Universitat Pompeu Fabra Barcelona
Department of Economics and Business

Economics Working Paper Series

Working Paper No. 1846

Tweeting for money: Social media and mutual fund flows

Javier Gil-Bazo and Juan F. Imbet

October 2022

Tweeting for Money: Social Media and Mutual Fund Flows*

Javier Gil-Bazo[§] Juan F. Imbet[‡]

October 12, 2022

Abstract

We investigate whether asset management firms use social media to persuade investors. Combining a database of almost 1.6 million Twitter posts by U.S. mutual fund families with textual analysis, we find that flows of money to mutual funds respond positively to tweets with a positive tone. Consistently with the persuasion hypothesis, positive tweets work best when they convey advice or views on the market and when investor sentiment is higher. Using a high-frequency approach, we are able to identify a short-lived impact of families' tweets on ETF share prices. Finally, we reject the alternative hypothesis that asset management companies use social media to alleviate information asymmetries by either lowering search costs or disclosing privately observed information.

Keywords: social media; Twitter; persuasion; mutual funds; mutual fund flows; machine learning; textual analysis.

JEL Classification: G11; G23; D83

*The authors thank Patrick Bolton, Victor DeMiguel, Ariadna Dumitrescu, Rich Evans, Javier Gomez Biscarri, Philip Hamill, Filippo Ippolito, Alberto Manconi, Andrea Polo, Emanuele Rizzo, Maria Tselika and Fernando Zapatero as well as seminar participants at CUNEF, Universitat Pompeu Fabra, Université Paris Dauphine, and conference participants at the 2021 International Conference of the French Finance Association, the 2021 Finance Forum of the Spanish Finance Association, the 2022 FMA Consortium on Asset Management and Fintech, and the 2022 British Accounting and Finance Association SWAG Annual Conference for helpful comments and suggestions. Javier Gil-Bazo acknowledges financial support from the Spanish Ministry of Science and Innovation (research grant PID2020-118541GB-I00), and Spanish Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2019-000915-S). Both authors acknowledge financial support from the FinTech & Digital Finance Chair at Dauphine-PSL Paris.

[§]Universitat Pompeu Fabra, Barcelona School of Economics, and UPF Barcelona School of Management, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain, e-mail: javier.gil-bazo@upf.edu.

[‡]Corresponding author. Université Paris Dauphine and PSL Research University, Pl. du Maréchal de Lattre de Tassigny, 75016 Paris, e-mail: juan.imbet@dauphine.psl.eu

1. Introduction

As of July 2022, more than 4.7 billion people in the world, 75.5% of the population aged 13 years and older, were using social media to communicate with others, entertain themselves, and stay informed.[1] The growing popularity of social media has raised concerns about their potential to misinform the public and manipulate individuals' opinions and behavior (e.g., Abramowitz, 2017; Aral and Eckles, 2019). In the context of financial markets, the online activities of some high-profile individuals have prompted investigations by the Securities and Exchange Commission (SEC).[2] Also, some firm managers have recognized social media as an opportunity to engage directly with their investors.[3] Indeed, if social media can be used to influence investors' perceptions, then companies issuing securities to the public have strong incentives to become active participants. Beyond anecdotal evidence, whether organizations use social media to persuade investors and the effectiveness of this strategy are open empirical questions. In this paper, we focus on the market for mutual funds and study whether asset management firms use social media to attract money from investors.

Mullainathan and Shleifer (2005) and Mullainathan, Schwartzstein, and Shleifer (2008) argue that individuals evaluate various propositions or objects by grouping them into categories. Persuaders take advantage of such associative thinking by creating *uninformative* messages that associate their product with a given category and its attributes. For example, asset management companies may want to associate themselves with enhancers of opportunities in times of high market returns and with trustworthy advisers in down markets. In this sense, social media is an ideal tool for persuasion for asset management firms, as it

[1] Data from https://datareportal.com/social-media-users.

[2] Mohamed (2021),"Big Short' investor Michael Burry says he'll stop tweeting after SEC regulators paid him a visit," Businessinsider.com, (https://markets.businessinsider.com/currencies/news/big-short-investor-michael-burry-stop-tweets-sec-regulators-visit-2021-3-1030222890); SEC (2018),"Elon Musk Charged With Securities Fraud for Misleading Tweets," (https://www.sec.gov/news/press-release/2018-219); Spichak (2021), "Elon Musk Hopes SEC Will Investigate Him over Dogecoin Tweets: 'It Would Be Awesome'," Newsweek, (https://www.newsweek.com/elon-musk-sec-investigation-dogecoin-bitcoin-cryptocurrency-tweets-1572290).

[3] Gladstone and Schwartzel (2021), "AMC Boss Adam Aron Basks in Meme-Stock Spotlight," Wall Street Journal (https://www.wsj.com/articles/amc-boss-adam-aron-basks-in-meme-stock-spotlight-11622799000)

allows them to communicate with current and prospective investors without the strict constraints imposed by mandatory information disclosures on the timing, content, and framing of information.[4] However, it is unclear whether they can succeed in their endeavor. The mutual fund market is highly regulated and abundant hard information is already available through mandatory disclosures, such as fund prospectuses and statements of additional information. More importantly, if investors understand the ability of asset management firms to strategically select and frame information, any attempts to influence investors could be self-defeating. This is particularly true for social media where, unlike advertising or mandatory disclosures, users are not passive receivers of messages. Instead, they can reply to the posts of firms and criticize them. Moreover, these interactions are observable by other users who can amplify messages by *liking*, reposting or posting them across different platforms. These features potentially limit the extent to which asset management firms can influence investor perception of their products.

The mutual fund industry is an ideal laboratory to study persuasion in financial markets, as thousands of actively managed mutual funds compete for investors' money in a context in which quality is hard to evaluate. Twitter is also particularly appropriate for our purposes given its rising popularity among investors. Indeed, several studies have provided evidence that Twitter activity can predict prices of stocks and other asset classes (Bollen, Mao, and Zeng, 2011; Ranco, Aleksovski, Caldarelli, Grčar, and Mozetič, 2015; You, Guo, and Peng, 2017; Gholampour and van Wincoop, 2017; Gu and Kurov, 2020; Bianchi, Gomez-Cram, Kind, and Kung, 2019; Bianchi, Gomez-Cram, and Kung, 2021). Also, the presence of asset management firms in Twitter has grown at a very fast pace in the last years. In our sample,

[4]Note, however, that advertisement and retail investor communication by asset management companies must comply with SEC rule 482 and FINRA rule 2210. In 2003, SEC rule 482 modified the Securities Act of 1933-Section 5 that stated that all fund advertisement must have information that is contained in the statutory prospectus. With rule 482, investment companies are allowed to include information not included in the statutory prospectus. This allows investment companies to include up-to-date information in rule-482 advertisements, such as information about current economic conditions that are not commonly included in a fund's prospectus. FINRA Rule 2210 governs communications with the public including communications with retail and institutional investors. The rule provides standards for the content, approval, recordkeeping and filing of communications with FINRA. The rule prohibits false, exaggerated, unwarranted, and misleading information communications, as well as projections of future performance.

the number of posts on Twitter (tweets) by all mutual fund families went from almost zero prior to 2009 to around 10,000 tweets per month in 2020.

To investigate whether asset management companies influence investors' decisions through social media communications, we build a database of Twitter posts by mutual fund families managing domestic equity funds in the U.S. between January 2009 and December 2020. We then employ supervised machine learning algorithms to classify tweets into positive or negative and compute the positiveness of the tone of asset management firms' tweets in a given month. In addition, we use an unsupervised topic modelling algorithm to study the content of these tweets. Finally, we merge these data with data on U.S. domestic equity mutual funds obtained from the CRSP Survivor-Bias-Free U.S. Mutual Fund database.

Our findings can be summarized as follows. First, 284 of 939 firms managing U.S. equity funds in our sample have a Twitter account and post at least one tweet during the sample period. Families that use Twitter tend to manage more assets, more funds, and funds in more investment categories than families that do not use Twitter, which suggests that economies of scale play a role in the decision to implement a social media strategy. Among those firms that use Twitter, intensity of usage is higher among younger firms and firms offering funds with better recent performance, lower fees, lower risk, and slower recent growth. Across all non index U.S. diversified equity funds, fund families that use Twitter manage in average 3.2 USD trillion which corresponds to 66% of the average total assets. Moreover, they operate in average 6417 different funds, or 54%.

Second, a more positive tone in a family's tweets in a given month predicts significantly higher flows to the family's funds in the following month. The increase in flows is economically significant. A one standard deviation increase in the positiveness of the tone of tweets is associated with an increase in assets under management of 9 basis points (bp) in the following month, or 7.25 USD million for the average family. This result is robust to different ways of modelling the flow-performance relationship and to the inclusion of time, fund, and fund family fixed effects, and survives after correcting for potential sample selection bias. Moreover, the link between activity on Twitter and mutual fund flows is not explained by fund families tweeting about a number of known events that may influence flows positively, such as a change in the fund's management, the presence of funds in the family with very recent stellar performance, or any other news about the fund family disseminated through social media by third parties.

Third, while these results suggest that fund families attract investors' money by tweeting with a more positive tone, it is unclear that persuasion is the mechanism that drives this association. To tackle this question, we test two different implications of the persuasion hypothesis. First, we exploit our topic classification of tweets and show that the positive link between the tone of tweets and flows is stronger for tweets classified as "financial advice" and "market commentary" than for tweets classified as "customer service." This evidence rules out the simple explanation that asset management firms use social media *only* as an additional channel for providing customer service. Second, the positive link between the tone of tweets and flows to equity funds is stronger in times of high investor sentiment, consistently with asset management companies using social media to persuade investors to take more risk in times of increased perceived opportunities.

Fourth, although in our tests we control for a variety of observable fund flow determinants as well as unobservable time-invariant fund and fund family characteristics, we cannot discard that tweets respond to unobservable news that, in turn, predict fund flows. To identify the effect of asset management companies' tweets on investor decisions, we combine a high-frequency approach with data on intraday ETF trading data. The analysis is based on the idea that asset management companies' positive tweets increase the demand not just for their open-end mutual funds but also for their ETFs. Since ETF shares are traded in an exchange, we expect a family's positive tweets to put upward pressure on the price of its ETFs. The fact that ETFs are traded intradaily allows us to isolate the effect of tweeting activity on prices using the high-frequency methodology of Bianchi, Gomez-Cram, Kind, and Kung (2019) and Bianchi, Gomez-Cram, and Kung (2021). More specifically, we study whether an ETF's price changes in a very short time window around the time when the ETF's family tweets. As explained by those authors, this high-frequency approach allows for clean identification of the effect of tweets on prices with the identifying assumption being that over such a short window of time no other relevant information affecting the security's price is released. Our results indicate that a positive tweet by an asset management company increases the price of its ETFs by 5 basis points. Consistently with the effect on ETF prices being driven by investor sentiment and not fundamentals, the price increase is fully reversed 35 minutes after the tweet. These findings provide support for the hypothesis that asset management companies' tweets affect fund investor decisions through social media.

Fifth, we consider and test two alternative explanations for our results. The first hypothesis is that asset management companies use social media to reduce search costs for investors. A reduction in search costs would increase flows of money coming from new investors but would not affect redemptions from existing investors, as those investors have already paid the search cost (Sirri and Tufano, 1998; Hortaçsu and Syverson, 2004). Also, a reduction in search costs would decrease the convexity of the flow-performance relation, since investors would be willing to consider a larger set of funds and not just recent winners (Huang, Wei, and Yan, 2007). Contrary to these predictions, we find that positive tweets not only increase inflows but also reduce outflows from existing investors. Also, the flow-performance relation does not become less convex following positive tweets.

We also investigate whether asset management companies use social media to convey to investors information that predicts future fund performance and is not available to the public, consistently with the model of Dumitrescu and Gil-Bazo (2016) of strategic communication by asset managers. Controlling for potential determinants of fund performance, we find that the positiveness of an asset management company's tweets does not predict future fund performance.

Therefore, the evidence does not support the notion that social media communications of asset management firms alleviate information asymmetries between mutual fund companies and investors by either reducing search costs or conveying new information to investors.

We contribute to the literature by unveiling a link between asset management companies' social media communications and mutual fund investors' decisions. Our paper is related to studies that investigate advertising in the mutual fund industry. Sirri and Tufano (1998) show that marketing effort, as proxied by fund fees, is associated with larger fund flows. Jain and Wu (2000) find that management companies tend to advertise funds with higher recent performance. Cronqvist (2006) investigates the content of mutual fund advertisements in Sweden and finds that most fund ads are not informative about fund quality and yet influence individuals' portfolio decisions. Gallaher, Kaniel, and Starks (2015) show that mutual fund families' advertising expenditures attract flows to the family's funds as well as to other funds in the industry, reduce redemptions, and increase the convexity of the flow-performance relationship. Our paper is also related to the study of Hillert, Niessen-Ruenzi, and Ruenzi (2018), who find a positive association between the tone of the text in shareholder letters from asset management companies and subsequent fund flows. Social media is fundamentally different from advertising or shareholder letters in that it allows users to interact with firms' messages and counteract or amplify their effect. Another important distinction is that social media communications take place in real time, which allows us to determine with absolute precision the time at which messages are released. In section 6 we exploit this feature to identify the effect of Twitter posts on ETF share prices.

