

Finfluencers*

Ali Kakhbod
UC BERKELEY

Seyed Kazempour
RICE UNIVERSITY

Dmitry Livdan
UC BERKELEY, CEPR

Norman Schürhoff
UNIVERSITY OF LAUSANNE, SFI, CEPR

August 5, 2023

*We thank Bloomberg and StockTwits for providing the data. 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 Jesse H. Jones School of Business, Rice University, Houston, TX, 77005. e-mail smkazempour@rice.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

Tweet-level data from a social media platform reveals high dispersion and systematic bias in the quality of advice by financial influencers, or “finfluencers”: 28% of finfluencers are skilled, generating 2.6% monthly abnormal returns, 16% are unskilled, and 56% have negative skill (“antiskill”) generating -2.3% monthly abnormal returns. Antiskilled finfluencers have more followers and more influence on retail trading than skilled finfluencers. The advice by antiskilled finfluencers creates overly optimistic beliefs most times and persistent swings in followers’ beliefs. Consequently, finfluencers cause excessive trading and inefficient prices such that a contrarian strategy yields 1.2% monthly out-of-sample performance.

JEL Classification: G12, G14, G41

Key words: Finfluencers, social media, mixture modeling, retail traders, homophily, belief bias

Financial influencers, commonly known as finfluencers, are individuals who provide unsolicited investment advice on social media platforms. Many finfluencers have large followings and their recommendations can have a significant impact on the investment decisions made by retail investors. The Securities and Exchange Commission (SEC) has been concerned about finfluencers, particularly because most of them provide investment advice or recommendations to the public without being registered as investment advisers or brokers. 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.¹ However, despite their growing influence, little is known about the quality of the unsolicited financial advice provided by individual finfluencers, the impact of finfluencers' advice on their follower base, trading activity, and asset prices.²

This paper assesses the quality of investment advice provided by different types of finfluencers, in contrast to the prior literature that has focused on the "average finfluencer." Using tweet-level data from StockTwits on over 29,000 finfluencers, we classify each finfluencer into three major groups: Skilled, unskilled, and antiskilled, defined as those with negative skill. While the average finfluencer has skill close to zero, we find that 28% of finfluencers provide valuable investment advice that leads to monthly abnormal returns of 2.6% on average, and only 16% of finfluencers are unskilled. The majority of finfluencers, 56%, are antiskilled and following their investment advice yields monthly abnormal returns of -2.3%. Unskilled and antiskilled finfluencers have more followers, more activity, and more influence on retail trading than skilled finfluencers.

The core of our analysis is the assessment of the finfluencers' quality for which distinguishing

¹See, 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](https://www.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](https://www.sec.gov/news/press-release/2022-221)).

²The SEC ([sec.gov/oiea/investor-alerts-and-bulletins/social-media-and-investment-fraud-investor-alert](https://www.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 influencers, particularly when the influencers have a financial interest in the products or services they are promoting. The SEC, for instance, advises investors to be cautious when considering investment advice from any source and to do their own research and due diligence before making any decisions. Investors can also check the registration status of investment advisers and brokers through the SEC's Investment Adviser Public Disclosure (IAPD) website.

between (anti)skill and luck is important. To resolve this issue we employ a mixture-modeling approach with multiple types and non-normal distributions which allows us to estimate the distribution of true skills or alphas among all users on StockTwits. Mixture modeling involves fitting a distribution that is a combination of multiple other distributions, known as components, to a set of data. We allow for three types of StockTwits users: skilled users with positive true alpha drawn from a mixture of exponential distributions, unskilled users with zero true alpha, and antiskilled users with a negative true alpha drawn from a mixture of negative exponential distributions. We use these assumptions to obtain a joint distribution for influencers' alpha that is a combination of exponentials for skilled users, a mass at zero for unskilled users, negative exponentials for antiskilled users, all combined with a Gaussian distribution for capturing luck.

We investigate how users' skill relates to their tweeting activity. We find that skilled influencers are less active than unskilled and antiskilled influencers. Users who tweet more frequently are less skilled in that a ten times increase in the total number of tweets posted by a user is associated with a 3.7% decrease in the probability of being skilled and a 0.08% decline in the monthly expected true alpha. Additionally, the tweet composition correlates with the degree of its informativeness as users posting more negative tweets tend to be more skilled. A one percent increase in the share of negative tweets is associated with a 0.01% increase in the expected true alpha and a 0.06% increase in the probability of being skilled. Finally, the users' skill is persistent. To study persistence, we split the sample into two halves and estimate users' skills separately in each half of the data. We find that a one percent increase in the expected true alpha over the first half of the data predicts a 0.09% increase in the expected true alpha over the second half.

Next, we dissect influencers' tweeting strategies, that is when and what they tweet, to check whether they possess unique skills or just follow commonly known investment behaviors including momentum, contrarian, return chasing, and herding.³ We find that skilled influencers are return-, social sentiment-, and news-contrarian. They also do not chase returns and do not herd on other users' tweets. A one percent increase in our measure of return chasing is associated with a 0.08%

³We 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 that are about stocks in the highest decile of overall positive tweeting volume over the past five days.

decrease in the probability of the user being skilled while a one percent increase in our measure of herding tendency is associated with a 0.09% decrease in the probability of being skilled. Antiskilled influencers ride return momentum and social sentiment momentum and tend to chase returns. A one percent increase in our measure of return chasing is associated with a 0.16% increase in the probability of the user being antiskilled. We also check if users with more negative tweets are more informed. The existing 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 as negative information is not incorporated into prices and, hence, stocks with higher short-selling constraints underperform. We test these hypotheses in our data using Markit’s measure of short-selling constraints for stocks and calculating the average decile of short-selling constraints for the user’s positive and negative tweets separately. The results show that users who tweet negatively about stocks with higher short-selling constraints are indeed more likely to be skilled. A one-unit increase in our measure of short-selling constraints among the user’s positive/negative tweets is associated with a 0.31% decrease in expected true alpha and a 0.89% decrease in the probability of being skilled.

Following our analysis of influencers’ tweeting activity and strategies, we shift our focus to the trading behavior of retail investors. The observed relation between tweeting activity and our measures of skill suggests social media users can and should use tweeting behavior to identify skilled influencers. However, a striking feature of the data is that more skilled influencers have *fewer* followers while less skilled influencers have *more* followers, with antiskilled influencers being the most popular, consistent with skill being effectively ignored (Golub and Jackson, 2012; Berk and Van Binsbergen, 2022; Pedersen, 2022). We find that retail investors are influenced differently by different types of StockTwits users. The advice by skilled influencers has little to no impact on retail order imbalances.⁴ However, the advice by antiskilled influencers strongly predicts retail order imbalances, and in a way that they follow the flawed advice leading to negative returns for retail investors. Moreover, in line with retail investors stubbornly following the crowd of antiskilled influencers, we find evidence for this “wisdom of the antiskilled crowd” in both in-sample and

⁴Retail order imbalances are computed following the approach in Boehmer, Jones, Zhang, and Zhang (2021).

out-of-sample tests.

To address the joint endogeneity of stock returns, retail order imbalances, and tweets we utilize a panel VAR specification that treats all variables as endogenous and interdependent, both in a contemporaneous and dynamic sense. We find that retail order imbalances and positive social sentiment by skilled influencers positively predict future returns whereas negative social sentiment by skilled influencers negatively predicts future returns. In line with our hypothesis, positive (negative) social sentiment by antiskilled influencers negatively (positively) predicts future returns.

Last, we explore the asset price distortions and aggregate belief biases introduced by following antiskilled influencers' advice. Following the advice by antiskilled influencers creates overly optimistic beliefs most of the time since their tweets tend to be bullish about most stocks, and overly pessimistic beliefs some of the time when their tweets are more pessimistic than the skilled influencers' tweets. Furthermore, the social media sentiment by antiskilled influencers is highly persistent and induces long swings in the magnitude of their followers' belief bias. As a result, one can earn 1.2% monthly out-of-sample buy-and-hold abnormal returns by trading against the antiskilled influencers' advice. When we combine these results with our additional findings that the influencers' skills are persistent but are not sufficient for influencers' survival, we conclude that social media platforms "sell the sizzle, not the steak." In other words, as long as there are antiskilled influencers "preaching" their message there will be investors who tend to like influencers' message and are willing to trade on it.

Literature review. Our results are important in a number of 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, i.e. charlatans, in equilibrium, providing a framework that aligns well with our observations that unskilled and antiskilled influencers, despite their lack of ability, command larger followings on StockTwits than their skilled counterparts. On the other hand, 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 influencers, but often neglect it. Our paper supports findings 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 adds to the literature on investor expectations. For instance, 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 most influencers on StockTwits, who also hold extrapolative beliefs. We extend this result by showing that a subset of influencers hold accurate beliefs about future stock returns, allowing them and their followers to potentially counterbalance the misguided actions of antiskilled influencers and their followers.⁵

Third, our paper draws on the literature on mixture modeling in asset pricing. Using a similar methodology, Harvey and Liu (2018) show that some mutual fund managers have true positive alphas. Similarly, our paper uncovers that despite the negative correlation between StockTwits' average sentiment and future returns, some influencers post informative tweets. However, unlike in the context of mutual fund management, we find that users do not flock to the most skilled influencers in the social media landscape.

In addition, our findings complement existing research on crowdsourcing platforms for stock prediction, such as Seeking Alpha. Chen, De, Hu, and Hwang (2014) find a positive correlation between the sentiment expressed in Seeking Alpha articles and future stock returns.⁶ Dim (2022) employs a mixture modeling methodology on Seeking Alpha articles and discovers that a considerable majority (56%) of its authors correctly predict stock returns. This finding is in stark contrast to our results, which indicate that the majority of StockTwits users are antiskilled, and only 28% can predict the direction of stock returns, leading to a negative correlation between sentiment and future returns. This divergence is, however, consistent with findings by Cookson, Lu, Mullins, and

⁵Despite antiskilled influencers' 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.

⁶Many more papers analyze the informativeness of social media and information crowdsourcing platforms. Interested readers can sample from the following list (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; Sprenger, Tumasjan, Sandner, and Welpe, 2014; Curtis, Richardson, and Schmardebeck, 2014; Azar and Lo, 2016; Bartov, Faurel, and Mohanram, 2018; Campbell, D'Adduzio, and Moon, 2021; Cookson, Engelberg, and Mullins, 2021).

Niessner (2022) who show that while attention is highly correlated across platforms, sentiment is not. Thus, our results underscore the importance of understanding the unique characteristics of each platform and the incentives faced by social media users.

Farrell, Green, Jame, and Markov (2022) demonstrate that the publication of Seeking Alpha articles within a trading day significantly increases retail trading activity compared to periods immediately before publication. Similarly, we find that tweet sentiment from the majority of influencers predicts retail order imbalance. The fact that retail order imbalance follows sentiment despite the disparity in the informativeness of the content highlights our key finding: users do not sufficiently distinguish between skilled and unskilled or antiskilled influencers. This observation suggests that platform dynamics and user behaviors significantly affect how investors receive and act upon information.

Finally, our paper contributes to the literature on the individual skills of influencers, their behavioral traits, and how these aspects influence their follower base. Cookson, Engelberg, and Mullins (2023) show that social media users follow influencers with similar beliefs and, as a result, live in their own bubbles, a phenomenon called information siloing.⁷ We complement their finding by demonstrating that social media users can identify skilled influencers, but often choose to follow antiskilled influencers who exhibit similar behavioral traits to their own. This behavior mirrors the sociological phenomenon of homophily, the tendency of individuals to associate and bond 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).⁸ In the context of our study, homophily manifests as social media users preferring to follow influencers who exhibit similar investment behaviors to their own, even if these behaviors lead to suboptimal investment outcomes.

⁷A broader literature, exemplified by Barber and Odean (2007) demonstrates several behavioral biases and traits of retail investors. In this paper, we choose to focus on influencers only.

⁸In contrast to homophily, echo chambers refer to situations where individuals are only exposed to information or viewpoints that confirm their existing beliefs and opinions and are sheltered from opposing perspectives. This can lead to the reinforcement of existing beliefs and result in limited exposure to diverse ideas and information. While homophily can contribute to the creation of echo chambers, the two concepts are different. Homophily refers to the preference for social connections with individuals with similar traits, while echo chambers refer to the phenomenon of information reinforcement within a particular group or community.

The rest of the paper is structured as follows. Section 1 describes the data and discusses the measurement of alpha for every StockTwits user in our sample. Section 2 addresses the type 1 and 2 errors of statistical tests on estimated alphas, measures true alpha for each finfluencer, and develops various measures of finfluencer skill. Section 3 documents the relation between skill and popularity of finfluencers. Section 4 tests if social media users can tell apart skilled from unskilled and antiskilled finfluencers. Section 5 explores the asset price distortions and aggregate belief biases introduced by following antiskilled finfluencers' advice. Section 6 concludes.

1 Data and Estimated Alphas

This section describes the data and discusses the measurement of alpha for every StockTwits user in our sample, which we call finfluencer if the user posts and not only follows others.

1.1 Data

Data sources. We collect data from several sources. We obtain tweet data from Bloomberg, finfluencer user-level data from StockTwits, 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, and TAQ to compute retail order imbalances. Our sample period covers July 13, 2013, through January 1, 2017.

The Bloomberg data contains for each tweet 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 that is 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. The social sentiment score by user i in stock j for its n th 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 positive/neutral/negative.

The Bloomberg data also contains news data. For each news story, it 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 n th 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 positive/neutral/negative. Comparing news to social sentiment, these statistics show that tweets are less likely negative than news.

We use the StockTwits API to collect user data for each user.⁹ The StockTwits data contain for each user the number of tweets, with a mean of 131.62, a minimum of 1, and a maximum of 615,145, the number of followers, number of other users being followed, number of stocks on the user’s watch list, number of investment ideas, and number of likes by other users as of the time of our download.

Matching and cleaning. The user name supplied by Bloomberg is the StockTwits user name displayed on the screen. We match the StockTwits user name supplied by 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. As a result, some users cannot uniquely be matched from Bloomberg to StockTwits and we pool or, alternatively, eliminate the duplicates.

The matching of returns and tweets is also important. We apply the following procedure: If a tweet was posted during trading hours, we match it to the same trading day. That is, day t will be the trading day. If a tweet was posted after hours, on holidays, or on 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 was posted.

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

$$SocSent_{i,j,t} = \max \left\{ 1, \min \left(1, \sum_{n=1}^{N_{i,j,t}} \mathbf{1}(SocSent_{i,j,t,n} = 1) - \sum_{n=1}^{N_{i,j,t}} \mathbf{1}(SocSent_{i,j,t,n} = -1) \right) \right\}, \quad (1)$$

⁹There were a total of 139,401 users as of February 2, 2018, when the data was collected. 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.

where $n = 1, \dots, N_{i,j,t}$ is the index of the tweet. The max and min operators are used to normalize $SocSent_{i,j,t}$ to the $[-1, 1]$ interval.

Stock abnormal returns. Abnormal returns are computed according to the following standard procedure. First, we calculate factor exposures $\beta_{j,t}$ for each stock j on trading day t by running daily regressions of excess returns on Fama/French factors over the year ending on the day t skipping the last month:

$$R_{j,t} - R_{f,t} = \alpha_{j,t} + \beta_{j,t}^{\theta} F_t + \epsilon_{j,t}, \quad \text{for days in } [t - 252, t - 21], \quad (2)$$

where F_t is a vector of Fama/French (one, three, or five) factors. Then, equipped with the estimated factor loadings from the first stage, $\hat{\beta}_{j,t}$, we calculate future abnormal returns for stock j over horizon H (e.g., 1, 2, 5, 10, or 20 days) using the following equation:

$$AbnRet_{j,t+1,t+H} = R_{j,t+1,t+H} - R_{f,t+1,t+H} - \hat{\alpha}_{j,t} - \hat{\beta}_{j,t}^{\theta} F_{t+1,t+H}. \quad (3)$$

Results are very similar if we estimate (2) and (3) without intercepts $\alpha_{j,t}$.

Computing influencer-level abnormal returns. We calculate user-specific abnormal returns, α_i , for each user i over different horizons $[t + 1, t + H]$, $H \in \{1, 2, 5, 10, 20\}$. We calculate the mean signed abnormal return and its standard error for every user in the data by running univariate regressions:¹⁰

$$SocSent_{i,j,t} - AbnRet_{i,j,t+1,t+H} = \alpha_i + \epsilon_{i,j,t+1,t+H}, \quad (5)$$

for all N_i stock-days for which $SocSent_{i,j,t} \neq 0$ and separately for all users $i = 1, \dots, I$ and multiple values of H . Equipped with user-specific abnormal returns $\tilde{\alpha}_i$, $i = 1, \dots, I$, over horizon H we now document mean signed returns and their t -stats.

