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

Wenting Song — Bank of Canada

Samuel Stern — University of Michigan

July 17, 2024

Abstract

This paper provides empirical evidence of the importance of firm attention to macroeconomic dynamics. We construct a text-based measure of attention to macroeconomic news and document that attention is polarized across firms and countercyclical. Differences in attention lead to asymmetric responses to monetary policy: expansionary monetary shocks raise market values of attentive firms more than those of inattentive firms, and contractionary shocks lower values of attentive firms by less. Attention also mitigates the effects of macroeconomic uncertainty on firm performance. In a quantitative rational inattention model that is calibrated with this new text-based measure, inattention drives monetary non-neutrality. As average attention varies over the business cycle, so does the efficacy of monetary policy.

JEL: D83, E44, E52

Keywords: Rational inattention, monetary policy, natural language processing

*Song (wenting.w.song@gmail.com): Bank of Canada. Stern (sternsa@umich.edu): University of Michigan. We are grateful to John Leahy, Pablo Ottonello, Linda Tesar, and Toni Whited for generous advice and guidance. We thank the Editor, four anonymous referees, Sushant Acharya, Hassan Afrouzi, Martin Eichenbaum, Xing Guo, Andrei Levchenko, Juan Herreño, Yulei Luo, Alistair Macauley, Albert Marcet, Michael McMahon, Oleksiy Kryvtsov, Laura Veldkamp, Choogryul Yang, and the participants of numerous seminars and conferences for useful comments and suggestions. We also thank Kevin MacLean and Joshua Slive for their generous help with the Canadian survey data and Landon Liu for excellent research assistance. The views expressed herein are those of the authors and not necessarily those of the Bank of Canada. First posted version: November 9, 2020.

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) consider firm managers who gather information to maximize value while facing cognitive costs of processing information. Inattention provides an intuitive microfoundation for monetary non-neutrality, 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 some of the first empirical evidence of the importance of firm attention to macroeconomic dynamics using a novel text-based measure. 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 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 study macroeconomic implications of firm attention. Empirically, we find that attention mitigates the effects of macroeconomic uncertainty on firms' long-term performance. Quantitatively, we use this measure to calibrate information costs in a general equilibrium model with rationally inattentive firms. Firm inattention generates monetary non-neutrality and is a source 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 with which managers discuss macro conditions. Our text-based classification of firm attention passes a number of sense checks: topic-specific attention is concentrated by industry; firms in more attentive industries adjust prices faster in response to monetary shocks; and attentive firms predict macro and firm-specific variables better in surveys.

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 referenced macroeconomic news rose notably during recessions. We also study potential drivers of firm attention which include firm characteristics and macroeconomic uncertainty.

Our main empirical result validates our text-based measure by testing for asymmetry in firm performance that is predicted by inattention models: following a macroeconomic shock, the responses of firms with greater information-processing capacity should be closer to the optimal response regardless of the shock's direction. Therefore, attentive firms should exhibit higher profit elasticities in response to positive shocks and lower elasticities in response to negative shocks as they make decisions 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 (CCM, 2020), and high-frequency monetary shocks (constructed as in Gürkaynak *et al.*, 2005; Gorodnichenko and Weber, 2016; Nakamura and Steinsson, 2018).

Consistent with the theoretical prediction, expansionary monetary shocks raise stock returns of attentive firms by nearly 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 is inconsistent with several common 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 produce an asymmetric response.

We then examine how attention affects firm performance under varying degrees of aggregate uncertainty. We construct an uncertainty index based on the Survey of Professional Forecasters and measure performance in three dimensions: profitability, financial performance, and survival. The resulting estimates show that attentive firms outperform their inattentive competitors under increased uncertainty. Interestingly, attention appears to weakly reduce performance in low-uncertainty environments, which may hint at the cost of attention.

Finally, we present a quantitative rational inattention model calibrated using our new measure to study the aggregate implications of incomplete firm attention. In the model, firms with heterogeneous information costs optimally trade off between the precision of their signals of aggregate demand and the cost of acquiring and processing information. Consistent with our empirical findings, attentive firms have higher output semi-elasticities to expansionary monetary shocks and lower semi-elasticities to contractionary shocks. We incorporate observed countercyclicality of firm attention to show that the efficacy of monetary policy declines as the share of attentive firms rises and more firms set prices closer to the optimum. This new interpretation of attention-dependent monetary policy implies that central banks should expect policy interventions 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[1] and consequently the need for empirical measures.[2] Existing research has successfully measured attention in lab experiments (Reutskaja *et al.*, 2011), field experiments (Bartoš *et al.*, 2016; Fuster *et al.*, 2018), and for individual consumers and banks (Macaulay, 2020; Weitzner and Howes, 2021). Our methodology complements those measures as well as survey-based evidence on firm expectations by Tanaka *et al.* (2019), Coibion *et al.* (2018), and Candia *et al.* (2021), and enables researchers to explore questions that lie outside the coverage of existing surveys.

Second, our findings 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 (Maćkowiak and Wiederholt, 2009; Afrouzi, 2020; Yang, 2022), asset prices (Van Nieuwerburgh and Veldkamp, 2009), discrete

[1] See, e.g., Coibion and Gorodnichenko (2015); Coibion *et al.* (2020); Malmendier and Nagel (2016).

[2] See Gabaix (2019) and Maćkowiak *et al.* (Forthcoming) for comprehensive surveys of existing measures of attention.

choices (Matějka and McKay, 2015; Caplin *et al.*, 2019), aggregate dynamics (Maćkowiak and Wiederholt, 2015; Afrouzi and Yang, 2021a), and reconciling micro and macro evidence (Auclert *et al.*, 2020). Our results estimate a substantial cost of information frictions in the US data, providing empirical 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 that are weaker in recessions than in expansions. Vavra (2014), McKay and Wieland (2021), 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.

Fourth, our paper relates to a broader and emerging literature that brings natural language processing techniques to economics. The seminal work of Loughran and McDonald (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 study financial constraints (Buehlmaier and Whited, 2018), central bank communication (Hansen *et al.*, 2018), firm-level political risk (Hassan *et al.*, 2019), inflation expectation formation (Larsen *et al.*, 2021), and uncertainty (Handley and Li, 2020). We contribute to this literature by constructing a set of keyword dictionaries based on macroeconomic news releases that correspond to nine macroeconomic topics. While this paper focuses on attention to monetary policy, our method for measuring attention and its effects can be generalized to other macroeconomic topics.

In a related paper, Flynn and Sastry (2022) independently and contemporaneously develop a text-based measure of firm attention to macroeconomic topics. They show, like we do, that their measure is countercyclical. Whereas we show that the stock prices of more attentive firms rise relative to less attentive firms in response to both positive and negative monetary shocks, they compare firms' labor market choices to those of a neoclassical model with full information and show the gap between model and firm behavior is negatively correlated with firm attention both over time and across firms. Together our papers present compelling evidence that our text-based measures contain information that is useful in predicting economic outcomes and that these predictions are consistent with interpreting these measures as measures of attention.

Road map The rest of the paper proceeds as follows: in Section 2 we describe our methodology for measuring attention and present evidence of the stylized facts listed above; in Section 3 we present a theoretical framework that incorporates attention and exposure to macro shocks and derive the predicted asymmetry; in Section 4 we outline an empirical strategy for testing the effects of attention on expected returns and present our results; in Section 5 we present the mitigating effects of attention on uncertainty; in Section 6 we construct a quantitative model of rational inattention and conduct policy counterfactuals; Section 7 concludes.

2. Textual Measure of Attention

This section presents our measure of firm attention, conducts preliminary validation exercises, and documents stylized facts about attention. We show that cross-industry patterns of our proposed measure and its correlation with price adjustment are consistent with predictions about attention behavior. We then highlight two key stylized facts: aggregate attention moves countercyclically over the sample period, and the majority of firms remain polarized between never and always paying attention. The section concludes with some reflections on the limitations and promise of textual analysis as a tool for measuring attention.

2.1. Data and methodology

10-K filings Our analysis uses all electronically available 10-K filings by publicly listed US companies between 1994 and 2019.[3] Under Regulation S-K, the US Securities and Exchange Commission (SEC) requires all public companies to disclose audited financial statements and a description of business conditions in these filings each year. Companies were phased into mandatory electronic filing between 1993 and 1996, meaning that our sample covers

[3]Our methodology is also well-suited for quarterly 10-Q filings. However, we exclude these filings because they are less descriptive and do not require audited financial statements. We start our sample in 1994, since fewer than ten 10-K filings are available electronically in 1993 at the beginning of the phase-in process, and end our sample in 2019 before the onset of the COVID-19 pandemic.

the universe of filers since 1996.[4] The final sample contains 203,050 documents submitted by 35,765 unique firms. Table 1 summarizes the length of these documents and unique vocabulary used by filers.

Table 1: 10-K filing length and vocabulary size

	N	Mean	Median	SD	5th Pctl	95th Pctl
Total word count	203,050	30,580	25,963	24,013	1,788	73,042
excl. stopwords	203,050	18,879	16,027	15,108	1,127	44,927
Unique word count	203,050	2,430	2,488	2,439	417	4,030
excl. stopwords	203,050	2,334	2,387	2,433	363	3,920

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

Textual measure of firm attention To measure firm attention, we employ dictionary-based frequency counts that identify when firms discuss any of the following nine macroeconomic topics: general economic conditions, output, inflation, 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 Online Appendix Table A.1.

[4] See SEC Release No. 33-7427 for more information about the phase-in process.

We then construct two measures of attention based on these keywords. Attention *prevalence*, d_{it}^k, indicates whether 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:

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

The total word count is generated by following the parsing strategy in Loughran and McDonald (2011): each text is stripped of all numbers and "stop words," such as articles, and then mapped onto a dictionary of all words that appear in our sample of 10-K filings.

We treat *prevalence* as our baseline measure of firm attention in the majority of the paper. Since both measures are closely related, this avoids presenting duplicate results. The prevalence measure is also less susceptible to contamination by changes unrelated to firm attention. For instance, *intensity* will decrease if a firm adds a new appendix to its filing despite no change in its discussion of any topics listed above. Nonetheless, *intensity* is essential for understanding the intensive margin of attention and countercyclical variation documented below.

Table A.2 in Online Appendix A reports summary statistics of the firms that are classified as attentive or inattentive according to our "general" topic. Firms that mention macro keywords tend to be larger—averaging \$7.6 billion in assets compared to \$2.9 billion among firms with no macro discussions—and older by just under four years on average. In contrast, average and median leverage appear fairly similar between the two groups of firms.

2.2. Sense check of the textual measure

This section uses cross-industry variation to test whether our prevalence measure is consistent with predictions of incomplete information models and then assesses how the measure relates to firms' forecast accuracy. We interpret the results as preliminary evidence that

Figure 1: Average industry attention by topic

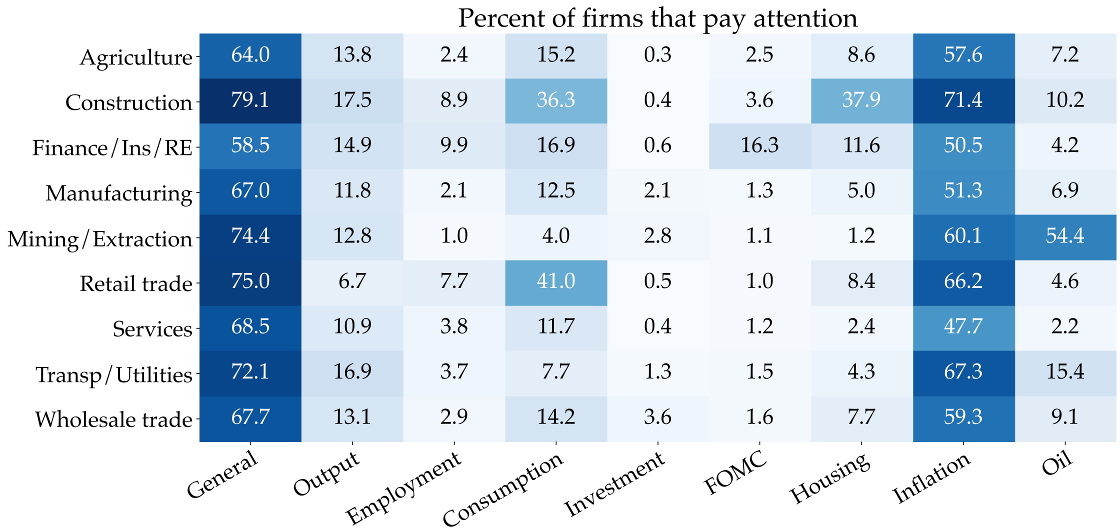

Notes: Heat map of the fraction of firms in an industry that discuss each macroeconomic topic. Industry is defined according to SIC. Darker color represents a higher fraction of firms that mention macro keywords.

the prevalence measure may capture firm attention before presenting firm-level evidence in Section 4.

