Should Retail Investors Listen to Social Media Analysts? Evidence from Text-Implied Beliefs*

Chukwuma Dim†

This version: September 2021

Latest version

Abstract

This paper uses machine learning to infer nonprofessional social media investment analysts' (SMAs) beliefs from opinions expressed about individual stocks. On average, SMA beliefs predict future abnormal returns and earnings surprise. However, there exists substantial heterogeneity in SMAs' ability to form beliefs that yield investment value. Some 13% high-skilled SMAs form beliefs that yield a sizeable one-week three-factor alpha of 61 bps, while the rest of the SMAs generate only 6 bps. Firm and industry specializations are the most distinctive characteristics of high-skill SMAs. When forming beliefs, SMAs extrapolate from past returns and herd on the consensus. Herding is, however, less pronounced in bad times and when the consensus is optimistic, but it is more pronounced when the consensus is correct ex-post. SMAs' behavioural bias does not result in systematically wrong beliefs.

Keywords: Nonprofessional analysts, Belief formation, Investor skill, Market efficiency, Herding, Extrapolation, Machine learning, Natural language processing

JEL: G11, G12, G14

*I am grateful to Grigory Vilkov and Francesco Sangiorgi for their support and suggestions. I also thank Zhe An (discussant), Fahiz Baba-Yara, Fousseni Chabi-Yo, Francesco Franzoni, Michalis Haliassos, Jonas Happel, Burton Hollifield, Emirhan Ilhan, Yigitcan Karabulut, Camelia Kuhnen, Matthias Lassak, Mario Milone (discussant), Altan Pazarbasi, Kim Peinjenburg, Zacharias Sautner, Nic Schaub, Onur Sefiloglu (discussant), Kelly Shue (discussant), Fabio Trojani (discussant), Raman Uppal, Maximilian Voigt, and Wei Wu (discussant) for helpful comments and suggestions. I thank participants of the Frankfurt School of Finance & Management Brownbag Seminar, Financial Intermediation Research Society Conference 2021, Northern Finance Association Conference 2021, NOVA Finance PhD Pitch Perfect, 11th Financial Markets and Corporate Governance Conference, 2nd LTI/Bank of Italy Workshop, Frankfurt Reading Group on Household Finance, 6th European Retail Investment Conference Doctoral Consortium, 3rd QMUL Economics and Finance Workshop, 37th International Conference of the French Finance Association, and 4th Annual Dauphine Finance PhD Workshop for their helpful comments.

†Frankfurt School of Finance & Management. ✉ c.dim@fs.de

1 Introduction

Social networks shape individuals' expectations and actions, as people rely on their networks for information. In financial markets, social media intensifies the role of social networks in belief formation by making it easier for individuals to interact and share investment ideas. For instance, retail investors' frenetic trading in the GameStop stock in January 2021 was primarily coordinated via social media, arguably leading to the 2,000% surge in the stock's price and the near-collapse of Melvin Capital, a hedge fund. Similarly, retail investors spurred via social media accounted for 23% of all US equity trading in January and February 2021, rendering their market footprint as large as that of all hedge funds and mutual funds combined.[1] These developments suggest that the interplay between social media and retail trading poses new challenges for financial markets, which require a deeper understanding of belief formation about asset prices on social media.

A crucial feature of investment-focused social media is the existence of influencers, i.e., non-professional *Social Media Investment Analysts* (SMAs), who publish stock-specific investment opinions and belief statements that shape the views and actions of many individual investors.[2] I examine these influencers' role in financial markets using two related strands of the theoretical literature as motivation. The model of Pedersen (2021) suggests that the overall rationality of an economy depends on the distribution of rational influencers in social networks, which implies that the distribution of SMAs with certain traits on social media, such as manner of belief formation, has relevant implications. Similarly, Bhamra et al.'s (2021) model of psychological distance implies that by influencing individual investors' psychological distance from firms, SMAs' views on social media shape individuals' investment decisions, with welfare consequences.[3] I, there-

[1] See, `https://www.ft.com/content/7a91e3ea-b9ec-4611-9a03-a8dd3b8bddb5`.

[2] SMAs are different from professional security analysts in several ways: Unlike professionals, most SMAs are individual investors who share their views on the stocks in which they invest. Hence, SMAs face different incentives than professionals. For example, SMAs are less likely to have a conflict of interest arising from business relationships with firms. They are less likely to have conflicts related to job security and the need to ingratiate themselves with managers and powerful investor groups. However, SMAs care about their reputation and popularity on social media, because it impacts their compensation. Farrell et al. (2020) show that, unlike professionals who cater to institutional investors, SMAs cater to retail investors, and their views inform retail trade.

[3] Psychological distance is a cognitive separation between oneself and other entities, such as persons, events, or times. Its dimensions include temporal, spatial, informational, and social distance (Baltatescu, 2014). SMAs' views on social media could bridge the social and informational distance between individuals and specific stocks by making the stocks more familiar, impacting beliefs and investment decisions.

fore, ask the following questions: Should investors listen to SMAs? How widespread is SMAs' ability to form correct beliefs? How do SMAs form the beliefs they impact on others?

I address these questions by examining beliefs about individual stock returns expressed by SMAs on social media, with two main contributions to the literature. First, I provide the first evidence on the cross-sectional distribution of social media influencers' ability to form correct beliefs about stocks. Existing studies on the informativeness of social media primarily focus on whether average opinions predict returns. However, individuals on social media operate in echo chambers instead of heeding the average opinion (e.g., Cookson et al., 2021), rendering the cross-sectional distribution of SMAs' ability crucial. Second, I document the roles of herding and extrapolation in social media influencers' beliefs about *individual* stocks. Existing research on belief formation about asset prices primarily focuses on expectations about the aggregate market, despite widespread evidence that individuals largely do not hold the market portfolio.[4] In contrast, I document how herding and extrapolation shape people's beliefs about the stocks in which they actually invest.

The major empirical challenge in studying beliefs about individual stocks is the difficulty in obtaining a reasonably large time-series and cross-sectional sample of belief statements. I overcome this challenge by applying natural language processing (NLP) and supervised machine learning (ML) to infer SMA beliefs about a large cross-section of stocks over a relatively long time period from views expressed by SMAs in a popular investment-focused social media platform, Seeking Alpha. More specifically, since 2018, most SMA opinion articles on Seeking Alpha are explicitly tagged with the author's belief about a stock using one of the following descriptions: "Very Bullish", "Bullish", "Neutral", "Bearish", or "Very Bearish". I use the subset of articles with explicit belief statements to train a relatively simple yet robust Support Vector Classifier ML model, which enables me to extract the beliefs implied by the rest of the unlabelled SMA articles on Seeking Alpha dating back to 2004.[5] I then use this large sample of stated and

[4] For evidence on under-diversification, see Ivković and Weisbenner (2005); Barber and Odean (2013).

[5] It is natural to apply NLP and ML in the context of beliefs, because people's beliefs can often be understood from their linguistic expressions. For instance, Tversky and Kahneman (1974) note that beliefs are usually expressed in statements such as "I think that ...", "Chances are...", "It is unlikely that ...", and so forth. Moreover, in the broader social science literature, Iyyer et al. (2014) show that people's political ideology can be reliably deduced from their linguistic expressions using supervised ML. Similarly, Gentzkow and Shapiro (2010) demonstrate that ML can be used to deduce the political slant of media outlets.

extracted beliefs for this paper's analyses. Simply put, the ML algorithm maps words in an investment thesis to the SMA's implicit belief – an approach reminiscent of the recent call for the application of ML in constructing proxies of beliefs from textual data to improve upon the limitations of survey data (see, Brunnermeier et al., 2021).

Several factors render Seeking Alpha an ideal setting to study belief formation on social media. First, Seeking Alpha has a long history, dating back to 2004. It is popular among retail investors who subscribe to the platform to gain access to SMAs' views or contribute their opinions. For instance, about 20 million people use Seeking Alpha monthly, and about 11,000 SMAs contributed views on roughly 7,200 firms between 2004 and 2019.[6] Second, the goal of Seeking Alpha is to provide opinions and analyses rather than news, primarily through individual investors who describe their approach to stock-picking (Seeking Alpha, 2006). Third, views expressed on Seeking Alpha are backed by an in-depth analysis of an investment thesis, which Seeking Alpha's editorial team checks for quality before publication. The second and third features imply that SMAs in this paper primarily refer to individual investors – relatively more sophisticated than the average retail investor – who share their beliefs about stocks on social media, potentially shaping the expectations and actions of other investors in their network.[7]

I begin the empirical analysis by examining whether, on average, SMA beliefs provide value-relevant information about individual stocks beyond that produced by professional analysts and the news media. I find that stock-level aggregated SMA beliefs predict both abnormal returns and earnings surprise after controlling for professional analyst recommendations, news sentiment, and a host of other variables. A standard deviation increase in belief bullishness predicts a 17 (26) basis point increase in the following one-week (three-month) three-factor abnormal returns and an increase in earnings surprise by 15% of its standard deviation.

The fact that SMAs' average beliefs predict earnings surprise and abnormal returns up to a three-month horizon suggests that the growing popularity of nonprofessional analysts on

[6]See `https://seekingalpha.com/page/about_us` for usage statistics. The number of SMAs and stocks include only single-ticker publications on US common stocks listed on the NYSE, NASDAQ, or AMEX stock exchange.

[7]SMAs' testimonials indicate that the main incentives for sharing their views on Seeking Alpha include feedback on investment thesis from other investors, increased visibility and recognition, which can result in professional opportunities, and direct monetary compensation from Seeking Alpha. Financial compensation is based on article-page views and can be as low as $40, suggesting that it is not the primary motivation of SMAs. For more details on the compensation scheme, see `https://seekingalpha.com/page/article_payments`.

social media is potentially beneficial to both investors and price discovery. However, a complete understanding of the relevance of SMAs to financial markets requires a deeper analysis of how much skill exists among SMAs whose views tend to inform retail trades. A desirable scenario is that the predictive ability of average SMA beliefs is primarily attributable to the fact that most SMAs are *truly* skilled in consistently producing and processing information that informs correct beliefs. However, it might instead be that most SMAs are merely lucky, with only a select few genuinely skilled.

To rigorously examine the distribution of SMAs' ability to form correct beliefs, I follow recent literature on financial professionals' skill (Crane and Crotty, 2020; Chen et al., 2017) and model SMAs' ability as a mixture distribution of multiple skill groups. Because a mixture model uses information in the cross-section of investor performance to reduce noise, it ameliorates the false discovery problem that often arises from the low signal-to-noise feature of abnormal returns. Guided by model selection criteria, I estimate a two-component mixture model, where one component comprises the low-skill SMAs and the other the high-skill type. About 87% of SMAs belong to the low-skill group, whose beliefs generate a moderate one-week three-factor alpha of six basis points (bps). Conversely, the remaining 13% high-skill SMAs form beliefs that generate a much larger abnormal return of 61 bps over the same horizon (31% annualized). Pulling the low-skill and high-skill SMAs, I estimate that roughly 56% of SMAs generate positive abnormal returns. To provide some context, Crane and Crotty (2020) use a similar setup and estimate the fraction of high-skill professional analysts to be 36%, with roughly 97% able to generate positive abnormal returns. Therefore, although SMAs as a group tend to add value, individual SMAs' skills are considerably limited, with substantial dispersion.

The dispersion in SMAs' ability to form correct beliefs suggests that investors can benefit from signals that reduce their search costs for valuable information on social media. Hence, I examine which SMA characteristics are associated with skill. Consistent with theoretical results on gains from specialization in information acquisition and cognitive capacity constraints (e.g., Van Nieuwerburgh and Veldkamp, 2010; Hirshleifer et al., 2011), I find that industry and firm specializations are the most distinctive characteristics of high-skill SMAs. In fact, SMAs who specialize in a few industries (firms) are 34 (31) percentage points (pp) more likely to be high-skill SMAs. Furthermore, SMAs who mostly have an investment position in the stock they have expressed a belief in are 12 pp more likely be high-skill. More popular SMAs also have a higher probability of being high-skill, suggesting that individuals follow skilled SMAs on Seeking Alpha.

Guided by the theoretical literature, I investigate channels through which SMAs form beliefs about stock returns, namely herding (e.g., Scharfstein and Stein, 1990; Bikhchandani et al., 1992; Banerjee, 1992) and extrapolation (e.g., Barberis et al., 2015). I find evidence that SMAs herd in stating their beliefs. Herding is, however, less pronounced during recessions and high market uncertainty, consistent with models on belief polarization in bad times (e.g., Cujean and Hasler, 2017). Furthermore, unlike professional analysts (see, Welch, 2000), SMAs herd less when the consensus is optimistic. A standard deviation increase in consensus optimism is associated with a reduction in herding by 32% of its unconditional value, indicating that SMAs prefer attention-grabbing deviation from the consensus, potentially to attract readership. Interestingly, herding is more predominant among SMAs when the consensus correctly predicts returns. This suggests that SMAs learn fundamental information from the consensus to improve their expectations, consistent with information-based herding (e.g., Banerjee, 1992).[8]

SMAs also extrapolate from past returns, with extrapolation stronger for stocks with more salient returns. For example, when a stock's price doubles over the past one-week period, the stock-level aggregate belief becomes more bullish by 40% of its standard deviation, with the influence of older weekly returns declining with time. This evidence, combined with earlier results on the informativeness of SMA beliefs, indicates that return extrapolation does not necessarily entail systematically biased beliefs. These results complement the work of Da et al. (2021), who document systematically biased beliefs for return extrapolators.

Taken together, my results suggest that, as a group, SMAs are sufficiently skilled to form beliefs that potentially improve price discovery and benefit investors, even though there is substantial dispersion in their skill distribution. In light of ongoing concerns regarding social media's growing influence over financial markets, it appears that some aspects of social media serve as veritable sources of information, which informs beliefs that can improve investment decisions.

[8]It is equally possible that SMAs independently follow the same fundamental information. Whatever the case, the result suggests that observational learning among SMAs would to some extent improve price discovery, as investors trade on SMAs' belief statements.

Related Literature: This paper contributes to several strands of research. First, the paper is related to the literature that explores whether investors' social media opinions about individual stocks are informative. Chen et al. (2014) show that the fraction of negative words in views expressed on Seeking Alpha predicts abnormal stock returns and earnings surprise. Avery et al. (2016) echo these results using individual investors' predictions on MotleyFool.[9] In contrast to these studies, which focus on the informativeness of *average* social media opinions, I provide the first evidence on the cross-sectional *distribution* of nonprofessional analysts' ability to form correct beliefs about stock returns, which enhances our understanding of social media's role in financial markets, especially in light of Cookson et al.'s (2021) evidence that individuals operate in echo chambers on investment-focused social media platforms. A related literature examines how nonprofessional analysts' views influence retail investors' trade and information environment, and professional analysts' informational role (Gomez et al., 2020; Farrell et al., 2020; Campbell et al., 2019; Drake et al., 2019). I contribute to this literature by examining whether retail investors should heed nonprofessional analysts in the first place. I demonstrate that SMAs as a group are valuable to investors, although substantial dispersion exists in their ability. I further document specific characteristics that can help investors identify skilled SMAs.

This paper is also related to the literature on individual investors' beliefs (e.g., Choi and Robertson, 2020; Giglio et al., 2019) and expectation formation (e.g., Kuchler and Zafar, 2019; Greenwood and Shleifer, 2014; Malmendier and Nagel, 2011). This literature largely relies on survey data of beliefs about aggregate outcomes. The very few exceptions include Bhamra et al. (2021), who extracts belief distortions about individual stocks from Finnish households' stock holdings, Cookson and Niessner (2020), Cookson et al. (2020), and Da et al. (2021), who study beliefs about a limited number of stocks using data posted on social media. My contribution to this literature lies in using NLP and ML techniques to infer beliefs about a much larger cross-section of stocks over a relatively long time period, enabling me to provide new insights on the role of herding and extrapolation in belief formation about individual stocks on social media.

[9] In a related paper, Antweiler and Frank (2004) find that messages posted on stock message boards on the internet predict the return and volatility of a narrow set of stocks. Das and Chen (2007) find similar results for volatility and volume.

This paper also speaks to the literature that studies whether retail investors' participation in financial markets improves market efficiency or introduces noise (e.g., Boehmer et al., 2020; Seasholes and Zhu, 2010; Barber and Odean, 2000; Kelley and Tetlock, 2013; Kaniel et al., 2012). In this paper, SMAs are also individual investors, as they disclose their investments on Seeking Alpha. In demonstrating that SMA beliefs contain value-relevant information, this paper aligns with the extant literature which argues that retail investors produce information that potentially improves price discovery. Unlike studies in this stream of literature that primarily examine retail trading, I study beliefs and the distribution of skill among individual investors, both of which offer new insights. It is noteworthy that investors' trades can differ from their beliefs, because trade can result from other reasons, such as liquidity needs, unrelated to beliefs.

Finally, I contribute to the growing body of literature that uses ML techniques to extract economic quantities from textual data (e.g., Gu et al., 2020; Chen et al., 2019; Ke et al., 2019; Manela and Moreira, 2017). By leveraging NLP and supervised ML techniques to infer beliefs, I deepen the range of economic research questions these techniques can address.

The paper proceeds as follows. Section 2 describes the data. Section 3 covers belief extraction using ML. Section 4 provides the results on the informativeness of stock-level aggregate SMA beliefs. Section 5 analyzes the cross-sectional distribution of SMAs' ability to form correct beliefs. Section 6 examines the role of herding and return extrapolation in SMAs' belief formation. Section 7 discusses robustness tests, and Section 8 concludes the paper.

2 Data

This section describes the data used in this study. The sample period, spanning from January 2005 to December 2019, was determined by the availability of SMAs' beliefs and opinions on Seeking Alpha.

2.1 Seeking Alpha Data

This paper's analysis relies on opinions and belief statements of SMAs on Seeking Alpha, a popular investment-focused social media platform launched in 2004.[10] Any registered user can contribute views on Seeking Alpha by submitting an opinion article with an extended analysis of an investment thesis (and, in more recent times, accompanied by an explicit belief statement) on specific stocks. However, the opinion article must pass through Seeking Alpha's editorial team, who checks for quality standards without interfering with the author's viewpoint. Most users of Seeking Alpha consume and comment on the views of a smaller subset of individuals who contribute opinion articles and belief statements. Hence, SMAs in this paper refer to this subset of individuals who contribute opinions that likely shape their followers' beliefs and actions.

To obtain SMAs' opinions and stated beliefs from Seeking Alpha, I develop a web-scraping algorithm to download all opinion articles (and associated belief statements where available) covering a single US common stock listed on either the NYSE, NASDAQ or AMEX stock exchange. I obtain the SMA's ID and disclosure of investment position in the stock, the stock ticker, and the publication date for each publication.[11] Furthermore, I retrieve all comments posted in response to the publication by other Seeking Alpha users.

In total, I downloaded 280,514 opinion articles and 7.3 million comments, contributed by roughly 11,000 SMAs and 300,000 users, covering about 7,200 stocks over the period spanning from January 2005 – December 2019. The SMA belief statement that accompanies each publication is tagged as one of the following: "Very Bullish", "Bullish", "Neutral", "Bearish", or "Very Bearish".[12] However, most SMA opinion articles published before 2018 did not explicitly state SMAs' beliefs: only 43% of publications prior to 2018 include an explicit belief statement com-

[10]See, https://seekingalpha.com/page/about_us. Very few opinions were contributed in 2004 after Seeking Alpha's launch. As a result, the analyses in this paper uses data beginning in January 2005.

[11]Most publications include a disclosure section, where the author discloses whether he/she has an investment position in the stock being written about. See Figures A1 and A2 for examples of these disclosures. I manually label a randomly selected 5,000 disclosures as either "Long position", "Short position", or "No position" and then employ this labelled sample to train a Support Vector Classifier ML model, as described in Section 3, which is used to extract the investment position stated in all other disclosures. Given the simplicity of this particular learning exercise, the trained model achieved an out-of-sample accuracy rate of 99%.

[12]One may wonder whether the stated beliefs are the true SMAs' beliefs or seek to mislead other investors and favourably drive prices. While deceit is a possibility, it is more likely that SMAs truthfully report their beliefs, because it is a sustainable way of building a strong reputation and followers on social media, which are important motivations for SMAs. For an example of an article on the reputational motive of SMAs, see the Wall Street Journal article "Retail Traders Wield Social media for Investing Fame".

pared to roughly 75% of publications after 2018. Therefore, to obtain a large cross-sectional and time-series sample of belief statements, I train a machine learning model, described in Section 3, to extract SMAs' underlying beliefs from the articles without explicit belief statements.

Figures A1 and A2 of the Appendix show sample SMA opinion articles where the authors explicitly state their beliefs about a stock as "Bullish" and "Bearish", respectively. These samples' narratives indicate that individuals who contribute beliefs and opinions to Seeking Alpha are generally more financially literate and sophisticated than the average retail investor.