¹⁰Alternatively, we have run multivariate regressions for all users $i = 1, \dots, I$ combined and multiple values H :

$$SocSent_{i,j,t} - AbnRet_{i,j,t+1,t+H} = \sum_{\iota=1}^I \alpha_{\iota} \mathbb{1}(\text{User } i = \iota) + \epsilon_{i,j,t+1,t+H}. \quad (4)$$

The results for the multivariate regression are very similar to (5).

Table 1: Summary Statistics of Users' Estimated Alphas ($\tilde{\alpha}_i$)

This table reports summary statistics of estimated alphas (α), their standard errors, and t -statistics. We calculate excess returns over the next 20 trading days using the Fama-French five-factor model. The estimated alpha (α) for each user is the average of signed adjusted returns after her tweets. Alphas and their standard errors are in percentage points.

Panel A: Distribution of α_i			
	α_i	S.E.	t -stat
Mean	-0.63	3.88	-0.90
S.D.	6.52	4.04	89.88
P10	-7.01	0.84	-2.22
P25	-2.82	1.45	-1.11
P50	-0.35	2.61	-0.16
P75	1.86	4.75	0.72
P90	5.44	8.42	1.66
N	29,477	29,477	29,477

Panel B: Significance of $H_0 : \alpha_i = 0$	
	Fraction of significant α_i
$p < 0.05$	19.8%
$p < 0.10$	25.7%

1.2 Estimated alphas across influencers

Table 1 reports users' estimated skills ($\tilde{\alpha}_i$) from specification (5) with $H = 20$ business days. The average user has a monthly estimated alpha of -0.63% (annualized: -7.56% per year). The median estimated alpha is -0.35% and hence also negative, meaning that most users post systematically anti-informative tweets. These results confirm the findings in previous papers that average social media users are systematically wrong in predicting stock returns (Giannini, Irvine, and Shu, 2018). However, Table 1 shows that the 75th percentile of estimated alpha is 1.86% per month, which is economically large.

Table 1 also shows that the standard errors of estimated alphas are large compared to the point estimates. The average (median) standard error is 3.88% (2.61%) monthly. However, despite the relatively large standard errors, some users have statistically significant estimated alphas. In column 3, the 10th percentile of t -statistics is -2.22, while the 90th percentile is 1.66. Panel B shows that the proportion of users for whom the p -value of the estimated alpha is less than 5% (10%) is 19.8% (25.7%). These numbers are larger than what we would expect if all users were

uninformed ($\alpha = 0$). The distribution in Table 1 shows that many StockTwits users achieve significant alphas, with either positive or negative signs. Thus, valuable information appears to be disclosed on StockTwits.

However, the issue is that the statistical tests have a size and power. If we use the t -stat threshold of 1.96, we know that 5% of users will appear with significant alpha (mean signed abnormal returns) even if the true alpha is zero. Hence, there are users with truly positive (or truly negative) alpha that we cannot detect when the t -stat is less than 1.96. While we can measure alpha for every user and calculate its t -stat, it is unclear how often the null of $\alpha = 0$ is falsely rejected or falsely accepted (type 1 and 2 errors).

2 Model of Finfluencer Skill

This section addresses the type 1 and 2 errors of statistical tests on estimated alphas, $\tilde{\alpha}_i$, measures true alpha for each finfluencer, and develops various measures of finfluencer skill.

2.1 Mixture modeling of finfluencer skill

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

$$\tilde{\alpha}_i = \alpha_i + \epsilon_i, \tag{6}$$

where $\epsilon_i \sim \mathcal{N}(0, \tilde{\sigma}_i^2)$ and $\tilde{\sigma}_i$ is the standard error of user i 's abnormal return in the data. It follows that the distribution of the observed skill can be calculated as 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 a mixture modeling methodology to estimate the distribution of α among users.

We motivate our model of true skills with the following economic assumptions. We assume there are three types of StockTwits users and they can consist of several subtypes:

1. Skilled users, whose true skill is positive: $\alpha_i > 0$.

2. Unskilled users, whose true skill is zero: $\alpha_i = 0$.

3. Antiskilled users, whose alpha is negative: $\alpha_i < 0$.

For the types of skilled and antiskilled users, respectively, we further assume there can be several subtypes with different levels of (anti)skill. Suppose 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 . Further, we assume that the skilled and antiskilled types are exponentially distributed, which is the maximum-entropy distribution having the greatest uncertainty consistent with the type constraints. Then, true skill α is distributed across finfluencers according to the finite mixture distribution

$$f(\alpha) = \mathbb{1}_{f\alpha > 0} g \sum_{k=1}^{K^+} \pi_k^+ g(\alpha; \mu_k^+) + \pi^0 \mathbb{1}_{f\alpha = 0} g + \mathbb{1}_{f\alpha < 0} g \sum_{k=1}^{K^-} \pi_k^- g(\alpha; \mu_k^-), \quad (7)$$

where $g(\alpha; \mu) = \frac{1}{\mu} \exp(-\frac{1}{\mu}\alpha)$ if $\mu > 0$ ($g(\alpha; \mu) = \frac{1}{|\mu|} \exp(\frac{1}{|\mu|}\alpha)$ if $\mu < 0$) is an exponential distribution with a mean of μ and

$$\begin{aligned} \sum_{k=1}^{K^+} \pi_k^+ + \pi^0 + \sum_{k=1}^{K^-} \pi_k^- &= 1, \\ \mu_k^+ &> 0 \text{ for } 1 \leq k \leq K^+, \\ \mu_k^- &< 0 \text{ for } 1 \leq k \leq K^-. \end{aligned} \quad (8)$$

In expression (7), μ_k^+ and μ_k^- are the expected abnormal returns of the positive and negative components $k = 1, \dots, K^+$ (K^-). π_k^+ , π_k^- , and π^0 denote the probability of positive, negative, and zero components, respectively.

Given that $\tilde{\alpha}_i = \alpha_i + \epsilon_i$, the distribution of estimated alphas, $\tilde{\alpha}_i$, can be calculated as the convolution of f and a mean-zero Normal distribution with standard deviation $\tilde{\sigma}_i$, i.e.,

$$G(\tilde{\alpha}_i; \tilde{\sigma}_i, \Theta) = (f \otimes \phi_{\tilde{\sigma}_i})(\tilde{\alpha}_i), \quad (9)$$

where \otimes 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 $\Theta = (\mu_1^+, \dots, \mu_{K^+}^+, \mu_1^-, \dots, \mu_{K^-}^-, \pi_1^+, \dots, \pi_{K^+}^+, \pi_1^-, \dots, \pi_{K^-}^-)$ is

the vector of parameters.¹¹ Therefore, the likelihood function can be written as

$$L(\tilde{\alpha}_1, \dots, \tilde{\alpha}_I; \tilde{\sigma}_1, \dots, \tilde{\sigma}_I, \Theta) = \prod_{i=1}^I G(\tilde{\alpha}_i; \tilde{\sigma}_i, \Theta). \quad (10)$$

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

2.2 The distribution of true alphas

We fit several distributions of this exponential family to the StockTwits data and find the results fit better than those of Gaussian mixture models. The best fit comes from a model with two exponential distributions for each of influencer types 1 and 3. The next section presents the results for this distribution. For the main results in this paper, we assume $K^+ = K^- = 2$. In the Appendix, we show 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% 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 estimated alphas 100 times, running our MLE estimation on each bootstrapped sample, and calculating the standard error of estimated parameters. Standard errors are relatively 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 ($t=5.51$).

¹¹Let 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:

$$h(x; \mu, \sigma) = \frac{1}{2\mu} \exp\left(\frac{\sigma^2}{2\mu^2} - \frac{x}{\mu}\right) \left(1 - \operatorname{erf}\left(\frac{\sigma}{2\mu} - \frac{x}{2\sigma}\right)\right),$$

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 (α_i)


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 three-factor model. The estimated alpha ($\hat{\alpha}$) 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.

	Mean alpha (%)	Fraction of users (%)
Skilled type 2	6.76 (0.49)	5.9 (0.8)
Skilled type 1	1.42 (0.14)	21.6 (1.2)
Unskilled	0.00 (0.00)	16.0 (2.9)
Antiskilled type 1	-1.06 (0.07)	45.6 (1.8)
Antiskilled type 2	-7.53 (0.29)	10.9 (0.7)
N		29,477
Log Likelihood		-86,385
AIC		172,786
BIC		172,806

Goodness of fit. Using the fitted distribution of true alphas, we perform the following steps to generate $N = 1,000$ samples of simulated $\tilde{\alpha}$'s, each with the same size as the original data (M).

1. Draw M observations from the fitted distribution of true alphas. Denote this vector by $a = [a_1, a_2, \dots, a_M]$.
2. Generate a sample of M standard errors by bootstrapping $[\tilde{\sigma}_1, \tilde{\sigma}_2, \dots, \tilde{\sigma}_M]$ with replacement. Denote this vector by $[s_1, s_2, \dots, s_M]$.
3. Generate a vector of estimation errors $e = [e_1, e_2, \dots, e_M]$ by drawing each e_i from a Normal distribution with a mean of zero and standard deviation of s_i .
4. Generate $[\tilde{a}_1, \tilde{a}_2, \dots, \tilde{a}_M]$ by adding a and e as in (6).
5. Calculate the vector of t -statistics $[t_1, t_2, \dots, t_M]$ through $t_i = \tilde{a}_i/s_i$.
6. Repeat steps one to five thousand times.

Panel A: Estimated and simulated alphas



Panel B: Estimated and simulated t -stats




Figure 1: Estimated and Simulated Alphas and Their t -Stats

In Panel A, the left plot shows histograms of estimated and simulated alphas. In Panel A, the right plot shows the average cdf of simulated alphas from the fitted model against estimated 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 estimated and simulated alphas.

After applying this procedure, we have $N = 1,000$ samples of estimated alphas and their standard errors, t -statistics, and the corresponding true alphas.

Figure 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 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 estimated 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 19.20%/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 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 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 1. We conclude that the fit with $K^+ = K^- = 2$ of the model alphas to the estimated alphas is very good.

2.3 Measures of finfluencer skill

An interpretation of the mixture modeling methodology is that it aggregates information to improve the signal-to-noise ratio of the data. Using the estimated distribution of true alphas, we can define measures of finfluencer skill, for example, the probability that a user is skilled, in addition to the user's expected alpha. We can then analyze the distribution and determinants of skill.

Using estimates from the mixture modeling methodology, we define four alternative measures of skill. For each user i , the probability of being skilled/antiskilled can be calculated as

$$\begin{aligned} \text{Pr}(\text{user } i \text{ skilled}) \quad \text{Pr}(\alpha_i > 0 \mid \tilde{\alpha}_i) &= \frac{\sum_{k=1}^{K^+} \pi_k^+ \eta(\alpha_i; \mu_k^+, \sigma_i)}{f_i(\alpha_i)}, \\ \text{Pr}(\text{user } i \text{ antiskilled}) \quad \text{Pr}(\alpha_i < 0 \mid \tilde{\alpha}_i) &= \frac{\sum_{k=1}^K \pi_k \eta(\alpha_i; \mu_k, \sigma_i)}{f_i(\alpha_i)}, \end{aligned} \tag{11}$$

where $\eta(\tilde{\alpha}_i; \mu, \tilde{\sigma}_i)$ is the convolution of a normal with mean zero and standard deviation of $\tilde{\sigma}_i$ and an exponential with a mean of μ evaluated at $\tilde{\alpha}_i$. In the denominator of (11), f_i is the distribution of $\tilde{\alpha}_i$.

We define the probability of being unskilled, $\text{Pr}(\alpha_i = 0 \mid \tilde{\alpha}_i) = 1 - \text{Pr}(\alpha_i > 0 \mid \tilde{\alpha}_i) - \text{Pr}(\alpha_i < 0 \mid \tilde{\alpha}_i)$, by subtracting the probabilities of being skilled and antiskilled from one.

The expected value of true skill α for any user i conditional on the measured skill $\tilde{\alpha}$ can be

Table 3: Distribution of Finfluencer Skill

This table reports descriptive statistics on alternative measures of influencer skill. The probability of being skilled, $\Pr(\alpha_i > 0 | \alpha_i)$, is defined in (11). The probability of being unskilled, $\Pr(\alpha_i = 0 | \alpha_i)$, and the probability of being antiskilled, $\Pr(\alpha_i < 0 | \alpha_i)$, are defined accordingly. The expected value of true alpha is defined in (12). FMM stands for Finite Mixture Models procedure.

	Skilled users $\Pr(\alpha_i > 0 \alpha_i)$	Unskilled users $\Pr(\alpha_i = 0 \alpha_i)$	Antiskilled users $\Pr(\alpha_i < 0 \alpha_i)$	True alpha $E[\alpha_i \alpha_i]$
Panel A: Distribution of $\Pr(\alpha_i \neq 0 \alpha_i)$ and $E[\alpha_i \alpha_i]$				
Mean	0.28	0.16	0.56	-0.57
S.D.	0.22	0.07	0.23	3.55
P10	0.04	0.03	0.26	-2.05
P25	0.13	0.14	0.45	-0.89
P50	0.24	0.17	0.57	-0.32
P75	0.34	0.20	0.69	0.15
P90	0.55	0.23	0.88	0.97
Panel B: Alternative classifications into skilled, unskilled, antiskilled influencers				
Classification based on FMM	0.28	0.16	0.56	
Classification based on $\Pr > 1/3$	0.26	0.01	0.86	
Classification based on max. Pr	0.18	0.01	0.81	
N	29,477	29,477	29,477	29,477

written as

$$E[\alpha_i | \tilde{\alpha}_i] = \frac{1}{f_i(\tilde{\alpha}_i)} \left(\int_{-1}^0 \alpha \phi(\tilde{\alpha}_i; \alpha, \tilde{\sigma}_i) \left(\sum_{k=1}^K \pi_k g(\alpha; \mu_k) \right) d\alpha + \int_0^1 \alpha \phi(\tilde{\alpha}_i; \alpha, \tilde{\sigma}_i) \left(\sum_{k=1}^{K^+} \pi_k^+ g(\alpha; \mu_k^+) \right) d\alpha \right), \quad (12)$$

where $\phi(\tilde{\alpha}_i; \alpha, \tilde{\sigma}_i)$ is a normal with a mean of α and standard deviation of $\tilde{\sigma}_i$.

Table 3 documents the descriptive statistics for the estimated user skill categories. The average probability that a user on StockTwits is skilled/unskilled/antiskilled is 28%/16%/56% with a standard deviation equal to 22%/7%/23%. The left subplot of Figure 2 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 first column of 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 2 shows.

Table 3 indicates that the distribution of $\Pr(\alpha_i = 0 \mid \tilde{\alpha}_i)$ is tight, and the left subplot of Figure 2 confirms this observation. The second column of Table 3 shows that the majority of StockTwits users have a low probability of being unskilled, as 99% of them have a probability of less than 1/3 of being unskilled. Column 3 of Table 3 shows that the vast majority of StockTwits users can be classified as antiskilled, as 86% of them have a probability of more than 1/3 of being antiskilled. Similarly, the left subplot of Figure 2 shows that the majority of users have a probability in excess of 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 18% of finfluencers as being skilled, 1% of finfluencers as being unskilled, and 81% of finfluencers as being antiskilled.

The last column of Table 3 demonstrates that the average monthly true alpha, $E[\alpha_i \mid \tilde{\alpha}_i]$, among finfluencers is equal to -57bps with a standard deviation of 3.55%, 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 -2.05% or less per month, while the top 10% of users generate alpha of 0.97% or more per month. Consequently, the right subplot of Figure 2 shows the distribution of true alphas among skilled and, respectively, antiskilled finfluencers (classified using the 1/3 rule). Most skilled influencers 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%.

