# Social contagion and asset prices: Reddit's self-organised bull runs

Semenova, Valentina and Winkler, Julian

Institute of New Economic Thinking at the Oxford Martin School, Department of Mathematics, University of Oxford, Department of Economics, University of Oxford

5 May 2021

# Social contagion and asset prices: Reddit's self-organised bull runs*

Valentina Semenova[†1,3] and Julian Winkler[‡2,3]

[1]*Department of Mathematics, University of Oxford*
[2]*Department of Economics, University of Oxford*
[3]*Institute of New Economic Thinking at the Oxford Martin School*

April 13, 2022

**Abstract**

This paper develops an empirical and theoretical case for how 'hype' among retail investors can drive large asset price fluctuations. We use text data from discussions on WallStreetBets (WSB), an online investor forum with over eleven million followers as of February 2022, as a case study to demonstrate how retail investors influence each other, and how social behaviours impact financial markets. We document that WSB users adopt price predictions about assets (bullish or bearish) in part due to the sentiments expressed by their peers. Discussions about stocks are also self-perpetuating: narratives about specific assets spread at an increasing rate before peaking, and eventually diminishing in importance – a pattern reminiscent of an epidemiological setting. To consolidate these findings, we develop a model for the impact of social dynamics among retail investors on asset prices. We find that the interplay between 'trend following' and 'consensus formation' determines the stability of price returns, with socially-driven investing potentially causing oscillations and cycles. Our framework helps identify components of asset demand stemming from social dynamics, which we predict using WSB data. Our predictions explain significant variation in stock market activity. These findings emphasise the role that social dynamics play in financial markets, amplified by online social media.

JEL codes: D91, G14, G41.

*[*]**Acknowledgements**: We are grateful to participants of the Royal Economic Society annual conference, the Alan Turing Institute interest group on economic data science, and the Rebuilding Macroeconomics seminar. We thank Rick Van der Ploeg, J. Doyne Farmer, Xiaowen Dong, Steve Bond, Francis DiTraglia, Kevin Sheppard, John Pougué-Biyong, Marteen Scholl, Will Wildi, Ilan Strauss, Jangho Yang, Torsten Heinrich, François Lafond, Matthias Winkler, José Moran, Farshad Ravasan, Pedro Bordalo, Cars Hommes, Stefan Zohren, Mirta Galesic, and Renaud Lambiotte for their helpful comments and insightful questions. We thank Baillie Gifford and the Institute for New Economic Thinking at the Oxford Martin School for funding our work at the University of Oxford.

[†]valentina.semenova@maths.ox.ac.uk
[‡]julian.winkler@economics.ox.ac.uk

# 1 Introduction

In investigating the stock market crash of May 28, 1962, the Securities and Exchange Commission (SEC) found that: 'investor "psychology" being what it is, the increasing decline in one or several issues can easily spread to others. Once the process becomes generally operative, the stage is set for a serious market break' (1963). The SEC concluded that large institutions acted as a balancing force during the collapse. The report pointed at retail traders as the *key players* behind the panic.

Over half a century later, we are again confronted with the consequences of investors' social behaviours. As online discussants on Reddit's 'WallStreetBets' (WSB) forum drove up the price of GameStop shares in January, 2021, retail investors regained a spotlight on the (virtual) trading floor. A key difference between 1962 and today is the internet, which offers both a coordination platform on an unprecedented scale, and a new datasource on investor narratives, interactions and psychology.

This paper sets out to reconcile observed behaviours on social media with economic theory. Economists have long deliberated to what extent social dynamics and human psychology play a role in economic decision-making (Shiller 1984, Black 1986), with Hirshleifer (2001) concluding that 'despite many empirical studies, scholarly viewpoints on the rationality of asset pricing have not converged'. Researchers have collected evidence of behaviours that undermine rational expectations theory and the Efficient Market Hypothesis, such as risk aversion (Kahneman & Tversky 1972, 1973, 1979), the perseverance of formed impressions (Ross et al. 1975, Anderson et al. 1980) and, more recently, peer effects (Bursztyn et al. 2014, Lahno & Serra-Garcia 2015, Banerjee et al. 2013). However, many questions still remain unanswered. Do these behaviours persist among investors evaluating opportunities in the 'wild'? If so, what are the impacts on financial markets? Should regulators be concerned about online forums turbo-charging 'irrational exuberance' (Shiller 2005)? We shed light on these questions, using data from the WSB investor forum as a case study for how retail investors behave, and hope to demonstrate concurrently the usefulness of online text data for understanding the social dynamics behind economic decisions.

This paper begins by documenting the existence of peer effects among retail investors. We find that the sentiments expressed by an investor's peers about an asset on WSB impact the future sentiment adopted by an investor, net of a set of asset, investor and stock market controls. We test for direct peer effects in investor decision-making in two ways. We evaluate the effect of peers on people who share their outlook about an asset multiple times, and interact with peers in-between. This approach allows us to precisely identify the change in individual sentiment due to peer interactions. We also pursue a network approach, leveraging interactions on WSB to identify a peer cohort. Both methods implement an Instrumental Variable (IV) approach to differentiate 'peer effects' from 'contextual effects' and 'correlated effects', and to tackle the 'common shock' problem (Zenou 2016).

Individuals appear to weight the opinions of others almost as heavily as their own prior opinion about an asset. The combination of both IV results suggests that a doubling in the odds of peers expressing bullish over bearish sentiments increases the odds of one given user to express bullish over bearish sentiment by 19%, on average. Asset returns also play a role in opinion formation – an increase in an asset's share price raises investor sentiment. Even though social contagion and the role that narratives play in investor decision-making have been heavily discussed in the literature, to the best of our knowledge, this is the first work documenting a causal relationship between an investor's sentiment and that of his peers, outside a controlled experimental setup (Bursztyn et al. 2014).

We also observe that asset interest is 'contagious' among peers. We follow the framework of Banerjee (1993) and Shiller (1984) in our analysis, and observe that their mechanism for information transmission fits the data well, explaining almost 28% of the variance in the sum of squares for the log-odds of an asset discussed on WSB in a given week.

This paper subsequently develops a theoretical framework for how social dynamics destabilise financial markets. Following the approach of Black (1986), Shefrin & Statman (1994), De Long et al. (1990), we divide the market into 'hype' investors, susceptible to social dynamics, and other traders who are not. Our key finding is that market stability depends on the interplay of two factors: the strength of social influence between investors, and the feedback between investor activity and asset returns. The extent to which hype investors pay attention to a historical, versus social, signal determines the presence of excess volatility and cycles in asset returns.

Our final section ties together our model, and empirical observations of market and WSB data. Our main result centers on predicting variation in 'consensus' and 'contagion' among WSB users unrelated to current price changes. This strategy works well because of the strong temporal persistence of sentiments and dedication to specific assets, due to peer effects and information transmission. Our estimates are economically and statistically significant in predicting changes in weekly average log returns, as well as changes in volatility and trading volumes. These results provide evidence for a causal relationship between social dynamics, proxied by WSB conversations, and financial markets.

The importance of peers and narratives in forming investor perspectives was famously highlighted by Shiller (1984), who provides statistical evidence of the greater volatility in stock prices than warranted by that of dividends. Since then, 'narrative economics' has played an increasingly important role in our understanding of investor decision-making and market moves (Shiller 2005, 2014, 2017, Banerjee et al. 2013, Hirshleifer 2020). Economists proposed impactful models for understanding how investors influence each other, with many studies focusing on information transmission in financial markets (Grossman & Stiglitz 1980, Barlevy & Veronesi 2000, Hellwig & Veldkamp 2009, Banerjee 1993, Cont & Bouchaud 2000). Simultaneously, psychologically-founded decision models have been introduced to explain deviations from rational expectations and expected utility theory, founded in 'prospect theory' (Kahneman & Tversky 1972, 1973, 1979, Gennaioli & Shleifer 2010, 2018, Bordalo et al. 2019). Despite the many important findings stemming from these areas of research, practical difficulties and a lack of data have largely restricted researchers to controlled laboratory experiments and a reliance on theoretical results, while their external validity unchecked (Barberis 2013). By leveraging new data, our work provides fresh empirical evidence of how heuristics, and peer effects in particular, affect investor decision-making.

Several impactful studies in the peer effects literature leverage naturally occurring variation in peers for their identification strategy. An area which pioneered many of these techniques investigates peer effects in the classroom (see Epple & Romano (2011), Sacerdote (2011) for a general overview, and Duflo et al. (2011) for a prominent example). Social networks are also an active area of study (see Bramoullé et al. (2020) for a recent review). Finally, techniques for correctly identifying peer impacts have shaped recent developments in this area of research (Angrist 2014, Blume et al. 2011). The present paper highlights how to transfer well-established techniques from this literature to social media data, thereby shedding light on investor psychology.

In a similar vein, Bursztyn et al. (2014) perform a field experiment with a financial brokerage in Brazil, where they study investment decisions made by peer pairs: the peers are offered a 'high stakes' investment opportunity (minimum investments were R$2,000 – around 50% of the median investor's monthly income) in a certain order to identify the effects of 'social learning' and 'social utility' in financial decision-making. Other related work investigates the diffusion of micro-finance decisions in a social network (Banerjee et al. 2013), the effect of peers on risk taking (Lahno & Serra-Garcia

2015), and the effect of social networks on saving (Breza & Chandrasekhar 2019). We distinguish ourselves from these and related works by studying a broader set of investors in a natural experiment, whose behaviours are unconstrained by the experimental setting. Our methods focus on exploiting the randomly occurring variation in investor groups to identify peer effects, rather than a controlled experiment. In this way, our research question is similar to Pool et al. (2015), who demonstrate that socially connected fund managers appear to hold similar stocks. In contrast, we are able to study a distinct class of retail investors active online. Moreover, our work contributes to the important question of the impact of social dynamics on financial markets, which is not tackled in the texts highlighted above.

Other works study the interplay between online forums and financial markets, as well as the spread of information in social networks. This paper differs from studies focusing on the spread of information through friend networks, such as Aral et al. (2009), Aral & Nicolaides (2017), since Reddit users are anonymous, without any explicitly defined friendship links. The anonymity within Reddit is crucial to the prominence of WSB: in contrast to the exercise in Banerjee et al. (2013), where information is transmitted via friendship networks, the mechanism by which information dissipates on WSB is much closer to the homogenous mixing conditions popular in traditional epidemiological models, and therefore closer in spirit to Banerjee (1993). Our work is distinct from the papers of Bailey et al. (2016), Sabherwal et al. (2011), Bollen et al. (2011), Kumar & Lee (2006), who focus on identifying one direct relationship between social activity and assets. In contrast, we focus on how investors decide strategies from observing their peers, and subsequently impact financial markets.

We present our results in five sections. The following section comprehensively describes the data source and relevant variables. Section 3 presents empirical evidence for investor social dynamics, including the identification of peer effects. Section 4 models how the observed social dynamics impact asset prices. Section 5 empirically evaluates the effect of retail investors on financial markets. Section 6 concludes.

## 2 What is *WallStreetBets*?

Reddit, launched in 2005, is a social news aggregation, web content rating, and discussion website. It was ranked as the $19^{\text{th}}$ most visited site globally in April 2021,[1] with over 430 million anonymous users by the end of 2019.[2] The website's contents are self-organised by subject into smaller subforums, 'subreddits', to discuss a unique, central topic.

Within subreddits, users publish titled posts (called 'submissions'), typically accompanied with a body of text or a link to an external website. These submissions can be commented and 'upvoted' or 'downvoted' by other users. A ranking algorithm raises the visibility of a submission with the amount of upvotes it receives, but lowers it with age. Therefore, the first submissions that visitors see are i) highly upvoted, and ii) recent. Comments on a submission, visible to anyone, are subject to a similar scoring system, and can, themselves, be commented on.

