain data on Greenbook forecasts from the Federal Reserve Bank of Philadelphia. Table A.8 show that our main results are robust to controlling for Greenbook forecast revisions.

Results robust to controlling for macro fluctuations While the high-frequency monetary shocks v_t are considered exogenous, we conduct additional robustness controlling for business-cycle fluctuations. Macro controls include: lagged real GDP growth, unemployment rate, and inflation, obtained from `FRED`. Column (1) of Table A.9 displays our baseline results without macro controls. Column (2) includes macro controls, controlling for aggregate fluctuations. Column (3) includes macro controls and their interactions with the monetary shock, controlling for differential firm sensitivity to aggregate fluctuations. Column (4) includes macro controls and their separate interactions with expansionary and contractionary monetary shocks, controlling for asymmetric firm sensitivity to aggregate fluctuations. Our main results are robust under all specifications.

Table A.6: Controlling for management quality

	(1)	(2)	(3)
Attention	-0.10		-0.10*
	(0.06)		(0.06)
Shock × Attn $\times \mathbb{1}_{v_t>0}$	2.48***		2.59***
	(0.86)		(0.91)
Shock × Attn $\times \mathbb{1}_{v_t<0}$	-7.53**		-8.18**
	(3.52)		(3.34)
Management		-0.04	-0.07
		(0.06)	(0.06)
Shock × Mgmt $\times \mathbb{1}_{v_t>0}$		1.24	1.53**
		(0.76)	(0.73)
Shock × Mgmt $\times \mathbb{1}_{v_t<0}$		-0.68	-2.54
		(2.88)	(2.58)
Observations	324154	324154	324154
R^2	0.038	0.038	0.038
Clustered SE	yes	yes	yes
Firm controls	yes	yes	yes
4-digit NAICS FE	yes	yes	yes
Wald Test p-value	0.008	0.528	0.003

Notes: Results from estimating the specification in (A.3), with variables as defined in the main text:

$$\begin{aligned} r_{it} = \delta_j + \delta_j v_t + \beta_{v+} v_t \mathbb{1}_{v_t>0} + \beta_{v-} v_t \mathbb{1}_{v_t<0} + \beta_d d_{it} + \beta_m m_{it} \\ + \beta_{dv+} d_{it} v_t \mathbb{1}_{v_t>0} + \beta_{dv-} d_{it} v_t \mathbb{1}_{v_t<0} + \beta_{mv+} m_{it} v_t \mathbb{1}_{v_t>0} + \beta_{mv-} m_{it} v_t \mathbb{1}_{v_t<0} \\ + \boldsymbol{\beta}'_{\mathbf{X}} (\mathbf{X}_t + \mathbf{X}_t v_t) + \varepsilon_{it}, \end{aligned}$$

Column (1) shows our baseline results, using only sample that overlaps with BoardEx data. Column (2) shows the effects of management quality. Column(3) the effects of attention controlling for management quality. Standard firm controls include age, size and leverage. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Table A.7: Controlling for exposure to monetary policy

	(1)	(2)	(3)
Shock $\times$ Attn $\times \mathbb{1}_{v_t>0}$	2.03***	2.03***	2.03***
	(0.73)	(0.72)	(0.72)
Shock $\times$ Attn $\times \mathbb{1}_{v_t<0}$	-5.99*	-5.99*	-5.94*
	(3.25)	(3.25)	(3.24)
Observations	572884	571708	568169
R^2	0.026	0.026	0.026
Clustered SE	yes	yes	yes
Firm controls	yes	yes	yes
4-digit NAICS FE	yes	yes	yes
Monetary sensitivity control	baseline model	CAPM model	FF3 model
Wald Test p-value	0.027	0.027	0.027

Notes: Results from estimating the baseline specification (1.4) with additional controls for monetary exposure, θ^{λ}_{it}, $\lambda \in$ {baseline, CAPM, FF3}, defined in Appendix A.2.2. Standard firm controls include age, size and leverage. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Table A.8: Controlling for Greenbook forecast revisions

	(1)	(2)	(3)	(4)
Shock $\times$ Attn $\times \mathbb{1}_{v_t>0}$	2.02***	1.88**	1.94***	1.94***
	(0.72)	(0.75)	(0.72)	(0.72)
Shock $\times$ Attn $\times \mathbb{1}_{v_t<0}$	-5.87*	-5.47	-5.71	-5.71
	(3.18)	(3.58)	(3.68)	(3.68)
Observations	575667	575667	575667	575667
R^2	0.026	0.026	0.026	0.026
Clustered SE	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes
Greenbook rev controls		rgdp	rgdp infl	rgdp infl unemp
Wald Test p-value	0.026	0.070	0.063	0.063

Notes: Results from estimating the baseline specification (1.4) with additional controls for Greenbook forecast revisions. Column (1) displays the baseline results from Table 1.2. Columns (2) - (4) adds Greenbook forecast revisions for real GDP, inflation, and unemployment iteratively. Standard firm controls include age, size and leverage. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Table A.9: Controlling for macroeconomic variables

	(1)	(2)	(3)	(4)
Shock × Attn $\times \mathbb{1}_{v_t>0}$	2.02***	2.06***	1.74**	1.74**
	(0.72)	(0.73)	(0.78)	(0.71)
Shock × Attn $\times \mathbb{1}_{v_t<0}$	-5.87*	-6.27*	-5.38	-7.31**
	(3.18)	(3.21)	(3.34)	(3.31)
Observations	575667	575667	575667	575667
R^2	0.026	0.028	0.028	0.028
Clustered SE	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes
Macro controls	no	yes	yes	yes
+ interactions	no	no	yes	no
+ asym interactions	no	no	no	yes
Wald Test p-value	0.026	0.021	0.060	0.014

Notes: Results from estimating the baseline specification (1.4) with an additional vector of macro control Z_{t-1}, where Z_{t-1} include lagged real GDP growth, unemployment rate, and inflation. Column (1) displays the baseline results from Table 1.2. Column (2) includes macro controls Z_{t-1},. Column (3) includes Z_{t-1} and $Z_{t-1}v_t$. Column (4) includes Z_{t-1} and $Z_{t-1}v_t\mathbb{1}_{v>0}$, and $Z_{t-1}v_t\mathbb{1}_{v<0}$. Standard firm controls include age, size and leverage. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

A.4 Additional results from textual analysis

This appendix contains a set of additional results using natural language processing to investigate the context in which firms discuss macro keywords in 10-K filings and to provide further validation of the text-based measures.

A.4.1 Itemized frequency search

10-K filings have standard formats and are organized in sections. We perform refined frequency counts for each of the section, or "items", to see where attention is concentrated in. Results of frequency counts of macroeconomic keywords by filing item are shown in Figure A.1 in the Appendix. Discussions of the macroeconomy are concentrated in Description of Business (Item 1), Risk Factors (Item 1A) and Management Discussion and Analysis of Financial Condition and Results of Operations (Item 7A).

Figure A.1: Firm attention by filing items

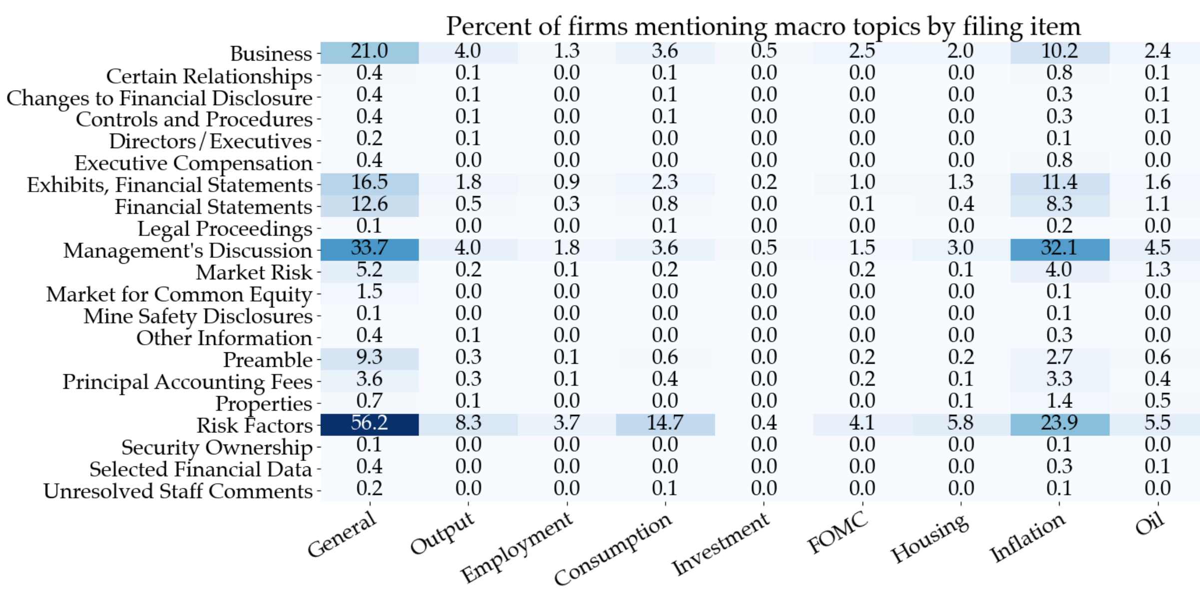

Notes: Heat map of firm attention by filing items. Each row represents a section ("item") of 10-K, and each column represents a macroeconomic topic. Darkness represents a higher fraction of firms that pay attention to a macroeconomic topic in an item.

Results in Figure A.1 show that firms pay attention to macro news to assess the impact on their business operations and risks, consistent with assumptions that firms mentioning a macroeconomic topic do so in order to incorporate the news into their decision making.