Overall, our results indicate that most StockTwits users are antiskilled. This is quite important since the content of the antiskilled users' tweets is informative in the sense of "do the opposite of what I say." Correspondingly, this finding explains why Giannini, Irvine, and Shu (2018) find a negative correlation between average StockTwits sentiment and future stock returns. Looking at the average sentiment hides, however, the fact that some finfluencers are informed on StockTwits. While Harvey and Liu (2018) find a similar result for mutual fund managers, it is in contrast with findings for other crowdsourcing platforms for predicting stock returns. Other papers have documented the informativeness of platforms such as ValueInvestorsClub.com, Estimote, and

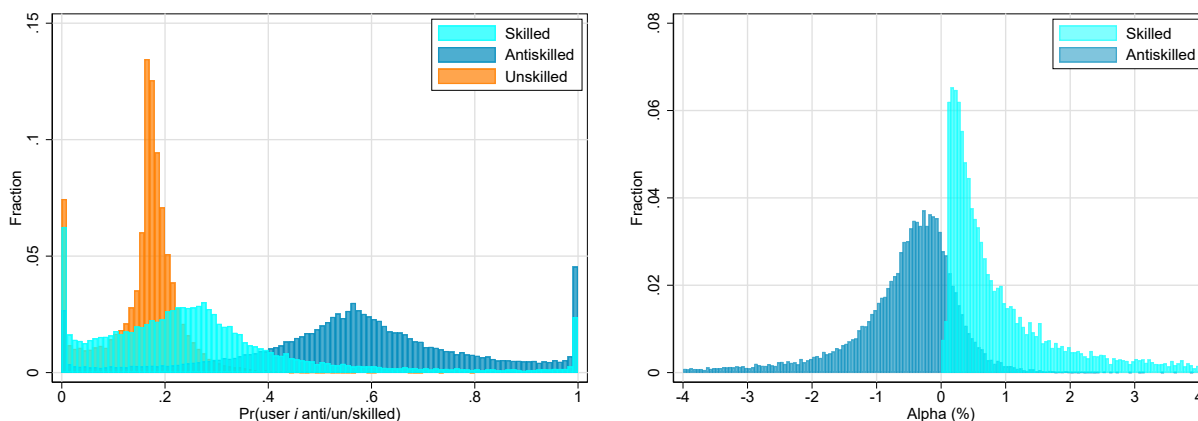


Figure 2: Distribution in Users’ Probability of Being Un/Anti/Skilled and True Alphas
 The plots show histograms of the probabilities of users being skilled, unskilled, and antiskilled, respectively, and the expected value of true skill, $E[\alpha_i | \alpha_i]$.

SumZero.com (Crawford, Gray, and Kern, 2017; Crawford, Gray, Johnson, and Price III, 2018; Jame, Johnston, Markov, and Wolfe, 2016). Chen, De, Hu, and Hwang (2014) find that the sentiment of the Seeking Alpha articles positively correlates with future stock returns. By contrast, Goutte (2020) finds that StockTwits users outperform Seeking Alpha users.

Contrasting our finding that 28% of StockTwits users are skilled with Dim (2022) who finds that approximately 56% of users on Seeking Alpha can predict stock returns correctly suggests that platforms with more curated users have more informative content. Consequently, the average sentiment of Seeking Alpha correlates positively with future returns, in contrast to what the literature has found for StockTwits.

In the rest of the paper, we will use these four skill measures to study which user characteristics explain influencers’ behavior and predict influencers’ skills.

3 Influencer Popularity

This section documents the relation between skill and popularity of influencers, which is important for assessing the quality of financial advice via social media platforms and the nature of competition among influencers. Do social media users follow more the skilled than the unskilled and antiskilled influencers? If so, we would expect the market mechanism to weed out unskilled and antiskilled

finfluencers over time and render the market for financial advice more efficient. Or, alternatively, are social media users more likely to follow finfluencers for reasons unrelated to their performance, such as behavioral traits and homophily? If so, we would expect unskilled and antiskilled finfluencers to survive and even grow in importance over time. In short, are more or less skilled finfluencers more likely to attract a large follower base? Finally, do retail investors adhere to the advice of finfluencers they follow, and which types of finfluencers have larger impact on retail order imbalances?

3.1 Do more skilled users have a larger follower base?

Given our split of finfluencers into skilled, unskilled, or antiskilled, we start by asking whether the crowd of StockTwits users can identify the skilled ones. If so, we would expect skilled users to have more followers than unskilled users, at least over the long term. An alternative hypothesis is that social media users like to follow finfluencers for reasons unrelated to their performance, such as behavioral traits and homophily (Currarini, Jackson, and Pin, 2009; Golub and Jackson, 2012). In this case, we may expect the opposite in that finfluencers with skill may have fewer followers than unskilled or antiskilled ones, while finfluencers with more followers are more likely unskilled or antiskilled. 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 may expect an ambiguous relation between skill and popularity (Benabou and Laroque, 1992).

To disentangle these three hypotheses, Figure 3 documents the univariate relation between the number of followers and our measures of users' skills. 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.

We next study more formally the effect of the market performance metric, alpha, on the finfluencer's follower count. In this analysis, our measurement of skill is based on each finfluencer's tweets in the time period 2013-2016, while the number of followers is measured as of February 2018.




Figure 3: Binscatter Plots of Skill versus Number of Followers

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

The lag of more than one year between alpha measurement and follower count should reduce any concern about reverse causality. To capture the effect of skill on popularity, we regress the number of followers on our measures of finfluencers' skill:

$$\text{Finfluencer's follower count}_i \text{ (measured out-of-sample)} = \alpha + \beta \text{ Skill}_i + \epsilon_i, \quad (13)$$

where the dependent variable is the log of one plus the finfluencer's follower count as of February 2018, and Skill_i are our skill measures (11) and (12). For comparison, we include a specification with the user-specific abnormal returns $\tilde{\alpha}_i$ measured by (4). Across specifications, the explanatory variables are the finfluencer's measured alpha in the data, $\tilde{\alpha}_i$, the expected value of alpha given its measurement in the data, $E[\alpha_i \mid \tilde{\alpha}_i]$, the probability that a user is skilled, $\Pr(\alpha_i > 0 \mid \tilde{\alpha}_i)$, the probability that a user is unskilled, $\Pr(\alpha_i = 0 \mid \tilde{\alpha}_i)$, or the probability that a user is antiskilled, $\Pr(\alpha_i < 0 \mid \tilde{\alpha}_i)$.

Table 4 reports the results when explaining finfluencer popularity (measured by the number of followers) by our measures of skill. The estimates show that neither finfluencers' measured alpha, $\tilde{\alpha}_i$, nor finfluencers' expected alpha given its measurement, $E[\alpha_i \mid \tilde{\alpha}_i]$, have a relationship with the follower count. By contrast, the probability that a user is skilled strongly negatively predicts popularity, with a coefficient $\beta = -0.80$ significant at 1%. Similarly, the probability that a user

Table 4: The Effect of Finfluencers' Alpha on Follower Count

This table reports the results of regressing the number of followers on influencers' measure of skill. The dependent variable is the log of one plus the influencer's follower count as of February 2018. The independent variables are: α_i is the influencer's measured alpha in the data, $E[\alpha_i | \alpha_i]$ is the expected value of alpha given its measurement in the data, $\Pr(\alpha_i > 0 | \alpha_i)$ is the probability that a user is skilled, $\Pr(\alpha_i = 0 | \alpha_i)$ is the probability that a user is unskilled, and $\Pr(\alpha_i < 0 | \alpha_i)$ 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.

	Influencer's follower count _i (measured out-of-sample)				
	(1)	(2)	(3)	(4)	(5)
α_i	0.00 (0.00)				
$E[\alpha_i \alpha_i]$		0.00 (0.00)			
Pr(user <i>i</i> skilled)			-0.80*** (0.06)		
Pr(user <i>i</i> unskilled)				3.79*** (0.23)	
Pr(user <i>i</i> antiskilled)					0.34*** (0.06)
Constant	2.53*** (0.01)	2.70*** (0.01)	2.92*** (0.02)	2.09*** (0.04)	2.50*** (0.04)
r ²	0.000	0.000	0.008	0.020	0.002
N	27,200	22,074	22,074	22,074	22,074

is unskilled ($\beta = 3.79$) and the probability that a user is antiskilled ($\beta = 0.34$) strongly positively predict popularity, all significant at 1%. This means skilled influencers have fewer followers than unskilled or antiskilled influencers.

These puzzling findings create the need to understand the economic forces behind the negative relation between the number of followers and skill measures. To narrow down the channel for why certain influencers are more popular than others, we next check if skill is persistent and if it affects influencer survival. The channels that this analysis helps to distinguish are whether social media users cannot correctly identify influencers' skills because skill is not long-lasting potentially due to reputation exploitation, whether they cannot correctly identify influencers' skills because tweeting patterns do not correlate with easily detectable determinants of skill, or whether they do not care about influencers' skills since they match with influencers based on other criteria such as their own behavioral traits and homophily.

Table 5: Persistence of Finfluencer Skill

The table reports the persistence of influencers' skill. The specification regresses $Skill_{i,post-2016}$ measured post-2016 on $Skill_{i,pre-2016}$ measured pre-2016. $Skill_i$ is one of the following five variables: the estimated alpha, α_i , the expected value of true alpha, $E[\alpha_i]$, and the probability of α_i being positive, zero, or negative. For each regression, the (in)dependent variable is calculated with tweets posted before 2016 and, respectively, in or after 2016 which falls in the middle of our sample period. Standard errors are robust to heteroskedasticity. * , ** , and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Finfluencer's skill (measured post-2016)				
	(1) α_i	(2) $E[\alpha_i \alpha_i]$	(3) $\Pr(\alpha_i > 0 \alpha_i)$	(4) $\Pr(\alpha_i = 0 \alpha_i)$	(5) $\Pr(\alpha_i < 0 \alpha_i)$
$Skill_{i,pre-2016}$	0.00 (0.02)	0.09 (0.03)	0.03 (0.01)	0.09 (0.01)	0.03 (0.01)
Constant	-0.29 (0.08)	-0.29 (0.03)	31.90 (0.47)	12.38 (0.22)	51.85 (0.80)
N	9,382	6,449	6,449	6,449	6,449

3.2 Skill persistence and influencer survival

To understand the economic forces behind the negative relation between popularity and skill, important questions that we now address are whether influencers' skills are persistent and how this affects influencer survival.¹² If influencers' skills are persistent, then social media users may not prioritize skill over other characteristics when deciding which influencers to follow. If they are not persistent, then it may not be surprising to observe a negative correlation between measures of influencers' skill and their follower count.

Persistence in (anti)skill. To address the question of skill persistence, we divide our sample into pre- and post-2016 periods and calculate each user's estimated alpha in each sub-sample separately. We then reestimate our MLE model separately on the pre-2016 and post-2016 data sub-samples and calculate the expected true alpha and the probability of each user being skilled, unskilled, or antiskilled. We choose 2016 because it falls in the middle of our sample period. We then have five variables describing each influencer's skill estimated over data subsamples, pre-2016 and post-2016 including 2016. To test the persistence of influencers' skills, we regress the estimates obtained

¹²The prior literature has studied this question in the context of professional analysts (Crane and Crotty, 2020), but not for non-professional or semi-professional influencers.

using the post-2016 data sample on the estimates obtained using the pre-2016 data sample:

$$\text{Skill}_{i,\text{post-2016}} = \alpha + \beta \text{Skill}_{i,\text{pre-2016}} + \epsilon_i, \quad (14)$$

where Skill_i is one of the following five variables: the estimated alpha, $\tilde{\alpha}_i$, the expected value of true alpha, $E[\alpha_i]$, and the probability of α_i being positive, zero, or negative. A statistically significant AR1 coefficient β would imply that influencers’ skills are persistent.

Table 5 reports the results of the persistence regressions. The number of observations in Table 5 is lower than in Table 4 because we now require data for both sub-samples. The drop in the number of observations suggests significant entry and exit in the market for influencers. The autoregressive coefficient β is small and insignificant when we measure influencers’ skill using the estimated alphas, $\tilde{\alpha}$. By contrast, measures of influencers’ skill derived from the MLE estimation show significant persistence. A 1% increase in the expected true alpha, $E[\alpha_i]$, in the pre-2016 data, is associated with a 9 bps increase in the expected true alpha in the post-2016 data. Similarly, a one percent increase in the probability of positive/negative alpha over the pre-2016 data is associated with a 0.03% increase in the same probability over the post-2016 data.

The finding that persistence is absent in measured alphas, but is present among expected true alphas is interesting. Not only is the autoregressive coefficient of the estimated alphas insignificant, but it is also much smaller. This observation suggests a sizeable error-in-variables (EIV) bias in measured alphas. Because the estimated alpha is a noisy measure of the true alpha, the magnitude of the autoregressive coefficient shrinks toward zero. The MLE estimation partially removes the estimation noise, thereby decreasing the EIV bias and increasing the magnitude of the autoregressive coefficient.

Influencer survival. Given the finding that our measures of influencers’ skill are persistent, we now check if skilled influencers are more likely to stay active, that is, “survive” despite the fact that they have fewer followers than unskilled and antiskilled influencers. 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 $\text{Influencer survival}_i$ is an indicator

Table 6: Finfluencer Survival

This table reports the determinants of influencers' 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 influencer is active in or after 2016, and zero otherwise. The independent variables are $\alpha_{i,\text{pre-2016}}$ is the influencer's measured alpha in the data before 2016, $E[\alpha_i | \alpha_{i,\text{pre-2016}}]$ is the expected value of alpha given its measurement in the data before 2016, $\Pr(\alpha_i > 0 | \alpha_{i,\text{pre-2016}})$ is the probability that a user is skilled, $\Pr(\alpha_i = 0 | \alpha_{i,\text{pre-2016}})$ is the probability that a user is unskilled, and $\Pr(\alpha_i < 0 | \alpha_{i,\text{pre-2016}})$ 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)	(5)
$\alpha_{i,\text{pre-2016}}$	0.00 (0.00)				
$E[\alpha_i \alpha_{i,\text{pre-2016}}]$		0.02*** (0.00)			
$\Pr(\text{user } i \text{ skilled } \alpha_{i,\text{pre-2016}})$			-0.08 (0.04)		
$\Pr(\text{user } i \text{ unskilled } \alpha_{i,\text{pre-2016}})$				1.47*** (0.13)	
$\Pr(\text{user } i \text{ antiskilled } \alpha_{i,\text{pre-2016}})$					-0.06 (0.04)
Constant	-0.24*** (0.01)	-0.18*** (0.01)	-0.17*** (0.01)	-0.40*** (0.02)	-0.16*** (0.02)
r ²	0.000	0.001	0.000	0.005	0.000
N	23,103	18,770	18,770	18,770	18,770

function equal to one if the finfluencer is active in or after 2016, and zero otherwise:

$$\text{Finfluencer survival}_i = \Phi(\alpha + \beta \text{ Skill}_{i,\text{pre-2016}}), \tag{15}$$

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, $E[\alpha_i | \tilde{\alpha}_{i,\text{pre-2016}}]$ is the expected value of alpha given its measurement in the data before 2016, $\Pr(\alpha_i > 0 | \tilde{\alpha}_{i,\text{pre-2016}})$ is the probability that a user is skilled, $\Pr(\alpha_i = 0 | \tilde{\alpha}_{i,\text{pre-2016}})$ is the probability that a user is unskilled, and $\Pr(\alpha_i < 0 | \tilde{\alpha}_{i,\text{pre-2016}})$ is the probability that a user is antiskilled.

The results in Table 6 show that skill is not a significant determinant of survival. First, the finfluencer's measured alpha, $\tilde{\alpha}_{i,\text{pre-2016}}$, is an insignificant determinant of survival. The expected value of alpha given its measurement in the data before 2016 eliminates some noise and indeed

positively correlates with survival. However, the economic magnitude is small. When we split skill into three types based on their probabilities in columns (3)-(5), only the probability of being unskilled statistically significantly predicts survival and the relation is positive. Column (3) shows that the probability of being skilled has no impact on the probability of survival.

Overall, the results so far suggest that finfluencer skill is persistent but despite this fact, skilled finfluencers are not more likely to “survive,” that is, stay active than unskilled and antiskilled finfluencers. Next, we investigate whether finfluencers and which type(s) have an economic impact by affecting retail trading.

3.3 Which finfluencers affect retail investor behavior?

A way for finfluencers to matter and have an economic impact is to affect retail trading in the direction of their tweets. Retail investors may be influenced differently by skilled vs. anti/unskilled finfluencers and the relation with the size of a finfluencer’s follower base is a priori not clear. While it may be that unskilled and antiskilled finfluencers have more followers, social media users may not necessarily invest based on their flawed advice.

