Finfluencers*

Ali Kakhbod
UC Berkeley

Seyed Kazempour
LSU

Dmitry Livdan
UC Berkeley, CEPR

Norman Schürhoff
University of Lausanne, SFI, CEPR

March 2025

*We thank Bloomberg and StockTwits for providing the data. We thank Tony Cookson, Chukwuma Dim, Mark Grinblatt, Ryan Israelsen, Amir Kermani, Eben Lazarus, Runjing Lu, and participants at the 2023 EFA Conference, 2024 Jackson Hole Finance Conference, 2024 SFI Research Days, 2024 NBER Conference on Big Data and Securities Markets, 2025 AFA Conference, and University of Calgary. Ali Kakhbod is with the Walter A. Haas School of Business, University of California, Berkeley, Berkeley, CA 94720. fax: (510) 643-1412; e-mail: akakhbod@berkeley.edu. Seyed Kazempour is with the E. J. Ourso College of Business, Louisiana State University, Baton Rouge, LA 70803. e-mail skazempour@lsu.edu. Dmitry Livdan is with the Walter A. Haas School of Business, University of California, Berkeley, Berkeley, CA 94720. fax: (510) 643-1412; e-mail: livdan@haas.berkeley.edu. Livdan is also a Research Fellow of the CEPR and has no conflict of interest to disclose. Norman Schürhoff is with the Faculty of Business and Economics and Swiss Finance Institute, University of Lausanne, CH-1015 Lausanne, Switzerland. Tel.: +41 (21) 692-3447; fax: +41 (21) 692-3435; e-mail: norman.schuerhoff@unil.ch. Schürhoff is also a Research Fellow of the CEPR. He gratefully acknowledges research support from the Swiss Finance Institute and the Swiss National Science Foundation under SNF project #100018_192584, "Sustainable Financial Market Infrastructure: Towards Optimal Design and Regulation" and has no conflict of interest to disclose.

Finfluencers

Abstract

The social media activity of financial influencers ("finfluencers") can propagate and amplify poor investment advice, especially if less skilled finfluencers are more active and their tweets attract more followers. Using tweet-level data from a popular stock-picking platform, we show most finfluencers are unskilled or "antiskilled," producing negative abnormal returns, while a minority demonstrate skill. Unskilled and antiskilled finfluencers are more engaging, post excessively optimistic tweets that precede price reversals, and attract larger followings than skilled finfluencers. Consistent with a model where social media prioritizes engagement over skill, this leads to the spread of false advice and distorted belief aggregation.

JEL Classification: G12, G14, G41

Key words: Finfluencers, social media, mixture modeling, belief bias, wisdom of the crowd

Financial influencers, or *finfluencers*, offer unsolicited investment advice on social media platforms. Many finfluencers have large followings, and their recommendations can significantly influence investors' investment decisions. The Securities and Exchange Commission (SEC) has expressed concern about finfluencers, particularly because most provide investment advice or recommendations without being registered as investment advisers or brokers.[1] Despite their growing influence, little is known about the quality of their advice, the strategies they employ, or the relationship between skill and popularity.

If the market for finfluencers functions efficiently, social media users should gravitate toward skilled finfluencers with more valuable information, leading to their dominance. More skilled finfluencers should hence have more followers than less skilled ones or at least attract more followers over time, and less skilled finfluencers should be driven out of the market. However, if engagement dynamics—such as homophily and entertainment aspects—drive attention, then unskilled or even antiskilled finfluencers (those whose recommendations systematically destroy value) may persist, distorting belief aggregation (Golub and Jackson, 2012; Hirshleifer, 2020) and asset pricing (Berk and Van Binsbergen, 2022; Pedersen, 2022).[2]

This paper empirically investigates the skill composition of finfluencers, their tweeting strategies, and how their advice contributes to or detracts from market efficiency. Using tweet-level data, we classify finfluencers into three skill types: skilled ($\alpha > 0$), unskilled ($\alpha = 0$), and antiskilled ($\alpha < 0$). Distinguishing skill from luck is central to this classification, as even poor forecasters may occasionally be right. To address this issue, we employ a mixture-modeling approach that assigns a probability to each finfluencer's skill level, incorporating an economic prior: skilled users have positive true alpha, unskilled users have zero true alpha, and antiskilled users have negative true alpha. Luck is captured through a Gaussian noise

[1] Under federal securities laws, individuals who provide investment advice for a fee or other compensation must register with the SEC or with a state securities regulator unless they qualify for an exemption. The SEC has taken action against individuals and firms that have violated these registration requirements, including those who have provided investment advice through social media. See, e.g., SEC press releases "SEC Obtains Emergency Asset Freeze, Charges California Trader with Posting False Stock Tweets," March 15, 2021 (sec.gov/news/press-release/2021-46?utm_medium=email&utm_source=govdelivery) and "SEC Charges Eight Social Media Influencers in $100 Million Stock Manipulation Scheme Promoted on Discord and Twitter," December 14, 2022 (sec.gov/news/press-release/2022-221).

[2] Regulators and industry groups caution investors against blindly following finfluencers, especially those with undisclosed financial incentives. The SEC (sec.gov/oiea/investor-alerts-and-bulletins/social-media-and-investment-fraud-investor-alert), state regulators (dfpi.ca.gov/2022/10/05/social-media-finfluencers-who-should-you-trust), and industry organizations (nasaa.org/64940/informed-investor-advisory-finfluencers) have issued guidance and warnings to investors about the potential risks of relying on financial advice from finfluencers, especially when the finfluencers have a financial interest in the products or services they promote.

component.[3]

By allowing for three skill groups, we can assess the effectiveness of social media as an information aggregation mechanism in financial markets. If not the most skilled but unskilled and even antiskilled finfluencers prevail, it raises questions about how investors process and aggregate information. We estimate the mixture model on tweet-level data from a popular social media platform, StockTwits, and measure the investment performance associated with following each finfluencer's advice. Using 72 million tweets from StockTwits by over 29,000 finfluencers and controlling for 36 million news stories, we find that the skill composition of finfluencers is heterogeneous. While the average finfluencer's skill level is close to zero, 28% of finfluencers provide valuable investment advice, leading to average monthly abnormal returns of up to 2.6%. Only 17% of finfluencers are unskilled, providing no outperformance relative to risk-adjusted returns. By contrast, the majority, 55%, are antiskilled, and following their advice results in monthly abnormal returns of -2.3%.[4]

It might seem odd that so many finfluencers are robustly classified as un- and antiskilled and, yet, persistently active on the platform. If users care about the quality of the information supplied by finfluencers, they will follow finfluencers with more valuable information. In that case, skilled finfluencers should have more followers than un- and antiskilled finfluencers, and the un/antiskilled ones should be driven out of the market so long as tweeting is not costless. In light of these findings, we develop a model of social media activity, user popularity, and aggregate belief formation that guides our empirical tests and potential policy recommendations.

Based on the model's predictions, we explore the relationships between finfluencers' skill and followings, activity, and the beliefs revealed by the tweets of the different types of finfluencers. If followers are misallocated towards less-skilled finfluencers, it can distort the "wisdom of the crowd", that is, the collective ability of platform users to aggregate diffuse information from a broad base of finfluencers. With the finfluencers' skill measures in hand,

[3]Mixture modeling involves fitting a distribution that combines several other distributions, known as components, to a dataset. Our assumptions allow us to derive and estimate the joint distribution for finfluencers' alpha that combines exponential distributions for skilled users, a mass at zero for unskilled users, negative exponential distributions for antiskilled users, plus a Gaussian distribution for luck.

[4]These findings differ from Cookson and Niessner (2020), which suggests that both professional and novice finfluencers exhibit some degree of skill. Cookson and Niessner (2020, p. 192) report that positive sentiment by professional finfluencers and negative sentiment by novice finfluencers both predict positive and negative abnormal returns, respectively, suggesting that both groups possess skill in identifying market outcomes. If novice finfluencers in Cookson and Niessner (2020) are a proxy for a mix of unskilled and antiskilled finfluencers, these results demonstrate the advantages of doing the performance-based type classification in our paper over conditioning on self-reported profile information.

we link each finfluencer's skill type to their follower base, user engagement, and tweeting strategies. Surprisingly but consistent with the model, we find that skilled finfluencers have fewer followers than unskilled and antiskilled finfluencers, and these relations are significant even out-of-sample. Thus, there is a misallocation of more followers to less skilled finfluencers.

User engagement is, as the model suggests, a potential explanation for the negative skill-follower relationship in the data. One way to improve user engagement and increase the follower count is to tweet more actively. Finfluencers who tweet more often may be considered experts with valuable information, allowing them to build a reputation (Benabou and Laroque, 1992). On the other hand, finfluencers who tweet more often may be more confident or a "charlatan" believing a large tweeting volume proxies for skill (Berk and Van Binsbergen, 2022). Thus, their tweets might be less informative. We therefore investigate how finfluencers' skill relates to their tweeting activity. Here we find that skilled finfluencers are less active than un- and antiskilled finfluencers, suggesting that activity is not a sign of skill but pure user engagement.[5]

To establish why un- and antiskilled finfluencers are more popular than skilled finfluencers, we explore the determinants of popularity by linking it to finfluencers' tweeting strategies: when and what do they tweet? This allows us to check whether they possess unique skills or just follow commonly known investment behaviors including momentum, contrarian, return chasing, and herding.[6] We find that skilled finfluencers are return-, social sentiment-, and news-contrarian. They also do not herd on other users' tweets. We also check if finfluencers with more negative tweets are more skilled. The literature has documented that short-sellers are informed (e.g., Engelberg, Reed, and Ringgenberg, 2012; Boehmer, Jones, and Zhang, 2008) and Miller (1977) suggests that imposing short-selling constraints leads to overpricing. Using Markit's measure of short-selling constraints for stocks, we show that users who tweet negatively about stocks with higher short-selling constraints are indeed more likely to be skilled. Contrary to skilled ones, antiskilled finfluencers ride return momentum and social sentiment momentum and tend to chase returns.

The observed relations between tweeting activity and strategies, and finfluencer skill suggest social media users can and should use tweeting behavior to identify skilled finfluencers.

[5] Furthermore, we find that finfluencers' skills are persistent. To study persistence, we separately train the model on the first year of data and the entire sample and look at the overlap.

[6] We define each user's return-chasing tendency as the percentage of her tweets that are either positive and about stocks in the highest decile of returns over the prior five trading days, or negative and about stocks in the lowest decile of returns over the prior five trading days. We define each user's herding tendency as the percentage of her positive tweets about stocks in the highest decile of overall positive tweeting volume over the past five days.

However, more skilled finfluencers have fewer followers while more active finfluencers have more followers. While skill can be identified, albeit imperfectly, by looking at finfluencers' followers, how active the finfluencer is, and what tweeting strategy the finfluencer pursues, social media users seem to use the available information insufficiently to screen skilled finfluencers, or they decide to ignore it. This behavior is consistent with theories of homophily that predict a reduction in the speed of learning and information diffusion (see, e.g., Golub and Jackson, 2012) and can, hence, lead to stock mispricing as in Berk and Van Binsbergen (2022) and Pedersen (2022).

Guided by the model, we explore the bias in aggregate beliefs and potential asset price distortions that arise from following the advice of the group of un/antiskilled finfluencers. For this to occur, it is not important that finfluencers have massive followings but that their tweets reflect widely held beliefs and that investors are slow learners about skill, consistent with our findings. Our empirical analysis reveals that skilled finfluencers generally maintain a more neutral stance, only occasionally expressing strongly positive or negative social sentiment. In contrast, unskilled and antiskilled finfluencers tend to be consistently overoptimistic and exhibit persistent belief swings. This pattern of tweeting behavior distorts the wisdom of the crowd in the model since these less skilled finfluencers receive more weight as they are more active and reflect the beliefs of larger followings than their skilled counterparts.

In the data, we examine the relationship between stock returns, finfluencers' tweeting activity, and sentiment, both past and future. Using panel vector autoregressions, which pool data across finfluencers and stocks, we document the lead-lag dynamics between stock returns and tweets while accounting for unobserved heterogeneity.[7] Our findings confirm that skilled finfluencers accurately predict future returns, whereas the more visible unskilled and antiskilled finfluencers make incorrect predictions. Consistent with the model, increased tweeting activity by un/antiskilled finfluencers coincides with price peaks, following price rises and anticipating subsequent price declines. By contrast, when pooling all finfluencers these predictive patterns wash out.

Our results emphasize the potential for misinformation and biased views in financial market environments where engagement and visibility, rather than content quality, drive follower growth. This provides another rationale for noise traders to matter (Kogan, Ross, Wang, and Westerfield, 2006). Consequently, our findings highlight the need for intervention strategies to promote high-quality content and reduce the influence of low-quality finfluencers. In counterfactual experiments, we use the model to guide the impact of these policy interven-

[7]Our findings are also consistent with those from out-of-sample portfolio tests.

tions. We find that improved transparency would help to better balance the follower growth rates between good and bad finfluencers. In turn, active quality tracking by the platform can help redirect follower growth toward higher-quality finfluencers. Last, curation processes and verification systems can render good finfluencers more competitive against bad finfluencers.

Literature review. Our results are important in several ways. First, our paper contributes to the literature on the survival of unskilled and antiskilled agents. Berk and Van Binsbergen (2022) examine the survival of unskilled and antiskilled experts, which they term "charlatans," in equilibrium. They provide a framework that aligns with our observations that unskilled and antiskilled finfluencers, despite their lack of ability, command larger followings on StockTwits than their skilled counterparts. Pedersen (2022) offers a theory of market dynamics in an environment populated by stubborn users who resist updating their beliefs. His theory provides a framework for the consequences of our main finding that social media users can access information that allows them to discern skilled finfluencers, but often neglect it. Our paper supports predictions by Pedersen (2022) by showing that social media users can identify which source of information is more reliable, but they ignore it.

Second, our paper relates to the literature on investor expectations.[8] Greenwood and Shleifer (2014) demonstrate a positive correlation between investor expectations and past market returns and a negative correlation with future returns. Our findings reveal a similar pattern among antiskilled finfluencers on StockTwits, many of whom exhibit tweeting behavior consistent with extrapolative beliefs. We show that a subset of finfluencers hold accurate beliefs about future stock returns, allowing them and their followers to partially counterbalance the misguided actions of antiskilled finfluencers and their followers.[9] Cookson, Lu, Mullins, and Niessner (2022) find that aggregate attention predicts negative returns while aggregate sentiment predicts positive returns. This also holds in our data when we pool finfluencers. However, once we split between skill and visible activity and sentiment, results differ. Consistent with our model, more positive tweets by skilled finfluencers anticipate positive stock returns while more positive tweets by more visible finfluencers anticipate negative stock returns. Skilled attention (i.e., higher activity by skilled finfluencers) precedes positive returns, and only visible attention (i.e., higher activity by more visible, less skilled

[8] In Banerjee (1992), once visible finfluencers promote a particular investment, others may follow, ignoring their own private information. This herding effect may amplify bad advice, where popularity supersedes accuracy, distorting collective beliefs. Bikhchandani, Hirshleifer, and Welch (1992) predict that users may abandon their own information in favor of mimicking the crowd.

[9] Despite antiskilled finfluencers' negative alpha, it is still possible for them to benefit their followers in a fashion similar to what Gennaioli, Shleifer, and Vishny (2015) suggest about professional managers.

finfluencers) predicts negative returns and the latter is not statistically significant.

Third, our paper draws on the literature identifying investment skill using mixture modeling. Harvey and Liu (2018) show that some mutual fund managers have true positive alphas. Our paper shows that despite the negative correlation between StockTwits' average sentiment and future returns, some finfluencers post informative tweets. However, unlike in the context of mutual fund management, users in the social media domain do not flock to the most skilled advisors. Our paper also contributes to the literature on the skills of individuals, their behavioral traits, and how these aspects influence their social networks.[10] We demonstrate that social media users can identify skilled finfluencers, but often choose to follow more active and visible finfluencers that are less skilled, even if this behavior leads to suboptimal investments and the activity of these finfluencers serves as a contrarian signal.[11]

Our findings contrast with existing research on social media platforms for stock prediction with curated users, such as Seeking Alpha. Curated users resemble semi-professional and professional financial analysts and advisors, in contrast to the non-curated users on StockTwits. Hence, it is a-priori unclear if StockTwits users are skilled. Chen, De, Hu, and Hwang (2014) find a positive correlation between the sentiment expressed in Seeking Alpha articles and future stock returns.[12] In contrast to platforms for curated users, our results indicate that the majority of non-curated StockTwits users are un/antiskilled, and only a

[10] A broader literature, e.g. Barber and Odean (2007), demonstrates behavioral biases of retail investors. While Cookson, Engelberg, and Mullins (2023) focus on followers (400,000 users), we focus on those 29,000 that are finfluencers. They find that self-declared bull vs. bear users follow users in the same category and, as a result, live in their bubbles, a phenomenon called information siloing, but they do not distinguish between finfluencers and users, split users by their skill types, or explore the impact of active user engagement.

[11] This behavior mirrors the sociological phenomenon of homophily, the tendency of individuals to associate with others who share similar characteristics or values (Lazarsfeld, Merton, et al., 1954; Kandel, 1978; McPherson, Smith-Lovin, and Cook, 2001). Homophily leads to positive assortative matching, which slows information diffusion (Currarini, Jackson, and Pin, 2009; Golub and Jackson, 2012). In our setting, homophily manifests as social media users preferring to follow finfluencers who exhibit similar investment behaviors to their own. In contrast to homophily, echo chambers refer to situations where individuals are only exposed to information or opinions that confirm their existing beliefs and are sheltered from opposing views. While homophily can contribute to the creation of echo chambers, the two concepts are different.

[12] Several papers analyze the informativeness of social media and information crowdsourcing platforms (Crawford, Gray, and Kern, 2017; Crawford, Gray, Johnson, and Price III, 2018; Jame, Johnston, Markov, and Wolfe, 2016; Ballinari and Behrendt, 2021; Giannini, Irvine, and Shu, 2018; Curtis, Richardson, and Schmardebeck, 2014; Azar and Lo, 2016; Bartov, Faurel, and Mohanram, 2018; Campbell, D'Adduzio, and Moon, 2021; Cookson, Engelberg, and Mullins, 2023). Our results differ in that we focus on the dispersion and systematic bias in the quality of advice provided by finfluencers, and their implications. Dim (2022) employs a Gaussian mixture modeling methodology on Seeking Alpha articles and discovers that a considerable majority, 56%, of its authors correctly predict stock returns. The Gaussian mixture model implies that a large portion of finfluencers are misclassified as (anti)skilled if the variance is large relative to the mean, while our approach directly models the economic prior.

minority can predict the direction of stock returns, leading to a negative relation between aggregate sentiment, tweet activity, and future returns.[13]

1 Data and Finfluencer Types

This section describes the data and outlines our approach to measuring the investment skill, or alpha, of active users on StockTwits, a social media platform dedicated to the exchange of stock-related investment ideas between investors free of charge.

1.1 Motivation

Social media has become an increasingly important channel for the dissemination of investment information, particularly among retail investors. Platforms like StockTwits allow users to share trading insights and stock recommendations in real time, potentially facilitating information aggregation and improving market efficiency. However, the quality of information shared on such platforms varies, and it is unclear whether social media improves price discovery or instead contributes to noise and mispricing.

StockTwits users can be broadly categorized into two groups: active users, who post tweets containing investment analysis and stock recommendations, and passive users, who consume content without contributing their own. While many StockTwits users engage solely by following others, we focus on the subset of active users who publicly share investment insights, trading strategies, or technical analysis. We refer to these active users as finfluencers to distinguish them from passive followers.[14] While not all active users necessarily seek to influence others or provide financial advice, their investment-related tweets affect the information environment on StockTwits.

[13] These findings align with the divergence documented in Cookson, Lu, Mullins, and Niessner (2022) who show that while attention is highly correlated across platforms, sentiment is not. Our results highlight the importance of curated vs. non-curated users and the incentives faced by social media users in a setting without a curator who warrants the reputation of its users, which is reflected in our model-based policy recommendations. We show that the sentiment from skilled finfluencers predicts returns correctly, but the sentiment from antiskilled ones predicts returns incorrectly. In contrast to attention signals based on pooling social media users, higher activity by skilled finfluencers precedes positive returns.

[14] StockTwits users that do not post do not aim to influence others or provide advice. Earlier papers have studied these users as interactions among peers (Cookson, Engelberg, and Mullins, 2023), while we focus on the subset of active users. We later refine the analysis by identifying active users with influence. In particular, we screen by follower base and volume of tweets. In the aggregate analysis, we focus on the most visible users as opposed to those who post once or twice and never again.