The WSB subreddit was created on January 31, 2012, and reached one million followers in March 2020.[3] As per a Google survey from 2016, the majority of WSB users are 'young, male, students that are inexperienced investors utilizing real money (not paper trading); most users have four figures in their trading account'.[4] The past conversation guidelines outlined by the moderators of WSB handily demonstrate the financial focus and whimsical tone of discussions:[5]

- Discussion about day trading, stocks, options, futures, and anything market related,
- Charts and technical analysis,

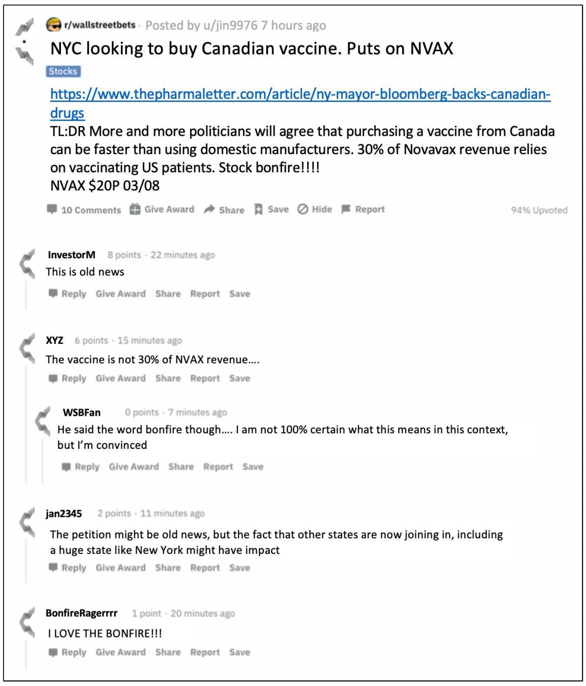


(a) A Typical Discussion on WSB

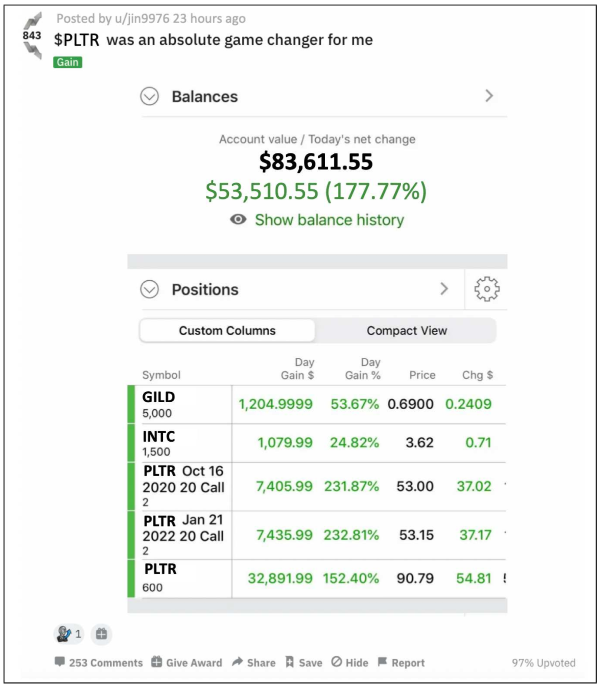


(b) A Sample Screenshot of User Profits

Figure 1: **What does WSB look like?** These snapshots display typical discussions on WSB. The exact text, usernames, and conversation details have been modified to protect user identities.

- Shower before posting,
- Some irresponsible risk taking,
- People sharing trades, ideas, observations.

Figure 1a displays a typical exchange on the WSB forum: individuals discuss stock-related news and their sentiments on whether this will affect stock prices in the future. In addition to market discussions, there is ample evidence of users pursuing the investment strategies encouraged in WSB conversations. Users post screenshots of their investment gains and losses, which subreddit moderators are encouraged to verify, as illustrated in Figure 1b. These observations are reminiscent of Shiller (2005) in his definition of an asset bubble (our emphasis):

> A situation in which news of price increases spurs investor enthusiasm which spreads by psychological contagion from person to person, in the process amplifying stories that might justify the price increases and bringing in a larger and larger class of investors, who, despite doubts about the real value of an investment, are drawn to it partly through envy of others' successes and partly through a *gambler's excitement*.

All posts made on Reddit, plus their metadata, can be queried via Reddit's API, as well as other sources. For this paper, we downloaded the entire history of WSB data using the PushShift API[6]. The only drawback of PushShift is that submissions are recorded at the time of their creation, and thus lack an up-to-date upvote/downvote count.

The full dataset consists of two parts. The first is a total of 1.4 million submissions, with their authors, titles, text and timestamps. The second is comprised of 16.5 million comments, with their authors, text, timestamp, and the identifier of the comment/submission they reply to. Figure 2 displays the evolution of WSB over time; submission and comment activity has grown exponentially

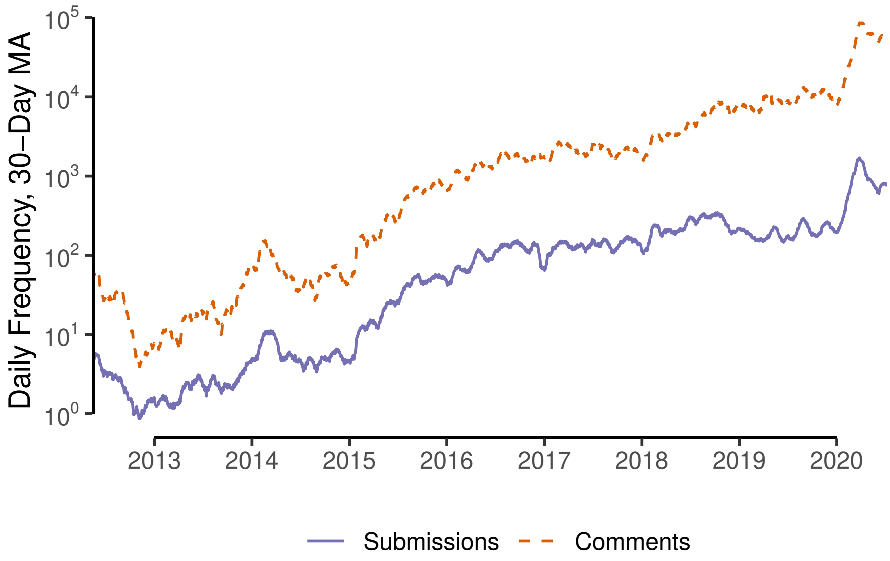


Figure 2: **Daily activity on WSB plotted on a logarithmic scale**; the daily submission and comment counts, averaged over 30 days, demonstrate a persistent exponential increase in activity on the WSB forum from 2015 to 2020, with a substantial jump in early 2020.

since 2015. Two jumps are notable: a smaller, seemingly idiosyncratic rise in early 2018, and a sharp spike during the COVID-19 pandemic.

Our dataset spans January, 2012 to July, 2020. Importantly, it does not include the events of the 2021 GameStop (GME) short squeeze. The decision to focus on this timeframe is intentional: before the GME short squeeze and widespread popularity of the forum, WSB received less attention from institutional investors, as well as less bot-activity. As such, our sample tracks retail investor discussions more precisely, without systematic external influence. Furthermore, ample research has emerged focusing exclusively on the GameStop short squeeze, such as Chohan (2021), Vasileiou et al. (2021), whereas our goal is to characterise investor behaviour, rather than examine a single event.

The following sections predominantly rely on submissions for text data, since they are substantially richer. Comments are used to trace user activity and, subsequently, the interactions between discussants. In order to understand how users discuss specific assets, we extract mentions of *tickers* from the WSB submissions' text data. A ticker is a short combination of capital letters, used to identify an asset on trading platforms. For example, 'AAPL' refers to shares in Apple, Inc. Appendix A.1 documents how tickers are extracted from submissions. Table 5 in Appendix A.1 displays the twenty tickers that feature most prominently in WSB conversations up to July, 2020. These are typically stocks in technology firms, such as AMD or FB. A handful of Exchange Traded Funds (ETF) are also present, notably the S&P 500 (SPY) and a leveraged gold ETF (JNUG). A small fraction of the 4,650 tickers we extract dominate the discourse on WSB: 90% of tickers are mentioned fewer than 31 times, and more than 60% are mentioned fewer than five times. Appendix A.1 documents the heavy-tailed nature of ticker discussions. In total, we are left with 111,765 submissions that mention one, unique ticker and were posted before July 1$^{st}$, 2020. These submissions have 1.9 million comments in total.

In addition to extracting tickers, we gauge whether submissions express an expectation for an asset's future price to rise, the *bullish* case, to fall, the *bearish* case, or to remain unpredictable/stable, the *neutral* case. Among other alternatives, we identify the sentiment using a supervised-learning approach, with a hand-labeled dataset of almost five thousand submissions for training, validation and testing. The sentiment model outputs a probability for each sentiment category, achieving 70% accuracy in categorising the manually labeled test set. Appendix A.2 discusses details of this Natural Language Processing (NLP) model.

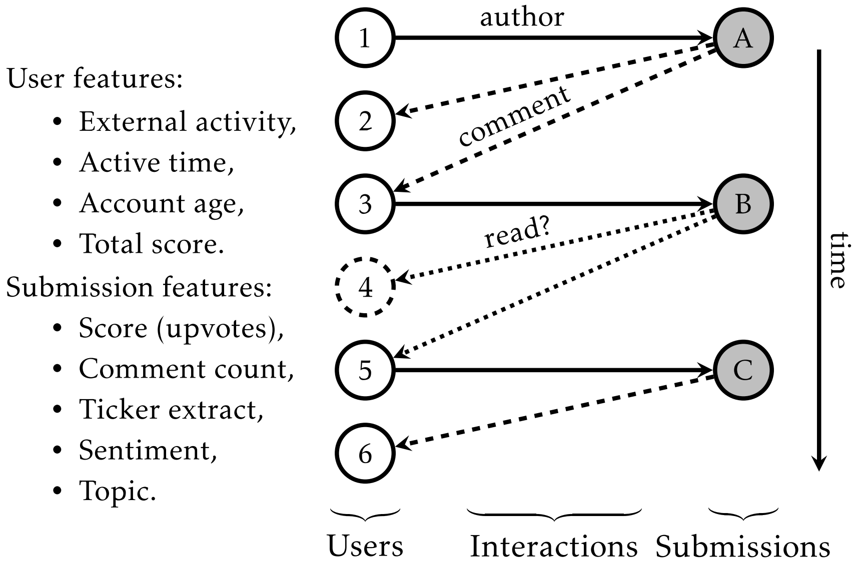


Figure 3: **Diagramatic representation of user interactions on WSB**; we encode how information on a specific asset propagates through WSB in a bipartite graph, with submissions on the right and users on the left. Dotted edges represent possible *read* interactions; a dashed line shows if the author comments on the submission. Solid edges denote users authoring submissions. Some of the relevant metadata available for each node type is listed on the left. Only users who author, or comment, on a submission are observed. Other users may not visibly interact with WSB content, but perhaps are exposed to it due to the time they are active on the forum, such as user **4**.

We estimate the spread of information among retail investors using sets of submissions mentioning the same asset. To illustrate the structure of our data, Figure 3 presents a bipartite graph of submissions, on the right, and users, on the left. We encode different types of relationships, such as node **1** authoring submission **A**. Other users, such as users **2** and **3**, comment on submission **A** and may, therefore, be influenced by it. We also consider the potential for authors to read and interact with submissions without commenting, such as the links from submission **B** to authors **4** and **5**. To proxy for these relationships, we consider the time at which authors are actively discussing the specific asset.

## 3 Social dynamics among retail investors

This section provides empirical evidence for the existence of peer effects among investors on WSB. We convey our main intuition using a game with strategic complementarities in retail investor decision-making (Zenou 2016). Strategic substitutes and complements were first proposed in Bulow et al. (1985) in the context of firms making production decisions – where an increase in production of one firm increases the marginal revenues of competitors. The framework was extended to information acquisition among agents in financial markets by Barlevy & Veronesi (2000), who argue, contrary to Grossman & Stiglitz (1980), that learning among investors can become a strategic complement.

Barlevy & Veronesi (2000) suggest that, as the fraction of 'informed' traders in an asset increases and prices, consequently, become more extreme, it may become harder, rather than easier, to identify the asset's payoff. We propose and empirically test a framework with strategic information complementarities among investors on WSB. These complementarities manifest in two ways: i) the sentiments expressed in relation to the future outlook of an asset, and ii) the choice of assets discussed.