2.2 Other data

I obtain stock returns, price, and market capitalization data from CRSP and firm fundamentals data from Compustat. I compute abnormal return $ABR_{k,t}(h)$ for firm k on day t for investment horizon h relative to three different benchmarks: the Capital Asset Pricing Model (CAPM), Fama and French (1993) three-factor model (FF3), and the Daniel et al. (1997) size/book-to-market/momentum characteristics-based benchmark (SBM). For the CAPM and FF3 benchmarks, I estimate betas for each stock using daily data over the trading-day window t-272 to t-21, where t is the belief publication day. Merging the CRSP/Compustat data with the SMA belief data reduces the number of observations to 236,250.

Data on professional stock analyst recommendations and forecasts of quarterly earnings per share are taken from the Institutional Brokers' Estimate System (IBES). I use analyst recommendations data to compute the number of recommendation upgrades and downgrades for firm k on day t. From the unadjusted detail history of analysts' earnings forecasts, I compute earnings surprise. Finally, I measure sentiment across a comprehensive set of cash flow relevant news events about a stock on a given day using the Event Sentiment Score (ESS) from RavenPack News Analytics. Appendix A.1 provides details on the construction of all variables.

3 Measuring SMA Beliefs

Recent applications of machine learning (ML) and natural language processing (NLP) techniques in finance and economics (e.g., Chen et al., 2019; Manela and Moreira, 2017) provide evidence that textual data can be used to generate important economic quantities. Inspired by these results, I use the subset of SMAs' opinion articles that includes an explicit belief statement (labelled samples) to train a supervised ML model to extract the underlying beliefs from all other articles that do not explicitly state the author's belief. The result is a large sample of beliefs that allows for a comprehensive study of SMAs' beliefs about individual stocks returns.

I begin by preprocessing the article text to reduce the vocabulary to terms potentially relevant for belief classification.[13] Next, the labelled data is randomly split into a training set and a test set, maintaining the proportion of the belief classes in the original data in both sets. The test set, comprising 30% of the labelled data, is used *only* for final model evaluation, *not for model training*. I vectorize the training text corpus, omitting terms in fewer than one percent of the text corpus, yielding an $N \times M$ document-term-matrix $\mathbf{X}$, with the element in row n and column m capturing the number of times unigram/bigram m appeared in article n.[14,15] Finally, I normalize $\mathbf{X}$ using the term frequency-inverse document frequency (tf-idf) algorithm so that its entries reflect the importance of each unigram/bigram (n-gram) to a particular article.

I use the normalized matrix $\mathbf{X}$ as the features in the machine learning classification task aimed at finding the optimal weights for combining the normalized n-gram frequencies to yield the best out-of-sample (OOS) classification of SMA beliefs. I adopt the linear Support Vector Classifier (SVC) machine learning algorithm for this exercise because it performs well in very high-dimensional feature spaces (e.g., Chen et al., 2019; Manela and Moreira, 2017; Frankel et al., 2016). Moreover, linear SVC is quickly trained on high-dimensional data, as only one hyperparameter c needs to be tuned. The output easily reflects the word combinations that matter for belief classification. Although penalized logistic regression has similar features as linear SVC, SVC produced better out-of-sample performance. Appendix A.2 provides a description of the SVC algorithm.

[13]The steps involve converting all words to lowercase and removing stopwords, words containing digits and punctuation. Finally, I lemmatize the texts so that words representing the same underlying concept are captured by the same word token while avoiding ambiguity.

[14]$\mathbf{X}$ is a high-dimensional sparse matrix with roughly 880,500 columns (features) and 101,400 rows (articles). The sparsity is because the feature space comprises all the unique unigrams/bigrams (n-grams) in the training text corpus, while only a small fraction of these n-grams appears in each article.

[15]Unigrams are single words, and bigrams are consecutive combinations of two words – they are both referred to as n-grams in this paper. I experimented with using only unigrams, or unigrams, bigrams and trigrams (consecutive combinations of three words) together. While unigrams produce slightly worse out-of-sample performance than unigrams and bigrams combined, adding trigrams to the latter yields similar performance.

To train the linear SVC, I first collapse the belief labels to three classes, setting the "Very Bullish" and "Bullish" labels to "Bullish", and setting the "Very Bearish" and "Bearish" labels to "Bearish"; the third label is "Neutral". This reduces the problem of imbalanced data, since together, the "Very Bullish" and "Very Bearish" beliefs account for only 3% of the labelled data. Nevertheless, the data remains highly imbalanced, with Bullish beliefs accounting for roughly 82% of the observations. Next, I search for the n-gram weights and the SVC's regularization hyperparameter c that produce the best OOS belief classification performance. I do so using grid search over a range of c values, five-fold cross-validation, and Area Under the Curve (AUC) of the Receiver Operating Characteristic (ROC) Curve as the model evaluation metric.[16,17] I use the AUC rather than the accuracy metric because accuracy can be uninformative on imbalanced data. For instance, a random classifier or a classifier that simply predicts the majority class can achieve high accuracy on imbalanced data. In contrast, the AUC metric is robust to class imbalance: Random predictions always produce an AUC of 0.5, regardless of data imbalance.

Panel A of Figure 1 shows the average AUC score achieved by the SVC model on the cross-validation test folds during training, indicating that the optimal hyperparameter is $c = 0.69$. Panel B shows the performance of the trained SVC model, based on hyperparameter $c = 0.69$, on the held-out test data (not used for training). The ROC curves for the different belief classes are all close to the top-left corner, indicating that the SVC model achieves a high true positive rate and low false positive rate for all belief classes. The area under these respective curves (AUC) is 0.94 for the Bullish and Neutral belief classes and 0.96 for the Bearish belief class. The macro- and micro-average AUC scores, which summarize the AUC across the different belief classes, are also comparable at 0.95 and 0.97, respectively. Moreover, the model achieves an accuracy score of 90% on the test data. These evaluation results suggest that the trained model reliably classifies the Bullish, Bearish, and Neutral belief classes.

[16] The AUC score ranges between 0 and 1, with 1 indicating a model that perfectly separates different classes.

[17] The grid search and five-fold cross-validation proceeds by randomly splitting the *training set* data into five equal groups. Then, for a given value of c, the SVC is trained on folds one to four, while the held-out fold five (test fold) is used to compute an OOS AUC score. The process is repeated until each of the five folds is used once as the test fold for the given c. The resulting five OOS AUC scores are averaged and saved, and the iteration moves to the next value for c. In the end, I select the value of c, which provides the highest average OOS AUC score, and use it to refit the model on the full training set data.

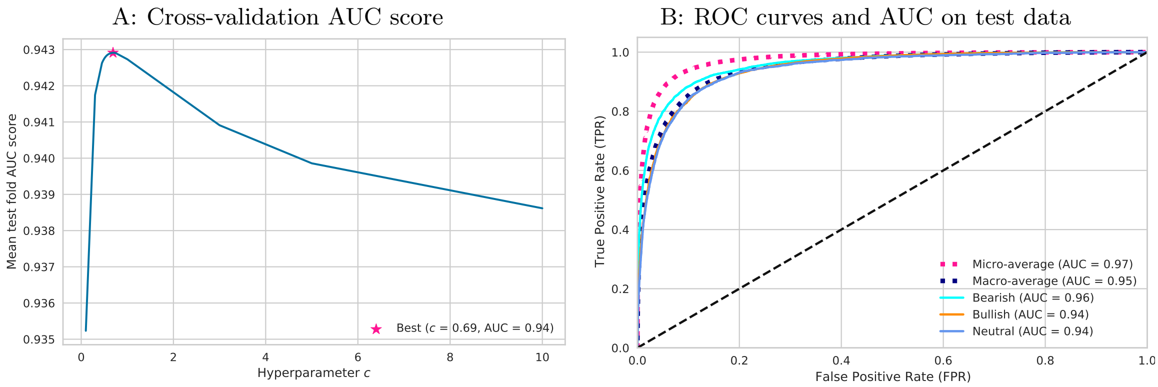

Figure 1: Model Evaluation using the AUC Metric. In Panel A, the figure shows the average AUC score achieved by the linear SVC model on the cross-validation test folds for different hyperparameter c values during training. Panel B shows the ROC curves of the final SVC model, based on the optimal $c = 0.69$, for the different belief classes on the *test data* not used for training.

Figure 2 provides additional validation and sanity check, showing the weights the trained model assigns to the top 100 n-grams for the Bullish and Bearish beliefs. Indeed, the model produces intuitive results, indicating that articles containing terms such as "overvalue", "neutral", "short", "avoid", "take profit", and "short opportunity" are less likely to be Bullish, while those containing terms such as "undervalue", "upside", "buy", "opportunity", and "bullish" are more likely to be Bullish. Similarly, articles containing terms such as "upside", "undervalue", "conservative", "opportunity", and "buy" are less likely to be Bearish, while those containing terms such as "overvalue", "short", "bearish", "short candidate", and "expensive" are more likely to be Bearish. Overall, the model's impressive performance on the test data and its intuitive n-gram weights provide strong evidence that the model can be used to reliably infer SMAs' implicit beliefs from article words with a high degree of accuracy.

Table 1, Panel A, shows summary statistics for the stated, extracted, and all SMA beliefs. The distribution of the belief classes in the subsample of extracted beliefs is comparable to the distribution in the subsample of stated beliefs, further revealing that the trained model produces reasonable results. Overall, bullish beliefs account for 81% of the stated and extracted beliefs (Column "All"), indicating that SMA beliefs are generally bullish. The rest of the paper, unless otherwise stated, uses the stated and extracted beliefs, which provide a larger sample size that is particularly helpful for the cross-sectional analysis of SMAs' beliefs. Panel B of Table 1 shows summary statistics for stock-level aggregated SMA beliefs (AB), firms' market capitalization,

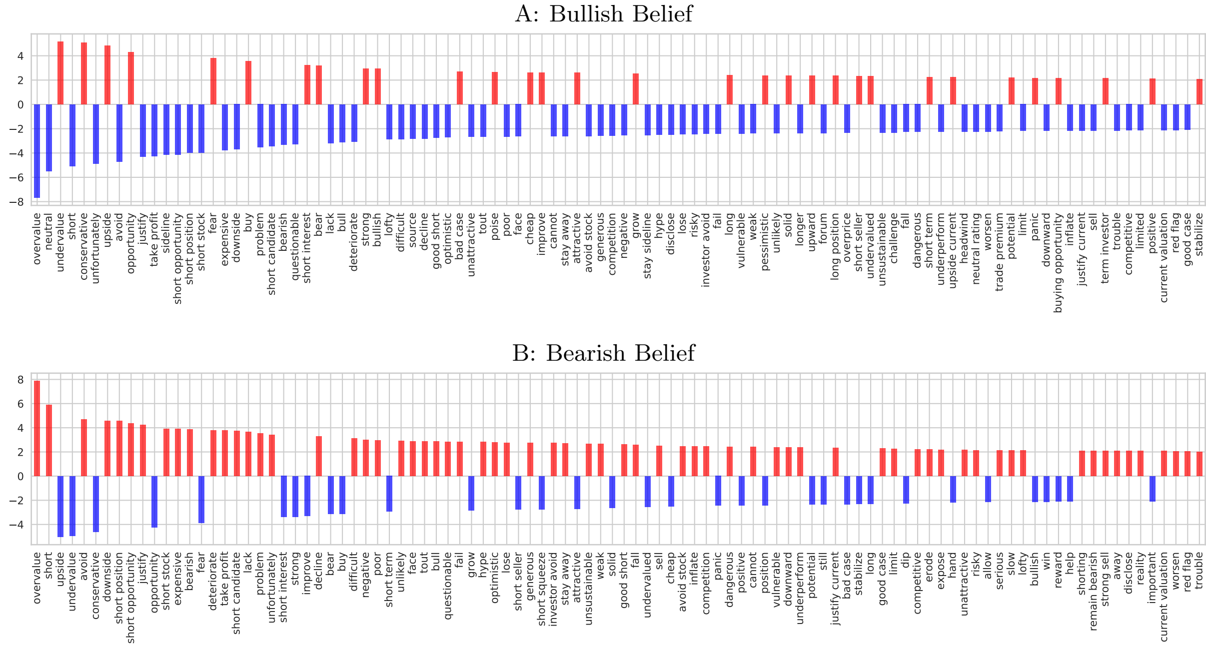

Figure 2: Most Relevant Terms for Extracting Beliefs. The figure shows the weights assigned by the trained SVC model to the top 100 n-grams for classifying Bullish and Bearish beliefs.

and abnormal returns (ABR) for the five and 21 trading-day horizons following belief statements. Average abnormal returns are positive, while average market capitalization is \$53.9 billion.

4 Are SMA Beliefs Informative on Average?

I examine the informativeness of average SMA beliefs by determining whether stock-level aggregate beliefs predict abnormal stock returns in subsection 4.1 and earnings surprise in subsection 4.2.

4.1 Evidence from Individual Stock Returns

To examine whether average SMA beliefs are informative about stock returns, I first aggregate beliefs about stock k on day t by subtracting the number of Bearish beliefs ($NBearish_{k,t}$) from the number of Bullish beliefs ($NBullish_{k,t}$) then normalize by the total number of beliefs ($NBelief_{k,t}$):

$$AB_{k,t} = \frac{NBullish_{k,t} - NBearish_{k,t}}{NBelief_{k,t}} \qquad \in \quad [-1, 1]. \tag{1}$$

Panel A: Proportion of Beliefs			
	Stated Belief	Extracted Belief	All
Bullish	0.82	0.79	0.81
Bearish	0.13	0.18	0.16
Neutral	0.05	0.03	0.04
Obs.	117,271	118,979	236,250

Panel B: Stock-level Variables						
			CAPM		FF3	
	AB	Mkt. Cap.	ABR(5)	ABR(21)	ABR(5)	ABR(21)
Mean	0.6506	53,938	0.0006	0.0013	0.0006	0.0014
SD	0.6571	95,058	0.0641	0.1201	0.0631	0.1184
10%	-0.3389	2,741	-0.0534	-0.1073	-0.0520	-0.1053
25%	0.6730	4,841	-0.0245	-0.0510	-0.0235	-0.0492
50%	0.9371	15,514	-0.0005	-0.0016	-0.0003	-0.0012
75%	0.9787	60,167	0.0235	0.0474	0.0228	0.0462
90%	0.9830	148,894	0.0540	0.1062	0.0526	0.1038

Table 1: Summary Statistics of SMA Beliefs, Abnormal Returns and Size. In Panel A, the table reports the proportion of each belief class in the subsample of SMA articles with an explicit belief statement (Stated Belief), in the subsample where beliefs are extracted using ML (Extracted Belief), and in the combined sample of stated and extracted beliefs (All). Panel B shows the time-series average of the cross-sectional summary statistics for the stock-level aggregate beliefs (AB), computed as the difference between the proportion of bullish and bearish beliefs normalized by the total number of beliefs as in Eq (1). ABR(h) is abnormal returns over the next h trading days, starting $t+1$, following belief publication on day t. Column headers indicate the benchmark used to compute abnormal returns: CAPM or the three-factor (FF3) model. Mkt. Cap. is the market capitalization (in million USD) on the belief publication day.

I then employ the following regression specification to assess whether stock-level aggregate SMA beliefs predict future abnormal returns:

$$ABR_{k,t+1\to t+1+h} = \beta_0 + \beta_1 AB_{k,t} + \mathbf{X}\,\mathbf{\Gamma} + \epsilon_{k,t}, \tag{2}$$

where $ABR_{k,t+1\to t+1+h}$ is the future abnormal return of stock k over horizon $h \in \{5,\ 21,\ 63\}$ trading days, with t the belief publication day, and the return calculation horizon starting from day $t+1$ to avoid potential complications arising from the time of day when beliefs were published. $AB_{k,t}$ is the aggregate belief about stock k on day t, which is increasing in belief bullishness. $\mathbf{X}$ captures the following control variables: current abnormal return ($ABR_{k,t}$); past abnormal returns ($ABR_{k,t-1}$, $ABR_{k,t-2}$, and $ABR_{k,t-h\to t-3}$); $Volatility_{k,t}$; the number of professional stock analysts upgrading and downgrading stock k on day t, respectively; cash flow news sentiment; and year-month fixed effects. All continuous right-hand side variables are standardized to unit variance.

	CAPM		FF3		SBM	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 5 days						
$AB_{k,t}$	0.0018	0.0018	0.0017	0.0017	0.0017	0.0017
	(7.54)	(7.46)	(7.48)	(7.36)	(7.51)	(7.42)
Controls	NO	YES	NO	YES	NO	YES
Obs.	194,172	194,172	194,172	194,172	194,172	194,172
R^2 (%)	0.31	0.42	0.30	0.42	0.30	0.42
Panel B: 21 days						
$AB_{k,t}$	0.0026	0.0024	0.0024	0.0022	0.0025	0.0023
	(4.77)	(4.62)	(4.73)	(4.49)	(4.78)	(4.65)
Controls	NO	YES	NO	YES	NO	YES
Obs.	193,521	193,521	193,521	193,521	193,521	193,521
R^2 (%)	0.30	0.37	0.31	0.37	0.30	0.35
Panel C: 63 days						
$AB_{k,t}$	0.0036	0.0034	0.0028	0.0026	0.0033	0.0031
	(2.69)	(2.55)	(2.31)	(2.15)	(2.71)	(2.58)
Controls	NO	YES	NO	YES	NO	YES
Obs.	192,650	192,650	192,650	192,650	192,650	192,650
R^2 (%)	0.49	0.53	0.48	0.52	0.49	0.54

Table 2: SMA Beliefs and Future Abnormal Stock Returns. The table reports results for the panel regression of future abnormal returns ($ABR_{k,t+1\to t+1+h}$) on stock-level aggregate SMA beliefs ($AB_{k,t}$), with t the belief publication day. Abnormal returns are computed relative to either the CAPM, the three-factor model (FF3), or the size/book-to-market/momentum characteristic-based benchmark (SBM) of Daniel et al. (1997). Panels A, B, and C report results for horizon $h = 5$, 21 and 63 trading days, respectively. All regressions include year-month fixed effects. Control variables included in columns (2), (4), and (6) are past abnormal returns: $ABR_{k,t}$; $ABR_{k,t-1}$; $ABR_{k,t-2}$; $ABR_{k,t-h\to t-3}$; $Volatility_{k,t}$; the number of professional stock analysts upgrading and downgrading stock k on day t; and cash flow news sentiment on day t. All right-hand side variables are normalized to unit variance. In parentheses are t-statistics based on standard errors clustered by firm and year-month.

Table 2 shows the results of the regression. The column headers indicate the benchmark used to compute abnormal returns: CAPM (Columns 1 and 2); three-factor model or FF3 (Columns 3 and 4); and the size/book-to-market/momentum characteristic-based benchmark or SBM (Columns 5 and 6). For the sake of brevity, coefficients for the control variables are suppressed in the table but are shown in Table A2 in the Appendix for the 63-day horizon. Panel A of Table 2 shows that SMA beliefs significantly predict future abnormal returns over the future five-day horizon. If beliefs become more bullish by one standard deviation, abnormal returns increase by 17-18 bps over the next five trading days, with t-statistic around 7.5. The coefficients are stable across different specifications and abnormal return benchmarks.

Panels B and C of Table 2 show results for the one-month and three-month future horizons. Interestingly, predictability does not reverse, even in three months. A standard deviation increase in the bullishness of beliefs is associated with approximate 26-36 bps increase in abnormal returns over the next three months, with the t-statistic between 2.2 and 2.7. For longer horizons – six-month and 12-month (untabulated) – the predictability is insignificant, and there is no reversal, suggesting that SMA beliefs contain value-relevant information rather than sentiment unrelated to fundamentals or noise.

One could argue that SMA beliefs appear informative, because they piggyback on other fundamental-relevant news that drives returns. I address this concern by directly controlling for the number of professional stock analysts upgrading and downgrading a given stock, as well as sentiment in cash flow news published about a stock on the day of belief publication in Columns (2), (4), and (6) of Table 2. The fact that the results remain significant after these and additional controls which account for momentum and volatility is strong evidence that SMA beliefs are informative about future abnormal returns beyond the information contained in major news outlets and the recommendations of professional analysts.[18]

Nevertheless, it is still not fully clear whether SMA beliefs' predictability of returns is because the expressed beliefs contain value-relevant information, or because SMAs' views generate trades that drive stock prices in the direction of their stated beliefs. Therefore, to disentangle the information and price impact hypotheses, in the next section, I examine whether SMA beliefs predict earnings surprise. Since it is implausible that SMA beliefs drive either firm earnings or the consensus analyst forecast, evidence that the beliefs predict earnings surprise will support the information hypothesis.