A key concern is that not all active users providing investment advice have relevant expertise. Some finfluencers possess genuine skill and contribute valuable information, while others provide recommendations that are uninformed or systematically biased. If social media serves as an efficient aggregator of dispersed private information, one would expect skilled contributors to be more influential over time. However, if content visibility is determined by factors unrelated to investment accuracy—such as user engagement or behavioral biases—then uninformed or biased opinions may be overrepresented, distorting belief formation rather than improving it. To formalize this distinction, we classify active users into three categories:

1. Skilled finfluencers: Users whose investment recommendations generate positive alpha ($\alpha_i > 0$), reflecting genuine predictive ability.
2. Unskilled finfluencers: Users whose recommendations are indistinguishable from noise, leading to zero alpha ($\alpha_i = 0$).
3. Antiskilled finfluencers: Users whose recommendations systematically generate negative alpha ($\alpha_i < 0$), thereby propagating false beliefs.

The last category is typically not considered in the literature, as these individuals are a priori unlikely to survive in the long run due to their poor advice; however, they are particularly relevant in the context of social media. There are several reasons why the tweets by antiskilled finfluencers may generate negative alpha. These reasons include their own biased beliefs, incentives to create engagement and attention, overconfidence, tendency to chase returns, potential conflicts of interest and incentives that do not align with investment performance, and their ability to influence the behavior of retail traders to their own advantage.[15]

Understanding the distribution and persistence of these finfluencer types has important implications for the role of social media in financial markets. If skilled finfluencers attract a larger following, their influence grows over time and social media could enhance information efficiency. Conversely, if unskilled or antiskilled finfluencers gain prominence due to factors unrelated to skill, then social media may lead to a failure of the "wisdom of the crowd" and contribute to market inefficiencies rather than mitigating them. The idea behind the wisdom of the crowd is that information is diffuse and dispersed among, in our setting, many finfluencers and needs to be aggregated to filter out noise. When only a subset of finfluencers

[15] Appendix C dissects finfluencers' tweeting strategies and shows which characteristics are associated with each skill group.

possess information, by aggregating beliefs uniformly or based on activity or follower count the correct beliefs are underrepresented in the aggregate.

Before formalizing this notion in Section 2, we discuss the data used to identify the composition of finfluencer types, show how to estimate the probabilities of finfluencers' skill types, and document the dominance of un- and antiskilled finfluencers.

1.2 Data

Data sources. StockTwits is a social media platform for sharing investment ideas between investors and is often described as the "Twitter for stocks" because it allows users to post tweets about specific stocks. Users can follow other users, see what they are saying, and interact with them. StockTwits is the most popular social trading platform for retail investors and has nowadays over 6 million registered users. A small group of users, which we term finfluencers, actively posts tweets. StockTwits finfluencers are not curated, as opposed to users on other social media platforms, which heightens activity and engagement with the platform and broadens the diversity of information and opinions present on the platform, but it also facilitates the spread of noise and misinformation since StockTwits finfluencers are not subject to the same editorial standards as curated users.

Our data is from several sources. We obtain tweet data from Bloomberg for the sample period from July 13, 2013, when Bloomberg started collecting the data to January 1, 2017, user-level data collected directly from StockTwits until February 2, 2018,[16] matched stock returns from CRSP, and factor returns from Ken French's website. In addition, we use Markit data for daily stock-level statistics on short interest and shorting costs.

Bloomberg provides tweet data for two platforms, StockTwits and X (formerly Twitter). We use tweet data for StockTwits delivered through Bloomberg because it comes with standardized social sentiment scores that are readily usable by investors and that are not subject to any biases we may introduce. The advantage of StockTwits over X is that Bloomberg provides the identities of the finfluencers on StockTwits, but not on X. As a result, we can measure finfluencer-specific skill for StockTwits only.

For each tweet, the Bloomberg data contains the time of the post, tweet content, stock ticker, and user name used to post the tweet. Bloomberg supplies a social sentiment score for each tweet based on a proprietary machine learning algorithm, the confidence level of the social sentiment score from 1/3 to 1, a relevance score from 0 to 1, and topic codes.

[16] StockTwits does not provide API access to us after this time.

The social sentiment score by user i in stock j for its nth tweet on the day t takes discrete values $SocSent_{i,j,t,n} \in \{-1, 0, 1\}$. Out of 72 million tweets, 11%/77%/12% are negative/neutral/positive. This distribution in social sentiment scores among finfluencers is similar to the sentiment in public news and more balanced between positive and negative sentiment than the self-declared sentiment of non-finfluencers in Dim (2022) and Cookson, Engelberg, and Mullins (2023).

The Bloomberg data also contains the news data and the corresponding news sentiment that we use to control for public sentiment. For each news story, Bloomberg reports the time of the release, news headline, stock ticker, and news source. Bloomberg supplies a news sentiment score for each story that is based on a proprietary machine learning algorithm, the confidence level of the news sentiment score from 1/3 to 1, a relevance score from 0 to 1, and topic codes. The news sentiment score in stock j for its nth news on the day t takes discrete values $NewsSent_{j,t,n} \in \{-1, 0, 1\}$. Out of 36 million news stories, 12%/59%/29% are negative/neutral/positive. Comparing news to social sentiment, these statistics show that tweets are less likely to be negative than news.

Table 1 provides descriptive statistics on user activity, follower base, and measured alphas. We use the StockTwits API to collect out-of-sample data for each user.[17] For each StockTwits user the data contains the number of tweets with a sample mean equal to 206, the minimum number of tweets equal to 1, and a maximum number of tweets equal to 321,154. The average number of followers in the data is 1,037 as of the time of our download, with a minimum of 0 a and maximum of 489,704.

Finfluencers' social sentiment scores. To compute a social sentiment score for each finfluencer i on each stock j each day t, we first need to match StockTwits user names with our Bloomberg data. The user name supplied by Bloomberg is the StockTwits user name displayed on the screen. We match the StockTwits user name from Bloomberg to the corresponding user name in StockTwits. While the user name is unique, the screen name is not. Therefore, the StockTwits screen name coincides in most but not all cases with the StockTwits user name from Bloomberg. As a result, some users cannot uniquely be matched from Bloomberg to StockTwits and we pool or eliminate the duplicates.

[17] There are a total of 139,401 users as of February 2, 2018, when the data was collected. Out of the total, we can match 105,535 StockTwits users to our Bloomberg data. Since many StockTwits users are inactive in posting tweets, we pool all users with total activity on StockTwits of fewer than 20 tweets or retweets. Since a user's StockTwits history can be longer than our sample period, we have users with fewer than 20 tweets in our sample. For 29,475 users we can measure activity levels and alphas, while for 22,072 of those we have the follower count.

Table 1: Summary Statistics

This table reports summary statistics of social sentiment in Panel A and finfluencers' activity levels, follower count, measured alphas $\tilde{\alpha}$, their standard errors, and t-statistics in Panel B. Finfluencers' activity levels and follower count are retrieved from StockTwits. We calculate excess returns over the next 20 trading days using the Fama-French five-factor model. The measured alpha $\tilde{\alpha}$ for each user is the average of signed adjusted returns after her tweets. Alphas and their standard errors are in percentage points.

	N	Mean	S.D.	Min	p10	p50	p90	Max
Panel A: User-stock-day statistics								
$SocSent^{+}_{i,j,t}$	7,842,522	0.39	1.08	0.00	0.00	0.00	1.00	365.00
$SocSent^{0}_{i,j,t}$	7,842,522	1.39	2.95	0.00	0.00	1.00	3.00	757.00
$SocSent^{-}_{i,j,t}$	7,842,522	0.20	0.75	0.00	0.00	0.00	1.00	365.00
$SocSent_{i,j,t}$	7,842,522	0.14	0.60	-1.00	-1.00	0.00	1.00	1.00
Panel B: User-level statistics								
Activity	29,475	206	2,653	2.00	15.00	52.00	289.00	321,154
Followers	22,072	1,037	8,977	0.00	1.00	10.00	134.50	489,704
Measured alpha $\tilde{\alpha}_i$	29,475	-0.67	7.48	-55.34	-7.76	-0.37	6.11	54.37
S.E. of $\tilde{\alpha}_i$	29,475	4.35	4.71	0.00	0.87	2.86	9.61	59.53
t-stat of $\tilde{\alpha}_i$	29,475	-0.54	58.55	-8,637	-2.18	-0.16	1.65	4,328

The matching of returns and tweets is also important. We apply the following procedure: If a tweet is posted during trading hours, we match it to the same trading day. That is, day t will be the trading day. If a tweet is posted after hours, on holidays, or weekends, we match it to the next trading day. In other words, day $t+1$ will be the trading day. That is, we match every tweet with the first trading-day closing after it is posted.

We aggregate all tweets by user i in stock j on day t into a single social sentiment score according to

$$\begin{aligned} SocSent^{+}_{i,j,t} &= \textstyle\sum_{n=1}^{N_{i,j,t}} \mathbb{1}(SocSent_{i,j,t,n} = 1), \\ SocSent^{-}_{i,j,t} &= \textstyle\sum_{n=1}^{N_{i,j,t}} \mathbb{1}(SocSent_{i,j,t,n} = -1), \\ SocSent_{i,j,t} &= \max\left\{-1, \min(1, SocSent^{+}_{i,j,t} - SocSent^{-}_{i,j,t})\right\}, \end{aligned} \tag{1}$$

where $n = 1, ..., N_{i,j,t}$ is the index of the tweet. $SocSent^{+}_{i,j,t}$ counts the positive tweets by finfluencer i in stock j on the day t, and $SocSent^{-}_{i,j,t}$ counts the negative ones. Table 1 provides user-stock-day level statistics, showing that $SocSent^{+}_{i,j,t}$ is 0.39 on average with a maximum of 365 while $SocSent^{-}_{i,j,t}$ is 0.20 on average with a maximum of 365. For comparison, $SocSent^{0}_{i,j,t} = \sum_{n=1}^{N_{i,j,t}} \mathbb{1}(SocSent_{i,j,t,n} = 0)$ capturing neutral sentiment is 1.39 on average with a maximum of 757. The max and min operators are used to normalize $SocSent_{i,j,t}$

to the $[-1, 1]$ interval. Table 1 shows that $SocSent_{i,j,t}$ is 0.14 on average with a standard deviation of 0.60.

Measuring naïve alphas for each finfluencer. To measure each finfluencer's alpha, $\tilde{\alpha}_i$, we compute finfluencer-level abnormal returns over horizon H based on their social sentiment scores (1). For finfluencer i, we calculate the naïve alpha $\tilde{\alpha}_i$ as the abnormal return obtained over different horizons $[t+1, t+H]$ depending on the finfluencer's social sentiment scores $SocSent_{i,j,t}$. Abnormal stock returns for stock j over different horizons, $AbnRet_{j,t+1,t+H}$, are computed using the standard procedure where we first calculate factor exposures for each stock and then subtract the expected returns over horizon H.

We call the measured alpha, $\tilde{\alpha}_i$, a "naïve" measure of skill because it does not account for type 1 and 2 errors. We calculate the mean signed abnormal return and its standard error, $\tilde{\sigma}_i$, for every user in the data by running univariate regressions:

$$SocSent_{i,j,t} \times AbnRet_{j,t+1,t+H} = \tilde{\alpha}_i + \epsilon_{i,j,t+1,t+H}, \tag{2}$$

for all N_i stock-days for which $SocSent_{i,j,t} \neq 0$ and separately for all users $i = 1, ..., I$ and multiple values of H. Equipped with user-specific abnormal returns $\tilde{\alpha}_i$, $i = 1, ..., I$, over horizon H we can compute mean signed returns and their t-stats. We can estimate equation (2) in different ways.[18] For the factors, we use the Fama-French one, three, and five-factor models. For the horizon, we use H = 1, 2, 5, 10, or 20 days. The results are generally comparable and do not materially depend on the factor model of returns.

Table 1, Panel B reports finfluencers' measured alphas, $\tilde{\alpha}_i$, from specification (2) with H = 20 business days. The average finfluencer has a monthly measured alpha of -0.67% (annualized: -8% per year). The median measured alpha is -0.37% and hence also negative, meaning that most finfluencers post systematically anti-informative tweets.[19] Table 1 also shows that the standard errors of measured alphas are large compared to the point estimates.

[18] As an alternative to (2), we have run multivariate regressions for all users $i = 1, ..., I$ combined and multiple values H: $SocSent_{i,j,t} \times AbnRet_{j,t+1,t+H} = \sum_{\iota=1}^{I} \tilde{\alpha}_i \times \mathbb{1}(\text{User } i = \iota) + \epsilon_{i,j,t+1,t+H}$. This specification has the advantage that it corrects standard errors $\tilde{\sigma}_i$ for contemporaneous correlation between finfluencers, while it has the disadvantage that it is very high-dimensional and leads to more noisy standard error estimates. Still, the results for the multivariate regression specification are similar to (2).

[19] These results are consistent with the findings in previous papers. They confirm that average social media users are systematically wrong in predicting stock returns (Giannini, Irvine, and Shu, 2018). We are therefore confident that our conventional approach to abnormal return computations (as documented in the prior section) yields both valid results and results that are comparable across time and social media platforms. This should eliminate potential concerns about external validity.

The average (median) standard error is 4.35% (2.86%) monthly. However, despite the relatively large standard errors, some users have statistically significant measured alphas.[20] Table 1 shows that the proportion of users for whom the p-value of the measured alpha is less than 5% (10%) is 19.5% (25.6%). These numbers are larger than what we would expect if all users were uninformed.

1.3 Identifying finfluencer types

The issue with taking measured alphas, $\tilde{\alpha}_i$, at face value as a proxy for finfluencer skill is that the statistical tests have a size and power. While we can measure $\tilde{\alpha}_i$ for every finfluencer and calculate its t-statistic, it is unclear without using a model how often the null of $\alpha = 0$ is falsely rejected or falsely accepted (type 1 and 2 errors).[21] In this section, we measure true skill, α_i, for each finfluencer by developing an empirical model that addresses the type 1 and 2 errors of statistical tests on measured alphas, $\tilde{\alpha}_i$. We start by describing the methodology to extract the true alphas and then describe several measures of finfluencer skill.

Mixture of finfluencer skills. Since the returns from following finfluencers' tweets are noisy, our naïve measure of skill, $\tilde{\alpha}_i$, is a noisy measure of finfluencers' true skills, α_i. The relation between α_i and $\tilde{\alpha}_i$ can be written as

$$\tilde{\alpha}_i = \alpha_i + \epsilon_i, \tag{3}$$

where we assume $\epsilon_i \sim \mathcal{N}(0, \tilde{\sigma}_i^2)$ and $\tilde{\sigma}_i$ is the standard error of finfluencer i's abnormal return in the data. It follows from (3) that the distribution of the measured skill is a convolution between the distributions of true skill and the error term ϵ_i. Following the literature on performance evaluation (Chen, Cliff, and Zhao, 2017; Harvey and Liu, 2018; Crane and Crotty, 2020; Dim, 2022), we employ the mixture modeling methodology while imposing three categories of StockTwits users to estimate the distribution of α among finfluencers. The idea behind the mixture modeling methodology is that it aggregates information to improve the signal-to-noise ratio of the data.

[20] One may be concerned that the users' measured alphas are specific to StockTwits. However, while we cannot fit the model on X users due to the lack of finfluencer identifiers in the X data, the distribution of measured alphas looks very similar for X.

[21] If one uses the t-stat threshold of 1.96, 5% of finfluencers will appear with significant alpha (mean signed abnormal returns) even if the true alpha is zero. Hence, there are finfluencers with truly positive (or truly negative) alpha that we cannot detect when the t-stat is less than 1.96.

For skilled and antiskilled users, we allow for several subtypes with different levels of (anti)skill. We assume there are K^+ (K^-) types of users with positive (negative) skills. Let π_k^+ be the share of skilled finfluencers of type k, π^0 the share of unskilled finfluencers, and π_k^- the share of antiskilled finfluencers of type k. Then, true skill α is distributed across finfluencers according to the finite mixture distribution

$$f(\alpha) = \mathbb{1}\{\alpha > 0\} \sum_{k=1}^{K^+} \pi_k^+ g(\alpha; \mu_k^+) + \pi^0 \mathbb{1}\{\alpha = 0\} - \mathbb{1}\{\alpha < 0\} \sum_{k=1}^{K^-} \pi_k^- g(\alpha; \mu_k^-), \tag{4}$$

where $g(\alpha; \mu)$ if $\mu > 0$ ($-g(\alpha; \mu)$ if $\mu < 0$) is a continuous distribution with a mean of μ and

$$\begin{aligned} \sum_{k=1}^{K^+} \pi_k^+ + \pi^0 + \sum_{k=1}^{K^-} \pi_k^- &= 1, \\ \mu_k^+ > 0 \text{ for } 1 \leq k \leq K^+&, \\ \mu_k^- < 0 \text{ for } 1 \leq k \leq K^-&. \end{aligned} \tag{5}$$

In expression (4), μ_k^+ and μ_k^- are the expected abnormal returns of the positive and negative components $k = 1, ..., K^+(K^-)$.

The standard assumption in the literature is that skilled and antiskilled types are normally distributed. This assumption does not naturally align with our economic prior that skilled types have positive true alpha ($\alpha_i > 0$) and antiskilled types have negative true alpha ($\alpha_i < 0$). The exponential and negative exponential distributions are the maximum-entropy distributions with the greatest uncertainty consistent with the type constraints (5). We therefore estimate both exponential-negative exponential and, alternatively, normal distributions for $g(\alpha; \mu)$. Our first choice is to assume that the skilled and antiskilled are a mixture of (negative) exponentially distributed types. In this case, $g(\alpha; \mu) \equiv \frac{1}{\mu} \exp(-\frac{1}{\mu}\alpha)$.[22] To check for robustness, we impose the more customary mixtures of normally distributed types. We then perform Kolmogorov-Smirnov tests which show that (negative) exponentially distributed types provide a better fit to our social media data.

Given that $\tilde{\alpha}_i = \alpha_i + \epsilon_i$, the distribution of measured alphas, $\tilde{\alpha}_i$, can be calculated as the

[22] The mixture of exponential distributions is a flexible distribution that has been used to model a wide variety of real-world phenomena. For example, it has been used to model the lifespans of electronic devices, the time between events in a queuing system, or the amount of rainfall in each period. It is a particular kind of Beta prime distribution, sometimes termed a Lomax distribution, and useful whenever the domain is one-sided.

convolution of f and a mean-zero Normal distribution with standard deviation $\tilde{\sigma}_i$

$$\mathcal{G}(\tilde{\alpha}_i; \tilde{\sigma}_i, \Theta) = (f * \phi_{\tilde{\sigma}_i})(\tilde{\alpha}_i), \tag{6}$$

where $*$ is the convolution operator, $\phi_{\tilde{\sigma}_i}$ denotes the Normal distribution function with a mean of zero and standard deviation of $\tilde{\sigma}_i$, and $\Theta = (\mu_1^+, ..., \mu_{K^+}^+, \mu_1^-, ..., \mu_{K^-}^-, \pi_1^+, ..., \pi_{K^+}^+, \pi_1^-, ..., \pi_{K^-}^-)$ is the vector of parameters. The likelihood function can be written as

$$\mathcal{L}(\tilde{\alpha}_1, ..., \tilde{\alpha}_I; \tilde{\sigma}_1, ..., \tilde{\sigma}_I, \Theta) = \prod_{i=1}^{I} \mathcal{G}(\tilde{\alpha}_i; \tilde{\sigma}_i, \Theta). \tag{7}$$

We use the maximum likelihood method to estimate the vector of parameters Θ.[23]

We fit several distributions to the StockTwits data and find the results for the exponential family to fit better than those of the Gaussian mixture model. The best fit comes from a model with two exponential distributions for each of the finfluencer categories. We present the results for this distribution and use $K^+ = K^- = 2$ for the main results in the paper. In the Appendix, we report the results of our estimation with alternative specifications.

Table 2 reports the results of our MLE estimation for the model with $K^+ = K^- = 2$. The first (second) positive exponential component has a mean of 1.42% (6.76%) per month and accounts for 21.6% (5.9%) of the population. The first (second) negative exponential component accounts for 45.6% (10.9%) of the population and has a mean of -1.06% (-7.53%). Overall, 27.5% of the population have positive true skills while 56.5% have negative skills. We identify 16.6% of the population with a true skill of zero. Moreover, we calculate the standard errors of all estimated parameters by bootstrapping (with replacement) the sample of measured alphas 100 times, running our MLE estimation on each bootstrapped sample, and calculating the standard error of estimated parameters. Standard errors are tight, which shows that all estimated parameters are statistically significant. The lowest t-statistic among the estimated parameters belongs to the probability of the zero component (t=5.51).

Fit and robustness. To assess the goodness of fit, we perform the bootstrap procedure described in Appendix A. We conclude from this simulation exercise that the fit with $K^{+/-} = 2$

[23]Let X be an exponential variable with mean μ and Y be a mean-zero Normal variable with standard deviation σ. Their sum $Z = X + Y$ is distributed as the convolution of a mean-zero Normal distribution with standard deviation σ and an exponential distribution with mean μ. The convolution has the following closed-form solution: $h(x; \mu, \sigma) = \frac{1}{2\mu}\exp(\frac{\sigma^2}{2\mu^2} - \frac{x}{\mu}) \times \left(1 - \text{erf}\left(\frac{\sigma}{\sqrt{2}\mu} - \frac{x}{\sqrt{2}\sigma}\right)\right)$, where erf is the error function. We use this closed-form solution to speed up our maximum likelihood estimation.