To address these questions, we test the relationship between different types of StockTwits users and the behavior of retail investors using lead-lag regressions. We split StockTwits users into antiskilled, unskilled, and skilled based on their respective probability given by (11). Our main variables of interest capture the sentiment of the tweets of different types of influencers. We split each finfluencer’s tweeting activity into the number of positive and, respectively, negative tweets in a given stock on a given day and then compute the average number of tweets weighted by each finfluencer’s probability of being of one of the three types:

$$\begin{aligned}
 \text{Positive sentiment by anti/un/skilled users}_{j,t} &= \frac{1}{I} \sum_{\text{all } i} \Pr(\text{user } i \text{ anti/un/skilled}) \text{ SocSent}_{i,j,t}^+, \\
 \text{Negative sentiment by anti/un/skilled users}_{j,t} &= \frac{1}{I} \sum_{\text{all } i} \Pr(\text{user } i \text{ anti/un/skilled}) \text{ SocSent}_{i,j,t}, \\
 & \hspace{15em} (16)
 \end{aligned}$$

where $\Pr(\text{user } i \text{ anti/un/skilled})$ are given by (11), $\text{SocSent}_{i,j,t}^+ = \sum_{n=1}^{N_{i,j,t}} 1(\text{SocSent}_{i,j,t,n} = 1)$

counts the positive tweets by finfluencer i in stock j on the day t , and

$$SocSent_{i,j,t} = \sum_{n=1}^{N_{i,j,t}} 1(SocSent_{i,j,t,n} = 1)$$

counts the negative ones.

To capture the impact of finfluencers on retail traders, we estimate the following lead-lag panel regressions with stock and day fixed effects:

$$\begin{aligned} \text{Retail order imbalance}_{j,t+1} &= \alpha_j + \alpha_t + \\ &+ \sum_{f \in \{2fa, u, sg\}} \beta_f^+ \text{ Positive sentiment by anti/un/skilled users}_{j,t} + \\ &+ \sum_{f \in \{2fa, u, sg\}} \beta_f \text{ Negative sentiment by anti/un/skilled users}_{j,t} + \gamma^\theta \mathbf{X}_{j,t} + \epsilon_{j,t+1}, \end{aligned} \quad (17)$$

where the set fa, u, sg refers to anti/un/skilled StockTwits users and controls $\mathbf{X}_{j,t}$ include average positive news sentiment in stock j on the day t and corresponding negative news sentiment, trading volume, retail order imbalances, and short-sales constraint index. The dependent variable, Retail order imbalance $_{j,t}$, is the retail order imbalance equal to a difference between the number of retail buy and sell orders in stock j on the day t .

Table 7 documents the results of specification (17). The coefficient estimates suggest that some types of StockTwits users impact retail trading, while others do not. The positive sentiment of antiskilled StockTwits users strongly predicts an increase in retail order imbalances on the next trading day, controlling for news sentiment, trading volume, past retail order imbalances, and stock-level short-sales constraints. In other words, positive tweeting activity by antiskilled finfluencers leads to an increase in retail buys relative to retail sales on the next trading day. Negative sentiment by antiskilled users strongly predicts a reduction in retail order imbalances on the next trading day. That is, negative tweeting activity by antiskilled finfluencers leads to a reduction in retail buys relative to retail sales on the next trading day. In terms of economic magnitudes, the results for positive and negative tweets are roughly symmetric. By contrast, neither positive nor negative tweeting by skilled users has a significant impact on retail order imbalances.

Table 7: Finfluencer Sentiment and Retail Order Imbalances

This table reports the determinants of retail order imbalances. Results are obtained from panel regressions with stock and day fixed effects. The independent variables of interest capture the tweet sentiment by different influencer types in stock j on the day t , which we compute by splitting the tweeting activity by each user into the number of positive and, respectively, negative tweets in a given stock on a given day and then compute the average number of tweets weighted by each user's probability of being of one of the three types. Standard errors are robust to clustering at the stock and day level. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Retail order imbalance $_{j,t+1}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Positive sentiment by antiskilled users $_{j,t}$	8.72*** (1.57)	6.99*** (1.48)			8.40*** (1.55)	6.68*** (1.46)
Positive sentiment by unskilled users $_{j,t}$	7.55 (4.07)	7.45 (3.81)			7.49 (4.04)	7.40 (3.78)
Positive sentiment by skilled users $_{j,t}$	2.54 (1.57)	1.72 (1.53)			2.19 (1.58)	1.41 (1.54)
Negative sentiment by antiskilled users $_{j,t}$			-8.49*** (2.50)	-7.31** (2.49)	-7.59** (2.50)	-6.53** (2.48)
Negative sentiment by unskilled users $_{j,t}$			-17.02* (7.82)	-18.58* (7.91)	-16.11* (7.83)	-17.74* (7.92)
Negative sentiment by skilled users $_{j,t}$			0.95 (2.56)	1.92 (2.56)	2.47 (2.59)	3.15 (2.59)
Positive news sentiment $_{j,t}$		-0.05 (0.06)		-0.04 (0.06)		-0.05 (0.06)
Negative news sentiment $_{j,t}$		-0.36 (0.21)		-0.36 (0.20)		-0.35 (0.20)
Volume $_{j,t}$		0.00 (0.00)		0.00 (0.00)		0.00 (0.00)
Retail order imbalance $_{j,t}$		0.07*** (0.00)		0.07*** (0.00)		0.07*** (0.00)
Short sale constraint $_{j,t}$		-3.80*** (1.13)		-3.80*** (1.13)		-3.79*** (1.13)
Stock & Day FEs	Yes	Yes	Yes	Yes	Yes	Yes
r ²	0.017	0.024	0.017	0.024	0.017	0.024
N	875,211	795,956	875,211	795,956	875,211	795,956

When combined with the results from Table 4 on influencer's popularity, these findings suggest that StockTwits users treat antiskilled influencers as "gurus", that is, they follow them, listen to their investment advice, and then act on it by trading in the advised direction. This behavior by StockTwits users raises an important question of why they ignore the experts or most skilled influencers in favor of gurus or antiskilled influencers? The next section addresses this question.

4 Influencer Skill and the Social Network

This section tests if social media users can tell apart skilled from unskilled and antiskilled influencers. To better understand why social media users do not follow the most skilled influencers, we check whether social media users can use observable characteristics to tell apart value-creating experts whom we associate with skilled influencers, from charlatans whom we associate with unskilled influencers, and gurus whom we associate with antiskilled influencers, and whether they do so. If influencers follow commonly known strategies with their tweets and these tweeting strategies correlate with users' skills, social media users can in principle use these characteristics to separate good from bad advice.

4.1 Dissecting influencers' tweeting strategies

We start by dissecting influencers' tweeting strategies depending on their skill. Doing this helps to understand the nature of information or skills held by influencers and what determines the positive or negative performance of different influencers. We use the measures of influencer skill from Section 2 to study whether influencers follow commonly known investment behaviors.

The measures of influencers' 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 influencers, are user-level characteristics such as the number of tweets, their tone, and the number of followers and likes. If influencers can be categorized by these characteristics or they use these observable characteristics to signal their type to other StockTwits users, then these characteristics should be informative about influencers' skills. A StockTwits user can control the first two characteristics and does not have full control over its follower base, but can create user attention based on tweeting activity.

User attention based on tweeting activity. StockTwits users are heterogeneous in their tweeting activity. It seems reasonable to expect that this heterogeneity affects the informativeness of their tweets. For example, one may think that users who tweet more often are more likely to be experts and have more valuable information. On the other hand, users who tweet more often




Figure 4: Binscatter Plots of Skill versus Tweeting Activity

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

are also more likely to be overconfident 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. Furthermore, the prior literature has documented that short sellers are informed (e.g., Engelberg, Reed, and Ringgenberg, 2012; Boehmer, Jones, and Zhang, 2008). Therefore, we might expect that users with more negative tweets are more informed.

To test these hypotheses, we relate our measures of skill to the number of tweets and the composition of the tweets, in particular, the fraction of tweets with a negative tone. Figure 4 documents the univariate relation between users’ 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 users’ tweeting activity and the probability of being antiskilled. By contrast, tweeting activity is very strongly negatively related to users’ 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 users or, in other words, that the tweeting activity exhibits decreasing identification power for skilled users.

Table 8 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 composition and the tweeting strategies of different influencers. The table reports the results of several sets of regressions of the form:

$$\text{Skill}_i = \alpha + \beta \quad \text{Tweeting Activity}_i / \beta \quad \text{TweetingStrategy}_i + \epsilon_i. \quad (18)$$

where Skill_i represents one of the following variables: (1) the estimated alpha ($\tilde{\alpha}$), (2) the expected value of true alpha ($E[\alpha]$), (3) the probability of α being positive, (4) the probability of α being negative. Across the different panels, we consider several popular tweeting strategies described below.

Panel A of Table 8 presents results for tweeting activity, *NumberTweets*, as well as for the fraction of negative tweets, *FractionNegative*, as explanatory variables of StockTwits users' skills. The composition of tweets, *FractionNegative*, is defined as the percentage of a influencer's non-neutral tweets that have a negative sentiment. The estimates show that an increase in tweeting activity does not have an economic effect on the measured alpha. A 10-times increase in the number of tweets increases the measured alpha by 8 bps per month. The point estimate is also not statistically significant. On the other hand, the expected true alpha also increases by 8 bps per month, and the point estimate is statistically significant at 1%. In agreement with the univariate results from Figure 4, the probability of being skilled decreases by 3.70% while the probability of being antiskilled increases by 1.26% when the number of tweets increases tenfold. Put together, users who tweet more frequently are less likely to be skilled, consistent with informed users tweeting less frequently. However, conditional on being skilled or antiskilled, the expected value of the frequent tweeter's skills is larger, implying that frequent tweeters have more experience in picking stocks.

Panel A also includes the estimates for the percentage of a influencer's non-neutral tweets that have a negative sentiment, *FractionNegative*, used as the explanatory variable. Consistent with the prior literature, we find that users with more negative tweets are more likely to be informed across all skill measures. A one-percent increase in the share of negative tweets is associated with a 3 bps increase in the monthly estimated alpha. The expected true alpha also increases by 1 bps

per month. The probability of being informed increases by 0.06%, while the probability of being antiskilled decreases by 0.09%. All of these estimates are significant at 1% and point to the same conclusion: StockTwits users with more negative tweets are more likely to post informative tweets.

Return chasing vs. contrarian behavior. The prior literature documents return chasing among retail traders (Barber and Odean, 2007). In our setup, we can ask if the tweets by all or some group(s) of users are motivated by return chasing. In particular, if antiskilled finfluencers' tweets chase returns, return chasing may contribute to these users' measured negative skill.

We measure each user's return-chasing tendency by the percentage of her 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 checks. We first regress measured and expected alphas on return chasing to test if return chasing is associated with better or worse performance.

Panel B of Table 8 reports the results of the return chasing tests. We find that a one percent increase in return chasing is associated with a 7 bps decrease in the estimated alpha, while the expected true alpha decreases by 4 bps. The probability of being skilled or antiskilled also changes with the tendency to chase returns. A one percent increase in return chasing tendency is associated with an 0.08% decrease in the probability of being skilled and a 0.16% increase in the probability of being antiskilled. Because the skilled, unskilled, and antiskilled components sum up to one, the probability of being unskilled also decreases by 0.08%. Overall, return chasing contributes to users being antiskilled.

Panel C of Table 8 reports results from regressing our measures of skill on contrarian tendency. It could be that skilled users follow a contrarian approach given that return chasing contributes to negative skill. We measure each user's contrarian tendency as the percentage of a user's 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. The results in Panel C show no significant association between contrarian tweeting and skill. In other words, users who post contrarian tweets do not exhibit higher skills.

Table 8: Dissecting Finfluencers' Tweeting Strategies

The table reports the results of several sets of regressions of the form:

$$\text{Skill}_i = \alpha + \beta \text{ Tweeting Activity}_i / \beta \text{ Tweeting Strategy}_i + \epsilon_i. \quad (19)$$

Skill_i represents one of the following variables: (1) the estimated alpha (α) (2) the expected value of true alpha ($E[\alpha]$) (3) the probability of α being positive (4) the probability of α being negative. The estimated alpha (α) for each user is the average of signed adjusted returns after her tweets. The other dependent variables are defined in expressions (11) and (12). All dependent variables are in percentage points. Tweeting activity is represented by *NumberTweets* 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 *FractionNegative* defined as the percentage of a finfluencer's non-neutral tweets that have a negative 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. *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. *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. Standard errors are robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)
	Alpha		Skilled	Antiskilled
	α_i	$E[\alpha_i \alpha_i]$	$\Pr(\alpha_i > 0 \alpha_i)$	$\Pr(\alpha_i < 0 \alpha_i)$
Panel A: Relationship between Number/Composition of Tweets and Users' Skill				
<i>NumberTweets_i</i>	0.08 (0.06)	0.08 (0.03)	-3.70 (0.23)	1.26 (0.25)
<i>FractionNegative_i</i>	0.03 (0.00)	0.01 (0.00)	0.06 (0.01)	-0.09 (0.01)
Panel B: Which Finfluencers Pursue Return Chasing?				
<i>ReturnChasing_i</i>	-0.07 (0.01)	-0.04 (0.01)	-0.08 (0.03)	0.16 (0.03)
Panel C: Which Finfluencers Pursue Contrarian Tweeting?				
<i>ContrarianTweet_i</i>	0.01 (0.01)	-0.01 (0.01)	0.06 (0.03)	0.01 (0.04)
Panel D: Tweeting about Short-Selling Constrained Stocks				
<i>SSI_i (Positive Tweets)</i>	-0.52 (0.04)	-0.31 (0.02)	-0.89 (0.10)	1.63 (0.11)
<i>SSI_i (Negative Tweets)</i>	0.43 (0.07)	0.18 (0.04)	1.76 (0.18)	-1.49 (0.17)
Panel E: Effect of Positive Herding on Users' Skill				
<i>PositiveHerding_i</i>	-0.03 (0.00)	-0.02 (0.00)	-0.09 (0.01)	0.11 (0.01)
Panel F: Effect of Negative Herding on Users' Skill				
<i>NegativeHerding_i</i>	0.04 (0.00)	0.02 (0.00)	0.07 (0.02)	-0.13 (0.02)
N	29,475	29,475	29,475	29,475

Short-sale constraints and tweet sentiment. Asset pricing theory suggests that risky assets are overpriced in a market with short-sale constraints (Miller, 1977). As a result, stocks with short-selling constraints tend to be overpriced. We ask whether influencers exploit this underpricing in their tweets. Due to this overpricing, we expect skilled users to post more negative tweets about stocks with tighter short-selling constraints. We use 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, 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 regress our measures of skill on short-selling constraints to test whether skilled social media users can exploit the overpricing of stocks with short-selling constraints.

Panel D of Table 8 reports the results of these regressions. A one decile increase in the short-selling constraints of positively (negatively) tweeted stocks are associated with a 0.52% (0.43) per month decrease (increase) in the user’s estimated alpha. Using expected true alphas, the same increase in short-selling constraints for positively (negatively) tweeted stocks results in a 0.31% (0.18%) decrease (increase) in skill. The probability of being skilled or antiskilled also changes with short-selling constraints. A one-decile increase in the short-selling constraints of positively tweeted stocks is associated with a 0.89% (1.63%) decrease (increase) in the user’s probability of being skilled. On the other hand, a one-decile increase in the short-selling constraints of negatively tweeted stocks is associated with a 1.76% (1.49%) increase (decrease) in the user’s probability of being antiskilled. Overall, these results show that exploiting short-selling constraints correctly contributes to users’ skills on both the negative and positive sides.

Herding and tweeting. An interesting question is whether herding affects the informativeness of users’ tweets. To quantify herding, we calculate the percentage of each user’s positive/negative tweets that are about stocks in the highest decile of positive/negative tweeting activity over the

past five days. Next, we regress our measures of skills on users' positive, *PositiveHerding*, and negative, *NegativeHerding*, herding tendencies.

Panel E of Table 8 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 3bp (2bp) decrease in estimated alpha (expected true alpha). Moreover, the probability of being skilled decreases by 0.09% while the probability of being antiskilled increases by 0.11%. Taken together, the results in Table 8 show that positive herding tendency is negatively correlated with users' 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 could also measure herding around negative tweets. Thus we repeat our regressions using an alternative definition of the independent variable that measures herding on negative tweets.