### 3.1 Peer effects due to strategic complementarities

**Utility framework with strategic complementarities** Suppose that investor $i$ derives the following utility, $U_i$, from adopting sentiment $\phi_i$ about an asset:

$$U_i(\phi_i) = \phi_i \mathbb{E}_i(r) - \gamma \phi_i^2 \mathbb{E}_i\left[r - \mathbb{E}_i(r)\right]^2 + u_i. \tag{1}$$

where $\mathbb{E}_i$ denotes $i$'s expectation, $r$ is the asset's return, $\gamma$ is a scalar, and $u_i$ is some idiosyncratic error. Utility is increasing in the asset's expected return, $\mathbb{E}_i(r)$, and decreasing in its expected forecast error. An investor's expectation includes some signal $b_i$, with information both public and private to $i$, and the expressed sentiments of investors they interact with:

$$\mathbb{E}_i(r) = g(b_i) + f(\bar{\phi}_{-i}), \tag{2}$$

where $\bar{\phi}_{-i}$ is the average, revealed peer sentiment on the same asset. We assume that functions $g(\cdot), f(\cdot)$ are monotonic and continuous, but are specifically interested in the form of $f(\cdot)$. Substituting into Eq. 1, utility can be rewritten as

$$U_i(\phi_i) = \phi_i g(b_i) + \phi_i f(\bar{\phi}_{-i}) - \gamma \phi_i^2 \sigma_i^2 + u_i,$$

where $\sigma_i^2$ is the variance in the privately-formed signal $g(b_i)$. This is a well-known formulation for strategic interactions between agents acting under quadratic loss (Zenou 2016).

Definition 1 establishes the notion of complementary sentiments in this context.

**Definition 1.** *Sentiments are complementary if investor $i$'s utility is increasing in the sentiments of others, such that $f'(\bar{\phi}_{-i}) > 0$.*

Leveraging the framework above, we rationalise the broad following of WSB by Hellwig & Veldkamp's (2009) main result, which we rewrite in Proposition 1. In addition to the quadratic loss assumed in Eq. 1, we remind the reader of the additional assumptions in their framework, adapted to the present context.

**Assumption 1.** *Investor $i$ can observe $j$'s sentiment, $\phi_j$, at a cost $C(\omega_{ij})$, where $\omega_{ij}$ takes value one if investor $i$ interacts with investor $j$, and zero otherwise.*

**Assumption 2.** *Investors update their sentiments $\phi_i$ according to Bayes' Law.*

**Proposition 1.** *If expressed sentiments are complementary, then the value of additional information is increasing in the information acquisition of other investors.*

*Proof.* See Proposition 1 by Hellwig & Veldkamp (2009).

□

Proposition 1 predicts that interactions between investors whose decisions include a social component (a group we call 'hype investors') provide them with additional expected utility. This, in turn, increases their propensity to invest according to their updated sentiment, under costly information acquisition. We empirically validate whether Proposition 1 holds in the context of WSB and, thereby, study the social aspects of investing among retail traders.

**Target independent variable** The target independent variable of interest for studying hype investor sentiment is the log-odds of bullish over bearish sentiment,

$$\Phi_{i,t} = \frac{1}{2}\log\left(\frac{P(\phi_{i,t} = +1)}{P(\phi_{i,t} = -1)}\right) = g(b_{i,t}) + f(\bar{\phi}_{-i,(t-1,t)}) + \varepsilon_{i,t}, \tag{3}$$

derived from the utility framework in Appendix B.1. One key addition is the time subscript, $t$. An author chooses a bullish over bearish strategy depending on: i) a signal $b_{i,t}$, and ii) the observed sentiments of peers, $\bar{\phi}_{-i,(t-1,t)}$.

**Testable propositions** The following corollary and proposition state testable implications of investors affecting each other directly, through complementarities in expressed sentiments, as well as using publicly available signals to form their expectations.

**Corollary 1.1.** *If expressed sentiments are complimentary, a marginal increase (decrease) in peer outlook about an asset will raise (lower) the future outlook of an investor about the asset.*

*Proof.* Follows from Proposition 1 and the assumption that $f(\cdot)$ is monotonic and increasing.

□

In addition to testing for peer effects, retail investor sentiments should consider an asset's historical performance. A rich literature, founded in prospect theory (Kahneman & Tversky 1972, 1973, 1979), discusses the effect of past observations on decision-making in a setting with risk (Bordalo et al. 2012, 2019, 2020, Gennaioli & Shleifer 2010). Among others, Wang et al. (2006) demonstrate a Granger-causal relationship between returns, volatility and sentiment. Therefore, our framework should control for the investor's observations of historical returns as part of the signal set $b_{i,t}$. We formalise this in Proposition 2.

**Proposition 2.** *$g(\cdot)$ is increasing in stock returns. A uniform, marginal increase (decrease) in an asset's returns will raise (lower) the future outlook of an investor about the asset. It will also indirectly increase (decrease) the outlook of an investor through increasing (decreasing) the outlook of his peers.*

*Proof.* Follows from the assumption that $g(\cdot)$ is monotonic and continuous, the fact that the investor is Bayesian and from derivations in Appendix B.1. Investors derive their utility from their sentiment and privately formed signal: $\phi_i g(b_{i,t})$. Given the investor is Bayesian, their estimate of future returns will include, and be positive in, the stock's historic returns. The uniform increase ensures that there is no concurrent increase in the asset's volatility.

□

Corollary 1.1 and Proposition 2 provide the main empirical implications for the proposed model that can be tested with WSB and market data. In subsequent sections, we argue that the data are consistent with Corollary 1.1 and Proposition 2: observed sentiments are influenced by previous stock returns as well as peer sentiments. WSB, as a platform, is a venue for hype investors to realise their strategic information complementarities, explaining its exponential growth in userbase.

### 3.2 Empirical strategy: consensus formation among investors

We empirically validate Corollary 1.1 and Proposition 2 using data from WSB and, thereby, substantiate the proposed framework for social dynamics in investor decision-making. We use two approaches with distinct specifications for 'peers', which we label i) the *Frequent Posters* approach, and

ii) the *Commenter Network* approach. We subsequently introduce different instruments to isolate variation in peer sentiment, according to both definitions, orthogonal to current user sentiments.

The *Frequent Posters* approach leverages random, temporal variation in our data to identify peer influence. 8,173 authors create at least two submissions about the same ticker. We quantify peer influence by identifying the impact of other authors who write posts about the same asset *between* an individual's two submissions. This allows us to control for the author's sentiment prior to exposure to his peers, in addition to market moves. The choice of peers and peer sentiments have random variation due to the nature of Reddit: the time at which authors become active on the forum is random and independent of other authors, since, unlike other social media platforms, users on Reddit do not follow each other and are not alerted to each other's activity. Put differently, users do not explicitly select their peers – therefore, we avoid the selection bias present on platforms where individuals are exposed only to their chosen group of peers, who all hold potentially similar views. Given the properties of Reddit, we can assume that users are exposed to the opinions of peers through a process similar to poisson sampling, where the probability of exposure to a given peer is independent, and determined by the probability of the user to be online when the peer's post is highly visible. Therefore, an author's 'peer group' is randomly assigned, and there is some naturally occurring variation in the sentiments of peers about specific assets.

One potential concern with the proposed approach is that an exogenous shock could affect the sentiments of peers and the individual who posts multiple times, a phenomenon commonly referred to as the 'common shock' problem in the literature. We employ an instrumental variable (IV) to estimate the sentiments of peers. The use of an IV has become the gold standard to control for endogeneity in peer effects (Angrist 2014, Zenou 2016). The goal is to select an IV which is a good predictor of the independent variable, but unrelated to the dependent variable. The approximation technique involves two steps: first, predicting the independent variable of choice (the First Stage), and, second, using the *predicted* independent variable to estimate its relationship with the dependent variable (the Second Stage).

The *Commenter Network* approach considers a submission-to-submission network, with an earlier submission exerting peer influence on a future submission if the author of the later submission commented on the earlier one. In network science terminology, this network is a *projection* of the bipartite graph displayed in Figure 3. The submission-to-submission network helps more precisely identify peers an author actually interacts with, however, has certain drawbacks in terms of identification. Here, we also control for market variables, and employ an IV to address endogeneity concerns. The methods employed in constructing the independent and dependent variables are discussed in Appendix B.2.

### 3.2.1 Identifying peer influence – Frequent Posters

The random temporal variation in Reddit users allows us to test for peer effects. We identify the change in an author's outlook on an asset due to the expressed views of their peers. Author $i$ initially expresses a view about an asset $j$, $\phi_{i,j,(t-1)}$, and, subsequently, creates a new post about the same asset at a later time, with an updated sentiment $\phi_{i,j,t}$ (where time $t$ is in event time). In the time between these posts, the author observes submissions by others on the same asset, in addition to market moves. This structure allows us to control for any public and private signal an author receives prior to exposure to his peers, by directly observing prior sentiment $\phi_{i,j,(t-1)}$.

Our target variable is the log-odds of an author expressing bullish over bearish sentiment, $\Phi_{i,j,t}$, as detailed in Appendix B.2. A semi-supervised learning technique, explained in Appendix A.2, estimates the probability a given submission is bullish, $P(\phi_{i,j,t} = +1)$, or bearish, $P(\phi_{i,j,t} = -1)$. These probabilities are then transformed into our target variable, $\Phi_{i,j,t}$, as per Eq. 3.

We first estimate the effect of average peer sentiment between an author's two submissions with the following linear model:

$$\Phi_{i,j,t} = \kappa\bar{\Phi}_{-i,j,(t-1,t)} + X_{i,j,t}\beta + \epsilon_{i,j,t}, \tag{4}$$

where the vector of control variables, $X_{i,j,t}$, is composed of stock-specific fixed effects, author $i$'s past sentiment, and stock log returns, both on day $t$ and the average of the five days preceding $t$, and the variance in log returns on the five days prior to day $t$; $\beta$ is a vector of corresponding coefficients. Even though peers appear randomly on the forum in this formulation (as discussed earlier in this section), an exogenous shock in the period $(t-1,t)$ may affect the views of both peers and the author in question simultaneously. For this reason, the Ordinary Least-Squares (OLS) estimates do not enable us to make causal inference about peer influence.

To tackle this issue, we use the historical views of peers as an IV for their views expressed within $(t-1,t)$. Our choice of IV is founded in psychology: Ross et al. (1975) find that 'once formed, impressions are remarkably perseverant and unresponsive to new input, even when such input logically negates the original basis for the impressions', with many later studies, such as Anderson et al. (1980), supporting these findings. Our approach is similar to that employed in the peer effects literature on educational outcomes. For example, Duflo et al. (2011) use a student's admission test scores to predict their ability. We estimate investor $k$'s sentiment (a peer of investor $i$) about asset $j$, $\Phi_{k,j,t}$, based on the sentiment they expressed previously, $\Phi_{k,j,t-1}$:

$$\Phi_{k,j,t} = \kappa_0\Phi_{k,j,t-1} + \epsilon^0_{k,j,t}, \tag{5}$$

where $\epsilon^0_{k,j,t}$ is an idiosyncratic error, and $\kappa_0$ a coefficient. Eq. 5 is estimated using a sample containing submissions by all authors who post multiple times. The F-statistic for this first stage estimate, presented in Table 1, supports that this is a strong instrument. Our choice of IV gives a good approximation for author sentiment, while allowing us to control for common shocks affecting the sentiments of peers and investor $i$ in the period $(t-1,t)$.

We subsequently use the predicted outlook of peers between an author's posts to estimate peer effects as our Second Stage regression. Equation 4 is estimated by the instrumental variable: sentiments expressed by peers before $(t-1)$ are used to estimate $\bar{\Phi}_{-i,j,(t-1,t)}$. We keep the same control variables described in our OLS approach. In the Second Stage, the coefficient estimated for $\hat{\bar{\Phi}}_{-i,j,(t-1,t)}$, the hat denoting predicted peer sentiment, reflects the causal effect of peer sentiment on an author's expressed sentiment about an asset. The overall method is illustrated in Figure 12 in Appendix B.2.