Our paper also contributes more generally to the literature on non-mandatory corporate disclosures (Kim and Verrecchia, 1991; Dye and Sridhar, 2004; Dye and Sridhar, 2004; Cornelli, Kominek, and Ljungqvist, 2013; Bertomeu and Marinovic, 2016), and to the recent literature of textual analysis in Finance and Accounting (see Loughran and McDonald, 2016 for a survey of the literature). The study of Blankespoor, Miller, and White (2014) is particularly relevant to our paper. These authors show that when public firms use Twitter to disseminate firm-initiated news, information asymmetries decline as evidenced by narrower bid-ask spreads. In contrast, our results suggest that Twitter does not help alleviate information asymmetries in the mutual fund industry.

2. Data

In this section, we present the data used in the analysis. We draw on two data sets, the CRSP Survivor-Bias-Free U.S. Mutual Fund database and a database of tweets from January 2009 to December 2020 posted by mutual fund families. From the former, we obtain information at the share class level on returns, assets under management, investment category, expenses, ticker, and age. Even though our Twitter database starts in 2009, we collect mutual fund data from 2006 so we can use three years of prior historical data to estimate risk-adjusted returns.

For our main analysis, we keep only non index U.S. diversified equity funds following CRSP style level one. To avoid discrepancies between the objective reported by CRSP and the Lipper classification of the fund we manually discard those funds whose Lipper objective class name does not correspond to an equity fund. Among these funds we discard less than one percent of the funds that are classified as municipal debt funds, money market funds, debt funds, and funds focusing on other securities.

To construct variables at the mutual fund level, we follow the same share aggregation procedure as Gil-Bazo and Ruiz-Verdú (2009). Total Net Assets (TNA) of a fund are the sum of the TNA for the fund's share classes. Returns and expense ratios are TNA-weighted averages across all share classes in the fund. The age of the fund is the age of the oldest share class in the fund. Following Berk and van Binsbergen (2016) and Pastor, Stambaugh, and Taylor (2015) we drop a fund's observations before the fund reaches USD 15 million in TNA.

To create some of our variables, we aggregate data at the fund family level based on the CRSP identifier mgmt_code. TNA at the family level is the sum of the TNA of each fund in the family, the age of the fund family is the age of the oldest fund in the family, and flows, expenses and returns are weighted averages across all funds in the family (based on the TNA of each fund in the family).[5]

[5] We compute flows at the family level using a weighted average of the flows of each fund in the family since we discard money market, bond, and hybrid funds from the sample before aggregation.

For a subsample of funds, and until 2016, we obtain data on inflows and outflows, as in Christoffersen, Evans, and Musto (2013) and Ha and Ko (2019). These data are reported in SEC's N-SAR form, Item 28, which includes cash-flow information on a monthly basis at the portfolio level.[6]

Given the findings of Barber, Huang, and Odean (2016) and Berk and van Binsbergen (2016) that investors appear to use the CAPM to evaluate mutual fund performance, throughout the paper we focus on CAPM alphas as a determinant of flows, although we test the robustness of our results to using the three-factor and four-factor models of Fama and French (1993) and Carhart (1997) to estimate performance.[7] We compute the risk-adjusted return, $\hat{\alpha}_{it}$, of fund i in month t as the intercept plus the residual of the CAPM model:

$$\hat{\alpha}_{it} = r_{it}^e - \hat{\beta}_{it} r_{mt}^e, \tag{1}$$

where r_{it}^e is the excess return of fund i at month t over the risk free rate and r_{mt}^e is the excess return of the market portfolio over the risk free rate. We obtain the monthly risk-free rate and the market portfolio return from Prof. Kenneth French's website and $\hat{\beta}_{it}$ is estimated for each fund and month t by running OLS rolling regressions of excess returns on market excess return over the three-year period ending in month $t-1$. If less than three years of data are available in a given window, we require the fund to have at least 30 months of data and run the regressions with the data available.

To construct the database of mutual fund families' Twitter accounts we obtain the names of all asset management companies in the CRSP database managing U.S. equity funds. Then, we perform a manual search through each one of the family names represented in the variable mgmt_name in the CRSP database and group similar names using the CRSP aggregation variable mgmt_code. Finally, we search for each family's Twitter account in the asset management company's website.

[6]We thank Yeonjeong Ha and Kwangsoo Ko for kindly sharing their data with us.

[7]Evans and Sun (2020) show that mutual fund flows have become more sensitive to three-factor abnormal returns since Morningstar changed its methodology to compute fund ratings to account for funds' investment style.

Once the list of Twitter accounts is collected, we web scrape all tweets from accounts that are active on two different dates: February 2018 and December 2021. It is important to notice that if a fund family was on Twitter and decided to cancel its Twitter account before February 2018, we would not be able to get information on that family's tweets. Similarly, for accounts that existed in February 2018 and were canceled before December 2021, we would have no information in that subperiod. The web-scraping procedure downloads tweets historically starting from the most recent tweet up to the first one. Web-scraping algorithms can get banned temporarily and the download procedure may stop prematurely. To ensure we download all information, we compare the earliest tweet obtained for each company with the true first tweet of the account as provided by Twitter.[8] Our database contains 1,592,486 tweets from 284 different usernames, from January 2009 to December 2020.

The procedure used to measure the positiveness of tweets is explained in detail in the Appendix and can be summarized as follows. To determine the tone of a tweet, we use a training sample with previously manually classified tweets. We consider both the occurrence of a word in the text and its Part of Speech (POS) as features.[9] To avoid any subjectivity in choosing the machine learning algorithm to classify the tweets, we use six different algorithms and select for each tweet the most voted label among them. If three algorithms classify a tweet as positive and three as negative, we consider the tweet to have a neutral tone. Using this voting scheme, all tweets in our sample are classified as either positive, neutral or negative. The approach also provides us with a measure of confidence in the classification. In particular, we define the confidence of classifying tweet k as c as:

$$w_k^c = \frac{\text{Number of algorithms that classify tweet } k \text{ with label } c}{\text{Total number of algorithms}} \tag{2}$$

We then define the positiveness of a family's tweets in month t as follows (we drop the fund

[8]The first tweet of any active account was found using the webpage https://discover.twitter.com/first-tweet, which is no longer available, although other websites provide the same service.

[9]Part of Speech is one of the grammatical groups, such as noun, verb, and adjective, into which words are divided depending on their use. Retrieved from: https://dictionary.cambridge.org/dictionary/english/part-of-speech

family subscript to save on notation):

$$\text{Positiveness}_t = \ln\left(\frac{1 + M_t^p}{1 + M_t^n}\right), \tag{3}$$

where $M_t^p(M_t^n)$ is the weighted count of positive (negative) tweets of that family in one month:

$$M_t^p = \sum_{k \in \mathcal{D}(t)} w_k^p x_k^p, \;\; M_t^n = \sum_{k \in \mathcal{D}(t)} w_k^n x_k^n, \tag{4}$$

where $D(t)$ is a monthly time interval, x_k^p (x_k^n) is an indicator variable that takes the value of 1 if tweet k at time t is positive (negative), and w_k^p (w_k^n) is the confidence in the tweet's positive (negative) label given the level of agreement among all classifiers for a particular tweet as in equation (2). Our measure of Positiveness is closely related to that employed by Antweiler and Frank (2004), but is more appropriate for handling Twitter accounts with zero tweets. Figure 1 displays the total number of tweets across all mutual fund families, as well as the weighted count of positive and negative tweets per month. The figure shows a sharp increase in Twitter usage by mutual fund families over the sample period, with a peak in 2016. As expected, positive tweets predominate over tweets classified as negative.

Out of 939 fund families in the final CRSP sample, 284 fund families tweet at least once during the sample period. This is the subsample we use in most of our analysis. Across all months in our sample, fund families that use Twitter collectively manage 3.2 USD trillion on average, which corresponds to 66% of the total assets of non-index U.S. diversified equity funds. Moreover, they operate 6,417 different funds on average, which is 54% of the total number of funds. To understand how this subsample differs from the rest, Table 1 presents descriptive statistics of both fund and family characteristics for the Twitter subsample and the full sample. At the fund level, there are no clear differences between funds managed by fund families in the Twitter subsample and funds in the entire sample. However, at the family level differences between fund families in both samples become more evident. Fund families in the Twitter subsample are on average older, manage more assets, more funds, and funds in more different investment categories.

We explore the content of the tweets in our sample using a Latent Dirichlet Allocation (LDA) topic modelling algorithm (Blei, Ng, and Jordan 2003). Under the LDA algorithm, tweets are represented as a mixture of topics, and a topic can be understood as a distribution of words (or lemmas). The LDA is a hierarchical Bayesian factor model for discrete data that has shown success in reducing the dimensionality of textual data in economic applications (Hansen, McMahon, and Prat 2018; Bandiera, Prat, Hansen, and Sadun 2020). We estimate the LDA algorithm with three topics over the entire sample of tweets after removing common stop-words in tweets, lemmatizing the remaining words, and considering common bigrams in the sample. We identify that the three topics estimated correspond to three different types of communication: i) *Customer Service*, ii) *Market Commentary*, and iii) *Financial Advice*. Figures 2, 3, and 4 present wordclouds of the three topics estimated in which the size of the word is proportional to its relevance in the topic. We label the first topic as *Customer Service* as the words *please, help, account, thank, sorry, number, hi* and *dm* (Direct Message) are the most common within the topic. We label the second topic as *Market Commentary* as the words *market, twitter, investment, investor, stock, global, equity, growth, bond* and *rate* are the most common within the topic. Finally, we label the third topic as *Financial Advice* as the words *need, business, help, ipo, retirement, plan, financial* and *work* are the most common within the topic.

To save on notation we define the set of topics as

$$\mathcal{T} = \{\text{Customer Service}, \text{Market Commentary}, \text{Financial Advice}\} \tag{5}$$

The LDA algorithm provides weights (or coordinates) across these three topics such that each tweet can be represented with the tuple

$$\theta = (\theta^j)_{j\in\mathcal{T}} \tag{6}$$

such that $\theta^j \geq 0\ \forall j \in \mathcal{T}$, $\sum_{j\in\mathcal{T}} \theta^j \leq 1$, and the remaining weight $1 - \sum_{j\in\mathcal{T}} \theta^j$ captures potential omitted topics not considered in the analysis, rounding errors in the estimation, and tweets that cannot be explained with the three topics. In our sample the average sum of weights equals 99.67 %. In unreported tests we estimate the model using two and four topics and find that in the first case the two topics resemble the customer service and market commentary identified in the paper, and in the second case the fourth topic has no clear separation from the other three topics. Moreover, we compare the *coherence score* of the trained models, a measure of out-of-sample performance, and conclude that three topics properly explain the distribution of tweets in our sample. Table 2 presents descriptive statistics of weights θ^j. On average, tweets have a larger weight towards the *Customer Service* topic with 40%, followed by a 34% towards *Market Commentary* and a 26% towards *Financial Advice*. We define a tweet k as conveying information about topic j if not all of its coordinates are zero and the weight on topic j is the largest.[10]

$$\tau_{kt}^{j} = \begin{cases} 1, & \text{if } j = \arg\max\{\theta^i\} \text{ and } \max\{\theta^i\} \neq 0 \\ 0, & \text{otherwise} \end{cases} \tag{7}$$

Using this definition we also compute topic specific measures of positiveness aggregating only over tweets classified with each of the three topics.

$$\text{Positiveness}_t^j = \ln\left(\frac{1 + M_t^{pj}}{1 + M_t^{nj}}\right) \; \forall j \in \mathcal{T} \tag{8}$$

where $M_t^{pj} (M_t^{nj})$ is the weighted count of positive (negative) tweets classified with topic j of a family in one month:

$$M_t^{pj} = \sum_{k \in \mathcal{D}(t)} w_k^p x_k^p \tau_{kt}^j, \;\; M_t^{nj} = \sum_{k \in \mathcal{D}(t)} w_k^n x_k^n \tau_{kt}^j, \tag{9}$$

We illustrate how our classification procedure works in practice by examining five tweets posted by asset management companies in our sample. The first example, presented in Figure

[10]We impose this restriction since around 0.3% of tweets cannot be explained by these three topics and have zero weights on all topics. These tweets tend to be short, use images, post urls or contain spelling mistakes which makes it improbable for the LDA algorithm to explain them.

5, corresponds to a post written by Northern Trust which is classified as negative by the six algorithms. The features (*challenge*, *Noun Singular*) together with (*growth*, *Noun Singular*) are informative enough to make all algorithms coincide with the classification. The second tweet presented in Figure 6 written by Amundi is classified as positive with a confidence of 1. In this tweet the features (*ranked*, *Adjective*), and (*top*, *Adjective*) are informative and make all six algorithms coincide with the classification. Finally we provide examples of the topic modelling of tweets. Figure 7 provides an example of a positive tweet classified as being related to *Customer Service*, Figure 8 provides an example of a negative tweet classified as being related to *Market Commentary*, and finally Figure 9 provides an example of a positive tweet classified as being related to *Financial Advice* together with their tone classification.