A.4.2 LDA: context of macro discussions

To enable automated context detection, we use the Latent Dirichlet Allocation (LDA) model to uncover topics firms tend to discuss in conjunction with macro news. LDA (Blei et al., 2003) is an unsupervised learning algorithm aimed at grouping words in documents into meaningful topics. We apply LDA to texts in earning filings within 20 words surrounding a macroeconomic keyword and set the number of topics to be 10.

Following Hansen et al. (2018), we pre-process texts of 10-K filings for LDA as follows: we remove numbers and words that are only one character. Then we lemmatize to combine different word forms (for example, "operated" and "operates" are lemmatized to "operate"). The advantage of lemmatizing over stemming is that the resulting LDA outputs are more friendly to interpret. Our corpus include words and bigrams which appear for at least 20 times. We filter out words that occur in less than 20 documents or more than 50% of the documents. Then we transform the texts through bag-of-words representation.

We model topics surrounding each of the nine macro categories for the attention measure, as well as an aggregate category containing keywords from all categories. Figures A.2 and A.3 visualize the LDA output surrounding keywords in all categories.Figure A.2 shows the heat map of LDA outputs. Each row represent a topic clustered by LDA, and the darkness of the cell within a topic represent the likelihood of a word to appear in the topic. Figure A.3 highlights the word cloud of selected topics in A.2.

Although LDA output does not label topics, it is natural to characterize some of the topics. Topic 1 relates to business operations, as firms discuss how macro conditions feed into into their daily operations; Topic 2 relates to demand, as firms track and gauge the aggregate demand; Topic 6 relate to financing costs, as firms pay attention to how monetary policy affect their financial costs, investment decisions, and portfolio holdings; Topic 10 relates to labor costs, as firms assess the tightness of the labor market. Rest of the topics relate to housing, currency, and risk factors.

Figure A.2: LDA output for texts surrounding all macro keywords

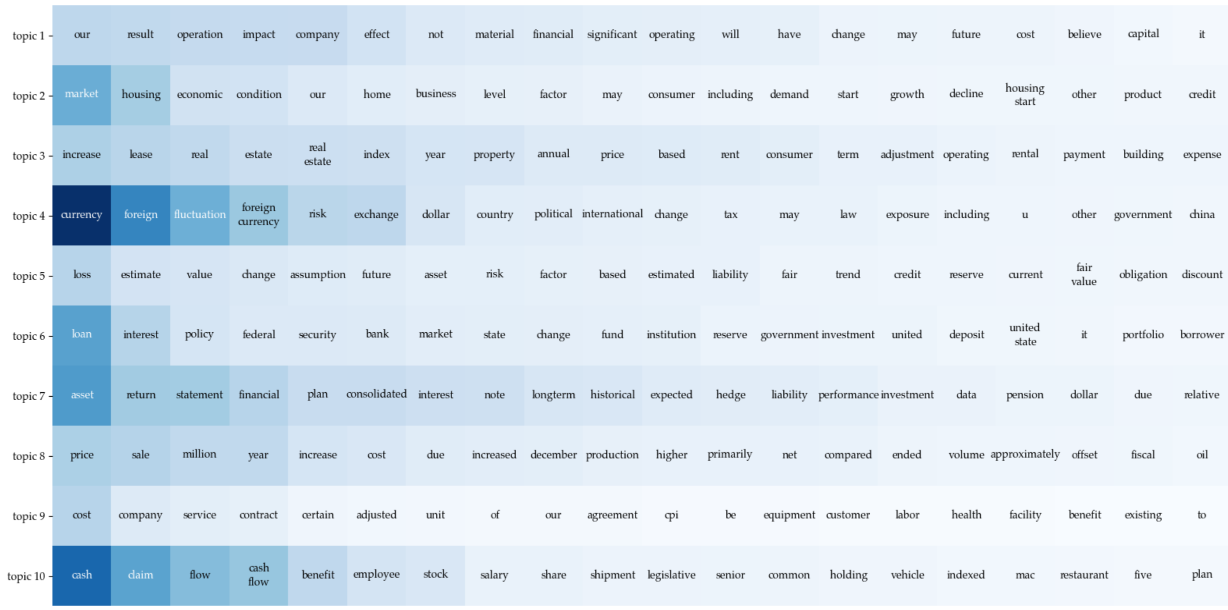

Figure A.3: LDA output for texts surrounding all macro keywords: Selected topics

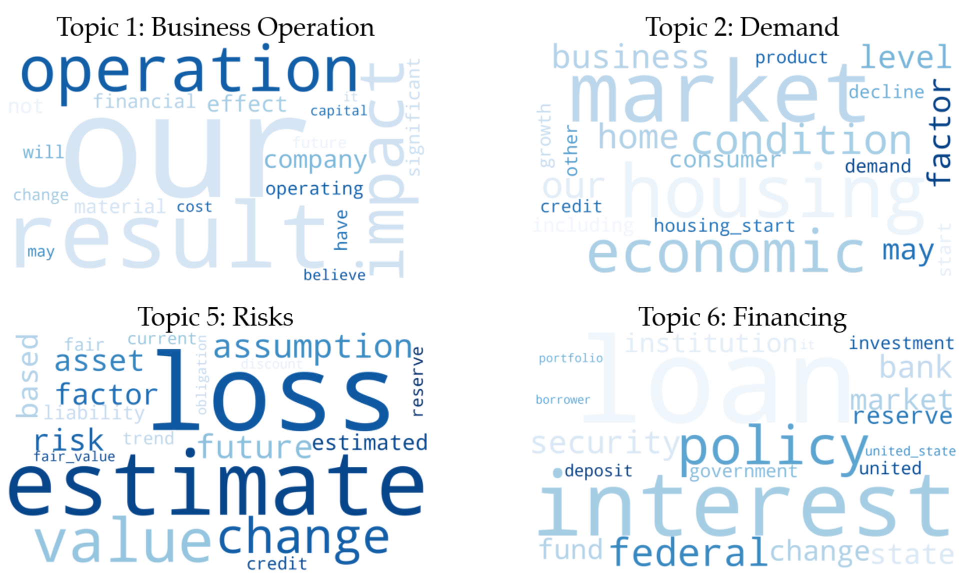

A.4.3 Lexical similarity

Our measure of lexical similarity is a Jaccard score, $J(y_{it}, y_{it-1})$, which measures the share of unique non-stop words that appear between the current year's 10-K (y_i) compared to the previous year's 10-K (y_{it-1}).

$$J(y_i, y_{it-1}) = \frac{|y_i \cap y_{it-1}|}{|y_i \cup y_{it-1}|}$$

The Jaccard score is bounded by the unit interval, and is decreasing with the "uniqueness" of the text. Figure A.4 reports the average Jaccard score for each section of 10-K filings.

Figure A.4: Lexical similarity by section of 10-K filings

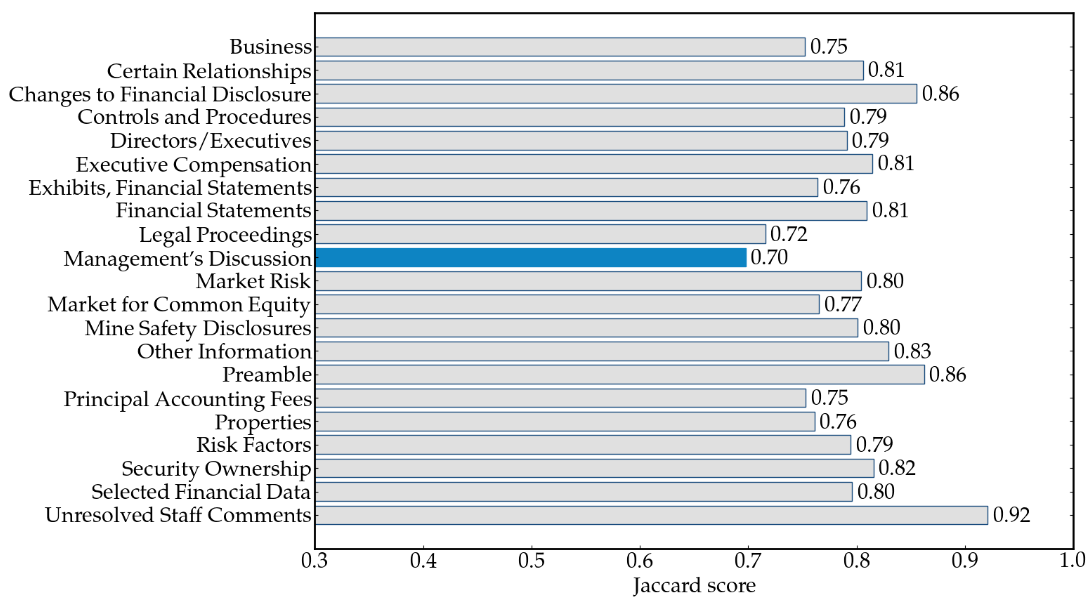

Notes: Average Jaccard scores for sections in 10-K filings. The Jaccard score is bounded by the unit interval. A high Jaccard score represents high lexical similarity between filings. The Management's Discussion section has the lowest level of lexical similarity in all 10-K sections.