Cross-industry patterns of prevalence measure We first check whether the prevalence measure for the nine topics listed above is concentrated among commonly associated industries. Figure 1 reports the share of firms in each industry that discuss each topic in their 10-K filings, where industry is defined using 2-digit SIC codes. Since each topic uses a different set of keywords, differences in the prevalence measure across topics may reflect the relative popularity of keywords. Therefore, results should be interpreted across industries rather than across topics.

By and large, the prevalence measure for each topic is highest within related industries: mining/extraction has the highest share of firms that discuss oil prices; retail trade has the highest share of firms that discuss consumption; and the financial sector has the highest prevalence on monetary policy (FOMC). While this cross-industry pattern is not unique to firm attention (for instance, a measure of profit exposure to each topic would produce the same pattern), it serves as a common sense check for our prevalence measure and suggests that textual analysis methods are capable of identifying firm attention.

Price adjustment following monetary shocks We next test whether industries with higher average prevalence adjust prices faster following monetary policy shocks, as predicted for attention in models of incomplete information (Mankiw and Reis, 2002; Maćkowiak and Wiederholt, 2009; Woodford, 2009). The association between prevalence and price response is estimated using the interaction between high-frequency monetary shocks—constructed as in Gorodnichenko and Weber (2016)[5]—and average industry prevalence in a local projection model (Jordà, 2005). Over an h-month horizon, our model takes the form

$$\log P_{s,t+h} - \log P_{s,t} = \alpha_s + \alpha_t + \beta_\nu^h \nu_t^{\mathrm{M}} + \beta_d^h d_{st} + \beta_{d\nu}^h d_{st} \nu_t^{\mathrm{M}} + \Gamma' Z_{st} + \varepsilon_{st}, \tag{1}$$

where $P_{s,t}$ is the BLS Producer Price Index (PPI, 2022) for industry s (4-digit NAICS) in month t, ν_t^{M} denotes the monetary shock in month t, d_{st} denotes average industry prevalence, and Z_{st} is a vector of controls including industrial production, a recession indicator, and industry size. We include industry and time fixed effects, $\{\alpha_s, \alpha_t\}$, and cluster standard errors by both industry and year. For ease of interpretation, monetary shocks are normalized so that positive values correspond to expansionary shocks. We exclude finance and utility industries as is common for estimating firm responses to monetary shocks.[6]

Figure 2 Panel (a) plots the estimated average price response, β_ν^h, and shows that prices rise in a hump-shaped manner following an expansionary monetary shock. At its peak, an unanticipated 25 basis point rate cut causes prices to rise by 1.8%. Panel (b) plots the marginal effect of average prevalence on an industry's price response, $\beta_{d\nu}^h$. Industries with a higher fraction of firms mentioning macro keywords raise prices faster in the first 10 months after a monetary shock, though the effect begins to decline after about seven months as the other industries catch up. This result is consistent with imperfect information models that predict faster price adjustment by attentive firms (e.g., Maćkowiak and Wiederholt, 2009).

Survey forecast accuracy The most direct test of our prevalence measure is whether it can predict a firm's forecasting accuracy. This can be implemented for a very limited subset of our sample that overlaps with a quarterly survey conducted by the Bank of Canada. The

[5]See Section 4.1 for a detailed description of their methodology.

[6]See, for example, Ottonello and Winberry (2020), Acharya *et al.* (2020), and Cloyne *et al.* (2023).

Figure 2: Prevalence measure and price adjustment

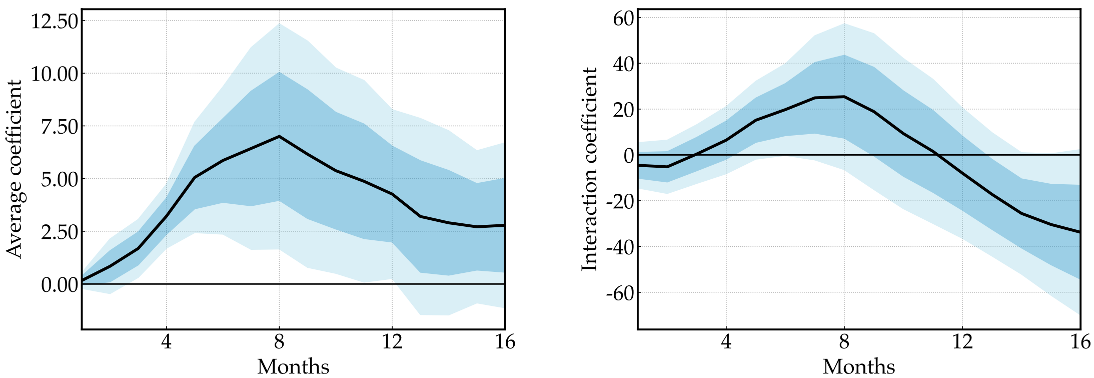

(a) Average responses to monetary policy **(b)** Marginal effects of attention

Notes: Panels (a) and (b) report the average and marginal coefficients, β_{ν}^{h} and $\beta_{d\nu}^{h}$, respectively, from estimating Equation (1) over months $h = 1, \cdots, 16$. We exclude finance and utility industries. Standard errors are double clustered by industry and year. Confidence intervals of 65% and 90% are reported. We have normalized the sign of monetary shocks so that positive shocks are expansionary.

Business Outlook Survey (BOS, 2022) began in 1997 and interviews senior managers about macroeconomic and firm-specific conditions. It covers 100 firms every quarter based on quota sampling by industry, region, and size (Amirault *et al.*, 2020).

Table 2: Attention and survey forecast accuracy

Panel A: Macro forecasts (2-year ahead inflation)

	General		**Inflation**		**Monetary**		
	Attn	Inattn	Attn	Inattn	Attn	Inattn	All
Avg accuracy	41%	31%	45%	32%	43%	40%	40%
N	131	13	92	52	14	130	144

Panel B: Micro forecasts (1-year ahead sales growth)

	General		**Consumption**		**Output**		
	Attn	Inattn	Attn	Inattn	Attn	Inattn	All
Avg accuracy	29%	12%	31%	26%	28%	27%	28%
N	121	8	49	80	39	90	129

Notes: This table reports average forecast accuracy by firm attention. Panel A reports the forecast accuracy of 2-year ahead inflation based on firm responses to BOS question 6.14. "General," "inflation," and "monetary" denote firm attention to the respective topics. Panel B reports the forecast accuracy of 1-year ahead sales growth based on firm responses to BOS question 2.6. "General," "consumption," and "output" denote firm attention to the respective topics.

Over the course of the survey, 137 firms in our sample have appeared in the BOS because they are either a US multinational with a presence in Canada or a Canadian firm listed in the US. Although these firms skew larger and are more likely to mention macro keywords than those in our full sample, they exhibit substantial variation in the prevalence measure across relevant topics.

The BOS includes two questions that pertain to firm forecasts. Question 6.14 asks about inflation expectations over the next two years, and question 2.6 asks about a firm's expected sales growth. The text for each is reproduced here:

> Question 6.14: Over the next 2 years what do you expect the annual rate of inflation to be based on the Canadian Consumer Price Index? (a) between 1–2%, (b) between 2–3%, (c) above 3%, (d) below 1%, (e) NA.
>
> Question 2.6: Over the next 12 months, is your firm's sales volume expected to increase (a) at a lesser rate, (b) the same rate, or (c) a greater rate, as over the past 12 months?

Since responses are multiple choice, we calculate the share of firms whose responses match the realized data and compare these shares between firms that are classified as attentive or inattentive using the prevalence measure. Responses to question 6.14 are compared to annual inflation over the next two years, from the OECD, and responses to question 2.6 are compared to sales volume in the next year, according to Compustat. We report forecast accuracy across all relevant economic topics: attention to general, inflation, and monetary news for forecasting inflation; and attention to general, output, and consumption news for forecasting sales.

Panel A in Table 2 shows that firms that are classified as attentive have more accurate inflation forecasts, and that the accuracy gap is highest for the inflation-specific prevalence measure: the accuracy rate of attentive firms was 13 percentage points higher than that of inattentive firms.

Panel B shows a similar pattern of accuracy when firms predict their own sales growth. Firms that are classified as attentive to aggregate demand (e.g., general, consumption, and output topics) forecast firm-specific demand better, which suggests that these firms translate macro information into better firm planning.

2.3. Stylized facts about firm attention

This section builds upon preliminary evidence that our text-based measures capture firm attention by summarizing how these measures vary over time and between firms, and then exploring potential drivers of firm attention. We document two key stylized facts: aggregate attention moves countercyclically over the sample period and the majority of firms remain polarized between never and always paying attention.

Countercyclical attention to economic conditions Both the share of firms that mention economic keywords and the intensity with which they are discussed vary countercyclically over our sample period. This is illustrated in Figure 3, which plots the annual average *prevalence* and *intensity* measures for the phrase "economic conditions," as well as detrended versions using an HP-filter ($\lambda = 400$).

Panel (a) shows that the share of firms mentioning "economic conditions" has steadily increased since 1994, with particularly rapid growth during the 2001 Recession and the Great Recession. Between 2008 and 2010, aggregate attention jumped by about 15 percentage points and remained elevated for the rest of the sample period. Average intensity in Panel (b) similarly spiked during the Great Recession but declined faster in subsequent years.

We point to these sharp dynamics around the 2001 Recession and the Great Recession as evidence of countercyclical attention, but we also acknowledge that our sample is limited and a longer time series is needed to confirm this pattern. In light of this, future sections use fluctuations in GDP growth and macroeconomic uncertainty—continuous measures related to business cycles—to provide further evidence of countercyclical attention and explore why attention increases in downturns.