Panel F Table 8 reports the results of regressing the skill measures on *NegativeHerding*. It shows that users who tweet more often about stocks in the top decile of negative tweeting activity are more likely to be skilled and less likely to be antiskilled. A one-percent increase in the negative herding measure is associated with a 0.07% increase (0.13% decrease) in the probability of being skilled (antiskilled). The estimated alpha and expected true alpha both increase with herding on negative tweets.

4.2 What finfluencer behaviors predict skill?

Table 8 has shown using uni/bivariate regressions that finfluencers follow some commonly known tweeting strategies. We now put these individual results together and ask if social media users can reasonably exploit finfluencer behavior to learn about their skills. If these skills can indeed be discerned by followers, then it is plausible that users on StockTwits are strategically selecting which finfluencers to follow based on their own behavioral traits or preferences and that they are not randomly or casually choosing who to follow. Instead, they are deliberately aligning themselves with finfluencers whose tweeting habits—and, by extension, whose skills and expertise—match their own financial goals, risk tolerance, or investing style.

We address the question of whether skill can be detected using a multivariate regression analysis

of the determinants of finfluencers' skill. Our dependent variables will again be the probability of being skilled, unskilled, and antiskilled:

$$\begin{aligned} \Pr(\alpha_i \in 0, j, \tilde{\alpha}_i) = & \alpha + \sum_{Event} \beta_{Event}^p \text{ Finfluencer}_i \text{ posts positive tweets after } Event + \\ & + \sum_{Event} \beta_{Event}^m \text{ Finfluencer}_i \text{ posts negative tweets after } Event + \gamma^T \mathbf{X}_i + \epsilon_i, \end{aligned} \quad (20)$$

where $Event \in \{Past \text{ returns, Social sentiment, News sentiment, Volatility, Retail order imbalance, Trading volume, Short-sale constraint}\}$ captures events that trigger finfluencers' tweeting activity and \mathbf{X}_i are characteristics of finfluencer i .

To detect finfluencer skill, we construct several variables that capture the tweeting behavior of different finfluencers. We proceed in two steps. In the first step, we construct stock-level events triggering tweets. The Appendix describes in detail the construction of the variables that we use to capture events triggering finfluencers' tweeting activity. We compute the event-based criteria for stock j on the day t by averaging over the past time window $[t - L - 1, t - 1]$. For the window length, we set $L = 20$. Alternatively, we have set $L = 1, 2, 5, 10$ and the results are unaffected. Denote the decile in which stock j falls on the day t according to any of the events by $Decile_{j,t}^{Event, L-1, t-1}$.

In the second step, we next link stock-level events to user-level events. We calculate for each finfluencer i the average decile of all stocks that i tweets positive (negative) about on a given day after the stock-level event $Event$ has occurred on the prior day. The user-level variable $Fin \text{ uencer posts positive tweets after } Event$ measures the average decile according to $Event$ of the stocks that finfluencer i tweets positive about, averaged across stocks and time:

$$\begin{aligned} \text{Finfluencer}_i \text{ posts positive tweets after } Event &= \frac{\sum_j \sum_t (Decile_{j,t}^{Event, L-1, t-1} \text{ if } SocSent_{i,j,t} > 0)}{\sum_j \sum_t \mathbb{1}(SocSent_{i,j,t} > 0)}, \\ \text{Finfluencer}_i \text{ posts negative tweets after } Event &= \frac{\sum_j \sum_t (Decile_{j,t}^{Event, L-1, t-1} \text{ if } SocSent_{i,j,t} < 0)}{\sum_j \sum_t \mathbb{1}(SocSent_{i,j,t} < 0)}. \end{aligned} \quad (21)$$

Table 9 summarizes the results for skilled finfluencers in columns 1 and 2, antiskilled finfluencers in columns 3 and 4, and unskilled finfluencers in columns 5 and 6. Across columns, we vary the specification. Standard errors are robust to heteroskedasticity.

Table 9: Detecting Finfluencer Skill

The table documents the determinants of predicting skilled, antiskilled, and unskilled influencers using multivariate regression analysis. Across columns, we vary the specification. The Appendix describes in detail the construction of the variables that we use to capture events triggering influencers' tweeting activity. Standard errors are robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Skilled $\Pr(\alpha_i > 0 \alpha_i)$		Antiskilled $\Pr(\alpha_i < 0 \alpha_i)$		Unskilled $\Pr(\alpha_i = 0 \alpha_i)$	
Fin influencer posts positive tweets after:						
Positive returns	-0.95*** (0.19)	-0.85*** (0.23)	0.66*** (0.20)	0.58* (0.23)	0.28*** (0.06)	0.27*** (0.07)
Positive social sentiment	-1.56*** (0.34)	-1.57*** (0.39)	1.79*** (0.35)	1.98*** (0.41)	-0.23* (0.11)	-0.41*** (0.12)
Negative social sentiment	1.59** (0.49)	1.58** (0.59)	-1.80*** (0.51)	-2.19*** (0.61)	0.21 (0.15)	0.61*** (0.19)
Positive news sentiment	-0.36 (0.40)	-0.36 (0.47)	-0.15 (0.41)	-0.11 (0.49)	0.51*** (0.12)	0.47** (0.15)
Negative news sentiment	1.29*** (0.38)	1.13* (0.45)	-1.70*** (0.39)	-1.62*** (0.47)	0.41*** (0.12)	0.49*** (0.14)
Volatility	0.06 (0.17)	-0.38 (0.21)	0.58** (0.18)	0.75*** (0.21)	-0.64*** (0.05)	-0.37*** (0.06)
Retail order imbalance	-0.38 (0.32)	-0.47 (0.39)	0.41 (0.34)	0.32 (0.41)	-0.03 (0.10)	0.16 (0.12)
Trading volume	-0.34 (0.42)	0.01 (0.50)	0.87* (0.44)	0.90 (0.52)	-0.53*** (0.14)	-0.91*** (0.16)
Short sale constraint	-0.33 (0.19)	-0.40 (0.23)	0.55** (0.20)	0.74** (0.24)	-0.22*** (0.06)	-0.34*** (0.07)
Fin influencer posts negative tweets after:						
Positive returns	1.69*** (0.28)	1.66*** (0.33)	-1.79*** (0.27)	-1.89*** (0.32)	0.10 (0.08)	0.23* (0.10)
Positive social sentiment	2.65*** (0.58)	2.90*** (0.68)	-3.25*** (0.59)	-3.66*** (0.69)	0.60** (0.18)	0.76*** (0.21)
Negative social sentiment	-3.10*** (0.79)	-3.08*** (0.90)	2.82*** (0.78)	3.00*** (0.91)	0.28 (0.24)	0.09 (0.27)
Positive news sentiment	-1.30* (0.59)	-1.54* (0.70)	0.90 (0.59)	1.10 (0.69)	0.40* (0.18)	0.44* (0.20)
Negative news sentiment	-0.47 (0.60)	-0.75 (0.70)	0.22 (0.60)	0.29 (0.70)	0.25 (0.17)	0.46* (0.20)
Volatility	0.22 (0.23)	-0.16 (0.27)	-0.03 (0.23)	0.15 (0.27)	-0.19** (0.07)	0.02 (0.08)
Retail order imbalance	0.31 (0.46)	0.11 (0.55)	-0.45 (0.46)	-0.27 (0.55)	0.14 (0.14)	0.16 (0.16)
Trading volume	1.12 (0.60)	1.78** (0.68)	-0.23 (0.59)	-0.57 (0.69)	-0.89*** (0.18)	-1.21*** (0.21)
Short sale constraint	0.89** (0.29)	0.97** (0.34)	-0.61* (0.29)	-0.70* (0.33)	-0.29*** (0.08)	-0.27** (0.09)
User activity		-1.88*** (0.16)		0.48** (0.17)		1.40*** (0.07)
No. of ideas		-0.28* (0.11)		0.17 (0.12)		0.11** (0.04)
No. of likes		0.18 (0.09)		-0.03 (0.10)		-0.14*** (0.03)
Watchlist size		0.06 (0.11)		-0.04 (0.12)		-0.02 (0.04)
No. of users followed		-0.22 (0.11)		0.14 (0.12)		0.08* (0.04)
r2	0.015	0.028	0.022	0.024	0.048	0.102
N	29,395	22,014	29,395	22,014	29,395	22,014

Detecting skilled finfluencers. Table 9, columns 1 and 2 reveal that skilled finfluencers are return contrarian; they make positive tweets after negative returns and negative tweets after positive returns. This suggests that skilled finfluencers may be good at identifying overreactions in the market, where a stock might be undervalued after bad news (leading to positive tweets) or overvalued after good news (resulting in negative tweets).

Skilled finfluencers are also social sentiment and news contrarian; they make fewer positive tweets when social sentiment is positive and more positive tweets after negative news. For instance, when the overall social sentiment towards a stock or market is positive, they tend to make fewer positive tweets, possibly reflecting a cautious attitude towards crowd behavior or potential market bubbles. Conversely, they make more positive tweets following negative news, perhaps seeing potential opportunities where others see only risk. This contrarian approach extends to negative sentiment as well. Skilled finfluencers make more negative tweets when social sentiment is positive, potentially warning their followers about overoptimistic evaluations. When sentiment is negative or after negative news, they make fewer negative tweets, possibly pointing out undervalued opportunities or questioning the crowd's pessimistic outlook.

The next finding is that skilled finfluencers tweet more frequently following periods of the high trading volume. This might indicate their active monitoring of market dynamics and willingness to provide timely input when there are significant market activities. Lastly, the ability to post negative tweets about stocks with short-sale constraints is another indicator of a skilled finfluencer. Short-selling constrained stocks usually come with higher risks and complexities, and a negative stance could indicate the finfluencer's understanding of these additional challenges and their ability to provide cautionary advice accordingly. Taken together, these characteristics provide valuable insights into the behaviors that might be indicative of a skilled finfluencer. By understanding these patterns, followers can better select which finfluencers to trust and follow, and other finfluencers can learn and improve their own practices.

Column 2 shows that finfluencers who post less frequently and have fewer ideas are more likely to be skilled. This might seem counterintuitive at first glance, but it could suggest that these influencers invest more time and effort in their market analysis before posting, which may result in

less frequent, but more accurate, advice.

Detecting antiskilled finfluencers. Columns 3 and 4 show that antiskilled finfluencers ride return and social sentiment momentum. In essence, they echo the existing market sentiment in their tweets, making positive posts following positive returns and negative posts after negative returns. This may suggest that these influencers simply go along with prevailing market trends, rather than analyzing or challenging them. Their commentary might lack depth and independent thought, and instead reflect a form of herd mentality. This momentum riding also applies to their response to social sentiment. When social sentiment is positive, antiskilled finfluencers are more likely to make positive tweets and less likely to make negative tweets. This further underscores their tendency to align with prevailing views, rather than offering a unique perspective or challenging conventional thinking.

The pattern continues when it comes to news sentiment. Antiskilled finfluencers are less likely to make positive tweets and more likely to make negative tweets in response to negative news. This shows a propensity to amplify prevailing sentiment, whether it's overly optimistic or overly pessimistic, rather than providing a balanced or contrarian viewpoint. Lastly, antiskilled finfluencers tend to make positive tweets even when market volatility is high or when stocks are subject to short-sale constraints, both of which typically signify higher risk. This may suggest a lack of understanding or disregard for the complexities and risks involved in these scenarios, which can potentially mislead their followers. These findings together paint a picture of antiskilled finfluencers as those who tend to go along with the crowd and avoid challenging the status quo, potentially missing out on nuanced analysis and balanced advice.

Column 4 demonstrates that more active finfluencers have a higher likelihood of being antiskilled. This might suggest that such influencers gain followers through high-profile but potentially reckless or overly simplistic market commentary. It is also possible that these influencers might prioritize gaining a large follower base over providing thoughtful, well-informed advice.

Detecting unskilled finfluencers. Columns 5 and 6 show that, when it comes to market news sentiment, the results indicate a correlation between the skills of a finfluencer and how they react to

“hot” stocks - those that are currently popular or making news. Specifically, those finfluencers who exhibit extreme reactions, whether positive or negative, to these trending stocks are more likely to be unskilled. This could suggest that they rely too heavily on the market’s overall sentiment or news headlines rather than conducting their own thorough analysis. Moreover, those finfluencers who frequently tweet about stocks with low trading volumes are also more likely to be unskilled. These low-volume stocks often lack the liquidity and market attention that larger, more frequently traded stocks have. Finfluencers focusing on these stocks might be less informed, using these lesser-known stocks as a way to appear unique or insightful, rather than providing solid advice based on well-analyzed information.

Short-sale constraints are also indicative of skill. Unskilled finfluencers tend to focus their tweets on stocks without short-sale constraints, potentially because these stocks are easier to analyze and speculate on. On the other hand, skilled finfluencers more often tweet negatively about stocks with short-sale constraints. This could be because they understand the additional risk involved in these stocks and caution their followers accordingly. In contrast, those finfluencers who show a positive bias towards short-sale constrained stocks, despite the inherent risk and complexity, are termed “antiskilled.” These individuals might be either downplaying or not understanding the risk involved, leading to potentially misleading information being disseminated to their followers.

Column 6 shows that very active finfluencers with many ideas are more likely to be uninformed or/and unskilled. Few likes by followers indicate a lack of skill. A higher number of users followed also indicates that the user is rather unskilled.

4.3 Which finfluencers’ tweeting strategies are most popular?

To narrow down the channel for why certain finfluencers are more popular than others despite the fact that skill is at least partially detectable, we now drill into the determinants of popularity by linking it to the alpha determinants used in the prior sections. We use the same characteristics to predict the users’ follower count out-of-sample as the tweeting strategies used to explain alpha in Table 8.

Table 10 reports the results from regressing the number of followers for each finfluencer on the

Table 10: Effect of finfluencers' tweeting patterns on follower count

This table reports the results of regressing the number of followers on users' 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 8. Standard errors are robust to heteroskedasticity. * , ** , and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Finfluencers' follower count _{<i>i</i>} (measured out-of-sample)				
	(1)	(2)	(3)	(4)	(5)
<i>ReturnChasing_i</i>	-0.013 (0.001)				0.000 (0.001)
<i>ContrarianTweet_i</i>	-0.021 (0.001)				-0.004 (0.001)
<i>NumberTweets_i</i>		0.683 (0.010)			0.828 (0.011)
<i>FractionNegative_i</i>		-0.001 (0.000)			-0.002 (0.000)
<i>PositiveHerd_i</i>			0.025 (0.004)		-0.005 (0.003)
<i>NegativeHerd_i</i>			-0.036 (0.004)		-0.021 (0.003)
<i>SSI_i (Positive Tweets)</i>				-0.020 (0.004)	-0.056 (0.004)
<i>SSI_i (Negative Tweets)</i>				-0.057 (0.006)	-0.063 (0.007)
Constant	1.275 (0.007)	0.208 (0.017)	1.290 (0.010)	1.240 (0.010)	0.463 (0.019)
N	22,027	22,027	22,027	22,027	22,027

characteristics of tweeting activity, i.e., return chasing, count and composition of tweets, herding, and short-selling constraints:

$$\text{Finfluencer's follower count}_i \text{ (measured out-of-sample)} = \alpha + \beta^T \text{TweetingStrategy}_i + \epsilon_i, \quad (22)$$

where the dependent variable is again the log of one plus the finfluencer's follower count as of February 2018, and $\text{TweetingStrategy}_i$ is the vector of tweeting/investment behaviors explored in Table 8.

Table 10 shows that the tendency to chase returns and post contrarian tweets negatively correlates with the finfluencer's follower count. However, the correlation fades away when we control for other finfluencer characteristics. On the other hand, users who tweet more often are more likely to have larger follower counts. A one percent increase in the total number of tweets is associated

with a 0.68% increase in followers. The correlation between the share of negative tweets and the number of followers is negative and significant but small in economic magnitude. 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 suggest that except for tweeting positively about stocks with high short-selling constraints, tweeting patterns that correlate with finfluencers' skills either do not predict the number of followers or predict it with the wrong sign, suggesting that social media users match with finfluencers based on their own behavioral traits. 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).

In summary, social media users tend to follow finfluencers with similar behavioral traits as their own. Retail investors also put their money where their finfluencer mouth is, especially those that follow antiskilled finfluencers. But this strategy is bound to lose money because the finfluencers that they are more likely to follow have negative predictive power. In the next section, we explore if one can exploit the wisdom of the skilled finfluencers to earn abnormal returns (both in-sample and out-of-sample) and if it is profitable to exploit the “wisdom” of the antiskilled finfluencers?