We then restrict the attention measures to keywords mentioned in low Jaccard score sections: Business (Item 1) and Managment's Discussion (Item 7). We exclude Legal Proceedings (Item 3) that has a low Jaccard score to avoid false positives from legal languages. Regression results with attention restricted to low lexical similarity 10-K sections are re-

Table A.10: Restricting attention to low lexical similarity 10-K sections

	(1)	(2)	(3)	(4)
Shock	4.13	4.13		
	(2.53)	(2.53)		
Attention		-0.03	-0.08*	-0.05
		(0.03)	(0.04)	(0.05)
Shock × Attn		0.02		
		(0.47)		
Shock × $\mathbb{1}_{v_t>0}$			4.55*	6.21**
			(2.62)	(2.53)
Shock × $\mathbb{1}_{v_t<0}$			-4.16	-1.45
			(4.29)	(4.36)
Shock × Attn ×$\mathbb{1}_{v_t>0}$			0.79	0.50
			(0.50)	(0.47)
Shock × Attn ×$\mathbb{1}_{v_t<0}$			-5.24**	-4.95**
			(2.10)	(2.09)
Observations	546596	546596	546596	409889
R^2	0.023	0.023	0.026	0.027
Clustered SE	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes
excl. ZLB	no	no	no	yes
Wald Test p-value			0.010	0.020

Notes: Results from variants of estimating the baseline specification in (1.4), restricting to 10-K items that discuss firm operations (Items 1 and 7). Standard errors are in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

ported in Table A.10.

A.5 Additional details for the stylized model

A.5.1 Approximation of firm profits in the stylized model

Under second-order approximation around the non-stochastic steady state, the log approximation of a firm's profits, denoted by $\hat{\pi}(s_t, a_t)$, is given by:

$$\begin{aligned}
\hat{\pi}(s_t, a_t) &= \pi(\bar{s}, \bar{a}) + \pi_s(\bar{s}, \bar{a})\bar{s}\hat{s}_t + \pi_a(\bar{s}, \bar{a})\bar{a}\hat{a}_t + \frac{1}{2}\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2\hat{s}_t^2 + \frac{1}{2}\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\hat{a}_t^2 + \pi_{sa}(\bar{s}, \bar{a})\bar{s}\bar{a}\hat{s}_t\hat{a}_t \\
&= \pi(\bar{s}, \bar{a}) + \pi_s(\bar{s}, \bar{a})\bar{s}\hat{s}_t + \frac{1}{2}\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2\hat{s}_t^2 + \frac{1}{2}\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\hat{a}_t^2 - \pi_{aa}(\bar{s}, \bar{a})\bar{a}\bar{s}\hat{a}_t\hat{s}_t \\
&= \pi(\bar{s}, \bar{a}) + \pi_s(\bar{s}, \bar{a})\bar{s}\hat{s}_t + \frac{1}{2}\left(\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2 - \pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\right)\hat{s}_t^2 + \frac{1}{2}\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2(\hat{a}_t - \hat{s}_t)^2
\end{aligned}$$

In the second line, $\pi_a(\bar{s}, \bar{a}) = 0$ because of optimal choice. In addition, the assumption that $a = s$ under full information yields $\pi_a(a, a) = 0 \ \forall a$, which implies $\pi_{sa}(\bar{s}, \bar{a}) = -\pi_{aa}(\bar{s}, \bar{a})$. The third line added and subtracted $\frac{1}{2}\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\hat{s}_t^2$ to complete squares and used the fact that $\bar{a} = \bar{s}$ in the steady state. The resulting expression is equation (1.1).

A.5.2 Proof of Proposition 1

Proof. We consider the responses of returns to an aggregate shock ε. Holding all else equal, that is, $\pi_{ss}^k(s, a) = \pi_{ss}(s, a)$ and $\pi_{aa}^k(s, a) = \pi_{aa}(s, a)$ for all firms k, we can show the following for heterogeneity in exposure and in attention.

1. **Exposure**: Let firms be heterogeneous in exposure and homogeneous in attention. Specifically, suppose firm i is more exposed to macro conditions than firm j, that is, $\pi_s^i > \pi_s^j > 0$. We consider how heterogeneity in exposure affects return elasticity for cases in which both firms are attentive and both are inattentive.

   (a) Case 1 (both firms attentive): When firms are both attentive, $\hat{a}_t = \hat{s}_t$. Then by equation (1.1) we can derive the return elasticity with respect to the aggregate shock to be:

$$\frac{\partial r_k}{\partial \varepsilon} = \frac{\partial \hat{\pi}_k}{\partial \varepsilon} = \pi_s^k(\bar{s}, \bar{a})\bar{s} + \left(\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2 - \pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\right)\varepsilon \quad \text{for firm } k = i, j.$$

Therefore, the return elasticity for firms i is larger for the return elasticity for firm j for all magnitudes of shocks

$$\frac{\partial r_i}{\partial \varepsilon} - \frac{\partial r_j}{\partial \varepsilon} = \pi_s^i(\bar{s}, \bar{a})\bar{s} - \pi_s^j(\bar{s}, \bar{a})\bar{s} > 0$$

because $\pi_s^i > \pi_s^j > 0$.

(b) Case 2 (both firms inattentive): When both firms are inattentive, the return elasticity with respect to the shock can be expressed as:

$$\frac{\partial r_k}{\partial \varepsilon} = \frac{\partial \hat{\pi}_k}{\partial \varepsilon} = \pi_s^k(\bar{s}, \bar{a})\bar{s} + \left(\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2 - \pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\right)\varepsilon$$
$$+ \pi_{aa}(\bar{s}, \bar{a})\bar{a}^2(f_k(\varepsilon) - \varepsilon)(f_k'(\varepsilon) - 1) \quad \text{for firm } k = i, j.$$

Since firms are only heterogeneous in exposure, the second and third term in the above expression for return elasticity is the same for both firms. Therefore:

$$\frac{\partial r_i}{\partial \varepsilon} - \frac{\partial r_j}{\partial \varepsilon} = \pi_s^i(\bar{s}, \bar{a})\bar{s} - \pi_s^j(\bar{s}, \bar{a})\bar{s} > 0$$

which is also independent of the magnitude of ε.

2. **Attention**: Now instead let firms be heterogeneous in attention and homogeneous in exposure, so the attentive firm i has $f_i'(\varepsilon) = 1$, the inattentive firm j has $f_j'(\varepsilon) < 1$, and both firms have $\pi_s^i = \pi_s^j$. The return elasticity for attentive and inattentive firms

can be expressed as:

$$\frac{\partial r_i}{\partial \varepsilon} = \pi_s(\bar{s},\bar{a})\bar{s} + \left(\pi_{ss}(\bar{s},\bar{a})\bar{s}^2 - \pi_{aa}(\bar{s},\bar{a})\bar{a}^2\right)\varepsilon \tag{A.4}$$

$$\frac{\partial r_j}{\partial \varepsilon} = \pi_s(\bar{s},\bar{a})\bar{s} + \left(\pi_{ss}(\bar{s},\bar{a})\bar{s}^2 - \pi_{aa}(\bar{s},\bar{a})\bar{a}^2\right)\varepsilon + \pi_{aa}(\bar{s},\bar{a})\bar{a}^2(f_j(\varepsilon)-\varepsilon)(f_j'(\varepsilon)-1) \tag{A.5}$$

since firms are homogenous in exposure: $\pi_s^i = \pi_s^j = \pi_s$. The relative magnitude of return elasticities between attentive and inattentive firms depends on the sign of the shock ε. Specifically, we consider three cases.

(a) Zero shock ($\varepsilon = 0$): Since $f(0) = 0$, (A.4) and (A.5) lead to:

$$\frac{\partial r_i}{\partial \varepsilon} = \pi_s(\bar{s},\bar{a})\bar{s} = \frac{\partial r_j}{\partial \varepsilon}$$

(b) Positive shock ($\varepsilon > 0$): Since $\varepsilon_t > f_j(\varepsilon_t) > 0$,

$$\frac{\partial r_j}{\partial \varepsilon} - \frac{\partial r_i}{\partial \varepsilon} = \underbrace{\pi_{aa}(\bar{s},\bar{a})\bar{a}^2}_{<0}\underbrace{(f_j(\varepsilon)-\varepsilon)}_{<0}\underbrace{(f_j'(\varepsilon)-1)}_{<0} < 0$$

(c) Negative shock ($\varepsilon < 0$) Since $\varepsilon_t < f_j(\varepsilon_t) < 0$,

$$\frac{\partial r_j}{\partial \varepsilon} - \frac{\partial r_i}{\partial \varepsilon} = \underbrace{\pi_{aa}(\bar{s},\bar{a})\bar{a}^2}_{<0}\underbrace{(f_j(\varepsilon)-\varepsilon)}_{>0}\underbrace{(f_j'(\varepsilon)-1)}_{<0} > 0$$

□

A.6 Additional details for the quantitative model

A.6.1 Approximation of firms' value function

A firms' value function for its operating profits can be expressed as

$$\begin{aligned}
V^{op} &= \max_{\{P_{it},\sigma^2_{i,t|t},\sigma^2_{i,t+1|t}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t\, \mathbb{E}\left[\Pi(P_{it},P_t,Q_t)|\sigma^2_{i,0|-1}\right] \\
&= \max_{\{P_{it},\sigma^2_{i,t|t},\sigma^2_{i,t+1|t}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t\, \mathbb{E}\left[\frac{\Pi(P_{it},P_t,Q_t)}{\Pi^*(P^*_{it},P_t,Q_t)}\Pi^*(P^*_{it},P_t,Q_t)|\sigma^2_{i,0|-1}\right] \\
&= \max_{\{P_{it},\sigma^2_{i,t|t},\sigma^2_{i,t+1|t}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \Pi^*(P^*_{it},P_t,Q_t)\, \mathbb{E}\left[L(P_{it},P_t,Q_t)|\sigma^2_{i,0|-1}\right]
\end{aligned}$$

where $\Pi(P_{it},P_t,Q_t)$ denotes the firm's operating profits, and $L(P_{it},P_t,Q_t) \equiv \frac{\Pi(P_{it},P_t,Q_t)}{\Pi^*(P^*_{it},P_t,Q_t)}$ denotes the loss from imperfect information relative to full-information profits $\Pi^*(P^*_{it},P_t,Q_t)$. The last equality follows the fact that L is homogeneous of degree 1.