Some models with endogenous firm attention predict the countercyclicality displayed in Figure 3. Maćkowiak and Wiederholt (2009) consider imperfect information firms that allocate attention between aggregate and idiosyncratic state variables. Increased aggregate uncertainty (itself countercyclical) induces these firms to shift attention toward aggregate

Figure 3: Time series of attention to "economic conditions"

(a) Prevalence (d_{it}) **(b)** Intensity (s_{it})

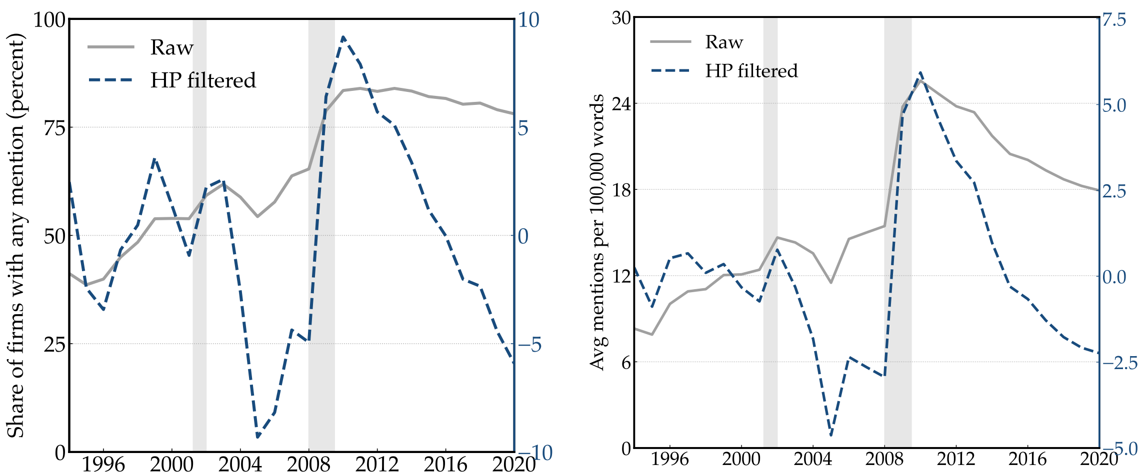

Notes: Time series of firm attention to the keyword "economic conditions." The left panel plots the prevalence measure and reports the share of firms that mention the keyword. The right panel plots the intensity measure and reports the average mentions of the keyword per 100,000 words. "Raw" refers to the unfiltered series, and "HP filtered" refers to the cyclical components of the HP-filtered series with smoothing factor $\lambda = 400$. Shares are reported in percent.

conditions. Chiang (2021) decomposes the impact of lower expected productivity on attention into income and substitution effects. Countercyclical attention emerges among goods-producing agents when their marginal utility rises faster than the returns to attention falls under lower productivity.

Polarization in firm attention Despite the countercyclical dynamics documented above, most firms in our sample are polarized between either always or never discussing economic conditions in their 10-K filings. Figure 4 illustrates this by plotting each firm's share of filings that mention the same key phrase, "economic conditions."[7] The resulting distribution in Panel (a) is heavily concentrated at each extreme, with about three quarters of firms taking values of either 0 or 1. This suggests that most variation in attention occurs across firms rather than within firms and countercyclical variation is caused by a relatively small subset of filers.

[7] In this section, we illustrate patterns of countercyclicality and polarization using attention to economic conditions. Online Appendix Section A.2 reports the times series and histograms of firm attention to all 9 macro topics, which show similar patterns.

Figure 4: Share of filings that mention "economic conditions"

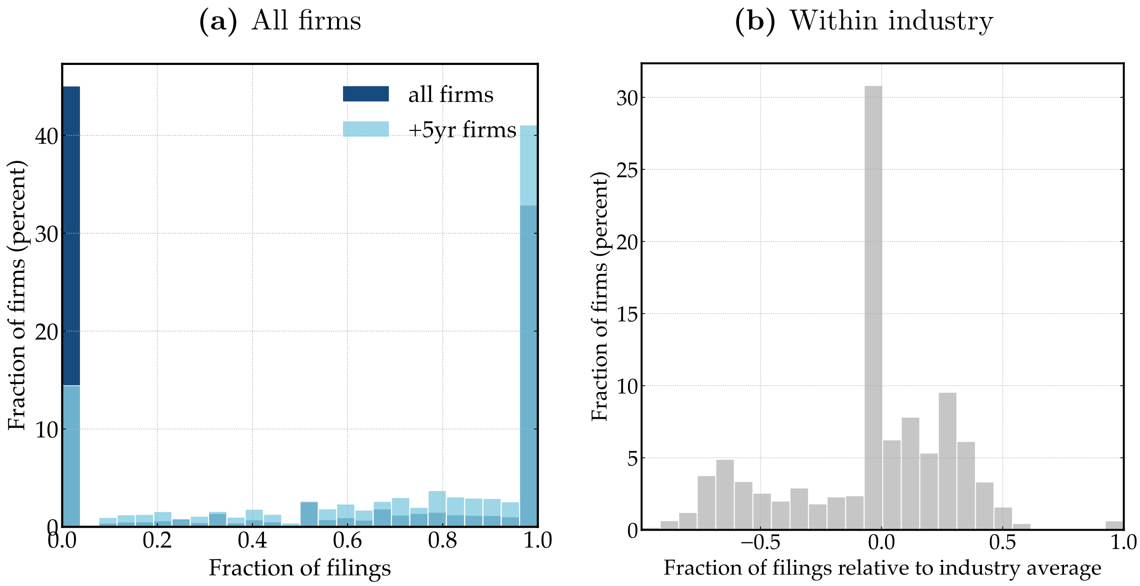

Notes: Histograms of the share of filings by a firm that mention "economic conditions." The left panel shows the histogram of the average fraction of filings that mention the keyword "economic conditions" over the sample period of 1994–2019. Dark blue bars correspond to the distribution of all firms, 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 percentages.

To test whether polarized attention is driven by firms with few filings, we overlay a second histogram in Panel (a) that restricts to firms with at least five years of data. This adjustment greatly reduces the share of firms that never pay attention, yet over half remain polarization between always- and never-attentive firms.

We also test whether cross-industry differences in attention are responsible for polarization. Panel (b) in Figure 4 demeans by industry, which explains approximately one quarter of the attention variation. The distribution now contains a large mass of firms around their industry average (i.e., industries with little attention dispersion), while the remaining firms form a bimodal pattern consistent with polarization.

The presence of any inattentive firms may be surprising given that most US macroeconomic data are readily available. However, this result is consistent with a broader interpretation of attention that includes information processing, communication, and optimal response in addition to information acquisition. As highlighted in Reis (2006), firms likely require significant resources and expertise to process, summarize, and forecast macroeconomic series into sufficient statistics that inform firm decision-making. This is consistent with plant-level evidence from Zbaracki *et al.* (2004). To this end, Abis and Veldkamp (2023) estimate a data production function that uses labor and capital inputs to produce knowledge from unstructured data.

Potential drivers of firm attention To understand what motivates attention and how attentive firms differ from their competitors, we turn to potentially related firm and macroeconomic factors. First, we estimate the relationship between attention and four firm variables—size, age, leverage, and productivity—both cross-sectionally and within-firm. We then examine how attention evolves alongside GDP growth and aggregate beliefs from the Survey of Professional Forecasters. Our findings suggest that size, age, productivity, economic growth, and aggregate uncertainty are important to understanding observed variation in attention.

The relationship between attention and firm characteristics is estimated with the following pair of regressions,

$$\text{Cross-firm variation:} \quad y_{i(j)t} = \alpha_j + \delta_t + \mathbf{X}_{it}\beta + \varepsilon_{it} \tag{2}$$

$$\text{Within-firm variation:} \quad y_{i(j)t} = \nu_i + \delta_t + \mathbf{X}_{it}\beta + \varepsilon_{it}, \tag{3}$$

where $y_{i(j)t} \in \{d_{i(j)t}, \log(s_{i(j)t})\}$ represents attention by firm i in industry j and year t, and $\mathbf{X}_{it}$ is a vector of firm variables including size, age, leverage, filing length, and productivity.[8] Note that our intensity measure, s_{it}, is logged for ease of interpretation. Equation (2) includes time and industry (4-digit NAICS) fixed effects to highlight cross-sectional variation, while Equation (3) uses firm fixed effects to isolate within-firm variation. The first model clusters standard errors by industry and year, and the second model clusters by firm and year.

Data on firm characteristics are from Compustat or firms' 10-K filings. Size is measured as the log of total assets, age as the number of years that a firm has appeared in Compustat, and leverage as the debt-to-asset ratio.[9] Productivity is estimated using the control func-

[8]Existing literature has found each of these characteristics to be relevant for the transmission of macroeconomic policy (Gertler and Gilchrist, 1994; Ottonello and Winberry, 2020; Cloyne *et al.*, 2023).

[9]We exclude observations with leverage greater than 100% (about 3% of the sample) since this ratio is susceptible to extreme values. Filing length is measured as the log of total words appearing in a firm's 10-K. Unlike the other firm variables, it is intended to control for information provision, which affects the likelihood and frequency of keywords, and is therefore not included in the results below.

Table 3: Firm characteristics and attention

	Cross-firm		Within-firm	
	d_{it}	$\log(s_{it})$	d_{it}	$\log(s_{it})$
Size (log total assets)	1.25***	4.25***	0.80	3.47***
	(0.22)	(0.56)	(0.48)	(0.85)
Age	0.00	0.42**	1.78***	3.33***
	(0.07)	(0.18)	(0.04)	(0.17)
Leverage	0.57	6.01*	-1.14	4.20
	(1.77)	(3.31)	(1.41)	(3.32)
Productivity (TFPR)	1.27***	2.89***	1.20***	0.78
	(0.34)	(0.73)	(0.32)	(0.67)
Observations	73101	55276	72283	54365
R^2	0.313	0.290	0.635	0.698
Industry FE	yes	yes	no	no
Firm FE	no	no	yes	yes

Notes: Columns (1) and (2) report estimates for $y_{i(j)t} = \alpha_j + \delta_t + \mathbf{X}_{it}\beta + \varepsilon_{it}$, while Columns (3) and (4) report estimates for $y_{i(j)t} = \nu_i + \delta_t + \mathbf{X}_{it}\beta + \varepsilon_{it}$. The outcome variable, $y_{i(j)t} \in \{d_{i(j)t}, \log(s_{i(j)t})\}$, represents attention by firm i in industry j and year t, and $\mathbf{X}_{it}$ is a vector of firm variables including size, age, leverage, filing length, and productivity. The first two columns include industry fixed effects (4-digit NAICS), and the second two columns include firm fixed effects. All four columns include year fixed effects and control for filing length (log words), though not reported. Outcome variables are scaled by 100 so that units for d_{it} are percentage points and those for $\log(s_{it})$ are percents.

tion approach from Olley and Pakes (1996) and implemented with GMM as in Wooldridge (2009).[10]

Results for this analysis are displayed in Table 3. The first two columns report estimates for Equation (2) and the second two columns do the same for Equation (3). Columns using our prevalence measure, d_{it}, capture changes in attention along the extensive margin, while those using our intensity measure, $\log(s_{it})$, restrict the sample to attentive firms and measure changes on the intensive margin. Both outcome variables are scaled by 100 so that units for d_{it} are percentage points and those for $\log(s_{it})$ are percents.

By and large, we find that larger, older, and more productive firms exhibit higher rates

[10] Firm output is measured as total sales (SALE) deflated by the BEA's implicit price deflator, and labor is defined as total number of employees. Firm capital is constructed using the perpetual inventory method where capital stock for each firm is initialized as Gross Property, Plants, and Equipment (PPEGT), and annual net investment in all subsequent years is defined as the change in Net Property, Plants, and Equipment (PPENT). The capital of each period is defined as the sum of capital from the previous period and net investment. Finally, nominal capital is deflated using the BEA's investment price deflator.

of attention. The association with size and productivity is strongest across firms, while the effect of age appears almost exclusively within-firm. Together, these results suggest that larger and more productive firms pay greater attention to aggregate conditions over time, while smaller and less productive competitors may never do so.

Next, we consider how attention varies with aggregate conditions and beliefs. We estimate the magnitude of countercyclical attention observed in Figure 3 by regressing our attention measures on annual real GDP growth and then see how attention comoves with three measures of aggregate uncertainty: the interquartile range of expectations, the consensus forecast error, and the absolute value of that error.[11] Each measure emphasizes a different dimension of uncertainty. The interquartile range captures disagreement among forecasters, the consensus forecast error considers how attention responds to positive or negative surprises differently, and the absolute forecast error isolates the accuracy of consensus beliefs regardless of error direction.[12]

The relationship between attention and aggregate variables is estimated with the following model,

$$y_{i(j)t} = \nu_i + \delta z_t + \beta x_{it} + \varepsilon_{it}, \tag{4}$$

where $y_{i(j)t} \in \{d_{i(j)t}, \log(s_{i(j)t)})\}$ again represents attention, z_t is our aggregate variable of interest, x_{it} is the 10-K filing length, and ν_i represents a firm fixed effect. Standard errors are clustered by both firm and year.