5 Belief Biases and the “Wisdom” of the Crowd

This section explores the asset price distortions and aggregate belief biases introduced by following antiskilled finfluencers' advice. In view of the previous sections' findings, one hypothesis could be that information is diffuse and dispersed among all finfluencers and needs to be aggregated to filter out noise. This is the idea behind the widely used term “wisdom of the crowd.” An alternative hypothesis is that not all finfluencers hold valuable information, but only a subset of skilled finfluencers are informed. A complementary hypothesis is that finfluencers catering to retail investors persistently provide flawed investment advice and one can earn abnormal returns doing the

opposite of their advice. To distinguish between these hypotheses, we next investigate if following the tweets by different groups of finfluencers in aggregate (i) leads to systemically biased beliefs across stocks and time, and (ii) generates profitable trading strategies and, if so, by which types of finfluencers.

5.1 Belief biases induced by antiskilled finfluencers

We can compute the stock-level and aggregate bias in beliefs resulting from the tweets of antiskilled finfluencers by comparing them to the tweets of unskilled finfluencers. The identifying assumption here is that unskilled finfluencers produce mostly noise and that any stock-specific and time-specific confounding factors are reflected in systematic patterns of their tweeting activity. Confounding factors can hence be filtered out by netting out unskilled finfluencers' average sentiment.

We perform the following steps to run our tests. We first calculate the aggregate measures of beliefs at the stock level as in Table 7, but instead of separating the positive and negative sentiments, we net them. To be more specific, we use the following formulas (for finfluencer i , stock j , day t):

$$\begin{aligned} Sent_{j,t}^{skilled} &= \frac{1}{I} \sum_{\text{all } i} \Pr(\text{user } i \text{ skilled}) \quad SocSent_{i,j,t}, \\ Sent_{j,t}^{antiskilled} &= \frac{1}{I} \sum_{\text{all } i} \Pr(\text{user } i \text{ antiskilled}) \quad SocSent_{i,j,t}, \end{aligned} \quad (23)$$

where $\Pr(\text{user } i \text{ skilled})$ and $\Pr(\text{user } i \text{ antiskilled})$ are given by expressions (11) and $SocSent_{i,j,t}$ is given by expression (1). The sentiment of unskilled finfluencers is defined similarly. To capture the belief bias induced by antiskilled users tweeting about stocks about which they are misinformed or faking their tweets, we define the belief bias relative to the sentiment of the unskilled finfluencers. We do this for the skilled and the antiskilled finfluencers in every stock j and every day t :

$$\begin{aligned} AbnSent_{j,t}^{skilled} &= Sent_{j,t}^{skilled} - Sent_{j,t}^{unskilled}, \\ AbnSent_{j,t}^{antiskilled} &= Sent_{j,t}^{antiskilled} - Sent_{j,t}^{unskilled}. \end{aligned} \quad (24)$$

Figure 5 plots the average abnormal social sentiment of skilled and antiskilled finfluencers by day (Panel A) and stock (Panel B) for the years 2015 and 2016. To construct daily averages we

Panel A: Abnormal social sentiment by day

Panel B: Abnormal social sentiment by stock

Figure 5: Abnormal Social Sentiment for Skilled and Antiskilled Fin uencers

The plot in Panel A shows the daily average abnormal social sentiment by skilled and antiskilled n uencers, respectively. The plot in Panel B shows the distribution of the average abnormal social sentiment by skilled and antiskilled n uencers, respectively, for each stock.

aggregate the abnormal social sentiment either by day or stock:

$$\begin{aligned} \text{AbnSent}_t^{\text{anti=skilled}} &= \frac{1}{J} \sum_{\text{all } j} \text{AbnSent}_{j;t}^{\text{anti=skilled}} ; \\ \text{AbnSent}_j^{\text{anti=skilled}} &= \frac{1}{T} \sum_{\text{all } t} \text{AbnSent}_{j;t}^{\text{anti=skilled}} ; \end{aligned} \tag{25}$$

Table 11: Abnormal Social Sentiment Revealed by the Tweets of Skilled and Antiskilled Finfluencers

This table reports descriptive statistics about the abnormal social sentiment revealed by the tweets of skilled and antiskilled influencers by day or stock. We measure abnormal social sentiment by skilled and antiskilled influencers relative to unskilled influencers. Standard errors are robust to clustering at the stock and day level. * , ** , and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	N	Mean	S.D.	Skewness	Kurtosis	p25	p50	p75
Panel A: Abnormal social sentiment revealed by skilled influencers								
Skilled influencers, by day	692	0.02	0.06	1.29	5.48	-0.01	0.00	0.04
Skilled influencers, by stock	4,580	0.02	0.07	1.15	16.99	-0.01	0.01	0.04
Panel B: Abnormal social sentiment revealed by antiskilled influencers								
Antiskilled influencers, by day	692	0.09	0.09	-0.22	7.22	0.05	0.07	0.11
Antiskilled influencers, by stock	4,580	0.08	0.07	1.89	23.66	0.04	0.08	0.12

The figure illustrates several intriguing patterns. 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 skilled influencers is centered at zero with several episodes when skilled influencers disseminate strongly positive social sentiment for extended periods of time and a few episodes when skilled influencers disseminate strongly negative social sentiment. By contrast, antiskilled influencers behave very differently. The daily average abnormal social sentiment of antiskilled influencers is significantly positive almost all the time. This implies antiskilled influencers in aggregate tend to tweet more positively than negatively, biasing their followers' beliefs upward. Antiskilled influencers' sentiment exhibits persistent swings and few spikes, in contrast to skilled influencers. Users that follow antiskilled influencers thus exhibit overly optimistic beliefs most of the time, overly pessimistic beliefs some of the time, and persistent swings in their belief bias. Panel B of Figure 5 demonstrates that antiskilled influencers are significantly positive about most stocks, while the fraction of stocks skilled influencers are positive about is almost the same as the fraction they are negative about.

Table 11 reports summary statistics for the abnormal social sentiment revealed by the tweets of skilled and antiskilled influencers. The statistics in Table 11 are consistent with Figure 5. The abnormal social sentiment revealed by skilled influencers in Panel A is close to zero on average with positive skewness and large kurtosis, both across time and stocks. By contrast, the abnormal social

sentiment revealed by antiskilled finfluencers in Panel B is positive on average and even at the 25% quantile. Its volatility over time is larger than that of skilled finfluencers. At the 75% quantile, antiskilled finfluencers' abnormal social sentiment exceeds 11%, both across time and stocks.

5.2 Belief biases and abnormal stock returns

Next, we investigate whether the tweeting activity of different finfluencers leads to inefficient prices directly through biased beliefs and indirectly through inducing more retail trading which has a price impact. To address the joint endogeneity of stock returns, retail order imbalances, and tweets we utilize different techniques. We start with a panel VAR specification that treats all variables as endogenous and interdependent, both in a contemporaneous and dynamic sense. We then perform a series of portfolio tests.

Panel VAR. We collect in vector $Y_{j,t}$ the six endogenous variables for return in every stock j and every day t , retail order imbalances (ROI), and positive (negative) tweets by skilled (antiskilled) finfluencers:

$$Y_{j,t} = \begin{pmatrix} Ret_{j,t} \\ ROI_{j,t} \\ Skilled_PosSent_{j,t} \\ Antiskilled_PosSent_{j,t} \\ Skilled_NegSent_{j,t} \\ Antiskilled_NegSent_{j,t} \end{pmatrix}. \quad (26)$$

For the variables in (26), we identify skilled and antiskilled finfluencers, respectively, as users in the highest and lowest deciles based on expected alpha, $E[\alpha_j | \tilde{\alpha}_i]$. The panel VAR specification for $Y_{j,t}$ is

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

with 6-dimensional error term $\epsilon_{j,t} \sim iid(0, \Sigma)$ and lag length L . We estimate (27) 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. Table A.1 in the Appendix summarizes the GMM estimation results.

Figure 6 shows the impulse response functions (IRF) of the six endogenous variables (Return, ROI, Skilled_PosSent, Antiskilled_PosSent, Skilled_NegSent, Antiskilled_NegSent) to unit shocks. Based on the GMM estimates with $L = 2$ and the Wold decomposition based on the order of the variables in (26), the IRFs show how $Y_{j,t+h}$, $h = 1, \dots, 6$, 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.

The first row in Figure 6 shows the impact of returns over the next 6 days of shocks to ROI and social sentiment. ROI and positive social sentiment by skilled finfluencers positively predict future returns whereas negative social sentiment by skilled finfluencers negatively predicts future returns. More surprisingly, positive (negative) social sentiment by antiskilled finfluencers negatively (positively) predicts future returns. The second row in Figure 6 shows that positive returns reinforce positive retail order imbalances on the following day. Similarly, positive sentiment by both skilled and antiskilled finfluencers encourages positive retail order imbalances over several days. Negative sentiment by both skilled and antiskilled finfluencers encourages negative retail order imbalances over the next day but the impact is weaker than for positive sentiment. The remaining rows decompose the impulse responses of the four different social sentiment variables. Past returns lead to more tweets in the same direction as the stock price movement, ROI shocks lead to higher tweeting activity, and past tweet activity leads to more tweeting activity in the future irrespective of the direction and the source. Overall, the panel VAR results suggest that (anti)skilled finfluencers (in)correctly predict future returns, and yet both types stipulate more retail order imbalances.

Belief biases and in-sample portfolio tests. Next, we investigate whether the belief biases can be exploited to earn abnormal returns in more classical portfolio sorts. There are several empirical choices to be made in constructing portfolios based on influencer tweets and, hence, there are several ways we can run portfolio tests using signals from StockTwits. Our baseline approach proceeds as follows:

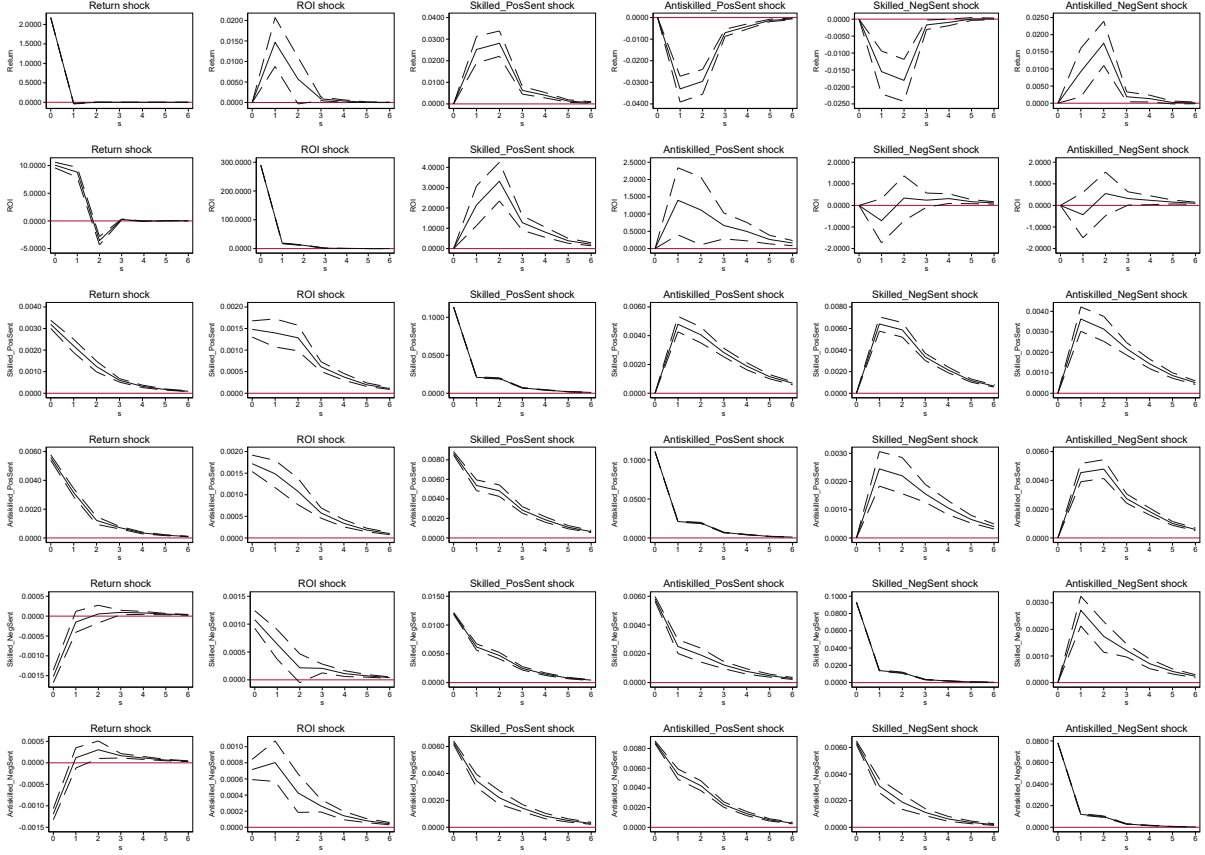


Figure 6: Impulse response functions

The plot shows the impulse response functions of the six endogenous variables (Return, ROI, Skilled_PosSent, Antiskilled_PosSent, Skilled_NegSent, Antiskilled_NegSent) to unit shocks. The specification is from (27) with $L = 2$.

1. We identify users in the highest and lowest deciles based on expected alpha, $E[\alpha \mid \tilde{\alpha}_i]$. Alternatively, we identify users in the highest deciles based on the probability of being skilled, $\Pr(\alpha > 0 \mid \tilde{\alpha}_i)$, and probability of being antiskilled, $\Pr(\alpha < 0 \mid \tilde{\alpha}_i)$. In both cases, we denote the two groups as skilled and antiskilled.
2. Every day t , we get a list of stocks that have been mentioned positively and negatively by each group over the past H days, where we vary $H = 1, 2, 5, 10, 20$. That is, a stock stays in the portfolio for H days if tweeted on days $t - H + 1, \dots, t$. Denote the group of stocks tweeted on the day t by $Tweet_t$ and over the past H days by $Tweet_{t - H + 1, t}$.
3. Every day t , we go long a portfolio of stocks that have been either (1) tweeted positively

by the skilled group, or (2) tweeted negatively by the unskilled group. Similarly, we short a portfolio of stocks that have been either (3) tweeted negatively by the skilled group, or (4) tweeted positively by the antiskilled group. This approach yields four legs of a composite strategy.

4. We calculate the time series of daily excess returns for each of the four portfolios. We compute buy-and-hold portfolio returns where we make the initial investment at the close of the day the tweets occur and hold the initial positions for H days. Portfolio returns for trades initiated based on tweets on the day t , $Tweet_t$, are

$$Ret_{t+1,t+H}^{bh} = \frac{1}{jTweet_t} \sum_{j \in Tweet_t} \prod_{h=1}^H (1 + AbnRet_{j,t+h}) - 1.$$

5. We construct the long-short returns by subtracting the returns of the short portfolio from those of the long portfolio.

Table 12 provides in-sample buy-and-hold portfolio returns with the reported numbers being multi-day returns $Ret_{t+1,t+H}^{bh}$ over the corresponding holding period. Across panels, we vary the influencers and tweet content. Panels A and C (B and D) report results for positive (negative) tweeting activity and Panels A and B (C and D) split results into skilled (antiskilled) influencers according to the procedure of variables construction described above. The portfolio returns show that skilled influencers' positive tweets predict positive returns over 1, 2, 5, 10, and 20-day horizons, reaching 2.3% over 20 days in the FF5 specification. Similarly, skilled influencers' negative tweets predict significant negative returns over 1, 2, 5, 10, and 20-day horizons, reaching -2.4% over 20 days. The results for antiskilled influencers are the exact opposite. The portfolio returns show that antiskilled influencers' positive tweets predict negative returns over 1, 2, 5, 10, and 20-day horizons, reaching -4.6% over 20 days. Antiskilled influencers' negative tweets predict significant positive returns over 10 and 20-day horizons, reaching 1.3% over 20 days. The results are consistent with the panel VAR in that the social sentiment of (anti)skilled influencers (in)correctly predicts returns over several days.