Under the second-order log approximation around the non-stochastic steady state, we can express the loss as:

$$\begin{aligned}
\frac{\Pi(P_{it},P_t,Q_t)}{\Pi^*(P^*_{it},P_t,Q_t)} &\approx \frac{\Pi(\bar{P},\bar{P},\bar{Q})}{\Pi(\bar{P},\bar{P},\bar{Q})} + (p_{it}-p^*_{it})\bar{P}\frac{\Pi_1(\bar{P},\bar{P},\bar{Q})}{\Pi(\bar{P},\bar{P},\bar{Q})} + (p^2_{it}-p^{*2}_{it})\frac{\bar{P}}{2}\frac{\Pi_{11}(\bar{P},\bar{P},\bar{Q})}{\Pi(\bar{P},\bar{P},\bar{Q})} \\
&- p_{it}p^*_{it}\bar{P}^2\frac{\Pi_1(\bar{P},\bar{P},\bar{Q})^2}{\Pi(\bar{P},\bar{P},\bar{Q})^2} + (p_{it}-p^*_{it})p_t\bar{P}^2\frac{\Pi_{12}(\bar{P},\bar{P},\bar{Q})}{\Pi(\bar{P},\bar{P},\bar{Q})} + (p_{it}-p^*_{it})q_t\bar{P}\bar{Q}\frac{\Pi_{13}(\bar{P},\bar{P},\bar{Q})}{\Pi(\bar{P},\bar{P},\bar{Q})} \\
&= 1 + (p^2_{it}-p^{*2}_{it})\frac{\bar{P}}{2}\frac{\Pi_{11}(\bar{P},\bar{P},\bar{Q})}{\Pi(\bar{P},\bar{P},\bar{Q})} + (p_{it}-p^*_{it})p_t\bar{P}^2\frac{\Pi_{12}(\bar{P},\bar{P},\bar{Q})}{\Pi(\bar{P},\bar{P},\bar{Q})} + (p_{it}-p^*_{it})q_t\bar{P}\bar{Q}\frac{\Pi_{13}(\bar{P},\bar{P},\bar{Q})}{\Pi(\bar{P},\bar{P},\bar{Q})} \\
&= 1 + (p_{it}-p^*_{it})^2\frac{\bar{P}}{2}\frac{\Pi_{11}(\bar{P},\bar{P},\bar{Q})}{\Pi(\bar{P},\bar{P},\bar{Q})}
\end{aligned}$$

where lowercase letters denote log deviations from the steady state. The second equality uses the fact that $\Pi_1 = 0$ from optimal choices. In addition, $\Pi_1(P^*_{it},P_t,Q_t) = 0$ implies $p^*_{it}\bar{P}\Pi_{11}(\bar{P},\bar{P},\bar{Q}) + p_t\bar{P}\Pi_{12}(\bar{P},\bar{P},\bar{Q}) + q_{it}\bar{Q}\Pi_{13}(\bar{P},\bar{P},\bar{Q}) = 0$, which leads to the third equality.

A firm's log operating value, v^{op}, can be decomposed into:

$$v^{op} = v^* + l$$

consisting of v^*, the full-information value, and l, the loss in firm value from imperfect information approximated as above.

A.6.2 Details for model calibration

Figure A.5: Sensitivity of simulated moments to ω_L

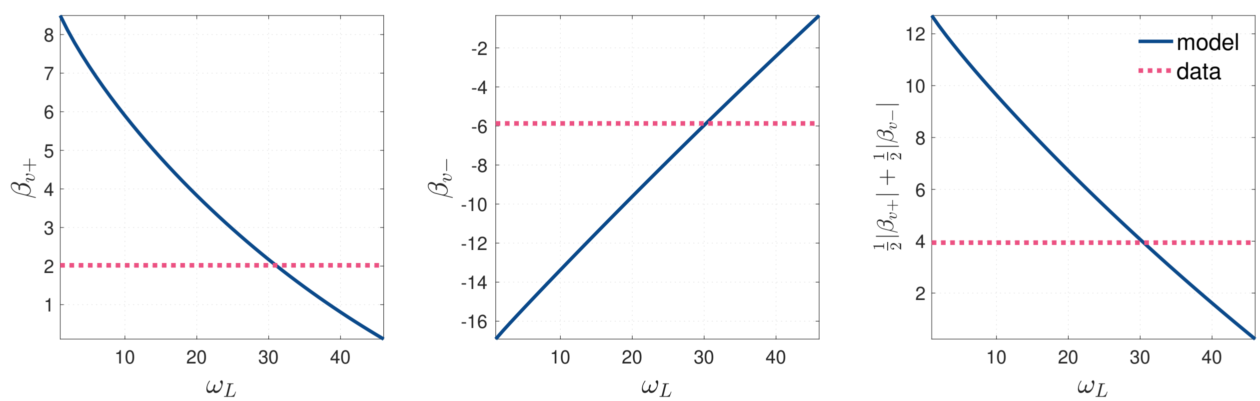

Notes: Calibration plots showing simulated moments for a range of costs of information parameters (ω_L). We simulate models for a panel of 100 firms and for 1000 periods with 100 periods burn-ins. Simulated moments are generated with regressions discussed in the text in Section 1.5:

$$r_{it} = c + \beta_1 \mathbb{1}_{v>0} + \beta_{v+} v_t \mathbb{1}_{v>0} + \beta_{v-} v_t \mathbb{1}_{v<0} + \beta_d d_{it} + \beta_{dv+} d_{it} \nu_{it} \mathbb{1}_{v>0} + \beta_{dv-} d_{it} \nu_{it} \mathbb{1}_{v<0} + \varepsilon_{it}$$

The left panel shows the sensitivity of simulated β_{v+} to the calibration of ω_L; the middle panel shows the sensitivity of β_{v-}; the right panel shows the sensitivity of $\frac{1}{2}|\beta_{v+}| + \frac{1}{2}|\beta_v|$, which we use to calibrate ω_L to match the empirical moment in the data.

A.6.3 Passthrough regressions

The passthrough of nominal interest rate change to nominal demand change is estimated with local projections (Jordà, 2005). We estimate the following model for horizons $h =$

Figure A.6: Passthrough of rates to nominal demand

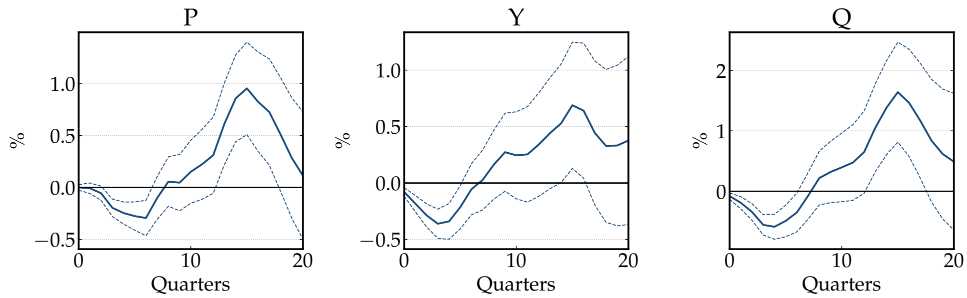

$1, 2, \cdots, 20$:

$$\Delta_h y_{t-1,t+h} = \alpha_h + \beta_h \varepsilon_t^{\mathtt{i}} + u_{th}$$

where y is the variable of interest, and $\varepsilon_t^{\mathtt{i}}$ is a shock to the nominal interest rate. Path of β_h informs the cumulative changes in the dependent variable in response to the interest rate shock.

The dependent variables are U.S. manufacturing output over the sample period of 1994 to 2019. We estimate the responses of manufacturing prices, real output and nominal output. We time aggregate high-frequency monetary policy shocks to quarterly to match the frequency of dependent variables. Figure A.6 shows the results of the local projection. A one percentage point expansionary shock to the interest rate leads to about 1.6 percent peak increase in nominal demand.

Appendix for Chapter 2

B.1 Additional tables and figures in Chapter 2

Table B.1: Daily returns of equity indices

	Release	**Nonrelease**	**All Days**
SP500 Ex-Financial			
Mean	-0.03	0.03	0.02
	(0.06)	(0.02)	(0.02)
Std Deviation	1.32	1.12	1.14
	(0.04)	(0.01)	(0.01)
Observations	486	5,048	5,534
SML			
Mean	0.03	0.03	0.03
	(0.07)	(0.02)	(0.02)
Std Deviation	1.58	1.39	1.41
	(0.05)	(0.01)	(0.01)
Observations	486	4,603	5,089
Russell			
Mean	0.02	0.02	0.02
	(0.08)	(0.02)	(0.02)
Std Deviation	1.70	1.46	1.48
	(0.05)	(0.01)	(0.01)
Observations	486	4,603	5,089

Notes: This table shows descriptive statistics of daily returns of equity indices (S&P 500 ex-financials, S&P Small Cap 600 and Russell 2000). Returns are computed as daily log differences. Mean and standard deviations are reported in percent. "Release Days" refer to days with earnings releases by financial intermediaries in the sample; "Nonrelease Days" refer to days without earnings releases; "All Days" include both release days and nonrelease days. Standard errors are in parentheses.