4.2 Evidence from Earnings Surprise

I use the following regression specification to test whether SMA beliefs are informative about firm earnings:

$$SUE_{k,t} = \beta_0 + \beta_1 AB_{k,t-30\rightarrow t-1} + \mathbf{X}\,\mathbf{\Gamma} + \epsilon_{k,t}, \tag{3}$$

[18] I conduct further robustness tests, discussed in Section 7, by examining only stated beliefs, subsample analysis, and alternative belief aggregation. These alternative tests support the main result.

where t and k denote earnings announcement day and firm, respectively, $SUE_{k,t}$ is earnings surprise, computed as quarterly earnings per share minus consensus (average) professional analyst forecast divided by the stock price at the end of the last quarter. $AB_{k,t-30 \to t-1}$ is aggregate SMA belief regarding firm k (from Eq. (1)), averaged over the 30-day period ending one day before the earnings announcement. $\mathbf{X}$ captures the following control variables: lag of earnings surprise, price-scaled standard deviation of earnings forecasts, cumulative stock return over the past 30-day window, log of market value as of the end of the last quarter, log of book-to-market ratio as of the most recent fiscal year-end, cash flow news sentiment averaged over the past 30-day window, and year-month and single-digit industry SIC code fixed effects. I winsorize earnings surprise at the 1st and 99th percentiles to mitigate the influence of outliers and normalize all continuous regressors to unit variance to enhance interpretation.

Table 3 shows the regression results. Depending on the set of controls, the coefficient of average SMA beliefs (AB) is between 0.06% and 0.11%. In terms of economic magnitude, earnings surprise has a mean of 0.07%, median of 0.05% and standard deviation of 0.4%. Hence, if average SMA belief become more bullish by one standard deviation, earnings surprise increases by between 15% to 28% of its standard deviation. The fact that SMA beliefs significantly predict future earnings surprise, even after controlling for several characteristics and cash flow news sentiment, strongly supports the idea that SMA beliefs expressed on social media reflect value-relevant information.

Overall, the results of this section suggest that earlier results on stock returns are unlikely to be driven purely by the price impact of trades following SMAs' belief statements, thereby reinforcing the view that SMAs on the aggregate hold value for investors. However, it is essential to note that this conclusion does not easily extend beyond the group of nonprofessionals considered SMAs in this paper. In particular, SMAs do not include the unsophisticated crowd in some social media segments who express views that are not backed by fundamental analyses.

While it is reassuring to find that SMA beliefs reflect value-relevant information, we still do not know whether the average informativeness of beliefs arises from luck, or because SMAs consistently form correct beliefs. In other words, are SMA beliefs informative because most SMAs are skilled, or is the result driven by a small fraction of SMAs skilled in forming correct

	(1)	(2)	(3)	(4)
AB[-30, -1]	0.0011	0.0008	0.0006	0.0006
	(4.04)	(3.50)	(2.74)	(2.77)
Lagged(SUE)		0.0085	0.0084	0.0084
		(9.44)	(9.38)	(9.38)
Forecast Dispersion		-0.0017	-0.0014	-0.0014
		(-2.87)	(-2.40)	(-2.40)
Ret[-30, -1]			0.0015	0.0015
			(3.80)	(3.85)
Size			0.0018	0.0018
			(6.43)	(6.42)
Book-to-Market			0.0003	0.0003
			(1.53)	(1.55)
News Sentiment				-0.0001
				(-0.61)
Obs.	19,046	18,679	18,679	18,679
R^2 (%)	2.30	15.88	16.69	16.69

Table 3: SMA Beliefs and Earnings Surprise. The table reports the results for the panel regression of price-scaled earnings surprise ($SUE_{k,t}$) on stock-level SMA beliefs averaged over the 30-day window from $t-30$ to $t-1$, i.e., AB[-30, -1], where t is the earnings announcement day. Lagged(SUE) is the lag of price-scaled earnings surprise. Forecast Dispersion is the price-scaled standard deviation of earnings forecasts; Ret[-30, -1] is the cumulative return of stock k over the 30-day window $t-30$ to $t-1$; Size is the log of market capitalization as of June of the previous calendar year; Book-to-Market is the log of book-to-market ratio as of the most recent fiscal year-end; News Sentiment is cash flow news sentiment averaged over the 30-day window $t-30$ to $t-1$. Regressors are normalized to unit-variance, and all regressions include year-month and single-digit industry SIC code fixed effects. Values in parentheses are t-statistics based on standard errors clustered by firm and year-month.

beliefs?[19] Understanding the cross-sectional distribution of skill among SMAs will shed more light on how investors should approach individual SMAs' views, as well as SMAs' role in price discovery.

5 Modeling Belief Formation Ability as a Mixture Distribution

To understand the cross-sectional distribution of SMAs' ability to form correct beliefs, I model SMAs' ability as arising from a mixture distribution of multiple skill groups, following the recent literature on financial professionals' skill (Crane and Crotty, 2020; Harvey and Liu, 2018; Chen et al., 2017).[20] Modeling SMA skill as a mixture distribution avoids common pitfalls that arise from the low signal-to-noise feature of estimated abnormal returns – the standard measure of

[19] The ability to form correct beliefs may arise from experience, expertise, information-processing capabilities, or literacy levels.

[20] Papers such as Bajgrowicz and Scaillet (2012) and Barras et al. (2010) use the False Discovery Rate (FDR) method to differentiate luck from skill. However, Andrikogiannopoulou and Papakonstantinou (2019) show that the FDR method suffers from significant power problems due to the low signal-to-noise feature of returns data.

unobservable skill. Noise in estimated abnormal returns can result in conventional significance tests at the SMA-level misjudging good luck for skill or bad luck for lack of skill. Such tests also suffer from low test power, which makes it challenging to separate skill from luck. Conversely, the mixture distribution model can use information from the cross-section of SMA performance to reduce noise, and it is not impeded by low test power.

The formulation of the mixture model in this paper follows Crane and Crotty (2020). Assume that there is an unknown number J of skill groups. For each group $j \in \{0, 1, 2, ..., J\}$, there is a fraction π_j of SMAs with true belief formation ability, captured by abnormal returns, centered on μ_j. The dispersion of true abnormal returns for SMAs in group j is driven by variations in true ability, arising from investor-specific traits. Let $\alpha_i = \mu_j + \omega_i$ denote true belief formation ability of SMA i, where ω_i captures individual-specific traits and is normally distributed with zero mean and variance σ_j^2. On the other hand, estimated ability, $\widehat{\alpha}_i$, is measured with noise, e_i, which is assumed to be independent of ω_i and normally distributed with zero mean and variance s_i^2 (i.e., s_i is the standard error of estimated alpha). Thus, the estimated abnormal performance of an SMA is $\widehat{\alpha}_i = \mu_j + \omega_i + e_i$. Setting $J = 2$, for illustration, the specifications boil down to a two-component distribution of belief formation ability with the following density function:

$$f(\widehat{\alpha}_i) = \pi_0 \cdot \phi(\widehat{\alpha}_i; \mu_0, \sigma_{i,0}) \ + \ \pi_1 \cdot \phi(\widehat{\alpha}_i; \mu_1, \sigma_{i,1}), \tag{4}$$

where $\phi(\widehat{\alpha}_i; \mu_j, \sigma_{i,j})$ is the normal density function with mean μ_j and variance $\sigma_{i,j}^2 = \sigma_j^2 + s_i^2$ evaluated at $\widehat{\alpha}_i$.[21] The log-likelihood function L for a sample of N estimated SMA belief-formation ability is

$$L(\widehat{\alpha}_1, \widehat{\alpha}_2, ..., \widehat{\alpha}_N | s_1, s_2, ..., s_N, \Theta) = \sum_{i=1}^{N} \log(f(\widehat{\alpha}_i)), \tag{5}$$

[21] Assuming that the component distributions are normal allows for the interpretability of the model parameters. For instance, one could view true skill α_j as the sum of several random investor characteristics, which approaches normal distribution under the central limit theorem. Moreover, although component distributions are assumed to be normal, the composite distribution is not necessarily normally distributed but, rather, is fit to best characterize the data.

where the parameter set $\Theta = \{\pi_0, \pi_1, \mu_0, \mu_1, \sigma_0, \sigma_1\}$ is estimated via maximum likelihood, subject to the constraints that $0 \leq \pi_0 \leq 1$, $\pi_1 = 1 - \pi_0$, and $\sigma_j \geq 0$. To select the number of components in the mixture model, I use the sample-size-adjusted Bayesian information criterion (BIC) (Sclove, 1987; Schwarz, 1978).

To take the mixture model to the data, I use the five-day abnormal return relative to the three-factor model to measure estimated SMA belief formation ability.[22] I proceed by first computing abnormal returns ABR_i^k over window $t+1$ to $t+6$ trading days for each belief statement by SMA i about stock k on day t. ABR_i^k is then signed by pre-multiplying it by +1 for bullish beliefs (long positions) and -1 for bearish beliefs (short positions). Neutral beliefs are excluded, because they do not provide a clear investment signal. Finally, I aggregate each SMA's abnormal return by averaging across all belief statements, n_i, by SMA i as follows:

$$\overline{ABR}_i = \frac{1}{n_i} \sum_{k=1}^{n_i} ABR_i^k. \tag{6}$$

The main analysis uses SMAs with at least five belief statements over the sample period, i.e., $n_i \geq 5$. $\overline{ABR}_i$ is an SMA's estimated belief formation ability, $\widehat{\alpha}_i$. Its standard error, s_i, is calculated by clustering on the belief statement day to account for correlation across belief statements on the same day, and stock, to account for correlation in belief statements on the same stock.

Table 4 presents the summary statistics for $\overline{ABR}_i$ and s_i. The average estimated SMA ability is 26 bps, with a median of 11 bps. The standard errors s_i, with an average (median) of 136 (94) bps, suggest that the estimated SMA-specific abnormal returns are considerably noisy. Furthermore, the skewness (1.8) and kurtosis (24.1) of $\overline{ABR}_i$ suggest that the estimated abnormal return is not normally distributed, and the Kolmogorov-Smirnov test strongly rejects normality at the 1% significance level. These statistics suggest that standard significance tests based on normality can yield incorrect inferences regarding SMAs' ability, validating the application of a mixture distribution model to isolate true belief formation skill among SMAs.

[22] I use five-day abnormal returns because results in Section 4.1 show that the return predictability of beliefs declines with the horizon. As discussed in the robustness Section 7, the results are robust to alternative benchmarks for computing abnormal returns.

	Mean	SD	P10	P25	P50	P75	P90	Skewness	Kurtosis	Obs.	Frac. +ve	K-S p-value
$\overline{ABR}_i$ (%)	0.26	2.57	-1.83	-0.72	0.11	0.99	2.54	1.84	24.1	4,190	0.55	0.00
s_i (%)	1.36	1.61	0.31	0.54	0.94	1.61	2.81	6.20	72.4			

Table 4: Summary Statistics for SMA-specific Abnormal Returns. The table reports the summary statistics for the SMA-specific average five-day abnormal returns $\overline{ABR}_i$ and its standard error s_i, both in percent. Abnormal returns for belief statements result from buying stocks with bullish beliefs and selling stocks with bearish beliefs. The benchmark return for each event is based on the three-factor model. Abnormal returns are aggregated to the SMA level by estimating the average across all of an SMA's belief statements. Standard errors s_i are clustered by publication date and stock. The reported "Frac. +ve" is the fraction of the SMA cross-section with positive estimated abnormal return. The K-S p-value is the p-value of a Kolmogorov-Smirnov test of the null hypothesis that the demeaned cross-sectional distribution of $\overline{ABR}_i$ is normally distributed.

5.1 How Many SMAs Have Belief Formation Ability?

Following the BIC model selection criteria, I estimate a two-component mixture model, where component $j = 0$ comprises the low-skill SMAs and $j = 1$ comprises the high-skill SMAs. Table 5 reports the parameter estimates. Columns (1) and (2) of Panel A show estimates corresponding to the low-skill and high-skill SMA groups, respectively. The estimated fraction of low-skill SMAs (i.e., π_0) is roughly 87%, with the skill distribution centred on an abnormal return of 6 bps, with a dispersion of 0.4%. Conversely, the fraction of high-skill SMAs (i.e., π_1) is 13%, with a much larger five-day average abnormal return of 61 bps (31% annualized) and a dispersion of 3.2%.

Panel B of Table 5 summarizes the overall distribution of abnormal returns implied by the mixture model. Importantly, roughly half of the SMAs (56%) have genuinely positive average abnormal returns following their belief statements, i.e. have the ability to form correct beliefs.[23] However, there is enormous dispersion in SMAs' true ability to form correct beliefs, as the cross-sectional standard deviation of true ability is 1.3%, which is roughly 51% of the estimated abnormal return's dispersion (i.e., 2.6% reported in Table 4, which includes variations in abnormal return attributable to luck/noise). The substantial heterogeneity in SMAs' true ability and the small fraction of high-skill SMAs point towards the difficulty investors might face in identifying skilled individual SMAs on social media. To provide some context, Crane and Crotty (2020) use a similar setup and estimate the fraction of high-skill professional analysts to

[23] Fraction positive is computed as $1-\left[\pi_0 \cdot \Phi(\frac{0-\mu_0}{\sigma_0}) + \pi_1 \cdot \Phi(\frac{0-\mu_1}{\sigma_1})\right]$. For a given percentile P, the corresponding quantile q is computed numerically by solving $P = \pi_0 \cdot \Phi\left(\frac{q-\mu_0}{\sigma_0}\right) + \pi_1 \cdot \Phi\left(\frac{q-\mu_1}{\sigma_1}\right)$, where $\Phi(\cdot)$ is the cumulative normal distribution function.

Panel A: Mixture Model Parameters		
	(1) Component 0	(2) Component 1
π	0.8719	0.1281
	(0.0419)	(0.0419)
μ	0.0006	0.0061
	(0.0002)	(0.0021)
σ	0.0040	0.0322
	(0.0014)	(0.0038)
$\sigma_{i,j}$	0.0142	0.0350
	(0.0010)	(0.0034)

Panel B: Mixture Return Distribution		
	Estimate	SE
Mean	0.0013	(0.0002)
SD	0.0132	(0.0003)
P10	-0.0056	(0.0007)
P25	-0.0024	(0.0004)
P50	0.0007	(0.0002)
P75	0.0038	(0.0004)
P90	0.0075	(0.0006)
Fraction positive	0.5608	(0.0164)
N	4,190	

Table 5: SMA Belief Formation Ability: Two-component Mixture model. The table reports the result for the two-component mixture model of belief formation skill using data for only SMAs with at least five belief statements. Panel A reports the estimates of the model parameters, where π is the fraction of low and high-type SMAs, μ is the mean of each group's true belief formation skill, σ is its dispersion, and $\sigma_{i,j}$ is the average dispersion of the estimated skill of each group. Each SMA's estimated abnormal return ($\overline{ABR}_i$) is computed relative to the three-factor model for all publications by the SMA, as in Eq. (6). The reported $\sigma_{i,j}$ is based on the cross-sectional average of $\overline{ABR}_i$'s standard error s_i. Hence, $\sigma_{i,j} = \sqrt{\sigma_j^2 + \overline{s}^2}$. Estimates in Panel A are used to compute statistics for the cross-sectional distribution of belief formation skill reported in Panel B. P(10) – P(90) are percentiles of the implied cross-sectional distribution of SMA skill. Standard errors (in parentheses) are computed as the standard deviation of the statistics from 1,000 bootstrap replications.

be 36%, with roughly 97% sufficiently skilled to generate positive abnormal returns. Therefore, although SMAs as a group tend to add value, individual SMAs' skill is relatively limited.

Have SMAs become better at forming correct beliefs over time? To answer this question, I re-estimate the two-component mixture model in two subsamples. The first subsample comprises the first half of the sample, 2005 – 2012, while the second comprises the second half, 2013 – 2019. Figure 3 plots the fraction of high-skill SMAs and their true ability in terms of average abnormal returns for the two subsamples. The fraction of high-skill SMAs increased from 8% in the first subsample to 13% in the second, with average abnormal return rising from 0.57% to 0.68%. This indicates that the fraction of high-skill SMAs on Seeking Alpha improved over time, suggesting that the rise in Seeking Alpha's popularity among investors over the years potentially benefits investors. However, there are at least two reasons why the fraction of skilled SMAs on Seeking Alpha might improve with time: SMAs might have learned from experience, or more highly skilled individuals might have joined the Seeking Alpha platform as it gained prominence. Indeed, analysis in the next section points towards the second channel.

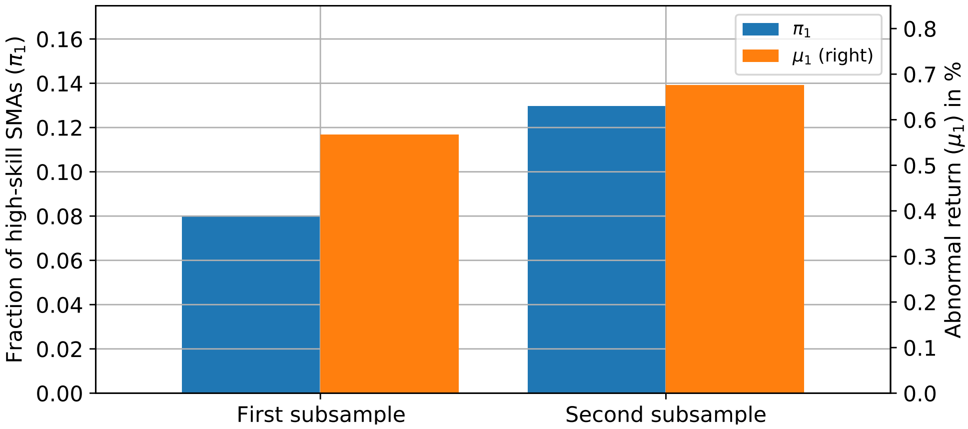

Figure 3: Distribution of SMAs' Ability in Two Subsamples. The figure shows the fraction of high-skill SMAs π_1 and their true ability in terms of five-day abnormal returns μ_1 in two subsamples. The first subsample comprises the period from 2005 – 2012, while the second subsample covers 2013 – 2019. The estimates are based on a two-component mixture model with only SMAs with at least five belief statements in each subsample. The number of SMAs in the first and second subsamples are 1,444 and 3,090, respectively.

One may wonder whether the results presented so far depend on the minimum number of belief statements required for computing the estimated skill ($\overline{ABR_i}$) and its standard error (s_i), or the number of components in the mixture model. Robustness tests discussed in Section 7 indicate that this is not the case. Taken together, the potential difficulty in identifying the few high-skill SMAs and the earlier evidence on average SMA beliefs' predictive power suggest that investors may be better off following SMAs' consensus beliefs rather than relying on the view of an individual SMA.

5.2 SMA Characteristics and Belief Formation Ability

Since SMAs are substantially heterogeneous in their ability to form correct beliefs, identifying SMA attributes associated with skill could reduce investors' search costs for valuable information on social media. I therefore incorporate SMA attributes observable on the Seeking Alpha platform in modelling the probability that an SMA is high- or low-skill. Following Crane and

Crotty (2020), I parameterize the probabilities using the logistic function:

$$\pi_{i,0} = \frac{1}{1+\exp(b_0 + b_1 x_i)}; \qquad \pi_{i,1} = 1 - \pi_{i,0}, \tag{7}$$

where x_i is a dummy variable that equals one if a certain SMA characteristic is above the cross-sectional median and zero otherwise. With this parameterization, the density function (4) has additional parameters b_0 and b_1, and the set of parameters in the maximum likelihood problem of Eq. (5) is now $\Theta = \{b_0, b_1, \mu_0, \mu_1, \sigma_0, \sigma_1\}$. Once these parameters are estimated, the probabilities $\pi_{i,0}$ and $\pi_{i,1}$ for SMAs with low ($x = 0$) and high ($x = 1$) values of the characteristic can be calculated using the estimates of b_0 and b_1.

I consider the following SMA characteristics computed over the entire sample period: industry specialization; firm specialization; workload; experience on Seeking Alpha; effort, proxied by the average number of words in opinion articles corresponding to each belief statement; popularity on Seeking Alpha; disagreements with other Seeking Alpha users, and; skin-in-the-game, proxied by the fraction of time an SMA discloses an investment position in the stock about which he/she expresses a belief. Appendix A.1 describes the construction of these characteristics..

Table 6 shows the results for SMAs' skill conditional on characteristics. The table reports, for each investor characteristic x, the fraction of high- and low-skill SMAs, average abnormal return, and its standard deviation, implied by the mixture model conditional on whether x is above or below its cross-sectional median.