Table 2: Estimating the Distribution of True Alphas

This table reports the results of fitting a mixture model with two exponentials on the $\alpha > 0$, two exponentials on $\alpha < 0$, and a mass at $\alpha = 0$. We calculate excess returns over the next 20 trading days using the Fama-French five-factor model. The measured alpha ($\tilde{\alpha}$) for each user is the average of signed adjusted returns after her tweets. The first column shows the mean of each component (μ's). The second column shows the weight of the component in the mixture (π's). The numbers in parentheses are standard errors of each estimate. To calculate the standard errors, we bootstrap the data 100 times with replacement, estimate the model for each bootstrapped sample, and calculate the standard deviation of the estimated parameters. All numbers are in percentages.

	μ_k (%)	π_k (%)
Skilled type 2	8.14	5.1
	(0.49)	(0.6)
Skilled type 1	1.49	23.5
	(0.10)	(1.0)
Unskilled	0.00	16.6
		(2.7)
Antiskilled type 1	-1.19	45.5
	(0.08)	(1.8)
Antiskilled type 2	-9.15	9.3
	(0.33)	(0.5)
N	29,475	
Log Likelihood	-88,878	
AIC	177,771	
BIC	177,838	

is tight. Internet Appendix IA explores the robustness of our estimates by providing alternative specifications for the distribution of true alphas. Table IA.1 reports parameter estimates for eight alternative model specifications. Panel A reports the estimated distribution of true alphas assuming one and three components for types 1 and 3. The likelihood value and the AIC and BIC criteria improve considerably by moving from one component to two while adding the third component does not improve the fit by much. Panel B reports the results of fitting mixture models over different horizons $H = 1, 2, 5, 10, 20$ with two exponentials on the $\alpha > 0$, two exponentials on $\alpha < 0$, and a mass at $\alpha = 0$. Results are generally consistent with Table 2, with the longer horizon better at separating (anti)skilled from unskilled finfluencers.

1.4 Distribution of finfluencer skill

Using the estimated distribution of true alphas, we can define different measures of finfluencer skill, including the finfluencer's expected alpha and the probability that a finfluencer is skilled. We then analyze the distribution and determinants of skill.

Distribution of skill types. Using estimates from the mixture modeling methodology, we define four measures of skill. In most of the analysis, we care about the probability that a finfluencer is un/anti/skilled and not directly about true alpha. For instance, a follower may want to identify skilled from un/antiskilled finfluencers. For such inference, the likelihood of having positive/negative alpha matters, not the expected value of alpha. Therefore, we compute the probability of being skilled/antiskilled/unskilled which can be calculated for each finfluencer i as

$$\begin{aligned} \Pr(\text{user } i \text{ skilled}) &\equiv \Pr(\alpha_i > 0 \mid \tilde{\alpha}_i) = \tfrac{\sum_{k=1}^{K^+} \pi_k^+ \eta(\tilde{\alpha}_i; \mu_k^+, \tilde{\sigma}_i)}{f_i(\tilde{\alpha}_i)}, \\ \Pr(\text{user } i \text{ antiskilled}) &\equiv \Pr(\alpha_i < 0 \mid \tilde{\alpha}_i) = \tfrac{\sum_{k=1}^{K^-} \pi_k^- \eta(\tilde{\alpha}_i; \mu_k^-, \tilde{\sigma}_i)}{f_i(\tilde{\alpha}_i)}, \\ \Pr(\text{user } i \text{ unskilled}) &\equiv \Pr(\alpha_i = 0 \mid \tilde{\alpha}_i) = 1 - \Pr(\alpha_i > 0 \mid \tilde{\alpha}_i) - \Pr(\alpha_i < 0 \mid \tilde{\alpha}_i), \end{aligned} \tag{8}$$

where $\eta(\tilde{\alpha}_i; \mu, \tilde{\sigma}_i)$ is the convolution of a Normal distribution with a mean of zero and standard deviation of $\tilde{\sigma}_i$ and an exponential distribution with a mean of μ evaluated at $\tilde{\alpha}_i$. In the denominator in (8), f_i is the distribution of $\tilde{\alpha}_i$. We obtain the probability of being unskilled by subtracting the probabilities of being skilled and antiskilled from one. Our last measure is the expected value of true alpha, $\mathbb{E}[\alpha_i \mid \tilde{\alpha}_i]$, for any finfluencer i conditional on the measured skill $\tilde{\alpha}_i$ which can be written as

$$\begin{aligned} \mathbb{E}[\alpha_i \mid \tilde{\alpha}_i] = \frac{1}{f_i(\tilde{\alpha}_i)} \Bigg(\int_{-\infty}^{0} \alpha \phi(\tilde{\alpha}_i; \alpha, \tilde{\sigma}_i) \left(-\sum_{k=1}^{K^-} \pi_k^- g(\alpha; \mu_k^-) \right) d\alpha \\ + \int_{0}^{\infty} \alpha \phi(\tilde{\alpha}_i; \alpha, \tilde{\sigma}_i) \left(\sum_{k=1}^{K^+} \pi_k^+ g(\alpha; \mu_k^+) \right) d\alpha \Bigg), \end{aligned} \tag{9}$$

where $\phi(\tilde{\alpha}_i; \alpha, \tilde{\sigma}_i)$ is a Normal distribution with a mean of α and standard deviation of $\tilde{\sigma}_i$. Expression (9) captures the best estimate of finfluencer skill given the data.

Table 3 documents descriptive statistics for the skill categories. The average probability that a user on StockTwits is skilled/unskilled/antiskilled is 29%/17%/55% with a standard

Table 3: Distribution of Finfluencer Skill

This table reports descriptive statistics on alternative measures of finfluencer skill. The probability of being skilled, $\Pr(\alpha_i > 0 \mid \tilde{\alpha}_i)$, is defined in (8). The probability of being unskilled, $\Pr(\alpha_i = 0 \mid \tilde{\alpha}_i)$, and the probability of being antiskilled, $\Pr(\alpha_i < 0 \mid \tilde{\alpha}_i)$, are defined accordingly. The expected value of true alpha is defined in (9). FMM stands for Finite Mixture Models procedure.

Panel A: Distribution of finfluencer skill

	N	Mean	S.D.	Min	p10	p50	p90	Max
Pr(user i skilled)	29,475	0.29	0.22	0.00	0.04	0.25	0.55	1.00
Pr(user i unskilled)	29,475	0.17	0.08	0.00	0.04	0.18	0.24	0.81
Pr(user i antiskilled)	29,475	0.55	0.23	0.00	0.25	0.55	0.87	1.00
True alpha $\mathbb{E}[\alpha_i \mid \tilde{\alpha}_i]$	29,475	-0.60	3.93	-48.63	-2.09	-0.32	1.02	41.91

Panel B: Finfluencer classification

	Classification based on		
	FMM	Pr > 1/3	max. Pr
Skilled	0.29	0.29	0.20
Unskilled	0.17	0.02	0.01
Antiskilled	0.55	0.85	0.79

deviation equal to 22%/8%/23%. The left subplot of Figure 1 shows histograms of the probabilities that users are skilled, unskilled, and antiskilled. The plot reveals that there exists a lot of dispersion in the probability of being a skilled or antiskilled StockTwits user. It is evident from the plot that less than 3% of StockTwits users are unambiguously skilled, and the second column of Panel B in Table 3 confirms that the majority of StockTwits users have a probability of less than 1/3 of being skilled. Skilled finfluencers deliver unambiguously positive returns, as the right subplot of Figure 1 shows.

Table 3 also indicates that the distribution of $\Pr(\alpha_i = 0 \mid \tilde{\alpha}_i)$ is tight, and the left subplot of Figure 1 confirms this observation. The second column of Panel B in Table 3 shows that the majority of StockTwits users have a low probability of being unskilled, as 98% of them have a probability of less than 1/3 of being unskilled. Panel B of Table 3 shows that the vast majority of StockTwits users can be classified as antiskilled, as 85% of them have a probability of more than 1/3 of being antiskilled. Similarly, the left subplot of Figure 1 shows that the majority of users have a probability over 50% of being antiskilled, while the right subplot of the same figure shows that almost 75% of antiskilled users deliver unambiguously negative returns. Finally, based on the maximum of the probabilities of being skilled, unskilled, or antiskilled, one can classify 20% of finfluencers as being skilled, 1% of

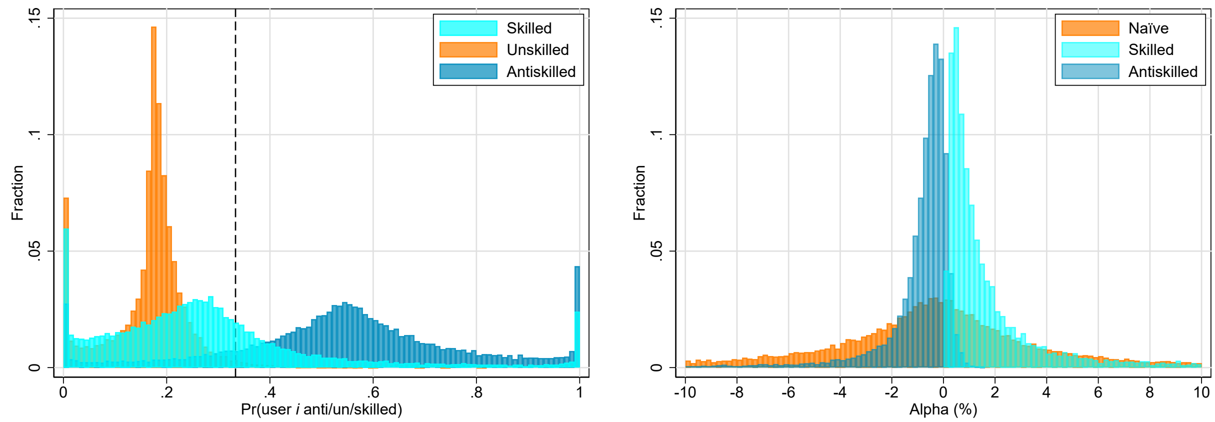

Figure 1: Distribution in Users' Probability of Being Un/Anti/Skilled and True Alphas
The plots show histograms of the probabilities of users being skilled, unskilled, and antiskilled, respectively, and the expected value of true skill, $\mathbb{E}[\alpha_i \mid \tilde{\alpha}_i]$.

finfluencers as being unskilled, and 79% of finfluencers as being antiskilled.

The last row of Panel A demonstrates that the average monthly true alpha, $\mathbb{E}[\alpha_i \mid \tilde{\alpha}_i]$, among finfluencers is equal to -60bps with a standard deviation of 3.93%, indicating a large dispersion in the true alpha among them. This dispersion is mainly due to the left tail of the distribution since the bottom 10% of users generate alpha of -2.09% or less per month, while the top 10% of users generate alpha of 1.02% or more per month. Consequently, the right subplot of Figure 1 shows the distribution of true alphas among skilled and, respectively, antiskilled finfluencers (classified using the 1/3 rule). Most skilled finfluencers have a true alpha of less than 4%, with a peak of 0.2%. Most antiskilled finfluencers have a true alpha of more than -4%, with a peak at -0.3%.

These results indicate that the majority of finfluencers on StockTwits are un/antiskilled. This is relevant since the content of the antiskilled tweets is informative in the sense of "do the opposite of what I say." It raises the question of which finfluencers are more popular, what tweeting patterns explain finfluencers' true alphas, whether skill levels are detectable, and whether finfluencers' tweets distort belief aggregation.

Persistence in (anti)skill. Given the large share of un/antiskilled finfluencers, it is important to first check how skill levels persist or change over time. We explore the dynamics in finfluencer composition by splitting the sample and re-estimating skill in both periods. We then check the skill of incumbent finfluencers over time as opposed to newly joined ones. We first use the first year of the sample to calculate each finfluencer's alpha and then use

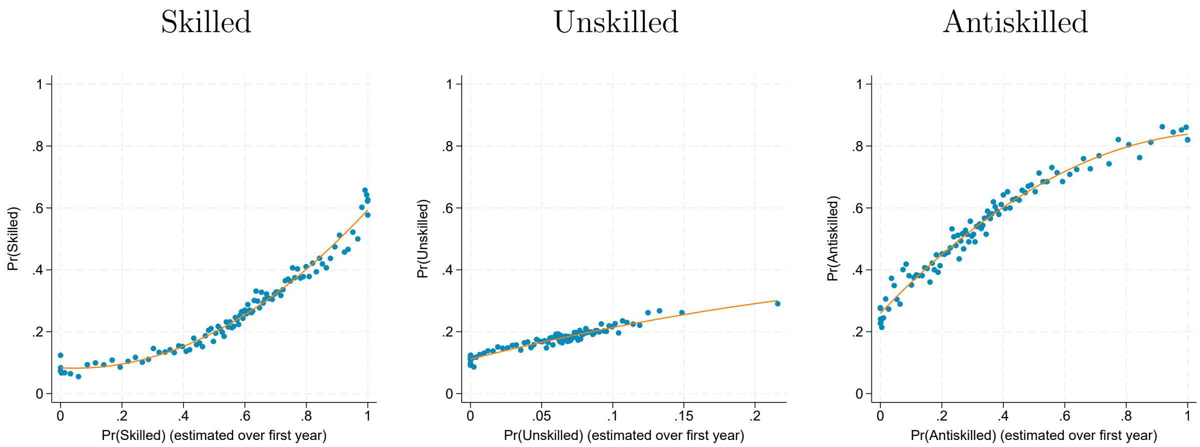

Figure 2: Persistence of Finfluencer Skill

The plots show binscatter plots of the probability of being skilled/unskilled/antiskilled estimated over the first year versus the probability of being skilled/unskilled/antiskilled estimated over the entire sample period.

these alphas to extract finfluencers' skills via mixture modeling to calculate the probabilities of being skilled, unskilled, and antiskilled. We then repeat this procedure using the whole sample and plot the probability of being skilled/unskilled/antiskilled estimated over the whole sample against the corresponding probability estimated over the first year in Figure 2. Left/middle/right plots show the results for skilled/unskilled/antiskilled finfluencers. All three plots are monotonically increasing functions implying all skill types are persistent. Appendix 1.4 documents statistics on finfluencer survival, finding that particularly the unskilled finfluencers are most likely to be present throughout the sample.

The next section develops a model of finfluencers that provides guidance on the relationship between the skill and popularity of finfluencers and the beliefs revealed by the social media activity of different types of finfluencers. These relations are important for assessing the overall quality of financial advice, the lead-lag relationship with returns and more generally the nature of competition in social media.

2 Model of Finfluencers

To motivate further analysis of the relationship between the finfluencers' skill, their social media behavior, and stock prices we consider the following stylized model of finfluencers.

2.1 Assumptions

Time is discrete, indexed by $t = 0, 1, 2, \ldots$. The economy has J stocks indexed by $j = 1, \ldots, J$. In period t, each stock j has a random return $\tilde{R}_{j,t+1}$ realized in period $t+1$. The economy is populated by two types of agents: a finite mass of finfluencers indexed by $i \in [0, 1]$ and infinitely many small followers/investors indexed by k. At every t, finfluencer i can tweet about *any* subset of stocks, possibly all stocks J. Let $\text{Tweet}_{i,j}(t)$ be an indicator variable equal to one if finfluencer i tweets about stock j in period t, and zero otherwise. Then the aggregate tweeting activity by finfluencer i in period t is her tweet volume $v_i(t) = \sum_{j=1}^{J} \text{Tweet}_{i,j}(t)$.

A period-t tweet by finfluencer i about stock j is the finfluencer's opinion captured by the sentiment $SocSent_{i,j,t} \in \{-1, 0, +1\}$ about the next period's return $\tilde{R}_{j,t+1}$. $SocSent_{i,j,t} = -1$ implies that finfluencer i is pessimistic about stock j at t, thinks $\tilde{R}_{j,t+1} < 0$. $SocSent_{i,j,t} = 0$ implies finfluencer i has a neutral opinion about stock j at t, thinks $\tilde{R}_{j,t+1} = 0$. Finally, $SocSent_{i,j,t} > 0$ means finfluencer i is bullish about stock j at t, thinks $\tilde{R}_{j,t+1} > 0$.

Motivated by our results from the previous section, we partition finfluencers into three types based on their true ability to forecast $\tilde{R}_{j,t+1}$, $\tau_i \in \{\text{Skilled}, \text{Unskilled}, \text{Antiskilled}\}$. A fraction $\Pr(\tau_i = S) = P_S$ of finfluencers is *Skilled* (type S). Skilled finfluencers' are informed and the content of their tweets is positively correlated with $\tilde{R}_{j,t+1}$, i.e., they tweet $SocSent_{i,j,t} = +1$ and $\tilde{R}_{j,t+1} > 0$ is likely. A fraction $\Pr(\tau_i = U) = P_U$ of finfluencers is *Unskilled* (type U). Unkilled finfluencers' are uninformed and the content of their tweets is uncorrelated with $\tilde{R}_{j,t+1}$, i.e., $SocSent_{i,j,t} = +1$ and $\tilde{R}_{j,t+1} > 0$ or $\tilde{R}_{j,t+1} = 0$ or $\tilde{R}_{j,t+1} < 0$ are all equally likely. A fraction $\Pr(\tau_i = A) = P_A$ of finfluencers is *Antiskilled* (type A). The content of antiskilled finfluencers' tweets is negatively correlated with $\tilde{R}_{j,t+1}$, i.e., they tweet $SocSent_{i,j,t} = +1$ and $\tilde{R}_{j,t+1} < 0$ is likely. Here we are agnostic about why this is the case, while Appendix C dissects finfluencers' tweeting strategies in the data. By construction, we have $P_S + P_U + P_A = 1$.

Followers' problem. Followers are social media users consisting of incumbents and entrants and both evolve endogenously over time. Followers derive utility from following finfluencer i based on their beliefs $\theta_{k,i}(t)$ about whether finfluencer i is truly skilled as opposed to unskilled or antiskilled, and the finfluencer's visibility, denoted $V_i(t)$. Visibility captures the degree to which finfluencers stimulate the followers' engagement via activity levels, tweeting strategies, and other entertainment aspects. The utility for a follower k derived from following and engaging finfluencer i at time t is

$$U^F_{ki}(t) = \theta_{k,i}(t) \cdot U_S + (1 - \theta_{k,i}(t)) \cdot U_B + \gamma V_i(t), \tag{10}$$

where the coefficient $\gamma \geq 0$ captures the followers' preference for visibility. The parameters U_S and U_B satisfy $U_S > U_B$.

Let $F_i(t)$ denote the number of followers of finfluencer i at time t. There is a constant inflow of followers per unit of time, $F_{\text{new}}(t)dt$. A new follower arriving at time t chooses finfluencer i with probability

$$P_{ki}(t) = \frac{U^F_{ki}(t)}{U^F_{ki}(t) + \max_{m \neq i}\{U^F_{km}(t)\}}, \tag{11}$$

where, similarly, $U^F_{km}(t)$ is the utility from following finfluencers m.[24] The growth rate of followers of finfluencer i is proportional to the total inflow of followers and the probability that new entrants choose finfluencer i,

$$\frac{dF_i(t)}{dt} = F_{\text{new}}(t) \times \int_{k \in \Omega_{\text{entrants}}} P_{ki}(t)dk, \tag{12}$$

where Ω_{entrants} is the set of new entrants.

Note that the followers' utility (10) that determines the choice probabilities and follower dynamics does not depend on the financial payoffs, captured by the signed returns from following finfluencer $i \in \{S, U, A\}$ recommendation $SocSent_{i,j,t}$ about stock j,

$$\Delta^{(i)}_{j,t+1} = SocSent_{i,j,t} \cdot R_{j,t+1}, \tag{13}$$

where $R_{j,t+1}$ is the period $t+1$ stock j's realized return, rendering them slow learners. Over many periods, the follower's realized performance from adopting i's recommendations reveals that skilled finfluencers on average generate positive $\Delta^{(i)}_{j,t+1}$, antiskilled finfluencers generate negative $\Delta^{(i)}_{j,t+1}$, and unskilled finfluencers generate zero $\Delta^{(i)}_{j,t+1}$. Therefore, over time, the followers' beliefs $\theta_{k,i}(t)$ should shift if they place sufficient weight on the realized market performance. However, a sufficiently large $\gamma V_i(t)$ in utility (10) can still dominate the long-run alpha evidence, letting un/antiskilled finfluencers remain popular if they generate high

[24]Results are robust to using the logit or softmax form, $P_{ki}(t) = \frac{\exp(\phi U^F_{ki}(t))}{\sum_m \exp(\phi U^F_{km}(t))}$, with precision parameter $\phi > 0$.

engagement. We therefore omit the market performance metric from the followers' utility to keep the model parsimonious.