**Credible estimation** We check whether our estimation strategy is credible, with respect to the three challenges presented by Zenou (2016) in estimating peer effects. The first lies is in distinguishing peer effects from contextual effects – the tendency of perspectives to vary with some observable characteristics of the group, rather than individuals influencing each other. Controls for asset price movements and ticker specific characteristics – the main sources of exogenous variation – address this concern. Second, the random, anonymous nature of WSB, as well as controlling for ticker-specific fixed effects, address the possibility for correlated effects. The specification with the IV addresses the common shock problem. A more rigorous, statistical analysis of our identification strategy is included with the results.

### 3.2.2 Identifying peer influence – Commenter Network

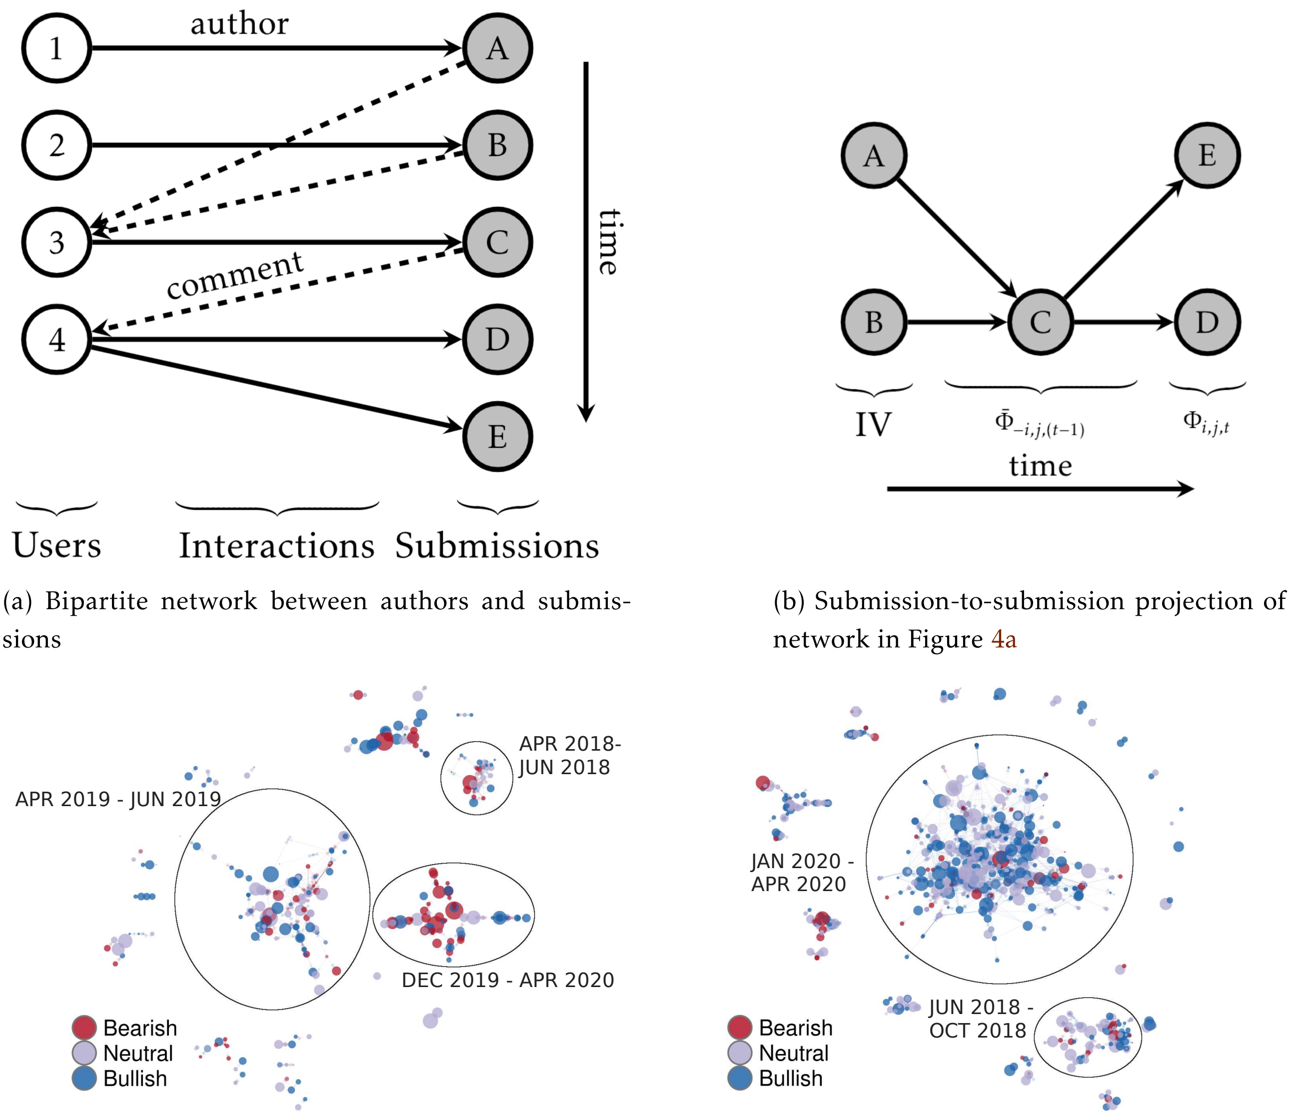


(a) Bipartite network between authors and submissions

(b) Submission-to-submission projection of network in Figure 4a

(c) Submission-to-submission network: DIS

(d) Submission-to-submission network: MSFT

Figure 4: **User networks in WSB conversations**; WSB data is summarised as a bipartite graph, illustrated in Figure 4a, where users (left) are linked to submissions (right) when they author the submission (solid edge) or comment on the submission (dashed edge). The resulting projection of submissions, in Figure 4b, tracks the propagation of sentiments $\Phi$. The submission-to-submission networks for two stocks in Figures 4c and 4d reveal that individuals post more submissions that are bullish(bearish) at times when the price of an asset increases(decreases) dramatically, with some visual evidence that similar sentiments tend to cluster.

WSB allows us to trace the interactions of users through a commenting network, even though there are no user friendship ties. We exploit a submission-to-submission interaction network for each asset, tracking which submissions in the past influence future submissions based on authors' commenting histories. This method offers a more precise way to identify an individual's peers by observing which individuals, and submissions, an author explicitly interacts with.

We build an example network in Figure 4a. User **3** comments on submissions **A** and **B** prior to creating their own submission, **C**. Therefore, the sentiments in **C** could be influenced by sentiments in **A** and **B**. User **4** comments on submission **C** before creating their own two submissions, **D** and **E**. Figure 4b visualises the resulting projection of this bipartite user-to-submission network, onto a submission-to-submission network. We place a directed edge from an earlier submission to a later one if the author of the later submission commented on the earlier submission. Submissions with the same author are not linked – an author's own previous submissions about a ticker are considered as a separate, independent variable in evaluating peer effects.

Two examples of submission-to-submission networks in our data are displayed in Figures 4c and 4d. Distinct temporal clusters emerge, as a certain asset gains and loses prominence on WSB. Some discussions appear fragmented: the DIS discussion in Figure 4c, for example, contains several smaller clusters, with distinct differences in overall sentiments. Others, such as the MSFT discussion in Figure 4d, contain a giant component where investors with different sentiments interact.

We use an IV approach to estimate peer influence. As the First Stage, we estimate the sentiments of neighbours to estimate an author's view. As indicated in Figure 4b, the sentiments in submissions **A**, **B** can be used to predict that of submission **C**. The *predicted* sentiment of **C** can then, in turn, be used to predict the sentiments of **D** and **E**. This choice of IV is well-established in the networks literature, and discussed by Zenou (2016). It controls for exogenous events that might affect both the individual and his peers. For the Second Stage of our network approach, the peer sentiment, $\hat{\Phi}_{-i,j,(t-1,t)}$, is the average *estimated* sentiment expressed within posts about the same ticker by a different author, that an individual has commented on prior to posting. We also modify our control for an author's past sentiment about the stock to account for authors who post for the first time: a dummy variable encodes whether the author's most recent previous post is bearish, neutral, bullish or missing.

The *Commenter Network* offers certain upsides, but also certain shortcomings, as compared to the *Frequent Posters* approach. The network method more precisely identifies the channels of influence between authors. However, the allocation of peers is no longer random, since the network structure is governed by a *choice* to comment on certain submissions over others. We use techniques outlined in Patacchini & Zenou (2016) and Bifulco et al. (2011) to address the three main challenges in correctly estimating peer effects through a network of interaction (Zenou 2016). The highlighted papers follow similar strategies; following their general approach allows us to paint a coherent picture of the influence of social interaction on retail investors' sentiments on WSB.

### 3.3 Results: Consensus Formation and Peer Effects

In this section, we present the Reduced Form, Second Stage, and First Stage regression estimates for both the *Frequent Posters* and *Commenter Network* approaches outlined above. The Reduced Form and Second Stage estimates, across both model specifications, show that peer sentiments directly impact an individual's sentiment about an asset, with the individual conforming to his peers.

Table 1 presents the results, with Panel A presenting OLS estimates for $\kappa$, from Eq. 4, using observed variation in peer sentiments, and Panel B.1 using predicted variation in peer sentiments. We relegate estimated coefficients for control variables to Appendix B.2. In the *Frequent Posters* approach, the estimated peer effects can be summarised as follows: an average estimate for $\kappa$ at 0.06, in the Reduced Form case, means that doubling in the odds of peers expressing bullish over bearish sentiments increases the odds of a given submission to be bullish, over bearish, by 4.2%, on average (we raise the log-odds estimate for $\Phi_{i,j,t}$ to an exponent). In the IV setting, an estimate of

Table 1: Peer influence in WSB sentiments

| | Frequent Posters | Network |
|---|---|---|
| | (1) | (2) |
| **Panel A**: Reduced Form – peer influence estimated using *observed* average sentiment of peers | | |
| | *Independent Variable* | |
| | Average peer sentiment, $\bar{\Phi}_{-i,j,(t-1)}$ (observed) | |
| Investor Sentiment ($\Phi_{i,j,t}$) | 0.06 (0.02) *** | 0.05 (0.01) *** |
| Author & asset controls ($X_{i,j,t}$) | Yes | Yes |
| Number of obs. | 11,129 | 24,963 |
| Adjusted $R^2$ | 0.06 | 0.06 |
| **Panel B.1**: Second Stage – peer influence estimated using *predicted* average sentiment of peers | | |
| | *Independent Variable* | |
| | Average peer sentiment, $\hat{\bar{\Phi}}_{-i,j,(t-1)}$ (predicted) | |
| Investor Sentiment ($\Phi_{i,j,t}$) | 0.19 (0.05) *** | 0.31 (0.07) *** |
| Author & asset controls ($X_{i,j,t}$) | Yes | Yes |
| Number of obs. | 11,122 | 24,963 |
| Adjusted $R^2$ | 0.06 | 0.06 |
| J-statistic | 0.00 | 0.00 |
| **Panel B.2**: First Stage – estimating peers' sentiments | | |
| | *Independent Variable* | |
| | Historical Sentiment of Peers | Sentiment of Neighbours' Neighbours |
| Sentiment of Peers | 0.33 (0.01) *** | 0.16 (0.01) *** |
| Author & asset controls ($X_{i,j,t}$) | No | No |
| Number of obs. | 19,814 | 27,472 |
| Adjusted $R^2$ | 0.11 | 0.02 |
| F-statistic | 2,430 | 467 |

*Notes*: This table presents the First Stage, Second Stage and Reduced Form OLS estimates for peer influence on WSB. In column (1), the First Stage is estimated using the initial sentiment expressed by an author about an asset to estimate his sentiment in the following post. In column (2), the First Stage is estimated using the sentiment of previous submissions that an author commented on, regarding the same asset. The Second Stage is estimated using the average predicted sentiment of peers. Ticker-level dummies, asset return and volatility controls, and the intercept are included in the Second Stage and Reduced Form estimates, but not shown here – the complete estimates are presented in Appendix B.2. Robust standard errors, clustered at the ticker level for Panels A and B.1, are presented in parentheses. Observations with incomplete data are dropped.
*** Significant at 1% level ** Significant at 5% level * Significant at 10% level

0.19 translates to an increase in the corresponding odds for bullish over bearish sentiments by 14.1%. In all cases, the robust standard errors, clustered at the ticker level, produce estimates statistically significant at the 1% level. The *Commenter Network* approach yields a similar result, with a higher estimated impact of peers in the IV setting. The higher estimates in the network setting may be attributable to the fact that individuals directly interact with their peers.