Finally, for some of our analysis we obtain intraday data on ETF prices and volumes from FirstRate Data. Prices are adjusted for both splits and dividends. Using the ticker code of the ETF and focusing on ETFs within fund families that have a Twitter account, we are able to match intraday data for 867 different ETF tickers in our sample.

3. Determinants of Twitter activity by mutual fund families

We start our analysis by investigating the determinants of Twitter usage by fund families. Although social media communication has low explicit costs, the implicit costs are non-trivial. Managing a social media communication strategy requires that social media managers coordinate with the marketing department and senior management in the process of setting goals and engaging with the public. In addition, social managers must create contents, promote the firm's social media presence, and carefully monitor and evaluate the whole process. Since such costs are likely to have a fixed component, we expect larger asset management firms to be more likely to use social media. To proxy for size, we use the family's total assets under management, the number of funds and the number of different categories (both in logs) in which families offer funds. Considering the potential benefits of social media, we conjecture that younger asset management firms have stronger incentives to use social media to gain visibility. Also, it seems plausible to think that asset management companies are more likely to communicate with investors when their funds have recently experienced better performance. We proxy for fund age and performance as the age of the oldest fund in the family and the funds asset-weighted average CAPM alpha over the previous 12 months of the funds' in the family. In addition to these variables of interest, we control for the characteristics of funds aggregated at the family level: asset-weighted average expense ratio; number of funds in the family that charge loads (in logs); asset-weighted average volatility of fund returns in the previous 12 months, and asset-weighted average flows.

We analyze both the extensive and the intensive margins of fund families' Twitter usage. More specifically, we employ two different dependent variables. The first variable, Twitter, is an indicator that equals one if the fund family has a Twitter account and uses it at least once in our sample period. The second variable, Number of Tweets, is computed for each family and month as the natural logarithm of 1 plus the number of Tweets posted by the family in that month.

We first estimate a cross-sectional linear probability model with the Twitter indicator as the dependent variable using the full sample. Explanatory variables are first computed for each family and month. We then compute their time-series means for each family. Estimation results are presented in columns (1)-(3) of Table 3. As expected, all three proxies for family size are positively and significantly associated with the family's presence in Twitter. In contrast, family age is not associated with having a Twitter account. Better fund performance and lower return volatility are weakly associated with a higher probability of using Twitter. There is no other significant association between fund characteristics and a Twitter account.

We then regress Number of Tweets on the same set of explanatory variables as in the previous regression, but defined at the family-month level, and lagged one month with respect to the dependent variable. In this case, we naturally restrict the sample to families with Twitter= 1. We include family and time fixed effects and compute robust standard errors clustered at the month level. Estimation results are presented in columns (4)-(6) of Table 3. Conditional on having a Twitter account, both younger and larger firms tend to tweet more frequently. As conjectured, families managing funds with higher recent performance tweet more actively. The number of tweets is also higher when the family manages less expensive funds, less risky funds, and funds with lower recent flows. One possible interpretation is that these funds are less salient and social media helps asset management companies gain visibility for them.

The results in this section suggest that economies of scale are a key determinant of social media usage by asset managers. Conditional on having presence on Twitter, its usage appears to respond not only to cost considerations but also to the potential benefits of social media for asset management firms: gaining visibility for younger firms and funds with better performance, less risk, lower fees, and slower recent growth.

4. Twitter activity and fund flows

In this section, we investigate whether mutual fund investors respond to the Twitter activity of asset management companies. More specifically, we study how the flows of money into the funds of a family are related to the family's number of tweets and tone of those tweets in the previous month, controlling for fund performance and other well-documented flow determinants. To perform the analysis, we restrict the sample to funds of fund families that tweet at least once between January 2009 and December 2020.

Following the literature, we compute net flows to fund i between month t and month $t+1$ as the growth rate in total net assets net of the fund's return:

$$\text{Flows}_{i,t+1} = \frac{\text{TNA}_{i,t+1} - \text{TNA}_{it}(1 + r_{i,t+1})}{\text{TNA}_{it}}, \tag{10}$$

where TNA_{it} is the total net assets of fund i at the end of month t, and $r_{i,t+1}$ is the fund's monthly return. To minimize the impact of outliers - mostly small funds with large percentage of inflows or outflows - we follow the literature and winsorize flows at the 1% level.

Like Sirri and Tufano (1998), we allow for a non-linear flow-performance relationship. To model dependence on performance, we employ two different approaches. First, we define the variable Rank_{it} as the ranking of fund i's CAPM alpha in the 12-month period ending in month t against all other funds in the same Lipper category, normalized to be between $1/N$ (lowest performing fund) and 1 (highest performing fund), where N denotes the number of funds in the corresponding category and month. Second, we use objective-adjusted abnormal return (OAR) as in Ha and Ko (2019). We compute OAR_{it} by standardizing the 12-month CAPM alpha to have zero mean and unit standard deviation across all funds in the same investment category.

For both $\text{Performance}_{it} \in \{\text{Rank}_{it}, \text{OAR}_{it}\}$, we compute the following variables:

$$\begin{aligned}
\text{Low Performance}_{it} &= \min(\text{Performance}_{it}, p20) \\
\text{Mid Performance}_{it} &= \min(\text{Performance}_{it} - \text{Low Performance}_{it}, p80 - p20) \\
\text{High Performance}_{it} &= \text{Performance}_{it} - \text{Mid Performance}_{it} - \text{Low Performance}_{it},
\end{aligned} \tag{11}$$

where $p20$, $p80$ denote the 20th and 80th percentiles, respectively, of either the cross-sectional distribution of performance rank or OAR.

To analyze the link between a fund family's Twitter activity and subsequent flows, we estimate the regression equation:

$$\begin{aligned}
\text{Flows}_{i,t+1} = \gamma_0 &+ \gamma_1 \times \text{Positiveness}_{it} + \gamma_2 \times \text{Number of Tweets}_{it} \\
&+ \gamma_3 \times \text{Low Performance}_{it} + \gamma_4 \times \text{Mid Performance}_{it} + \gamma_5 \times \text{High Performance}_{it} \\
&+ \Gamma \times X_{it} + \delta_{t+1} + \lambda_i + \mu_{cat} + \theta_{fam} + \nu_{i,t+1},
\end{aligned} \tag{12}$$

where $\text{Flows}_{i,t+1}$ is in %. Positiveness is calculated as in equation (3) and the Number of Tweets is the log of the number of tweets of the fund's family (plus 1) or those tweets. Low, Mid, and High Performance are calculated using alternatively Rank or OAR based on the fund's 12-month CAPM alpha, as in equation (11).[11] Following the large literature on the determinants of fund flows, the vector of lagged controls, X_{it}, includes the natural logarithm of the fund's total net assets, the fund's expense ratio, the fund's age (log of months since inception), contemporaneous flows to funds in the same investment category, return volatility in the previous 12 months, flows to the fund in the previous month. We also control for family size (log of assets under management) and family age (age of the family's oldest fund). δ_{t+1}, λ_i, μ_{cat}, and θ_{fam} denote month, fund, investment category, and family fixed effects, respectively.[12] Finally, $\nu_{i,t+1}$ denotes the error term. We estimate equation (12) using pooled OLS and compute robust standard errors clustered at the month, fund family, and month-fund family levels.

Table 4 presents the results. In columns (1)-(3) we report results for OAR. In column (1), we include Number of Tweets as the only measure of Twitter activity. The coefficient on this variable is positive and statistically significant at the 1% level. In column (2), we include only Positiveness. The coefficient on Positiveness is also positive and statistically significant at the 1% level. In column (3) we regress flows on both Number of Tweets and Positiveness and find that the estimated coefficient on Positiveness remains positive and statistically significant at the 5% level, while the coefficient on Number of Tweets becomes smaller and statistically insignificant. These results are consistent with mutual fund investors' responding to fund families' Twitter activity, and particularly to the tone of those tweets. The loss of significance of Number of Tweets is explained by the fact that this variable is positively correlated with Positiveness so in column (1) the coefficient on Number of Tweets captures the association between Flows and Positiveness.[13]

In columns (4) to (6), we report estimation results when we use Rank instead of OAR. The estimated coefficients on Positiveness and Number of Tweets and their standard errors are virtually identical to those in columns (1) to (3), which implies that the association between Twitter activity and subsequent Flows is robust to the alternative way of modelling

[11]As explained above, we choose the CAPM to evaluate fund performance because this is the model that best explains fund flows (Barber, Huang, and Odean, 2016; Berk and van Binsbergen, 2016).

[12]In our sample, some funds change investment categories and fund families through time.

[13]The correlation between Positiveness and Number of Tweets is 74%.

the flow-performance relationship.

Consistently with the literature, we find a convex relation between flows and performance for both ways of modelling non-linearity. Also, fund size, flows to funds in the same category, and volatility are all negatively associated with flows. Flows are persistent as evidenced by the positive and significant coefficient on lagged flows. Finally, younger funds and larger families capture more flows.

In unreported results, we repeat the analysis using the three-factor and four-factor models of Fama and French (1993) and Carhart (1997), respectively, to compute both Rank and OAR. Our conclusions are qualitatively and quantitatively similar.

Therefore, mutual fund investors' decisions appear to be influenced by asset management companies' Twitter activity, and more specifically by the tone of their tweets. Note that this association between fund flows and Positiveness cannot be driven by fund or family time-invariant characteristics that potentially determine both the tone of families' tweets and fund flows. It is not driven either, by larger or younger companies' tendency to tweet more frequently.

In terms of the economic significance of the association between flows and Positiveness, using the estimated coefficient of columns (2) and (5) of Table 4, a one standard deviation increase in Positiveness (0.75 for the Twitter subsample) corresponds to an increase in subsequent flows of 0.06% (= 0.08% $\times$ 0.75), which for the average fund in the Twitter sample corresponds to an increase of USD 693,870 per month (= 0.06%$\times$ USD 1,156.45 million). To gauge the economic impact of Twitter posts' tone for *all* funds in the family, we need to estimate the marginal effect of positiveness on flows at the family level. In Table 5 we estimate a version of equation (12) where all variables are collapsed at the family-month level and we exclude Number of Tweets. In column (1) we use OAR and include time fixed effects but not family fixed effects.[14] The estimated coefficient on Positiveness$_{it}$, 0.10%, is statistically significant at the 1% level.[15] In column (2) we include also fund family fixed

[14]Naturally, the regression equation does not include fund fixed effects or investment objective fixed effects.

[15]One possible reason why the estimated increase in percentage flows is larger for the average family than for the average individual fund is that families with fewer funds, which are underrepresented in fund-level regressions, benefit more from positive tweets.

effects and find a slightly larger estimated coefficient on Positiveness, 0.12%. Using this number, a one standard deviation increase in Positiveness$_{it}$ is associated with an increase of 0.09% (=0.12% $\times$ 0.75) in family flows, which given the average assets under management per family of USD 8.05 billion represents an increase in assets of USD 7.25 million. In columns (3) and (4) we use Rank to model the flow-performance relationship and find a slightly stronger association between Positiveness and Flows, both with and without family fixed effects.

Although these findings suggest that asset management companies' social media activity influences investors' decisions, we also consider the possibility that their Tweets simply reflect publicly available information that impacts fund flows and is not included in the regression. To explore this possibility, we reestimate the flow regression equation controlling for the following information: i) a change in the fund's management company in the previous month; ii) a change in the fund's portfolio manager in the previous month; iii) the number of tweets by third parties that mention the fund family in the previous month; iv) the Positiveness of those third-party tweets; and v) the presence of funds in the family with stellar performance, as proxied by the fraction of funds in the family with monthly CAPM alpha in the top 5% of their investment category in the previous month. Changes in the fund's management could be perceived by investors as a positive signal of future returns (Khorana, 2001). Third-party social media mentions of the family proxy for any news that affects the fund family and that is disseminated through the same channel as the family's tweets. Finally, the presence of funds with stellar performance in the family has been documented to attract disproportionate flows and generate positive spillover effects (Nanda, Wang, and Zheng, 2004).

Results in Table 6 indicate that the association between Positiveness and fund flows remains positive, similar in magnitude, and statistically significant at the 1% level after controlling for changes in management, family-related news or stellar performance. Interestingly, a change in the fund's manager and a larger number of external tweets are followed by lower flows. As expected, the number of funds with stellar performance is strongly associated with larger subsequent flows.

To end this section, we address a potential sample-selection bias in our tests due to the fact that we only include in the flow regressions those fund families that tweet at least once in our sample. In particular, we perform a two-step Heckman procedure where in the first step, we estimate a probit model of a fund belonging in the sample with fund and fund family characteristics, and then include the Mill's ratio as an explanatory variable in the second step. Table 7 reports the results of this estimation. Column (1) contains the results of the probability model in which performance is modelled linearly. Columns (2) and (4) contain the results of the base specification used in the paper as a comparison, and columns (3) and (5) report the results of estimating the same model including the Mill's ratio obtained from column (1). We find that the coefficient of positiveness and R^2 of both specifications remains unchanged which alleviates the concern of a potential sample-selection bias in our results.

5. Testing the persuasion hypothesis

While the results in the previous section are suggestive that social media activity attracts investors' money, they do not reveal the mechanism. In this section, we address this issue, and more specifically, we test for two different implications of the persuasion hypothesis.