Supporting theoretical results on gains from specialization in information acquisition (e.g., Van Nieuwerburgh and Veldkamp, 2010), I find that the most distinctive characteristics for separating high- and low-skill SMAs are industry and firm specializations. For instance, SMAs specializing in a few industries (above median specialization) have a 41% probability of being considered high skill compared to 0.5% for SMAs that cover many industries. The model-implied average abnormal return is 15 bps lower for SMAs with less industry specialization. Similarly, SMAs with a lower workload (below median average publications per year) have a 24% probability of being considered high skill compared to 7% for SMAs with more publications per year, consistent with models of limited attention and cognitive capacity constraints (e.g.,

Panel A								
	Specialization							
	Industry		Firm		Workload		Skin-in-the-game	
	$x = 0$	$x = 1$	$x = 0$	$x = 1$	$x = 0$	$x = 1$	$x = 0$	$x = 1$
Mean	0.0006	0.0021**	0.0008	0.0019***	0.0018	0.0010***	0.0009	0.0018***
SD	0.0139	0.0153***	0.0135	0.0136	0.0134	0.0133	0.0134	0.0133
π_0	0.9944	0.5895**	0.9566	0.7302***	0.7570	0.9267**	0.9373	0.7911***
π_1	0.0056	0.4105**	0.0434	0.2698***	0.2430	0.0733**	0.0627	0.2089***

Panel B								
	Effort		Disagreement		Popularity		SA Experience	
	$x = 0$	$x = 1$	$x = 0$	$x = 1$	$x = 0$	$x = 1$	$x = 0$	$x = 1$
Mean	0.0011	0.0015***	0.0016	0.0010***	0.0012	0.0015***	0.0015	0.0010**
SD	0.0133	0.0132	0.0132	0.0133	0.0133	0.0132	0.0132	0.0133
π_0	0.9136	0.8368***	0.8207	0.9244***	0.9035	0.8462**	0.7977	0.9163**
π_1	0.0864	0.1632***	0.1793	0.0756***	0.0965	0.1538**	0.2023	0.0837**

Table 6: SMA Characteristics and Belief Formation Ability. The table reports the results for the cross-sectional distribution of SMAs' belief formation ability from a two-component mixture model, where the proportion of SMAs in each component depends on an SMA characteristic, as shown in Eq. (7). The sample includes only SMAs with at least five belief statements. The table reports the estimated proportion (π) of the low- and high-type SMAs for the below- ($x = 0$) and above-median ($x = 1$) SMAs for a given characteristic. Also shown in the table are the conditional mean of SMAs' true ability (Mean) and its standard deviation (SD), implied by the mixture model. Column headers indicate the SMA characteristics described in Appendix A.1. ***, **, * indicate the statistical significance of the one-sided test of the difference between groups at the 1%, 5% and 10% significance levels, respectively. The p-values are based on bootstrap distribution with 1,000 bootstrap replications.

Hirshleifer et al., 2011). These results indicate that SMAs who specialize less and have a heavier workload are likely less able to effectively process pieces of information to obtain precise signals.

SMAs who have stayed longer on the Seeking Alpha platform (SA experience) have a 12 pp lower probability of being considered high skill than those with below-median years of Seeking Alpha experience. This suggests that the earlier results in Section 5.1, showing a higher fraction of high-skill SMAs on Seeking Alpha in the second half of the sample, are driven by more skilled individuals who joined Seeking Alpha as the platform became more popular over time.

SMAs who often have an investment in the stock they express belief in (above median Skin-in-the-game) have a 14 pp higher probability of being considered high skill, with 9 bps higher performance. To the extent that SMAs' truthfully disclose their investments, this result suggests that more skin in the game motivates more diligent information acquisition and processing, leading to superior performance. Consistent with this view, Campbell et al. (2019) show that having an investment position in a stock does not impair the informativeness of opinions expressed by nonprofessional analysts. Table 6 further shows that SMAs who invest more effort (write longer articles), who are more popular (receive more comments on their publications), and whose views other investors tend to disagree less with are also more likely to be considered high skill.

Overall, the analyses indicate that specific SMA characteristics, particularly firm and industry specialization, can help investors identify skilled SMAs on social media. Furthermore, there are differences in how nonprofessional and professional analyst characteristics relate to skill. For instance, while the literature shows that professional analysts who issue more sell recommendations tend to be more skilled (e.g., Barber et al., 2006), I do not find that to be the case for SMAs.

5.3 Can Investors Profitably Trade on SMAs' Beliefs?

We have seen that average SMAs' beliefs are informative about stock returns. However, there is substantial heterogeneity in SMAs' ability to form correct beliefs, which raises the question: Can investors rely on individual SMAs' beliefs to form profitable portfolios? I address this question by examining simple transaction-based, calendar-time portfolios, as in Jeng et al. (2003) and Seasholes and Zhu (2010). Precisely, I form buy-and-hold Bullish and Bearish portfolios by putting a unit of stock k in the Bullish (Bearish) portfolio on day t whenever an SMA publishes a bullish (bearish) belief about stock k on day $t-1$. The position is then held for period $h \in \{21, 63, 126\}$ trading days. Daily returns for each portfolio are value-weighted based on the number of units of each stock in the portfolio on day t and the closing stock prices on day $t-1$.

Compared to daily rebalancing, using buy-and-hold portfolios for the analysis allows for realistic portfolio strategies, implementable at moderate trading costs in practice.[24] Similarly, adding stocks to portfolios one day after the belief statement ensures that investors have sufficient time to observe SMAs' beliefs and trade on them. It further avoids complications surrounding the time of day when SMAs express beliefs.

[24] Furthermore, buy-and-hold portfolios correct for noisy prices that can bias portfolio-based tests (see, Blume and Stambaugh, 1983).

	No. Stocks	Exc. Ret. (%)	Alpha (%)			
			CAPM	FF3	Carhart	FF5+MOM
Panel A: 21 days						
Bullish	460	1.205	0.444	0.320	0.321	0.321
		(4.71)	(2.99)	(2.11)	(2.12)	(2.19)
Bearish	117	0.669	-0.152	-0.279	-0.218	-0.131
		(1.63)	(-0.62)	(-1.38)	(-1.10)	(-0.55)
Bullish-Bearish		0.536	0.596	0.599	0.539	0.453
		(1.79)	(1.89)	(2.02)	(1.82)	(1.32)
Panel B: 63 days						
Bullish	902	1.079	0.471	0.303	0.304	0.329
		(3.58)	(3.29)	(1.88)	(1.88)	(2.07)
Bearish	262	0.603	-0.043	-0.221	-0.183	-0.063
		(1.51)	(-0.18)	(-1.14)	(-1.00)	(-0.33)
Bullish-Bearish		0.476	0.514	0.524	0.487	0.392
		(1.60)	(1.63)	(1.69)	(1.62)	(1.26)
Panel C: 126 days						
Bullish	1247	1.254	0.487	0.303	0.305	0.319
		(4.50)	(3.42)	(1.94)	(1.95)	(2.04)
Bearish	401	0.959	0.138	-0.056	-0.011	0.104
		(2.27)	(0.47)	(-0.27)	(-0.06)	(0.53)
Bullish-Bearish		0.295	0.349	0.359	0.316	0.215
		(0.96)	(1.08)	(1.16)	(1.06)	(0.71)

Table 7: Belief-based Portfolio Performance. The table reports average monthly excess returns (Exc. Ret.) and alphas in per cent for portfolios based on bullish and bearish SMA beliefs. For each SMA belief statement on stock k on day $t-1$, a unit of the stock is added to the Bullish portfolio if the belief is bullish or to the Bearish portfolio if the belief is bearish on day t. The position is then held for 21 (Panel A), 63 (Panel B) and 126 (Panel C) trading days. "Bullish-Bearish" is the difference portfolio that is long the Bullish portfolio and short the Bearish portfolio. The daily portfolio returns are cumulated to monthly frequency to compute excess returns and alphas. The column "No. Stocks" indicates the average number of unique stocks in each portfolio daily. Alphas are relative to CAPM, three-factor (FF3), four-factor (Carhart), and five-factor plus momentum (FF5+MOM) benchmarks indicated in the column headers. Reported in parentheses are t-statistics based on the Newey and West (1987) method. The SMA belief data range between January 2005 - December 2019.

Table 7 summarizes the performance of the Bullish and Bearish portfolios. The Bearish portfolio has fewer stocks than the Bullish portfolio, which is not surprising given that SMA beliefs are generally more bullish. For the 21-trading-day holding period, the Bullish (Bearish) portfolio has an average excess return of 1.2% (0.54%) per month. Furthermore, while the Bullish portfolio has a positive and significant alpha of between 0.32% to 0.44% per month depending on the benchmark, the Bearish portfolio has negative alphas. The difference portfolio yields a four-factor alpha of 0.60% per month (7.2% annualized). The portfolios' performance is similar in magnitude for the longer holding periods of 63 and 126 trading days, but the alphas are less significant. For instance, while the Bearish portfolio's alphas remain negative for these holding periods, the alphas of the Bullish portfolio range between 0.30% and 0.48%, with t-statistics between 1.88 and 2.07 for the more stringent factor benchmarks.[25]

Considering that individual investors, instead of institutional investors, primarily follow SMAs on social media, it is interesting to observe that a portfolio of only long positions with a decent holding period, and hence modest portfolio turnover and transaction cost, yields significant abnormal returns. Therefore, investors can trade profitably on SMAs' beliefs and can even do better if they can identify the few highly skilled SMAs documented earlier.

6 What Informs Social Media Analysts' Beliefs?

I now examine channels through which SMAs form beliefs about stock returns by testing two theory-inspired hypotheses on belief formation. Section 6.1 explores the existence and nature of herding among SMAs, and Section 6.2 examines the role of return extrapolation in beliefs.

6.1 Evidence on Herding

Social media can serve as a coordination mechanism for mutual imitation, i.e., herding, because it quickens information transmission and enhances individuals' ability to observe peers' actions.[26] Hence, SMAs may herd in stating their beliefs on social media, in line with the reputational herding and information cascade models (e.g., Banerjee, 1992; Bikhchandani et al., 1992; Scharfstein and Stein, 1990). On the other hand, since SMAs face different incentives from professional investment analysts, it is equally possible that herding is less pervasive among them. Since herding might intensify mispricing if it is based on little or no information or promote price discovery if it is caused by fundamental information, I examine the existence and nature of herding among SMAs to shed further light on their role in financial markets.

To test for the existence of herding among SMAs, I adopt the herding test developed in Welch (2000), which is applicable in settings where choices are discrete (e.g., Bullish, Neutral and Bearish belief statements). Let θ represent a parameter that measures whether the SMA

[25] I conduct robustness tests, discussed in Section 7, by excluding penny stocks and microcap stocks and find qualitatively similar results.

[26] Herding requires a coordination mechanism which can be either a widely spread rule to coordinate based on some signal, or a direct ability to observe other decision makers' decisions (Devenow and Welch, 1996).

belief transition probability matrix $\mathbf{P}$ depends on the observed consensus C. $\mathbf{P}$ is a 3-by-3 matrix whose elements, $p_{i,j}$, capture the probability of an SMA moving from a previous belief statement, row i, to a new belief statement, column j. For example, transitioning from a Bearish to a Bullish belief. If $\theta = 0$, then $\mathbf{P}$ is independent of the consensus – the null hypothesis. Conversely, $\theta > 0$ indicates a tendency for belief statements to follow the consensus, while $\theta < 0$ indicates a tendency to avoid the consensus. $\mathbf{P}$ is then defined as a function of θ and the consensus C as:

$$p_{i,j}(\theta, C) \equiv \overline{p}_{i,j} \left\{ \frac{\left[1 + (j - C)^2\right]^{-\theta}}{D_i} \right\}, \tag{8}$$

where index i, $j \in \{1, 2, 3\}$ denote belief states: 1 represents Bearish, 2 Neutral, and 3 Bullish; the denominator $D_i = \sum_{j=1}^{3} \overline{p}_{i,j} \left[1 + (j - C)^2\right]^{-\theta}$ ensures that rows of the transition matrix sum to 1; $\overline{p}_{i,j}$ is the unconditional transition probability from i to j, estimated from the historical SMA belief revisions; and C is the target towards which SMAs may herd, i.e., consensus belief about a given stock on a given day. The estimator $\hat{\theta}$ is obtained by maximizing the log-likelihood function over the sample period, and statistical inference is based on the likelihood ratio test. Appendix A.3 describes the estimation and inference procedure.[27]

To implement the test, I use SMAs' belief revisions not older than one year over the period spanning from January 2006 – December 2019.[28] For an SMA's belief revision about stock k on day t, the consensus is computed as the equal-weighted or characteristic-weighted average of other SMAs' beliefs on the same stock over the past six-month period ending $t-2$. The window for estimating the consensus ends at $t-2$ to ensure that SMAs had sufficient time to observe the consensus. More so, the consensus is estimated only if there are at least two belief statements by other SMAs over the six-month period.[29] The motivation for the characteristic-weighted consensus is to see if SMAs herd more towards views of peers with qualities identified to be associated with skill in Section 5.2. To compute the characteristic-weighted consensus, SMA characteristics – industry specialization, effort, and popularity (defined in Appendix A.1) – are

[27]Using Monte-Carlo simulations, Welch (2000) shows that the test is neither mechanically driven by the discrete/limited number of choices nor by the fact that the target (consensus) itself is the outcome of prior transitions.

[28]I use belief revisions not older than one year to avoid stale beliefs. The sample starts from 2006 due to very few belief revisions in 2005.

[29]Robustness Section 7 shows that the results are robust to estimating the consensus over an alternative window.

computed as of the last calendar month using data over the past one year. Missing values for each characteristic are replaced with the median value.

Table 8, Panel A, shows the estimated herding coefficient $\hat{\theta}$ and the associated χ^2 p-value for the equal-weighted consensus and characteristic-weighted consensus beliefs. The estimated herding coefficient is around 0.27 and is significant, regardless of the consensus weighting scheme. This suggests that SMAs herd towards the consensus when stating their beliefs. Furthermore, the fact that the estimated herding coefficient is insensitive to the consensus weighting scheme indicates that the views of more skilled SMAs do not disproportionately influence beliefs.

Panel A: Estimated Herding Coefficient		
Concensus is	$\hat{\theta}_0$	χ^2 $p-value$
Equal-weighted	0.267	0.000
Specialization-weighted	0.257	0.000
Effort-weighted	0.269	0.000
Popularity-weighted	0.268	0.000

Panel B: Probability of Hitting Target								
	Herding Coefficient $\hat{\theta}_0$							
Target	-10	-1	0	0.15	0.25	0.5	1	10
1 (Bearish)	0.000	0.037	0.159	0.193	0.218	0.292	0.469	1.000
2 (Neutral)	0.000	0.019	0.037	0.041	0.044	0.052	0.072	0.975
3 (Bullish)	0.000	0.481	0.804	0.835	0.854	0.892	0.941	1.000

Table 8: Herding among SMAs. The table reports results for the herding test of Eq. (8) in Panel A and the economic significance of herding in Panel B. Panel A shows the estimated herding coefficient $\hat{\theta}$ and χ^2 p-value for different targets (consensus estimates). Panel B shows the probability of a belief revision hitting a hypothetical Bearish, Neutral or Bullish target for different values of $\hat{\theta}$. If $\hat{\theta} = -\infty$, the target will always be avoided. If $\hat{\theta} = 0$, the probability of hitting the target is equal to the unconditional probability of hitting the target. If $\hat{\theta} = \infty$, the target will always be hit. Values in Panel B were produced using the unconditional transition matrix and hypothetical values for $\hat{\theta}$.

Panel B of Table 8 shows the economic implications of the estimated herding coefficient. The column with $\hat{\theta} = 0$ captures the unconditional probabilities of an SMA stating a Bearish, Neutral or Bullish revision. Focusing on the column with $\hat{\theta} = 0.25$, it follows that the estimated herding coefficients reported in Panel A imply that herding increases the probability of hitting a hypothetical Bullish (Bearish) belief target by about 6 (7) pp. This indicates a moderate level of herding among SMAs, which is interesting given that SMAs can deviate from the consensus in an attempt to attract attention and readership.

6.1.1 When is Herding more Pronounced?

The nature of SMAs' herding has relevant implications for market efficiency. For instance, if SMAs herd when the consensus is wrong and many investors follow their views, prices could be driven away from fundamentals, with potential costs to the real economy through inefficient capital allocation. Therefore, to understand the nature of SMAs' herding, I modify Eq. (8) by making θ a function of some variable y: $\theta(y) = \theta_0 + \theta_1 y$. If $\theta_1 = 0$, herding does not depend on y. Conversely, $\theta_1 > 0$ indicates that herding increases with y, while $\theta_1 < 0$ indicates that herding decreases with y.

Table 9 shows the results for the estimated $\hat{\theta}_0$ and $\hat{\theta}_1$ and the associated χ^2 p-values. Panel A shows the results, where y is either the indicator variable for NBER recession (Recession), the indicator variable for whether market volatility exceeds its sample median (High Volatility), or consensus optimism (CO), measured as $CO = C - 2$, where C is the equal-weighted consensus. Because beliefs are labelled as 3 (Bullish), 2 (Neutral), and 1 (Bearish), $CO > 0$ ($CO < 0$) implies an optimistic (pessimistic) consensus. Panel B shows the results, where y is a measure of consensus correctness (CC), quantified as consensus optimism times future stock return: $CC = CO \times Ret(h)$, where $Ret(h)$ is the future horizon h (starting $t+1$) return of a stock. CC is positive when the consensus is correct ex-post, that is an optimistic (pessimistic) consensus is followed by a positive (negative) future stock return.[30] Appendix A.1 provides more details about the variables.

Panel A of Table 9 indicates that $\hat{\theta}_1$ is negative and significant for recession and high volatility periods, suggesting that SMAs herd less during economic downturns and episodes of high market uncertainty. The $\hat{\theta}_1$ coefficient of -0.15 for recession implies that during recessions, SMAs' herding falls by roughly 54% of its unconditional value. Therefore, during economic downturns, SMAs tend to rely more on their private information, consistent with the theoretical literature that predicts polarization of beliefs in bad times due to some agents giving little weight to news (e.g., Cujean and Hasler, 2017). It also tallies with the nature of herding among professional analysts, who tend to herd more in good times (Welch, 2000). On the other hand, unlike professional analysts, SMAs herd less when the consensus is more optimistic. In terms of economic

[30]The results are robust to using abnormal returns relative to the CAPM or the three-factor model.

Panel A:	Recession		High Volatility		Cons. Optimism (CO)	
State of Economy	$\hat{\theta}_0$	$\hat{\theta}_1$	$\hat{\theta}_0$	$\hat{\theta}_1$	$\hat{\theta}_0$	$\hat{\theta}_1$
Estimate	0.274	-0.145	0.286	-0.053	0.422	-0.212
χ^2 $p-value$	0.000	0.000	0.000	0.002	0.000	0.000
Panel B:	$CO \times Ret(5)$		$CO \times Ret(63)$		$CO \times Ret(126)$	
Cons. Correctness	$\hat{\theta}_0$	$\hat{\theta}_1$	$\hat{\theta}_0$	$\hat{\theta}_1$	$\hat{\theta}_0$	$\hat{\theta}_1$
Estimate	0.267	0.461	0.265	0.110	0.269	0.012
χ^2 $p-value$	0.000	0.001	0.000	0.010	0.000	0.694

Table 9: SMAs' Herding Conditional on Economic States. The table reports results for the herding test, conditional on the realization of some variable y, such that the estimated herding coefficient $\hat{\theta} = \hat{\theta}_0 + \hat{\theta}_1 y$. Panel A shows the results when y is either the indicator for NBER recession (Recession), the indicator for market volatility higher than its sample median (Market Volatility), and consensus optimism (CO), measured as the equal-weighted consensus minus 2. Panel B reports the results when y is consensus correctness, measured as $CO \times Ret(h)$, where $Ret(h)$ is the future stock return over horizon $h \in \{5, 63, 126\}$ trading days.

magnitude, with $\hat{\theta}_1$ of -0.21 and a standard deviation of 0.39 for consensus optimism, a standard deviation increase in consensus optimism reduces the herding coefficient by about 8.2 pp – a 31% decline relative to its unconditional value. This result supports the view that SMAs favour more attention-grabbing deviation from the consensus to attract readership.

Panel B of Table 9 shows that across different future return horizons $h \in \{5, 63, 126\}$ trading days, the incremental herding coefficient $\hat{\theta}_1$ for consensus correctness is positive. This implies stronger herding when the consensus is correct, i.e., when the consensus is optimistic (pessimistic), and the future return turns out to be positive (negative).[31] In terms of economic magnitude, the standard deviation of $CO \times Ret(h)$ for the five-day and 63-day horizons are 5% and 17%, which implies that herding increases by between 7% and 9% of its unconditional value for a standard deviation increase in consensus correctness. The fact that incremental herding when the consensus is correct remains positive, even for the six-month future return horizon, suggests that SMAs herd on fundamental information incorporated in the consensus rather than sentiment. The implication, therefore, is that instead of herding irrationally (Simonsohn and Ariely, 2008) or herding based on no information (Scharfstein and Stein, 1990), SMAs tend to learn fundamental information from the belief statements of their peers to improve forecasts, consistent with information-based herding models (e.g., Banerjee, 1992).[32]

[31] The statistical significance of the incremental herding coefficient decreases with the horizon, turning insignificant for the 126-day and 252-day horizons.