The estimated coefficients in columns (1) and (2) of Panel B.1 suggest that an exogenous increase in average peer outlook appears to increase an investor's own future view about an asset. These findings demonstrate that the data are consistent with Corollary 1.1. As a result, we conclude that the data support a model where strategic complementarities govern the investment decisions of retail traders on WSB. Estimates for asset price movements also uphold Proposition 2. These results are discussed further, alongside the complete regression estimates, in Appendix B.2.

**Support for identification** One potential concern is that individuals who post multiple times about the same asset, or those who comment on others' submissions, may differ from the rest of the population on the forum. If this were the case, our findings would not allow us to draw valid conclusions about the overall population of investors. We provide evidence that sentiments expressed by our samples are similarly distributed to those of the overall population in Appendix B.2.

A second concern is whether our proposed independent variables – asset price movements, ticker fixed effects and author historical sentiments – are effective controls for unobserved ticker characteristics. If our controls in the *Frequent Posters* formulation are valid, then a randomly selected cohort of individuals who post on the same ticker *before* the author's first post, should have no effect on the sentiments expressed in dependent submissions. Similarly, if our controls are useful in the *Commenter Network* formulation, a random rewiring of the network should yield no effect. The results are detailed in Appendix B.2: no statistically significant correlation of the randomly selected cohorts is produced. This provides further evidence that unobserved factors influencing within-ticker variation in both peer composition and author sentiment are not confounding.

A final concern with our approach is that historical peer sentiments or sentiments expressed by a neighbour's neighbours are systematically correlated with an author's opinion about an asset – in other words, that our instruments are not exogenous. We regress the residuals from our second stage on our instrumental variable and our endogenous regressors: a procedure known as the Sargan-Hansen test. J-statistics close to zero in both the *Frequent Posters* case and the *Commenter Network* case imply that the residuals are not correlated with our IV: historical sentiments of peers and those of an author's neighbours' neighbours are exogenous and provide valid instruments for estimation.

### 3.4 Contagion dynamics and the origin of bull runs

WSB users adopt the investment strategies of their peers in a specific asset, but do they also mimic each other's choice of assets? Previous research by Shiller (1984) and Banerjee (1993) emphasises the role that word-of-mouth information transmission plays in an investment context. In one empirical example, Banerjee et al. (2013) investigate the diffusion of microfinance by studying the uptake of financial policies by villagers in India, under peer influence, and the spread of awareness via influential leaders. In the WSB context, we would similarly expect awareness about specific assets to spread from one user to another. The emphasis of this section is not on identifying a causal relationship, but rather understanding the dynamics which govern asset interest among investors. These insights, combined with our understanding of how investors adopt each other's sentiments, allow us to paint a more complete picture of retail investors' decision-making and the resultant stock market dynamics.

We propose that there are two key parts to the contagion dynamics we seek to model. First, an individual must become aware of a stock, before discussing its prospects on WSB. The number of new, aware users discussing a stock is a function of the number of unaware users, and the likelihood that they become informed after interacting with aware users. Banerjee et al. (2013) tackles this in a similar way to the contagion rates used in epidemiological contexts. Second, if stock $j$ captures a user's attention in period $t$, the user will continue discussing it in the next time period $t+1$ if the stock continues to offer interesting investment prospects. The percentage of currently informed investors, as well as the stock's volatility and returns, all play a role in determining whether the stock holds an investor's attention.

We formalise these dynamics in a model for the log-odds of an author posting about stock $j$ over a baseline:

$$l(a_{j,t}) = \log\left(\frac{a_{j,t}}{s_t}\right) = h_1(a_{j,t-1}(1-a_{j,t-1})) + h_2(a_{j,t-1}) + h_3(\bar{r}_{j,t-1}) + h_4(\sigma^2_{j,t-1}) + \zeta_{j,t}. \tag{6}$$

The baseline is the probability $s_t$ of posting about a stock that is not widely discussed within the forum – a stock that is mentioned in fewer than 31 submissions within our sample. Our framework resembles that of Section 3.2 – individuals become interested in an asset because of their peers and because of a public signal of the asset's performance. The function $h_1(\cdot)$ captures the rate of independent mixing between investors aware of stock $j$, $a_{j,t-1}$, with unaware investors, $1-a_{j,t-1}$, creating newly informed investors. $h_2(\cdot)$ captures the rate at which aware investors remain aware and engaged in discussions between $t-1$ and $t$, versus seeking out other opportunities. Intuitively, functions $h_1(\cdot)$ and $h_2(\cdot)$ capture how *popular* an asset has been in the recent past. The latter terms control for the asset's perceived profitability and riskiness. $h_3(\cdot)$ is a 'quality of signal' term capturing how well the asset has performed in the past, and $h_4(\cdot)$ is a 'noise of signal' term with the asset's recent volatility. We propose that the log-odds of posting about asset $j$ in time period $t$ are increasing in $h_1(\cdot), h_3(\cdot)$ and decreasing in $h_2(\cdot), h_4(\cdot)$.

### 3.4.1 Empirical Strategy

To measure the impact of profitability and popularity of past discussions about a stock on the log-odds of it being discussed in the future, we produce OLS estimates for the linear model

$$l(a_{j,t}) = ca_{j,t-1}(1-a_{j,t-1}) + da_{j,t-1} + \beta_1\bar{r}_{j,t-1} + \beta_2\sigma^2_{j,t-1} + X_j\beta_4 + \zeta_{j,t}, \tag{7}$$

where $a_{j,t-1}$ is the share of all active investors who post about ticker $j$ at times $t-1$: $a_{j,t} \in [0,1]$ for all $j$ and $t$, $\bar{r}_{j,t-1}$ is the average log-return in $t-1$, and $\sigma^2_{j,t-1}$ is the variance of the same log-returns (these variables are mostly consistent with Section 3.2, and outlined separately in Appendix B.2). $X_j$ is a vector of stock dummies, and $t$ denotes time, in weeks. The choice to aggregate over weeks is done to address sparsity of submissions, especially pre-2017. Two more choices are made to tackle this: i) we restrict ourselves to a sample that spans January 2012 to July 2020, and ii) we categorise stocks mentioned fewer than 31 times since January 2012 into an 'other stocks' group, which forms our benchmark $s_t$. We also consider a different formulation where we test for the direct impact of historical peer sentiments and the interactions between historical sentiments and returns / volatility: $\bar{\phi}_{j,t-2}\bar{r}_{j,t-1}$, $\bar{\phi}_{j,t-2}$ and $\bar{\phi}_{j,t-2}\sigma^2_{j,t-1}$. This formulation allows us to evaluate whether WSB users are more likely to discuss a stock if the predictions of their peers have been accurate in the past.

### 3.4.2 Results

Our results, presented in Table 2, demonstrate that WSB users follow each other in their choice of investment instruments. There is strong evidence that the homogeneous mixing property explains

Table 2: Stocks discussed on WSB

| | *Dependent variable:* $l(a_{j,t})$ | | | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| $a_{j,t-1}(1-a_{j,t-1})$ | 83.49*** (8.20) | 100.20*** (9.15) | 46.30*** (5.33) | 57.90*** (5.13) |
| $a_{j,t-1}$ | −48.01*** (7.04) | −62.37*** (7.83) | −24.06*** (3.94) | −33.73*** (3.95) |
| $\bar{r}_{j,t-1}$ | 1.24*** (0.39) | | 1.36*** (0.42) | |
| $\sigma^2_{j,t-1}$ | −2.15*** (0.60) | | −0.96* (0.54) | |
| $\bar{\phi}_{j,t-2}\bar{r}_{j,t-1}$ | | 0.56 (1.09) | | 1.59 (1.10) |
| $\bar{\phi}^2_{j,t-2}\sigma^2_{j,t-1}$ | | −5.14** (2.19) | | −1.71 (1.53) |
| Constant | −3.89*** (0.01) | −3.88*** (0.02) | | |
| Ticker FE | No | No | Yes | Yes |
| Number of obs. | 13,184 | 6,429 | 13,184 | 6,429 |
| Adjusted $R^2$ | 0.28 | 0.36 | 0.10 | 0.14 |
| F-Statistic | 1,293.50 | 919.07 | 428.65 | 318.28 |

*Notes*: This table presents OLS estimates for the log-odds of users discussing stock $j$ in week $t$, over a collection of stocks that are mentioned fewer than 31 times. Explanatory variables include: the lag in the share of authors discussing $j$, $a_{j,t-1}$, the interaction with the share of authors not discussing $j$, $a_{j,t-1}(1-a_{j,t-1})$, as well as the lag in stock $j$'s weekly average log-return, $\bar{r}_{j,t-1}$, and variance, $\sigma^2_{j,t-1}$. In columns (2) and (4), the average log-return is multiplied by the two period lag in the average sentiment expressed among WSB submissions on stock $j$, $\bar{\phi}_{j,t-2}$, and the variance in log-returns by the same sentiment's square, $\bar{\phi}^2_{j,t-2}$. Columns (3) and (4) include stock-specific fixed effects. Accompanying standard errors, displayed in brackets, are clustered at the stock level, and calculated in the manner of MacKinnon & White (1985).
*** Significant at 1% level ** Significant at 5% level * Significant at 10% level

the uptake of new assets: using estimates in column (1), an increase in the share of authors discussing stock $j$ from 0.1 to 0.2 increases the ratio of authors discussing $j$ over 'other stocks' in the following week by approximately threefold. This is contrasted by an increase to 0.2 to 0.3, which prompts a decline in the ratio of authors discussing $j$ over 'other stocks' by 50% – the difference is driven by the large negative coefficient on $a_{j,t-1}$. This is strongly reminiscent of epidemic contagion models, adapated to the spread of narratives (Banerjee 1993, Shiller 2017).

When we consider the impacts of stock-specific variables in isolation, presented in columns (1),(3) in Table 2, volatility and returns appear to be leading factors for authors deciding what asset to discuss. Average, historical returns are statistically significant at the 1% level in columns (1) and (3), indicating that discussion sizes are stimulated by large, notably positive, returns. Examining the coefficient in column (3), a stock that experienced a 5% greater return in one week is the subject of about 7% more submissions than usual. Volatility appears to play a greater role in our formulation without ticker-specific effects, with its significance declining from column (1) to (3) – a factor perhaps explained by the choice of hype investors to overlook recent volatility in certain assets, but not others. Our alternative formulation presented in columns (2) and (4), estimating the effect of the correctness and consistency of past WSB predictions in an asset, appears to have limited significance for better explaining asset interest.

In summary, the composition of discussions exhibits strong temporal persistence. This is emphasised not just by the statistical significance of the model's coefficients, but also by our ability to explain 28% in the variation of assets discussed overall, and 10% of within-ticker variation. This is not only because WSB serves to turbo-charge the spread of new stock opportunities within the hype investor community, but also because the forum reinforces associated narratives when existing opportunities exhibit large, consistent and, from the users' perspective, predictable returns.

# 4 A system for social dynamics and asset prices

Section 3 studies social dynamics among retail investors, using data from WSB as a case study. An outstanding question is whether the observed behaviours have implications for the financial markets. Debates about the importance of investor psychology date back to Shiller (1984) and Black (1986), who labels 'noise' traders as those who form expectations deviating from 'rational' rules. These discussions sparked a new body of research, separating market participants into two categories: i) a fully rational type, and ii) those who deviate from rational expectations (Shefrin & Statman 1994, Shleifer & Summers 1990, De Long et al. 1990). Authors argue for specific behaviours in each investor category to, ultimately, derive predictions for market stability and price dynamics.

In subsequent years, researchers scrutinised the interactions of heterogeneous agents in stock markets, following many, sometimes adaptable, rules – see Lux (1998), but also the comprehensive review by Hommes (2021). In relation to social dynamics specifically, investor crowding and coordination is famously explored in Kirman's (1993) ant study. Avery & Zemsky (1998) discuss the asset price impact of herd behaviour, depending largely on the price charged by 'a competitive market maker' and the perceived informativeness of trades during periods of herding. Coordination among groups of investors in a type of network is discussed in Cont & Bouchaud (2000), albeit at much shorter time scales than explored within this paper.