The resulting estimates are reported in Table 4. Columns 1 and 5 show a strong, countercyclical pattern of attention, while Columns 2 and 6 show a strong positive association between forecaster disagreement and firm attention. Column 4 suggests that higher rates of attentive firms are associated with negative economic surprises, while Column 7 suggests

[11] Given a series, x_t, and a sample of one-period-ahead forecasts, $\hat{x}_{it}$, the interquartile range is defined as $\mathrm{IQR(\hat{x}_{it}) = P_{75}(\hat{x}_{it}) - P_{25}(\hat{x}_{it})}$, where P_Y is the Yth percentile of the forecast sample. The consensus forecast error is defined as $\mathrm{FE(\hat{x}_{it}) = x_t - P_{50}(\hat{x}_{it})}$.

[12] Data on macroeconomic expectations are from the Survey of Professional Forecasters, which has been administered on a quarterly basis since 1968. In each installment, a panel of respondents forecast several economic indicators up to one year in the future. We focus on one-quarter-ahead forecasts for real GDP growth and the unemployment rate. Uncertainty is constructed for each series at a quarterly frequency, standardized over the sample period, and then averaged into an annual composite uncertainty index. Forecast errors for unemployment are inverted so that positive values correspond to positive economic surprises.

Table 4: Aggregate variables and attention

	Prevalence, d_{it}				Intensity, $\log(s_{it})$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
rGDP growth	-0.59***				-5.14***			
	(0.17)				(1.23)			
IQR index		3.29***				14.90**		
		(0.71)				(6.54)		
Abs(FE) index			1.15				24.19***	
			(0.81)				(4.52)	
FE index				-1.58***				-7.16
				(0.24)				(4.26)
Observations	129416	129416	129416	129416	99041	99041	99041	99041
R^2	0.651	0.651	0.651	0.651	0.749	0.747	0.752	0.745
Firm FE	yes	yes	yes	yes	yes	yes	yes	yes
Quadratic Trend	yes	yes	yes	yes	yes	yes	yes	yes

Notes: This table reports estimates of β from the model in (4): $y_{i(j)t} = \nu_i + \delta z_t + \beta x_{it} + \varepsilon_{it}$, where z_t represents either real GDP growth or one of three uncertainty indices: the IQR index, Abs(FE) index, or FE index. The outcome variable $y_{i(j)t}$ represents the attention of firm i in year t, x_{it} is the 10-K filing length, and ν_i represents a firm fixed effect. The first four columns use attention prevalence, d_{it}, as the outcome variable, and the last four columns use log intensity, $\log(s_{it})$. Standard errors are clustered by both firm and year.

that firms pay greater attention when consensus forecasts are less accurate.

2.4. Limitations and promise of textual measures

Boilerplate language is a key concern when using regulatory filings to measure firm attention. 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 (Cao *et al.*, 2020). Moreover, firms likely save time and resources by revising previous filings rather than starting from scratch each year.

The methods used above cannot distinguish between authentic attention to macroeconomic conditions and *cheap talk* references or recycled language that does not reflect current management practices. We address this shortcoming in Online Appendix C.1 by measuring the diversity in filing language with a Jaccard score of lexical similarity and testing whether our main findings are robust in the most linguistically diverse 10-K sections. Table A.8 confirms that our key findings are not driven by the most repetitive and standardized sections.

Even greater measurement error may arise from misidentifying firms as inattentive because they do not mention a certain keyword or discuss economic conditions when such conditions pose a financial risk.[13] *False negatives* can result from methodological limitations or variation in the amount of information that firms choose to disclose. It is worth noting that firm managers are obligated to disclose any material risk factors under SEC Regulation S-K.[14] Those who track inflation, unemployment, or any other topic because they are considered material risk factors are obligated to disclose this to their investors. 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. Online Appendix C.2 uses a Latent Dirichlet Allocation (LDA) unsupervised model to categorize words adjacent to each keyword and produces nine unlabeled "topics" in which keywords appear. Online Appendix C.3 uses the itemized structure of 10-K filings to identify sections that contain the most keywords. This analysis shows that keywords typically appear in sections that discuss firm risk factors (Item 1A) and business operations (Item 7A).

3. Illustrative Framework

This section derives a testable implication of firm attention to address a key identification challenge for our text-based measure: whether it captures *exposure* rather than attention to macroeconomic conditions. We present a stylized model in which firms are heterogeneous in both attention and exposure to an aggregate state variable and then consider how firm outcomes vary with each source of heterogeneity. The model predicts contradictory responses to aggregate shocks under varying attention and exposure, which guides our empirical design

[13] One reviewer compared such firms to smoke detectors, which are (ideally) always on but only beep in the presence of smoke.

[14] Konchitchki and Xie (2022) show that firms are subject to ligitation for undisclosed macroeconomic risks.

in Section 4. The model environment is kept minimal to highlight key mechanisms before Section 6 incorporates more realistic assumptions.

Environment The model 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[15]

$$\hat{\pi}(\hat{s}, \hat{a}) = \pi_s(\bar{s}, \bar{a})\bar{s}\hat{s} + \frac{1}{2}\left(\pi_{ss}(\bar{s}, \bar{a})\bar{s}^2 - \pi_{aa}(\bar{s}, \bar{a})\bar{a}^2\right)\hat{s}^2 + \frac{1}{2}\pi_{aa}(\bar{s}, \bar{a})\bar{a}^2(\hat{a} - \hat{s})^2, \tag{5}$$

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 profit function is concave in the own action, $\pi_{aa} < 0$.

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 actions are more sensitive 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 is consistent with that in rational inattention models such as Sims (2003), which yields a

[15]Under this approximation, $\pi_a(s, a)$ drops out because of the first-order condition and assumption that $a^* = s$ at the optimum. Appendix D.1 contains detailed derivations of the approximation.

steady-state Kalman gain between 0 and 1.[16]

***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 the stock return, 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

$$r = \hat{v} - \hat{v}_{-1}. \tag{6}$$

where lowercase $\hat{v} \equiv \log V - \log \bar{V}$ 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:*

(i) ***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*

[16] In our illustrative framework, a firm's actions are a deterministic function of the aggregate state s, whereas in rational inattention models, there is noise in an agent's signals, which leads to both a Kalman gain between 0 and 1 and noise in the agent's actions conditional on the state.

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

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

(ii) ***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 for 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*

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

Proof. See Online Appendix D.2. ■

Figure 5 illustrates the predictions from Proposition 1. In Panel (a), firms are heterogeneous in their exposure 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 correctly identifies firm attention and then estimate the cost of inattention based on the difference in return elasticities for positive and negative shocks.

Figure 5: Model predictions for exposure vs. attention

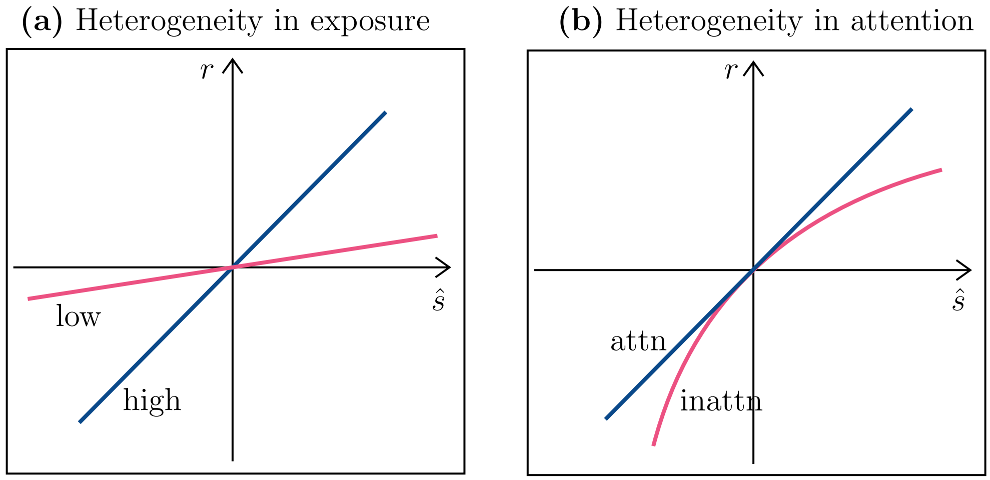

Notes: Illustration of model predictions of return elasticity with respect to aggregate shocks. Verticle axes represent conditional realized return, and horizontal axes represent the magnitude of shocks. The left panel shows return elasticity for firms that are highly exposed to macro conditions (*high*)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

4. Asymmetric Response to Monetary Shocks

We now test the hypothesis that attentive firms respond better to aggregate shocks using a high-frequency identification strategy. Shocks are constructed as plausibly exogenous monetary policy surprises following FOMC announcements, and resulting changes in firm value are measured using stock prices. We use our *prevalence* measure to estimate the relative performance of attentive firms and then test whether they fare better following both positive and negative shocks.[17] Results in this section serve the dual purpose of validating our text-based attention measure and quantifying the expected benefits of attention to economic conditions.

Stock prices are a particularly informative outcome variable because they are forward-looking and similarly high frequency as our monetary shocks. The cumulative effect of a rate surprise on expected future profits will be reflected quickly in a firm's stock price. By restricting to a narrow window around the shock, we isolate this price effect while avoiding

[17]This testable implication from Section 3 works for any aggregate shock with a related attention measure. We use high-frequency monetary shocks as "proof of concept" because they are familiar and well-identified. See Ramey (2016) for a comprehensive survey of alternative aggregate shocks.

other confounding factors. In comparison, a firm's investment and hiring decisions will be smoothed over a longer horizon and any low-frequency response is confounded by other factors that influence these choices. These limitations are exacerbated by the low statistical power of high-frequency monetary shocks, preventing precise estimates of investment and hiring responses.[18]

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.

4.1. Data

Monetary policy shocks are constructed using the high-frequency identification strategy developed by Cook and Hahn (1989) and Gürkaynak *et al.* (2005) and used more recently in Gorodnichenko and Weber (2016), Nakamura and Steinsson (2018), 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

$$\nu_t = \frac{D}{D-\tau}(ff^0_{t+\Delta t^+} - ff^0_{t-\Delta t^-}), \tag{7}$$

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 15 minutes before and 45 minutes after the announcement, D is the number

[18]See Nakamura and Steinsson (2018) for further discussion of this "power problem."

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 (2018) for monetary shocks from 1994 to 2014. For easier interpretation of our empirical results, we normalize the sign of the monetary shock so that a positive shock is expansionary (corresponding to a decrease in interest rates).

Firm outcome and control variables are constructed using CRSP and Compustat data (CCM, 2020). Daily stock returns are measured as the open-to-close change in stock prices on the day of an FOMC announcement. Firm size, age, and industry controls are constructed as described in Section 2.3.

Firm attention is measured using the *prevalence* measure d_{it}, described in Section 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 precludes the possibility that firms are identified as attentive to the FOMC announcement that inspired their attention.