[32] An alternative explanation, which I cannot differentiate, is that SMAs independently follow the same fundamental information. There is also the possibility that the consensus moves prices. I, however, favour the

6.2 Evidence on Return Extrapolation

Survey evidence and theoretical models on belief formation suggest that investors extrapolate from past returns when forming expectations about future returns (e.g., Greenwood and Shleifer, 2014; Barberis et al., 2015). However, most empirical studies on return extrapolation focus on the aggregate stock market due to data limitations. The recent work of Da et al. (2021) provides evidence on extrapolation in belief formation about individual stocks, demonstrating that extrapolated beliefs predict stock returns with the wrong sign. On the contrary, the analyses in previous sections of this paper indicate that SMAs' beliefs correctly provide value-relevant information, raising the following question: Do SMAs extrapolate while being correct on average?

To test whether SMAs extrapolate from past returns, i.e., whether they form expectations of future individual stock returns by relying more (less) heavily on recent (distant) stock returns, I regress SMA beliefs about stock k on the stock's past non-overlapping weekly returns over the past three months, as follows:

$$AB_{k,t} = \beta_0 + \sum_{\tau=1}^{12} \beta_\tau Ret(\tau)_{k,t} + \mathbf{X}\,\mathbf{\Gamma} + \epsilon_{k,t}, \tag{9}$$

where $AB_{k,t}$ is SMAs' aggregate beliefs about stock k on day t computed from Eq. (1); $Ret(\tau)_{k,t}$ is stock k's past τ'th non-overlapping one-week (five trading days) return, with the most recent return window ending two days before the belief statement day t.[33] That is, $Ret(1)_{k,t} = Ret_{k,t-6\to t-2}$, $Ret(2)_{k,t} = Ret_{k,t-11\to t-7}$, and so on. $\mathbf{X}$ is a vector of control variables that might influence beliefs, namely lagged belief, stock k's cash flow news sentiment averaged over the past week, and professional analysts' consensus forecast of quarterly earnings as of the last calendar month. The regression includes year-month fixed effects to absorb common time trends in SMAs' beliefs.

Figure 4 plots the coefficient estimates and 95% confidence intervals for the lagged weekly returns $Ret(\tau)$. The most recent one-week return has the largest influence on beliefs, with a

fundamental information story because, as shown in Section 4.2, aggregated beliefs contain value-relevant information.

[33] Calculation of past weekly returns and control variables ends at least two days before the belief publication day to ensure that the variables were observable by SMAs and hence could guide belief formation.

relatively tight 95% confidence interval. More so, the effect of past returns generally declines with time, consistent with return extrapolation in SMAs' belief formation. In terms of economic magnitude, the coefficient of 0.26 for the most recent one-week return and the SMA beliefs' standard deviation of 0.66 implies that the stock-level aggregate belief becomes significantly more bullish by roughly 40% of its standard deviation when a stock's price doubles over the past one-week period. Conversely, the influence of older returns is much lower, becoming indistinguishable from zero by two months.[34]

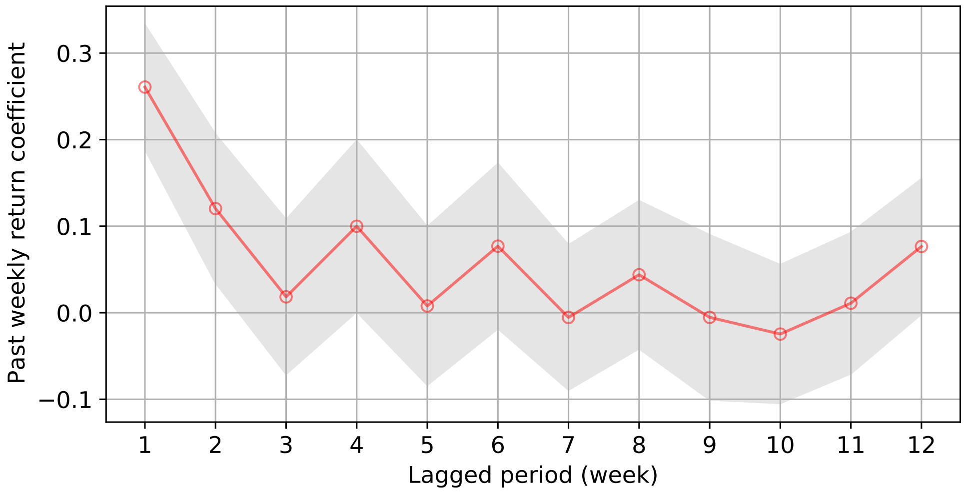

Figure 4: Return Extrapolation in SMAs' Belief Formation. The figure shows coefficient estimates and 95% confidence intervals (grey area) from regressing stock-level aggregate SMA belief $AB_{k,t}$ on stocks' past non-overlapping weekly (five trading days) returns, as described in Eq. (9). The regression controls for the lag of $AB_{k,t}$, cash flow news sentiment, professional analysts' consensus forecasts of quarterly earnings, and year-month fixed effects. Standard errors are clustered by stock and year-month.

SMAs' return extrapolation might differ across stocks depending on the salience of returns. For example, extrapolation might be stronger for stocks that belong to popular indices, such as the S&P 500 index, which receive substantial investor attention. Similarly, extrapolation could be stronger for stocks that receive more media coverage, making past returns easy to recall.[35] Therefore, to examine whether SMAs' return extrapolation depends on salience, I use an exponential decay specification (as in Greenwood and Shleifer, 2014; Da et al., 2021) to

[34]Table A10 in the Appendix shows additional results without control variables as well as subsample analysis, all supporting the extrapolation hypothesis.

[35]Salience bias is the tendency to irrationally overweight more readily available information, resulting in the overestimation of the risk of salient events based on vividness, proximity, or emotional impact (Tversky and Kahneman, 1973).

summarize the degree of extrapolation conditional on stocks' salience characteristics as follows:

$$
\begin{gathered}
AB_{k,t} = \lambda_0 + \lambda_{1,s} \cdot \sum_{\tau=1}^{12} \mathbb{1}_s \cdot w_{s,\tau} Ret(\tau)_{k,t} + \lambda_{1,ns} \cdot \sum_{\tau=1}^{12} \mathbb{1}_{ns} \cdot w_{ns,\tau} Ret(\tau)_{k,t} + \mathbf{X}\,\mathbf{\Gamma} + \epsilon_{k,t}, \\
w_{s,\tau} = \frac{\lambda_{2,s}^{\tau-1}}{\sum_{i=1}^{12} \lambda_{2,s}^{i-1}}, \quad w_{ns,\tau} = \frac{\lambda_{2,ns}^{\tau-1}}{\sum_{i=1}^{12} \lambda_{2,ns}^{i-1}}, \quad 0 \leq \lambda_{2,s} < 1 \text{ and } 0 \leq \lambda_{2,ns} < 1,
\end{gathered}
\tag{10}
$$

where $AB_{k,t}$ is SMAs' aggregate belief about stock k on the belief publication day t computed from Eq. (1); $Ret(\tau)_{k,t}$ and $\mathbf{X}$ are the past τ'th non-overlapping weekly return and vector of control variables, respectively defined under Eq. (9). Subscripts s and ns index salient and nonsalient returns, respectively. Hence, $\mathbb{1}_s$ is an indicator function that equals one if stock k's return is salient. Conversely, $\mathbb{1}_{ns}$ is an indicator function that equals one if stock k's return is nonsalient. I use two proxies for the salience of a stock on day t. First, stocks in the S&P 500 index as of $t-2$ are considered to have salient returns on day t. Alternatively, a stock is defined to have salient returns on day t if the number of cash flow news published about the stock over the past week ending $t-2$ is above the monthly median.

In Eq. (10), $\lambda_{2,j}$ governs the importance of recent returns relative to older returns in shaping beliefs about type-j stocks. $\lambda_{2,j}$ closer to zero suggests that recent returns primarily influence beliefs compared to older returns. In contrast, $\lambda_{2,j}$ close to one indicates that SMAs give about the same weight to older and recent returns. On the other hand, $\lambda_{1,j}$ captures the overall extent to which SMAs' beliefs respond to past returns of type-j stocks. Da et al. (2021) show that $\lambda_{1,j}(1-\lambda_{2,j})$ is an appropriate measure that summarizes the degree of extrapolation, with higher values indicating greater extrapolation.

Table 10 shows the results for the nonlinear least squares estimation. Columns (1) and (2) show the estimated coefficients for stocks that do not belong to the S&P500 index and those that do, respectively. Similarly, Columns (3) and (4) show coefficients for stocks with below and above median news coverage, respectively. For both measures of salience, the estimated coefficients are significantly different from zero, and λ_1 (λ_2) is higher (lower) for stocks with more salient returns (Columns (2) and (4)) compared to low salience stocks (Columns (1) and (3)), suggesting stronger extrapolation for stocks with salient returns. For example, focusing on Columns (3)

	(1)	(2)	(3)	(4)
	In S&P 500 Index		News Coverage	
	No	Yes	Low	High
λ_0	0.323		0.326	
	(0.012)		(0.012)	
λ_1	0.895	1.407	0.863	0.961
	(0.108)	(0.453)	(0.136)	(0.153)
λ_2	0.798	0.618	0.953	0.636
	(0.039)	(0.141)	(0.044)	(0.067)
Obs.	142,021	142,021	142,021	142,021
$\lambda_1(1-\lambda_2)$	0.181	0.538	0.040	0.350

Table 10: Return Extrapolation and Salience: Nonlinear Least Squares. The table reports the results for SMAs' return extrapolation based on nonlinear least squares Eq. (10). Columns (1) and (2) present results for the case where a stock's return is considered salient if the stock belongs (Yes) to the S&P 500 index as of $t-2$, where t is the belief publication day, and nonsalient otherwise (No). Columns (3) and (4) present results for the case where a stock's return is considered salient if the number of cash flow news published about the stock over the past week is above (High) the monthly median, and nonsalient otherwise (Low). Regressions include the following control variables: lagged belief, cash flow news sentiment, and professional analysts' consensus forecast of quarterly earnings per share. Standard errors are shown in parentheses.

and (4), where salience is proxied by news coverage, $\lambda_2 = 0.64$ in Column (4) indicates that the most recent one-week return has about 20 times the influence of the 8th-week return in shaping SMAs' beliefs about salient stocks. Conversely, for nonsalient stocks, $\lambda_2 = 0.95$ in Column (3) indicates that the most recent one-week return only has about twice the influence of the 8th-week return. The last row of Table 10 summarizes the degree of extrapolation, indicating that overall extrapolation is much stronger for stocks with salient returns.

To summarize, the analyses in this section indicate that SMAs extrapolate from past returns in belief formation, and extrapolation tends to be more pronounced for larger stocks and stocks with more news coverage, corroborating existing studies (e.g., Alok et al. 2020; Tversky and Kahneman 1973) which suggest that salience fuels behavioural biases. However, notwithstanding the extrapolation, previous sections show that SMAs' beliefs are informative about future stock returns and earnings surprises, implying that the beliefs are unlikely to be systematically biased. Hence, these results complement that of Da et al. (2021), who document systematically biased beliefs for return extrapolators. A potential explanation for the divergent findings is that unlike the individuals in the sample of Da et al. (2021), SMAs in this paper tend to be more sophisticated than the average retail investor and are therefore better equipped to combine information from past returns with other fundamental-relevant information when forming expectations. In fact, the article that accompanies each SMA belief statement precisely aims to provide fundamental information supporting the stated belief.

7 Robustness

Informativeness of Aggregated SMA beliefs. Although several model validation exercises demonstrate that the trained ML model classifies beliefs with a high degree of accuracy out-of-sample and has intuitive feature weights, some doubts might remain regarding whether the ML model somehow drives the results in this paper. To address this concern, I examine the ability of SMA beliefs to predict future abnormal returns using only the explicitly stated beliefs, i.e., ignoring beliefs extracted using ML. Although this leads to smaller sample size, Table IA1 in the Internet Appendix shows that aggregate SMA beliefs' predictability of abnormal returns holds. In particular, the coefficient of $AB_{k,t}$ is significant for all horizons and is comparable to the results in Table 2 of the main analysis.

I conduct additional robustness tests using only the stated beliefs but now averaging the beliefs over the past two weeks and one month, respectively. Precisely, for each firm-date observation of $AB_{k,t}$, I obtain $\overline{AB}_{k,t}$ by averaging $AB_{k,t}$ over the window $t-h$ to t such that $AB_{k,t} = \overline{AB}_{k,t}$ if there are no SMA beliefs about stock k over the past h days. Panel A and Panel B of Table IA2 in the Internet Appendix show the results for $h = 14$ and 30 days, respectively. Averaging SMA beliefs over the past reduces the significance of the abnormal return predictability for the CAPM benchmark and future 63-day horizon. Nonetheless, overall, the results are qualitatively comparable to those in Table 2. Conducting similar exercises with the stated and extracted beliefs combined produces similar insights, as Table IA3 in the Internet Appendix shows.

I further verify that the informativeness of the stock-level aggregate SMA beliefs holds in different subsamples. Panel A of Table IA4 in the Internet Appendix shows results for re-estimating regression (2) using only data for the first half of the sample 2005 – 2012, while Panel B shows results for the second half 2013 – 2019. Overall, aggregate SMA beliefs predict future abnormal returns in both subsamples but with less statistical significance for the 63-day horizon.

Alternative Mixture Model Setup and Data Requirement. The analysis of the cross-sectional distribution of SMAs' ability to form correct beliefs was based on a two-component mixture model using only SMAs that have at least five belief statements over the entire sample. To ensure that this specific setup does not drive the results, Table A3 in the Appendix reports the results for the constrained two-component mixture model, where the true abnormal return of SMAs in the low-skill group (Component 0) is restricted to zero, such that the variation in their estimated abnormal return is purely attributable to luck. We still see that most SMAs (roughly 73%) belong to the low-skill group. Similarly, Table A4 considers an alternative setup with three components in the mixture model. Again, most SMAs (66%) belong to Component 0 (lowest skill group), while Component 1 (medium) and Component 2 (highest skill group) comprise 29% and 5% of the SMAs, respectively. Overall, the three-component mixture model indicates that 60% of the SMAs have positive average abnormal returns, comparable to the 56% reported in the main analysis using the two-component mixture model. Nevertheless, the BIC model selection criteria suggest that the two-component mixture model best fits the data.

Table A5 further reports the results for the two-component mixture model using data for only SMAs with at least 10 (instead of five) belief statements over the sample period. Although this more stringent data requirement reduces the number of SMAs in the cross-section by about half, we still find a qualitatively similar result as in the main analysis: Most SMAs are low skill, with much lower abnormal returns.

Finally, Tables IA5 and IA6 of the Internet Appendix show the counterparts of Tables 5 and A4, respectively, where I now use CAPM instead of the three-factor model as the benchmark for SMAs' estimated abnormal returns. The fraction of high-skill SMAs and the proportion of SMAs with positive abnormal returns are very similar in both settings, indicating that the results are robust to alternative benchmarks. Overall, the robustness analysis shows that alternative modelling and estimation choices do not significantly influence earlier results on the cross-sectional distribution of SMAs' ability to form correct beliefs.

Transaction-based Calendar-time Portfolios Excluding Small Stocks. To demonstrate that small stocks do not drive the profitability of portfolios formed on SMAs' beliefs, I redo the portfolio analysis excluding penny stocks (price less than \$5) and microcap stocks (market capitalization less than the 2nd NYSE decile), respectively. Precisely, a stock is not included in a portfolio on day t if on day $t-1$, when beliefs are expressed, its price is less than \$5 or its market capitalization is less than the 2nd NYSE decile. Once a unit of a stock is in a portfolio, it is held until the end of the holding period, regardless of whether the stock failed to meet the inclusion criteria on certain dates during the holding period. Table A6 shows the results for the exclusion of Penny stocks, while Table A7 shows the results for the exclusion of microcap stocks. In both cases, the Bullish portfolio has a positive alpha ranging between 0.31% and 0.50% per month across different factor benchmarks and holding periods, with lower statistical significance for the longer holding periods. Conversely, the Bearish portfolio generally has negative alphas, corroborating the results in the main analysis.

Herding Test with Consensus based on Alternative Window. The main analysis on SMAs' herding is based on consensus computed over the past six-month period $t-180$ to $t-2$. Tables A8 and A9 of the Appendix show the robustness results, where the consensus is instead computed over the past three-month period $t-90$ to $t-2$, indicating that the window used in computing the consensus does not significantly influence the results. In particular, the estimated herding coefficient is around 0.26, as in the main analysis, with the results on conditional herding also qualitatively similar to the main analysis.

Extrapolation in Subsamples. Finally, Table A10 of the Appendix shows that the analysis on SMAs' return extrapolation holds across subsamples. Columns (3) and (4) show the results for the first half (2005 – 2012) and the second half (2013 – 2019) of the sample, respectively. In both cases, the two most recent one-week returns have the strongest and most statistically significant influence on SMAs' beliefs, with the influence of older one-week returns mostly statistically insignificant.

8 Conclusion

This paper uses natural language processing and machine learning to infer nonprofessional social media investment analysts' (SMAs) beliefs about a large cross-section of stocks from opinions expressed on the investment social media platform, Seeking Alpha. The paper then studies the informativeness of SMA beliefs about individual stock returns and the cross-sectional distribution of SMAs' ability to form correct beliefs about stock returns. Finally, the paper tests two theory-inspired channels through which SMAs form beliefs: herding and return extrapolation.

The analyses show that, on average, SMA beliefs contain value-relevant information. However, there exists substantial heterogeneity in SMAs' ability to form beliefs that yield investment value. For example, about half of SMAs form beliefs that generate positive abnormal returns, while high-skilled SMAs (only 13%) state beliefs that yield a sizeable one-week three-factor abnormal return of 61 bps (31% annualized). The analysis of portfolios formed on SMA beliefs indicates that investors can profitably trade on SMAs' belief statements. The 63-day holding period portfolio of Bullish (Bearish) beliefs yields an annualized three-factor alpha of 3.6% (-2.6%). These results suggest that SMAs as a group create value for investors, although skill is relatively limited at the individual SMA level. However, specific SMA characteristics can help investors identify skilled SMAs on social media.

The paper then documents that SMAs herd on the consensus and extrapolate from past returns when forming beliefs about future stock returns, with extrapolation stronger for stocks with salient returns. Nevertheless, the behavioural biases in SMAs' belief formation do not result in systematically wrong beliefs.

In summary, SMAs produce value-relevant information that can benefit investors, even though their belief formation process is not entirely consistent with rational models. In light of ongoing concerns surrounding social media's growing influence over financial markets, this paper suggests that some aspects of social media serve as veritable sources of information that inform beliefs which can improve investment decisions, with important welfare implications.

References

Alok, S., Kumar, N., and Wermers, R. (2020). Do fund managers misestimate climatic disaster risk. *The Review of Financial Studies*, 33(3):1146–1183.

Andrikogiannopoulou, A. and Papakonstantinou, F. (2019). Reassessing false discoveries in mutual fund performance: Skill, luck, or lack of power? *The Journal of Finance*, 74(5):2667–2688.

Antweiler, W. and Frank, M. Z. (2004). Is all that talk just noise? the information content of internet stock message boards. *The Journal of Finance*, 59(3):1259–1294.

Avery, C. N., Chevalier, J. A., and Zeckhauser, R. J. (2016). The "caps" prediction system and stock market returns. *Review of Finance*, 20(4):1363–1381.

Bajgrowicz, P. and Scaillet, O. (2012). Technical trading revisited: False discoveries, persistence tests, and transaction costs. *Journal of Financial Economics*, 106(3):473–491.

Baltatescu, S. (2014). *Psychological Distance*, pages 5145–5146. Springer Netherlands, Dordrecht.

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817.

Barber, B. M., Lehavy, R., McNichols, M., and Trueman, B. (2006). Buys, holds, and sells: The distribution of investment banks' stock ratings and the implications for the profitability of analysts' recommendations. *Journal of accounting and Economics*, 41(1-2):87–117.

Barber, B. M. and Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The journal of Finance*, 55(2):773–806.

Barber, B. M. and Odean, T. (2013). The behavior of individual investors. In *Handbook of the Economics of Finance*, volume 2, pages 1533–1570. Elsevier.

Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2015). X-capm: An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115(1):1–24.