A key feature scrutinised in this literature is the multiple equilibria that emerge in the presence of strategic complementarities. In Kirman's (1993) model, the rate at which ants conform in their exploitation of one food source over another determines whether a colony exploits two food sources simultaneously (one stable equilibrium), or one food source at a time under the threat of a shock, which causes a regime switch (two stable equilibria, and one unstable). Barlevy & Veronesi (2000) and Hellwig & Veldkamp (2009) also discuss how strategic learning complementarities among investors can lead to multiple equilibria.

Section 3 presents evidence that retail investors have the tendency to imitate each other's strategies online, showing that strategic learning complementarities exist among hype investors. The appearance of multiple equilibria in this setting and potential asset demand instability is, therefore, of particular relevance, especially in the aftermath of the infamous GameStop short-squeeze. This section is dedicated to better understanding the stock market dynamics that emerge when a certain fraction of social investors are present.

## 4.1 Modelling financial markets with social dynamics

We model the market impact of the social behaviours we observe in Section 3 by distinguishing between two types of agents – 'hype investors', who are susceptible to social forces, and 'other investors', who are not. Our model includes one endogenous variable, namely consensus in hype investors' strategies, and abstracts away from other sources of price variation.

Consider a market for one asset with two types of participants: 'hype' investors, who buy quantity $Y$, and 'non-hype' investors, who buy quantity $S$. Their total demand is equal to the number of shares outstanding, $Q = Y + S$. We propose a model in discrete time $t$ where market participants observe the behaviour of others, as well as the asset's returns, before updating their demand for the asset to maximise their expected utility. The goal is to study the stability of a market with social dynamics, given that hype investors are able to move price. The key endogenous variable is sentiment, which governs the propensity for an average hype investor to buy or sell the asset, thus shifting demand $Y$.

We first define the composition of hype investor demand in accordance with our observations in Section 3. Hype investor demand $Y_t$ is composed of i) the average buying/selling intensity across all active hype investors, $\phi_t \in [-1, 1]$, determined by individual sentiments, ii) the average purchasing power of an individual hype investor $M/p_t$, where $M$ is the capital held by an individual hype investor, and iii) the total number of hype investors $N$:

$$Y_t = \frac{M}{p_t} N \phi_t. \tag{8}$$

The key dynamics of interest are the strategic coordination observed among hype investors at a given point in time, captured in $\phi_t$.

We formalise a functional form for individual hype investor's buying / selling preferences $\phi_{i,t+1}$ using the framework established in Section 3. The utility an investor derives from a specific sentiment is a function of other hype investors' average buying intensity $\phi_t$, asset returns $r_t$, as well as an error $\epsilon_i$:

$$U_i(\phi_{i,t+1} = +1) = \epsilon_i^+ + \alpha\phi_t + \beta r_t - \gamma r_t^2, \quad U_i(\phi_{i,t+1} = -1) = \epsilon_i^- - \alpha\phi_t - \beta r_t - \gamma r_t^2. \tag{9}$$

where $\beta$ is a trend following component, formulated in Proposition 2, $\alpha$ is the extent of consensus among hype investors from Corollary 1.1, and $\gamma$ captures aversion to large swings in price. For simplicity, these parameters are assumed common to all investors $i$ in both the bullish and bearish case. We abstract away from cognitive biases that may, for example, drive greater herding behaviour in a downturn. Section 3 and Appendix B.2 offer estimates for these parameters using WSB sentiment data.

Here, $\epsilon_i^+$ and $\epsilon_i^-$ refer to a random, exogenous element in individual preferences. We make Assumption 3 about $\epsilon_i^+$ and $\epsilon_i^-$, which will allow us to find an expression for aggregate sentiment.

**Assumption 3.** *$\epsilon_i^+$ is independent and identically distributed, following a type-I Extreme Value (EV) distribution. The same holds for $\epsilon_i^-$.*

Assumption 3 is standard in the literature, and justifiable in our present context since investors make a choice governed by a maximisation process. Echoing Bouchaud (2013), individuals may be prone to error, either in the measurement sense, or due to random heterogeneity in preferences. The type-I EV distribution, also known as the Gumbel distribution, is one of three possible limiting distributions for the maximum of random variables. The overall shape is governed by a scale and location parameter. We denote the scale parameter by $2\lambda$, but set the location parameter to zero for simplicity.[7] Appendix C.1 discusses the validity of this assumption by fitting type-I EV distributions to WSB sentiments, where we find a value for location parameter $2\lambda$ of around two.

Using Assumption 3, we arrive at a standard 'quantal response' function from Eq. 9 to determine the likelihood of a bullish / bearish investment:

$$\phi_{i,t+1}^D = \begin{cases} +1, & \text{with probability } 0 < \frac{\exp[(\alpha\phi_t + \beta r_t)/\lambda]}{\exp[(\alpha\phi_t + \beta r_t)/\lambda] + 1} < 1, \\ -1, & \text{with probability } 0 < \frac{1}{\exp[(\alpha\phi_t + \beta r_t)/\lambda] + 1} < 1. \end{cases} \tag{10}$$

With slight abuse of notation, we define $\phi^D_{i,t+1}$ as the individual buy/sell decision taken by investor $i$. Given that $\lambda$ comes from the scale parameter of the error distributions, here it specifies the amount of noise in the decision-making process. Appendix C.1 details the derivation of the above expression. Work not shown here suggests that varying the noise component in our model does not alter our theoretical findings in a meaningful way.

In Proposition 3, we use Eq. 10 to give an expression for the evolution of the aggregate buying intensity by hype investors, as a function of returns and sentiments in the previous time step.

**Proposition 3.** *The average buying intensity across N hype investors is*

$$\phi_{t+1} = tanh\left[\frac{\beta r_t + \alpha\phi_t}{\lambda}\right]. \tag{11}$$

*Proof.* Aggregate buying intensity $\phi_{t+1}$ follows from Eq. 10;

$$\phi_{t+1} = \frac{1}{N}\sum_{i=1}^{N} \mathrm{P}(\phi^D_{i,t+1} = +1) - \mathrm{P}(\phi^D_{i,t+1} = -1) = \frac{\exp\left[(\alpha\phi_t + \beta r_t)/\lambda\right] - 1}{\exp\left[(\alpha\phi_t + \beta r_t)/\lambda\right] + 1} = \tanh\left[\frac{\beta r_t + \alpha\phi_t}{\lambda}\right].$$

□

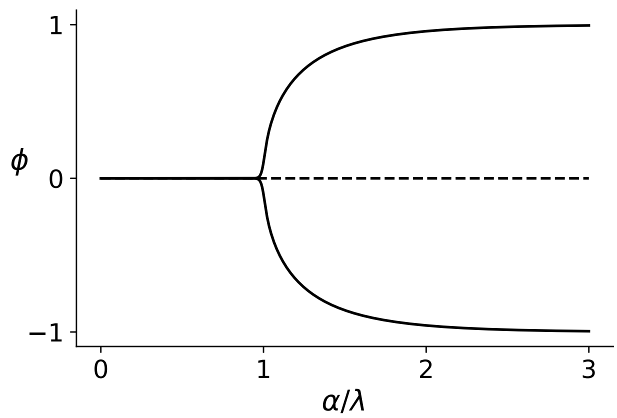


Figure 5: **Steady states of the hyperbolic tangent function**; the function bifurcates at $\alpha/\lambda = 1$. When $\alpha/\lambda < 1$, $\phi = 0$ is the unique, stable steady state, which loses stability once $\alpha/\lambda > 1$. Two new stable steady states emerge at that threshold, $\phi^+$ and $\phi^-$, for which we plot the numerical solution.

The function in Eq. 11 is known as the *hyperbolic tangent* function. Its properties have received significant attention in the area of strategic decision-making (Brock & Durlauf 2001, Bouchaud 2013), due to its varied behaviour for different values of $\alpha/\lambda$. When $\alpha/\lambda$ is between zero and one, the function produces one stable steady state for $\phi_{t+1} = \phi_t$ at zero. Two stable steady states, and one unstable one at zero, emerge when $\alpha/\lambda$ is greater than one, as shown in Figure 5. When the social signal $\alpha\phi_t$ is high relative to noise $\lambda$, investor coordination is effective and all hype investors converge on a similar investment strategy. On the other hand, if $\alpha/\lambda$ is low, then noise overwhelms the social signal – coordination fails and the zero steady state is the only one that is ever reached.

Figure 6 provides a visualisation of our hyperbolic tangent function, with respect to $\phi_t$, for three values of $\alpha/\lambda$. The points where the plotted function for $\phi_{t+1}$ crosses the 45 degree line constitute a steady state, since at these points $\phi_{t+1} = \phi_t$. When $\alpha/\lambda < 1$, the function crosses the 45 degree line in one place, corresponding to a single steady state at zero. Two additional steady states emerge

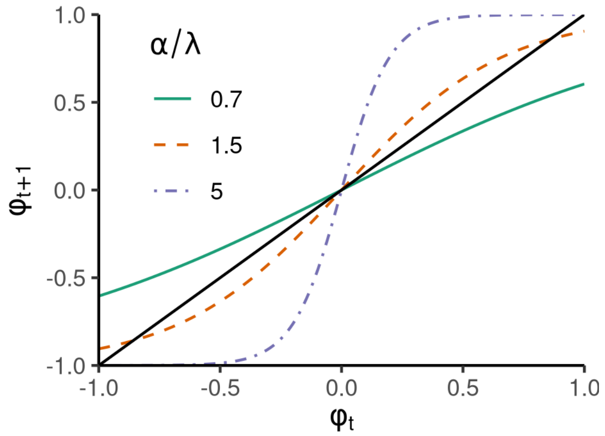


Figure 6: **Hyperbolic tangent function in $\phi_t$**; the hyperbolic tangent function, $\phi_{t+1} = \tanh(\alpha\phi_t/\lambda)$, is plotted with three different values of $\alpha/\lambda$. The lowest in solid green, with $\alpha/\lambda$ at 0.7, only crosses the 45 degree line (solid black) once at the origin. For values of 1.5, in dashed orange, and five, in dot-dashed purple, the hyperbolic tangent function crosses the 45 degree line in three places. When $\alpha/\lambda > 1$ the zero steady state is unstable whereas the two additional steady states present are stable.

when $\alpha/\lambda > 1$, as the plotted function crosses the 45 degree line in three places. The non-zero steady states lack a closed-form solution, so we denote the positive steady state by $\phi^+$, and the negative one by $\phi^-$. In a one-dimensional system, this function produces the well-known *pitchfork bifurcation*, presented in Figure 5. The significance of this and other types of bifurcations, such as the transcritical bifurcation, are discussed in Hommes (2013).

An additional step is necessary to link the pressure exerted by hype investors to market moves. We complete our model by introducing an expression for the effect of hype investor asset demand on returns, $r_{t+1} = \Delta p_{t+1}/p_t$, where $\Delta$ is the difference operator ($\Delta x_{t+1} \equiv x_{t+1} - x_t$). Recall that, under market clearing, the shares held by hype and non-hype investors must equal outstanding shares $Q$. For our purposes, we make a simplifying assumption that non-hype investors keep the nominal value invested in the asset fixed: $\Delta p_t S_t = 0$.

**Assumption 4.** *Non-hype investors keep the nominal value invested in the asset fixed,* $\Delta(p_t S_t) = 0$.

The key reasoning behind Assumption 4 is that it enables us to study the emergent price dynamics due to social dynamics in isolation. Shiller (1984) demonstrates that fully 'rational investors' (a category of investor we liken to non-hype investors in this paper) would anticipate persistent shifts in hype investors' demand for a stock, potentially justifying their willingness to pay a price different from that derived by discounting a stream of dividends. Our assumption is, therefore, somewhat reasonable as it implies that non-social investors hold the asset at every price point, and simply trade to keep their exposure constant. Using Assumption 4, we derive a simple expression for the dynamics of asset price returns, which we state in Proposition 4.