Finfluencers' problem. Followers care about finfluencers' visibility, $V_i(t)$. Finfluencer i's visibility at time t depends on the number of followers, $F_i(t)$, and the engagement effort (or just the engagement) the finfluencer needs to exert, $e_i(t)$, so that $V_i(t) = \kappa \cdot e_i(t) \cdot F_i(t)$. We define the engagement by finfluencer i as

$$e_i(t) = \underbrace{e_0}_{\text{Baseline engagement}} + \alpha_1 v_i(t) + \alpha_2 \sum_{j=1}^{J} |SocSent_{i,j,t}| \cdot \text{Tweet}_{i,j}(t), \tag{14}$$

where $e_0 > 0$, $\alpha_1 > 0$, and $\alpha_2 > 0$. The engagement per tweet above the baseline engagement is $\alpha_1 + \alpha_2 |SocSent_{i,j,t}|$ with the engagement α_1 required to generate a tweet and engagement α_2 required to generate non-neutral tweet content. Finfluencer $i \in \{S, U, A\}$ chooses her optimal engagement by maximizing her utility function

$$U_i(t) = \pi_i \cdot F_i(t) \cdot \Big[f\big(v_i(t)\big) + g\big(\{SocSent_{i,j,t}, \text{Tweet}_{i,j}(t)\}\big)\Big] - \frac{c_i}{2} e_i(t)^2. \tag{15}$$

$\pi_i > 0$ in utility (15) is the type-dependent reward per follower capturing how each follower's clicks, likes, and re-shares translate to direct or indirect benefits for the finfluencer. $c_i > 0$ in utility (15) is the type-dependent cost parameter for total engagement $e_i(t)$. The function $f(\cdot)$ captures the reward for the total tweeting activity and increases in tweet volume $v_i(t)$. Finally, the function $g(\cdot)$ can incorporate how *positive* or *negative* tweets about a stock's payoff yield additional ephemeral engagement. For instance, consistently bullish tweets might yield more short-run excitement from retail traders, while contrarian or negative tweets may do so only if they are skilled contrarian calls. In general, we can let

$$\begin{aligned} g\big(\{SocSent_{i,j,t}, \text{Tweet}_{i,j}(t)\}\big) = \sum_{j=1}^{J} \Big[& \beta_S \mathbb{1}_{\{\tau_i = S\}} SocSent_{i,j,t} \\ & + \beta_A \mathbb{1}_{\{\tau_i = A\}} SocSent_{i,j,t} + \beta_U \mathbb{1}_{\{\tau_i = U\}} SocSent_{i,j,t} \Big] \cdot \text{Tweet}_{i,j}(t). \end{aligned} \tag{16}$$

Maximizing utility (15) with respect to $\{v_i(t), SocSent_{i,j,t}\}$ yields finfluencer i's optimal tweeting strategy $\{SocSent^*_{i,j,t}, \text{Tweet}^*_{i,j}(t)\}_{j=1}^{J}$ and tweet volume $v^*_i(t)$. These optimal policies yield the following optimal engagement $e_i^*(t)$

$$e_i^*(t) = \frac{\pi_i}{c_i} \cdot F_i(t) \cdot \Phi\big(v_i^*(t), \{SocSent_{i,j,t}^*, \text{Tweet}_{i,j}^*(t)\}\big), \tag{17}$$

for some function Φ. In the proofs, we assume that each finfluencer's optimal tweeting volume and tweeting strategy cannot diverge to infinity (nor vanish to zero in a way that destroys visibility). Formally, we require function Φ to have lower bound $\underline{\Phi}$ and upper bound $\overline{\Phi}$ such that $0 < \underline{\Phi} \leq \Phi_i(t) \leq \overline{\Phi} < \infty$ for all t. One can ensure this, for instance, by imposing concavity in the finfluencer's utility so that tweet volume does not explode, or by explicitly bounding it. Relation (17) yields the following expression for the optimal visibility

$$V_i^*(t) = \kappa \cdot \frac{\pi_i}{c_i} \cdot F_i(t)^2 \cdot \Phi\big(v_i^*(t), \{SocSent_{i,j,t}^*, \text{Tweet}_{i,j}^*(t)\}\big). \tag{18}$$

This completes the model formulation. We proceed to discuss the model's implications, delegating all other details to Internet Appendix IB.

2.2 Model implications

The model raises several questions: How efficient is social media at weeding out un- and antiskilled finfluencers? Which finfluencers are more popular and influential, and why? What are finfluencers' tweeting strategies and how do they attract larger followings? What are the aggregate beliefs (and biases) associated with the most visible finfluencers? How are different types of finfluencers' activity and sentiment related to returns?

The following proposition proved in Internet Appendix IB summarizes our main model result on the number of followers.

Proposition 1. *Fix two finfluencers, i_1 and i_2, with type-dependent parameters (π_{i_1}, c_{i_1}) and (π_{i_2}, c_{i_2}). Suppose that initially $\frac{\pi_{i_1}}{c_{i_1}} F_{i_1}(t') > \frac{\pi_{i_2}}{c_{i_2}} F_{i_2}(t')$ for some t'. Then there exists a time $T > t'$ such that for all $t > T$ $F_{i_1}(t) > F_{i_2}(t)$, i.e., if finfluencer i_1 has a strictly higher reward-to-cost ratio and a large enough follower base at some date t', then i_1 eventually dominates i_2 in follower count.*

Let $\tau_{i_1} = A$ and $\tau_{i_2} = S$. Proposition 1 demonstrates that given the higher reward-to-cost ratio for the antiskilled finfluencer, the antiskilled finfluencer(s) will eventually surpass the skilled finfluencer(s) in follower count, provided she has a sufficiently large initial follower base or favorable engagement dynamics.

With the constant inflow of new followers, the belief updating process continues but the dominance of the antiskilled finfluencer will still skew the aggregate belief over time. The aggregate visibility-weighted fraction of skilled finfluencers $\bar{\theta}(t)$ is equal to

$$\bar{\theta}(t) = \frac{P_S \cdot V_S(t)}{P_S \cdot V_S(t) + P_A \cdot V_A(t) + (1 - P_S - P_A) \cdot V_U(t)}, \tag{19}$$

where $V_i(t)$ is the equilibrium visibility of finfluencer $i \in \{S, A, U\}$ at time t. In the long run, the visibility of antiskilled finfluencers will dominate due to their higher engagement levels and growth rates, as established previously. This leads to a skewed perception among followers regarding the quality of finfluencers. The following proposition proved in Internet Appendix IB summarizes these results.

Proposition 2. *Consider a continuum of finfluencers, some skilled* (S)*, some unskilled* (U)*, and some antiskilled* (A)*, with type-dependent parameters* $\{\pi_S, \pi_U, \pi_A\}$ *and* $\{c_S, c_U, c_A\}$*. Suppose* $\frac{\pi_A}{c_A} > \frac{\pi_S}{c_S}$ *and that antiskilled finfluencers have a sufficiently large initial follower base or sufficiently strong short-run sentiment payoff. Then in the limit* $t \to \infty$ *the visibility-weighted sentiment of stock* j *is dominated by antiskilled finfluencers*

$$\underbrace{\sum_{i:\tau_i = A} w_{i,j}(t) SocSent_{i,j,t}}_{\textit{antiskilled part}} \quad \gg \quad \underbrace{\sum_{i:\tau_i = S} w_{i,j}(t) SocSent_{i,j,t}}_{\textit{skilled part}}, \tag{20}$$

where $w_{i,j}(t) = \frac{V_i(t) \cdot Tweet_{i,j}(t)}{\sum_m V_m(t) \cdot Tweet_{m,j}(t)}$*. Consequently, the aggregate sentiment is biased toward incorrect signals, and* $\bar{\theta}(t)$*—the perceived fraction of skilled finfluencers—as measured by visibility weighting, converges to zero.*

Proposition 2 shows that, over time, the aggregate belief $\bar{\theta}(t)$ about the fraction of skilled finfluencers decreases and eventually approaches zero. This leads to a significant bias in aggregate beliefs, where followers underestimate the true proportion of good finfluencers (or the probability of a finfluencer being skilled) on the platform.

The model offers guidelines along two dimensions: which variables are relevant and which relationships should be inspected. It emphasizes the importance of the influencer's type, tweeting strategy, and popularity. The influencer's type is not directly observable and has to be extracted from granular tweet data. All other parameters and characteristics from the model are directly observable in the data. The following predictions from the model can be tested in our data.

1. **Persistent Misallocation.** It follows from Proposition 1 that if antiskilled finluencers have a higher reward-to-cost ratio and a larger follower base than skilled finfluencers, they would eventually dominate the media platform despite monetary losses from following their recommendations. We test these predictions in Sections 1.4 and 3.

2. **Abnormal Sentiment.** *Abnormal social sentiment* for stock j is the difference between the *visibility-weighted* sentiment of a certain group (e.g. "visible" or "antiskilled") and that of *skilled* finfluencers. Proposition 2 shows that this difference may remain strictly positive or negative for extended periods. We quantify this aggregate bias empirically in the first part of Section 4.

3. **Predicting Returns.** By the model's construction, skilled types tweet more accurately leading to a correct sign for future returns, $R_{j,t+1}$. Antiskilled types do the opposite. Hence, measuring each group's sentiment over time should reveal that skilled sentiment predicts *positive* future returns while the aggregated "visible" sentiment from high-engagement, antiskilled types *negatively* predicts future returns. We test these predictions using panel vector-autoregressions that capture the joint endogeneity of returns, activity, and sentiment in the second part of Section 4.

We now explore these predictions in the data.

3 Finfluencer Popularity and Follower Engagement

Proposition 1 raises an important question of which finfluencers are more popular and influential. To address this question, we explore which categories of finfluencers have more followers and are more active and engaging. We also explore what tweeting strategies finfluencers pursue to create user engagement, and how engagement relates to followings.

3.1 Finfluencer popularity: Do more skilled finfluencers have a larger follower base?

If finfluencers provide solid financial advice and followers value this advice, we would expect skilled finfluencers to have a larger follower base than un/antiskilled finfluencers. In this case, the market for financial advice by finfluencers works efficiently. In this regard, the large share of un- and antiskilled finfluencers documented in Section 1.3 appears surprising. However, according to Proposition 1 if antiskilled finluencers have a higher reward-to-cost ratio and a larger follower base than skilled finfluencers, they would eventually dominate the media platform despite monetary losses from following their recommendations. Skilled finfluencers can also have fewer followers than unskilled or antiskilled ones if social media users like to follow finfluencers for reasons unrelated to the quality of their advice, due to behavioral traits and homophily (Currarini, Jackson, and Pin, 2009; Golub and Jackson, 2012), or/and if they have a preference for visibility and engagement, and un- and antiskilled finfluencers would be more active and engaged with their followers. Yet another alternative is that, if finfluencers build a reputation by revealing valuable information and stop doing so once they have acquired a large body of followers, we expect an ambiguous relation between skill and popularity and our skill measures to not be persistent (Benabou and Laroque, 1992).

Given our split of finfluencers into skilled, unskilled, and antiskilled, we start by asking whether the crowd of StockTwits users follows the skilled finfluencers. If so, we expect skilled finfluencers to have more followers than unskilled finfluencers, at least longer term. Figure 3 documents the relation between the number of followers and our measures of finfluencers' skill. We measure finfluencers' followers by the log of overall follower count in February 2018 after the tweet sample has ended. Finfluencers' follower counts are thus measured out-of-sample. The left binscatter plot shows that the follower count is *negatively* related to finfluencers' probability of being skilled. By contrast, the right binscatter plot shows a strong positive relation between finfluencers' followers and the probabilities of being antiskilled. This means skilled finfluencers have fewer followers than unskilled or antiskilled finfluencers. These relations are highly statistically significant, even if we measure followers out-of-sample after our tweet sample has ended.

3.2 Follower engagement: Tweeting activity and finfluencer skill

One way to create user engagement and increase follower count is to tweet more actively. It seems reasonable to expect that tweeting activity affects user attention and is related to the informativeness of tweets. Finfluencers who tweet more often may be more likely to be experts, allowing them to build a reputation for having more valuable information. On the other hand, finfluencers who tweet more often may be more confident or a "charlatan" who believes that a large tweeting volume proxies for skill. Thus, their tweets might be less informative. Ultimately, how informed frequent tweeters are is an empirical question. In the model, user engagement is endogenous and important since the visibility of antiskilled

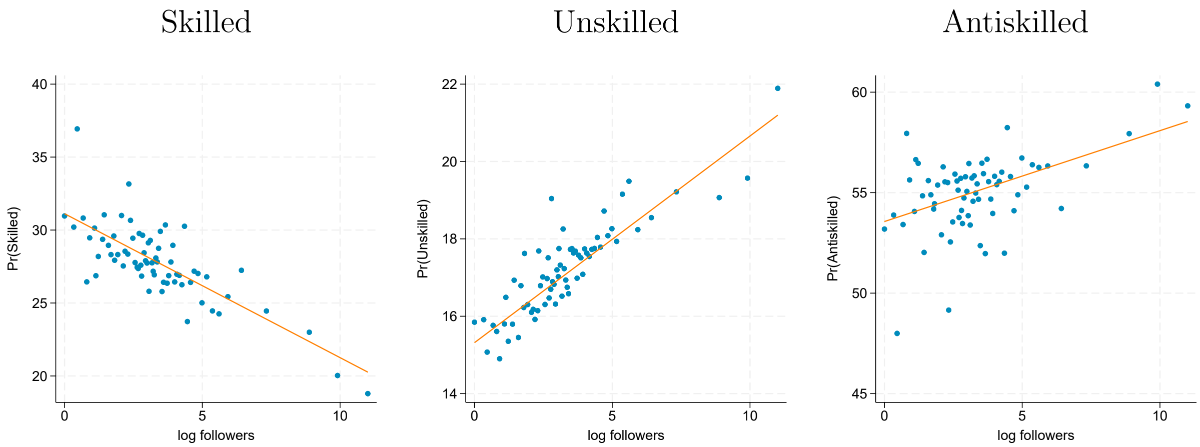

Figure 3: Finfluencer Skill versus Number of Followers

The plots show binscatter plots of the probabilities of users being skilled, unskilled, and antiskilled, respectively, versus the natural logarithm of the number of followers.

finfluencers will dominate the visibility of skilled finfluencers if the former are more engaged.

Figure 4 documents the relation between finfluencers' tweeting activity and our measures of their skill. Tweeting activity is captured by *log activity* defined as the log of one plus the total number of positive and negative tweets the user has posted. The right binscatter plot shows a strong positive relation between finfluencers' tweeting activity and the probability of being antiskilled. By contrast, tweeting activity is very strongly negatively related to finfluencers' probability of being skilled. These findings indicate that the marginal power of the additional tweet in identifying the user's skill declines with the number of tweets for skilled finfluencers or, in other words, that the tweeting activity exhibits decreasing identification power for skilled finfluencers.

3.3 Link between finfluencer popularity and follower engagement

While Figure 3 has documented a counter-intuitive but consistent with Proposition 1 relation between the number of followers and our measures of finfluencers' skill, Figure 4 has documented that the same relations hold between finfluencers' tweeting activity and our measures of their skill. Figure 5 now explores the resulting relation between the number of followers and finfluencers' tweeting activity. Consistently across skill groups and in agreement with the model, finfluencers with more followers are more active, suggesting that user engagement and followings are positively related to one another for all skill groups.

To better understand why un/antiskilled finfluencers are more popular than skilled ones,

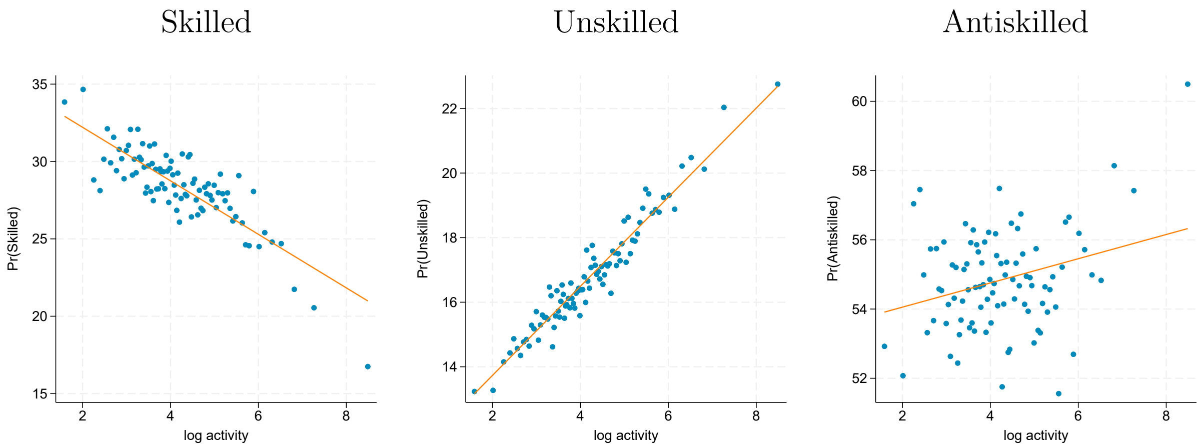

Figure 4: Finfluencer Skill versus Tweeting Activity

The plots show binscatter plots of the probabilities of finfluencers being skilled, unskilled, and antiskilled, respectively, versus tweeting activity. We measure tweeting activity by the natural logarithm of the number of tweets.

we now drill into the determinants of popularity by linking it to finfluencers' tweeting activity and strategies. We relate various characteristics of tweeting strategies to the finfluencers' out-of-sample follower count. Table 4 reports results from regressing the number of followers for each finfluencer on tweeting activity, the fraction of positive tweets, and various characteristics of tweeting strategies, including return chasing, the composition of tweets, herding, and short-selling constraints:

$$\text{Finfluencer's follower count}_i \text{ (measured out-of-sample)} = \alpha + \beta \cdot \mathit{TweetingActivity}_i + \\ + \delta \cdot \mathit{FractionPositive}_i + \gamma^\top \mathit{TweetingStrategy}_i + \epsilon_i, \quad (21)$$

where the dependent variable is again the log of one plus the finfluencer's follower count as of February 2018, $\mathit{TweetingActivity}_i$ is the log of one plus the total number of positive and negative tweets by finfluencer i, and $\mathit{TweetingStrategy}_i$ is one of several tweeting/investment behaviors. Appendix C where we dissect finfluencers' tweeting strategies provides detailed definitions of each tweeting strategy that we consider in specification (21).

Results from Table 4 confirm results from Figure 5. More active tweeting and more positive messages/tweets attract followers. A one percent increase in the total number of tweets is associated with a 0.72-0.74% increase in followers. The correlation between the share of positive tweets and the number of followers is positive and significant but small in economic magnitude. Table 4 also shows that both the tendency to chase returns and

Table 4: Effect of finfluencers' tweeting patterns on follower count

This table reports the results of regressing the number of followers on finfluencers' tweeting characteristics. The dependent variable is the log of one plus the user's follower count as of February 2018. The independent variables are the same as in Table C.1. Standard errors are robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Finfluencer's follower count$_i$ (out-of-sample)				
	(1)	(2)	(3)	(4)	(5)
$TweetingActivity_i$	0.72***				0.74***
	(0.01)				(0.01)
$FractionPositive_i$	0.00***				0.00***
	(0.00)				(0.00)
$ReturnChasing_i$		-0.00***			-0.00**
		(0.00)			(0.00)
$ContrarianTweet_i$		-0.01***			-0.00***
		(0.00)			(0.00)
$PositiveHerding_i$			0.00***		-0.00*
			(0.00)		(0.00)
$NegativeHerding_i$			-0.01***		-0.01***
			(0.00)		(0.00)
SSI_i *(Positive Tweets)*				-0.01***	-0.03***
				(0.00)	(0.00)
SSI_i *(Negative Tweets)*				-0.00	-0.01***
				(0.00)	(0.00)
Constant	-0.03	1.34***	1.52***	1.29***	0.57***
	(0.02)	(0.01)	(0.02)	(0.01)	(0.03)
N	19,593	19,593	19,593	19,593	19,593

post contrarian tweets correlates negatively with the finfluencer's follower count. Moreover, herding on positive tweets is positively correlated with the follower count, but the sign switches when we control for other user characteristics, and its magnitude shrinks. Herding on negative tweets is negatively correlated with the follower count. Finally, tweeting about stocks with higher short-selling constraints negatively correlates with the number of followers regardless of the tweet sentiment.