Barras, L., Scaillet, O., and Wermers, R. (2010). False discoveries in mutual fund performance: Measuring luck in estimated alphas. *The Journal of Finance*, 65(1):179–216.

Bhamra, H., Uppal, R., and Walden, J. (2021). Psychological distance and deviations from rational expectations. *Working Paper*.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026.

Blume, M. E. and Stambaugh, R. F. (1983). Biases in computed returns: An application to the size effect. *Journal of Financial Economics*, 12(3):387–404.

Boehmer, E., Jones, C. M., Zhang, X., and Zhang, X. (2020). Tracking retail investor activity. *Journal of Finance, Forthcoming.*

Brunnermeier, M., Farhi, E., Koijen, R. S. J., Krishnamurthy, A., Ludvigson, S. C., Lustig, H., Nagel, S., and Piazzesi, M. (2021). Review Article: Perspectives on the Future of Asset Pricing. *The Review of Financial Studies*.

Campbell, J. L., DeAngelis, M. D., and Moon, J. R. (2019). Skin in the game: Personal stock holdings and investors' response to stock analysis on social media. *Review of Accounting Studies*, 24(3):731–779.

Chen, H., De, P., Hu, Y. J., and Hwang, B.-H. (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. *The Review of Financial Studies*, 27(5):1367–1403.

Chen, M. A., Wu, Q., and Yang, B. (2019). How valuable is fintech innovation? *The Review of Financial Studies*, 32(5):2062–2106.

Chen, Y., Cliff, M. T., and Zhao, H. (2017). Hedge funds: The good, the bad, and the lucky. *Journal of Financial and Quantitative Analysis (JFQA)*, 52(3):1081–1109.

Choi, J. J. and Robertson, A. Z. (2020). What matters to individual investors? evidence from the horse's mouth. *The Journal of Finance*.

Cookson, J. A., Engelberg, J., and Mullins, W. (2021). Echo chambers. *Available at SSRN 3603107*.

Cookson, J. A., Engelberg, J. E., and Mullins, W. (2020). Does partisanship shape investor beliefs? evidence from the covid-19 pandemic. *The Review of Asset Pricing Studies*, 10(4):863–893.

Cookson, J. A. and Niessner, M. (2020). Why don't we agree? evidence from a social network of investors. *The Journal of Finance*, 75(1):173–228.

Crane, A. and Crotty, K. (2020). How skilled are security analysts? *The Journal of Finance*, 75(3):1629–1675.

Cujean, J. and Hasler, M. (2017). Why does return predictability concentrate in bad times? *The Journal of Finance*, 72(6):2717–2758.

Da, Z., Huang, X., and Jin, L. J. (2021). Extrapolative beliefs in the cross-section: What can we learn from the crowds? *Journal of Financial Economics*, 140(1):175–196.

Daniel, K., Grinblatt, M., Titman, S., and Wermers, R. (1997). Measuring mutual fund performance with characteristic-based benchmarks. *The Journal of Finance*, 52(3):1035–1058.

Das, S. R. and Chen, M. Y. (2007). Yahoo! for amazon: Sentiment extraction from small talk on the web. *Management Science*, 53(9):1375–1388.

Devenow, A. and Welch, I. (1996). Rational herding in financial economics. *European Economic Review*, 40(3-5):603–615.

Drake, M., Moon, J., Twedt, B., and Warren, J. (2019). Are social media analysts disrupting the information content of sell-side analysts' reports. *Available at SSRN 3456801*.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3 – 56.

Farrell, M., Green, T. C., Jame, R., and Markov, S. (2020). The democratization of investment research and the informativeness of retail investor trading. *Available at SSRN 3222841*.

Frankel, R., Jennings, J., and Lee, J. (2016). Using unstructured and qualitative disclosures to explain accruals. *Journal of Accounting and Economics*, 62(2-3):209–227.

Gentzkow, M. and Shapiro, J. M. (2010). What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71.

Giglio, S., Maggiori, M., Stroebel, J., and Utkus, S. (2019). Five facts about beliefs and portfolios. CESifo Working Paper, No. 7666.

Gomez, E., Heflin, F., Moon, J., and Warren, J. (2020). Can financial analysis on social media help level the playing field among investors? evidence from seeking alpha. *Georgia Tech Scheller College of Business Research Paper 18-45*.

Greenwood, R. and Shleifer, A. (2014). Expectations of Returns and Expected Returns. *The Review of Financial Studies*, 27(3):714–746.

Gu, S., Kelly, B., and Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5):2223–2273.

Harvey, C. R. and Liu, Y. (2018). Detecting repeatable performance. *The Review of Financial Studies*, 31(7):2499–2552.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). Unsupervised learning. In *The Elements of Statistical Learning*, pages 485–585. Springer, New York, NY.

Hirshleifer, D., Lim, S. S., and Teoh, S. H. (2011). Limited investor attention and stock market misreactions to accounting information. *The Review of Asset Pricing Studies*, 1(1):35–73.

Ivković, Z. and Weisbenner, S. (2005). Local does as local is: Information content of the geography of individual investors' common stock investments. *The Journal of Finance*, 60(1):267–306.

Iyyer, M., Enns, P., Boyd-Graber, J., and Resnik, P. (2014). Political ideology detection using recursive neural networks. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1113–1122.

Jeng, L. A., Metrick, A., and Zeckhauser, R. (2003). Estimating the returns to insider trading: A performance-evaluation perspective. *Review of Economics and Statistics*, 85(2):453–471.

Kaniel, R., Liu, S., Saar, G., and Titman, S. (2012). Individual investor trading and return patterns around earnings announcements. *The Journal of Finance*, 67(2):639–680.

Ke, Z. T., Kelly, B. T., and Xiu, D. (2019). Predicting returns with text data. Technical report, National Bureau of Economic Research.

Kelley, E. K. and Tetlock, P. C. (2013). How wise are crowds? insights from retail orders and stock returns. *The Journal of Finance*, 68(3):1229–1265.

Kuchler, T. and Zafar, B. (2019). Personal experiences and expectations about aggregate outcomes. *The Journal of Finance*, 74(5):2491–2542.

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

Malmendier, U. and Nagel, S. (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?*. *The Quarterly Journal of Economics*, 126(1):373–416.

Manela, A. and Moreira, A. (2017). News implied volatility and disaster concerns. *Journal of Financial Economics*, 123(1):137 – 162.

Newey, W. K. and West, K. D. (1987). A simple, positive-semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.

Pedersen, L. H. (2021). Game on: Social networks and markets. *Available at SSRN 3794616*.

Scharfstein, D. S. and Stein, J. C. (1990). Herd behavior and investment. *The American Economic Review*, pages 465–479.

Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464.

Sclove, S. L. (1987). Application of model-selection criteria to some problems in multivariate analysis. *Psychometrika*, 52(3):333–343.

Seasholes, M. S. and Zhu, N. (2010). Individual investors and local bias. *The Journal of Finance*, 65(5):1987–2010.

Seeking Alpha (2006). About seeking alpha. https://seekingalpha.com/amp/article/5824-about-seeking-alpha. Accessed: 2020-09-30.

Simonsohn, U. and Ariely, D. (2008). When rational sellers face nonrational buyers: evidence from herding on ebay. *Management Science*, 54(9):1624–1637.

Tversky, A. and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, 5(2):207–232.

Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.

Van Nieuwerburgh, S. and Veldkamp, L. (2010). Information acquisition and under-diversification. *The Review of Economic Studies*, 77(2):779–805.

Welch, I. (2000). Herding among security analysts. *Journal of Financial Economics*, 58(3):369–396.

A Appendix

A.1 Variable Description

Stock-level and market variables

- $ABR_{k,t+1 \to t+1+h}$: stock k abnormal return relative to either the CAPM, Fama and French (1993) three-factor model (FF3), or the Daniel et al. (1997) size/book-to-market/momentum characteristics-based benchmark (SBM), where t is the SMA belief statement day. For the CAPM and three-factor benchmarks, betas for each stock are estimated using daily data over the trading-day window t-272 to t-21.

- $Volatility_{k,t}$: the sum of squared daily returns in the calendar month before day t.

- $AB_{k,t}$: stock-level aggregate SMA belief about stock k on day t computed as the number of Bullish beliefs ($NBullish_{k,t}$) minus number of Bearish beliefs ($NBearish_{k,t}$) divided by the total number of beliefs ($NBelief_{k,t}$): $(NBullish_{k,t} - NBearish_{k,t}) / NBelief_{k,t}$.

- $Downgrade_{k,t}$: number of professional stock analysts that downgraded stock k on day t. If there are no downgrades, the value is set to zero.

- $Upgrade_{k,t}$: number of professional stock analysts that upgraded stock k on day t. If there are no upgrades, the value is set to zero.

- $SUE_{k,t}$: earnings surprise for firm k computed as the difference between Actual EPS and the average forecasts across analysts (consensus estimate) divided by the stock price at the end of the last quarter. To avoid stale forecasts, I use only forecasts published within the 30 days ending one day before the earnings announcement day t.

- $Forecast\ Dispersion_{k,t}$: Dispersion of analyst forecasts of firm k's earning per share computed as the standard deviation of the analyst forecasts over the 30-day window ending one day before earnings announcement day t scaled by the stock price at the end of the previous quarter.

- $New\ Sentiment_{k,t}$: average news sentiment across a comprehensive set of cash flow relevant news events about stock k on day t computed as the Event Sentiment Score (ESS) from RavenPack News Analytics divided by 100. ESS ranges between 0 and 100, where 50 indicates neutral sentiment, values above 50 indicate positive sentiment and values below 50 indicate negative sentiment. I use only news events with a relevance score of at least 75 to focus on news events that mostly relate to a specific firm. If there are no news events for a stock on a given day, $News\ Sentiment_{k,t}$ is set to its neutral value of 0.5. Table A1 lists the cash flow relevant news categories.

- Size: is the log of stock k's market capitalization as of June of the previous calendar year.

- Book-to-Market: is the log of stock k's book-to-market ratio as of the most recent fiscal year-end.

- Equal-weighted consensus: For an SMA's belief revision about a stock on day t, the consensus is computed as the average of other SMAs' beliefs on the same stock over the past 6-month period ending $t-2$. The consensus is estimated if there are at least two belief statements by other SMAs over the 6-month period. Bearish, Neutral and Bullish beliefs are assigned the values 1, 2, and 3, respectively.

- Characteristic-weighted consensus: for an SMA's belief revision about a stock on day t, the characteristic-weighted consensus is computed as the weighted average of other SMAs' beliefs on the same stock over the past 6-month period ending $t-2$, where the weights correspond to each SMA's value for a given characteristic x computed as of last calendar month using data over the past one year. Missing values for each characteristic are replaced with the median value, and the weights are normalized to sum to 1. The consensus is estimated if there are at least two belief statements by other SMAs over the past 6-month period. Bearish, Neutral and Bullish beliefs are assigned the values 1, 2, and 3, respectively.
- Consensus Optimism (CO): a stock's equal-weighted consensus C on day t minus 2, i.e., $CO = C - 2$. Since beliefs are labeled as 3 (Bullish), 2 (Neutral) and 1 (Bearish), $CO > 0$ implies an optimistic consensus, while ($CO < 0$) implies a pessimistic consensus.
- Consensus Correctness (CC): a stock's consensus optimism (CO) on day t times its future return, i.e., $CC = CO \times Ret(h)$, where $Ret(h)$ is the future horizon h return of the stock starting $t+1$ for belief revision on day t. CC is positive when the consensus is correct ex-post, that is an optimistic (pessimistic) consensus is followed by a positive (negative) future stock return.
- Recession: indicator variable for NBER recession dates.
- High Volatility: indicator variable that equals 1 if the volatility of the stock market return exceeds its sample median and 0 otherwise. Market volatility is computed as of the last calendar month using one year of daily data. The market return is the market factor plus risk-free rate from Kenneth French's website.

SMA-level variables

- $\overline{ABR_i}$: SMA i's estimated abnormal return computed as the average of signed abnormal returns across all belief statements by SMA i. To sign the abnormal returns, I premultiply $ABR_{k,t+1\to t+1+h}$ by +1 following a bullish belief about stock k and by -1 following a bearish belief.
- s_i: the standard error of SMA i's estimated abnormal return ($\overline{ABR_i}$) calculated by clustering on the belief statement day, to account for correlation across belief statements on the same day, and stock level, to account for correlation in belief statements on the same stock.
- Industry specialization: proxied by the one divided by the number of unique SIC industries across which an SMA expressed beliefs over the sample period.
- Firm specialization: proxied by one divided by the average number of unique firms an SMA expressed beliefs about in a given year.
- Workload: proxied by the average number of belief statements published per year.
- Skin-in-game: proxied by the fraction of time an SMA discloses an investment position in the stock about which he/she expresses a belief. For each SMA's publication about a stock, I create a dummy variable that equals 1 if the SMA discloses a "long" or "short" position in the stock and 0 if there is no disclosure or "no position" is disclosed. Finally, the indicator variable is averaged over the sample period for each SMA.

- Effort: proxied by the average number of words in opinion articles corresponding to each belief statement.
- Disagreement: The average of the absolute difference between the fraction of negative words in an SMA's opinion article and the average fraction of negative words in the comments posted within two days of article publication. The negative word list is from the Loughran and McDonald (2011) dictionary.
- Popularity: proxied by the average number of comments on each belief statement by an SMA within two days of publication.
- SA Experience: captures an SMA's experience on the Seeking Alpha (SA) platform, computed as the number of years between an SMA's first and last publication on Seeking Alpha.

A.2 Description of the Support Vector Classifier ML Algorithm

To get a glimpse of how linear SVC works, consider a simple binary classification task. The SVC algorithm solves the following optimization problem:[36]

$$\underset{w_0,\mathbf{w}}{min} \sum_{n=1}^{N} [1 - y_n(w_0 + \mathbf{w} \cdot \mathbf{x}_n)]^+ \ + \ c \cdot (\mathbf{w} \cdot \mathbf{w}), \tag{A1}$$

where y is the vector of binary labels consisting of entries -1 and +1 used in place of the actual class labels – for instance, -1 could denote "Bearish Belief" while +1 denotes "Bullish Belief"; w_0 and $\mathbf{w}$ are weights to be optimized on; c is a scalar regularization hyperparameter; $[1 - y_n(w_0 + \mathbf{w} \cdot \mathbf{x}_n)]^+$ is the "hinge" loss function, with the superscript "+" indicating that only positive parts should be considered; and N is the number of observations. As shown in Hastie et al. (2009), for a given regularization hyperparameter c, the solution to (A1) is a weighted average of the regressors:

$$\hat{\mathbf{w}} = \sum_{n=1}^{N} \hat{\alpha}_n y_n \mathbf{x}_n, \tag{A2}$$

where only some of the N observation weights $\hat{\alpha}_n$ are nonzero. These observations with nonzero $\hat{\alpha}_n$, called the support vectors, are the ones relatively close to the separating hyperplane, i.e.,

[36] Since a full treatment of the SVC algorithm is beyond the scope of this paper, I present only an illustration. See Hastie et al. (2009) for a formal exposition.

points that are most difficult to classify. With $\hat{\mathbf{w}}$, SVC predicts the classes using the decision function:

$$\hat{G}(\mathbf{x}_n) = sign\left(\hat{w}_0 + \hat{\mathbf{w}}\cdot\mathbf{x}_n\right). \tag{A3}$$

SVC, therefore, selects a relatively small number of observations for identifying the coefficients $\hat{\mathbf{w}}$. This, in turn, implies that instead of estimating M coefficients for the large feature space, only a relatively small set of parameters (selecting the support vectors and $\hat{\alpha}_n$) is estimated, yielding nonzero coefficients for only the important features.

While the preceding illustration is for binary classification, it extends to multiclass classification as in this paper. In the multiclass context, a *one-vs-rest* binary SVC is trained for each class to separate the class from all other classes, resulting in as many binary models as there are classes. To make a prediction, all binary classifiers are run on the data and the class label of the classifier that has the highest score (by evaluating $\hat{w}_0 + \hat{\mathbf{w}}\cdot\mathbf{x}_n$) is returned as the prediction.

A.3 Description of Herding Test

This section provides additional details on the herding test used in the paper based on Welch (2000). Eq. (8) specifies how the transition probability depends on θ and a target, namely the consensus (C). The estimate for θ is obtained by maximizing the log-likelihood function given the observed transitions $i_n \to j_n$, each with its own target C_n, over all observations $n \in [1, N]$:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \sum_n \log\left[p_{i,j}(\theta, C_n)|_{\{i_n, j_n\}}\right], \tag{A4}$$

where $p_{i,j}(\theta, C_n)|_{\{i_n,j_n\}}$ is the transition probability in Eq. (8) evaluated at each realized n transition and their associated target C_n in the data. Therefore, any chosen θ parameter implies for each observation n a probability vector p_i. As Welch (2000) notes, under the assumption that draws are independent, the likelihood function is the probability of empirically observing the full sequence of transitions.

Statistical inference on the estimator $\hat{\theta}$ is based on the likelihood ratio test. The likelihood ratio statistic is the ratio of the probability of the data for a given constant transition probability versus a transition probability that varies with C according to θ:

$$2\left\{\sum_n \log\left[p_{i,j}(\theta, C_n)|_{\{i_n,j_n\}}\right] - \sum_n \log\left[p_{i,j}(0)|_{\{i_n,j_n\}}\right]\right\} \quad \sim \quad \chi_1^2. \tag{A5}$$

Since $p_{i,j}(0)$ is not observable, it is assumed that the $\mathbf{P}(0)$ matrix under the null is the empirically observed transition matrix. This assumption is conservative in that it is correct under the null and biases the test against the herding hypothesis. For the setup where θ depends on some variable y: $\theta(y) = \theta_0 + \theta_1 y$, the likelihood ratio test is performed by comparing the restricted model (with $\theta = \theta_0$) against the unrestricted model (with $\theta(y) = \theta_0 + \theta_1 y$).

For the asymptotic properties of the described maximum-likelihood estimation, see Welch (2000) who also documents excellent small sample properties of the maximum-likelihood ratio.

A.4 Additional Tables and Figures

Author: Dana Blankenhorn
Timestamp: 2012-12-18T15:47:05Z
Ticker: MSFT
Belief statement: Bullish
Article URL: https://seekingalpha.com/article/1070611-hope-for-microsoft-cloud

Microsoft (NASDAQ:MSFT) has become the company tech investors love to hate. The company's P/E of 14.69, which is about average for an S&P 500 stock, masks the extraordinary charges in the second quarter, when the company wrote off its aQuantive ad network. Had the company even had an average quarter for June, you would have trailing years' earnings of $2.50/share, a P/E lower than that of Ford (NYSE:F). Plus, there's over $66 billion in cash on the balance sheet – take that out and the price is a snip. But investors aren't buying that story anymore. The stock is back to the levels of the first of the year, before Windows 8 was rolled out. Since its peak in April it's down 16%. And did I mention the 3.38% yield? That said, the reason investors are fleeing Microsoft makes sense. They see the ads for Windows 8, they've gone into the stores, and they've seen the Microsoft products languishing on the shelves. They've read the tweets and seen the reviews. They think Windows 8 does something that rhymes with trucks, and that it's not rolling out as expected. I could argue that most software is seen in a negative light when it first comes out, that even Apple (NASDAQ:AAPL) has its share of negative reviews, and that a new user interface always takes some time to learn and get used to. But instead I'm going to talk about the cloud. When I write about cloud, I seldom mention Microsoft Azure. But Azure is a pretty good cloud, with decent price performance. And according to a recent survey by Forrester Research, which tracks corporate computing professionals, it's getting strong reviews there. As per the survey, Azure is currently doing as well as Google (NASDAQ:GOOG) cloud services, and a full 20% expect their usage of it to grow over the next year. The reason: It's easy to set up and easy to use. If you know Windows, you're halfway there – you can bring all your existing tools and skills to the party. This is important stuff, because the next big step for the cloud market is the move from cloud infrastructures to cloud platforms. If Microsoft can make the jump, with anything like its current market share as a cloud platform and if its languages and other tools can be seen as the easy way to build cloud applications, that is a huge leg up in the market. And hidden within Windows 8 are all the tools you need – both you and your employees – to build and deploy cloud applications on Azure. So you've got a dirt cheap stock with a leg up on the cloud market of 2013. Is that worth an implied P/E of 7 when you also get $66 billion in cash? At this point, Microsoft is as discounted as it is going to get. It's beyond cheap here.

Disclosure: I am long MSFT, AAPL, GOOG. I wrote this article myself, and it expresses my own opinions. I am not receiving compensation for it (other than from Seeking Alpha). I have no business relationship with any company whose stock is mentioned in this article.

Figure A1: Sample SMA Article with Bullish Belief. The figure shows an opinion article published in Seeking Alpha where the author explicitly states a bullish belief about a stock.