Table B.2: Daily changes in bond option-adjust spreads

	Release	**Non-Release**	**All Days**	**N Bonds**
AAA				
Mean	-0.07	0.03	0.02	293
	(0.05)	(0.03)	(0.02)	
Std Deviation	5.96	10.38	10.02	
	(0.03)	(0.02)	(0.02)	
CCC				
Mean	1.20	1.80	1.74	3308
	(0.29)	(0.10)	(0.09)	
Std Deviation	110.09	106.81	107.17	
	(0.20)	(0.07)	(0.06)	
All bonds				
Mean	1.07	1.60	1.55	3601
	(0.26)	(0.09)	(0.08)	
Std Deviation	104.11	100.86	101.21	
	(0.18)	(0.06)	(0.06)	

Notes: This table shows descriptive statistics of daily changes in option-adjusted spreads for nonfinancial constituent bonds in ICE BofA's AAA and CCC & Lower indices of U.S. corporate bonds. Mean and standard deviation of daily changes in option-adjusted spreads are reported in basis points. The number of nonfinancial constituent bonds (CUSIP) in each index are reported. "Release Days" refer to days with earning releases by financial intermediaries in the sample; "Nonrelease Days" refer to days without earnings releases; "All Days" include both release days and nonrelease days. Standard errors are in parentheses.

Table B.3: Bond holdings by intermediary

Financial Intermediary	**Mean**	**Std Dev**	**Min**	**Max**
J.P. Morgan Chase	9,782	31,692	0	351,996
Goldman Sachs	3,963	14,337	0	254,385
Ameriprise Financial	2,486	8,389	0	113,540
Northern Trust	910	4,674	0	88,840
Wells Fargo	734	2,875	0	38,253
Citicorp	624	3,029	0	66,300
Bank of New York Mellon	588	2,647	0	48,695
Morgan Stanley	244	1,234	0	28,555
Merrill Lynch	17	276	0	8,901
U.S. Bancorp	7	55	0	2,000
Bank of America	2	45	0	1,580
All	1,760	11,337	0	351,996

Notes: This table shows descriptive statistics on bond holdings by financial intermediaries. The set of bonds includes bonds rated CCC or lower in ICE issued by firms with at least 10 bonds outstanding.

Table B.4: Summary statistics for event and nonevent days

	Financial Intermediaries		Nonfinancial Firms	
	Event	Nonevent	Event	Nonevent
Mean of Weighted ΔP	-0.00	-0.02	-0.04	0.01
	(0.13)	(0.04)	(0.04)	(0.02)
Std Dev of Weighted ΔP	2.98	2.34	1.31	1.20
	(0.09)	(0.03)	(0.03)	(0.02)
Covariance with SP500	2.76	2.07	1.45	1.29
Covariance with CCC	-0.17	-0.10	-0.05	-0.06
Observations	492	3,012	877	2,347

Notes: This table shows summary statistics for identification through heteroskedasticity: weighted daily stock-price changes for event days and nonevent days. Nonevent days are defined as days without earning releases and are at least 2 days away from event days. S&P500 excludes financial firms. Standard errors are in parentheses.

Figure B.1: Construction of financial shocks

(a) Median Positive Shock (Inside Trading Hours)

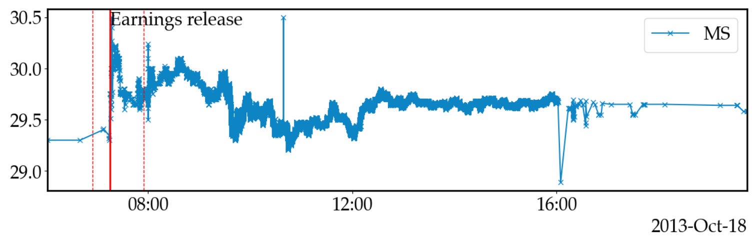

(b) Median Negative Shock (Inside Trading Hours)

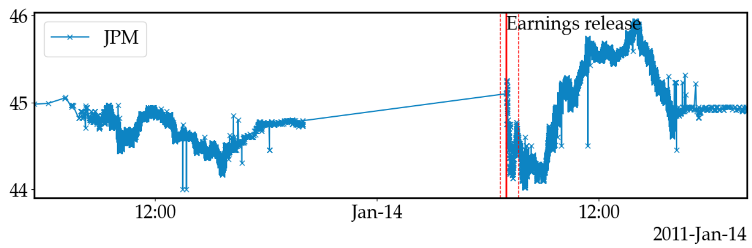

(c) Median Positive Shock (Outside Trading Hours)

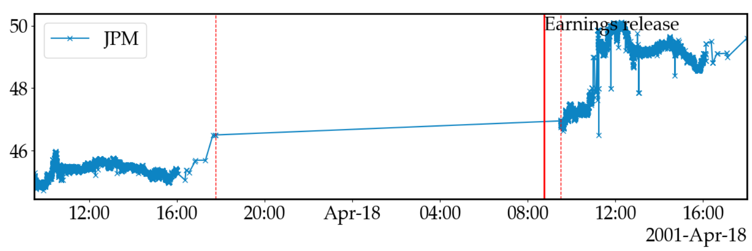

(d) Median Negative Shock (Outside Trading Hours)

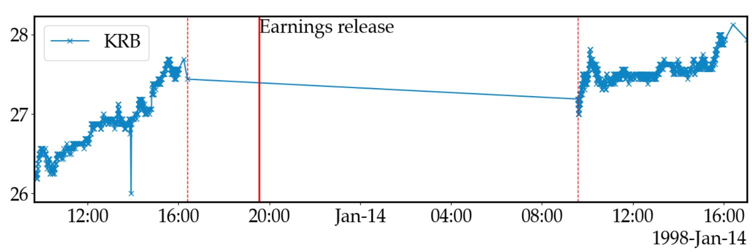

Figure B.2: The effect of financial shocks on financial sector's net worth

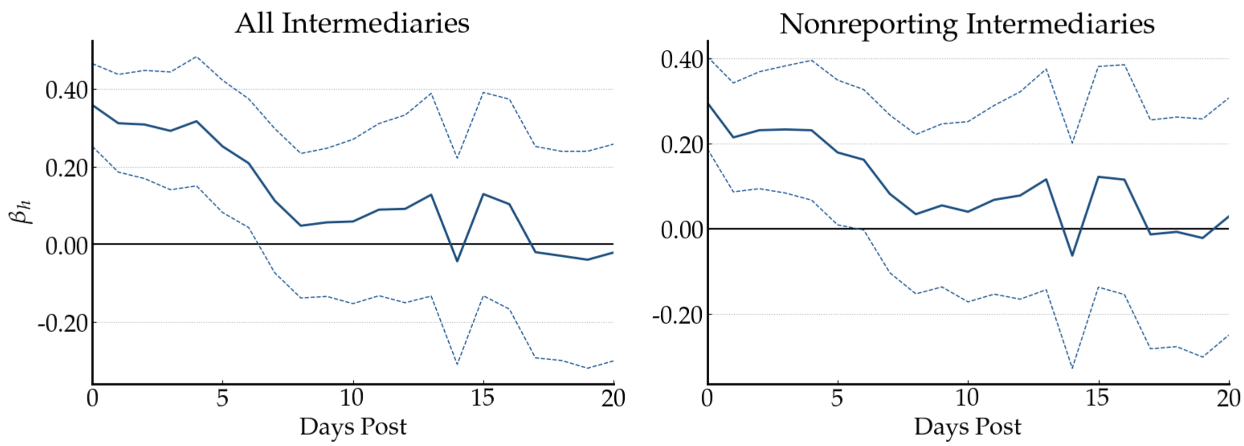

Notes: The figures show the cumulative responses of financial intermediaries' market capitalization to individual unweighted financial shocks. The left panel shows the market capitalization responses from all financial intermediaries in our sample in response to a financial shock. The middle panel shows the market capitalization response from the intermediary that reports the earning underlying the financial shock. The right panel shows the market capitalization response from all remaining nonreporting intermediaries.

Figure B.3: Placebo tests: financial shocks and nonevent days

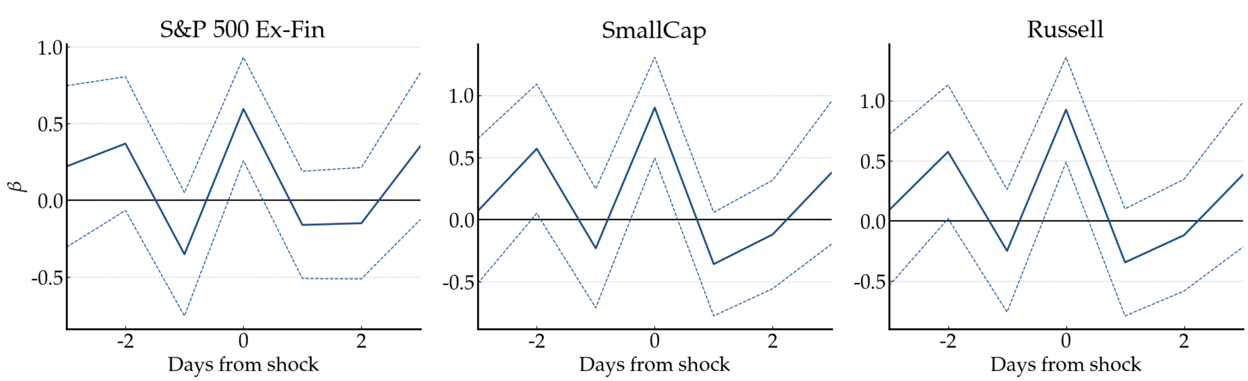

Notes: The figures show placebo tests with nonevent days. Specifications take the form: $\Delta \log y_{t+j} = c + \beta \nu_t + \varepsilon_t$. Changes in the dependent equity indices are constructed using alternative dates $j = -3, \cdots, 3$ around the event date, with $j = 0$ corresponding to the event date of earnings releases.