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

$$\begin{aligned} r_{it} = {} & \beta_d d_{it} + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+} \nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-} \nu_t \mathbb{1}_{\nu_t<0} \\ & + \beta_{d\nu_+}(d_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{d\nu_-}(d_{it}\nu_t \mathbb{1}_{\nu_t<0}) + \delta_j + \delta_j \nu_t + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}, \end{aligned} \tag{8}$$

where d_{it} is the attention prevalence, ν_t is the monetary policy shock, $\mathbb{1}_{\nu_t>0}$ indicates positive monetary policy shocks, $\mathbb{1}_{\nu_t<0}$ indicates negative monetary policy shocks, and X_t is a vector 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 dummies and firm controls to capture the average effects of industry and 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 $\beta_{d\nu_+}$ and $\beta_{d\nu_-}$. The theoretical framework in Section 3 hypothesizes $\beta_{d\nu_+}$ to be positive and $\beta_{d\nu_-}$ 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 $H_0 : \beta_{d\nu_+} = \beta_{d\nu_-}$.

4.3. Empirical results

Our baseline results are reported in Table 5. In the first column, we estimate the effect of high-frequency monetary shocks without our attention measures and find that a 25 basis point expansionary monetary shock is associated with about a 1.4% increase in stock prices. This result is consistent with existing estimates from Gorodnichenko and Weber (2016) and Nakamura and Steinsson (2018). 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 3, which remains agnostic as to the average interaction over the entire range of monetary shocks.

The main results from Equation (8) 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 at 5% significance. Column 4 shows that this result is not driven by outsized monetary surprises during the Great Recession nor unconventional monetary policy at the zero lower bound by ending the sample in 2007.

The *asymmetric* response to positive and negative shocks is inconsistent with alternative interpretations of the textual measure that predict a symmetric effect. The foremost alternative discussed in Section 3 is that the textual measure identifies exposure to monetary

Table 5: Baseline results

	(1) Average	(2) Exposure	(3) Attention	(4) excl. ZLB
Shock	5.55***	3.58		
	(1.16)	(2.34)		
Attention		-0.02	-0.07	-0.04
		(0.05)	(0.06)	(0.06)
Shock × Attn		0.94		
		(0.68)		
Shock × $\mathbb{1}_{\nu_t>0}$			3.80	5.40**
			(2.48)	(2.53)
Shock × $\mathbb{1}_{\nu_t<0}$			-4.04	-1.42
			(3.64)	(3.67)
Shock × Attn × $\mathbb{1}_{\nu_t>0}$			1.87***	1.47**
			(0.66)	(0.68)
Shock × Attn × $\mathbb{1}_{\nu_t<0}$			-6.16**	-6.12*
			(3.02)	(3.11)
Observations	603940	603940	603940	458794
R^2	0.018	0.022	0.025	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			0.017	0.031

Notes: 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). Column (1) reports the average effect of monetary shocks from estimating $r_{it} = \delta_j + \beta_\nu \nu_t + \Gamma' X_t + \varepsilon_{it}$. Column (2) estimates the exposure model $r_{it} = \delta_j + \delta_j' \nu_t + \beta_\nu \nu_t + \beta_d d_{it} + \beta_{d\nu}(d_{it}\nu_t) + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}$. Column (3) estimates the baseline attention model Equation (8): $r_{it} = \beta_d d_{it} + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+} \nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-} \nu_t \mathbb{1}_{\nu_t<0} + \beta_{d\nu_+}(d_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{d\nu_-}(d_{it}\nu_t \mathbb{1}_{\nu_t<0}) + \delta_j + \delta_j \nu_t + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}$, where ν_t is the monetary shock, d_{it} is the prevalence attention measure, δ_j is an industry fixed effect, $\delta_j' \nu_t$ is its interaction with the shock, and X_t contains firm-level controls of size, age and leverage. The vector $X_t \nu_t$ contains the interactions between firm controls and the shock. Column (4) re-estimates Equation (8) on the sample ending in 2007 to exclude the zero lower bound period following the Great Recession. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

shocks rather than attention. Any such symmetric effect would also appear in the interaction coefficient $\beta_{d\nu}$ in Column 2, which is only weakly positive. Online Appendix B.2 further shows that directly estimating and controlling for exposure to monetary shocks leaves our main findings unchanged.

Suboptimal responses to monetary policy by inattentive firms reported in Table 5, together with the large fraction of inattentive firms documented in Figure 4, provide some of the first direct evidence of the empirical consequences of firm inattention in the US. We estimate that inattentive firm returns rise by nearly 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 5.6%.

Alternative sources of asymmetry We now consider alternative explanations for the asymmetric price response documented above. Each explanation is tested by augmenting our baseline model to include interaction terms for a confounding variable, c_{it}, that match those for firm attention, d_{it}. The resulting "horse-race" model takes the form

$$\begin{aligned} r_{it} = &\delta_j + \delta_j \nu_t + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+} \nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-} \nu_t \mathbb{1}_{\nu_t<0} + \left[\beta_d d_{it} + \beta_{d\nu_+}(d_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{d\nu_-}(d_{it}\nu_t \mathbb{1}_{\nu_t<0})\right] \\ &+ \left[\beta_c c_{it} + \beta_{c\nu_+}(c_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{c\nu_-}(c_{it}\nu_t \mathbb{1}_{\nu_t<0})\right] + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}, \end{aligned} \tag{9}$$

where, as in the baseline specification, we control for industry fixed effects, industry-specific responses to monetary shocks, a vector of firm controls and their interaction with monetary shocks. If the main result, $\beta_{d\nu_-} < 0 < \beta_{d\nu_+}$, holds true, then we rule out c_{it} as a confounding source of asymmetry.

The first factor considered is productivity. Van Nieuwerburgh and Veldkamp (2006) present a model in which higher productivity increases learning as well as production. If productivity determines both information acquisition and the response to aggregate shocks, it could explain the asymmetric result found above. Productivity is constructed as above in Section 2.3.

Management quality is another potential confounder that could explain both attention and firm performance. Effective managers who capitalize on opportunities during expansionary shocks and mitigate losses from contractionary shocks will generate the same asymmetric performance pattern documented in our main results. We approximate a firm's management quality using the share of board members who hold a graduate degree since existing research documents a strong relationship between education and management quality (Bloom

Table 6: Controlling for alternative explanations of asymmetry

	Productivity (LTFP)	Mgmt Quality	Profit (ROA)	Filing Length
Shock $\times \mathbb{1}_{\nu_t>0}$	6.24**	-0.21	3.38	-2.38
	(2.55)	(2.16)	(2.43)	(4.95)
Shock $\times \mathbb{1}_{\nu_t<0}$	-1.39	-10.34**	-4.54	21.52
	(3.78)	(4.08)	(3.58)	(13.85)
Attention	-0.12	-0.09	-0.07	-0.08*
	(0.09)	(0.06)	(0.06)	(0.05)
Shock $\times$ Attn $\times \mathbb{1}_{\nu_t>0}$	2.46**	2.29***	1.85***	1.68***
	(1.20)	(0.79)	(0.66)	(0.57)
Shock $\times$ Attn $\times \mathbb{1}_{\nu_t<0}$	-6.82***	-8.11**	-6.09**	-5.51*
	(1.90)	(3.23)	(3.00)	(2.79)
Control Var	0.04***	-0.05	0.04*	-0.01
	(0.01)	(0.06)	(0.02)	(0.04)
Control $\times$ Shock $\times \mathbb{1}_{\nu_t>0}$	-0.08	1.80**	-1.37	0.66*
	(0.14)	(0.75)	(1.60)	(0.38)
Control $\times$ Shock $\times \mathbb{1}_{\nu_t<0}$	-0.18	-3.56	-9.02**	-2.70*
	(0.22)	(2.76)	(4.22)	(1.56)
Observations	386094	337927	603065	603940
R^2	0.026	0.039	0.026	0.026
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	no
Wald test p-value: Attention	0.000	0.003	0.018	0.022
Wald test p-value: Control	0.473	0.071	0.114	0.042

Notes: This table augments Column (3) of Table 5 to control for four potential confounding sources of asymmetry. The estimated regression has the form $r_{it} = \delta_j + \delta_j \nu_t + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+} \nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-} \nu_t \mathbb{1}_{\nu_t<0} + \beta_d d_{it} + \beta_{d\nu_+}(d_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{d\nu_-}(d_{it}\nu_t \mathbb{1}_{\nu_t<0}) + \beta_c c_{it} + \beta_{c\nu_+}(c_{it}\nu_t \mathbb{1}_{\nu_t>0}) + \beta_{c\nu_-}(c_{it}\nu_t \mathbb{1}_{\nu_t<0}) + \Gamma_1' X_t + \Gamma_2' X_t \nu_t + \varepsilon_{it}$, where c_{it} represents the alternative "control" variable. As with attention, the control variable is interacted with both positive and negative monetary shocks. All other features of the model specification remain unchanged from Table 5. The four control variables considered are (1) firm productivity estimated as in Olley and Pakes (1996), (2) management quality approximated with board member educational attainment, (3) profit measured as earnings before extraordinary items over total assets, and (4) filing length measured as the log word count of the 10-K filing. The final two rows report p-values of Wald tests for $H_0 : \beta_{d\nu_+} = \beta_{d\nu_-}$ and $H_0 : \beta_{c\nu_+} = \beta_{c\nu_-}$, respectively. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

and Van Reenen, 2010).[19] Data on the educational attainment is from BoardEx (BoardEX, 2021), which covers publicly traded US firms.

[19] Graduate degrees include MBA, MS, MSC, MA, JD, MD, MPA, MSE, PHD, and any degree names that include "master" or "doctor."

The third variable considered is a firm's financial performance measured using return on assets (ROA). Managers may feel compelled to cite macroeconomic conditions when explaining recent performance, and a tendency for well-performing firms to cite such conditions could generate the asymmetry observed.

Finally, we control for the length of a firm's SEC filing as a measure of its preference for information provision. Longer filings—measured using log word count—offer more opportunities for managers to mention macro keywords and signal commitment to due diligence. If thorough due diligence engenders investor confidence, then stocks should perform better following either positive or negative monetary shocks.

Table 6 reports the estimates for Equation (9) using each factor described above: productivity, management quality, profitability, and filing length. As in our baseline results, attentive firms experience a larger increase in market value following an expansionary shock and a smaller contraction following a contractionary shock. The estimates are statistically significant, with similar magnitudes as those in Table 5. While some of the control variables (e.g., filing length) also display an asymmetric effect on firms' responses to monetary shocks, the explanatory power of firm attention remains under all specifications. All four Wald tests for $H_0 : \beta_{d\nu_+} = \beta_{d\nu_-}$ are rejected at 5% significance. In Online Appendix Table A.3, we show that these results are also robust to excluding zero-lower-bound periods.

Additional robustness checks Further robustness analysis pertaining to the identification of high frequency monetary shocks can be found in Online Appendix B. Appendix B.3 controls for the information effect of FOMC announcements using Greenbook forecast revisions, and Appendix B.4 tests whether aggregate conditions confound the estimated effect of high frequency shocks. In each case, our main results remain robust.

5. Attention, Performance, and Aggregate Uncertainty

This section explores how attention affects firm performance under varying levels of aggregate uncertainty. One implication of our illustrative model is that the performance gap between attentive and inattentive firms widens with the magnitude of nominal demand shocks. Returns to attention should therefore increase in periods of greater uncertainty and larger shocks.[20]

We test this prediction by estimating the interaction effect between attention and uncertainty on firm performance. Aggregate uncertainty is measured using the interquartile range of quarterly forecasts for real GDP, inflation, and unemployment from the Survey of Professional Forecasters (SPF, 2023). Each series is standardized over our sample period (1994–2019) and then averaged into a composite uncertainty index.

Firm performance is measured along three dimensions: profitability, financial performance, and survival. Profitability is measured as a firm's return on assets (ROA), which we construct using earnings before extraordinary items over total assets. Financial performance is measured as return on equity (ROE) using earnings before extraordinary items over market capitalization. Finally, survival is defined as whether a firm remains in operation in the next year. Each variable is constructed using annual Compustat data, and ROA and ROE are winsorized at 1%.