Author: Benjamin Clark
Timestamp: 2014-11-03T06:02:53Z
Ticker: MMM
Belief statement: Bearish
Article URL: https://seekingalpha.com/article/2629295-moderngraham\-quarterly-valuation-of-3m-company

Summary: MMM is suitable for Enterprising Investors following the ModernGraham approach. According to the ModernGraham valuation model, the company is overvalued at the present time. The market is implying 6.96% earnings growth over the next 7-10 years, a rate of growth which is unsupported by the company's recent results.
3M Company (NYSE:MMM) has seen quite a run-up in price over the last five years, and for many investors that alone presents a reason to turn away. However, Benjamin Graham, the father of value investing, taught that looking at the price cannot be the sole factor in investment decisions as the most important aspect to consider is whether the company is trading at a discount relative to its intrinsic value. It is through a thorough fundamental analysis that the investor is able to make a determination about a potential investment's merits. Here is a look at how 3M Company fares in the ModernGraham valuation model. The model is inspired by the teachings of Benjamin Graham and considers numerous metrics intended to help the investor reduce risk levels. The first part of the analysis is to determine whether the company is suitable for the very conservative Defensive Investor or the less conservative Enterprising Investor, who is willing to spend a greater amount of time conducting further research. In addition, Graham strongly suggested that investors avoid speculation in order to remove the subjective elements of emotion. This is best achieved by utilizing a systematic approach to analysis that will provide investors with a sense of how a specific company compares to another company. By using the ModernGraham method one can review a company's historical accomplishments and determine an intrinsic value that can be compared across industries. 3M Company is not a great opportunity for Defensive Investors, as the company has a low current ratio and high PEmg and PB ratios. However, the less conservative Enterprising Investor has no initial concerns and is willing to proceed to the next part of the analysis, which is a determination of the company's intrinsic value. When estimating the intrinsic value, it is critical to consider the company's historical earnings results in combination with a review of the market's implied estimate for further growth. In this case, the company has grown its EPSmg (normalized earnings) from $5.14 in 2010 to an estimated $6.71 for 2014. While this is a strong level of demonstrated growth, it does not quite support the market's implied estimate for earnings growth of 6.96% over the next 7-10 years. In order to reach that growth rate, the company would need to achieve higher growth than it has in the recent past. The ModernGraham valuation model therefore returns an estimate of intrinsic value falling below the current price, indicating the company is overvalued at the present time. Be sure to check out previous ModernGraham valuations of 3M Company for more perspective!

Disclosure: I/we have no positions in any stocks mentioned, but may initiate a short position in NKE over the next 72 hours. I wrote this article myself, and it expresses my own opinions. I am not receiving compensation for it (other than from Seeking Alpha). I have no business relationship with any company whose stock is mentioned in this article.

Figure A2: Sample SMA Article with Bearish Belief. The figure shows an opinion article published in Seeking Alpha where the author explicitly states a bearish belief about a stock.

acquisition, acquisition-regulatory-approval, acquisition-regulatory-scrutiny, analyst-ratings-change, antitrust-investigation, antitrust-settlement, antitrust-suit, bankruptcy, business-combination, business-contract, buybacks, conference, conference-call, congressional-testimony, corruption, credit-rating-change, credit-rating-outlook, credit-rating-watch, debt-restructuring, discrimination, dividend, dividend-guidance, earnings, earnings-estimate, earnings-guidance, earnings-per-share, earnings-per-share-guidance, earnings-revision, executive-appointment, executive-compensation, executive-firing, executive-resignation, executive-salary, executive-scandal, executive-shares-options, expenses, expenses-guidance, fraud, going-private, government-contract, index-delisting, index-listing, insider-buy, insider-sell, insider-trading-lawsuit, investment, layoffs, legal-issues, market-entry, market-guidance, merger, merger-regulatory-approval, merger-regulatory-scrutiny, ownership, partnership, patent, patent-infringement, product-discontinued, product-outage, product-price, product-recall, product-release, product-resumed, production-outlook, project-abandoned, protest, public-offering, regulatory-investigation, regulatory-product-application, regulatory-product-approval, regulatory-product-review, regulatory-product-warning, revenue, revenue-guidance, revenue-volume, rights-issue, same-store-sales, sell-registration, settlement, spin-off, stock-splits, tax-evasion, trading, unit-acquisition, unit-acquisition-regulatory-approval, workers-strike.

Table A1: Cashflow-Relevant News Event Categories The table reports news event categories in the RavenPack database deemed relevant for firm cash flow. When aggregating sentiment across cash-flow news for a given firm, I focus only on these event categories.

	CAPM		FF3		SBM	
	(1)	(2)	(3)	(4)	(5)	(6)
$AB_{k,t}$	0.0036	0.0034	0.0028	0.0026	0.0033	0.0031
	(2.69)	(2.55)	(2.31)	(2.15)	(2.71)	(2.58)
$ABR_{k,t}$		0.0019		0.0021		0.0015
		(1.77)		(2.07)		(1.52)
$ABR_{k,t-1}$		-0.0018		-0.0018		-0.0020
		(-1.56)		(-1.56)		(-1.76)
$ABR_{k,t-2}$		-0.0000		0.0001		0.0005
		(-0.04)		(0.06)		(0.43)
$ABR_{k,t-63\to t-3}$		0.0004		0.0017		-0.0010
		(0.23)		(1.00)		(-0.56)
$Volatility_{k,t}$		-0.0032		-0.0027		-0.0033
		(-1.55)		(-1.38)		(-1.73)
$Upgrade_{k,t}$		0.0014		0.0014		0.0015
		(3.12)		(3.18)		(3.29)
$Downgrade_{k,t}$		-0.0003		-0.0004		-0.0004
		(-0.56)		(-0.69)		(-0.74)
$News\ Sentiment_{k,t}$		0.0011		0.0009		0.0010
		(1.72)		(1.56)		(1.69)
Obs.	192,650	192,650	192,650	192,650	192,650	192,650
R^2 (%)	0.49	0.53	0.48	0.52	0.49	0.54

Table A2: SMA Belief and Future Three Months Abnormal Returns The table reports results for the regression of future 63 trading days stock abnormal returns ($ABR_{k,t+1\to t+64}$) on stock-level aggregate SMA belief $AB_{k,t}$, where t is the belief publication day. Abnormal return for stock k is computed relative to either the CAPM, Fama and French (1993) three-factor model (FF3) or a value-weighted portfolio of firms with similar size, book-to-market and momentum characteristics (SBM) as in Daniel et al. (1997). All regressions include year-month fixed effects. Additional control variables included in columns (2), (4) and (6) are past abnormal returns: $ABR_{k,t}$; $ABR_{k,t-1}$; $ABR_{k,t-2}$; $ABR_{k,t-63\to t-3}$; $Volatility_{k,t}$ measured as the sum of squared daily returns in the calendar month prior to day t; the number of professional stock analysts upgrading ($Upgrage_{k,t}$) and downgrading ($Downgrade_{k,t}$) stock k on day t; and stock k's news sentiment on day t. Values in parentheses are t-statistics based on standard errors clustered by both stock and year-month.

Panel A: Mixture Model Parameters		
	(1) Component 0	(2) Component 1
π	0.7333	0.2667
	(0.0220)	(0.0220)
μ	0.0000	0.0041
	(NA)	(0.0009)
σ	0.0000	0.0237
	(NA)	(0.0019)
$\sigma_{i,j}$	0.0136	0.0273
	(0.0002)	(0.0017)

Panel B: Mixture Return Distribution		
	Estimate	SE
Mean	0.0011	(0.0002)
SD	0.0124	(0.0002)
Fraction positive	0.1519	(0.0128)
N	4,190	

Table A3: SMA Belief Formation Ability: Constrained Two-component Mixture Model. The table reports the result for the constrained two-component mixture model of belief formation skill using data for only SMAs with at least five belief statements. Panel A reports estimates of model parameters, where π is the fraction of low and high-type SMAs, μ is the mean of each group, σ is the dispersion of each group's true belief formation skill, and $\sigma_{i,j}$ is average dispersion of estimated ability of each group. The mean and dispersion of true ability for the first skill group (Component 0), μ_0 and σ_0 respectively, are constrained to zero. Each SMA's estimated abnormal return ($\overline{ABR_i}$) is computed relative to the three-factor model for all publications by the SMA as in Eq. (6). The reported $\sigma_{i,j}$ is based on the cross-sectional average of $\overline{ABR_i}$'s standard error s_i. Hence $\sigma_{i,j} = \sqrt{\sigma_j^2 + \bar{s}^2}$. Estimates in Panel A are used to compute statistics for the cross-sectional distribution of belief formation skill reported in Panel B. SD is the standard deviation, and "Fraction positive" is the fraction of the distribution with positive ability. Standard errors (in parentheses) are computed as the standard deviation of the statistics from 1,000 bootstrap replications.

Panel A: Mixture Model Parameters			
	(1) Component 0	(2) Component 1	(3) Component 2
π	0.6583	0.2927	0.0489
	(0.1809)	(0.1735)	(0.0406)
μ	0.0007	0.0010	0.0115
	(0.0004)	(0.0006)	(0.0044)
σ	0.0020	0.0118	0.0456
	(0.0057)	(0.0054)	(0.0064)
$\sigma_{i,j}$	0.0138	0.0180	0.0476
	(0.0035)	(0.0027)	(0.0059)
Panel B: Mixture Return Distribution			
	Estimate	SE	
Mean	0.0013	(0.0002)	
SD	0.0107	(0.0003)	
P10	-0.0058	(0.0009)	
P25	-0.0013	(0.0003)	
P50	0.0007	(0.0001)	
P75	0.0029	(0.0003)	
P90	0.0088	(0.0009)	
Fraction positive	0.6015	(0.0324)	
N	4,190		

Table A4: SMA Belief Formation Ability: Three-component Mixture model. The table reports the result for the three-component mixture model of belief formation skill using data for only SMAs with at least five belief statements. Panel A reports estimates of model parameters, where π is the fraction of low and high-type SMAs, μ is the mean of each group, σ is the dispersion of each group's true belief formation skill, and $\sigma_{i,j}$ is average dispersion of estimated ability of each group. Each SMA's estimated abnormal return ($\overline{ABR_i}$) is computed relative to the three-factor model for all publications by the SMA as in Eq. (6). The reported $\sigma_{i,j}$ is based on the cross-sectional average of $\overline{ABR_i}$'s standard error s_i. Hence $\sigma_{i,j} = \sqrt{\sigma_j^2 + \bar{s}^2}$. Estimates in Panel A are used to compute statistics for the cross-sectional distribution of belief formation skill reported in Panel B. P(10) — P(90) are percentiles of the implied cross-sectional distribution of SMA skill. "Fraction positive" is the fraction of the distribution with positive ability. Standard errors (in parentheses) are computed as the standard deviation of the statistics from 1,000 bootstrap replications.

Panel A: Mixture Model Parameters		
	(1) Component 0	(2) Component 1
π	0.9290	0.0710
	(0.0915)	(0.0915)
μ	0.0007	0.0073
	(0.0002)	(0.0043)
σ	0.0028	0.0259
	(0.0022)	(0.0070)
$\sigma_{i,j}$	0.0111	0.0281
	(0.0016)	(0.0063)

Panel B: Mixture Return Distribution		
	Estimate	SE
Mean	0.0011	(0.0002)
SD	0.0105	(0.0004)
P10	-0.0033	(0.0005)
P25	-0.0013	(0.0003)
P50	0.0007	(0.0001)
P75	0.0028	(0.0003)
P90	0.0049	(0.0005)
Fraction positive	0.5937	(0.0238)
N	2,656	

Table A5: SMA Belief Formation Ability: Two-component Mixture model. The table reports the result for the two-component mixture model of belief formation skill using data for only SMAs with at least 10 belief statements. Panel A reports estimates of model parameters, where π is the fraction of low and high-type SMAs, μ is the mean of each group, σ is the dispersion of each group's true belief formation skill, and $\sigma_{i,j}$ is average dispersion of estimated ability of each group. Each SMA's estimated abnormal return ($\overline{ABR}_i$) is computed relative to the three-factor model for all publications by the SMA as in Eq. (6). The reported $\sigma_{i,j}$ is based on the cross-sectional average of $\overline{ABR}_i$'s standard error s_i. Hence $\sigma_{i,j} = \sqrt{\sigma_j^2 + \overline{s}^2}$. Estimates in Panel A are used to compute statistics for the cross-sectional distribution of belief formation skill reported in Panel B. P(10) — P(90) are percentiles of the implied cross-sectional distribution of SMA skill. "Fraction positive" is the fraction of the distribution with positive ability. Standard errors (in parentheses) are computed as the standard deviation of the statistics from 1,000 bootstrap replications.

	No. Stocks	Exc. Ret. (%)	Alpha (%)			
			CAPM	FF3	Carhart	FF5+MOM
Panel A: 21 days						
Bullish	415	1.201	0.441	0.317	0.317	0.317
		(4.70)	(2.99)	(2.10)	(2.11)	(2.18)
Bearish	106	0.651	-0.168	-0.293	-0.235	-0.147
		(1.56)	(-0.65)	(-1.38)	(-1.13)	(-0.59)
Bullish-Bearish		0.550	0.609	0.610	0.552	0.463
		(1.78)	(1.87)	(1.99)	(1.80)	(1.31)
Panel B: 63 days						
Bullish	805	1.080	0.473	0.305	0.305	0.329
		(3.59)	(3.30)	(1.88)	(1.88)	(2.07)
Bearish	235	0.610	-0.034	-0.212	-0.175	-0.054
		(1.53)	(-0.14)	(-1.07)	(-0.94)	(-0.28)
Bullish-Bearish		0.471	0.507	0.517	0.480	0.384
		(1.57)	(1.59)	(1.65)	(1.58)	(1.22)
Panel C: 126 days						
Bullish	1101	1.255	0.490	0.305	0.306	0.320
		(4.51)	(3.44)	(1.94)	(1.96)	(2.04)
Bearish	359	0.965	0.147	-0.048	-0.004	0.112
		(2.30)	(0.50)	(-0.22)	(-0.02)	(0.57)
Bullish-Bearish		0.290	0.343	0.353	0.310	0.208
		(0.94)	(1.05)	(1.12)	(1.02)	(0.68)

Table A6: Belief-based Portfolio Performance Excluding Penny Stocks. The table reports average monthly excess returns (Exc. Ret.) and alphas in per cent for portfolios based on bullish and bearish SMA beliefs. The sample of stocks excludes penny stocks with (dollar rounded) price below $5. For each SMA belief statement on a stock, a unit of the stock is added to the Bullish portfolio if the belief is bullish or to the Bearish portfolio if the belief is bearish. The position is then held for 21 (Panel A), 63 (Panel B) and 126 (Panel C) trading days. "Bullish-Bearish" is the difference portfolio that is long the Bullish portfolio and short the Bearish portfolio. The column "No. Stocks" indicate the average number of unique stocks in each portfolio daily. Alphas are relative to CAPM, three-factor (FF3), four-factor (Carhart), and five-factor plus momentum (FF5+MOM) benchmarks indicated in column headers. Reported in parentheses are t-statistics based on the Newey and West (1987) method. The SMA belief data ranges from January 2005 to December 2019.

	No. Stocks	Exc. Ret. (%)	Alpha (%)			
			CAPM	FF3	Carhart	FF5+MOM
Panel A: 21 days						
Bullish	345	1.224	0.462	0.330	0.331	0.333
		(4.76)	(3.00)	(2.04)	(2.06)	(2.12)
Bearish	92	0.644	-0.180	-0.321	-0.262	-0.178
		(1.47)	(-0.64)	(-1.36)	(-1.15)	(-0.65)
Bullish-Bearish		0.579	0.642	0.651	0.594	0.510
		(1.69)	(1.79)	(1.90)	(1.75)	(1.31)
Panel B: 63 days						
Bullish	634	1.112	0.505	0.330	0.331	0.358
		(3.63)	(3.20)	(1.82)	(1.82)	(1.98)
Bearish	199	0.602	-0.044	-0.230	-0.194	-0.074
		(1.46)	(-0.17)	(-1.09)	(-0.97)	(-0.35)
Bullish-Bearish		0.510	0.549	0.560	0.525	0.432
		(1.53)	(1.56)	(1.61)	(1.55)	(1.22)
Panel C: 126 days						
Bullish	840	1.291	0.528	0.337	0.339	0.355
		(4.57)	(3.32)	(1.89)	(1.91)	(1.96)
Bearish	299	0.981	0.160	-0.043	-0.001	0.117
		(2.32)	(0.54)	(-0.20)	(-0.02)	(0.59)
Bullish-Bearish		0.310	0.368	0.380	0.339	0.238
		(0.95)	(1.07)	(1.14)	(1.05)	(0.72)

Table A7: Belief-based Portfolio Performance Excluding Microcap Stocks. The table reports average monthly excess returns (Exc. Ret.) and alphas in per cent for portfolios based on bullish and bearish SMA beliefs. The sample of stocks excludes microcap stocks with (dollar rounded) market capitalization below the NYSE 2nd decile. For each SMA belief statement on a stock, a unit of the stock is added to the Bullish portfolio if the belief is bullish or to the Bearish portfolio if the belief is bearish. The position is then held for 21 (Panel A), 63 (Panel B) and 126 (Panel C) trading days. "Bullish-Bearish" is the difference portfolio that is long the Bullish portfolio and short the Bearish portfolio. The column "No. Stocks" indicate the average number of unique stocks in each portfolio daily. Alphas are relative to CAPM, three-factor (FF3), four-factor (Carhart), and five-factor plus momentum (FF5+MOM) benchmarks indicated in column headers. Reported in parentheses are t-statistics based on the Newey and West (1987) method. The SMA belief data ranges from January 2005 to December 2019.

Panel A	Estimated Herding Coefficient	
Concensus is	$\hat{\theta}_0$	$\chi^2\ p-value$
Equal-weighted	0.265	0.000
Specialization-weighted	0.251	0.000
Effort-weighted	0.266	0.000
Popularity-weighted	0.265	0.000

Panel B	Probability of Hitting Target							
	Herding Coefficient $\hat{\theta}_0$							
Target	-10	-1	0	0.15	0.25	0.5	1	10
1 (Bearish)	0.000	0.037	0.159	0.193	0.218	0.292	0.469	1.000
2 (Neutral)	0.000	0.019	0.037	0.041	0.044	0.052	0.072	0.975
3 (Bullish)	0.000	0.481	0.804	0.835	0.854	0.892	0.941	1.000

Table A8: Herding among SMAs. The table reports results for the herding test of Eq. (8) where the consensus belief is computed over the past three-month period. Panel A shows the estimated herding coefficient $\hat{\theta}$ and χ^2 p-value for different targets (consensus estimates). Panel B shows the probability of a belief revision hitting a hypothetical Bearish, Neutral or Bullish target for different values of $\hat{\theta}$. If $\hat{\theta} = -\infty$, the target will always be avoided. If $\hat{\theta} = 0$, the probability of hitting the target is equal to the unconditional probability of hitting the target. If $\hat{\theta} = \infty$, the target will always be hit. Values in Panel B were produced using the unconditional transition matrix and hypothetical values for $\hat{\theta}$.

Panel A:	Recession		High Volatility		Cons. Optimism (CO)	
State of Economy	$\hat{\theta}_0$	$\hat{\theta}_1$	$\hat{\theta}_0$	$\hat{\theta}_1$	$\hat{\theta}_0$	$\hat{\theta}_1$
Estimate	0.272	-0.142	0.283	-0.047	0.433	-0.230
$\chi^2\ p-value$	0.000	0.000	0.000	0.006	0.000	0.000
Panel B:	$CO \times Ret(5)$		$CO \times Ret(63)$		$CO \times Ret(126)$	
Cons. Correctness	$\hat{\theta}_0$	$\hat{\theta}_1$	$\hat{\theta}_0$	$\hat{\theta}_1$	$\hat{\theta}_0$	$\hat{\theta}_1$
Estimate	0.266	0.460	0.264	0.114	0.268	0.012
$\chi^2\ p-value$	0.000	0.001	0.000	0.006	0.000	0.676

Table A9: SMAs' Herding Conditional on Economic States. The table reports results for herding test conditional on the realization of some variable y, where the consensus belief is computed over the past three-month period. The estimated conditional herding coefficient $\hat{\theta} = \hat{\theta}_0 + \hat{\theta}_1 y$. If $\hat{\theta}_1 < 0$, herding is decreasing in y. If $\hat{\theta}_1 > 0$, herding is increasing in y. Panel A shows results where y is either the indicator for NBER recession (Recession), indicator for market volatility higher than its sample median (Market Volatility), and Consensus Optimism (CO) measured as the equal-weighted consensus minus 2. Panel B reports results when y is Consensus Correctness measured as $CO \times Ret(h)$, where $Ret(h)$ is future stock return over horizon $h \in \{5, 63, 126\}$ trading days.