Table B.5: Placebo tests: effect of shocks to nonfinancial firms

	SP500	**SmallCap**	**Russell**	**Obs**
Baseline	0.334	0.064	0.135	877
	(0.220)	(0.256)	(0.268)	
Narrow window only	−0.205	−0.557*	−0.513	546
	(0.272)	(0.330)	(0.346)	
Macro controls	0.338	0.067	0.137	877
	(0.219)	(0.256)	(0.268)	

Notes: This table shows placebo tests with HF shocks generated with nonfinancial firms in Dow Jones. Shock construction and regression specifications follow those for financial shocks. Firms include 3M, Alco, Philip Morris, Apple, AT&T, Bethlehem Steel, Boeing, Caterpillar, Chevron, Cisco, Coca-Cola, Dow, Dupont, Eastman Kodak, Exxon, FW Woolworth, General Electric, General Motors, Goodyear, Hewlett-Packard, Home Depot, Intel, IBM, International Paper, Johnson & Johnson, Kraft, McDonald's, Merck, Microsoft, Nike, Pfizer, Procter & Gamble, Sears, Texaco, Union Carbide, United Technologies, UnitedHealth, Verizon, Visa, Walgreens, Walmart, Walt Disney, and Westinghouse.

B.2 Additional exercises of HF financial shocks

B.2.1 Financial shocks and surprise earnings

In this section, we document the connection between HF shocks and the surprise component of intermediaries' earnings. Figure B.4 reports the binned scatter plot between the unweighted HF shocks and earnings surprises of the reporting intermediaries. Following the post-earnings-announcement-drift literature (Chordia and Shivakumar, 2006), earnings surprises are measured using the standardized unexpected earnings, defined as the difference between the reported earnings per share and the consensus forecast, normalized by the standard error of analysts' forecast errors. We observe a positive relationship between HF shocks and earnings surprises, suggesting the HF movements of an intermediary's stock prices around the earnings announcement are associated with its earnings outcome.

Figure B.4: Financial shocks and earnings surprises

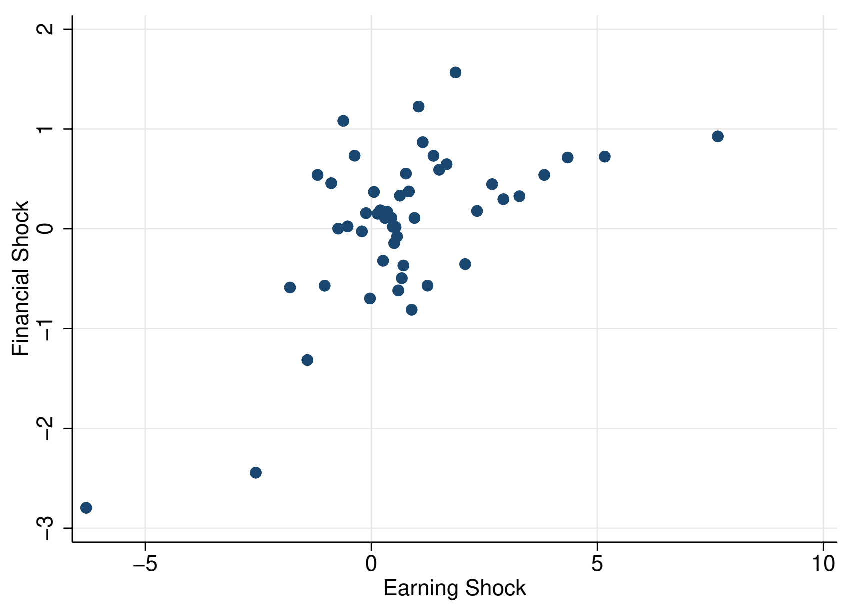

Notes: This figure shows a binned scatter plot between financial shocks and earnings surprises with 50 bins. Financial shocks are unweighted and constructed as described in the main text. Earnings surprises are measured as standardized unexpected earnings, defined as surprises in earnings per share normalized by the standard errors of analysts' forecast errors.

B.2.2 Predictability of financial shocks

In this section, we use a state-of-art machine-learning model to provide evidence suggesting that the HF financial shocks are not predictable using macroeconomic and financial variables available prior to the shock. We use two sets of predictors. The first macro panel includes a large panel of 126 monthly macroeconomic series constructed by McCracken and Ng (2016) and available through FRED-MD. The second financial panel is of higher daily frequency and includes stock prices of the financial intermediaries in our sample, as well as SP500 and VIX.

Our main forecasting model is random forests (Breiman, 2001), which produce an averaged prediction from a large collection of regression trees. Random forests incorporate nonlinearity and multi-way interactions between predictors, which makes it useful for macroeconomic and financial forecasting (Gentzkow et al., 2019). The random-forest predictor is defined as

$$\hat{f}_{\mathrm{rf}}^{B} = \frac{1}{B}\sum_{b=1}^{B} T(x;\Theta_b),$$

which averages the forecasts of B regression trees $T(x;\Theta_b)$, where x is the set of predictors and Θ_b characterizes the parameters in the bth tree.[1]

As Gentzkow et al. (2019) argues, the benefits of regression trees from nonlinearity and high-order interactions lessens with high-dimensional predictors, so we first perform variable selection with elastic net (Zou and Hastie, 2005), an implementation of soft thresholding regularization that drops uninformative predictors using penalized regressions. The elastic net estimator is defined by

$$\hat{\beta}_{\mathrm{EN}} = \arg\min_{\beta}\left\{\frac{1}{2}\sum_{i=1}^{N}(y_i-\beta_0-\sum_{j=1}^{p}x_{ij}\beta_j)^2+\lambda\left(\frac{1}{2}(1-\alpha)\|\beta\|_{l_2}^2+\alpha\|\beta\|_{l_1}\right)\right\},$$

[1] Please see Hastie et al. (2009) for a comprehensive exposition of trees and random forests.

Table B.6: Out-of-sample R^2 of predictions of financial shocks

	Macro	**Financial**
Random Forest	−9.7%	−16.2%
Random Walk		24.2%

Notes: This table reports the out-of-sample R^2 of random-forest forecasts based on a large panel of macroeconomic and financial variables compared against out-of-sample R^2 of random-walk forecasts based on stock returns one day before the shock. The out-of-sample R^2 is defined as $R^2_{\text{oos}} = 1 - \frac{\Sigma_t(y_t - \hat{y}_{m,t})^2}{\Sigma_t(y_t - \bar{y}_t)^2}$ where $\bar{y}_t$ is the rolling-mean forecast computed on a window matching the model-estimation window, and $\hat{y}_{m,t}$ is the forecast from the model. Negative numbers indicate the forecast underperforms the rolling historical mean of the series.

which minimizes the sum or regression residuals and a penalty term, a weighted average of LASSO and ridge. Following Borup and Schütte (2020), we set $\alpha = 0.5$ for an equal weight between LASSO and ridge regressions and tune the penalty parameter λ so that the elastic net selects the 20 best predictors.

We then use random forests to form predictions using 48-month rolling windows for macro predictors and 48-day rolling windows for financial predictors. To assess forecastability, we compare the predictions from random forests against those from a random walk, formed with the stock returns one day before the financial shock, converted to match the size of the 60-minute shock window. The metric for evaluating forecastability is the out-of-sample R^2 (Campbell and Thompson, 2008), defined as

$$R^2_{\text{oos}} = 1 - \frac{\Sigma_t(y_t - \hat{y}_{m,t})^2}{\Sigma_t(y_t - \bar{y}_t)^2},$$

where $\bar{y}_t$ is the rolling-mean forecast computed on a window matching the model-estimation window, and $\hat{y}_{m,t}$ is the forecast from the model. R^2_{oos} lies in the range $(-\infty, 1]$ with negative numbers indicating the model underperforms the historical mean of the series.

Assessments of forecastability of the financial shocks by macroeconomic and financial predictors are shown in Table B.6. Random-forest forecasts with both the macro and financial predictors have negative R^2_{oos}, suggesting worse performance than historical rolling-mean forecasts. The results also suggest incorporating panels of macro and financial variables does not help forecast the HF financial shocks compared to a random walk.

Appendix for Chapter 3

C.1 Additional tables and figures in Chapter 3

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

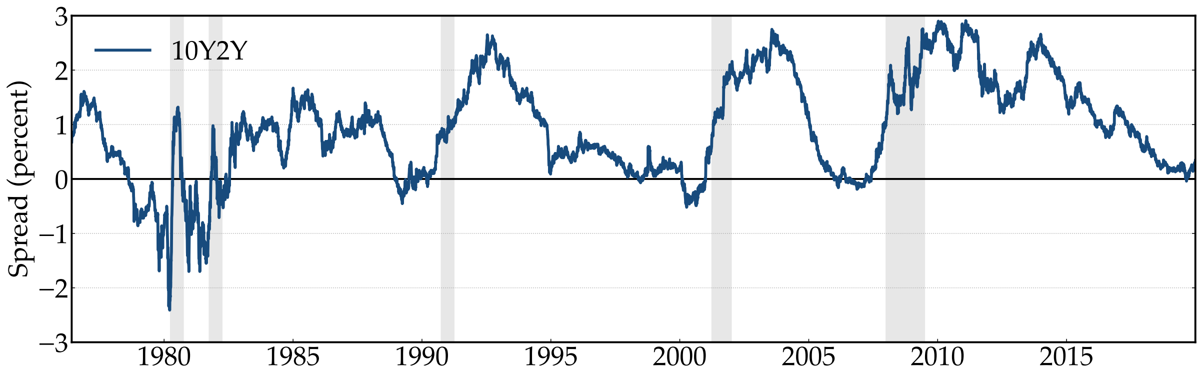

Notes: Yield curve and recessions in the US for 1976–2019. The blue solid line displays the spread between 10-year treasury yield and 2-year treasury yield ("10Y2Y"). Recession dates as classified by NBER are shaded in grey.