We start by investigating which types of tweets drive the association between tone and flows. Through this analysis we intend to rule out the possibility that asset management companies use social media simply as an additional channel for providing customer service and this service. To the extent that investors value customer service, they could decide to invest with families that use Twitter to help their clients. That explanation is plausible but inconsistent with the persuasion hypothesis. In contrast, tweets related to *Financial Advice* and *Market Commentary* are more likely to exert more influence on investors' views of their products and their future performance (Gennaioli, Shleifer, and Vishny, 2015). To find out which type of tweets attracts investors' flows, we rerun the flow regressions with the measure of Positiveness calculated only for tweets of one specific topic at a time. To facilitate the comparison across specifications, we standardize the Positiveness variable within each topic.

Table 8 reports the results. We find that for all three topics, a more positive tone of tweets is positively associated with subsequent flows. However, the coefficients for Positiveness of tweets related to *Market Commentary* and *Financial Advice* are larger than for Positiveness of tweets related to *Customer Service*. Also, they are statistically significant at the 1% level, while the estimated coefficient for Positiveness related to customer service is only significant at the 10% level. Therefore, the positive link between social media activity of asset management companies and flows is not just the consequence of investors' appreciating the value of customer service provided through social media.

Second, persuasion requires that the message conveyed to investors be adapted to different circumstances. For instance, in times of low investor sentiment, asset managers will want to associate their products to safety and trust, whereas in times of high investor sentiment, asset managers will favor the idea of seizing investment opportunities (Mullainathan, Schwartzstein, and Shleifer, 2008). Therefore, it seems plausible that in the former case, successful persuasion will lead investors towards safer products such as money market or bond funds. In contrast, in times of high investor sentiment, persuasion will attempt to drive investors' money to riskier products, such as equity funds. Since our sample consists of only equity funds, we would expect persuasion to be more effective in times of high investor sentiment. To test this hypothesis, in Table 9, we repeat the analysis including an interaction of Positiveness with the two sentiment measures of Baker and Wurgler (2006).[16] For both measures, the interaction term is statistically significant at the 1% level. It is also economically significant, a one standard deviation increase in aggregate sentiment from its unconditional mean is associated with an increase in the effect of positiveness on flows from 0.02 to 0.096 for the non-orthogonal sentiment measure, and from 0.013 to 0.10 for the orthogonal version.

[16] We download from Jeffrey Wurgler's website a measure of investor sentiment based on the first principal component of five sentiment proxies as well as a different measure, where each of the proxies has first been orthogonalized with respect to a set of six macroeconomic indicators.

6. Identifying the effect of fund families' tweets

The findings reported in the previous two sections are consistent with asset management companies influencing mutual fund investor decisions through persuasion. Although in our tests we control for a variety of observable fund flow determinants as well as unobservable time-invariant fund and fund family characteristics, we cannot discard that tweets respond to unobservable news that, in turn, predict fund flows. In this section we combine a high-frequency test with intraday ETF price data to identify the effect of asset management companies tweets on investor decisions.

The analysis is based on the idea that asset management companies' positive tweets increase the demand not just for their open-end mutual funds but also for their ETFs. Since ETF shares are traded in an exchange, we expect a family's positive tweets to put upward pressure on the price of its ETFs. The main advantage of using ETFs over open-end mutual funds for the purpose of this analysis is that ETFs are traded intradaily, which allows us to identify the effect of tweeting activity on prices using the high-frequency methodology of Bianchi, Gomez-Cram, Kind, and Kung (2019) and Bianchi, Gomez-Cram, and Kung (2021). More specifically, we study whether an ETF's price changes in a very short time window around the time when the ETF's family tweets. As explained by those authors, this high-frequency approach allows for clean identification of the effect of tweets on prices with the identifying assumption being that over such a short window of time no other relevant information affecting the security's price is released.

To perform the test, we follow closely Bianchi, Gomez-Cram, and Kung (2021). First, we use the timestamp in Eastern Time of every tweet in our sample to determine two 10-minute windows: one from 15 to 5 minutes *before* the tweet and another one from 5 to 15 minutes *after* the tweet. Second, for each one of the family's ETFs, we identify the price of the last trade in the [-15,-5] minute window and the price of the first trade in the [+5,+15] minute window as depicted below:[17]

[17] We exclude the observation from the analysis if either the last price in the window before the tweet or the first price in the window after the tweet have zero associated volume in our dataset

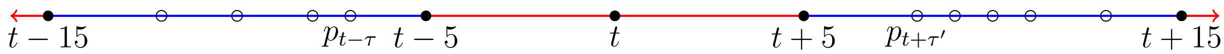

Finally, we compute the change in log prices and regress it on the tweet's Positiveness.[18] To maintain consistency with the rest of the paper, Positiveness is defined at the tweet level as a variable that takes the values 1, 0, and -1, depending on whether the tweet is classified as positive, neutral or negative.[19] Following Bianchi, Gomez-Cram, Kind, and Kung (2019) and Bianchi, Gomez-Cram, and Kung (2021) we include fund fixed effects. That is, we estimate the regression equation:

$$\Delta p_{i,t} = a + b \times \text{Positiveness}_{i,t} + \delta_i + \epsilon_{i,t},$$

where $\Delta p_{i,t}$ is the change in ETF i's log price between the last trade of the [-15,-5] minute window and the first trade of the [+5,+15] minute window around a tweet by i's family taking place at time t.

Note that under the persuasion hypothesis, the increase in demand following tweets with a positive tone is purely driven by investor sentiment about the family's products and not by fundamentals. Consequently, we expect the effect on prices to be temporary, as any price deviations from fundamentals (the ETF's net asset value) will be arbitraged away by authorized participants. To investigate whether price changes are short-lived, we also consider 10-minute windows that start later after the tweet. If the price change is temporary, it should fade as we look further into the future. If, on the other hand, the price increase is permanent, we should observe it even at longer horizons. More specifically, we consider four different 10-minute windows starting at $t + 15$, $t + 25$ and $t + 35$ minutes and recalculate $\Delta p_{i,t}$ accordingly.

Estimation results are reported in Table 10. In column (1) we study the effect of family tweets on its ETF share prices between 5 minutes before the tweet and 5 minutes after each tweet. The estimated coefficient is positive and statistically significant at the 1% significance

[18] We winsorize price changes at the 1% level.

[19] As a robustness test we have multiplied this variable by our measure of confidence and obtained almost identical results (unreported).

level, which indicates that family tweets with a positive tone push up the prices of their ETFs relative to neutral tweets. Based on the estimated coefficient, a positive tweet increases ETF share prices by 5 basis points. In column (2) we move forward the window after the tweet by 10 minutes and estimate an even larger and statistically significant impact of 11 basis points. Such large effect could be explained by the fact that we are giving more time for investors to react. Also, the impact of one tweet is likely amplified over time through retweets and cross-posting in social media. When we study price changes 25 minutes after the tweet, column (3), the estimated coefficient goes down to a marginally significant 5%. Finally, the estimated coefficient in column (4) is zero, which suggests that the effect of a tweet is completely reversed 35 minutes after the tweet.

The results in this section give further credence to the hypotheses that asset management companies' tweets affect fund investor decisions and that this effect is not explained by the release of fundamental information.

7. Alternative hypotheses

In this section, we consider and test two alternative explanations to the persuasion hypothesis. Both explanations are built on the possibility that asset management companies use social media to alleviate information asymmetries.

Our first hypothesis is that asset management companies use social media to reduce search costs for investors. This can be achieved by facilitating investors' access to information that is already available but difficult to locate for investors, such as information about fund offerings, fees, or past performance. Building on the work of Sirri and Tufano (1998), Hortaçsu and Syverson (2004), and Huang, Wei, and Yan (2007), we posit that a reduction in search costs has two potential effects. First, it increases inflows of money from *new* investors, as the benefits of learning about a family's funds outweight the costs for a larger number of investors. However, a reduction in search costs does not affect redemption decisions by *existing* investors, since these investors have already paid the search cost. Second, a reduction in search costs decreases the convexity of the flow-performance relation. When search costs are high, investors pay the cost to learn only about funds with extreme outperformance. As search costs decrease, it becomes optimal for investors to consider a larger set of funds with good performance. As a consequence, flows become less sensitive to recent performance in the high-performance region and the overall flow-performance relation is less convex (Huang, Wei, and Yan, 2007).

To investigate the search-cost hypothesis, we first test whether the tweets of asset management companies increase net flows by increasing inflows or reducing outflows. More specifically, we define Inflows and Outflows for fund i in month $t+1$ as:

$$\text{Inflows}_{i,t+1} = \frac{\text{New Sales}_{i,t+1}}{\text{TNA}_{i,t}}$$
$$\text{Outflows}_{i,t+1} = \frac{\text{Redeemed Cash}_{i,t+1}}{\text{TNA}_{i,t}}$$

As argued by Ha and Ko (2019) inflows and outflows are simultaneously determined by investors' rebalancing strategies. To tackle the mutual dependence between inflows and outflows we follow closely the methodology proposed by Ha and Ko (2019). In particular, we first estimate the following OLS regressions for inflows and outflows separately:

$$\text{Inflows}_{i,t+1} = a + \sum_{s=0}^{11} b_s \times \text{Inflows}_{i,t-s} + c \times X_{it} + \nu_{i,t+1},$$
$$\text{Outflows}_{i,t+1} = a + \sum_{s=0}^{11} b_s \times \text{Outflows}_{i,t-s} + c \times X_{it} + \nu_{i,t+1},$$

where X_{it} contains the same controls used in the flow regressions. We then regress inflows and outflows on the fitted values on outflows and inflows, respectively, and keep the residuals. Finally, residual inflows and residual outflows are regressed on Positiveness and performance (Rank and OAR).

In Table 11, we show estimation results for inflows in columns (1) and (2) and for outflows in columns (3) and (4). The coefficient of inflows on Positiveness is positive and significant at the 1% level for both OAR and Rank. The magnitude of the association is at least as large as that estimated for net flows. As for the results in columns (3) and (4) we find a negative and statistically significant association between Positiveness and outflows. These results indicate that social media communications appear to increase net flows not only by fostering purchases of new shares but also by deterring investors from redeeming old shares, which is not consistent with the search-cost hypothesis.

To investigate whether asset management companies' tweets reduce the sensitivity of flows to performance in the good performance region, in Table 12 we extend the specification of equation (12) to include interactions of Positiveness with performance in each of the three performance regions. Columns (1) and (3) present the base specification without interactions, and columns (2) and (4) present the results with interactions between Positiveness and performance. We find that positiveness increases the sensitivity of flows to performance in the low-performance region. Using OAR to measure performance, a one standard deviation increase in Positiveness increases the slope of flows with respect to performance in the low performance region by 0.06 (=0.08 $\times$ 0.75). This is a small fraction of the estimated slope in column (1) for that region (0.330). More importantly, Positiveness has no significant impact on the sensitivity of flows to performance in the high performance region. Therefore, Positiveness does not seem to induce investors to consider funds with good but not extremely good recent performance, as we would expect if it helped reduce search costs for investors.

A second alternative explanation for the positive link between social media activity and flows is that asset management companies use social media to convey to investors performance-relevant information that is *not publicly available*. More specifically, the model of Dumitrescu and Gil-Bazo (2016) of strategic communication by asset managers predicts that asset management companies will choose to communicate only privately observed information that is favorable for future fund performance. Since in equilibrium, such communications are truthful, not previously known, and favorable to the fund, they should have a positive impact on flows. But such communications should also possess predictive ability with respect to future fund performance beyond and above public information. To test this prediction of the theory, we regress one-month ahead performance on Positiveness while controlling for past performance and fund and family characteristics that have been documented in the literature to predict performance. One difficulty that arises with this test is the fact that net performance is partially determined by investors' reaction. In particular, if there are diseconomies of scale in asset management, fund performance will deteriorate as money flows into funds that are expected to outperform (Berk and Green, 2004). To account for that possibility, we control for lagged assets under management as well as recent fund flows. Therefore, we estimate the regression equation:

$$\hat{\alpha}_{i,t+1} = \rho_0 + \rho_1 \times \alpha_{i,t-3:t} + \rho_2 \times \text{Positiveness}_{i,t} + P \times X_{it} + \delta_{t+1} + \lambda_i + \mu_{cat} + \theta_{fam} + \nu_{i,t+1}, \tag{13}$$

where $\alpha_{i,t-3:t}$ is the fund's abnormal return in the previous three months. X_{it} is a vector of control variables that includes: fund size; expense ratio; flows (in month t); portfolio turnover; 12-month return volatility; an indicator variable that equals one if the fund charges loads; fund age; family size and family age. δ_{t+1}, λ_i, μ_{cat}, and θ_{fam} denote time, fund, investment category, and family fixed effects, respectively. Standard errors are clustered at the fund, month, and fund-month levels.

Table 13 shows the estimation results. In columns (1) and (2) we use CAPM alphas both as the dependent variable and as a control (in this case, measured over the previous three months), whereas in columns (3) and (4) we use Carhart's (1997) four-factor alphas. In all cases, we include time fixed effects to account for market-wide changes that affect funds' risk-adjusted performance. In columns (1) and (3) we do not include fund fixed effects, so we are effectively testing whether positive tweets predict funds' future outperformance *relative to their peers*. In columns (2) and (4), we include fund fixed effects, so we are testing whether funds perform better in the periods following more positive tweets *relative to other periods*. The coefficients on Positiveness are negative and insignificant in all specifications. This implies that funds in families that post more positive tweets do not outperform in the following month either in the cross-section or in the time-series. In unreported results, we repeat our analysis by aggregating performance and control variables at the fund family level and reach identical conclusions.