	(1)	(2)	(3)	(4)
	Full sample		1st half	2nd half
Ret[-6, -2]	0.316	0.261	0.207	0.308
	(8.66)	(7.13)	(3.47)	(6.85)
Ret[-11, -7]	0.140	0.120	0.235	0.070
	(3.46)	(2.76)	(3.10)	(1.35)
Ret[-16, -12]	0.088	0.018	0.126	-0.029
	(2.00)	(0.40)	(1.86)	(-0.51)
Ret[-21, -17]	0.174	0.100	0.216	0.053
	(3.91)	(1.99)	(2.83)	(0.82)
Ret[-26, -22]	0.137	0.008	0.138	-0.059
	(3.07)	(0.17)	(1.76)	(-1.11)
Ret[-31, -27]	0.148	0.077	0.113	0.066
	(3.37)	(1.58)	(1.64)	(0.99)
Ret[-36, -32]	0.075	-0.005	0.073	-0.031
	(1.75)	(-0.13)	(1.15)	(-0.57)
Ret[-41, -37]	0.092	0.044	0.184	-0.033
	(2.05)	(1.01)	(2.89)	(-0.60)
Ret[-46, -42]	0.051	-0.005	0.058	-0.021
	(1.08)	(-0.11)	(0.81)	(-0.35)
Ret[-51, -47]	0.074	-0.025	0.142	-0.106
	(1.66)	(-0.60)	(2.50)	(-1.97)
Ret[-56, -52]	0.070	0.011	0.141	-0.054
	(1.68)	(0.27)	(1.95)	(-1.06)
Ret[-61, -57]	0.096	0.077	0.093	0.075
	(2.40)	(1.92)	(1.40)	(1.53)
Lagged(Belief)		0.124	0.116	0.126
		(15.21)	(10.94)	(12.90)
Earnings Estimate		0.045	0.034	0.049
		(4.03)	(2.42)	(3.76)
News Sentiment		0.028	0.032	0.026
		(9.36)	(6.48)	(8.03)
Obs.	192,627	142,021	42,772	99,249
R^2 (%)	1.94	5.90	4.91	6.41

Table A10: Return Extrapolation: Stock-level Linear Regression. The table reports results for panel regression of SMAs' stock-level aggregate belief ($AB_{k,t}$) on past non-overlapping weekly (5 trading days) returns as in Eq. (9). Ret[$-j$, $-i$] is stock k's return computed over window $t-j$ to $t-i$ trading days, where t is the belief publication day. Columns (1) and (2) show results for the full sample. Column (3) shows the result for the first half of the sample 2005 – 2012, while Column (4) shows the result for the second half 2013 – 2019. Lagged(Belief) is the lag of $AB_{k,t}$, Earnings Estimate is stock k's professional analyst consensus forecast of quarterly earnings per share as of last calendar month. News Sentiment is stock k's cash flow news sentiment averaged over the past week ending $t-2$. All regressions include year-month fixed effects. Reported in parenthesis are t-statistics based on standard errors clustered by stock and year-month.

Internet Appendix

for

Should Retail Investors Listen to Social Media Analysts?
Evidence from Text-Implied Beliefs

By Chukwuma Dim

This version: September 2021

IA Robustness and Additional Results

	CAPM		FF3		SBM	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 5 days						
$AB_{k,t}$	0.0029	0.0028	0.0029	0.0029	0.0029	0.0028
	(7.70)	(7.71)	(7.97)	(7.97)	(7.93)	(7.92)
Controls	NO	YES	NO	YES	NO	YES
Obs.	104,441	104,441	104,441	104,441	104,441	104,441
R^2 (%)	0.74	0.85	0.62	0.73	0.64	0.74
Panel B: 21 days						
$AB_{k,t}$	0.0038	0.0035	0.0038	0.0036	0.0037	0.0034
	(4.53)	(4.40)	(4.85)	(4.65)	(4.85)	(4.70)
Controls	NO	YES	NO	YES	NO	YES
Obs.	104,154	104,154	104,154	104,154	104,154	104,154
R^2 (%)	1.57	1.69	1.07	1.19	1.18	1.30
Panel C: 63 days						
$AB_{k,t}$	0.0038	0.0037	0.0039	0.0038	0.0040	0.0039
	(2.12)	(2.10)	(2.28)	(2.28)	(2.48)	(2.46)
Controls	NO	YES	NO	YES	NO	YES
Obs.	103,771	103,771	103,771	103,771	103,771	103,771
R^2 (%)	1.75	1.82	1.24	1.30	1.29	1.35

Table IA1: SMAs' Stated Beliefs and Future Abnormal Returns. The table reports results for the panel regression of future abnormal returns ($ABR_{k,t+1\to t+1+h}$) on stock-level aggregate SMAs' *stated* beliefs $AB_{k,t}$ based on Eq. (2) in the main text. Only explicitly stated beliefs (i.e., excluding beliefs inferred with ML) is used to compute $AB_{k,t}$ from Eq. (1). Abnormal return for firm k over horizon $t+1$ to $t+1+h$, with t being the belief publication day, is computed relative to either the CAPM, Fama and French (1993) three-factor model (FF3) or a value-weighted portfolio of firms with similar size, book-to-market and momentum characteristics (SBM) as in Daniel et al. (1997). Panel A reports results for horizon $h = 5$ trading days, Panel B for $h = 21$ trading days, and Panel C for $h = 63$ trading days. All regressions include year-month fixed effects. Control variables included in columns (2), (4) and (6) are past abnormal returns: $ABR_{k,t}$, $ABR_{k,t-1}$, $ABR_{k,t-2}$, $ABR_{k,t-h\to t-3}$; $Volatility_{k,t}$, the number of professional stock analysts upgrading and downgrading stock k on day t respectively; and cash flow news sentiment. All right hand side variables are normalized to unit variance. Values in parentheses are t-statistics based on standard errors clustered by both firm and year-month.

	CAPM		FF3		SBM	
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Two weeks $t-14$ to t average belief ($\overline{AB}_{k,t}$)					
ABR horizon: 5 days						
$\overline{AB}_{k,t}$	0.0031	0.0030	0.0031	0.0031	0.0031	0.0031
	(7.85)	(7.89)	(8.09)	(8.12)	(8.29)	(8.31)
Controls	NO	YES	NO	YES	NO	YES
Obs.	104,441	104,441	104,441	104,441	104,441	104,441
R^2 (%)	0.76	0.87	0.64	0.75	0.67	0.77
ABR horizon: 21 days						
$\overline{AB}_{k,t}$	0.0041	0.0039	0.0042	0.0039	0.0041	0.0038
	(4.33)	(4.20)	(4.61)	(4.41)	(4.70)	(4.56)
Controls	NO	YES	NO	YES	NO	YES
Obs.	104,154	104,154	104,154	104,154	104,154	104,154
R^2 (%)	1.58	1.70	1.09	1.20	1.20	1.31
ABR horizon: 63 days						
$\overline{AB}_{k,t}$	0.0040	0.0038	0.0041	0.0040	0.0043	0.0041
	(1.88)	(1.85)	(2.03)	(2.02)	(2.22)	(2.20)
Controls	NO	YES	NO	YES	NO	YES
Obs.	103,771	103,771	103,771	103,771	103,771	103,771
R^2 (%)	1.75	1.82	1.24	1.30	1.29	1.35
	Panel B: One month $t-30$ to t average belief ($\overline{AB}_{k,t}$)					
ABR horizon: 5 days						
$\overline{AB}_{k,t}$	0.0030	0.0030	0.0030	0.0030	0.0031	0.0030
	(7.29)	(7.39)	(7.43)	(7.50)	(7.79)	(7.87)
Controls	NO	YES	NO	YES	NO	YES
Obs.	104,441	104,441	104,441	104,441	104,441	104,441
R^2 (%)	0.75	0.86	0.63	0.74	0.66	0.76
ABR horizon: 21 days						
$\overline{AB}_{k,t}$	0.0044	0.0042	0.0045	0.0042	0.0044	0.0041
	(4.52)	(4.41)	(4.73)	(4.56)	(4.88)	(4.76)
Controls	NO	YES	NO	YES	NO	YES
Obs.	104,154	104,154	104,154	104,154	104,154	104,154
R^2 (%)	1.59	1.71	1.10	1.21	1.21	1.32
ABR horizon: 63 days						
$\overline{AB}_{k,t}$	0.0043	0.0041	0.0044	0.0042	0.0045	0.0043
	(1.93)	(1.89)	(2.07)	(2.05)	(2.24)	(2.22)
Controls	NO	YES	NO	YES	NO	YES
Obs.	103,771	103,771	103,771	103,771	103,771	103,771
R^2 (%)	1.75	1.82	1.25	1.30	1.30	1.36

Table IA2: Average SMAs' Stated Beliefs and Future Abnormal Returns. The table reports results for the panel regression of future abnormal returns (ABR) on stock-level aggregate SMAs' *stated* beliefs averaged over the past two weeks $t-14$ to t (Panel A) and over the past one month $t-30$ to t (Panel B) denoted $\overline{AB}_{k,t}$, with t being a belief statement day. The regression is based on Eq. (2) in the main text. Only explicitly stated beliefs (i.e., excluding beliefs inferred with ML) is used to compute $AB_{k,t}$ from Eq. (1). Abnormal returns for firm k is computed over horizon $t+1$ to $t+1+h$, for $h \in \{5, 21, 63\}$ trading days. Abnormal return is relative to either the CAPM, three-factor model (FF3) or the size/book-to-market/momentum characteristics-based benchmark (SBM) as in the main text. All regressions include year-month fixed effects. Control variables included in columns (2), (4) and (6) are defined under Table IA1. All right hand side variables are normalized to unit variance. Values in parentheses are t-statistics based on standard errors clustered by both firm and year-month.

	CAPM		FF3		SBM	
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Two weeks $t-14$ to t average belief ($\overline{AB}_{k,t}$)					
ABR horizon: 5 days						
$\overline{AB}_{k,t}$	0.0020	0.0020	0.0019	0.0018	0.0020	0.0019
	(7.42)	(7.45)	(7.20)	(7.19)	(7.52)	(7.57)
Controls	NO	YES	NO	YES	NO	YES
Obs.	194,172	194,172	194,172	194,172	194,172	194,172
R^2 (%)	0.32	0.43	0.31	0.43	0.32	0.43
ABR horizon: 21 days						
$\overline{AB}_{k,t}$	0.0029	0.0027	0.0026	0.0025	0.0028	0.0026
	(4.83)	(4.70)	(4.71)	(4.50)	(4.83)	(4.73)
Controls	NO	YES	NO	YES	NO	YES
Obs.	193,521	193,521	193,521	193,521	193,521	193,521
R^2 (%)	0.31	0.37	0.32	0.38	0.30	0.36
ABR horizon: 63 days						
$\overline{AB}_{k,t}$	0.0039	0.0037	0.0030	0.0028	0.0036	0.0034
	(2.44)	(2.32)	(2.04)	(1.90)	(2.49)	(2.37)
Controls	NO	YES	NO	YES	NO	YES
Obs.	192,650	192,650	192,650	192,650	192,650	192,650
R^2 (%)	0.50	0.54	0.48	0.52	0.50	0.54
	Panel B: One month $t-30$ to t average belief ($\overline{AB}_{k,t}$)					
ABR horizon: 5 days						
$\overline{AB}_{k,t}$	0.0018	0.0018	0.0017	0.0017	0.0018	0.0018
	(6.63)	(6.66)	(6.49)	(6.48)	(6.85)	(6.90)
Controls	NO	YES	NO	YES	NO	YES
Obs.	194,172	194,172	194,172	194,172	194,172	194,172
R^2 (%)	0.31	0.42	0.30	0.42	0.30	0.42
ABR horizon: 21 days						
$\overline{AB}_{k,t}$	0.0029	0.0027	0.0027	0.0025	0.0028	0.0026
	(4.64)	(4.51)	(4.53)	(4.32)	(4.64)	(4.55)
Controls	NO	YES	NO	YES	NO	YES
Obs.	193,521	193,521	193,521	193,521	193,521	193,521
R^2 (%)	0.32	0.37	0.32	0.38	0.30	0.35
ABR horizon: 63 days						
$\overline{AB}_{k,t}$	0.0039	0.0037	0.0030	0.0028	0.0035	0.0033
	(2.24)	(2.11)	(1.88)	(1.74)	(2.25)	(2.13)
Controls	NO	YES	NO	YES	NO	YES
Obs.	192,650	192,650	192,650	192,650	192,650	192,650
R^2 (%)	0.50	0.54	0.48	0.52	0.49	0.54

Table IA3: Average SMA Beliefs and Future Abnormal Returns. The table reports results for the panel regression of future stock abnormal returns (ABR) on stock-level aggregate SMAs' beliefs averaged over the past two weeks $t-14$ to t (Panel A) and over the past one month $t-30$ to t (Panel B), with t being a belief publication day. The regression is based on Eq. (2) in the main text. Abnormal returns for firm k is computed over horizon $t+1$ to $t+1+h$, for $h \in \{5, 21, 63\}$ trading days. Abnormal returns are relative to either the CAPM, three-factor model (FF3) or the size/book-to-market/momentum characteristics-based benchmark (SBM) as in the main text. All regressions include year-month fixed effects. Other Control variables included in columns (2), (4) and (6) are defined under Table IA1. All right-hand side variables are normalized to unit variance. Values in parentheses are t-statistics based on standard errors clustered by both firm and year-month.

	CAPM		FF3		SBM	
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Subsample 2005 - 2012					
ABR horizon: 5 days						
$AB_{k,t}$	0.0018	0.0018	0.0017	0.0017	0.0018	0.0017
	(3.95)	(3.77)	(3.67)	(3.51)	(3.73)	(3.54)
Controls	NO	YES	NO	YES	NO	YES
Obs.	55,873	55,873	55,873	55,873	55,873	55,873
R^2 (%)	0.88	1.01	0.90	1.01	0.87	1.01
ABR horizon: 21 days						
$AB_{k,t}$	0.0026	0.0025	0.0029	0.0027	0.0027	0.0027
	(2.69)	(2.68)	(3.10)	(3.05)	(2.86)	(2.92)
Controls	NO	YES	NO	YES	NO	YES
Obs.	55,666	55,666	55,666	55,666	55,666	55,666
R^2 (%)	0.67	0.69	0.70	0.74	0.67	0.71
ABR horizon: 63 days						
$AB_{k,t}$	0.0031	0.0031	0.0030	0.0029	0.0034	0.0035
	(1.44)	(1.48)	(1.40)	(1.38)	(1.53)	(1.61)
Controls	NO	YES	NO	YES	NO	YES
Obs.	55,413	55,413	55,413	55,413	55,413	55,413
R^2 (%)	1.01	1.06	1.06	1.11	1.04	1.13
	Panel B: Subsample 2013 - 2019					
ABR horizon: 5 days						
$AB_{k,t}$	0.0018	0.0018	0.0017	0.0017	0.0017	0.0017
	(6.51)	(6.72)	(6.50)	(6.72)	(6.60)	(6.88)
Controls	NO	YES	NO	YES	NO	YES
Obs.	138,299	138,299	138,299	138,299	138,299	138,299
R^2 (%)	0.25	0.41	0.24	0.41	0.24	0.40
ABR horizon: 21 days						
$AB_{k,t}$	0.0025	0.0023	0.0022	0.0019	0.0024	0.0021
	(3.98)	(3.80)	(3.62)	(3.35)	(4.02)	(3.84)
Controls	NO	YES	NO	YES	NO	YES
Obs.	137,855	137,855	137,855	137,855	137,855	137,855
R^2 (%)	0.39	0.55	0.38	0.53	0.36	0.50
ABR horizon: 63 days						
$AB_{k,t}$	0.0037	0.0034	0.0028	0.0025	0.0032	0.0029
	(2.39)	(2.24)	(2.00)	(1.83)	(2.51)	(2.34)
Controls	NO	YES	NO	YES	NO	YES
Obs.	137,237	137,237	137,237	137,237	137,237	137,237
R^2 (%)	0.66	0.78	0.64	0.76	0.63	0.74

Table IA4: SMA Beliefs and Future Abnormal Returns: Subsample Regressions. The table reports results for two subsample panel regressions of future stock abnormal returns (ABR) on stock-level aggregate SMA beliefs. The regressions are based on Eq. (2) in the main text. Panel A reports results for the 2005 - 2012 subsample, and Panel B reports results for the 2013 - 2019 subsample. Abnormal returns for firm k is computed over horizon $t+1$ to $t+1+h$, with $h \in \{5, 21, 63\}$ trading days and t the belief publication day. Abnormal returns are relative to either the CAPM, three-factor model (FF3) or the size/book-to-market/momentum characteristics-based benchmark (SBM) as in the main text. All regressions include year-month fixed effects. Other Control variables included in columns (2), (4) and (6) are defined under Table IA1. All right-hand side variables are normalized to unit variance. Values in parentheses are t-statistics based on standard errors clustered by both firm and year-month.

Panel A: Mixture Model Parameters		
	(1) Component 0	(2) Component 1
π	0.8845	0.1155
	(0.0471)	(0.0471)
μ	0.0006	0.0065
	(0.0002)	(0.0023)
σ	0.0044	0.0343
	(0.0017)	(0.0038)
$\sigma_{i,j}$	0.0146	0.0370
	(0.0013)	(0.0034)
Panel B: Mixture Return Distribution		
	Estimate	SE
Mean	0.0012	(0.0002)
SD	0.0136	(0.0003)
P10	-0.0060	(0.0005)
P25	-0.0027	(0.0003)
P50	0.0007	(0.0002)
P75	0.0040	(0.0003)
P90	0.0078	(0.0005)
Fraction positive	0.5532	(0.0133)
N	4,188	

Table IA5: SMA Belief Formation Ability: Two-component Mixture model. The table reports the result for the two-component mixture model of belief formation skill using data for only SMAs with at least five belief statements and measuring an SMA's estimated abnormal return ($\overline{ABR_i}$) relative to the CAPM benchmark as in Eq. (6). Panel A reports estimates of model parameters, where π is the fraction of low and high-type SMAs, μ is the mean of each group, σ is the dispersion of each group's true belief formation skill, and $\sigma_{i,j}$ is average dispersion of estimated ability of each group. The reported $\sigma_{i,j}$ is based on the cross-sectional average of $\overline{ABR_i}$'s standard error s_i. Hence $\sigma_{i,j} = \sqrt{\sigma_j^2 + \bar{s}^2}$. Estimates in Panel A are used to compute statistics for the cross-sectional distribution of belief formation skill reported in Panel B. P(10) — P(90) are percentiles of the implied cross-sectional distribution of SMA skill. "Fraction positive" is the fraction of the distribution with positive ability. Standard errors (in parentheses) are computed as the standard deviation of the statistics from 1,000 bootstrap replications.

Panel A: Mixture Model Parameters			
	(1) Component 0	(2) Component 1	(3) Component 2
π	0.5978	0.3488	0.0534
	(0.1367)	(0.1267)	(0.0414)
μ	0.0006	0.0009	0.0112
	(0.0003)	(0.0004)	(0.0042)
σ	0.0019	0.0103	0.0457
	(0.0063)	(0.0042)	(0.0059)
$\sigma_{i,j}$	0.0141	0.0174	0.0478
	(0.0042)	(0.0017)	(0.0053)

Panel B: Mixture Return Distribution		
	Estimate	SE
Mean	0.0012	(0.0002)
SD	0.0107	(0.0003)
P10	-0.0067	(0.0008)
P25	-0.0015	(0.0004)
P50	0.0006	(0.0002)
P75	0.0028	(0.0004)
P90	0.0093	(0.0008)
Fraction positive	0.5870	(0.0388)
N	4,188	

Table IA6: SMA Belief Formation Ability: Three-component Mixture model. The table reports the result for the three-component mixture model of belief formation skill using data for only SMAs with at least five belief statements and measuring an SMA's estimated abnormal return ($\overline{ABR}_i$) relative to the CAPM benchmark as in Eq. (6). Panel A reports estimates of model parameters, where π is the fraction of low and high-type SMAs, μ is the mean of each group, σ is the dispersion of each group's true belief formation skill, and $\sigma_{i,j}$ is average dispersion of estimated ability of each group. The reported $\sigma_{i,j}$ is based on the cross-sectional average of $\overline{ABR}_i$'s standard error s_i. Hence $\sigma_{i,j} = \sqrt{\sigma_j^2 + \overline{s}^2}$. Estimates in Panel A are used to compute statistics for the cross-sectional distribution of belief formation skill reported in Panel B. P(10) — P(90) are percentiles of the implied cross-sectional distribution of SMA skill. "Fraction positive" is the fraction of the distribution with positive ability. Standard errors (in parentheses) are computed as the standard deviation of the statistics from 1,000 bootstrap replications.