These results provide important insights into the inner workings of user engagement in the data not fully represented in our stylized model. They suggest that un/antiskilled finfluencers create user engagement in two ways, by tweeting more, captured by the model, and following tweeting strategies that are popular with their follower base, missing from the model. When we explore the tweeting strategies that finfluencers pursue to attract user attention, we find that skilled finfluencers are return-, social sentiment-, and news-

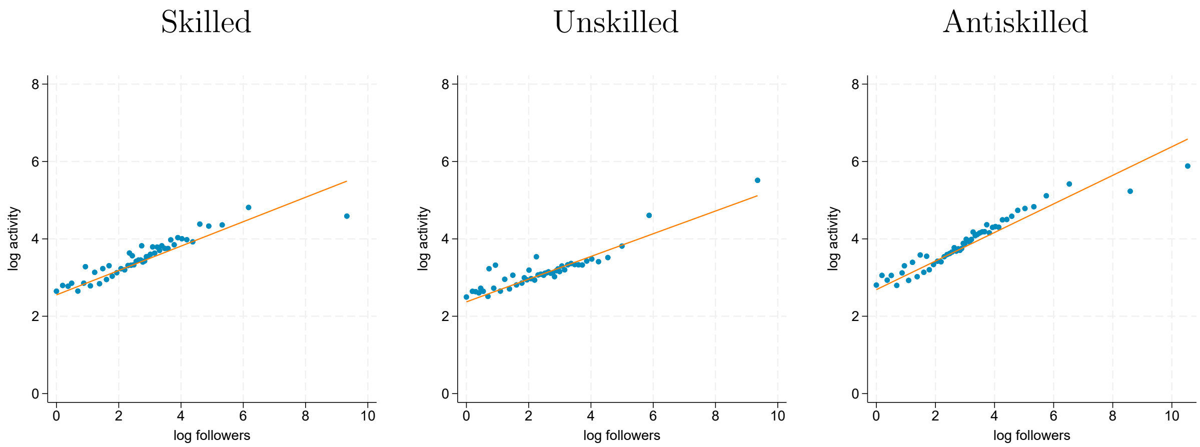

Figure 5: Finfluencer Activity versus Number of Followers

The plots show binscatter plots of tweeting activity versus the natural logarithm of the number of followers. We measure tweeting activity by the natural logarithm of the number of tweets. We classify finfluencers as skilled, unskilled, and antiskilled, depending on the estimated probabilities.

contrarian, that is they go against the consensus of social sentiment and news coverage and recommend stocks that have gone down and, hence, tweets by skilled finfluencers provide valuable information about underpriced stocks. By contrast, antiskilled chase returns and ride return and social sentiment momentum. This means that they recommend stocks that have already gone up, which leads to buying stocks at inflated prices and pushing prices even further up. Taken together, this suggests antiskilled finfluencers systematically push past winners.

An implication of these findings is that skill could be identified, even if imperfectly, by looking at finfluencers' followers, based on how active is the finfluencer and what tweeting strategy the finfluencer pursues. However, social media users seem to insufficiently use the available information to pick skilled finfluencers to follow, or they decide to ignore it. This inefficient matching between the finfluencers and their followers can be explained by Proposition 1 according to which given the higher reward-to-cost ratio for the antiskilled finfluencer, the antiskilled finfluencers may surpass the skilled finfluencers in follower count, provided they have a sufficiently large initial follower base or favorable engagement dynamics. Alternative explanations for the observed matching between finfluencers and users are that antiskilled finfluencers' tweeting strategies are most popular with their followers and that (anti)skill is not persistent or finfluencers first build a reputation and later exploit it (Benabou and Laroque, 1992).

Overall, the results suggest that, even though finfluencer skill is persistent, skilled finfluencers are not more popular and less likely to survive (that is, stay active on StockTwits in the long term) than unskilled and antiskilled finfluencers. In the next section, we investigate whether the apparent misallocation of more followers to less skilled finfluencers can induce significant belief biases among the follower base and, if so, in which direction.

4 Failure in the Wisdom of the Crowd

The findings in the previous sections raise the question of what aggregate beliefs are induced by following un/anti/skilled finfluencers and how stock returns, finfluencer activity, and sentiment are jointly related.

4.1 Beliefs induced by skilled vs. visible finfluencers

Based on the model, we assess the aggregate beliefs resulting from the tweets of various types of finfluencers. We approximate these beliefs by first aggregating the tweets at the level of all visible, unskilled, and antiskilled finfluencers for each stock day. We then compare the group-level beliefs to the tweets of skilled finfluencers as a reference group. We then have a measure of abnormal belief, or "belief bias." The assumption here is that skilled finfluencers produce the most accurate information and that any stock-specific and time-specific confounding factors are uncorrelated with systematic patterns in their tweeting activity. The underlying idea is that true information can be filtered out by netting out skilled finfluencers' average sentiment. This takes care of potential concerns, for instance, about the fact that our sample period experienced positive average market returns.

It remains to determine the weight of each finfluencer in aggregating social sentiment across tweets on a given day. The model suggests that visibility, $V_i(t) = \kappa \cdot \frac{\pi_i}{c_i} \cdot F_i(t)^2 \cdot V_i^*(t) = \kappa \cdot \frac{\pi_i}{c_i} \cdot F_i(t)^2 \cdot \Phi(\cdot)$, is increasing in the number of followers, $F_i(t)$, and since $F_i(t) \propto \frac{c_i}{\pi_i} \cdot e_i^*(t)$ the visibility is a quadratic function of the engagement, $V_i(t) \propto \kappa \cdot \frac{c_i}{\pi_i} \cdot e_i^*(t)^2$. Therefore for each finfluencer i we use her social sentiment $SocSent_{i,j,t}$ aggregated across all her tweets on StockTwits as a measure of visibility $Visibility_i$.

We compute proxies for aggregate finfluencer-induced beliefs, which we call $VisibleAbnSent_t$, $UnskilledAbnSent_t$, and $AntiskilledAbnSent_t$, respectively. First, we aggregate for each

Panel A: Visibility-minus-skill weighted abnormal social sentiment by day

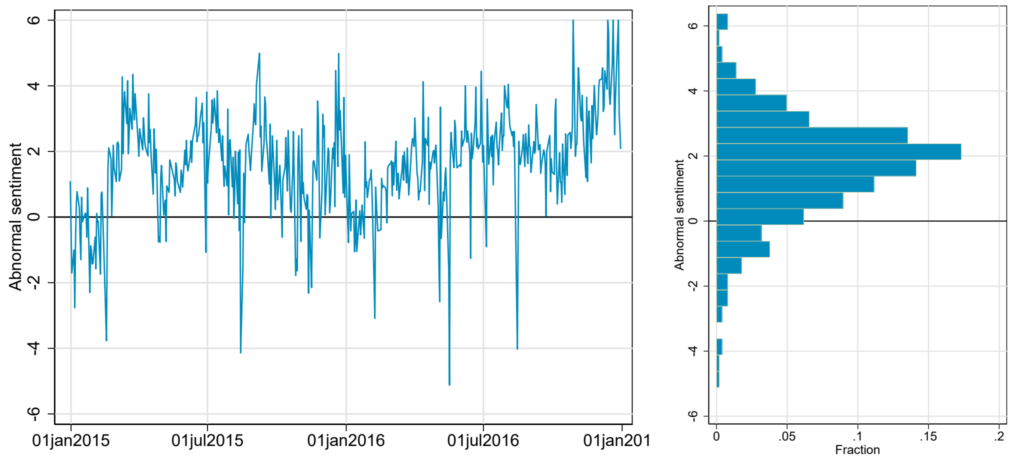

Panel B: Finfluencer type-weighted abnormal social sentiment by day

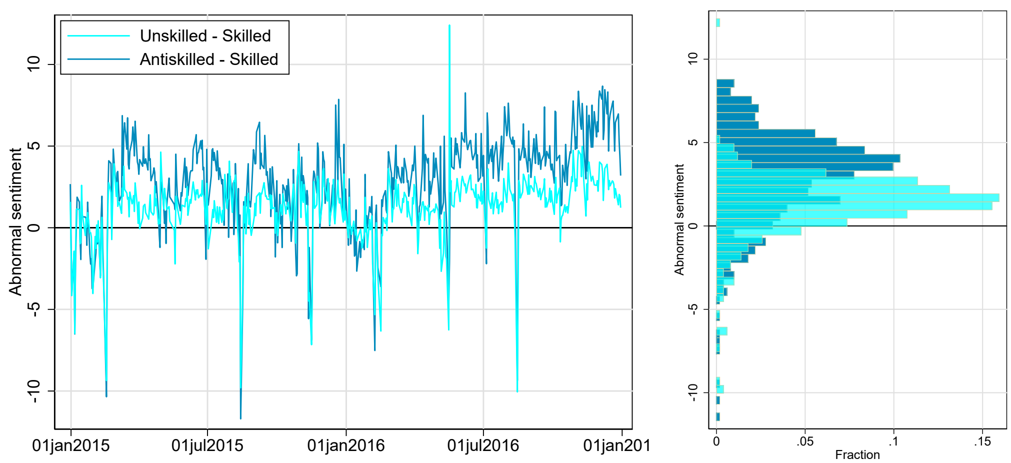

Figure 6: Abnormal Social Sentiment

The plot in Panel A shows the daily average visibility minus skill-weighted average social sentiment. The plot in Panel B shows the daily average social sentiment by unskilled and antiskilled finfluencers, respectively, net of the daily average social sentiment by skilled finfluencers.

stock j and day t across finfluencers i:

$$\begin{aligned} VisibleSent_{j,t} &= \frac{\sum_{\text{all } i} Visibility_i \times SocSent_{i,j,t}}{\sum_{\text{all } i} Visibility_i \times Active_{i,j,t}}, \\ Un/Anti/SkilledSent_{j,t} &= \frac{\sum_{\text{all } i} \Pr(\text{user } i \text{ un/anti/skilled}) \times SocSent_{i,j,t}}{\sum_{\text{all } i} \Pr(\text{user } i \text{ un/anti/skilled}) \times Active_{i,j,t}}, \end{aligned} \tag{22}$$

where $Visibility_i$ is our measure of visibility for each finfluencer i, Pr(user i un/anti/skilled) are given by expressions (8) and $SocSent_{i,j,t}$ is given by expression (1). To control for finfluencers' temporal focus, $Active_{i,j,t}$ is an indicator if finfluencer i tweets about stock j on

Table 5: Abnormal Social Sentiment Revealed by the Tweets of More Visible Finfluencers

This table reports descriptive statistics about the abnormal social sentiment revealed by the tweets of visible and un/antiskilled relative to skilled finfluencers by day.

	N	Mean	S.D.	Min	p10	p50	p90	Max
$VisibleAbnSent_t$ (%)	692	1.42	1.72	-5.12	-0.60	1.41	3.33	16.17
$UnskilledAbnSent_t$ (%)	692	0.99	2.41	-10.05	-1.78	1.21	3.08	14.32
$AntiskilledAbnSent_t$ (%)	692	2.47	2.99	-11.70	-1.21	2.67	5.87	13.96

day t. To capture the belief bias induced by antiskilled finfluencers tweeting about stocks about which they are misinformed or faking their tweets, we define the belief bias relative to the sentiment of the skilled finfluencers. We do this for the more visible, unskilled, and antiskilled finfluencers in every stock j and day t

$$\begin{aligned} VisibleAbnSent_{j,t} &= VisibleSent_{j,t} - SkilledSent_{j,t}, \\ Un/AntiskilledAbnSent_{j,t} &= Un/AntiskilledSent_{j,t} - SkilledSent_{j,t}. \end{aligned} \tag{23}$$

To construct daily averages we aggregate the abnormal social sentiment

$$\begin{aligned} VisibleAbnSent_t &= \tfrac{1}{J}\textstyle\sum_{\text{all } j} VisibleAbnSent_{j,t}, \\ Un/AntiskilledAbnSent_t &= \tfrac{1}{J}\textstyle\sum_{\text{all } j} Un/AntiskilledAbnSent_{j,t}. \end{aligned} \tag{24}$$

Figure 6 plots the average abnormal social sentiment of visible (Panel A) and unskilled and antiskilled finfluencers (Panel B) for the years 2015 and 2016. The figure illustrates a clear pattern. The left subplot of Panel A plots the time series of the daily average abnormal social sentiment, while the right subplot of Panel B shows its distribution. Both subplots show that the abnormal social sentiment of more visible finfluencers is centered above zero with several episodes when more visible finfluencers disseminate strongly positive social sentiment for extended periods and a few episodes when skilled finfluencers disseminate strongly negative social sentiment.

Panel B reveals a similar picture when we split unskilled and antiskilled finfluencers. Unskilled and antiskilled finfluencers behave similarly, with antiskilled disseminating more pronounced positive tweets. The daily average abnormal social sentiment of antiskilled finfluencers is significantly positive almost all the time. This implies antiskilled finfluencers in aggregate tend to tweet more positively than negatively, biasing their followers' beliefs upward. Antiskilled finfluencers' sentiment exhibits persistent swings and few spikes, in contrast to skilled finfluencers. Finfluencers that follow antiskilled finfluencers thus exhibit overly optimistic beliefs most of the time, overly pessimistic beliefs some of the time, and persistent swings in their belief bias.

Table 5 reports summary statistics for the abnormal social sentiment revealed by the tweets of unskilled and antiskilled finfluencers. The statistics in Table 5 are consistent with Figure 6. The abnormal social sentiment revealed by more visible finfluencers is 1.42% higher than the control group of skilled finfluencers. Relative to skilled finfluencers, Table 5 shows that the social sentiment revealed by unskilled (antiskilled) finfluencers is 0.99% (2.47%) higher on average. Un- and antiskilled finfluencers are thus overly optimistic even beyond the market uptrend.

4.2 Joint relation between returns, activity, and sentiment

Next, we test how finfluencers' tweeting activity and tweets' sentiment relate to past and future returns. We employ panel vector-autoregressions (PVAR) to explore the lead-lag relationship between stock returns and tweets. Stock returns and tweets likely affect one another contemporaneously and over time. PVARs allow us to account for the endogeneity among the variables. By pooling the data across finfluencers and stocks, PVARs exploit more information than traditional VARs, providing more efficient estimates of the dynamic interactions while controlling for unobserved heterogeneity.

Similar to the aggregate finfluencer-induced beliefs, we compute stock-day level activity for each stock j and day t as

$$\begin{aligned} VisibleActivity_{j,t} &= \frac{\sum_{\text{all } i} Visibility_i \times Activity_{i,j,t}}{\sum_{\text{all } i} Visibility_i \times Active_{i,j,t}}, \\ Un/Anti/SkilledActivity_{j,t} &= \frac{\sum_{\text{all } i} \Pr(\text{user } i \text{ un/anti/skilled}) \times Activity_{i,j,t}}{\sum_{\text{all } i} \Pr(\text{user } i \text{ un/anti/skilled}) \times Active_{i,j,t}}, \end{aligned} \tag{25}$$

where $Activity_{i,j,t}$ is the sum of $SocSent^{+}_{i,j,t}$, $SocSent^{-}_{i,j,t}$, and $SocSent^{0}_{i,j,t}$.

Our main variables of interest are stock returns $Ret_{j,t}$, stock-specific finfluencer activity, $SkilledActivity_{j,t}$ and $VisibleActivity_{j,t}$, respectively, and the sentiment of the stock-specific tweets of different types of finfluencers, $SkilledSent_{j,t}$ and $VisibleSent_{j,t}$. We collect in vector $Y_{j,t}$ the endogenous variables for return in every stock j and everyday t, social media activity, and the tweets by skilled (visible) finfluencers

$$Y_{j,t} = \begin{pmatrix} Ret_{j,t} \\ SkilledActivity_{j,t} \\ VisibleActivity_{j,t} \\ SkilledSent_{j,t} \\ VisibleSent_{j,t} \end{pmatrix}. \tag{26}$$

For the variables in (26), we identify skilled and visible finfluencers, respectively, as in the prior section. The panel VAR specification for $Y_{j,t}$ is

$$Y_{j,t} = \alpha_j + \sum_{l=1}^{L} A_l Y_{j,t-l} + \epsilon_{j,t}, \tag{27}$$

with 4-dimensional error term $\epsilon_{j,t} \sim iid(0, \Sigma)$ and lag length L. We estimate (26) using a system GMM estimation (Arellano and Bover, 1995) with the lags as instruments. We control for stock-level fixed effects by forward-mean-differencing, also known as Helmert transformation. The Helmert transformation preserves the orthogonality between the variables and their lags which is essential for the system GMM.

The PVAR's results are best summarized by the impulse response functions (IRF) of the four endogenous variables ($Return$, $SkilledActivity$, $VisibleActivity$, $SkilledSent$, $VisibleSent$) to unit shocks displayed in Figure 7. Based on the GMM estimates with $L = 1$ and the Wold decomposition based on the order of the variables in (26), the IRFs show how $Y_{j,t+h}$, $h = 1, ..., 6$, reacts to a unit innovation in the disturbance term $\epsilon_{j,t}$ holding all other shocks constant. The confidence bands of the IRF are generated in Monte Carlo simulations with 1,000 draws.

Figure 7 in the first (second) row shows the impact on returns (activity) over the next 6 days of shocks to returns, activity, and social sentiment. To set the stage, the reference finding is by Cookson, Lu, Mullins, and Niessner (2022) who show that pooled attention predicts negative next-day returns. This finding holds in our data but when we split activity and sentiment between skilled and visible, respectively, results differ. Higher finfluencer activity follows positive stock returns (column 1, rows 2 and 3). However, skilled activity precedes positive stock returns (row 1, column 2) while only visible activity precedes negative stock returns (row 1, column 3), though the latter is insignificant once we control for skilled activity. Therefore, higher skilled activity indicates prices going up for at least 2 consecutive

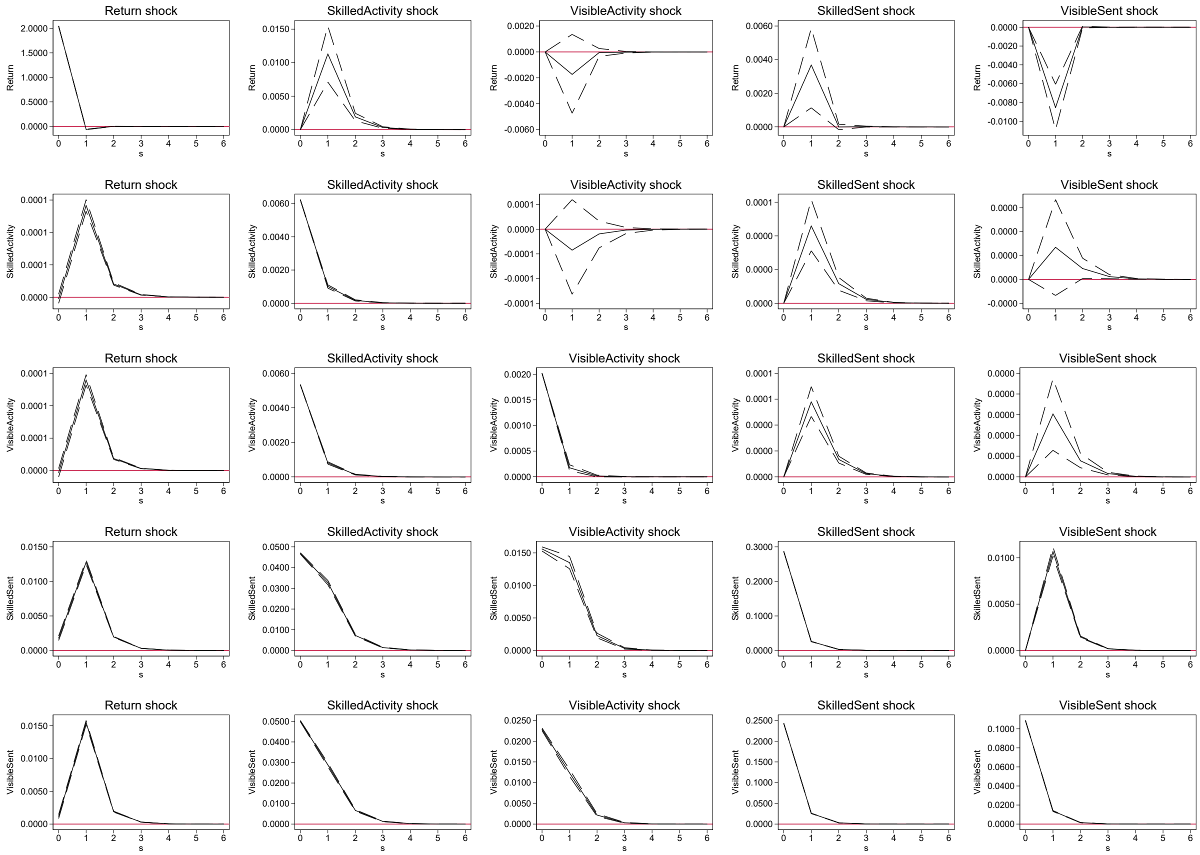

Figure 7: Impulse response functions

The plot shows the impulse response functions of the four endogenous variables ($Return$, $SkilledActivity$, $VisibleActivity$, $SkilledSent$, $VisibleSent$) to unit shocks. The specification is from (27) with $L = 1$.

days, while higher visible activity indicates that price momentum has peaked.