Table C.1: Limiting the number of outlets in user timelines

(a) Narrative: "recession"

Dependent variable: Consumer sentiment	(1) Δs_{it}	(2) Δd_{it}	(3) Δs_{it}	(4) Δd_{it}
$\mathbb{1}(\theta_d^{\texttt{rec}} > \bar{\theta}^{\texttt{rec}})$	-0.02*	-0.10*		
	(0.01)	(0.05)		
$\theta_d^{\texttt{rec}}$			-0.02*	-0.11*
			(0.01)	(0.07)
Constant	0.00	-0.00	0.00	0.00
	(0.00)	(0.02)	(0.00)	(0.02)
Observations	227	227	227	227
R^2	0.014	0.014	0.015	0.012

(b) Narrative: "overblown"

Dependent variable: Consumer sentiment	(1) Δs_{it}	(2) Δd_{it}	(3) Δs_{it}	(4) Δd_{it}
$\mathbb{1}(\theta_d^{\texttt{ovb}} > \bar{\theta}^{\texttt{ovb}})$	0.00	0.02		
	(0.01)	(0.04)		
$\theta_d^{\texttt{ovb}}$			-0.00	-0.01
			(0.01)	(0.05)
Constant	-0.00	-0.02	-0.00	-0.01
	(0.00)	(0.02)	(0.00)	(0.02)
Observations	227	227	227	227
R^2	0.001	0.002	0.000	0.000

Notes: Results from estimating the specification in (3.10) and (3.11) while limited the number of different news outlets appearing in a user's timeline over a two-month period to be 4. Details are as described in the main text in Section 3.4.5. Standard errors are in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

C.2 Additional details for narratives

The main text describes two narratives:

$$G^1: \quad a \to s \leftarrow \theta \to y$$
$$G^2: \quad a \to s \to y$$

This appendix makes further functional-form assumptions to illustrate how G^1, G^2, and their implied policies can both be stable in the steady state. Suppose the long-run probability of the current policy $p(a = 1) = \alpha$, and of the fundamentals is $p(\theta = 1) = \delta \in (\varepsilon, 1 - \varepsilon)$ for ε small, which is independent of a. The output is a direct results of the fundamentals, $p(y = \theta|\theta) = 1, \forall\theta$, and the yield curve is influeced by both the economic policy and the fundamentals, $p(s = 1|a, \theta) = a + (1 - a)\theta$. Assume that the policy which place weight d on the action $a = 1$ gives the payoff $u(d, y) = y - \frac{1}{2}(d - \varepsilon)^2$. The payoff implies that an expansion $y = 1$ is the ideal outcome, and that the optimal policy is $d = \varepsilon$, deviating as infrequently from the current policy as possible. The environment follows Eliaz and Spiegler (2020).

First, consider a steady-state with only the narrative G^1, which correctly predicts $y = 1$ with probability δ independent of the action. The expected payoff from the narrative is given by:

$$U(G^1, d; \alpha) = \delta - \frac{1}{2}(d - \varepsilon)^2$$

which induces the ideal policy $d = \varepsilon$ and an expected utility of δ.

Now consider the entry of a new narrative G^2, which creates an artificial causal structure between policy, yield curve inversion, and recession. The expected utility under this narrative is given by:

$$
\begin{aligned}
U(G^2, d; \alpha) &= \sum_y p_{G^1}(y|a)u(a, y) = \sum_y \sum_a p(s|a)p(y|s)u(a, y) \\
&= (1 \cdot d \cdot \frac{\delta}{\delta + \alpha(1-\delta)} + \delta \cdot (1-d) \cdot \frac{\delta}{\delta + \alpha(1-\delta)})(1 - \frac{1}{2}(d-\varepsilon)^2) \\
&= \delta \cdot \frac{\delta + d(1-\delta)}{\delta + \alpha(1-\delta)} - \frac{1}{2}(d-\varepsilon)^2 \\
&\approx \delta + (1-\delta)d - \frac{1}{2}d^2
\end{aligned}
$$

where the last line uses the fact that $\alpha = \varepsilon \approx 0$ under the steady-state with a single narrative G^1. The optimal policy under the subjective belief $d = 1 - \delta$ leads to an expected utility of $\delta + \frac{1}{2}(1-\delta)^2 > \delta$, which is greater than the steady-state utility under the single narrative.

The economy reaches a new steady state when the long-run probability of the current policy is α^* so that $U(G^2, d^2; \alpha^*) = \delta$ with $d^2 = \arg\max_d U(G^2, d; \alpha^*)$. In the new steady state, both narratives-policy pairs, (G^1, ε) and (G^2, d^2), yield the same expected utility and can co-exist.

Bibliography

Abis, Simona and Laura Veldkamp, "The Changing Economics of Knowledge Production," *Available at SSRN 3570130*, 2020.

Adrian, Tobias, Erkko Etula, and Tyler Muir, "Financial intermediaries and the cross-section of asset returns," *The Journal of Finance*, 2014, *69* (6), 2557–2596.

Afrouzi, Hassan, "Strategic Inattention, Inflation Dynamics, and the Non-neutrality of Money," 2020.

_ **and Choongryul Yang**, "Dynamic Rational Inattention and the Phillips Curve," 2021.

Anderson, Gareth and Ambrogio Cesa-Bianchi, "Crossing the credit channel: credit spreads and firm heterogeneity," 2020.

Angeletos, George-Marios and Jennifer La'O, "Sentiments," *Econometrica*, 2013, *81* (2), 739–779.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub, "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," Technical Report, National Bureau of Economic Research 2020.

Bénabou, Roland, Armin Falk, and Jean Tirole, "Narratives, imperatives, and moral reasoning," Technical Report, National Bureau of Economic Research 2018.

Bernanke, Ben S, "The real effects of disrupted credit: evidence from the global financial crisis," *Brookings Papers on Economic Activity*, 2018, *2018* (2), 251–342.

_ **and Kenneth N Kuttner**, "What explains the stock market's reaction to Federal Reserve policy?," *The Journal of finance*, 2005, *60* (3), 1221–1257.

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, Andrew Ng, and Michael Jordan, "Latent Dirichlet Allocation," *Journal of Machine Learning Research*, 2003, *3* (Jan), 993–1022.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Diagnostic expectations and credit cycles," *The Journal of Finance*, 2018, *73* (1), 199–227.

Borup, Daniel and Erik Christian Montes Schütte, "In search of a job: Forecasting employment growth using Google Trends," *Journal of Business & Economic Statistics*, 2020, pp. 1–15.

Breiman, Leo, "Random forests," *Machine learning*, 2001, *45* (1), 5–32.

Brunnermeier, Markus K and Yuliy Sannikov, "A Macroeconomic Model with a Financial Sector," *The American Economic Review*, February 2014, *104* (2), 379–421.

Buehlmaier, Matthias and Toni Whited, "Are Financial Constraints Priced? Evidence from Textual Analysis," *The 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.

Campbell, John Y and Samuel B Thompson, "Predicting excess stock returns out of sample: Can anything beat the historical average?," *The Review of Financial Studies*, 2008, *21* (4), 1509–1531.

Candia, Bernardo, Olivier Coibion, and Yuriy Gorodnichenko, "The Inflation Expectations of US Firms: Evidence from a new survey," Technical Report, National Bureau of Economic Research 2021.

Caplin, Andrew and Daniel Spulber, "Menu Costs and the Neutrality of Money," *The Quarterly Journal of Economics*, 1987, *102* (4), 703–725.

_ , Mark Dean, and John Leahy, "Rational Inattention, Optimal Consideration Sets, and Stochastic Choice," *The Review of Economic Studies*, 2019, *86* (3), 1061–1094.

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.

Chiang, Yu-Ting, "Strategic Uncertainty over Business Cycle," 2021.

Chodorow-Reich, Gabriel, "The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis," *The Quarterly Journal of Economics*, 2014, *129* (1), 1–59.

Chordia, Tarun and Lakshmanan Shivakumar, "Earnings and price momentum," *Journal of financial economics*, 2006, *80* (3), 627–656.

Christiano, Lawrence J, Martin S Eichenbaum, and Mathias Trabandt, "Understanding the great recession," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 110–67.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Cloyne, James, Clodomiro Ferreira, Maren Froemel, and Paolo Surico, "Monetary policy, corporate finance and investment," Technical Report, National Bureau of Economic Research 2018.

Cochrane, John H and Monika Piazzesi, "The fed and interest rates-a high-frequency identification," *American Economic Review*, 2002, *92* (2), 90–95.

Coibion, Olivier and Yuriy Gorodnichenko, "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 2015, *105* (8), 2644–78.

_ , _ , and Saten Kumar, "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review*, 2018, *108* (9), 2671–2713.

_ , _ , and Tiziano Ropele, "Inflation Expectations and Firm Decisions: New Causal Evidence," *The Quarterly Journal of Economics*, 2020, *135* (1), 165–219.

Cook, Timothy and Thomas Hahn, "The effect of changes in the federal funds rate target on market interest rates in the 1970s," *Journal of Monetary Economics*, 1989, *24* (3), 331–351.

Coval, Joshua and Erik Stafford, "Asset fire sales (and purchases) in equity markets," *Journal of Financial Economics*, 2007, *86* (2), 479–512.