**Proposition 4.** *Given Eq. 8 for hype investor demand, and Assumption 4 for non-hype investor demand, asset price returns can be modeled as:*

$$r_{t+1} = \underbrace{\frac{MN}{p_t Q}}_{\text{(1) Capacity}} \times \underbrace{\Delta\phi_{t+1}}_{\text{(2) Demand}} . \tag{12}$$

*Proof.* Eq. 12 follows from the first difference of the market capitalization:

$$\Delta(p_{t+1}Q) = \Delta(p_{t+1}Y_{t+1}) + \Delta(p_{t+1}S_{t+1}) = MN\Delta\phi_{t+1}, \tag{13}$$

where equality holds because $\Delta(p_{t+1}S_{t+1}) = 0$ by Assumption 4, and Eq. 8 for $p_{t+1}Y_{t+1}$. Rearranging, and substituting $r_{t+1} = \Delta p_{t+1}/p_t$ yields the desired result.

□

Component (1) in Eq. 12, labelled 'capacity', is a market depth term, measuring how much market power hype investors hold, as a share of the total market capitalization of the asset. Term (2) captures the change in their aggregate buying intensity, as they respond to lagged returns and sentiments.

### 4.2 Market stability in the presence of social dynamics

Propositions 3 and 4 imply that, in the absence of any fundamental news, a market with hype investors is governed by the following dynamic system:

$$\phi_{t+1} = \tanh\left[\frac{\beta r_t + \alpha\phi_t}{\lambda}\right], \tag{14}$$

$$r_{t+1} = C(\phi_{t+1} - \phi_t), \tag{15}$$

where we simplify our initial expression for returns by fixing the capacity for hype investors $C = MN/p_tQ$, since our key dynamics of interest are the strategic coordination observed among hype investors at a given point in time, $\phi_t$, and its impact on asset returns.

**Assumption 5.** *Capacity $C = MN/p_tQ$ remains fixed.*

We consider a shift in the number of hype investors $N$ to be exogenous for the purposes of this section. We also hold the ratio $M/p_t$ constant: this implies that the balances hype investor have invested in the asset, $M$, increase and decrease proportionally with the asset's price. Any changes in price are met with proportional profits / decreases in capital $M$ – potentially due to gains from investments or moving capital elsewhere, allowing us to keep the ratio $M/p_t$ fixed. Assumption 5 allows us to study the impact of social dynamics in isolation.

Our goal is to gain insight into how the parameters governing trend-following, capacity and consensus affect price stability in our dynamic system, and to analytically determine regimes in which hype investors actively destabilise markets. Eqs. 14 and 15 are two non-linear, simultaneous difference equations, with sentiment $\phi_t$ and returns $r_t$ as state variables. They produces three steady states. In any steady state, returns $r_{t+1} = r_t$ must be zero. The steady state values for sentiments $\phi$ are governed by $\alpha/\lambda$, as per Figure 5: one at zero, then two more at $\{\phi^+, \phi^-\}$ when $\alpha/\lambda > 1$. This yields three steady states for the system as whole, which we note as (sentiment, return) tuples: i) $(0,0)$, ii) $(\phi^+, 0)$, and iii) $(\phi^-, 0)$.

Determining the steady states in a dynamic system is critical for understanding how the system evolves over time. As the name implies, when a dynamic system is in its steady state, it will remain there: in our case, sentiments and asset returns will remain stable. However, an important question still remains of how the system will behave if it is initiated outside of a steady state. For example, if hype investors observe an asset with non-zero returns or if a rumour rips through the hype investor population raising their individual assessments of the asset as an investment prospect. To tackle this, we examine the stability of each steady state individually.

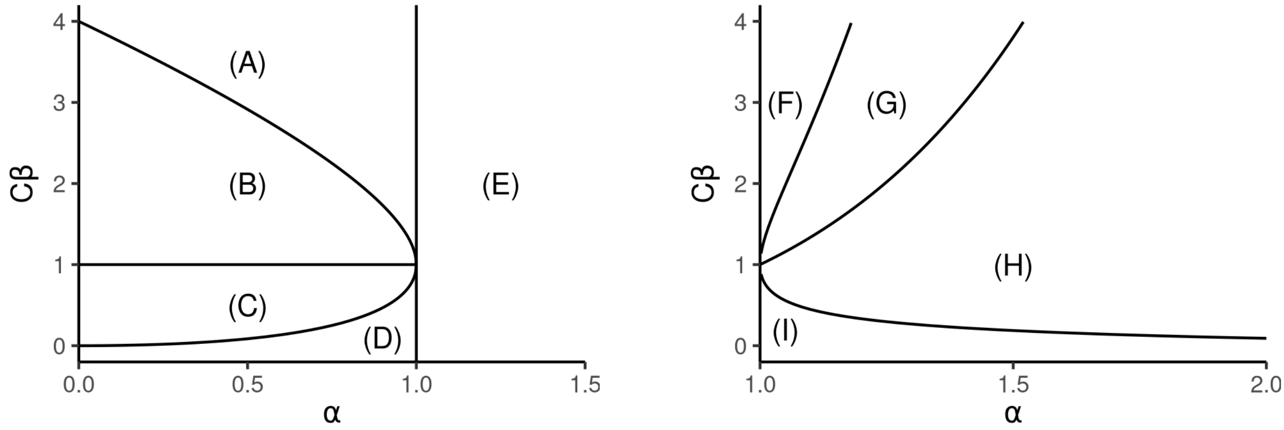


(a) $(0,0)$ (b) $(\phi^+,0)$ and $(\phi^-,0)$

Figure 7: **Stability regions of the steady states**; depending on the parameters, the dynamic system in Eqs. 14-15 for hype investor sentiment and asset price returns displays qualitatively different behaviours. The regions for different behaviours around steady state $(0,0)$ are marked on the left, and those for steady states $(\phi^+,0)$ and $(\phi^-,0)$, which coincide, on the right. The eigenvalues at $(0,0)$ are: outside the unit circle and real in (A), outside the unit circle and complex in (B), inside the unit circle and complex in (C), inside the unit circle and real in (D), inside and outside the unit circle in (E). The eigenvalues at $(\phi,0)$ are: outside the unit circle and real in (F), outside the unit circle and complex in (G), inside the unit circle and complex in (H), inside the unit circle and real in (I).

We show that the key governing parameters of our dynamic system are: i) the consensus parameter $\alpha$ – the degree to which investors' sentiments mimic those of their peers, and ii) the capacity parameter $C$. The degree of trend following $\beta$ appears significant, however we show that it has the same effect on the system's stability as capacity $C$. The noise parameter $\lambda$ can also be normalised to one without changing our results – we set $\lambda$ to one in what follows. We examine the stability of each steady state in turn, using the eigenvalues of the Jacobian evaluated at the steady states.

Figure 7 shows the different steady states that emerge within our system from the Jacobian, for different parameter values of $\alpha$ and $C\beta$ (with the precise derivation in Appendix C.2). Qualitatively different behaviours emerge within the different stability regions of Figure 7 – from cyclicality in Region A to direct convergence to the steady states in Region I. We show two examples of the behaviours we observe in Figures 8a-8b, which plot phase portraits of the system under two sets of parameters. In Figure 8a, all points for $\phi_t$ and $r_t$ converge to point $(0,0)$, given that $\alpha/\lambda < 1$. In contrast, the points in Figure 8b either converge to $(\phi^+,0)$ or $(\phi^-,0)$, depending on the initial value, since $\alpha/\lambda > 1$. As can be seen from the direction of the arrows, the precise combination of parameters will determine the stability of those steady states. Figure 8b posses a steady state at $(0,0)$, but no point, except for the steady state itself, converges to it since it is unstable. In other regions, we also observe cyclicality in returns and sentiments – Appendix C.2 contains plots of phase portraits in each of our stability regions of interest.

### 4.3 Model predictions

In the presence of hype investors, the interplay of 'trend following', 'capacity' and 'consensus formation' determine market stability. We consider the scenario closest to our observations in Section 3, where a lower rate of trend following and a social component are present. In this scenario, we

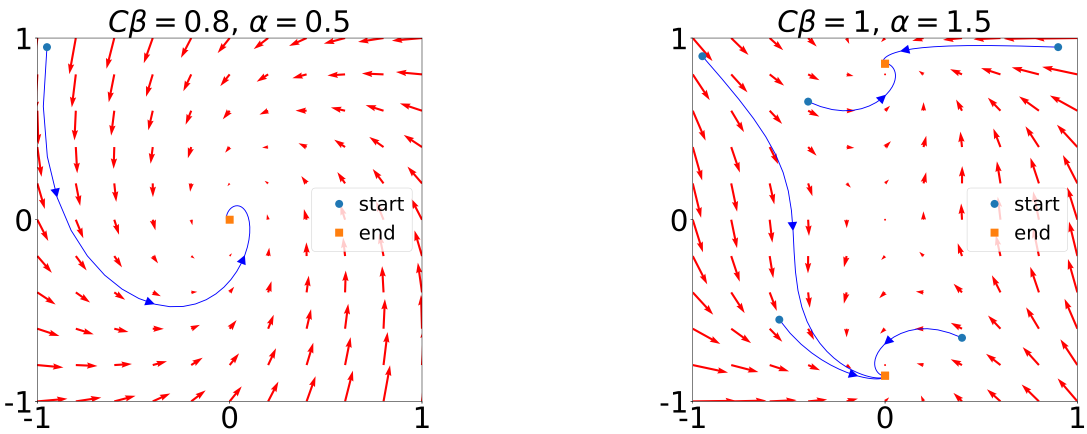


(a) One stable steady state.

(b) One unstable and two stable steady states.

Figure 8: **Phase portraits for sentiments and returns**; returns, on the horizontal axis, and sentiments, on the vertical axis, settle at one steady state in the left diagram, for a value of $\alpha < 1$. On the right, they settle at two steady states, depending on initial values, given $\alpha > 1$. Note that $\lambda = 1$, and hence does not feature. The trajectories for the system at different starting points are simulated for 200 time steps using a continuous time, Ordinary Differential Equation (ODE) approximation.

would expect oscillatory dynamics, as in Figure 8a. The lower trend-following component dampens an initial exogenous shock to price, however, over-shoots the steady state and returns oscillate before they converge. The social component influences the duration of oscillations we observe since, once consensus is reached, hype investors may be slower to change their minds, due to the impact of peer influence, than if they were simply following stock trends. In other regions, we also observe cyclicality in returns and sentiments.

An outstanding question is what oscillations and cyclicality mean for the distribution of returns. We simulate the dynamic system in Eqs. 14-15, by fixing $\alpha$ at 0.1 and 1.1, setting $\beta$ at 1 and varying $C$. We initialise sentiment $\phi_0$ to be zero, and randomly draw one hundred values for returns $r_0$ from a normal distribution with mean zero and variance 0.25. For a set of values of $C$, ranging from zero to four, we iterate forward $\phi_t$ and $r_t$ one thousand times. The final values for $\phi_t$ and $r_t$ are then plotted as a function of $C$ in Figure 9.

We observe that for values of $C$ below one, the system is stable around its unique steady state. For values of $C$ above one, sentiments, and concurrently returns, display quasi-periodic dynamics (in the clouded regions), interspersed with stable cycles (for example when $C$ is around 2). Stronger consensus dampens these instabilities, to an extent. This is seen in Figures 9c and 9d where $\alpha$ is set to 1.1. The effect of hype investors' larger capacity, in the region where $C$ is somewhat larger than one, is dominated by strong consensus, and the system settles at its steady state values, with either positive or negative sentiment. Appendix C.3 provides additional bifurcation diagrams which vary consensus parameter $\alpha$ instead of capacity $C$.