Our regression model takes the form

$$y_{it} = \alpha_j + \beta d_{it} + \delta \sigma_t + \gamma d_{it} \cdot \sigma_t + \Gamma' Z_{it} + \varepsilon_{it}, \tag{10}$$

where y_{it} represents one of the three performance variables defined above, d_{it} is our binary attention measure, σ_t is aggregate uncertainty, α_j captures industry fixed effects with 4-digit NAICS, and Z_{it} is a vector of firm controls including size, age, and 10-K filing length (as previously defined). Standard errors are clustered by both year and industry. We extend the model to future outcomes, $y_{i,t+h}$, to capture any lagged effects of attention on performance.

Results from this analysis are reported in Table 7. On average, aggregate uncertainty reduces profitability, financial performance, and the probability of survival, which is consistent with existing models of uncertainty (e.g., Bloom *et al.*, 2007).[21] Attention to macroeconomic conditions, however, mitigates the negative effects of uncertainty: in periods of high uncertainty, attentive firms have higher profitability, better financial performance, and a higher probability of survival. Interestingly, the first row in Table 7 suggests that attention reduces

[20] See Appendix D.3 for an extended illustrative framework that incorporates time-varying uncertainty.

[21] See Leahy and Whited (1996) for a general discussion of firm decisions under uncertainty.

Table 7: Effects of attention on firm performance under uncertainty

	ROE		ROA		Surivival	
	(1) Impact	(2) Peak	(3) Impact	(4) Peak	(5) Impact	(6) Peak
Attention (general)	-0.02*	-0.02**	-0.02*	-0.03**	-0.01	-0.02**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Uncertainty (SPF IQR)	-0.03***	-0.02*	-0.05***	-0.04*	-0.01	-0.03***
	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.01)
Attention × Uncertainty	0.03***	0.03***	0.05***	0.06***	0.01	0.03**
	(0.01)	(0.01)	(0.02)	(0.02)	(0.01)	(0.01)
Observations	104507	92023	110267	97180	111637	66813
R^2	0.163	0.156	0.247	0.236	0.034	0.028
Clustered SE	yes	yes	yes	yes	yes	yes
Firm controls	yes	yes	yes	yes	yes	yes
4-digit NAICS FE	yes	yes	yes	yes	yes	yes

Notes: The table reports results from estimating (10), $y_{it}^h = \alpha_j + \beta d_{it} + \delta \sigma_t + \gamma d_{it} \cdot \sigma_t + \Gamma' Z_{it} + \varepsilon_{it}$, for horizons $h = 1, \cdots, 5$. The dependent variables y_t include (i) profitability measured with ROA (i.e., net income over total assets), (ii) financial performance measured with ROE (i.e., net income over equity), and (iii) an indicator variable for firm survival. Independent variables include the prevalence attention to general economic conditions, d_{it}; macroeconomic uncertainty, σ_t^2, measured as the interquartile range of quarterly growth rate forecasts for real GDP and unemployment from the SPF; the interaction between attention and uncertainty; industry fixed effects δ_j; and firm controls, Z_{it}. We standardize the interquartile range of each series over our observed sample period, take the absolute average deviation each quarter, and then average these quarterly values each year. The on-impact effect corresponds to the estimates for $h = 1$. The peak effect corresponds to the largest estimated marginal effect over the 5-year horizon. Standard errors are clustered at the shock level and reported in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

firm performance under low uncertainty, consistent with models of imperfect information in which firms face a cost of attention and reap the benefit in states with large realized shocks (such as Reis, 2006).[22] Online Appendix Table A.7 further interacts attention with recession indicators and shows that attention improves firm performance mainly by reducing uncertainty.

Section 4 showed that attentive firms respond better to monetary shocks. This section finds that these same firms outperform less attentive competitors under elevated aggregate

[22] Related to our findings that attentive firms appear to be "better opportunists" with state-dependent outperformance, Ahnert *et al.* (2021) find that banks with better information technology are more productive and spur better job creation and innovation. Furthermore, Kwon *et al.* (2022) find that industry concentration rises with investment intensity in information technology and research and development. Attentive firms with better information-processing technologies are better equipped to react to evolving macroeconomic conditions, which may have contributed to the rise of "superstar" firms (Autor *et al.*, 2020).

uncertainty. Together, they paint a picture of attentive firms as more responsive to evolving macroeconomic conditions and highlight the benefits gained for their diligence.

6. Quantitative Model

Our attention measure can inform model-based analysis in addition to the new empirical findings above. This section presents a quantitative model in which inattention to aggregate conditions drives monetary non-neutrality. Both the rate of attentive firms and the cost of inattention are calibrated using the prevalence measure presented in Section 2. This model demonstrates the importance of attention by showing that the efficacy of monetary policy depends on aggregate attention when firms face information frictions.

6.1. Model environment

We start with a canonical dynamic general-equilibrium model with rationally inattentive firms as in Maćkowiak and Wiederholt (2009) and Afrouzi and Yang (2021a). Time is discrete and infinite. The economy consists of a representative household, heterogeneous firms, and a central bank. Households and the central bank have full information about the economy, while firms pay a cost proportional to information obtained (measured using Shannon mutual information as in Sims, 2003). Firms differ ex-ante in their marginal costs of information, which is motivated by the heterogeneity documented in Section 2.3.

Household A representative household consumes a bundle of goods over the continuum of varieties $i \in [0, 1]$ and supplies labor, N_t, in a competitive labor market with wage, W_t. In addition to the wage income, the household has access to a one-period bond, D_t, with the interest rate ι_t and receives firms' profits, Π_t. The household maximizes its life-time utility:

$$\max_{\{C_{it},D_t,N_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (\log C_t - \psi N_t), \tag{11}$$

$$\text{s.t.} \int_0^1 P_{it} C_{it} di + D_t \leq W_t N_t + (1 + \iota_t) D_{t-1} + \Pi_t,$$

where consumption, C_t, is aggregated over each good type i with a CES aggregator, $C_t = \left(\int_0^1 C_{it}^{\frac{\varepsilon-1}{\varepsilon}} di\right)^{\frac{\varepsilon}{\varepsilon-1}}$, and ε is the elasticity of substitution. Let $Q_t \equiv P_t C_t$ denote nominal aggregate demand. The household's optimal choices are given by the following three conditions

$$C_{it} = C_t(P_{it}/P_t)^{-\varepsilon}; \qquad 1 = \beta(1+\iota_t)\,\mathbb{E}_t(Q_t/Q_{t+1}); \qquad W_t = \psi Q_t. \tag{12}$$

Central bank The central bank targets the aggregate money supply, $P_t C_t$, similar to Caplin and Spulber (1987) and Gertler and Leahy (2008). Nominal aggregate demand, therefore, follows

$$\Delta \log Q_t = \rho \Delta \log Q_{t-1} + \nu_t, \quad \nu_t \sim N(0, \sigma_\nu^2). \tag{13}$$

Firms There is a unit measure of monopolistically competitive firms, indexed by $i \in [0,1]$. Firms operate a decreasing-returns-to-scale production technology with labor as its only input: $Y_{it} = N_{it}^{\gamma}$. They take wage and demand as given and can flexibly set prices, P_{it}, based on their information set in period t. After setting prices, they hire labor from a competitive labor market to produce the realized level of demand induced by their prices.

Firms are assumed to be *rationally inattentive*, meaning that they do not observe shocks to aggregate demand and endogenously acquire information about Q_t. In each period, firm i starts with their information set from the previous period, S_i^{t-1}, and selects the stochastic process for their new signal, s_{it}, from a set of available signals, $\mathcal{S}^t$, that vary in cost and precision. These signals satisfy the properties outlined in Definition 3.

***Definition* 3 (set of available signals).** *The set of available signals, $\mathcal{S}^t$, consists of all signal processes satisfying the following three properties:*

i. $\mathcal{S}^t$ is rich: for any posterior distribution on $\{Q_t\}_{t\geq 0}$, there is a set of signals $S^t \in \mathcal{S}^t$ that generate that posterior;

ii. Signals do not expire over time: $\mathcal{S}^t \subset \mathcal{S}^{t+h}$ for $h \geq 0$;

iii. Signals contain no information about future shocks: $S_t \perp Q_{t+h}$ for $S_t \in \mathcal{S}^t$ and $h \geq 1$.

We assume that the cost of information is linear in the Shannon mutual information:

$$2\omega_i \cdot \mathcal{I}(Q_t; s_{it}|S_i^{t-1}), \tag{14}$$

where the Shannon mutual information, $\mathcal{I}(Q_t; s_{it}|S_i^{t-1})$, measures the expected reduction in uncertainty about aggregate demand from observing the signal.[23] A more precise signal requires a higher flow of mutual information and is therefore more expensive.

The information represented by $\mathcal{I}(\cdot)$ can be thought of as a firm's *attention* to the economy: for each unit of mutual information (or nat), a firm pays a marginal cost, ω_i, and reduces its expected uncertainty about aggregate demand. A firm's information set evolves according to $S_i^t = S_i^{t-1} \cup s_{it}$.

Firms are ex-ante heterogeneous in their information-processing technology and face either high or low marginal costs of attention

$$\omega_i \in \{\omega_H, \omega_L\}. \tag{15}$$

A fraction $\theta \in (0,1)$ of firms are assumed to have low information-processing costs, while all remaining firms face high costs.

Firms maximize expected profits by choosing the stochastic process of the set of signals to observe over time, $\{s_{it} \in \mathcal{S}_{it}\}_{t\geq 0}$, and a pricing strategy, $P_{it}(S_i^t)$, that depends on its information set at time t containing realizations of current and past signals. The firm's problem is given by

$$\max_{\{s_{it}\in\mathcal{S}_{it}, P_{it}(S_i^t)\}_{t\geq 0}} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \frac{1}{P_t C_t}\Big(\underbrace{(P_{it}Y_{it} - W_t N_{it})}_{\text{operational profits}} - \underbrace{2\omega_i \mathcal{I}(Q_t; s_{it}|S_i^{t-1})}_{\text{information costs}}\Big) \,\Big|\, S_i^{-1}\right] \tag{16}$$

$$\begin{aligned} \text{s.t.} \quad & Y_{it} = Y_t(P_{it}/P_t)^{-\varepsilon} && \text{(demand for goods)} \\ & Y_{it} = N_{it}^{\gamma} && \text{(production technology)} \\ & S_i^t = S_i^{t-1} \cup s_{it}, && \text{(evolution of information)} \end{aligned}$$

[23] Formally, the Shannon mutual information between random variables X and Y is defined as $\mathcal{I}(X;Y) = \int_Y \int_X p(x,y) \log \frac{p(x,y)}{p(x)p(y)} dxdy$, which measures the difference between conditional and marginal entropies. See Cover and Thomas (2006) for details.

where Y_t is the aggregate output, P_t is the aggregate price index, P_{it} is firm i's price, Y_{it} is the demand for the firm's good, N_{it} is the firm's labor demand.

Equilibrium Given the exogenous process for aggregate demand, $\{\Delta \log Q_t\}_{t\geq 0}$, the equilibrium consists of an allocation for the household, $\Omega^H = \{C_t, D_t, N_t, (C_{it})_{i\in[0,1]}\}_{t\geq 0}$, allocations for every firm $i \in [0,1]$ given their initial information sets S_i^{-1}, $\Omega_i^F = \{s_{it} \in \mathcal{S}_{it}, P_{it}, N_{it}, Y_{it}\}_{t\geq 0}$, a set of prices $\{\iota_t, P_t, W_t\}_{t\geq 0}$, and a stationary distribution over firms' states such that

i. Given the set of prices and firms' allocations, the household's allocation solves the problem in Equation (11);

ii. Given the set of prices and the household's allocation, firms' allocations solve the problem in Equation (16);

iii. All markets clear, that is, for $t \geq 0$ and $i \in [0,1]$, $D_t = 0$, $Y_{it} = C_{it}$, $Y_t = C_t$, and $N_t = \int_0^1 N_{it} di$.