These results do not support the alternative explanation that positive tweets convey private information that is relevant for future fund performance beyond and above observable fund characteristics.

8. Conclusions

Social media provides asset managers with a powerful tool to communicate with investors avoiding the constraints of traditional mandatory disclosures. We show that asset management companies' social media activity and particularly, the tone of their posts, predicts subsequent flows of money to their equity funds, not explained by performance or fund characteristics. Consistently with the persuasion hypothesis, Tweets with a positive tone attract more money from investors when they contain financial advice and market views and when investors perceive more opportunities in capital markets. That is, social media activity works best precisely when we would expect persuasion to be most effective. A high-frequency test that exploits intraday ETF trading data provides support for the hypothesis that asset management companies' tweets have an impact on investor decisions through social media. The effect is short-lived, consistently with tweets not conveying fundamental information. Our tests strongly reject the alternative hypothesis that tweets by asset management companies reduce information asymmetries by lowering search costs or conveying privately observed information to investors.

Clients of mutual fund management firms could benefit a great deal from the enhanced, more frequent, and easier-to-access information that social media can provide. Instead, the results of our paper suggest that asset management companies use social media to influence investors' perception of the value of their asset management services and increase assets under management. In other words, in the market for mutual funds, the incentives to persuade investors seem to dominate the incentives to inform investors.

REFERENCES

Abramowitz, Michael J. Stop the manipulation of democracy online. *New York Times*, December 11, 2017.

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

Aral, Sinan and Eckles, Dean. Protecting elections from social media manipulation. *Science*, 365(6456):858–861, 2019.

Baker, Malcom and Wurgler, Jeffrey. Investor sentiment and the cross-section of stock returns. *The Journal of Finance*, 61(4):1645–1680, 2006.

Bandiera, Oriana, Prat, Andrea, Hansen, Stephen, and Sadun, Raffaella. CEO Behavior and Firm Performance. *Journal of Political Economy*, 128(4):1325–1369, 2020.

Barber, Brad M, Huang, Xing, and Odean, Terrance. Which factors matter to investors? Evidence from mutual fund flows. *Review of Financial Studies*, 29(10):2600–2642, 2016.

Berk, Jonathan B. and Green, Richard. Mutual fund flows and performance in rational markets. *Journal of Political Economy*, 112(6):1269–1295, 2004.

Berk, Jonathan B. and van Binsbergen, Jules. Assessing asset pricing models using revealed preference. *Journal of Financial Economics*, 119(1):1–23, 2016.

Bertomeu, Jeremy and Marinovic, Iván. A theory of hard and soft information. *The Accounting Review*, 91(1):1–20, 2016.

Bianchi, Francesco, Gomez-Cram, Roberto, Kind, Thilo, and Kung, Howard. Threats to central bank independence: High-frequency identification with twitter. *Working Paper*, 2019.

Bianchi, Francesco, Gomez-Cram, Roberto, and Kung, Howard. Using social media to identify the effects of congressional viewpoints on asset prices. *Working Paper*, 2021.

Blankespoor, Elizabeth, Miller, Gregory S, and White, Hal D. The role of dissemination in market liquidity: Evidence from firms' use of Twitter™. *The Accounting Review*, 89(1): 79–112, 2014.

Blei, David M., Ng, Andrew Y., and Jordan, Michael I. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.

Bollen, Johan, Mao, Huina, and Zeng, Xiaojun. Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1):1–8, 2011.

Carhart, Mark. On persistence in mutual fund performance. *Journal of Finance*, 52(1): 57–82, 1997.

Christoffersen, Susan E. K., Evans, Richard, and Musto, David K. What Do Consumers' Fund Flows Maximize? Evidence from Their Brokers' Incentives. *Journal of Finance*, 68 (1):201–235, February 2013.

Cornelli, Francesca, Kominek, Zbigniew, and Ljungqvist, Alexander. Monitoring managers: Does it matter? *Journal of Finance*, 68(2):431–481, 2013.

Cronqvist, Henrik. Advertising and portfolio choice. *Working Paper*, 2006.

Dumitrescu, Ariadna and Gil-Bazo, Javier. Information and investment under uncertainty. *Economics Letters*, 148(C):17–22, 2016.

Dye, Ronald A. and Sridhar, S. Reliability-relevance trade-offs and the efficiency of aggregation. *Journal of Accounting Research*, 42(1):51–88, 2004.

Evans, Richard B and Sun, Yang. Models or stars: The role of asset pricing models and heuristics in investor risk adjustment. *The Review of Financial Studies*, 2020.

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

Gallaher, Steven T, Kaniel, Ron, and Starks, Laura T. Advertising and mutual funds: From families to individual funds. *Working Paper*, 2015.

Gennaioli, Nicola, Shleifer, Andrei, and Vishny, Robert. Money doctors. *Journal of Finance*, 70(1):91–114, 2015.

Gholampour, Vahid and van Wincoop, Eric. What can we learn from euro-dollar tweets? *Working Paper*, 2017.

Gil-Bazo, Javier and Ruiz-Verdú, Pablo. The relation between price and performance in the mutual fund industry. *Journal of Finance*, 64(5):2153–2183, 2009.

Gu, Chen and Kurov, Alexander. Informational role of social media: Evidence from twitter sentiment. *Journal of Banking and Finance*, 121:105969, 2020.

Ha, Yeonjeong and Ko, Kwangsoo. Misspecifications in the fund flow-performance relationship. *Journal of Financial Intermediation*, 38:69–81, 2019.

Hansen, Stephen, McMahon, Michael, and Prat, Andrea. Transparency and deliberation within the fomc: A computational linguistics approach*. *The Quarterly Journal of Economics*, 133(2):801–870, 2018.

Hillert, Alexander, Niessen-Ruenzi, Alexandra, and Ruenzi, Stefan. Mutual fund shareholder letters: Flows, performance, and managerial behavior. *Working Paper*, 2018.

Hortaçsu, Ali and Syverson, Chad. Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds. *Quarterly Journal of Economics*, 119(2):403–456, 2004.

Huang, Jennifer, Wei, Kelsey D, and Yan, Hong. Participation costs and the sensitivity of fund flows to past performance. *Journal of Finance*, 62(3):1273–1311, 2007.

Jain, Prem C. and Wu, Joanna Shuang. Truth in mutual fund advertising: Evidence on future performance and fund flows. *Journal of Finance*, 55(2):937–958, 2000.

Khorana, Ajay. Performance changes following top management turnover: Evidence from open-end mutual funds. *Journal of Financial and Quantitative Analysis*, 36(3):371–393, 2001.

Kim, O and Verrecchia, Re. Trading volume and price reactions to public announcements. *Journal of Accounting Research*, 29(2):302–321, 1991.

Loughran, Tim and McDonald, Bill. Textual analysis in accounting and finance: A survey. *Journal of Accounting Research*, 54(4):1187–1230, 2016.

Mullainathan, Sendhil and Shleifer, Andrei. Persuasion in finance. *Working Paper*, 2005.

Mullainathan, Sendhil, Schwartzstein, Joshua, and Shleifer, Andrei. Coarse Thinking and Persuasion. *The Quarterly Journal of Economics*, 123(2):577–619, 05 2008.

Nanda, Vikram, Wang, Z Jay, and Zheng, Lu. Family values and the star phenomenon: Strategies of mutual fund families. *Review of Financial Studies*, 17(3):667–698, 2004.

Pastor, Lubos, Stambaugh, Robert, and Taylor, Lucian A. Scale and skill in active management. *Journal of Financial Economics*, 116(1):23–45, 2015.

Ranco, Gabriele, Aleksovski, Darko, Caldarelli, Guido, Grčar, Miha, and Mozetič, Igor. The effects of twitter sentiment on stock price returns. *PloS one*, 10(9):e0138441, 2015.

Sirri, Erik R. and Tufano, Peter. Costly search and mutual fund flows. *Journal of Finance*, 53(5):1589–1622, 1998.

You, Wanhai, Guo, Yawei, and Peng, Cheng. Twitter's daily happiness sentiment and the predictability of stock returns. *Finance Research Letters*, 23:58–64, 2017.

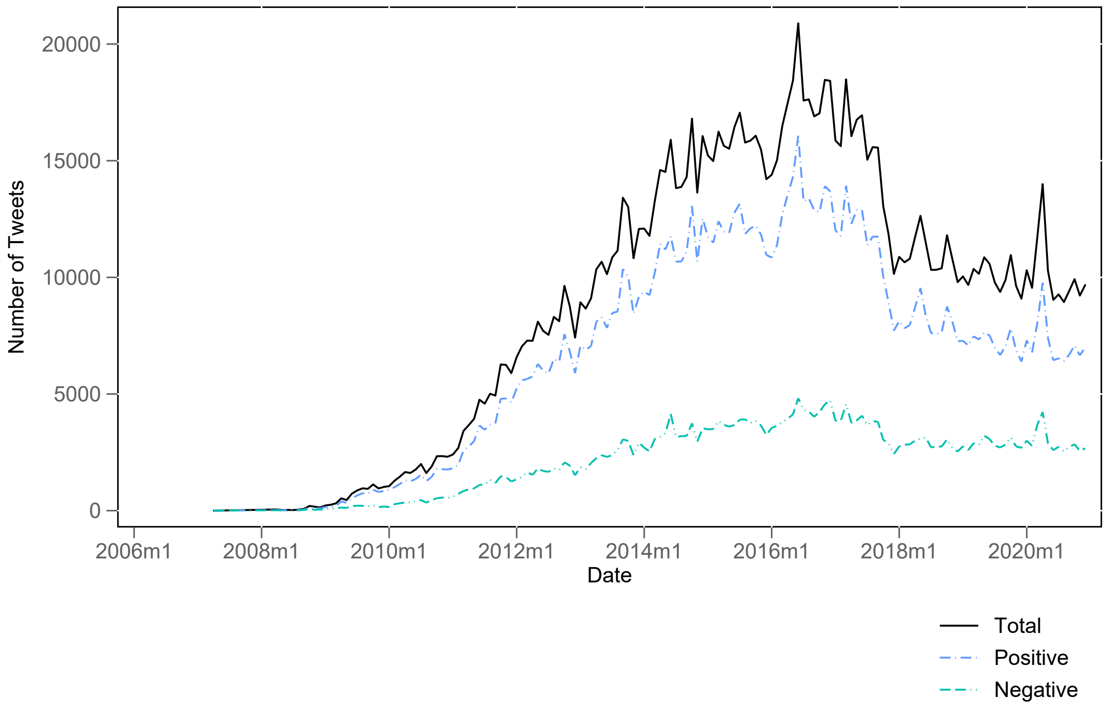

Figure 1: Evolution of tweets by fund families through time. The figure shows the number of tweets by fund families per month. The solid black line shows the total number of tweets obtained based on the fund family identifier mgmt_cd for the entire CRSP database. The dot-dashed line and the dashed line represent out of the entire sample the number of tweets classified as positive and negative, respectively.

Figure 2: First topic estimated using LDA. The topic corresponds to lemmas commonly used in customer service.

Figure 3: Second topic estimated using LDA. The topic corresponds to lemmas commonly used in market commentary.

Figure 4: Third topic estimated using LDA. The topic corresponds to lemmas commonly used in financial advice.

Figure 5: Example of a tweet classified as negative with a confidence of 1. The tweet was written by asset management company Northern Trust (@NorthernTrust) on October 1 2013.

Figure 6: Example of a tweet classified as positive with a confidence of 1. The tweet was written by asset management company Amundi (@Amundi_ENG) on September 21 2012.

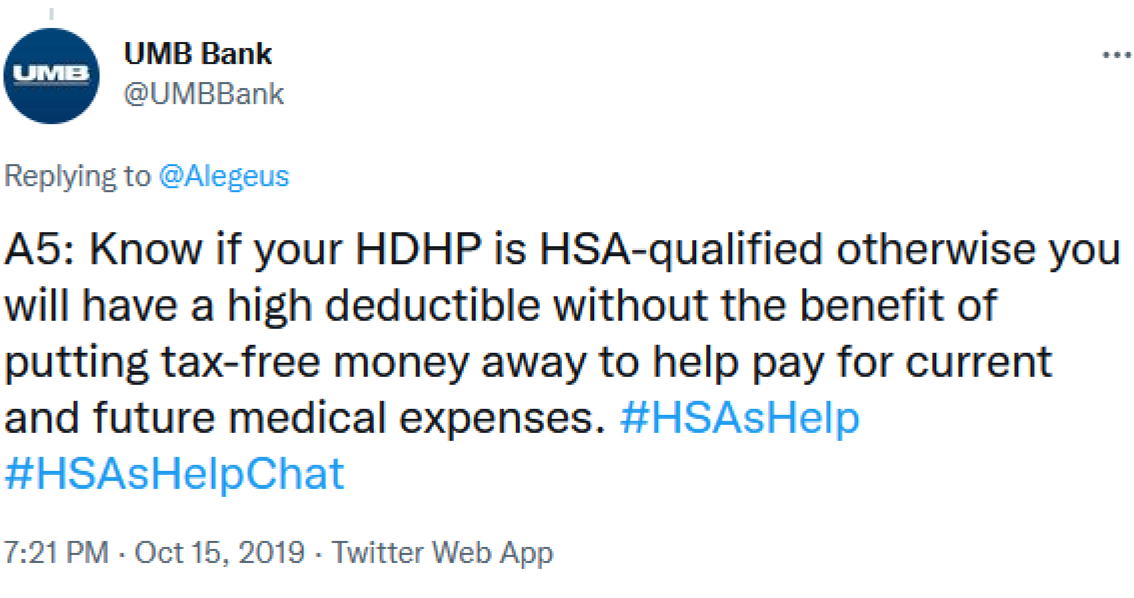

Figure 7: Example of a positive tweet with topic loadings of 0.93 for customer service, 0.048 for market commentary and 0.016 for financial advice.