Table 12: Portfolio Returns

The table documents buy-and-hold portfolio returns. The reported numbers are multi-day buy-and-hold returns over the corresponding holding period $[t + 1, t + H]$ and $H \in \{1, 2, 5, 10, 20\}$. Standard errors are robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1) Raw	(2) FF1	(3) FF3	(4) FF5
Panel A: Positive tweets by skilled investors				
1-day abnormal return	0.211*** (0.057)	0.169*** (0.038)	0.146*** (0.034)	0.151*** (0.033)
2-day abnormal return	0.432*** (0.085)	0.342*** (0.055)	0.310*** (0.047)	0.319*** (0.047)
5-day abnormal return	0.998*** (0.141)	0.767*** (0.090)	0.695*** (0.074)	0.725*** (0.074)
10-day abnormal return	1.999*** (0.195)	1.476*** (0.125)	1.319*** (0.102)	1.359*** (0.101)
20-day abnormal return	3.598*** (0.248)	2.553*** (0.160)	2.318*** (0.133)	2.329*** (0.134)
Panel B: Negative tweets by skilled investors				
1-day abnormal return	-0.097 (0.067)	-0.145** (0.049)	-0.181*** (0.047)	-0.181*** (0.046)
2-day abnormal return	-0.155 (0.098)	-0.263*** (0.071)	-0.329*** (0.067)	-0.330*** (0.067)
5-day abnormal return	-0.336* (0.160)	-0.600*** (0.116)	-0.727*** (0.110)	-0.726*** (0.110)
10-day abnormal return	-0.299 (0.221)	-0.875*** (0.165)	-1.176*** (0.163)	-1.193*** (0.160)
20-day abnormal return	-0.791** (0.290)	-1.922*** (0.220)	-2.340*** (0.219)	-2.423*** (0.215)
Panel C: Positive tweets by antiskilled investors				
1-day abnormal return	-0.275*** (0.058)	-0.320*** (0.041)	-0.362*** (0.038)	-0.360*** (0.038)
2-day abnormal return	-0.443*** (0.087)	-0.544*** (0.059)	-0.625*** (0.054)	-0.623*** (0.054)
5-day abnormal return	-0.894*** (0.127)	-1.158*** (0.083)	-1.331*** (0.075)	-1.334*** (0.075)
10-day abnormal return	-1.545*** (0.167)	-2.087*** (0.108)	-2.359*** (0.098)	-2.367*** (0.098)
20-day abnormal return	-3.058*** (0.228)	-4.128*** (0.150)	-4.565*** (0.136)	-4.593*** (0.134)
Panel D: Negative tweets by antiskilled investors				
1-day abnormal return	0.132 (0.083)	0.096 (0.070)	0.066 (0.067)	0.069 (0.068)
2-day abnormal return	0.249* (0.123)	0.169 (0.101)	0.128 (0.097)	0.125 (0.097)
5-day abnormal return	0.547** (0.179)	0.286* (0.141)	0.196 (0.128)	0.193 (0.128)
10-day abnormal return	1.138*** (0.257)	0.572** (0.201)	0.340 (0.176)	0.388* (0.175)
20-day abnormal return	2.684*** (0.364)	1.574*** (0.295)	1.229*** (0.260)	1.256*** (0.248)

Additional portfolio tests. As a robustness check, we dynamically readjust the portfolio every day to account for the varying number of stocks being tweeted about by adjusting the initial positions for how many stocks are in each portfolio. We compute dynamic portfolio returns where we rebalance the initial positions for H days. Portfolio returns over $[t + 1, t + H]$ are

$$Ret_{t+1,t+H}^{dy} = \frac{H}{\sum_{j \in Tweet_{t+1,t}} 1} \sum_{j \in Tweet_{t+1,t}} AbnRet_{j,t+1}.$$

In Panel A of Table 13, the reported numbers are dynamically rebalanced returns $Ret_{t+1,t+20}^{dy}$ over a 20-day holding period. The results are broadly in line with Table 12. The main differences are that positive tweets by antiskilled finfluencers now produce even larger negative returns of -6% over 20 days, while the positive returns following negative tweets by antiskilled finfluencers are statistically insignificant.

In another robustness check, in Panels B and C of Table 13 we document in-sample portfolio returns using the probability of (anti)skill as a sorting variable. Here we identify skilled finfluencers as users in the highest decile of $\Pr(\alpha > 0 | \tilde{\alpha}_i)$ and antiskilled finfluencers as users in the highest decile of $\Pr(\alpha < 0 | \tilde{\alpha}_i)$. The reported numbers in Panel B are cumulative abnormal returns $Ret_{t+1,t+20}^{bh}$ over a 20-day holding period based on the sorting variable. The reported numbers in Panel C are dynamically rebalanced abnormal returns $Ret_{t+1,t+20}^{dy}$. The results are again broadly in line with Table 12 but the alphas are overall smaller in magnitude. The reason is that probability-based sorting is noisier than expectation-based sorting. For positive (negative) tweets by skilled finfluencers, the monthly alpha in Panel B becomes 0.52% (-0.17%). Positive tweets by antiskilled finfluencers again predict negative returns, now of -0.63%. The main difference with Table 12 is that negative tweets by antiskilled finfluencers produce 1% raw returns (column 1) but once we control for market movements the alpha becomes negative. The alphas in Panel C are similar to Panel B with the main difference being that fewer are statistically significant.

Last, we repeat the tests except we identify the skilled and antiskilled groups in the first part of the data (pre-2016) and run our portfolio tests in the second part of the data (post-2016).

1. We identify users in the highest and lowest deciles based on $E[\alpha | \tilde{\alpha}_{i,pre-2016}]$, or $\Pr(\alpha > 0 |$

$\tilde{\alpha}_{i,\text{pre-2016}}$) and $\Pr(\alpha < 0 \mid \tilde{\alpha}_{i,\text{pre-2016}})$ calculated in the pre-2016 period. We again denote these two groups as skilled and antiskilled.

2. For every day post-2016, we get a list of stocks that have been mentioned positively and negatively by each group over the past H days.
3. For every day post-2016, we go long a portfolio of stocks that have been either (1) tweeted positively by the skilled group, or (2) tweeted negatively by the unskilled group. Similarly, we short a portfolio of stocks that have been either (3) tweeted negatively by the skilled group, or (4) tweeted positively by the antiskilled group.
4. We calculate the time series of daily excess returns for the four portfolio legs and subtract the returns of the short portfolios from those of the long portfolios.

Panel D of Table 13 summarizes the results. It uses notations akin to the ones used in Table 12 but for skills measured using a pre-2016 sample. The panel provides out-of-sample buy-and-hold portfolio returns with the reported numbers being 20-day returns over the corresponding holding period. The out-of-sample results are generally weaker than the in-sample tests in Table 12 and Table 13, Panels A-C. The out-of-sample portfolio returns show that influencers identified as being skilled before 2016 do not significantly predict returns in 2016. By contrast, the performance of antiskilled influencers is persistent. Antiskilled influencers' positive tweets predict negative returns over all horizons, reaching -1.24% over 20 days. Antiskilled influencers' negative tweets also predict significant negative out-of-sample returns, reaching -1.05% over 20 days.

The out-of-sample portfolio results are quite interesting when combined with the findings from Table 5 that the influencers' skills are persistent but are not sufficient for influencers' survival according to Table 6. They indicate that the message is more important than the messenger. That is as long as there are any antiskilled influencers "preaching" their message the investors like their message and are willing to trade on it.

6 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 can use the tweeting behavior of finfluencers to identify their skills. However, instead of following more skilled influencers, social media users follow unskilled and antiskilled finfluencers, which we define as finfluencers whose tweets generate negative alpha. Antiskilled finfluencers ride return and social sentiment momentum, which coincide with the behavioral biases of retail investors who trade on antiskilled finfluencers' flawed advice.

These results are consistent with homophily in behavioral traits between social media users and finfluencers shaping finfluencer's follower networks and limiting competition among finfluencers, resulting in the survival of un- and antiskilled finfluencers despite the fact that they do not provide valuable investment advice.

Investing contrarian to the tweets by antiskilled finfluencers yields abnormal out-of-sample returns, which we term the "wisdom of the antiskilled crowd." These findings shed light on the quality of finfluencers' unsolicited financial advice and the competition among and economic incentives faced by finfluencers which the SEC has been concerned about.

Table 13: Portfolio Returns: Robustness Checks

The table documents in-sample portfolio returns using alternative portfolio constructions. In Panel A, the reported numbers are returns over a 20-day holding period with dynamic rebalancing, $Ret_{t+1,t+20}^{dy}$. In Panel B, the reported numbers are buy-and-hold returns over a 20-day holding period with the probability of (anti)skill as a sorting variable. In Panel C, the reported numbers are returns over a 20-day holding period with dynamic rebalancing, $Ret_{t+1,t+20}^{dy}$, with the probability of (anti)skill as a sorting variable. In Panel D, the reported numbers are buy-and-hold returns over a 20-day holding period during the post-2016 period with the expected alpha computed pre-2016 as a sorting variable. Standard errors are robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1) Raw	(2) FF1	(3) FF3	(4) FF5
Panel A: 20-day dynamically rebalanced returns $Ret_{t+1,t+20}^{dy}$ based on $E[\alpha_j \alpha_i]$				
Positive tweets by skilled investors	4.114*** (1.073)	3.201*** (0.669)	2.788*** (0.581)	2.842*** (0.572)
Negative tweets by skilled investors	-1.877 (1.230)	-2.909*** (0.835)	-3.504*** (0.788)	-3.458*** (0.775)
Positive tweets by antiskilled investors	-4.451*** (1.090)	-5.382*** (0.713)	-6.095*** (0.658)	-6.075*** (0.651)
Negative tweets by antiskilled investors	2.231 (1.441)	1.339 (1.129)	0.947 (1.053)	1.009 (1.054)
Panel B: 20-day buy-and-hold returns $Ret_{t+1,t+20}^{bh}$ based on $\Pr(\alpha_i > 0 \alpha_i)$				
Positive tweets by skilled investors	1.600*** (0.163)	0.691*** (0.068)	0.512*** (0.050)	0.519*** (0.050)
Negative tweets by skilled investors	0.911*** (0.199)	-0.027 (0.124)	-0.196 (0.117)	-0.171 (0.119)
Positive tweets by antiskilled investors	0.557*** (0.156)	-0.355*** (0.056)	-0.614*** (0.041)	-0.634*** (0.040)
Negative tweets by antiskilled investors	0.998*** (0.191)	0.046 (0.098)	-0.263*** (0.074)	-0.254*** (0.072)
Panel C: 20-day dynamically rebalanced returns $Ret_{t+1,t+20}^{dy}$ based on $\Pr(\alpha_i > 0 \alpha_i)$				
Positive tweets by skilled investors	1.788* (0.800)	0.997** (0.321)	0.692** (0.240)	0.701** (0.242)
Negative tweets by skilled investors	0.345 (0.895)	-0.529 (0.480)	-0.865* (0.424)	-0.855* (0.423)
Positive tweets by antiskilled investors	0.600 (0.788)	-0.215 (0.268)	-0.563** (0.189)	-0.570** (0.187)
Negative tweets by antiskilled investors	0.994 (0.885)	0.141 (0.431)	-0.201 (0.348)	-0.202 (0.349)
Panel D: 20-day buy-and-hold returns $Ret_{t+1,t+20}^{bh}$ based on $E[\alpha_j \alpha_{i,pre-2016}]$				
Positive tweets by skilled investors	2.743*** (0.393)	0.771*** (0.221)	-0.288 (0.184)	-0.307 (0.183)
Negative tweets by skilled investors	2.745*** (0.468)	0.684* (0.297)	-0.393 (0.262)	-0.384 (0.261)
Positive tweets by antiskilled investors	2.156*** (0.466)	-0.032 (0.304)	-1.309*** (0.268)	-1.243*** (0.265)
Negative tweets by antiskilled investors	1.995*** (0.458)	-0.113 (0.329)	-1.180*** (0.318)	-1.045** (0.314)

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.
- 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.
- Boehmer, Ekkehart, Charles M. Jones, and Xiaoyan Zhang (2008). “Which Shorts Are Informed?” In: *The Journal of Finance* 63.2, pp. 491–527.
- Boehmer, Ekkehart, Charles M. Jones, Xiaoyan Zhang, and Xinran Zhang (2021). “Tracking retail investor activity”. In: *The Journal of Finance* 76.5, pp. 2249–2305.
- 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 Engelberg, and William Mullins (2021). “Echo chambers”. In: *Working Paper*.
- 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*.
- 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.
- Farrell, Michael, T. Clifton Green, Russell Jame, and Stanimir Markov (2022). “The democratization of investment research and the informativeness of retail investor trading”. In: *Journal of Financial Economics* 145.2, pp. 616–641.
- 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.
- Goutte, Maud (2020). “When the ‘Dumb’ Crowd Beats the ‘Smart’ Crowd”. In: *Working Paper*.
- 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.
- Jame, Russell, Rick Johnston, Stanimir Markov, and Michael C. Wolfe (2016). “The value of crowd-sourced 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.
- 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.
- Sprenger, Timm O., Andranik Tumasjan, Philipp G. Sandner, and Isabell M. Welpe (2014). “Tweets and trades: The information content of stock microblogs”. In: *European Financial Management* 20.5, pp. 926–957.

Appendix

This Appendix describes the construction of variables describing tweeting behavior and reports additional empirical results discussed in the main body of the manuscript.

To detect influencer skill, we construct variables that capture the tweeting behavior of different influencers. We proceed in two steps. We first construct stock-level events triggering tweets. We then calculate for each influencer i the average decile of all stocks that i tweets positive (negative) about after the stock has satisfied an event-based criterion over the past time window $[t - L + 1, t - 1]$. For the window length, we set $L = 20$. Alternatively, we have let $L = 1, 2, 5, 10$. We use the following stock-level event-based criteria to capture triggers that cause influencers to post tweets.

Step 1: Stock-level events triggering tweets.

1. *Stock j on the day t is in the highest (lowest) decile of past returns.* We measure past returns by the lagged 20-day Fama-French 5-factor abnormal return.¹
2. *Stock j on the day t is in the highest (lowest) decile of past positive/negative social sentiment.* Bloomberg measures social sentiment using a proprietary machine learning algorithm and reports a social sentiment score on a discrete scale, $SocSent_{i,j,t,n} \in \{-1, 0, 1\}$, with an associated confidence level between 1/3 to 1. Out of about 72 million tweets in our sample, 11%/77%/12% are positive/neutral/negative. The variables $SocSent_{j,t}^p/SocSent_{j,t}^n/SocSent_{j,t}^m$ count the number of positive/neutral/negative tweets about stock j on the day t . We measure past social sentiment by the average fraction of positive (negative) tweets over the prior L days:

$$\begin{aligned} SocSent_{j,t}^{p\%} &= \frac{1}{L} \sum_{s=t-L+1}^{t-1} \frac{SocSent_{j,s}^p}{SocSent_{j,s}^p + SocSent_{j,s}^n + SocSent_{j,s}^m}, \\ SocSent_{j,t}^{m\%} &= \frac{1}{L} \sum_{s=t-L+1}^{t-1} \frac{SocSent_{j,s}^m}{SocSent_{j,s}^p + SocSent_{j,s}^n + SocSent_{j,s}^m}. \end{aligned} \quad (A1)$$

Alternatively, we have computed past social sentiment by the average number of positive/negative tweets over the prior L days.

3. *Stock j on the day t is in the highest (lowest) decile of past positive/negative news sentiment.* Bloomberg measures news sentiment using a proprietary machine learning algorithm and reports a sentiment score on a discrete scale, $NewsSent_{j,t,n} \in \{-1, 0, 1\}$, with an associated confidence level between 1/3 to 1. Out of 36 million news stories, 12%/59%/29% are positive/neutral/negative. The variables $NewsSent_{j,t}^p/NewsSent_{j,t}^n/NewsSent_{j,t}^m$ count the number of positive/neutral/negative news stories about stock j on the day t . We measure news social sentiment by the average fraction of positive (negative) news over the prior L

¹Alternatively, we have measured past returns in different ways, by the L -day cumulative close-to-close return, the CAPM return, the Fama-French 3-factor abnormal return, and the Fama-French 5-factor abnormal return.

days:

$$\begin{aligned} NewsSent_{j,t}^{p\%} &= \frac{1}{L} \sum_{s=t-L+1}^t \frac{NewsSent_{j,s}^p}{NewsSent_{j,s}^p + NewsSent_{j,s}^n + NewsSent_{j,s}^m}, \\ NewsSent_{j,t}^{m\%} &= \frac{1}{L} \sum_{s=t-L+1}^t \frac{NewsSent_{j,s}^m}{NewsSent_{j,s}^p + NewsSent_{j,s}^n + NewsSent_{j,s}^m}. \end{aligned} \quad (A2)$$

Alternatively, we have computed past news sentiment by the average number of positive/negative news stories over the prior L days.