Solution We approximate a firm's flow profits with second-order log approximations around the full-information steady state.[24] A firm's total value under log approximation, v, is decomposed 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 firm value under imperfect information.

Let lowercase letters denote log deviations from the steady state. The imperfect-information value is given by

$$\tilde{v} = \max_{\{s_{it}\in\mathcal{S}_{it}, p_{it}(S_i^t)\}_{t\geq 0}} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \Big(-B(p_{it} - p_t^*)^2 - 2\omega_i \mathcal{I}(q_t; s_{it}|S_i^{t-1})\Big) \Big| S_i^{-1}\right] \tag{17}$$

$$\text{s.t.} \quad p_t^* = \alpha p_t + (1-\alpha) q_t$$

$$S_i^t = S_i^{t-1} \cup s_{it},$$

[24] Appendix E.1 contains details of the approximation. Log-quadratic approximation is a common simplifying assumption in rational inattention models (see, e.g., Maćkowiak and Wiederholt, 2009; Afrouzi and Yang, 2021a). Sims (2003) shows the optimal distribution under Gaussian priors and quadratic payoffs is Gaussian, so log-quadratic approximation of the profit function greatly reduces the dimensionality the problem.

where $\alpha \in (0, 1)$ and $B > 0$ are constants that depend on non-information-friction parameters and relate to the degree of strategic complementarity and the curvature of the profit function, respectively.[25] p_t^* denotes the optimal price under perfect information. Since prices are fully flexible, price setting with perfect information is a static problem (see Online Appendix E.2).

The imperfect information problem in (17) is solved numerically based on the algorithm for dynamic rational inattention problems (`DRIPs`) developed in Afrouzi and Yang (2021a). Appendix E.3 provides detailed information on its implementation.

6.2. Calibration

Model parameters are divided into two sets: those that govern information frictions and all remaining parameters. In the first set, the share of attentive firms and relative cost of information between firms are calibrated to match two empirical moments using our text-based measure of attention. The cost of information among attentive firms, ω_L, is set near zero so that attentive firms have nearly full information. The second set of non-information parameters are calibrated to external sources or estimates using quarterly data on US output from the Bureau of Economic Analysis. A summary of all model parameters can be found in Online Appendix Table A.9.

For non-information parameters, we calibrate the model quarterly and set the discount rate to be $\beta = 0.96^{1/4}$. The stochastic process for aggregate demand, $\{\rho, \sigma_\nu\}$, is estimated using quarterly US nominal manufacturing output between 1994 and 2019. Restricting to the manufacturing sector is consistent with the within-sector results presented in our empirical analysis. The elasticity of substitution is set to $\varepsilon = 10$, implying a steady-state markup of 11%, and the disutility of labor is set to $\psi = 0.90$. Finally, we set returns to scale $\gamma = 0.93$ according to the estimate by Basu and Fernald (1997) for the US manufacturing sector.

For information parameters, the share of firms with low information costs, θ, is set to 65% to match the share of attentive firms in Figure 3. As in Maćkowiak *et al.* (2009) and Afrouzi and Yang (2021b), attention depends inversely on the ratio between attention costs and the curvature of the profit function, ω_i/B. A firm pays greater attention when the information cost is low or when its incentives to pay attention are high. We focus on calibrating

[25] See Equations (24) and (25) in the appendix.

information parameters and therefore fix the curvature of profit function, which depends on non-information parameters.

We set ω_L close to zero to reflect the assumption that firms with low information costs have nearly full information. The relative cost of information for high-cost firms, $\omega_H - \omega_L$, is calibrated to match the heterogeneous responses to monetary shocks estimated in Table 5. Stock returns in the model are defined as the log change in a firm's value, $r_{it} = \log V_{it} - \log \mathbb{E}_{t-1}(V_{it})$.[26] To connect Shannon mutual information, $\mathcal{I}_{it}$, in the model with the text-based attention measure, we assume that the frequency of macro keywords in 10-K filings is strictly increasing in firm attention. This allows us to match the cross-sectional distribution of firm attention without explicitly modelling the writing process of 10-K filings. Since our main empirical analysis uses the prevalence attention measure, we define a corresponding indicator variable, $d_{it} = \mathbb{1}(\mathcal{I}_{it} > \bar{\mathcal{I}}_t)$, for firms whose attention is above the cross-sectional mean in a given period. 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. With the simulated data, we estimate

$$r_{it} = c + \beta_1 \mathbb{1}_{\nu_t>0} + \beta_{\nu_+}\nu_t \mathbb{1}_{\nu_t>0} + \beta_{\nu_-}\nu_t \mathbb{1}_{\nu_t<0} + \beta_d d_{it} + \boldsymbol{\beta_{d\nu_+}} d_{it}\nu_{it}\mathbb{1}_{\nu_t>0} + \boldsymbol{\beta_{d\nu_-}} d_{it}\nu_{it}\mathbb{1}_{\nu_t<0} + \varepsilon_{it}.$$

We set $\omega_H - \omega_L$ to target the elasticity $\frac{1}{2}|\hat{\beta}_{d\nu_+}| + \frac{1}{2}|\hat{\beta}_{d\nu_-}|$ from Column 3 in Table 5, which measures the relative stock return losses of firms that do not pay attention. Online Appendix Figure A.7a shows how the parameter is identified. As ω_H increases and the gap between ω_L and ω_H widens, the simulated elasticity monotonically increases, implying greater heterogeneity between attentive and inattentive firms. The resulting calibration for $\omega_H - \omega_L$ is 1.13 per nat.

[26] A firm's value function in (16) can be expressed in recursive form as

$$V(S_i^{t-1}) = \max_{\{s_{it}\in\mathcal{S}_{it}, P_{it}(S_i^t)\}_{t\geq 0}} \mathbb{E}_t\left[\frac{1}{P_t}\Big((P_{it}Y_{it} - W_t N_{it}) - 2\omega_i \mathcal{I}(Q_t; s_{it}|S_i^{t-1})\Big) + \beta\Lambda_{t,t+1}V(S_i^t)\,\Big|\, S_i^{t-1}\right]$$
$$\text{s.t.}\quad Y_{it} = Y_t(P_{it}/P_t)^{-\varepsilon},\quad Y_{it} = N_{it}^{\gamma},\quad S_i^t = S_i^{t-1}\cup s_{it}.$$

Discussion of the calibration strategy One primary challenge to calibrating a rational inattention model is that information costs—which determine the degree of information frictions—are unobserved in the data. Existing studies have successfully calibrated rational inattention parameters by matching moments related to aggregate consumption dynamics and monetary policy responses, while other have calibrated these parameters using survey data (Luo, 2008; Maćkowiak and Wiederholt, 2015, 2023). Our calibration strategy differs in three main ways. First, we allow heterogeneity in information costs and therefore have two parameters for information costs, ω_L and ω_H, instead of a single parameter. We focus on calibrating the heterogeneity in attention costs, $\omega_H - \omega_L$, to study its implication for monetary transmission. In doing so, we set ω_L close to zero, which implies that our calibrated model provides a lower bound on the degree of monetary non-neutrality arising from inattention.

Second, our calibration makes use of the text-based attention measure instead of macro data or survey data. The measure informs attention at the granular firm level, but the tradeoff is that 10-K filings do not have a direct model counterpart. To connect the concept of Shannon mutual information with our text-based attention measure, we need to assume that the frequency of macro keywords in 10-K filings is strictly increasing in firm attention. This allows us to use the textual measure to discipline the cross-sectional distribution of firm attention.

Lastly, the relative cost of attention, $\omega_H - \omega_L$, is calibrated by targeting a micro elasticity (namely, the relative stock return losses of inattentive firms in response to monetary shocks) rather than macro moments. This is possible because our proposed attention measure is available for a large number of firms over a long sample period. It is well-known since Mehra and Prescott (1985) that standard macro models, including ours, are not designed to match the unconditional cross section of stock returns. However, our target moment in Table 5 is the *conditional* responses of firm values to monetary shocks, with stock returns, r_{it}, in Equation (8) capturing log changes in a firm's value. Online Appendix Table A.10 shows that our model matches heterogenous responses of firms' values to monetary shocks through relative information costs.

For further robustness, Online Appendix E.5 implements an alternative calibration strategy that targets industry-level price adjustment estimates from Figure 2. It finds that atten-

Figure 6: Aggregate responses to expansionary monetary shock

(a) Inflation **(b)** Output

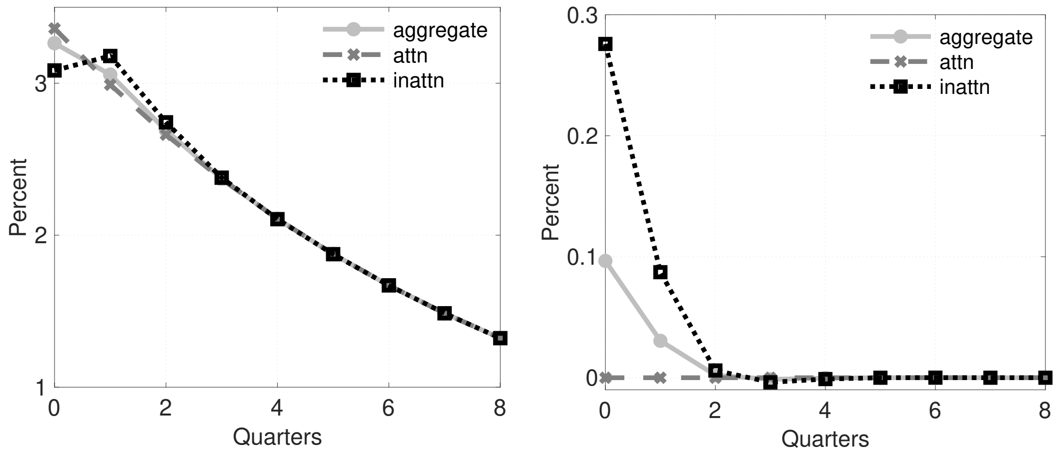

Notes: The figures report impulse responses in percent deviations from the perfect-information steady state of inflation and output for the aggregate economy, attentive firms, and inattentive firms.

tion remains quantitatively important for the transmission of monetary policy.

6.3. Attention and the efficacy of monetary policy

Figure 6 plots the aggregate responses to a one standard deviation expansionary shock to nominal aggregate demand growth. Panel (a) shows that inattentive firms under-adjust prices, reflecting partial incorporation of noisy signals about demand. Attentive firms track aggregate demand better than inattentive firms and exhibit more responsive prices.[27]

Panel (b) shows that inattentive firms are responsible for increased output following an expansionary shock. Firms set prices and commit to supply the quantity demanded. Since attentive firms raise prices by more in response to an expansionary shock, their output responds by less.

The grey solid lines represent aggregate inflation and output responses, driven by attention costs and the share of attentive firms.[28] Monetary non-neutrality increases with the

[27] Online Appendix Figure A.8 shows individual firms' impulse responses for prices, profits, attention, and stock returns (including full-information returns, imperfect-information returns, and total returns) in response to both expansionary and contractionary monetary shocks.

[28] To compare the aggregate responses with standard benchmarks, we convert the nominal aggregate demand shock to the nominal interest rate shock used in Christiano *et al.* (2005) by estimating the passthrough of the interest rate on the nominal aggregate demand in Appendix E.7. The right scale of Appendix Figure A.10 show that in response to a 25 basis point interest rate cut, output increases by 0.1% on impact, in line with the impact responses of 0.1% in Christiano *et al.* (2005) and smaller than the peak responses of 0.5%.