The fourth (fifth) row in Figure 7 shows the impact on tweets by skilled (visible) finfluencers over the next 6 days of shocks to returns, activity, and social sentiment. Once again, the reference finding is by Cookson, Lu, Mullins, and Niessner (2022) who show that pooled sentiment predicts positive next-day returns. When we split into skilled and visible, we find as expected more positive tweets by skilled finfluencers precede positive stock returns (row 1, column 4). However, unlike the result by Cookson, Lu, Mullins, and Niessner (2022), more positive tweets by more visible, un/antiskilled finfluencers precede negative stock returns (row 1, column 5). The first column of Figure 7 also shows that both skilled and visible finfluencers' tweets become more optimistic in response to positive past returns. In summary, the impulse responses from PVAR (26) confirm that skilled finfluencers correctly

Table 6: Panel VAR

This table reports the results from GMM estimation of the panel VAR specification (27). Standard errors are robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. The number of observations is 2,115,166.

Panel A: Specification (27) with skilled and visible finfluencers

	$Ret_{j,t}$	$SkilledActivity_{j,t}$	$VisibleActivity_{j,t}$	$SkilledSent_{j,t}$	$VisibleSent_{j,t}$
$Ret_{j,t-1}$	-0.030***	0.000***	0.000***	0.006***	0.008***
	(0.002)	(0.000)	(0.000)	(0.000)	(0.000)
$SkilledActivity_{j,t-1}$	2.348***	0.181***	0.055***	-0.316	-0.390*
	(0.776)	(0.033)	(0.019)	(0.254)	(0.213)
$VisibleActivity_{j,t-1}$	-0.583	-0.023	0.085***	5.523***	4.821***
	(0.910)	(0.029)	(0.017)	(0.281)	(0.239)
$SkilledSent_{j,t-1}$	0.079***	0.000**	0.000**	0.007***	-0.017***
	(0.013)	(0.000)	(0.000)	(0.003)	(0.002)
$VisibleSent_{j,t-1}$	-0.079***	0.000	0.000***	0.098***	0.127***
	(0.014)	(0.000)	(0.000)	(0.003)	(0.003)

Panel B: Specification (27) with all finfluencers pooled

	$Ret_{j,t}$	$Activity_{j,t}$	$Sent_{j,t}$
$Ret_{j,t-1}$	-0.030***	0.002***	0.474***
	(0.002)	(0.000)	(0.115)
$Activity_{j,t-1}$	-0.029***	0.551***	-77.049***
	(0.014)	(0.060)	(18.219)
$Sent_{j,t-1}$	0.000	0.000	0.461***
	(0.000)	(0.000)	(0.057)

predict future returns while visible or un/antiskilled finfluencers do so incorrectly.

Table 6 reports the estimated regression coefficients A_1 from PVAR specification (27). All coefficients in predicting returns (first column) are statistically significant at the 1% level except for visible activity which is not statistically significant. Returns decline with past returns and the abnormal social sentiment of more visible finfluencers, but increase with skilled finfluencers' activity and sentiment. Finfluencers' activity increases with past returns, finfluencers' activity, and the social sentiment of skilled finfluencers. Sentiment increases with past returns, past visible activity, and sentiment.

Overall, our results show that skilled finfluencers are more neutral most time and only occasionally disseminate strongly positive or negative social sentiment. By contrast, un- and antiskilled finfluencers tend to be overoptimistic and have persistent belief swings. This tweeting behavior by un/antiskilled finfluencers distorts the wisdom of the crowd, that is, the ability to aggregate diffuse information dispersed across a large number of finfluencers. These findings are consistent with our model highlighting the potential for misinformation and skewed perceptions in environments where engagement and visibility drive follower growth, rather than the quality of content. Therefore these results underscore the potential for intervention strategies to promote high-quality content and mitigate the dominance of low-quality finfluencers. The next section uses our model to motivate several policy interventions.

4.3 Policy interventions

Given the model insights into the finfluencer activity and followers' dynamics of belief formation and growth, several policies can be proposed to mitigate the dominance of bad finfluencers and enhance the overall quality of content on the platform. Each policy is analyzed for its expected impact and potential unintended consequences.

Quality-based visibility boosts. Platforms can increase the visibility (γ) of good finfluencers based on content quality metrics. By increasing γ for good finfluencers, their visibility $V_G(t)$ can become more comparable with $V_B(t)$. This shift helps to balance the follower growth rates between good and bad finfluencers. However, it is worth noting that bad finfluencers might respond by further increasing their engagement levels to maintain dominance, potentially leading to an arms race in visibility and engagement. This reaction could increase the overall noise on the platform and lower content quality.

Penalizing low-quality content. Platforms can impose penalties on bad finfluencers for low-quality content, such as reduced visibility (for instance, through inferior placement on the platform or unfavorable sort order) or temporary suspension. In other words, introducing penalties effectively reduces $V_B(t)$, making it more challenging for bad finfluencers to compete with good finfluencers. This shift can help redirect follower growth towards higher-quality finfluencers. However, it is worth noting that if the penalties are perceived as overly harsh or unfair, they might provoke a backlash from users who follow bad finfluencers, potentially driving these users to other platforms.

Transparency and verification. Platforms can also implement a curation process or verification system where finfluencers can, for instance, earn badges or certifications based on content quality and adherence to platform guidelines. A verified status can enhance U_G, making good finfluencers more competitive against bad finfluencers even if $V_B(t)$ remains

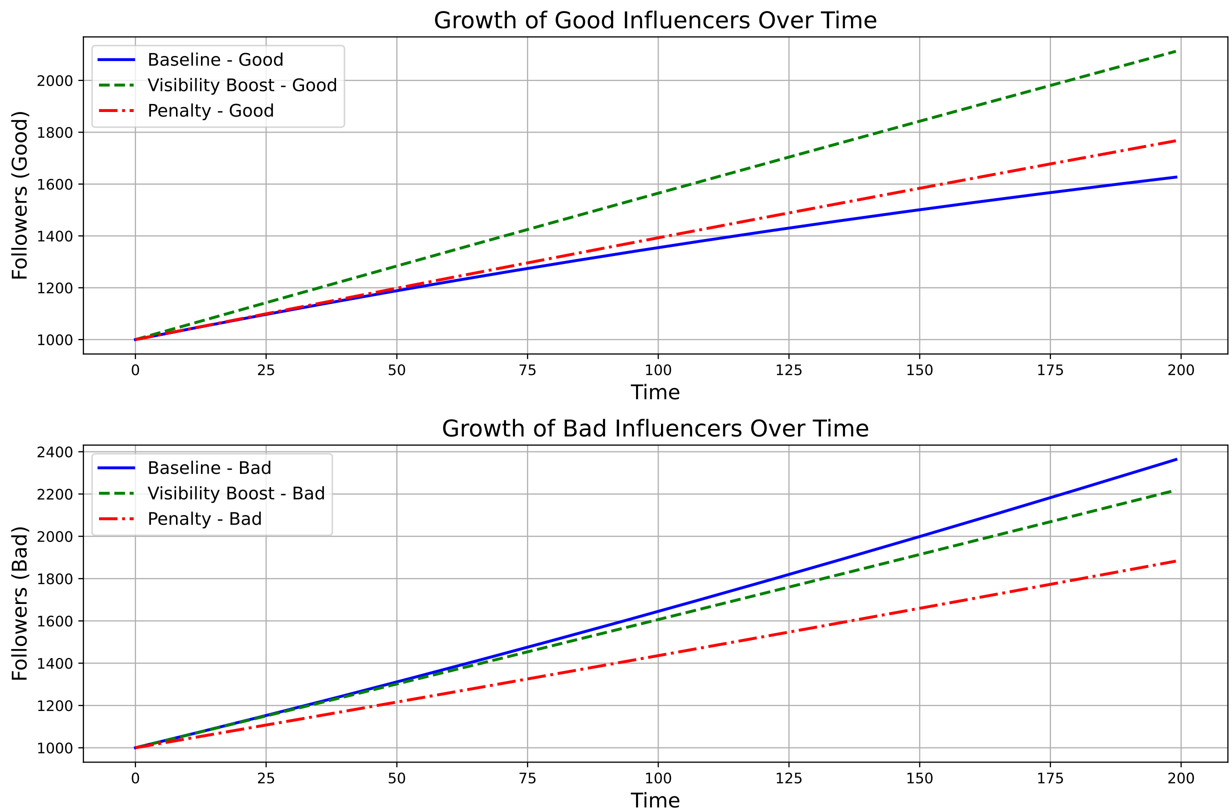

Figure 8: Policy interventions

The top plot shows the growth of good finfluencers over time. The baseline (solid blue) shows natural growth. *'Visibility Boost'* (dashed green) increases γ_G, enhancing growth. *'Penalizing Low-Quality Content'* (dash-dotted red) decreases γ_B, indirectly benefiting good finfluencers. The bottom plot shows the growth of bad finfluencers over time. The baseline (solid blue) shows natural growth. *'Visibility Boost'* (dashed green) increases γ_G, with minimal effect on bad finfluencers. *'Penalizing Low-Quality Content'* (dash-dotted red) decreases γ_B, significantly reducing growth. Parameters: $U_G = 1.5, U_B = 1, k = 1, F_{new} = 10, \gamma_G = 0.01, \gamma_B = 0.01, \pi_G = 2, \pi_B = 3, c_G = 1, c_B = 1$, $\gamma_G^{boost} = 0.02$, and $\gamma_B^{penalty} = 0.005$.

high. However, it is worth noting that verification processes need to be robust and fair to avoid any perception of bias or favoritism, which could undermine the system's credibility.

Figure 8 explores several policy interventions. The baseline is the solid blue line showing natural growth. The top plot shows the growth of good finfluencers over time. The bottom plot shows the growth of bad finfluencers over time. The policy experiment *'Visibility Boost'* increases γ_G, enhancing the growth of the good type's followers, which is depicted by the dashed green line. By contrast, *'Visibility Boost'* has minimal effect on bad finfluencers. *'Penalizing Low-Quality Content'* decreases γ_B, indirectly benefiting good finfluencers, which is depicted by the dash-dotted red line. The reason is that *'Penalizing Low-Quality Content'* significantly reduces the growth of the bad type's followers, depicted in the bottom plot.

In sum, implementing these policies can help counteract the dominance of bad finfluencers and improve overall content quality on the platform. By understanding and addressing these dynamics, platforms can foster a healthier content environment and ensure that high-quality finfluencers are appropriately recognized and followed.

5 Conclusion

Social media has gained great importance in recent years for sharing and acquiring information. An important question is whether competition among users of social media platforms is such that followers can easily identify skilled financial influencers, so-called finfluencers, and drive out unskilled finfluencers from the market for social information. We find that the answer is no.

Social media users could use the tweeting behavior of finfluencers to identify their skills. However, instead of following more skilled finfluencers, social media users follow unskilled and antiskilled finfluencers, which we define as finfluencers whose tweets generate negative alpha. Un/antiskilled finfluencers are more active and ride return and social sentiment momentum, which coincide with the behavioral biases of many retail investors.

These results are consistent with slow learning by social media users and active follower engagement shaping finfluencer's follower networks and limiting competition among finfluencers, resulting in the failure of the "wisdom of the crowd" due to the long-term survival of un/antiskilled finfluencers even though they do not provide valuable investment advice.

Our findings shed light on the quality of finfluencers' unsolicited financial advice and the competition among and economic incentives faced by finfluencers which regulators have been concerned about. We consider several policy interventions to correct the aggregate belief biases.

References

Arellano, Manuel and Olympia Bover (1995). "Another look at the instrumental variable estimation of error-components models". In: *Journal of Econometrics* 68.1, pp. 29–51.

Azar, Pablo D. and Andrew W. Lo (2016). "The wisdom of Twitter crowds: Predicting stock market reactions to FOMC meetings via Twitter feeds". In: *The Journal of Portfolio Management* 42.5, pp. 123–134.

Ballinari, Daniele and Simon Behrendt (2021). "How to gauge investor behavior? A comparison of online investor sentiment measures". In: *Digital Finance* 3.2, pp. 169–204.

Banerjee, Abhijit V. (1992). "A simple model of herd behavior". In: *The Quarterly Journal of Economics* 107.3, pp. 797–817.

Barber, Brad M. and Terrance Odean (2007). "All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors". In: *The Review of Financial Studies* 21.2, pp. 785–818.

Bartov, Eli, Lucile Faurel, and Partha S. Mohanram (2018). "Can Twitter help predict firm-level earnings and stock returns?" In: *The Accounting Review* 93.3, pp. 25–57.

Benabou, Roland and Guy Laroque (1992). "Using Privileged Information To Manipulate Markets: Insiders, Gurus, and Credibility". In: *The Quarterly Journal of Economics* 107, pp. 921–58.

Berk, Jonathan B. and Jules H. Van Binsbergen (2022). "Regulation of Charlatans in High-Skill Professions". In: *The Journal of Finance* 77.2, pp. 1219–1258.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992). "A theory of fads, fashion, custom, and cultural change as informational cascades". In: *Journal of Political Economy* 100.5, pp. 992–1026.

Boehmer, Ekkehart, Charles M. Jones, and Xiaoyan Zhang (2008). "Which Shorts Are Informed?" In: *The Journal of Finance* 63.2, pp. 491–527.

Campbell, John L., Jenna D'Adduzio, and James Moon (2021). "Everyone Has an Opinion: The Informativeness of Social Media's Response to Management Guidance". In: *Working Paper.*

Chen, Hailiang, Prabuddha De, Yu Hu, and Byoung-Hyoun Hwang (2014). "Wisdom of crowds: The value of stock opinions transmitted through social media". In: *The Review of Financial Studies* 27.5, pp. 1367–1403.

Chen, Yong, Michael Cliff, and Haibei Zhao (2017). "Hedge Funds: The Good, the Bad, and the Lucky". In: *Journal of Financial and Quantitative Analysis* 52.3, pp. 1081–1109.

Cookson, J. Anthony, Joseph E. Engelberg, and William Mullins (2023). "Echo Chambers". In: *The Review of Financial Studies* 36.2, pp. 450–500.

Cookson, J. Anthony, Rinjing Lu, William Mullins, and Marina Niessner (2022). "The Social Signal". In: *Working Paper*.

Cookson, J. Anthony and Marina Niessner (2020). "Why Don't We Agree? Evidence from a Social Network of Investors". In: *The Journal of Finance* 75.1, pp. 173–228.

Crane, Alan and Kevin Crotty (2020). "How skilled are security analysts?" In: *The Journal of Finance* 75.3, pp. 1629–1675.

Crawford, Steven, Wesley Gray, Bryan R. Johnson, and Richard A. Price III (2018). "What motivates buy-side analysts to share recommendations online?" In: *Management science* 64.6, pp. 2574–2589.

Crawford, Steven S., Wesley R. Gray, and Andrew E. Kern (2017). "Why do fund managers identify and share profitable ideas?" In: *Journal of Financial and Quantitative Analysis* 52.5, pp. 1903–1926.

Currarini, Sergio, Matthew O. Jackson, and Paolo Pin (2009). "An Economic Model of Friendship: Homophily, Minorities, and Segregation". In: *Econometrica* 77.4, pp. 1003–1045.

Curtis, Asher, Vernon J. Richardson, and Roy Schmardebeck (2014). "Investor attention and the pricing of earnings news". In: *Handbook of Sentiment Analysis in Finance, Forthcoming*.

Dim, Chukwuma (2022). "Should retail investors listen to social media analysts? evidence from text-implied beliefs". In: *Working Paper*.

Engelberg, Joseph E., Adam V. Reed, and Matthew C. Ringgenberg (2012). "How are shorts informed?: Short sellers, news, and information processing". In: *Journal of Financial Economics* 105.2, pp. 260–278.

Ferrara, Emilio and Zeyao Yang (Nov. 2015). "Measuring Emotional Contagion in Social Media". In: *PLOS ONE* 10.11, pp. 1–14.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny (2015). "Money Doctors". In: *The Journal of Finance* 70.1, pp. 91–114.

Giannini, Robert, Paul Irvine, and Tao Shu (2018). "Nonlocal disadvantage: An examination of social media sentiment". In: *The Review of Asset Pricing Studies* 8.2, pp. 293–336.

Golub, Benjamin and Matthew O. Jackson (2012). "How Homophily Affects the Speed of Learning and Best-Response Dynamics". In: *The Quarterly Journal of Economics* 127.3, pp. 1287–1338.

Greenwood, Robin and Andrei Shleifer (2014). “Expectations of Returns and Expected Returns”. In: *The Review of Financial Studies* 27.3, pp. 714–746.

Harvey, Campbell R. and Yan Liu (2018). “Detecting Repeatable Performance”. In: *The Review of Financial Studies* 31.7, pp. 2499–2552.

Hirshleifer, David (2020). “Presidential Address: Social Transmission Bias in Economics and Finance”. In: *The Journal of Finance* 75.4, pp. 1779–1831.

Jame, Russell, Rick Johnston, Stanimir Markov, and Michael C. Wolfe (2016). “The value of crowdsourced earnings forecasts”. In: *Journal of Accounting Research* 54.4, pp. 1077–1110.

Kandel, Denise B. (1978). “Homophily, Selection, and Socialization in Adolescent Friendships”. In: *American Journal of Sociology* 84.2, pp. 427–436.

Kogan, Leonid, Stephen A. Ross, Jiang Wang, and Mark M. Westerfield (2006). “The Price Impact and Survival of Irrational Traders”. In: *The Journal of Finance* 61.1, pp. 195–229.

Lazarsfeld, Paul F., Robert K. Merton, et al. (1954). “Friendship as a social process: A substantive and methodological analysis”. In: *Freedom and control in modern society* 18.1, pp. 18–66.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook (2001). “Birds of a Feather: Homophily in Social Networks”. In: *Annual Review of Sociology* 27, pp. 415–444.

Miller, Edward M. (1977). “Risk, uncertainty, and divergence of opinion”. In: *The Journal of Finance* 32.4, pp. 1151–1168.

Pedersen, Lasse Heje (2022). “Game on: Social networks and markets”. In: *Journal of Financial Economics* 146.3, pp. 1097–1119.

Appendix

A Goodness of fit

To assess the goodness of fit, we perform the following procedure using the fitted distributions.

1. Draw M observations $a = [a_1, a_2, ..., a_M]$ from the fitted distribution of true alphas.

2. Generate a sample of M standard errors by bootstrapping $[\tilde{\sigma}_1, \tilde{\sigma}_2, ..., \tilde{\sigma}_M]$ with replacement. Denote this vector by $[s_1, s_2, ..., s_M]$.

3. Generate a vector of estimation errors $e = [e_1, e_2, ..., e_M]$ by drawing each e_i from a Normal distribution with a mean of zero and standard deviation of s_i.

4. Generate $[\tilde{a}_1, \tilde{a}_2, ..., \tilde{a}_M]$ by adding a and e as in (3).

5. Calculate the vector of t-statistics $[t_1, t_2, ..., t_M]$ through $t_i = \tilde{a}_i / s_i$.

6. Repeat steps one to five $1,000$ times.

After applying this procedure, we have $1,000$ samples of simulated $\tilde{\alpha}$'s, each with the same size M as the original data, and their standard errors, t-statistics, and the corresponding true alphas.