4. *Stock j on the day t is in the highest (lowest) decile of past absolute price movements.* We measure past absolute price movements by the average absolute close-to-close return over the past L days.
5. *Stock j on the day t is in the highest (lowest) decile of past retail order imbalances.* We measure past retail order imbalances by the average volume of retail purchases minus retail sales over the past L days, divided by the stock's market capitalization. We follow the method of Boehmer, Jones, Zhang, and Zhang (2021) to measure retail trading activity. The source of the data is TAQ.
6. *Stock j on the day t is in the highest (lowest) decile of past share turnover.* We measure past share turnover by the average trading volume divided by the stock's market capitalization over the past L days.
7. *Stock j on the day t is in the highest (lowest) decile of past short-sale constraints.* We capture short-sale constraints by the Markit indicator $dcbs_{j,t}$ ranging from 1 (unconstrained) to 10 (most constrained), averaged over the past L days. Alternatively, we have used a dummy variable that indicates $dcbs_{j,t} \in [2, 10]$.

Step 2. User-level events triggering tweets.

We next link stock-level events to user-level events. We denote the decile in which stock j falls on the day t according to any of the above events by $Decile_{j,t}^{Event} \in [1, 10]$. The user-level variable *Fin uencer posts positive tweets after Event* measures the average decile according to *Event* of the stocks that finfluencer i tweets positive about, averaged across stocks and time:

$$\begin{aligned} \text{Finfluencer}_i \text{ posts positive tweets after } Event &= \frac{\sum_j \sum_t (Decile_{j,t}^{Event} \cdot \mathbb{1}_{SocSent_{i,j,t} > 0})}{\sum_j \sum_t \mathbb{1}_{SocSent_{i,j,t} > 0}}, \\ \text{Finfluencer}_i \text{ posts negative tweets after } Event &= \frac{\sum_j \sum_t (Decile_{j,t}^{Event} \cdot \mathbb{1}_{SocSent_{i,j,t} < 0})}{\sum_j \sum_t \mathbb{1}_{SocSent_{i,j,t} < 0}}, \end{aligned} \quad (A3)$$

with *Event* \in {Returns, Social sentiment, News sentiment, Volatility, Retail order imbalance, Trading volume, Short-sale constraint}.

To give an example, suppose finfluencer i tweets positively about stock j on the day t . To capture news coverage as an event triggering positive tweets, we first calculate the number of positive news stories on Bloomberg for each stock over the $L = 20$ days ending on the day $t - 1$.

Denote this variable by $NewsSent^p$. We then calculate which decile of $NewsSent^p$ stock j belongs to on the day t . Our user-level variable is the average of this decile for all positive tweets of user i . Similarly, to capture social media an event triggering positive tweets, we first calculate the number of positive tweets from all StockTwits users reported on Bloomberg for each stock over the $L = 20$ days ending on the day $t - 1$. Denote this variable by $SocSent^p$. We then calculate which decile of $SocSent^p$ stock j belongs to on the day t . Our user-level variable is the average of this decile for all positive tweets of user i .

Table A.1: Panel VAR

The table reports coefficient estimates for the panel VAR in (27). Standard errors are reported in parenthesis, * , ** , and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1)		(2)	
	$L = 1$		$L = 2$	
Panel A: Return _t				
Return _{t-1}	-0.02***	(0.00)	-0.02***	(0.00)
ROI _{t-1}	0.00***	(0.00)	0.00***	(0.00)
Skilled_PosSent _{t-1}	0.29***	(0.03)	0.26***	(0.03)
Antiskilled_PosSent _{t-1}	-0.32***	(0.03)	-0.30***	(0.03)
Skilled_NegSent _{t-1}	-0.20***	(0.04)	-0.17***	(0.04)
Antiskilled_NegSent _{t-1}	0.16***	(0.05)	0.12***	(0.05)
Return _{t-2}			0.00	(0.00)
ROI _{t-2}			0.00	(0.00)
Skilled_PosSent _{t-2}			0.25***	(0.03)
Antiskilled_PosSent _{t-2}			-0.24***	(0.03)
Skilled_NegSent _{t-2}			-0.20***	(0.04)
Antiskilled_NegSent _{t-2}			0.22***	(0.05)
Panel B: ROI _t				
Return _{t-1}	3.06***	(0.21)	3.69***	(0.25)
ROI _{t-1}	0.06***	(0.00)	0.06***	(0.00)
Skilled_PosSent _{t-1}	25.74***	(4.73)	18.97***	(4.81)
Antiskilled_PosSent _{t-1}	15.17***	(5.15)	13.45***	(5.16)
Skilled_NegSent _{t-1}	-3.05	(6.55)	-7.23	(6.58)
Antiskilled_NegSent _{t-1}	-0.53	(7.91)	-5.51	(8.01)
Return _{t-2}			-2.08***	(0.22)
ROI _{t-2}			0.04***	(0.00)
Skilled_PosSent _{t-2}			22.38***	(4.57)
Antiskilled_PosSent _{t-2}			6.76	(4.99)
Skilled_NegSent _{t-2}			3.91	(6.33)
Antiskilled_NegSent _{t-2}			6.39	(7.66)
Panel C: Skilled_PosSent _t				
Return _{t-1}	0.00***	(0.00)	0.00***	(0.00)
ROI _{t-1}	0.00***	(0.00)	0.00***	(0.00)
Skilled_PosSent _{t-1}	0.19***	(0.00)	0.17***	(0.00)
Antiskilled_PosSent _{t-1}	0.05***	(0.00)	0.04***	(0.00)
Skilled_NegSent _{t-1}	0.08***	(0.00)	0.07***	(0.00)
Antiskilled_NegSent _{t-1}	0.06***	(0.00)	0.05***	(0.00)
Return _{t-2}			0.00***	(0.00)
ROI _{t-2}			0.00***	(0.00)
Skilled_PosSent _{t-2}			0.13***	(0.00)
Antiskilled_PosSent _{t-2}			0.02***	(0.00)
Skilled_NegSent _{t-2}			0.04***	(0.00)
Antiskilled_NegSent _{t-2}			0.02***	(0.00)

Continued.

Table A.1: Panel VAR—continued

The table reports coefficient estimates for the panel VAR in (27). Standard errors are reported in parenthesis. $*$, $**$, and $***$ indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	(1) $L = 1$		(2) $L = 2$	
Panel D: Antiskilled_PosSent _t				
Return _{t-1}	0.00***	(0.00)	0.00***	(0.00)
ROI _{t-1}	0.00***	(0.00)	0.00***	(0.00)
Skilled_PosSent _{t-1}	0.04***	(0.00)	0.03***	(0.00)
Antiskilled_PosSent _{t-1}	0.21***	(0.00)	0.18***	(0.00)
Skilled_NegSent _{t-1}	0.03***	(0.00)	0.02***	(0.00)
Antiskilled_NegSent _{t-1}	0.07***	(0.01)	0.06***	(0.01)
Return _{t-2}			0.00	(0.00)
ROI _{t-2}			0.00	(0.00)
Skilled_PosSent _{t-2}			0.01***	(0.00)
Antiskilled_PosSent _{t-2}			0.13***	(0.00)
Skilled_NegSent _{t-2}			0.01***	(0.00)
Antiskilled_NegSent _{t-2}			0.04***	(0.00)
Panel E: Skilled_NegSent _t				
Return _{t-1}	0.00	(0.00)	0.00	(0.00)
ROI _{t-1}	0.00	(0.00)	0.00***	(0.00)
Skilled_PosSent _{t-1}	0.04***	(0.00)	0.04***	(0.00)
Antiskilled_PosSent _{t-1}	0.02***	(0.00)	0.01***	(0.00)
Skilled_NegSent _{t-1}	0.16***	(0.00)	0.15***	(0.00)
Antiskilled_NegSent _{t-1}	0.04***	(0.00)	0.03***	(0.00)
Return _{t-2}			0.00	(0.00)
ROI _{t-2}			0.00	(0.00)
Skilled_PosSent _{t-2}			0.01***	(0.00)
Antiskilled_PosSent _{t-2}			0.00	(0.00)
Skilled_NegSent _{t-2}			0.10***	(0.00)
Antiskilled_NegSent _{t-2}			0.01***	(0.00)
Panel F: Antiskilled_NegSent _t				
Return _{t-1}	0.00	(0.00)	0.00	(0.00)
ROI _{t-1}	0.00***	(0.00)	0.00***	(0.00)
Skilled_PosSent _{t-1}	0.02***	(0.00)	0.02***	(0.00)
Antiskilled_PosSent _{t-1}	0.04***	(0.00)	0.04***	(0.00)
Skilled_NegSent _{t-1}	0.03***	(0.00)	0.02***	(0.00)
Antiskilled_NegSent _{t-1}	0.17***	(0.01)	0.15***	(0.01)
Return _{t-2}			0.00	(0.00)
ROI _{t-2}			0.00	(0.00)
Skilled_PosSent _{t-2}			0.00	(0.00)
Antiskilled_PosSent _{t-2}			0.01***	(0.00)
Skilled_NegSent _{t-2}			0.00	(0.00)
Antiskilled_NegSent _{t-2}			0.10***	(0.00)

Internet Appendix

A Alternative Specifications for the Distribution of True Alphas


Table [IA.1](#) 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. Figures – to – (– to –) show how close the estimated distribution and the data are for $K^+ = K^- = 1(3)$.

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

This table reports the results of fitting mixture models with one, two, and three components for skilled and antiskilled users. Means and probabilities are reported in percentage points.

	(1) $K^+ = K^- = 1$		(2) $K^+ = K^- = 2$		(3) $K^+ = K^- = 3$	
	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)
Skilled type 3					7.41	4.4
Skilled type 2			6.76	5.9	2.75	7.2
Skilled type 1	4.16	15.3	1.42	21.7	0.99	18.9
Unskilled	0.00	56.5	0.00	16.0	0.00	1.4
Antiskilled type 1	-4.33	28.3	-1.06	45.6	-0.44	35.5
Antiskilled type 2			-7.53	10.9	-1.82	24.1
Antiskilled type 3					-8.38	8.5
N	29,477		29,477		29,477	
Log Likelihood	-86,981		-86,385		-86,363	
AIC	173,971		172,786		172,750	
BIC	173,981		172,806		172,780	

Panel A: Estimated and simulated alphas



Panel B: Estimated and simulated t -stats

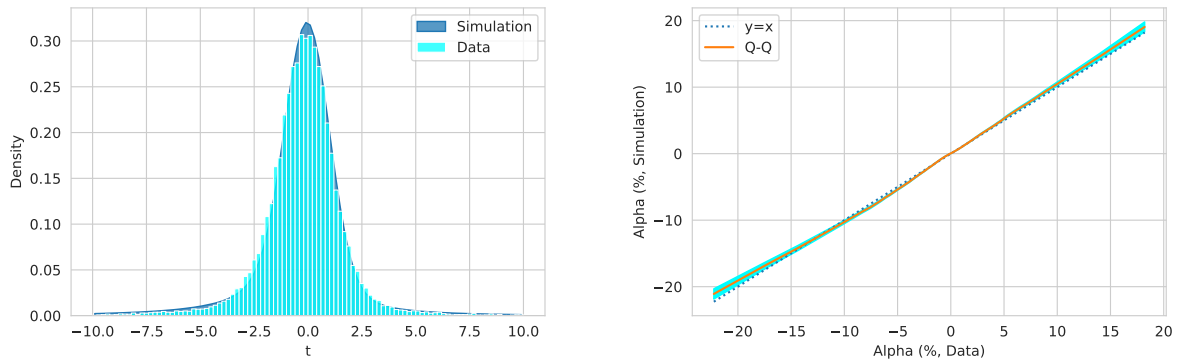
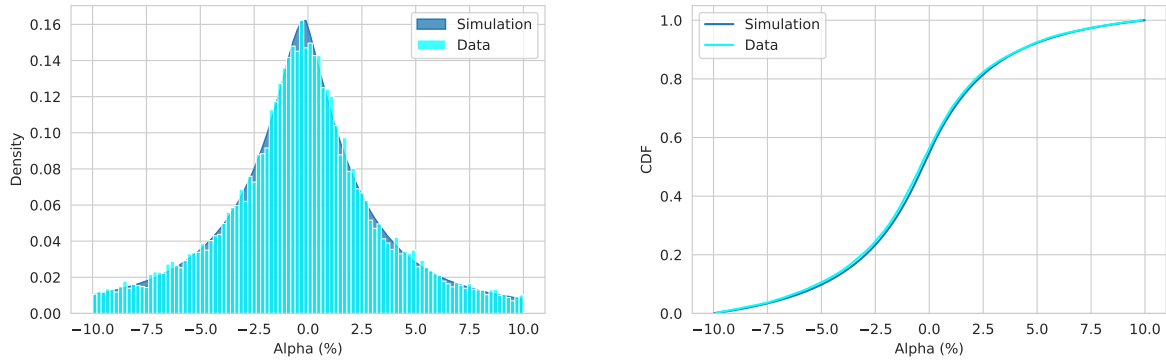


Figure IA.1: Estimated and Simulated Alphas and Their t -Stats From the Model with $K^+ = K^- = 1$

In Panel A, the left plot shows histograms of estimated and simulated alphas. In Panel A, the right plot shows the average cdf of simulated alphas from the fitted model against estimated 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 estimated and simulated alphas.

Panel A: Estimated and simulated alphas



Panel B: Estimated and simulated t -stats




Figure IA.2: Estimated and Simulated Alphas and Their t -Stats From the Model with $K^+ = K^- = 3$

In Panel A, the left plot shows histograms of estimated and simulated alphas. In Panel A, the right plot shows the average cdf of simulated alphas from the fitted model against estimated 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 estimated and simulated alphas.

Table IA.2: Estimating the Distribution of True α 's: Robustness to Different Horizons

This table reports the results of fitting mixture models 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 estimated alpha ($\hat{\alpha}$) 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.

	(1) $H = 1$		(2) $H = 2$		(3) $H = 5$		(4) $H = 10$		(5) $H = 20$	
	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)
Skilled type 1	2.66	2.4	3.21	1.7	3.43	4.2	5.18	4.5	6.76	5.9
Skilled type 2	0.34	17.6	0.68	16.0	0.68	18.9	1.18	18.4	1.42	21.7
Unskilled	0.00	49.4	0.00	48.6	0.00	25.9	0.00	23.5	0.00	16.0
Antiskilled type 1	-0.28	26.5	-0.42	29.1	-0.46	42.5	-0.64	42.6	-1.06	45.6
Antiskilled type 2	-2.34	4.2	-2.69	4.6	-3.63	8.4	-4.81	11.0	-7.53	10.9
N	30,720		30,329		30,175		30,054		29,477	
Log Likelihood	44,484		53,597		66,482		77,202		86,385	
BIC	-88,886		-107,112		-132,882		-154,322		-172,688	
AIC	-88,953		-107,179		-132,949		-154,389		-172,754	

Table IA.3: Estimating the Distribution of True α 's: Using Fama-French Three-Factor Model

This table reports the results of fitting mixture models with two exponentials on the $\alpha > 0$, two exponentials on the $\alpha < 0$, 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 three-factor model. The first number at the top of the columns shows the horizon of future returns. The estimated alpha ($\hat{\alpha}$) 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.

	(1) $H = 1$		(2) $H = 2$		(3) $H = 5$		(4) $H = 10$		(5) $H = 20$	
	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)	Mean alpha (%)	Fraction of users (%)
Skilled type 1	2.61	2.5	2.93	2.1	3.44	4.4	5.17	4.6	7.35	5.3
Skilled type 2	0.34	17.8	0.60	17.9	0.71	18.3	1.22	19.4	1.67	21.5
Unskilled	0.00	50.7	0.00	46.3	0.00	27.8	0.00	27.5	0.00	19.6
Antiskilled type 1	-0.30	25.0	-0.44	29.1	-0.53	41.5	-0.87	38.7	-1.19	42.4
Antiskilled type 2	-2.32	4.1	-2.68	4.6	-3.77	7.9	-5.13	9.9	-7.71	11.2
N	30,720		30,329		30,175		30,054		29,477	
Log Likelihood	44,774		53,966		67,118		78,148		87,644	
BIC	-89,466		-107,850		-134,153		-156,214		-175,205	
AIC	-89,532		-107,916		-134,219		-156,280		-175,271	