Duffee, Gregory R, "The relation between treasury yields and corporate bond yield spreads," *The Journal of Finance*, 1998, *53* (6), 2225–2241.

Eliaz, Kfir and Ran Spiegler, "A model of competing narratives," *American Economic Review*, 2020, *110* (12), 3786–3816.

Farhi, Emmanuel and Iván Werning, "Monetary policy, bounded rationality, and incomplete markets," *American Economic Review*, 2019, *109* (11), 3887–3928.

Flynn, Joel P and Karthik Sastry, "Attention cycles," 2021.

Friedman, Jerome, Trevor Hastie, Robert Tibshirani et al., *The elements of statistical learning*, Vol. 1, Springer series in statistics New York, 2001.

Gabaix, Xavier, "Behavioral Inattention," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 2, Elsevier, 2019, pp. 261–343.

_ and Ralph SJ Koijen, "Granular instrumental variables," *Available at SSRN 3368612*, 2020.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, "Text as data," *Journal of Economic Literature*, 2019, *57* (3), 535–74.

_ , **Jesse M Shapiro, and Michael Sinkinson**, "Competition and ideological diversity: Historical evidence from us newspapers," *American Economic Review*, 2014, *104* (10), 3073–3114.

Gertler, Mark and John Leahy, "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, 2008, *116* (3), 533–572.

_ **and Nobuhiro Kiyotaki**, "Financial Intermediation and Credit Policy in Business Cycle Analysis," *Handbook of Monetary Economics*, 2010, *3* (11), 547–599.

_ **and Simon Gilchrist**, "Monetary policy, business cycles, and the behavior of small manufacturing firms," *The Quarterly Journal of Economics*, 1994, *109* (2), 309–340.

_ **and** _ , "What happened: Financial factors in the great recession," *Journal of Economic Perspectives*, 2018, *32* (3), 3–30.

_ **and** _ , "The Channels of Financial Distress During the Great Recession: Some Evidence on the Aggregate Effects," 2019.

Gilchrist, Simon, Vladimir Yankov, and Egon Zakrajšek, "Credit market shocks and economic fluctuations: Evidence from corporate bond and stock markets," *Journal of monetary Economics*, 2009, *56* (4), 471–493.

Go, Alec, Richa Bhayani, and Lei Huang, "Twitter sentiment classification using distant supervision," *CS224N project report, Stanford*, 2009, *1* (12), 2009.

Gorodnichenko, Yuriy and Michael Weber, "Are Sticky Prices Costly? Evidence from the Stock Market," *American Economic Review*, 2016, *106* (1), 165–99.

Gürkaynak, Refet S, Brian P Sack, and Eric T Swanson, "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements," 2004.

_ , **Brian Sack, and Eric T Swansonc**, "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 2005.

Handley, Kyle and J Frank Li, "Measuring the Effects of Firm Uncertainty on Economic Activity: New Evidence from One Million Documents," Technical Report, National Bureau of Economic Research 2020.

Hansen, Stephen, Michael McMahon, and Andrea Prat, "Transparency and Deliberation within the FOMC: A Computational Linguistics Approach," *The Quarterly Journal of Economics*, 2018, *133* (2), 801–870.

Harvey, Campbell R, "The real term structure and consumption growth," *Journal of Financial Economics*, 1988, *22* (2), 305–333.

Hassan, Tarek, Stephan Hollander, Ahmed Tahoun, and Laurence van Lent, "Aggregate and Idiosyncratic Political Risk: Measurement and Effects," *Available at SSRN 2838644*, 2016.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman, *The elements of statistical learning: data mining, inference, and prediction*, Springer Science & Business Media, 2009.

He, Zhiguo and Arvind Krishnamurthy, "A model of capital and crises," *The Review of Economic Studies*, 2011, *79* (2), 735–777.

_ **and** _ , "Intermediary asset pricing," *American Economic Review*, 2013, *103* (2), 732–70.

_ **and** _ , "Intermediary asset pricing and the financial crisis," *Annual Review of Financial Economics*, 2018, *10*, 173–197.

_ **, Bryan Kelly, and Asaf Manela**, "Intermediary asset pricing: New evidence from many asset classes," *Journal of Financial Economics*, 2017, *126* (1), 1–35.

Hébert, Benjamin and Jesse Schreger, "The costs of sovereign default: Evidence from argentina," *American Economic Review*, 2017, *107* (10), 3119–45.

Herreño, Juan, "The Aggregate Effects of Bank Lending Cuts," Technical Report, Working Paper, Columbia University 2020.

Huber, Kilian, "Disentangling the effects of a banking crisis: Evidence from German firms and counties," *American Economic Review*, 2018, *108* (3), 868–98.

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.*

Kermack, William Ogilvy and Anderson G McKendrick, "A contribution to the mathematical theory of epidemics," *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 1927, *115* (772), 700–721.

Khwaja, Asim Ijaz and Atif Mian, "Tracing the impact of bank liquidity shocks: Evidence from an emerging market," *American Economic Review*, 2008, *98* (4), 1413–42.

Kuttner, Kenneth N, "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market," *Journal of monetary economics*, 2001, *47* (3), 523–544.

Larsen, Vegard and Leif Anders Thorsrud, "Business cycle narratives," 2019.

Levy, Ro'ee, "Social media, news consumption, and polarization: Evidence from a field experiment," *American economic review*, 2021, *111* (3), 831–70.

Loughran, Tim and Bill McDonald, "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *The Journal of Finance*, 2011, *66* (1), 35–65.

Lucas, Robert, "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 1972, *4* (2), 103–124.

Luo, Yulei, Jun Nie, Gaowang Wang, and Eric R Young, "Rational inattention and the dynamics of consumption and wealth in general equilibrium," *Journal of Economic Theory*, 2017, *172*, 55–87.

Mackowiak, Bartosz and Mirko Wiederholt, "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 2009, *99* (3), 769–803.

Malmendier, Ulrike and Stefan Nagel, "Learning from Inflation Experiences," *The Quarterly Journal of Economics*, 2016, *131* (1), 53–87.

Mankiw, Gregory and Ricardo Reis, "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.

Matějka, Filip and Alisdair McKay, "Rational Inattention to Discrete Dhoices: A New Foundation for the Multinomial Logit Model," *American Economic Review*, 2015, *105* (1), 272–98.

McCaulay, Alistair, "Cyclical Attention to Saving," Technical Report, Working Paper 2020.

McCracken, Michael W and Serena Ng, "FRED-MD: A monthly database for macroeconomic research," *Journal of Business & Economic Statistics*, 2016, *34* (4), 574–589.

McKay, Alisdair and Johannes Wieland, "Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy," Technical Report, National Bureau of Economic Research 2019.

Miranda-Agrippino, Silvia and Giovanni Ricco, "The transmission of monetary policy shocks," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 74–107.

Morelli, Juan, Pablo Ottonello, and Diego Perez, "Global Banks and Systemic Debt Crises," Working paper 2019.

Nakamura, Emi and Jón Steinsson, "High Frequency Identification of Monetary Non-neutrality: The Information Effect," *Quarterly Journal of Economics*, 2018.

_ and _ , "Identification in macroeconomics," *Journal of Economic Perspectives*, 2018, *32* (3), 59–86.

Nieuwerburgh, Stijn Van and Laura Veldkamp, "Information Immobility and the Home Bias Puzzle," *The Journal of Finance*, 2009, *64* (3), 1187–1215.

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.

Ottonello, Pablo and Thomas Winberry, "Financial heterogeneity and the investment channel of monetary policy," *Econometrica*, 2020, *88* (6), 2473–2502.

Ramey, Valerie, "Macroeconomic Shocks and Their Propagation," in "Handbook of Macroeconomics," Vol. 2, Elsevier, 2016, pp. 71–162.

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.

Reis, Ricardo, "Inattentive Producers," *The Review of Economic Studies*, 2006, *73* (3), 793–821.

Reutskaja, Elena, Rosemarie Nagel, Colin Camerer, and Antonio Rangel, "Search Dynamics in Consumer Choice under Time Pressure: An Eye-tracking Study," *American Economic Review*, 2011, *101* (2), 900–926.

Rigobon, Roberto, "Identification through heteroskedasticity," *Review of Economics and Statistics*, 2003, *85* (4), 777–792.

_ **and Brian Sack**, "The impact of monetary policy on asset prices," *Journal of Monetary Economics*, 2004, *51* (8), 1553–1575.

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.

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, "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Siriwardane, Emil, "Limited investment capital and credit spreads," *Journal of Finance*, 2019, (forthcoming).

Spiegler, Ran, "Bayesian networks and boundedly rational expectations," *The Quarterly Journal of Economics*, 2016, *131* (3), 1243–1290.

Tanaka, Mari, Nicholas Bloom, Joel David, and Maiko Koga, "Firm Performance and Macro Forecast Accuracy," *Journal of Monetary Economics*, 2019.

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a String: US Monetary Policy is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 2016, *8* (4), 43–74.

Vavra, Joseph, "Inflation Dynamics and Time-varying Volatility: New evidence and an Ss Interpretation," *The Quarterly Journal of Economics*, 2014, *129* (1), 215–258.

Woodford, Michael, "Information-constrained State-dependent Pricing," *Journal of Monetary Economics*, 2009, *56*, S100–S124.

Zbaracki, Mark, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen, "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," *Review of Economics and Statistics*, 2004, *86* (2), 514–533.

Zou, Hui and Trevor Hastie, "Regularization and variable selection via the elastic net," *Journal of the royal statistical society: series B (statistical methodology)*, 2005, *67* (2), 301–320.