Figure A.1 reports the results of several approaches to gauge the goodness of fit. First, we calculate the average pdf and cdf of the simulated samples and plot them against the pdf and cdf of the data. Panel A of Figure A.1 shows the results. The distribution of simulated alphas is close to the distribution of alphas estimated from the data. To quantify the closeness of the distributions, we run Kolmogorov-Smirnov tests between the measured alphas from the data and the simulated alphas from each of the simulated samples, using the null hypothesis that the two distributions are equal. The KS test rejects the null at 10%/5%/1% significance levels for 14.6%/7.40%/0.70% of simulations. Second, we calculate the average pdf of the simulated t-statistics and plot them against the pdf of t-statistics in the data. Panel B of Figure A.1 shows that t-statistics from simulated data are distributed similarly to t-statistics from the data. Another way to visualize the closeness of the two distributions is the Q-Q plot. We calculate the percentiles $(1\%, 2\%, ..., 99\%)$ of each simulated sample of alphas. We plot the mean of the n-th percentiles from the simulated samples against the n-th percentile

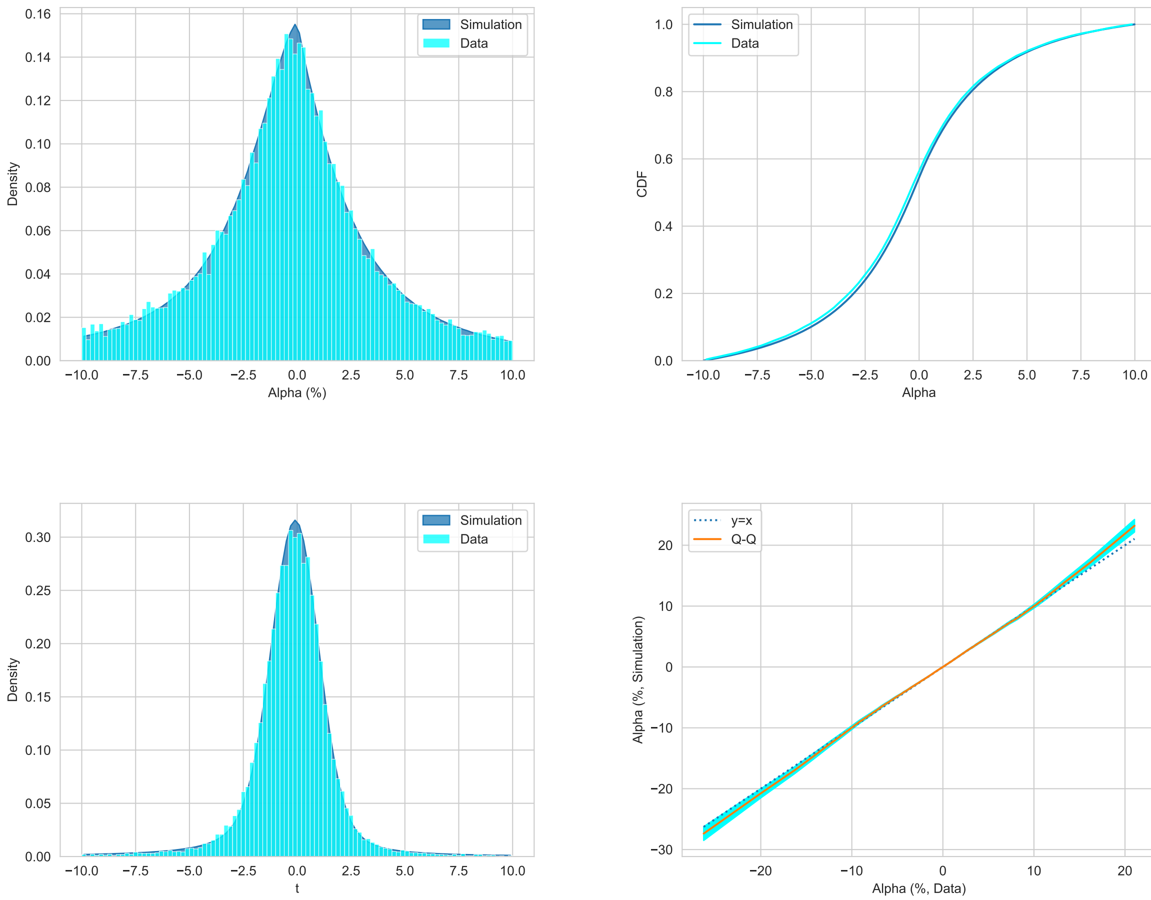

Figure A.1: Measured and Simulated Alphas and Their t-Stats

In Panel A, the left plot shows histograms of measured and simulated alphas. In Panel A, the right plot shows the average cdf of simulated alphas from the fitted model against measured alphas from the data. In Panel B, the left plot shows histograms of the estimated and simulated t-stats. In Panel B, the right plots show a Q-Q plot of the measured and simulated alphas.

from the data to get a Q-Q plot. We also calculate the 95% confidence intervals for each percentile and plot them around the Q-Q plot line on the right subplot of Panel B in Figure A.1. We conclude that the fit with $K^+ = K^- = 2$ is tight.

B Finfluencer survival

Given the finding that our measures of finfluencers' skill are persistent, we check if skilled finfluencers are more likely to stay active, that is, "survive" even though they have fewer

Table B.1: Finfluencer Survival

This table reports the determinants of finfluencers' survival. The results are obtained from Probit regressions. For each regression, the (in)dependent variable is calculated with tweets posted in or after 2016 (before 2016). The dependent variable equals one if the finfluencer is active in or after 2016, and zero otherwise. The independent variables are $\Pr(\alpha_i > 0 \mid \tilde{\alpha}_{i,\text{pre-2016}})$ which is the probability that a user is skilled, $\Pr(\alpha_i = 0 \mid \tilde{\alpha}_{i,\text{pre-2016}})$ which is the probability that a user is unskilled, and $\Pr(\alpha_i < 0 \mid \tilde{\alpha}_{i,\text{pre-2016}})$ which is the probability that a user is antiskilled. Standard errors are robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Finfluencer survival$_i$			
	(1)	(2)	(3)	(4)
Pr(user i skilled | $\tilde{\alpha}_{i,\text{pre-2016}}$)	-0.08			
	(0.04)			
Pr(user i unskilled | $\tilde{\alpha}_{i,\text{pre-2016}}$)		1.31***		1.39***
		(0.12)		(0.13)
Pr(user i antiskilled | $\tilde{\alpha}_{i,\text{pre-2016}}$)			-0.07	0.08
			(0.04)	(0.04)
Constant	-0.17***	-0.41***	-0.15***	-0.46***
	(0.01)	(0.02)	(0.02)	(0.04)
r2	0.000	0.005	0.000	0.005
N	18,770	18,770	18,770	18,770

followers than unskilled and antiskilled finfluencers. We address this question using Probit regressions. For each regression, the (in)dependent variable is calculated with tweets posted in or after 2016 (before 2016). The dependent variable Finfluencer survival$_i$ is an indicator function equal to one if the finfluencer is active in or after 2016, and zero otherwise

$$\text{Finfluencer survival}_i = \Phi(\alpha + \beta \times \text{Skill}_{i,\text{pre-2016}}), \tag{B1}$$

where Φ is the Normal cdf and Skill$_i$ is one of the following five variables: $\tilde{\alpha}_{i,\text{pre-2016}}$ is the finfluencer's measured alpha in the data before 2016, $\mathbb{E}[\alpha_i \mid \tilde{\alpha}_{i,\text{pre-2016}}]$ is the expected value of alpha given its measurement in the data before 2016, $\Pr(\alpha_i > 0 \mid \tilde{\alpha}_{i,\text{pre-2016}})$ is the probability that a user is skilled, $\Pr(\alpha_i = 0 \mid \tilde{\alpha}_{i,\text{pre-2016}})$ is the probability that a user is unskilled, and $\Pr(\alpha_i < 0 \mid \tilde{\alpha}_{i,\text{pre-2016}})$ is the probability that a user is antiskilled. The results in Table B.1 show that skill does not improve survival. First, the finfluencer's measured alpha, $\Pr(\alpha_i > 0 \mid \tilde{\alpha}_{i,\text{pre-2016}})$, is an insignificant determinant of survival. Only the probability of being unskilled statistically significantly predicts survival and the relation is positive.

C Dissecting finfluencers' tweeting strategies

Table C.1 reports results from multivariate regressions explaining StockTwits users' skills by observable characteristics of their tweeting activity, captured by the same variable used in Figure 4, and the potential tweeting strategies pursued by different types of finfluencers. The table reports the results of regressions of the form

$$\text{Skill}_i = \alpha + \beta \cdot \textit{TweetingActivity}_i + \delta \cdot \textit{FractionPositive}_i + \gamma^\top \textit{TweetingStrategy}_i + \epsilon_i, \quad \text{(C2)}$$

where Skill_i represents one of the following variables: (1) the probability of α being positive, (2) the probability of α being zero, and (3) the probability of α being negative. Across the different specifications, we consider several potential motives for tweeting and tweeting strategies.

Finfluencers may cater to users who enjoy positive messages. This can help attract and retain followers who seek encouragement and optimism in financial content, making followers more likely to trust and follow the finfluencers' advice.[1] In addition, the literature has documented that short sellers are informed (e.g., Boehmer, Jones, and Zhang, 2008; Engelberg, Reed, and Ringgenberg, 2012). Therefore, we might expect finfluencers with more negative tweets to be more skilled. To test these hypotheses, we relate our measures of skill to the number of tweets and the composition of the tweets, particularly the fraction of tweets with a negative tone.

Another way for finfluencers to create user engagement is to follow certain tweeting strategies popular with investors/users. To explore this channel, we dissect finfluencers' tweeting strategies for each skill type. Doing this also helps to understand the nature of information or skills held by finfluencers and what determines the potential belief biases and abnormal performance of different finfluencers. We use the finfluencers' skill measures from Section 1.3 to study whether finfluencers follow commonly known investment behaviors.

The prior literature has documented several stylized patterns among investors finfluencers may exploit or cater to create engagement.[2]

[1] Research in psychology, marketing, and consumer behavior finds that people are generally drawn to positive messages and positive messaging can significantly impact audience engagement and behavior. For instance, Ferrara and Yang (2015) analyze Twitter data and find that positive sentiment is associated with higher levels of engagement, including likes, retweets, and replies.

[2] The measures of finfluencers' skills computed in the previous sections are not directly observable in the data by StockTwits users. Directly observable by StockTwits users, and thus potentially more relevant for distinguishing skilled from unskilled finfluencers, are user-level characteristics such as the number of tweets, their tone, and the number of followers and likes. If finfluencers can be categorized by these characteris-

- Return chasing vs. contrarian behavior: Retail traders tend to chase returns (Barber and Odean, 2007). In our setup, we can ask if the tweets by all or some groups of finfluencers are consistent with return chasing. In particular, if antiskilled finfluencers' tweets chase returns, return chasing may contribute to these finfluencers' measured negative skill. In turn, contrarian behavior tends to be associated with skilled investors because going against prevailing market trends requires sound market understanding.

- Positive vs. negative herding: Herding behavior can affect the informativeness and engagement of finfluencers' tweets.

- Short-sale constraints: Asset pricing theory suggests that risky assets are overpriced in a market with short-sale constraints (Miller, 1977). As a result, short-sale-constrained stocks tend to be overpriced. We ask whether finfluencers exploit this mispricing in their tweets. Due to this mispricing, we expect skilled finfluencers to post more negative tweets about stocks with tighter short-selling constraints.

Table C.1 starts with the results for tweeting activity, *TweetingActivity*, as well as for a fraction of positive tweets, *FractionPositive*, as explanatory variables of finfluencers' skills. The composition of tweets, *FractionPositive*, is defined as the percentage of a finfluencer's non-neutral tweets that have a positive sentiment. In agreement with the univariate results from Figure 4, the probability of being skilled decreases by 3.34% while the probability of being unskilled (antiskilled) increases by 2.45% (0.88%) when the number of tweets increases tenfold. Put together, finfluencers who tweet more frequently are less likely to be skilled, consistent with informed finfluencers tweeting less frequently.

The table also includes the estimates for the percentage of a finfluencer's non-neutral tweets that have a positive sentiment, *FractionPositive*, used as the explanatory variable. Consistent with the prior literature, finfluencers with more positive tweets are less likely to be skilled. A one-percent increase in the share of positive tweets is associated with a 6bps decrease in the probability of being skilled, while the probability of being antiskilled increases by 8bps. All estimates are significant at 1% and point to the same conclusion: StockTwits users with more positive messages/tweets are more likely to post anti-informative tweets.

We measure each user's return-chasing tendency by the percentage of tweets that are either positive and about the highest decile of prior week returns or negative and about the lowest decile of prior week returns. To test the return chasing hypothesis, we perform two

tics or they use these observable characteristics to signal their type to other StockTwits users, then these characteristics are informative about finfluencers' skills.

checks. We first regress measured and expected alphas on return chasing to test if return chasing is associated with better or worse performance. Table C.1 also reports the results of the return chasing tests. We find that the probability of being skilled or antiskilled changes with the tendency to chase returns. A one percent increase in return chasing tendency is associated with a 3bps decrease in the probability of being skilled and a 3bps increase in the probability of being antiskilled. Because the skilled, unskilled, and antiskilled components sum up to one, the probability of being unskilled remains unchanged. Overall, return chasing contributes to finfluencers being antiskilled.

We measure each user's contrarian tendency as the percentage of tweets that are either positive and about the lowest decile of prior week returns or negative and about the highest decile of prior week returns. Table C.1 reports results from regressing our measures of skill on contrarian tendency, together with the other determinants. It could be that skilled finfluencers follow a contrarian approach given that return chasing contributes to negative skill. The results show a weak positive association between contrarian tweeting and skill, significant at 10%. Finfluencers who post contrarian tweets indeed exhibit higher skills.

To quantify herding, we calculate the percentage of each finfluencer's positive/negative tweets that are about stocks in the highest decile of positive/negative tweeting activity over the past five days. We include in our regressions of skill both the finfluencers' positive, *PositiveHerding*, and negative, *NegativeHerding*, herding tendencies. Table C.1 reports the results of regressing the skill measures on *PositiveHerding*. It shows that a one-percent increase in positive herding tendency is associated with a 2bps increase in the probability of being unskilled, while the probability of being (anti)skilled is not significantly affected. Taken together, the results in Table C.1 show that positive herding tendency is weakly negatively related to finfluencers' skills. Anecdotal evidence shows that herding behavior on social media is associated with positive sentiment. The meme stock episode in 2021 is one such example. However, one can also measure herding around negative tweets. Thus, we include in our regressions an alternative definition of the independent variable that measures herding on negative tweets, with similar results.[3]

Last, we use the Markit short-selling index to measure the short-selling constraints of individual stocks. The Markit index is a number between 1 and 20 with 1 representing no short-selling constraints and 20 representing maximum short-selling constraint. Every day,

[3] Table C.1 also reports the results of regressing the skill measures on *NegativeHerding*. It shows that finfluencers who tweet more often about stocks in the top decile of negative tweeting activity are less likely to be unskilled and more likely to be skilled. A one-percent increase in the negative herding measure is associated with a 2bps increase in the probability of being unskilled.

we sort stocks into deciles based on the average of their Markit index over the past five days. For each user, we calculate two variables representing the average decile of the Markit index for all stocks that she tweeted positively and negatively. These two variables are our measures of short-selling constraints for positive and negative tweets. We include them in our regressions to test whether skilled social media users can exploit the overpricing of stocks with short-selling constraints. The last rows of Table C.1 report the coefficients. A one-decile increase in the short-selling constraints of positively tweeted stocks is associated with a 0.38% (0.36%) increase (decrease) in the user's probability of being antiskilled (unskilled). On the other hand, a one-decile increase in the short-selling constraints of negatively tweeted stocks is associated with a 0.25% (0.20%) increase (decrease) in the user's probability of being skilled (unskilled). Overall, these results show that exploiting short-selling constraints correctly contributes to finfluencers' skills on the negative side.

Table C.1: Dissecting Finfluencers' Tweeting Strategies

The table reports the results of regressing skill on tweeting activity or certain tweeting strategies. Skill represents one of the following variables: (1) the probability of α being positive (2) the probability of α being zero (3) the probability of α being negative. The dependent variables are defined in expressions (8). All dependent variables are in percentage points. Tweeting activity is defined as the log of one plus the total number of positive and negative tweets the user has posted. The composition of tweets is represented by *FractionPositive* defined as the percentage of a finfluencer's non-neutral tweets with a positive sentiment. The rest of the explanatory variables proxy for tweeting strategies. *ReturnChasing* is defined as the percentage of user's tweets that are either (1) positive and about stocks in the highest decile of returns over the past week, or (2) negative and about stocks in the lowest decile of returns over the past week. *ContrarianTweet* is defined as the percentage of user's tweets that are either (1) positive and about stocks in the lowest decile of returns over the past week, or (2) negative and about stocks in the highest decile of returns over the past week. *PositiveHerding* is the percentage of the user's positive tweets that are about stocks in the top decile of positive tweeting activity over the past five days. *NegativeHerding* is defined in a similar way for negative tweets. *SSI (Positive Tweets)* represents the average decile of short-selling constraints for stocks positively tweeted by the user. Short-selling constraints are measured using the Markit short-selling index for the stock over the past five trading days. *SSI (Negative Tweets)* is defined in a similar way for negative tweets. Standard errors are robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Pr(user i skilled)	Pr(user i unskilled)	Pr(user i antiskilled)
$TweetingActivity_i$	-3.34***	2.45***	0.88***
	(0.29)	(0.13)	(0.32)
$FractionPositive_i$	-0.06***	-0.02***	0.08***
	(0.01)	(0.00)	(0.01)
$ReturnChasing_i$	-0.03***	0.00	0.03***
	(0.01)	(0.00)	(0.01)
$ContrarianTweet_i$	0.02*	0.01*	-0.02**
	(0.01)	(0.00)	(0.01)
$PositiveHerding_i$	-0.03	0.02**	0.01
	(0.02)	(0.01)	(0.02)
$NegativeHerding_i$	0.03	-0.02**	-0.02
	(0.02)	(0.01)	(0.02)
SSI_i *(Positive Tweets)*	-0.01	-0.36***	0.38***
	(0.08)	(0.03)	(0.08)
SSI_i *(Negative Tweets)*	0.25***	-0.20***	-0.05
	(0.08)	(0.02)	(0.08)
Constant	36.10***	17.06***	46.84***
	(0.80)	(0.27)	(0.83)
N	19,593	19,593	19,593

Internet Appendix

IA Alternative Specifications for the Distribution of True Alphas

Table IA.1 reports parameter estimates for alternative model specifications. Panel A reports the estimated distribution of true alphas assuming one and three components for types 1 and 3. The likelihood value and the AIC and BIC criteria improve considerably by moving from one component to two. However, adding the third component does not improve the fit by much. We also repeat our tests of goodness-of-fit for these alternative models. In KS tests, the model with $K^+ = K^- = 1$ is rejected at the 10%/5%/1% level for 100%/100%/98.2% of simulations. For the model with $K^+ = K^- = 3$, the KS tests reject the null hypothesis at the 10%/5%/1% level for 6.20%/2.50%/0.30% of simulations. Panel B reports the results of fitting mixture models over different horizons H with two exponentials on the $\alpha > 0$, two exponentials on $\alpha < 0$, and a mass at $\alpha = 0$. We calculate excess returns over the next 1, 2, 5, 10, and 20 trading days using the Fama-French five-factor model.

IB Omitted proofs

We first have the following remarks.

Remark 1 (Boundedness of Tweeting Strategy $\Phi_i(t)$)**.** *In the proofs, we assume that each finfluencer's chosen tweeting volume and sentiment strategy cannot diverge to infinity (nor vanish to zero in a way that destroys visibility). Formally, if*

$$\Phi_i(t) = \textit{some function of } \big[v_i(t), \{SocSent_{i,j,t},\, Tweet_{i,j}(t)\}\big],$$

we require $0 < \underline{\Phi} \leq \Phi_i(t) \leq \overline{\Phi} < \infty$ *for all* t*. One can ensure this, for instance, by imposing concavity in the finfluencer's utility so that tweet volume does not explode, or by explicitly bounding* $v_i(t)$*. This ensures that*

$$V_i(t) = \kappa \frac{\pi_i}{c_i} \big[F_i(t)\big]^2 \times \Phi_i(t)$$

still reflects the main growth rate in $\frac{\pi_i}{c_i} \cdot [F_i(t)]^2$*. A more general unbounded* $\Phi_i(t)$ *could, in principle, overshadow the reward–cost ratio effect, but we rule that out by assumption.*

Remark 2 (Short-Run Overshadowing of Skill by Visibility)**.** *Even though skilled finfluencers deliver better investment performance in the long run, new followers can heavily weight the entertainment or social aspect captured by $\gamma V_i(t)$. Over short horizons, if a bad (un/antiskilled) finfluencer has very high visibility, the immediate utility $\gamma V_i(t)$ may outweigh the difference between U_S and U_B, thus slowing or reversing the process of learning about actual skill. This tension underlies our main finding that high engagement can sustain large followings for low-quality finfluencers.*

Remark 3 (Pairwise vs. Continuum of Finfluencers)**.** *In the statement of Proposition 1, we compare two specific finfluencers i_1 and i_2. The same argument extends naturally when there are multiple competing finfluencers in each category. In that broader case, we can consider any pair of finfluencers with different π_i/c_i ratios and show that the one with a higher ratio and sufficiently large follower base eventually outgrows the other. Proposition 2 then uses this logic in the continuum, showing how the entire set of high–π_i/c_i (antiskilled) agents dominates the entire set of skilled ones if $\frac{\pi_A}{c_A} > \frac{\pi_S}{c_S}$.*

Remark 4 (Unskilled vs. Antiskilled in the Bias Result)**.** *While Proposition 2 specifically highlights $\frac{\pi_A}{c_A} > \frac{\pi_S}{c_S}$ (i.e., antiskilled overshadowing skilled), the same argument applies if unskilled agents also have $\frac{\pi_U}{c_U} > \frac{\pi_S}{c_S}$. In that scenario, unskilled finfluencers can similarly capture the lion's share of attention and followers, thereby skewing aggregate sentiment away from skilled fundamentals. The model thus uniformly explains why "bad" finfluencers (whether truly antiskilled or merely unskilled with high engagement) can dominate on social media platforms.*

Proof of Proposition 1 (Dominance of High-Engagement Finfluencers). We prove that if finfluencer i_1 has a strictly higher ratio $\frac{\pi_{i_1}}{c_{i_1}}$ and a sufficiently large initial follower base at some time t', then $F_{i_1}(t)$ eventually exceeds $F_{i_2}(t)$ for all large t.