Table 8: Attention and monetary non-neutrality

	Least attentive	Baseline	Most attentive
Fraction of attentive firms (θ)	56%	65%	73%
Output response	0.12%	0.09%	0.07%

Notes: Dependence of output responses on the fraction of attentive firms in the economy. Output responses are calculated as percent deviations from the steady state in response to a 25 basis point rate cut. Calibration for the least and most attentive economy is described in the main text.

degree of inattention in the economy. Since we assume attentive firms face near-zero attention costs ($\omega_L \approx 0$), the impulse responses in Figure 6 provide a lower bound on the output responses to monetary shocks and an upper bound on the inflation responses.

A key implication of Panel (b) is that the aggregate output response to monetary policy increases with the share of inattentive firms. To illustrate the quantitative scope of the effect, we exogenously vary the fraction of attentive firms and compare output responses in our baseline calibration against two alternatives, $\check{\theta} = 56\%$ and $\hat{\theta} = 73\%$, which correspond to the minimum and maximum fraction of attentive firms over the sample period.

Table 8 reports the aggregate responses to monetary policy change as the fraction of attentive firms in the economy changes. The response of output growth to monetary policy is 5 basis points (or 42%) weaker in the most attentive calibration compared to the least attentive calibration. This suggests that expansionary policy in the depth of a recession when more firms are paying attention will be weaker than a preemptive interest rate (i.e., leaning against the wind) when aggregate attention is lower. This pattern is consistent with existing studies on the state dependency of monetary policy (e.g., Tenreyro and Thwaites, 2016). Similarly, monetary tightening imposes a smaller contractionary effect on output when more firms are attentive to monetary news, which highlights the importance of clear central-bank communication (highlighted, e.g., by Haldane *et al.*, 2021).

7. 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' deviations from full-information rational expectations (FIRE). To discipline models without FIRE, researchers require an understanding of firms' information sets and expectation-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 captures 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 in 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.

Data availability statement

The data and code underlying this research is available on Zenodo at https://dx.doi.org/10.5281/zenodo.11529484.

References

ABIS, S. and VELDKAMP, L. (2023). The changing economics of knowledge production. *The Review of Financial Studies*, p. hhad059.

ACHARYA, V. V., CROSIGNANI, M., EISERT, T. and EUFINGER, C. (2020). Zombie credit and (dis-)inflation: Evidence from europe.

AFROUZI, H. (2020). Strategic inattention, inflation dynamics, and the non-neutrality of money. *Manuscript*.

— and YANG, C. (2021a). Dynamic rational inattention and the phillips curve.

— and — (2021b). Selection in information acquisition and monetary non-neutrality.

Ahnert, T., Doerr, S., Pierri, N. and Timmer, Y. (2021). *Does IT Help? Information Technology in Banking and Entrepreneurship*. Tech. rep., International Monetary Fund.

Amirault, D., Rai, N. and Martin, L. (2020). *A Reference Guide for the Business Outlook Survey*. Tech. rep., Bank of Canada.

Auclert, A., Rognlie, M. and Straub, L. (2020). *Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model*. Tech. rep., National Bureau of Economic Research.

Autor, D., Dorn, D., Katz, L. F., Patterson, C. and Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics*, **135** (2), 645–709.

Bartoš, V., Bauer, M., Chytilová, J. and Matějka, F. (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review*, **106** (6), 1437–75.

Basu, S. and Fernald, J. G. (1997). Returns to scale in us production: Estimates and implications. *Journal of Political Economy*, **105** (2), 249–283.

Bloom, N., Bond, S. and Van Reenen, J. (2007). Uncertainty and investment dynamics. *Review of Economic Studies*, **74** (2), 391–415.

— and Van Reenen, J. (2010). Why do management practices differ across firms and countries? *Journal of Economic Perspectives*, **24** (1), 203–24.

BoardEX (2021). BoardEx [dataset]. Wharton Research Data Services, wrds-www.wharton.upenn.edu, accessed January 23, 2021.

BOS (2022). Business Outlook Survey [dataset]. Bank of Canada, accessed October 9, 2022.

Buehlmaier, M. and Whited, T. (2018). Are financial constraints priced? evidence from textual analysis. *Review of Financial Studies*, **31** (7), 2693–2728.

Candia, B., Coibion, O. and Gorodnichenko, Y. (2021). *The Inflation Expectations of US Firms: Evidence from a New Survey*. Tech. rep., National Bureau of Economic Research.

Cao, S., Jiang, W., Yang, B. and Zhang, A. L. (2020). *How to Talk When a Machine is Listening: Corporate Disclosure in the Age of AI*. Tech. rep., National Bureau of Economic Research.

Caplin, A., Dean, M. and Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *Review of Economic Studies*, **86** (3), 1061–1094.

— and Spulber, D. (1987). Menu costs and the neutrality of money. *Quarterly Journal of Economics*, **102** (4), 703–725.

CCM (2020). CRSP/Compustat merged database [dataset]. Wharton Research Data Services, wrds-www.wharton.upenn.edu, accessed October 6, 2020.

Chiang, Y.-T. (2021). Strategic uncertainty over business cycles. *Manuscript*.

Christiano, L., Eichenbaum, M. and Evans, C. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, **113** (1), 1–45.

Cloyne, J., Ferreira, C., Froemel, M. and Surico, P. (2023). Monetary policy, corporate finance, and investment. *Journal of the European Economic Association*, p. jvad009.

Coibion, O. and Gorodnichenko, Y. (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, **105** (8), 2644–78.

—, — and Kumar, S. (2018). How do firms form their expectations? new survey evidence. *American Economic Review*, **108** (9), 2671–2713.

—, — and Ropele, T. (2020). Inflation expectations and firm decisions: New causal evidence. *Quarterly Journal of Economics*, **135** (1), 165–219.

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

Cover, T. and Thomas, J. (2006). *Elements of Information Theory*. Wiley-Interscience.

Flynn, J. P. and Sastry, K. (2022). Attention cycles. *Manuscript*.

Fuster, A., Perez-Truglia, R., Wiederholt, M. and Zafar, B. (2018). Expectations with endogenous information acquisition: An experimental investigation. *Review of Economics and Statistics*, pp. 1–54.

Gabaix, X. (2019). Behavioral inattention. In *Handbook of Behavioral Economics: Applications and Foundations 1*, vol. 2, Elsevier, pp. 261–343.

Gertler, M. and Gilchrist, S. (1994). Monetary policy, business cycles, and the behavior of small manufacturing firms. *Quarterly Journal of Economics*, **109** (2), 309–340.

— and Leahy, J. (2008). A phillips curve with an *ss* foundation. *Journal of Political Economy*, **116** (3), 533–572.

Gorodnichenko, Y. and Weber, M. (2016). Are sticky prices costly? evidence from the stock market. *American Economic Review*, **106** (1), 165–99.

Gürkaynak, R. S., Sack, B. and Swanson, E. T. (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*.

Haldane, A., Macaulay, A. and McMahon, M. (2021). The three E's of central-bank communication with the public. In E. Pastén and R. Reis (eds.), *Independence, Credibility, and Communication of Central Banking*.

Handley, K. and Li, J. F. (2020). *Measuring the Effects of Firm Uncertainty on Economic Activity: New Evidence from One Million Documents*. Tech. rep., National Bureau of Economic Research.

Hansen, S., McMahon, M. and Prat, A. (2018). Transparency and deliberation within the FOMC: A computational linguistics approach. *Quarterly Journal of Economics*, **133** (2), 801–870.

Hassan, T. A., Hollander, S., Van Lent, L. and Tahoun, A. (2019). Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics*, **134** (4), 2135–2202.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, **95** (1), 161–182.

Konchitchki, Y. and Xie, J. (2022). Undisclosed material inflation risk. *Manuscript*.

Kwon, S. Y., Ma, Y. and Zimmermann, K. (2022). 100 years of rising corporate concentration. *Manuscript*.

Larsen, V. H., Thorsrud, L. A. and Zhulanova, J. (2021). News-driven inflation expectations and information rigidities. *Journal of Monetary Economics*, **117**, 507–520.

Leahy, J. V. and Whited, T. M. (1996). The effect of uncertainty on investment: Some stylized facts. *Journal of Money, Credit and Banking*, **28** (1), 64–83.

Loughran, T. and McDonald, B. (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-ks. *Journal of Finance*, **66** (1), 35–65.

Luo, Y. (2008). Consumption dynamics under information processing constraints. *Review of Economic Dynamics*, **11** (2), 366–385.

Macaulay, A. (2020). Cyclical attention to saving. In *UniCredit Foundation Working Papers Series 140*.

Maćkowiak, B., Matějka, F. and Wiederholt, M. (Forthcoming). Forthcoming, rational inattention: A review. *Journal of Economic Literature*.

—, Moench, E. and Wiederholt, M. (2009). Sectoral price data and models of price setting.

Journal of Monetary Economics, **56**, S78–S99.

— and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, **99** (3), 769–803.

— and Wiederholt, M. (2015). Business cycle dynamics under rational inattention. *Review of Economic Studies*, **82** (4), 1502–1532.

— and Wiederholt, M. (2023). Rational inattention and the business cycle effects of productivity and news shocks.

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *Quarterly Journal of Economics*, **131** (1), 53–87.

Mankiw, G. and Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *Quarterly Journal of Economics*, **117** (4), 1295–1328.

Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, **105** (1), 272–98.

McKay, A. and Wieland, J. F. (2021). Lumpy durable consumption demand and the limited ammunition of monetary policy. *Econometrica*, **89** (6), 2717–2749.

Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, **15** (2), 145–161.

Nakamura, E. and Steinsson, J. (2018). High frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics*.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, **64** (6), 1263–1297.

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

PPI (2022). Producer Price Index [dataset]. U.S. Bureau of Labor Statistics, www.bls.gov/ppi, accessed April 5, 2022.

Ramey, V. (2016). Macroeconomic shocks and their propagation. In *Handbook of Macroeconomics*, vol. 2, Elsevier, pp. 71–162.

Reis, R. (2006). Inattentive producers. *Review of Economic Studies*, **73** (3), 793–821.

Reutskaja, E., Nagel, R., Camerer, C. and Rangel, A. (2011). Search dynamics in consumer choice under time pressure: An eye-tracking study. *American Economic Review*, **101** (2), 900–926.

Sims, C. (2003). Implications of rational inattention. *Journal of Monetary Economics*, **50** (3), 665–690.

SPF (2023). Survey of Professional Forecasters [dataset]. Federal Reserve Bank of Philadelphia, www.philadelphiafed.org, accessed November 5, 2023.

Tanaka, M., Bloom, N., David, J. and Koga, M. (2019). Firm performance and macro forecast accuracy. *Journal of Monetary Economics*.

Tenreyro, S. and Thwaites, G. (2016). Pushing on a string: Us monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, **8** (4), 43–74.

Van Nieuwerburgh, S. and Veldkamp, L. (2006). Learning asymmetries in real business cycles. *Journal of Monetary Economics*, **53** (4), 753–772.

— and Veldkamp, L. (2009). Information immobility and the home bias puzzle. *Journal of Finance*, **64** (3), 1187–1215.

Vavra, J. (2014). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *Quarterly Journal of Economics*, **129** (1), 215–258.

Weitzner, G. and Howes, C. (2021). Bank information production over the business cycle. *Manuscript*.

Woodford, M. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics*, **56**, S100–S124.

Wooldridge, J. M. (2009). On estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters*, **104** (3), 112–114.

Yang, C. (2022). Rational inattention, menu costs, and multi-product firms: Micro evidence and aggregate implications. *Journal of Monetary Economics*, **128**, 105–123.

Zbaracki, M., Ritson, M., Levy, D., Dutta, S. and Bergen, M. (2004). Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *Review of Economics and Statistics*, **86** (2), 514–533.