Figure 8: Example of a negative tweet with topic loadings 0.03 for customer service, 0.93 for market commentary and 0.033 for financial advice.

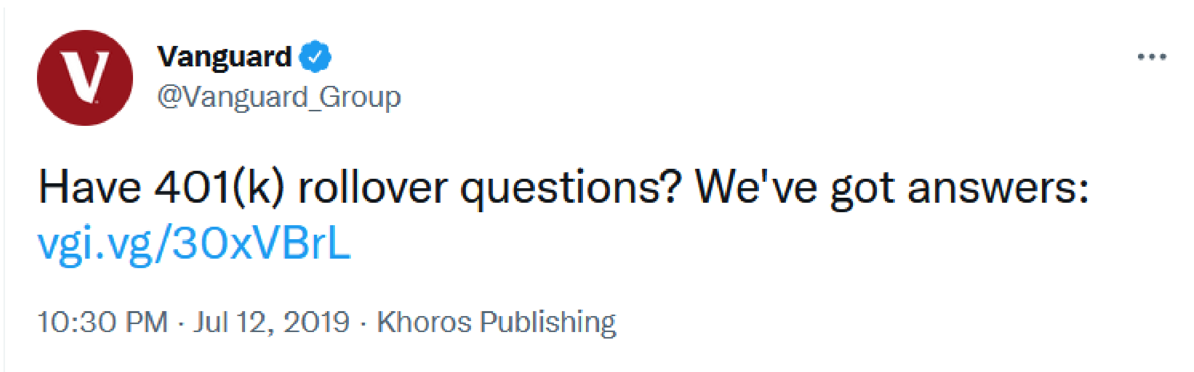

Figure 9: Example of a positive tweet with topic loadings 0.09 for customer service, 0.06 for market commentary and 0.84 for financial advice

Table 1: Descriptive Statistics, Fund and Fund Family characteristics

This table contains summary statistics of fund and fund family characteristics for two samples. The first sample, Twitter Subsample, consists of fund families managing U.S. equity funds that have tweeted at least once from January 2009 to December 2020. The Full Sample includes all families managing U.S. equity funds in the same period. The first set of rows show descriptive statistics for variables computed at the fund-month level, while the second set of rows show descriptive statistics for variables computed at the fund family level. All variables are defined in the Appendix.

	Twitter Subsample						Full Sample					
Variable	mean	median	s.d.	p1	p99	N	mean	median	s.d.	p1	p99	N
Fund Level												
Positiveness	0.78	0.77	0.75	-0.16	2.53	1212755	0.45	0.00	0.69	0.00	2.39	2099985
Performance %	-0.17	-0.09	2.20	-6.01	5.37	1068237	-0.17	-0.10	2.54	-6.22	5.61	1805722
Age (in years)	11.80	10.08	9.33	0.50	47.00	1212755	11.56	10.00	8.94	0.42	42.50	2099985
Flows %	0.32	-0.32	6.69	-17.35	28.82	1199360	0.24	-0.40	6.98	-19.55	31.07	2071800
Expense ratio	0.01	0.01	0.01	0.00	0.02	1212755	0.01	0.01	0.01	0.00	0.02	2099985
Total Net Assets (USD Millions)	1156.45	153.10	5886.72	15.80	17713.70	1212755	884.21	134.00	4650.09	15.70	12861.90	2099985
Average Front Load	0.04	0.04	0.01	0.01	0.05	220112	0.04	0.04	0.01	0.01	0.05	289131
Average Back Load	0.02	0.02	0.01	0.00	0.02	137451	0.02	0.02	0.01	0.00	0.02	240483
Turnover	0.67	0.38	3.04	0.00	5.10	1212755	0.59	0.30	2.54	0.00	4.93	2099985
Flows Category	-0.00	-0.01	0.10	-0.17	0.22	1212667	-0.00	-0.01	0.71	-0.16	0.19	2099785
Flows %	0.32	-0.32	6.69	-17.35	28.82	1199360	0.24	-0.40	6.98	-19.55	31.07	2071800
Fund Family Level												
Positiveness	0.59	0.00	0.75	-0.51	2.51	22379	0.20	0.00	0.52	0.00	2.19	65877
Number of tweets per month	30.30	2.00	70.24	0.00	322.00	22379	10.29	0.00	43.38	0.00	191.00	65877
Age (Years)	30.79	22.50	24.97	2.67	90.08	22379	23.17	18.00	20.12	2.58	87.17	65877
Family Size (USD Billions)	8.05	8.03	2.65	2.93	14.05	22379	6.79	6.53	2.52	2.83	13.08	65877
Number of Funds	52.03	12.00	87.14	1.00	415.00	22379	24.25	4.00	58.15	1.00	271.00	65877
Number of Investment Categories	13.72	6.00	15.92	1.00	69.00	22379	7.49	3.00	11.63	1.00	55.00	65877
CAPM monthly alpha %	-0.00	-0.00	0.01	-0.02	0.02	22379	-0.00	-0.00	0.01	-0.02	0.02	65877
Funds with loads %	25.72	14.94	29.32	0.00	100.00	22379	28.99	12.65	35.48	0.00	100.00	65877
Expense Ratio	1.05	1.00	1.32	0.11	2.21	22379	1.14	1.10	0.85	0.11	2.47	65877
Volatility %	4.01	3.62	2.09	0.84	10.68	22379	4.25	3.80	2.36	0.77	11.55	65877
Unique fund families						284						939

Table 2: Descriptive Statistics, Topic Loadings and Topic Specific Positiveness

This table contains summary statistics of topic loadings for *customer service*, *market commentary*, and *financial advice*, as well as the correlation between positiveness across topics. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	mean	sd	min	max
Customer Service	.397	.369	0	.9799
Market Commentary	.337	.368	0	.9800
Financial Advice	.2638	.352	0	.9800
Observations	1522205			

	Pos. Customer Service	Pos. Market Commentary	Pos. Financial Advice
Pos. Customer Service	1		
Pos. Market Commentary	0.570***	1	
Pos. Financial Advice	0.527***	0.661***	1

Table 3: Determinants of Twitter Activity

This table shows estimation results for regressions of Twitter activity on family characteristics. Columns (1) to (3) display results for cross-sectional regressions of a dummy variable that equals 1 if a fund family has tweeted at least once in our sample, on family characteristics. Columns (4) to (6) provide results of running a regression of the natural logarithm of one plus the number of tweets (# of Tweets) posted by a fund family in a given month on fund family characteristics. In columns (1) to (3) explanatory variables are averaged across time for each family. In columns (4) to (6) the unit of observation is family-month and family characteristics are lagged one month. In columns (1) to (3) robust standard errors are presented in parentheses, while in columns (4) to (6) robust standard errors are clustered at the time level. All variables are defined in the Appendix. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	Twitter			Number of Tweets		
	(1)	(2)	(3)	(4)	(5)	(6)
Family Age	-0.01	-0.01	-0.00	-0.25***	-0.14**	-0.11**
	(0.02)	(0.02)	(0.02)	(0.06)	(0.06)	(0.06)
Family Size	0.05***			0.21***		
	(0.01)			(0.01)		
Number of Funds		0.11***			0.16***	
		(0.02)			(0.03)	
Number of Inv. Categories			0.14***			0.09**
			(0.03)			(0.03)
CAPM alpha	446.58	467.23	465.36	602.71***	639.18***	644.20***
	(295.35)	(290.89)	(293.56)	(185.22)	(201.19)	(202.61)
Expense Ratio	-1.00	-4.16	-3.94	-4.27***	-4.43***	-4.41***
	(3.32)	(3.16)	(3.19)	(0.26)	(0.27)	(0.28)
Funds with loads	0.04	0.04	0.04	-0.06	-0.11**	-0.13**
	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)	(0.05)
Volatility	-1.59*	-1.66*	-1.63*	-7.33***	-7.75***	-7.91***
	(0.83)	(0.85)	(0.85)	(1.18)	(1.22)	(1.23)
Lagged flows	44.25	49.36	46.40	-6.92***	-5.87**	-5.67**
	(35.00)	(32.75)	(33.96)	(2.24)	(2.71)	(2.65)
Observations	814	814	814	22139	22139	22139
Adjusted R^2 (%)	7.88	8.35	7.31	70.39	70.18	70.15
Time FE	No	No	No	Yes	Yes	Yes
Family FE	No	No	No	Yes	Yes	Yes
Estimation	Cross	Cross	Cross	Panel	Panel	Panel
Sample	Full	Full	Full	Twitter	Twitter	Twitter

Table 4: Flows, Positiveness, and Number of Tweets

This table shows estimation results for regressions of mutual fund flows (in %) on Number of Tweets, Positiveness, and control variables. Robust standard errors, clustered at the month, fund family, and month-fund family levels, are shown in parentheses. All variables are defined in the Appendix.***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Positiveness		0.08*** (0.02)	0.06** (0.03)		0.08*** (0.02)	0.06** (0.03)
Number of Tweets	0.03*** (0.01)		0.01 (0.01)	0.03*** (0.01)		0.01 (0.01)
Low OAR	0.33*** (0.03)	0.33*** (0.03)	0.33*** (0.03)			
Mid OAR	0.66*** (0.02)	0.66*** (0.02)	0.66*** (0.02)			
High OAR	0.77*** (0.04)	0.77*** (0.04)	0.77*** (0.04)			
Low Rank				2.72*** (0.23)	2.73*** (0.23)	2.72*** (0.23)
Mid Rank				1.59*** (0.06)	1.59*** (0.06)	1.59*** (0.06)
High Rank				8.74*** (0.46)	8.75*** (0.46)	8.75*** (0.46)
Size	-0.71*** (0.03)	-0.71*** (0.03)	-0.71*** (0.03)	-0.71*** (0.03)	-0.71*** (0.03)	-0.71*** (0.03)
Flows to the same category	-2.27*** (0.26)	-2.28*** (0.26)	-2.28*** (0.26)	-2.27*** (0.26)	-2.27*** (0.26)	-2.27*** (0.26)
Volatility	-11.12*** (1.18)	-11.12*** (1.18)	-11.12*** (1.18)	-9.07*** (1.18)	-9.06*** (1.18)	-9.07*** (1.18)
Expense ratio	1.96 (2.23)	2.18 (2.23)	2.10 (2.23)	2.35 (2.25)	2.57 (2.26)	2.49 (2.25)
Age	-1.94*** (0.11)	-1.95*** (0.11)	-1.95*** (0.11)	-1.98*** (0.11)	-1.98*** (0.11)	-1.98*** (0.11)
Family Size	0.17*** (0.05)	0.16*** (0.05)	0.17*** (0.05)	0.17*** (0.05)	0.17*** (0.05)	0.17*** (0.05)
Family Age	-0.27*** (0.09)	-0.27*** (0.09)	-0.27*** (0.09)	-0.29*** (0.09)	-0.28*** (0.09)	-0.29*** (0.09)
Past Flows	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)
Constant	12.92*** (0.92)	12.98*** (0.92)	12.95*** (0.92)	12.36*** (0.92)	12.42*** (0.92)	12.40*** (0.92)
Observations	455071	455071	455071	455071	455071	455071
Adjusted R^2 (%)	12.28	12.28	12.28	12.22	12.23	12.23
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Inv. Category FE	Yes	Yes	Yes	Yes	Yes	Yes
Fund FE	Yes	Yes	Yes	Yes	Yes	Yes
Fund Family FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 5: Flows and Positiveness (Fund-family level)