Step 1: Setup and definitions.
Consider two specific finfluencers i_1 and i_2, each with:

$$(\pi_{i_1}, c_{i_1}) \quad \text{and} \quad (\pi_{i_2}, c_{i_2}),$$

where π_i is the reward per follower for finfluencer i, and c_i is the cost parameter in her engagement utility. Let

$$F_i(t) = \text{the number of followers of finfluencer } i \text{ at time } t.$$

The premise is that at some time t' we have

$$\frac{\pi_{i_1}}{c_{i_1}} F_{i_1}(t') > \frac{\pi_{i_2}}{c_{i_2}} F_{i_2}(t'). \qquad (*)$$

We wish to show there exists $T > t'$ such that for all $t > T$, $F_{i_1}(t) > F_{i_2}(t)$.

Step 2: Relation between follower growth and visibility.

By construction, the visibility $V_i(t)$ for finfluencer i at time t satisfies

$$V_i(t) = \kappa \cdot e_i(t) \cdot F_i(t),$$

and optimal engagement $e_i(t)$ is proportional to $\left(\frac{\pi_i}{c_i}\right)F_i(t)$. In particular, a standard outcome of the finfluencer's utility maximization is

$$e_i(t) = \text{constant factor} \times F_i(t) = \frac{\pi_i}{c_i}F_i(t) \quad \text{(up to a multiplicative constant } \kappa\text{)},$$

so that, *up to an overall constant* κ, we have

$$V_i(t) = \kappa \cdot \left(\frac{\pi_i}{c_i}F_i(t)\right) \cdot F_i(t) = \kappa \cdot \frac{\pi_i}{c_i}F_i(t)^2.$$

(We ignore minor factors like $\alpha_1 v_i(t)$ or α_2 in the proof since they scale similarly, but one can incorporate them as well; they do not affect the dominance result.)

Step 3: Utility impact on new followers.
The key determinant of follower growth is that each new follower (arriving in a mass $F_{\text{new}}(t)$ at each date t) chooses a finfluencer i with probability

$$P_{k,i}(t) = \frac{U^F_{k,i}(t)}{\sum_m U^F_{k,m}(t)},$$

where

$$U^F_{k,i}(t) = \underbrace{\theta_{k,i}(t)U_S + (1-\theta_{k,i}(t))U_B}_{\text{type-based payoff}} + \underbrace{\gamma V_i(t)}_{\text{visibility payoff}}.$$

Since $U_S > U_B$, skilled finfluencers can have a slight advantage if followers believe strongly that i is skilled, but *the dominant effect* in large t can come from $\gamma V_i(t)$. Hence,

$$\frac{dF_i(t)}{dt} = F_{\text{new}}(t) \times \int_{k \in \Omega_{\text{entrants}}} P_{k,i}(t)dk.$$

In a simplified pairwise comparison between i_1 and i_2, if $V_{i_1}(t) > V_{i_2}(t)$ by a large margin, then $P_{k,i_1}(t)$ can approach 1, implying i_1 quickly accumulates most new followers.

Step 4: Showing finfluencer i_1 eventually outgrows finfluencer i_2.
Define the function

$$G(t) = \frac{\frac{\pi_{i_1}}{c_{i_1}}F_{i_1}(t)^2}{\frac{\pi_{i_2}}{c_{i_2}}F_{i_2}(t)^2} = \frac{\pi_{i_1}/c_{i_1}}{\pi_{i_2}/c_{i_2}} \times \left(\frac{F_{i_1}(t)}{F_{i_2}(t)}\right)^2.$$

Observe that

$$V_{i_1}(t) > V_{i_2}(t) \iff G(t) > 1.$$

From the premise $(*)$, at time t',

$$\frac{\pi_{i_1}}{c_{i_1}} F_{i_1}(t') > \frac{\pi_{i_2}}{c_{i_2}} F_{i_2}(t') \Longrightarrow G(t') = \frac{\frac{\pi_{i_1}}{c_{i_1}} F_{i_1}(t')^2}{\frac{\pi_{i_2}}{c_{i_2}} F_{i_2}(t')^2} > 1.$$

Hence at t', we already have $V_{i_1}(t') > V_{i_2}(t')$. Because new followers place substantial weight on $\gamma V_i(t)$, finfluencer i_1 will capture a larger fraction of entrants as long as $V_{i_1}(t)$ remains greater.

Claim. Once $G(t) > 1$ at some $t \geq t'$, it remains > 1 at all later times, and in fact $G(t)$ grows over time (unless it is already unbounded).
Reasoning. For $t \geq t'$, since $V_{i_1}(t) > V_{i_2}(t)$, the probability $P_{k,i_1}(t)$ that a new follower chooses i_1 can stay close to or above $1/2$, whereas $P_{k,i_2}(t)$ is below $1/2$. As a result,

$$\frac{dF_{i_1}(t)}{dt} \geq \frac{1}{2} F_{\text{new}}(t), \qquad \frac{dF_{i_2}(t)}{dt} \leq \frac{1}{2} F_{\text{new}}(t).$$

Thus $F_{i_1}(t)$ grows at least as fast as $F_{i_2}(t)$. In fact, because $\frac{\pi_{i_1}}{c_{i_1}} > \frac{\pi_{i_2}}{c_{i_2}}$ *and* $F_{i_1}(t)$ starts out bigger up to that proportionality, the margin in $V_{i_1}(t)$ vs. $V_{i_2}(t)$ generally *widens*.

Concretely, if we define

$$R(t) = \frac{F_{i_1}(t)}{F_{i_2}(t)},$$

then for large t (assuming $R(t) \geq 1$), the difference in new inflows makes $R(t)$ grow. Hence $R(t)$, and therefore $G(t) = \left(\frac{\pi_{i_1}/c_{i_1}}{\pi_{i_2}/c_{i_2}}\right) R(t)^2$, diverges away from 1 as t increases, ensuring $V_{i_1}(t) > V_{i_2}(t)$ persistently.

Step 5: Concluding dominance.
Since $V_{i_1}(t) > V_{i_2}(t)$ for all sufficiently large t, the net inflow of followers for i_1 strictly exceeds that for i_2 (the exact fraction can be near 1 if $V_{i_1}(t) \gg V_{i_2}(t)$). Therefore,

$$\lim_{t \to \infty} \left[F_{i_1}(t) - F_{i_2}(t) \right] > 0,$$

and in fact one can pick some finite $T > t'$ such that $F_{i_1}(t) > F_{i_2}(t)$ for all $t > T$. This proves finfluencer i_1 *eventually dominates* i_2 in follower count whenever $\frac{\pi_{i_1}}{c_{i_1}} F_{i_1}(t') > \frac{\pi_{i_2}}{c_{i_2}} F_{i_2}(t')$.

This completes the proof of the Proposition. □

Proof of Proposition 2. We wish to show that, under the assumptions

$$\frac{\pi_A}{c_A} > \frac{\pi_S}{c_S} \quad \text{and sufficiently large initial scale or short-run payoff for antiskilled,}$$

the *antiskilled* finfluencers eventually dominate the visibility measure and hence bias the aggregate sentiment. Formally, for each stock j, the visibility-weighted sentiment of type A

will exceed (and in the limit vastly exceeds) that of type S. In addition, the fraction of total visibility contributed by skilled finfluencers converges to zero.

Step 1: Comparing growth of antiskilled vs. skilled.

Recall that each finfluencer i has a flow of new followers governed by

$$\frac{dF_i(t)}{dt} = F_{\text{new}}(t) \times \bar{P}_i(t),$$

where

$$\bar{P}_i(t) = (\text{average probability that a newly arriving follower chooses finfluencer } i).$$

From the model, this choice probability depends on the perceived utility

$$U_{ki}^F(t) = \underbrace{\theta_{k,i}(t)U_S + (1-\theta_{k,i}(t))U_B}_{\text{type-based payoff}} + \underbrace{\gamma V_i(t)}_{\text{direct visibility}},$$

with $U_S > U_B$. In the long run, if visibility differences become large, the term $\gamma V_i(t)$ will dominate. Now,

$$V_i(t) = \kappa \cdot \frac{\pi_i}{c_i}\big[F_i(t)\big]^2 \times \Phi_i(t),$$

where $\Phi_i(t)$ collects any further dependence on the finfluencer's tweeting strategy (e.g. volume $v_i(t)$, sentiments $\{\text{SocSent}_{i,j,t}\}$, etc.). For the argument, it suffices that $\Phi_i(t)$ remains bounded away from zero and does not diverge to infinity.[1]

Let $i = S$ denote a skilled finfluencer, and let $i = A$ denote an antiskilled finfluencer. Suppose $\frac{\pi_A}{c_A} > \frac{\pi_S}{c_S}$. We claim that, if an antiskilled finfluencer A has a sufficiently large initial base $F_A(t')$ or strong short-run advantage, then *asymptotically* $F_A(t)$ grows faster than any skilled finfluencer's $F_S(t)$.

Intuition. For large $F_i(t)$, visibility $V_i(t) \approx \kappa(\pi_i/c_i)\big[F_i(t)\big]^2$ is the primary driver of new-follower inflows. Because $\pi_A/c_A > \pi_S/c_S$, if $F_A(t)$ is not too small relative to $F_S(t)$ at some point in time, then $\big(\pi_A/c_A\big)\big[F_A(t)\big]^2$ will catch up to and surpass $\big(\pi_S/c_S\big)\big[F_S(t)\big]^2$. From that point onward, the probability that new followers join the antiskilled finfluencer (rather than a skilled one) remains close to 1, further widening the gap.

Formal argument. For large t, we can approximate (suppressing lower-order terms),

$$\frac{dF_A(t)/dt}{F_A(t)} \approx F_{\text{new}}(t)\frac{\gamma V_A(t)}{\gamma V_A(t) + (\text{visibility of all other types})} \approx \frac{\gamma F_{\text{new}}(t)\kappa\frac{\pi_A}{c_A}\big[F_A(t)\big]^2}{\gamma\kappa\frac{\pi_A}{c_A}\big[F_A(t)\big]^2 + \dots},$$

[1] In other words, each type has some stable or bounded approach to choosing volume/sentiment; we do not require that $\Phi_i(t)$ be constant, only that it does not vanish or explode to invalidate the main growth rates. One can ensure this by standard concavity assumptions in $f(\cdot)$ and $g(\cdot)$ in the finfluencer's utility, or by bounding tweet volume.

while for a skilled finfluencer S,

$$\frac{dF_S(t)/dt}{F_S(t)} \approx \frac{\gamma F_{\text{new}}(t)\kappa \frac{\pi_S}{c_S}\left[F_S(t)\right]^2}{\gamma\kappa\frac{\pi_A}{c_A}\left[F_A(t)\right]^2 + \dots}.$$

If $\frac{\pi_A}{c_A} > \frac{\pi_S}{c_S}$ and $F_A(t)$ is not initially negligible, then the ratio

$$\frac{V_A(t)}{V_S(t)} = \frac{\frac{\pi_A}{c_A}\left[F_A(t)\right]^2}{\frac{\pi_S}{c_S}\left[F_S(t)\right]^2}$$

tends to increase over time, so $F_A(t)$ eventually outstrips $F_S(t)$. Hence, a critical mass of antiskilled finfluencers (or even one sufficiently advantaged) captures most new followers. This establishes the *dominance in total follower counts.*

Step 2: Visibility dominance and implication for weighted sentiment.

Once $F_A(t) \gg F_S(t)$ for antiskilled vs. skilled finfluencers (and remains so for large t), we have

$$V_A(t) = \kappa\frac{\pi_A}{c_A}[F_A(t)]^2 \gg \kappa\frac{\pi_S}{c_S}[F_S(t)]^2 = V_S(t).$$

As a result, $\sum_{i:\tau_i=A} V_i(t)$ dwarfs $\sum_{i:\tau_i=S} V_i(t)$.

By definition, for each stock j, the *visibility-weight* of finfluencer i who tweets about j is

$$w_{i,j}(t) = \frac{V_i(t)\text{Tweet}_{i,j}(t)}{\sum_m V_m(t)\cdot\text{Tweet}_{m,j}(t)}.$$

Since $\sum_{i:\tau_i=A} V_i(t) \gg \sum_{i:\tau_i=S} V_i(t)$, almost all of the denominator is contributed by type A. Consequently, even if not all antiskilled finfluencers tweet about stock j, *some* subset of them that does tweet about j will collectively exceed the corresponding subset of skilled tweeters. Formally, for large t,

$$\sum_{i:\tau_i=A} w_{i,j}(t)\text{SocSent}_{i,j,t} \quad\gg\quad \sum_{i:\tau_i=S} w_{i,j}(t)\text{SocSent}_{i,j,t}.$$

Hence the *weighted sum of sentiments*—the measure used to construct, for example, "abnormal sentiment"—is dominated by antiskilled content.

Step 3: Convergence of $\bar{\theta}(t)$ to 0.

Finally, let $\bar{\theta}(t)$ be the perceived fraction of skilled finfluencers under a visibility-based weighting. One natural way to define it is:

$$\bar{\theta}(t) = \frac{\sum_{i:\tau_i=S} V_i(t)}{\sum_m V_m(t)} = \frac{V_S^{(\text{total})}(t)}{V_S^{(\text{total})}(t) + V_A^{(\text{total})}(t) + \dots}.$$

Since $V_A^{(\text{total})}(t)$ grows faster than $V_S^{(\text{total})}(t)$ (by the argument in Step 1), the ratio converges to zero:

$$\lim_{t\to\infty} \bar{\theta}(t) = 0,$$

implying that the fraction of total visibility accounted for by skilled finfluencers vanishes in the limit, no matter that skilled agents have fundamentally more accurate signals. The higher reward–cost ratio of antiskilled users (π_A/c_A) and their sufficient initial foothold ensure they dominate in followers, engagement, and thus measured visibility.

Step 4: Biased or incorrect signals.

Because type A (antiskilled) finfluencers systematically produce negative alpha calls (or systematically the wrong sign of future returns), the dominance of their weighted sentiment implies the *aggregate* or "consensus" on social media is biased away from fundamentals. Skilled finfluencers' correct signals are eventually outweighed. Empirically, this corresponds to observing (i) persistently overoptimistic sentiment from the high-engagement antiskilled group, and (ii) the overall average sentiment being a poor or contrarian predictor of returns.

Conclusion. Under the condition $\frac{\pi_A}{c_A} > \frac{\pi_S}{c_S}$ and sufficiently large initial scale of antiskilled finfluencers, we have shown that, as $t \to \infty$,

$$\sum_{\tau_i=A} V_i(t) \gg \sum_{\tau_i=S} V_i(t), \quad \text{and hence} \quad \bar{\theta}(t) = \frac{\sum_{\tau_i=S} V_i(t)}{\sum_m V_m(t)} \longrightarrow 0.$$

For each stock j, the visibility-weighted sentiment is dominated by type A, causing the aggregate belief about returns to be systematically distorted. This completes the proof of the proposition. □

Linking policies to the results

In this section, we explain how the policies fit into the model framework (Propositions 1 and 2).

1. Quality-Based Visibility Boost. The policy of *increasing visibility for good (skilled) finfluencers* can be implemented in the extended model by allowing the parameter γ (which multiplies $V_i(t)$ in the follower's utility) to differ between skilled and bad types. For example:

$$U_{k,i}^F(t) = \left[\theta_{k,i}(t)U_S + (1-\theta_{k,i}(t))U_B\right] + \gamma_i V_i(t),$$

where

$$\gamma_S > \gamma_B.$$

If the platform increases γ_S (or equivalently multiplies the visibility $V_i(t)$ by a factor greater than 1 *only for skilled finfluencers*), it makes $U_{k,i}^F(t)$ larger for skilled types at any given follower base $F_i(t)$, thus attracting more new followers to them.

- *Link to Proposition 1.* By boosting the term $\gamma_S V_S(t)$, you effectively increase the speed at which skilled types can accumulate new followers. This can prevent a scenario in which $\frac{\pi_B}{c_B} F_B(t)^2$ dwarfs $\frac{\pi_S}{c_S} F_S(t)^2$ over time, thereby mitigating the bad finfluencer's dominance result.

- It is worth noting that if a "bad" type simultaneously raises her own engagement $e_B(t)$ in response, one could see an *arms race*. In other words, if bad types increase their effort in response to having lower γ_B, they can partially offset the penalty.

2. Penalizing Low-Quality Content. This policy amounts to reducing the effective visibility of antiskilled (or low-quality) finfluencers— *e.g.*, demoting them in recommendation feeds or search results. It can be done by lowering γ_B *below* its baseline or imposing a multiplier < 1 on $V_B(t)$.

- *Interpretation in the model.* If previously $U^F_{k,B}(t) = \theta_{k,B}(t)U_S + \cdots + \gamma V_B(t)$, then "penalizing" is akin to forcing $\gamma_B < \gamma$ or directly capping $V_B(t)$ in the follower's utility.

- *Connection to Proposition 2.* By lowering the visibility advantage of bad finfluencers, we make it less likely that $\frac{\pi_B}{c_B} F_B(t)^2$ alone can propel them to long-run dominance. This weakens the main condition $\frac{\pi_B}{c_B} > \frac{\pi_S}{c_S}$ in practice, because the *effective* reward-to-cost ratio for bad types is reduced.

3. Transparency and Verification. This policy increases the *intrinsic* payoff to following a skilled finfluencer, for instance by letting them wear a "verified" badge or curation label that assures users "this is a reliable source." In the model, that is captured by raising the skilled type's U_S or, in some forms, shifting followers' prior beliefs $\theta_{k,S}(t)$ upward.

- We might write

$$U^F_{k,i}(t) = \theta_{k,i}(t)(U_S + \delta_{\text{verified}}) + (1 - \theta_{k,i}(t))U_B + \gamma V_i(t),$$

  where $\delta_{\text{verified}} > 0$ is a "credibility premium." Now skilled finfluencers produce higher base utility for the follower.

- By making U_S effectively bigger, skilled finfluencers have an edge in the fraction of new followers even at early stages when their $V_S(t)$ might be smaller. Hence, the threshold needed for bad finfluencers to overtake them (Proposition 1) gets higher, and it becomes less likely that $\bar{\theta}(t)$ converges to 0 in Proposition 2.

Table IA.1: Robustness: Alternative Specifications of the Mixture Model

Panel A reports the results of fitting mixture models with one, two, and three components for skilled and antiskilled finfluencers. Means and probabilities are reported in percentage points. Panel B reports the results of fitting mixture models over different horizons H with two exponentials on the $\alpha > 0$, two exponentials on $\alpha < 0$, and a mass at $\alpha = 0$. We calculate excess returns over the next 1, 2, 5, 10, and 20 trading days using the Fama-French five-factor model. The first number at the top of the columns shows the horizon of future returns. The measured alpha ($\tilde{\alpha}$) for each user is the average of signed adjusted returns after her tweets. For each horizon, the first column shows the mean of each component (μ's), and the second column shows the weight of the component in the mixture (π's). Means and probabilities are in percentage points.

Panel A: Model estimates for different number of sub-groups K

	(1)		(2)		(3)	
	$K^+ = K^- = 1$		$K^+ = K^- = 2$		$K^+ = K^- = 3$	
	μ_k (%)	π_k (%)	μ_k (%)	π_k (%)	μ_k (%)	π_k (%)
Skilled type 3					8.36	4.8
Skilled type 2			8.14	5.1	1.53	23.4
Skilled type 1	4.53	15.1	1.49	23.5	0.11	6.8
Unskilled	0.00	56.3	0.00	16.6	0.00	0.0
Antiskilled type 1	-4.89	26.6	-1.19	45.5	-0.45	29.7
Antiskilled type 2			-9.15	9.3	-1.81	27.8
Antiskilled type 3					-10.11	7.6
N	29,475		29,475		29,475	
Log Likelihood	-89,600		-88,878		-88,858	
AIC	179,207		177,771		177,740	
BIC	179,240		177,838		177,839	

Panel B: Model estimates for different forecast horizon H

	(1)		(2)		(3)		(4)		(5)	
	$H = 1$		$H = 2$		$H = 5$		$H = 10$		$H = 20$	
	μ_k	π_k	μ_k	π_k	μ_k	π_k	μ_k	π_k	μ_k	π_k
Skilled type 2	4.28	1.1	4.08	1.4	4.76	2.6	6.54	3.5	8.14	5.1
Skilled type 1	0.43	17.2	0.68	17.0	0.82	19.6	1.26	19.2	1.49	23.5
Unskilled	0.00	53.6	0.00	48.7	0.00	31.5	0.00	29.4	0.00	16.6
Antiskilled type 1	-0.34	25.4	-0.47	29.7	-0.59	40.5	-0.85	39.9	-1.19	45.5
Antiskilled type 2	-2.99	2.7	-3.60	3.2	-5.01	5.8	-6.29	8.0	-9.15	9.3
N	30,721		30,329		30,175		30,054		29,475	
Log Likelihood	-46,891		-56,247		-69,004		-79,626		-88,878	
BIC	93,798		112,510		138,023		159,269		177,771	
AIC	93,865		112,577		138,090		159,335		177,838	