This table shows estimation results for regressions of fund family flows (in %) on Positiveness and control variables. Robust standard errors, clustered at the month, fund family, and month-fund family levels, are shown in parentheses. All variables are defined in the Appendix. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	(1)	(2)	(3)	(4)
Positiveness	0.10***	0.12***	0.13***	0.14***
	(0.03)	(0.04)	(0.03)	(0.04)
Low OAR - Family	0.82***	0.70***		
	(0.11)	(0.12)		
Mid OAR - Family	0.27***	0.27***		
	(0.08)	(0.08)		
High OAR - Family	0.69***	0.79***		
	(0.12)	(0.12)		
Low Rank - Family			-6.07***	-5.61***
			(0.80)	(0.83)
Mid Rank - Family			2.41***	2.17***
			(0.22)	(0.22)
High Rank - Family			6.01***	8.63***
			(1.39)	(1.48)
Family Size	-0.01	-0.27***	-0.00	-0.25***
	(0.01)	(0.07)	(0.01)	(0.07)
Volatility - Family	4.24*	-5.51	0.53	-5.55
	(2.36)	(4.04)	(2.32)	(4.01)
Expense Ratio - Family	-19.64**	20.29	-27.80***	27.17
	(9.37)	(23.30)	(9.30)	(23.14)
Age - Family	-0.22***	-0.29**	-0.23***	-0.33**
	(0.04)	(0.13)	(0.04)	(0.14)
Flows Category - Family	0.64**	-0.56	0.55**	-0.56
	(0.25)	(0.59)	(0.25)	(0.58)
Flows - Family(t)	0.31***	0.26***	0.32***	0.26***
	(0.02)	(0.02)	(0.02)	(0.02)
Constant	1.38***	3.57***	1.59***	3.54***
	(0.22)	(0.85)	(0.23)	(0.86)
Observations	22140	22137	22140	22137
Adjusted R^2 (%)	16.34	19.41	16.18	19.36
Time FE	Yes	Yes	Yes	Yes
Fund Family FE	No	Yes	No	Yes

Table 6: Flows, Positiveness and Fundamental Information

This table shows estimation results for regressions of mutual fund flows (in %) on Number of Tweets, Positiveness, variables that capture changes in fundamental information, and control variables. Robust standard errors, clustered at the month, fund family, and month-fund family levels, are shown in parentheses. All variables are defined in the Appendix. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	(1)	(2)	(3)	(4)	(5)
Positiveness	0.07***	0.07***	0.08***	0.07***	0.07***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Change Family	-0.28				
	(0.20)				
Change Manager		-1.11**			
		(0.55)			
Number of external tweets			-0.08*		
			(0.04)		
External Positiveness				-0.01	
				(0.06)	
Fraction Top Funds					0.49***
					(0.15)
Low Rank	2.70***	2.69***	2.70***	2.69***	2.68***
	(0.23)	(0.23)	(0.23)	(0.23)	(0.23)
Mid Rank	1.59***	1.59***	1.59***	1.59***	1.59***
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
High Rank	8.71***	8.70***	8.72***	8.71***	8.66***
	(0.46)	(0.46)	(0.46)	(0.46)	(0.46)
Size	-0.71***	-0.71***	-0.71***	-0.71***	-0.71***
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Flows to the same category	-2.36***	-2.36***	-2.37***	-2.36***	-2.36***
	(0.27)	(0.27)	(0.27)	(0.27)	(0.27)
Volatility	-8.88***	-8.89***	-8.89***	-8.88***	-8.91***
	(1.18)	(1.18)	(1.18)	(1.18)	(1.18)
Expense ratio	2.78	2.77	2.71	2.77	2.74
	(2.29)	(2.29)	(2.29)	(2.29)	(2.28)
Age	-1.98***	-1.98***	-1.98***	-1.98***	-1.97***
	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)
Family Size	0.15***	0.15***	0.15***	0.15***	0.15***
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Family Age	-0.28***	-0.28***	-0.27***	-0.28***	-0.28***
	(0.09)	(0.09)	(0.09)	(0.09)	(0.09)
Past Flows	0.14***	0.14***	0.14***	0.14***	0.14***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	455071	455071	455071	455071	455071
Adjusted R^2 (%)	12.19	12.2	12.19	12.19	12.2
Time FE	Yes	Yes	Yes	Yes	Yes
Fund Family FE	Yes	Yes	Yes	Yes	Yes
Fund FE	Yes	Yes	Yes	Yes	Yes

Table 7: Positiveness and Endogenous Twitter Presence

This table shows estimation results for a two-step Heckman correction procedure. Column 1 estimates a probit model on variable Twitter? which takes the value of one if a fund family tweets at least once in our sample and 0 otherwise. Columns 3-5 presents regressions of mutual fund flows (in %) on Positiveness, the Mill's ratio of Column 1. All variables are defined in the Appendix. Robust standard errors, clustered at the month, fund family, and month-fund family levels, are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	Twitter	Flows(t+1)			
	(1)	(2)	(3)	(4)	(5)
Positiveness		0.08***	0.08***	0.08***	0.08***
		(0.02)	(0.02)	(0.02)	(0.02)
Mills Ratio			0.89***		0.89***
			(0.24)		(0.24)
Low OAR		0.33***	0.33***		
		(0.03)	(0.03)		
Mid OAR		0.66***	0.66***		
		(0.02)	(0.02)		
High OAR		0.77***	0.77***		
		(0.04)	(0.04)		
Low Rank				2.72***	2.72***
				(0.23)	(0.23)
Mid Rank				1.59***	1.59***
				(0.06)	(0.06)
High Rank				8.76***	8.74***
				(0.46)	(0.46)
Size	0.16***	-0.71***	-0.64***	-0.71***	-0.64***
	(0.00)	(0.03)	(0.03)	(0.03)	(0.03)
Volatility	-4.65***	-11.12***	-13.62***	-9.07***	-11.58***
	(0.06)	(1.18)	(1.31)	(1.18)	(1.31)
Expense ratio	81.42***	2.16	11.01**	2.54	11.41**
	(0.29)	(2.22)	(4.79)	(2.25)	(5.14)
Age	-0.28***	-1.95***	-2.05***	-1.98***	-2.08***
	(0.00)	(0.11)	(0.11)	(0.11)	(0.11)
Family Size	-0.02***	0.17***	0.15***	0.17***	0.16***
	(0.00)	(0.05)	(0.05)	(0.05)	(0.05)
Family Age	0.68***	-0.27***	0.05	-0.28***	0.04
	(0.00)	(0.09)	(0.12)	(0.09)	(0.12)
Past Flows	0.00***	0.14***	0.14***	0.14***	0.14***
	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
Flows to the same category		-2.35***	-2.35***	-2.34***	-2.34***
		(0.27)	(0.27)	(0.27)	(0.27)
Performance	0.00				
	(0.00)				
Constant	-3.73***	12.96***	10.67***	12.41***	10.11***
	(0.02)	(0.92)	(1.09)	(0.92)	(1.09)
Observations	945034	455071	455071	455071	455071
Adjusted R^2 (%)	15.76	12.28	12.29	12.23	12.23
Time FE	No	Yes	Yes	Yes	Yes
Inv. Category FE	No	Yes	Yes	Yes	Yes
Fund FE	No	Yes	Yes	Yes	Yes
Fund Family FE	No	Yes	Yes	Yes	Yes

Table 8: Positiveness of Tweets Disaggregated by Topic

This table shows estimation results for regressions of mutual fund flows (in %), on Positiveness computed only for tweets for each topic, and control variables. All variables are defined in the Appendix. Robust standard errors, clustered at the month, fund family, and month-fund family levels, are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	(1)	(2)	(3)
Positiveness Customer Service	0.03*		
	(0.01)		
Positiveness Market Commentary		0.05***	
		(0.01)	
Positiveness Financial Advice			0.04***
			(0.01)
Low OAR	0.33***	0.33***	0.33***
	(0.03)	(0.03)	(0.03)
Mid OAR	0.66***	0.66***	0.66***
	(0.02)	(0.02)	(0.02)
High OAR	0.77***	0.77***	0.77***
	(0.04)	(0.04)	(0.04)
Size	-0.71***	-0.71***	-0.71***
	(0.03)	(0.03)	(0.03)
Flows to the same category	-1.96***	-1.97***	-1.97***
	(0.26)	(0.26)	(0.26)
Volatility	-11.11***	-11.10***	-11.10***
	(1.18)	(1.18)	(1.18)
Expense ratio	2.22	2.22	2.17
	(2.23)	(2.23)	(2.23)
Age	-1.95***	-1.94***	-1.95***
	(0.11)	(0.11)	(0.11)
Family Size	0.17***	0.17***	0.17***
	(0.05)	(0.05)	(0.05)
Family Age	-0.26***	-0.27***	-0.27***
	(0.09)	(0.09)	(0.09)
Past Flows	0.14***	0.14***	0.14***
	(0.01)	(0.01)	(0.01)
Constant	12.99***	12.99***	12.99***
	(0.92)	(0.92)	(0.92)
Observations	455071	455071	455071
Adjusted R^2 (%)	12.27	12.28	12.28
Time FE	Yes	Yes	Yes
Inv. Category FE	Yes	Yes	Yes
Fund FE	Yes	Yes	Yes
Fund Family FE	Yes	Yes	Yes

Table 9: Positiveness and Investor Sentiment

This table shows estimation results for regressions of mutual fund flows (in %) on Number of Tweets, Positiveness, control variables, and two sentiment measures. All variables are defined in the Appendix. Robust standard errors, clustered at the month, fund family, and month-fund family levels, are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	(1)	(2)	(3)
Positiveness	0.08***	0.12***	0.06**
	(0.02)	(0.04)	(0.03)
Sentiment		-0.04	
		(0.10)	
Positiveness × Sentiment		0.38***	
		(0.11)	
Sentiment Orth.			-0.08
			(0.08)
Positiveness × Sentiment Orth.			0.37***
			(0.09)
Low OAR	0.33***	0.41***	0.41***
	(0.03)	(0.03)	(0.03)
Mid OAR	0.66***	0.62***	0.62***
	(0.02)	(0.02)	(0.02)
High OAR	0.77***	0.77***	0.77***
	(0.04)	(0.04)	(0.04)
Size	-0.71***	-0.86***	-0.86***
	(0.03)	(0.04)	(0.04)
Flows to the same category	-2.28***	-1.61***	-1.62***
	(0.26)	(0.23)	(0.23)
Volatility	-11.12***	-11.71***	-11.99***
	(1.18)	(0.94)	(0.91)
Expense ratio	2.18	25.07	24.64
	(2.23)	(17.05)	(17.06)
Age	-1.95***	-2.58***	-2.58***
	(0.11)	(0.09)	(0.09)
Family Size	0.16***	0.21***	0.23***
	(0.05)	(0.06)	(0.06)
Family Age	-0.27***	-0.34***	-0.35***
	(0.09)	(0.10)	(0.10)
Past Flows	0.14***	0.14***	0.14***
	(0.01)	(0.01)	(0.01)
Constant	12.98***	16.82***	16.67***
	(0.92)	(0.90)	(0.89)
Observations	455071	370216	370216
Adjusted R^2 (%)	12.28	12.68	12.68
Time FE	Yes	No	No
Inv. Category FE	Yes	Yes	Yes
Fund FE	Yes	Yes	Yes
Fund Family FE	Yes	Yes	Yes

Table 10: High-Frequency ETF Price Responses to Fund Family Tweets

In all regressions, the dependent variable $\Delta p_{i,t}$ is the log change in ETF i's price between the last trade of the $[-15, -5]$ minute window and the first trade of the $[+5, +15]$, $[+15, +25]$, $[+25, +35]$, or $[+35, +45]$ minute window around the tweet hapenning at time t. $\Delta p_{i,t}$ is winsorized at the 1% level and expressed in basis points. Positiveness is defined at the tweet level and equals 1 if the tweet from ETF i's fund family is positive, 0 if it is neutral and -1 if it is negative. We exclude observations for which the traded volume associated with the last price in the window before the tweet or the first price in the window after the tweet is zero. In all regressions, we include fund fixed effects. Standard errors are reported in parenthesis and clustered at the fund family and day level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

$$\Delta p_{i,t} = a + b \times \text{Positiveness}_{it} + \delta_i + \epsilon_{i,t}$$

	Time Window after the Tweet			
	[+5,+15]	[+15, +25]	[+25, +35]	[+35, +45]
	(1)	(2)	(3)	(4)
Positiveness	0.05***	0.11***	0.05*	0.00
	(0.02)	(0.02)	(0.02)	(0.03)
Constant	-0.17***	-0.23***	-0.13***	-0.01
	(0.02)	(0.02)	(0.03)	(0.03)
Observations	1789392	1756007	1723952	1672237
Adjusted R^2 (%)	.07	.07	.07	.06
Fund FE	Yes	Yes	Yes	Yes

Table 11: Inflows, Outflows, and Positiveness

This table shows estimation results for regressions of residual inflows and residual outflows (both in %) on Positiveness and fund performance. In a first stage (not reported) Inflows (Out- flows) are regressed on 12 lags of the variable and controls: Size, Flows to the same category, Volatility, Expense Ratio, Age, one-month lagged flows to the fund, Family size, and Family age. In a second stage (not reported) we regress Inflows (Outflows) on the fitted values of Ouflows (Inflows) estimated in the first stage. Fitted residuals from the regression are then regressed on Positiveness and the three performance variables. All variables are defined in the Appendix. OLS robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	Inflows		Outflows	
	(1)	(2)	(3)	(4)
Positiveness	0.11***	0.08***	-0.06***	-0.05***
	(0.02)	(0.02)	(0.01)	(0.01)
Low OAR	0.26***		-0.58***	
	(0.03)		(0.04)	
Mid OAR	0.66***		-0.31***	
	(0.03)		(0.02)	
High OAR	0.94***		-0.08**	
	(0.07)		(0.04)	
Low Rank		1.81***		-4.53***
		(0.26)		(0.27)
Mid Rank		1.88***		-0.61***
		(0.09)		(0.06)
High Rank		9.35***		-1.28**
		(1.03)		(0.50)
Observations	44943	44943	44943	44943
Adjusted R^2 (%)	4.43	3.96	2.53	2.43

Table 12: Positiveness and Flow-Performance Sensitivity

This table shows estimation results for regressions of mutual fund flows (in %) on Number of Tweets, Positiveness, and control variables. All variables are defined in the Appendix. Robust standard errors, clustered at the month, fund family, and month-fund family levels, are shown in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	(1)	(2)	(3)	(4)
Positiveness	0.081***	0.123***	0.080***	-0.047
	(0.022)	(0.047)	(0.022)	(0.048)
Low OAR	0.330***	0.270***		
	(0.030)	(0.041)		
Mid OAR	0.658***	0.622***		
	(0.023)	(0.034)		
High OAR	0.772***	0.805***		
	(0.039)	(0.057)		
Positiveness × Low OAR		0.080**		
		(0.036)		
Positiveness × Mid OAR		0.041		
		(0.029)		
Positiveness × High OAR		-0.041		
		(0.053)		
Low Rank			2.724***	2.162***
			(0.227)	(0.316)
Mid Rank			1.590***	1.575***
			(0.059)	(0.087)
High Rank			8.766***	9.222***
			(0.459)	(0.701)
Positiveness × Low Rank				0.712***
				(0.270)
Positiveness × Mid Rank				0.015
				(0.073)
Positiveness × High Rank				-0.526
				(0.621)
Size	-0.706***	-0.708***	-0.712***	-0.714***
	(0.029)	(0.029)	(0.029)	(0.029)
Expense ratio	2.199	2.194	2.582	2.557
	(2.227)	(2.228)	(2.256)	(2.254)
Past Flows	0.142***	0.142***	0.143***	0.143***
	(0.005)	(0.005)	(0.005)	(0.005)
Volatility	-11.122***	-11.163***	-9.069***	-9.087***
	(1.184)	(1.184)	(1.184)	(1.185)
Family Size	0.165***	0.165***	0.171***	0.170***
	(0.046)	(0.046)	(0.046)	(0.046)
Family Age	-0.268***	-0.270***	-0.283***	-0.284***
	(0.091)	(0.091)	(0.091)	(0.091)
Flows to the same category	-2.378***	-2.376***	-2.385***	-2.384***
	(0.270)	(0.270)	(0.270)	(0.270)
Age	-1.950***	-1.946***	-1.985***	-1.981***
	(0.105)	(0.105)	(0.105)	(0.105)
Constant	12.962***	12.945***	12.403***	12.506***
	(0.924)	(0.924)	(0.924)	(0.924)
Observations	455071	455071	455071	455071
Adjusted R^2 (%)	12.29	12.29	12.23	12.23
Time FE	Yes	Yes	Yes	Yes
Investment Category FE	Yes	Yes	Yes	Yes
Fund FE	Yes	Yes	Yes	Yes
Fund Family FE	Yes	Yes	Yes	Yes

Table 13: Positiveness and Future Fund Performance

Note: This table shows estimation results of regressions of fund's monthly alpha (in %) on Positiveness, past performance, and control variables. Robust standard errors clustered at the family, month, and family-month levels are presented in parentheses. All variables are defined in the Appendix. ***, **, and * denote statistical significance at the 1%, 5%, and 10% significance level, respectively.

	Dependent Variable			
	α_{t+1}	α_{t+1}	α^{4F}_{t+1}	α^{4F}_{t+1}
	(1)	(2)	(3)	(4)
Positiveness	-0.000	-0.003	-0.003	-0.006
	(0.006)	(0.007)	(0.005)	(0.006)
$\alpha_{t-3:t}$	0.021***	0.012***		
	(0.003)	(0.003)		
$\alpha^{4F}_{t-3:t}$			0.008**	-0.001
			(0.004)	(0.004)
Size	-0.004*	-0.086***	-0.004**	-0.094***
	(0.002)	(0.006)	(0.002)	(0.005)
Expense ratio	-4.493***	0.913	-4.696***	1.156
	(0.866)	(0.913)	(0.788)	(0.911)
Past Flows	0.000	-0.001**	0.001***	0.000
	(0.001)	(0.001)	(0.000)	(0.000)
Turnover	-0.006***	-0.007*	-0.006***	0.001
	(0.002)	(0.004)	(0.002)	(0.003)
Number of Funds	0.000*	0.000*	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
Category Share	0.194**	-0.002	0.201***	0.067
	(0.098)	(0.275)	(0.076)	(0.228)
% of new funds	0.000	0.000	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Number of tweets	0.004	0.004	-0.001	-0.001
	(0.003)	(0.003)	(0.002)	(0.003)
Volatility	-2.573*	-1.794	-0.428	0.115
	(1.346)	(1.575)	(1.331)	(1.571)
Family Size	-0.057***	-0.027**	-0.074***	-0.041***
	(0.011)	(0.014)	(0.009)	(0.011)
Family Age	-0.024	-0.040*	-0.011	-0.024
	(0.019)	(0.022)	(0.016)	(0.018)
Charges loads	0.001	0.002	0.006	0.007
	(0.007)	(0.042)	(0.006)	(0.036)
Age fund	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Retail	-0.012	-0.002	-0.015*	-0.005
	(0.011)	(0.012)	(0.009)	(0.009)
Constant	0.786***	0.889***	0.797***	0.922***
	(0.157)	(0.196)	(0.129)	(0.160)
Observations	505979	505958	505979	505958
Adjusted R^2 (%)	8.67	8.58	6.18	6.11
Time FE	Yes	Yes	Yes	Yes
Investment Category FE	Yes	Yes	Yes	Yes
Fund FE	No	Yes	No	Yes
Fund Family FE	Yes	Yes	Yes	Yes

Appendix: Variable Definitions, Bag of Words and Part of Speech

A1. Fund Variables

- Flows: $\text{Flows}_{i,t+1} = \frac{\text{TNA}_{i,t+1} - \text{TNA}_{i,t}(1+r_{i,t+1})}{textTNA_{i,t}}$ are net flows to fund i between month t and month $t+1$, defined as the growth rate in total net assets net of the fund's return where TNA_{it} is the total net assets of fund i at the end of month t and $r_{i,t+1}$ is the fund's monthly return. To minimize the impact of outliers - mostly small funds with large percentage of inflows or outflows - we follow the literature and winsorize flows at the 1% level.
- Performance: Performance is the monthly CAPM alpha computed from the asset-weighted average return of all share classes of the fund.
- Age: Age is the number of years (in logs) since the inception of the oldest share class in the fund.
- Expense ratio is the asset-weighted average across all share classes of the fund, expressed in decimal units.
- Total Net Assets is the sum of total net assets of all share classes of the fund.
- Front-end and back-end loads denote the asset-weighted average of the maximum loads across all share classes.
- Turnover denotes the annual turnover of the fund's portfolio.
- Flows to the same category are computed as the percentage of contemporaneous net flows to all funds with the same Lipper class investment objective.
- Funds with loads is the percentage of funds in the family that charge either front-end or back-end load.
- Expense Ratio is the Asset weighted average of the expense ratios of all funds within each fund family.
- Volatility is the 12 month rolling volatility of returns.

- Low, Mid, and High Performance are computed using objective adjusted 12-month CAPM alphas (normalized to have zero average and unit standard deviation across all funds with the same Lipper class investment objective.
- Low, Mid, and High Ranks are computed using ranks based on the 12-month CAPM alpha of each fund under the same Lipper class investment objective.
- Size is the natural logarithm of the total net assets under management of a fund.
- Performance is defined in two ways: The CAPM abnormal return compounded over the last three months $\alpha_{t-3:t}$ and the four-factor abnormal return compounded over the last three months $\alpha_{t-3:t}$.
- Charges Loads is a dummy that equals 1 if the fund charges either front-end or back-end loads.
- Change Family indicates if the fund has changed fund family (measured as mgmt_cd) in the last month.
- Change Manager indicates if the fund has changed manager in the last month.

A2. Fund Family Variables

- Family Flows: Weighted average of the Flows in % across the funds of a family in the sample.
- Family Age: Family Age is the natural logarithm of the age of the oldest fund in the family in years
- Family size is the natural logarithm of all the total net assets managed by the company in the sample in USD millions.
- Family CAPM alpha: The CAPM alpha is the asset-weighted average 12-month CAPM alpha across funds in the family.
- Family Expense Ratio: Expense Ratio is is the asset-weighted average of expense ratios across funds in the family.
- Funds with loads denotes the percentage of funds in the family that charge front-end or back-end loads.

- Volatility is the asset-weighted average of each fund's 12-month rolling volatility of returns.
- Inv. Categories denotes the log of the number of different Lipper investment categories across all funds in the family.
- Number of Tweets is computed as the natural logarithm of one plus the number of tweets posted by the fund's fund family.
- Low, Mid and High OAR are weighted averages across all funds in the sample belonging to the fund family.
- Low, Mid and High Rank are weighted averages across all fund in the sample belonging to the fund family.
- The number of external tweets is computed as the natural logarithm of the number of external mentions retrieved from the fund.
- External positiveness is computed using external tweets.
- Fraction of top funds is the fraction of funds in the fund family that are in the top 5% performance percentile in the month across all funds in the same Lipper category.
- Number of Funds is the total number of funds managed by the family (in logs).
- Fraction retail is the value weighted average of a dummy that takes the value of 1 if a fund is classified as retail.

A3. Bag of Words and Part of Speech

In this section we explain the bag-of-words and part of speech approach used to classify the tweets in our sample into positive, neutral and negative tweets. Since in most of our sample, tweets are limited by the 140-character length restriction, just the presence of words might not be informative enough to classify them.[20] At the same time the *informality* of communications increases the use of words with less defined tonal categories, and provides less informative features to predict labels. To obtain a more informative set of features from tweets, we consider both the appearance of a word as well as its role in the sentence - also

[20]Twitter changed to 240 the character limit of tweets only starting in 2017.

known as Part of Speech (POS).

We use six different algorithms: Naive-Bayes classifier, Multinomial-Naive-Bayes classifier, Bernoulli Naive-Bayes classifier, Stochastic Gradient Descent, Support Vector Machines, and Logistic Regression. We then consider a voting scheme that consists of classifying each tweet with the most voted label among the different algorithms. If three algorithms classify a tweet as positive and three as negative, we consider the tweet to have a neutral tone. The procedure also provides us with a measure of agreement between the algorithms.

We start by extracting two important features of each tweet: Words, and Part of Speech (POS). We start by applying a tokenization based on regular expressions to automatically split the text into words. We proceed by using a POS tagger (an algorithm that tags each word with its more likely POS) to identify the role of each word within the sentence.

To illustrate the procedure we use the following example:

Figure 10: Example of a financial tweet posted by Bloomberg @business on September 27 2017, 14:00.

The first step consists of tokenizing the tweet to isolate its components. We use a special tokenization procedure to account for hyperlinks, emoticons, and punctuation. The tokenization splits the tweet as follows:

Fidelity, Chairman, and, CEO, Abigail, Johnson, discusses, the, struggle, between, active and, passive, investing, https://www.bloomberg.com/...

Once the tweet is tokenized, we apply a POS tagger which applies an optimization algorithm that maximizes the likelihood of tuples of the form $(token, POS)$ to appear in a sentence.

After the POS tagger is applied, the tweet becomes:

$$\underbrace{\text{Fidelity}}_{\text{proper noun, singular}}\ \underbrace{\text{Chairman}}_{\text{proper noun, singular}}\ \underbrace{\text{and}}_{\text{coordinating conjunction}}\ \underbrace{\text{CEO}}_{\text{proper noun, singular}}\ \underbrace{\text{Abigail}}_{\text{proper noun, singular}}$$

$$\underbrace{\text{Johnson}}_{\text{proper noun, singular}}\ \underbrace{\text{discusses}}_{\text{verb, 3rd person, singular, present}}\ \underbrace{\text{the}}_{\text{determiner}}\ \underbrace{\text{struggle}}_{\text{noun singular}}$$

$$\underbrace{\text{between}}_{\text{preposition/subordinating conjunction}}\ \underbrace{\text{active}}_{\text{adjective}}\ \underbrace{\text{and}}_{\text{coordinating conjunction}}\ \underbrace{\text{passive}}_{\text{adjective}}\ \underbrace{\text{investing}}_{\text{noun singular}}$$

$$\underbrace{\text{https://www.bloomberg.com/...}}_{\text{noun singular}}$$

After tokenizing and extracting the POS of every tweet in our database we proceed to extract the most common features. We do this by calculating the frequency of each tuple $(token, POS)$ and select the most common 4000 features.

The machine learning algorithms we use are supervised algorithms. They require an initial set of tweets labelled whether they are positive or negative. The algorithms then find common patters which are applied to classify unlabelled tweets. We train our algorithms with a sample of 10,000 tweets manually classified by two research assistants (undergraduate students in economics and management science respectively). To ensure the training sample has enough tweets from all possible categories in both dimensions we randomly select them from the accounts of the Financial Times (FT) and The Wall Street Journal (WSJ). We use this source of tweets instead of those posted by fund families since negative information is less likely to be disclosed.