To gain more insight into asset volatility, we re-run the simulation with the same procedure as for Figure 9, but add a small stochastic term to returns in each iteration, drawn from a normal distribution with mean zero and variance $10^{-6}$. After running the system one thousand times for each chosen value of $\alpha$ (0.7 and 1.1) and $r_0$ of 0, we compute the standard deviation from the history of returns $r_t$ for each run, after discarding the first 50 values, then compute the average standard deviation for a given $\alpha$. The resulting estimates are plotted in Figure 10 with the solid green line

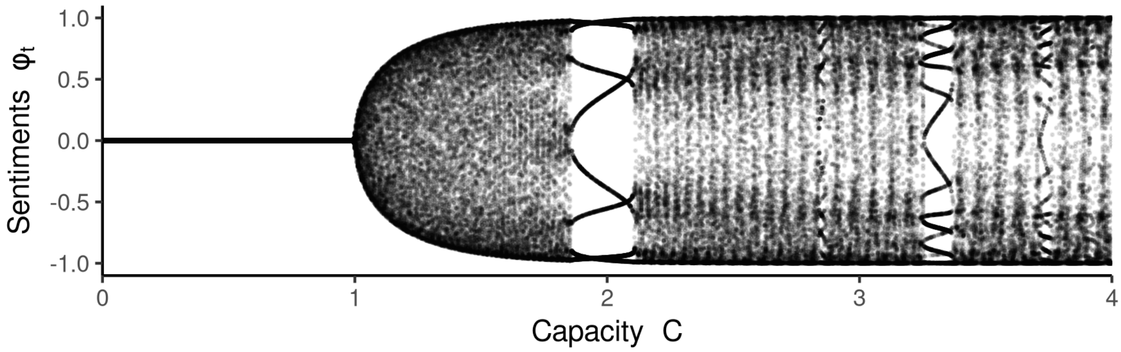


(a) $\alpha = 0.7$

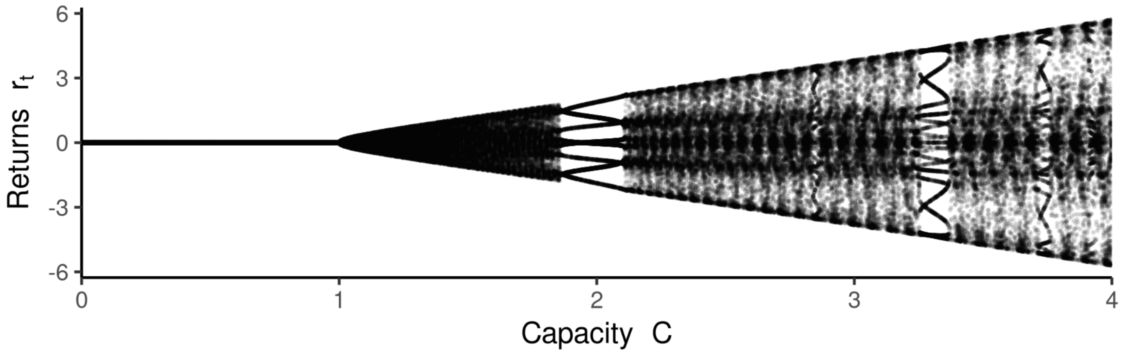


(b) $\alpha = 0.7$

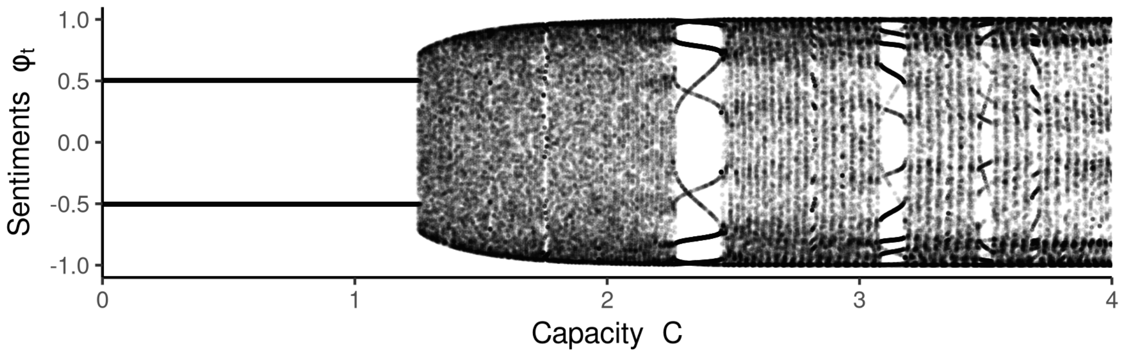


(c) $\alpha = 1.1$

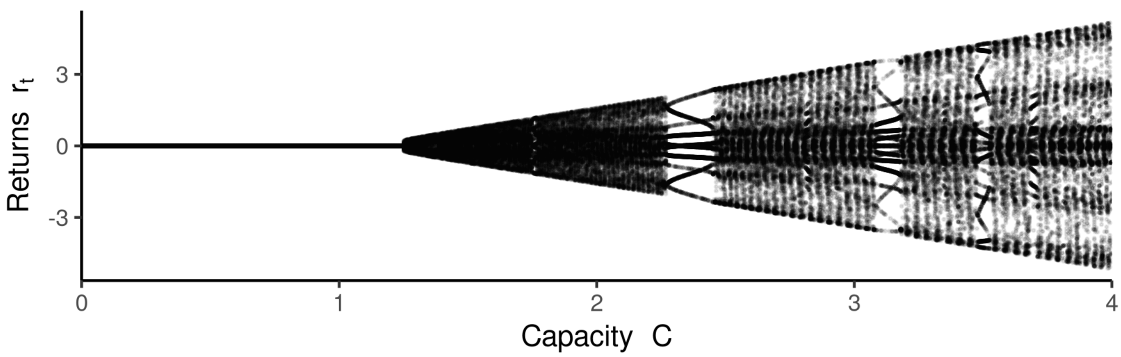


(d) $\alpha = 1.1$

Figure 9: **Bifurcation diagrams with respect to capacity parameter** $C$; the final values from one thousand iterations of the dynamic system in Eqs. 14-15, are plotted, for two values of $\alpha$ (0.7 and 1.1) and a range of values for $C$ between zero and four, keeping $\beta = 1$. One hundred initial values for $r$ are drawn from a normal distribution with mean zero and variance 0.25, with $\phi$ initiated at zero.

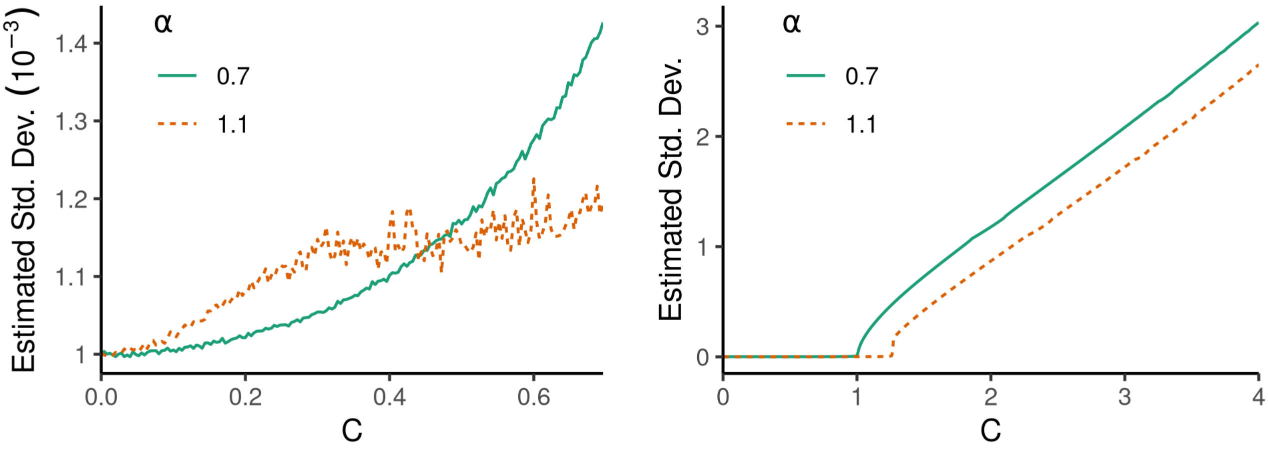


Figure 10: **Standard deviation of returns as a function of capacity**; the dynamic system in Eqs. 14-15 is simulated for 1000 steps, with initial values $\phi_0 = 0, r_0 = 0$, and parameters $\beta = \lambda = 1$, $\alpha = 0.7$ (in solid green) or $\alpha = 1.1$ (in dotted orange). Given the scale of the observed standard deviations, $C$ is drawn from a range of values between zero to 0.7 for the left chart, and zero to four for the right. A stochastic term is added to $r_t$ in each iteration, drawn from a normal distribution with mean zero and variance $10^{-6}$. The standard deviation of $r_t$ is computed separately for 100 runs of this procedure (after discarding the first 50 values) by value of $\alpha$. The average standard deviation of the 100 runs for each value of $\alpha$ is plotted as a function of parameter $C$.

denoting estimates when $\alpha$ equals 0.7, and the dotted orange line for $\alpha$ equal to 1.1; the left-hand figure highlights the region of lower volatility where $C < 1$.

On the left of Figure 10, the estimated standard deviations are more dispersed in the scenario where $\alpha = 1.1$, due to the complicated dynamics under three steady states. When $\alpha$ is set to 0.7, the effect of capacity on volatility is, on average, positive and monotonic. On the right of Figure 10, we observe that standard deviation in returns increases linearly with $C$ after a certain threshold. The threshold is higher for $\alpha$ of 1.1 than 0.7, perhaps highlighting that after a certain point, greater coordination among hype investors can result in greater stability.

The model is heavily simplified and fails to account for a number of important stabilising factors in real markets. It, therefore, should not be interpreted as a precise forecast of returns. However, it makes predictions on the qualitative relationship between social contagion, consensus, and asset prices. We leverage our observations from Section 3 to tie our findings to the markets. In Section 3, estimates for $\alpha$ are roughly around 0.2, and those for $\beta$ around 0.85. Given hype investor's limited capital[8], the dynamics we observe likely fall in regions (C) and (D) of Figure 7. On average, consensus is not strong enough for hype investors to sustain a fixed positive, or negative, position in perpetuity, and their capacity is unlikely large enough for prices to go out of control (unless we consider the special circumstances of the GameStop short-squeeze). As a result, in instances where WSB is able to garner large capacity, sentiments cycle around the zero steady state, increasing asset volatility. Periods of strong, positive sentiment are followed by negative returns, as consensus dissipates and investors close positions. We seek to substantiate these claims in Section 5.

# 5 Has WSB destabilised markets?

Our model provides motivation for testing the impact of consensus and contagion among WSB users on stock prices. From Section 3, we observe that certain changes of stock discussions on WSB can be explained by the temporal dynamics of hype investors' behaviours. Specifically, Section 3.4 shows that the popularity of stocks is predictable, and Section 3.2 demonstrates that the strategies (buy or sell) among WSB users are persistent. The goal of this section is to exploit our findings in Section 3 in order to establishes a causal relationship between social dynamics, proxied by WSB activity, and a set of stock market variables. We do not argue that WSB activity alone moves the markets, but rather that WSB gives us a meaningful sample of socially-driven retail investor activities, interests and behaviours. While our work in Section 3 guides our strategy for estimating hype investor activity, Section 4 establishes the structural relationship we expect to see between social dynamics and stock market variables.

Section 4.1 proposes that asset prices are affected by both coordination and consensus among hype investors. We observe this through the changes in behaviour of returns when varying capacity $C$, or the consensus parameter $\alpha$. Inspired by these findings, we empirically evaluate the following relationship motivated by Eq. 12:

$$r_t = w_1 \underbrace{\left(\frac{\phi_t}{pQ} \times \Delta A_t\right)}_{\Delta\Omega_t} + w_2 \underbrace{\left(\frac{A_t}{pQ} \times \Delta\phi_t\right)}_{\Delta\chi_t}, \tag{16}$$

where we no longer assume that the number of hype investors $N$ is constant and replace it by variable $A_t$; purchasing power $M$, from Eq. 12, is assumed constant. We separately construct a measure for contagion, denoted by $\Omega_t$, and consensus, by $\chi_t$, using historical data from WSB.

Our model for returns in the presence of hype investors from Section 4 allows us to make several predictions. Our main hypothesis is that a stock's price volatility increases when it becomes more popular or when consensus emerges on WSB, due to the cyclicality and oscillatory dynamics discussed in the previous section. This logic can be extended further to returns: oscillations and cyclicality imply that stock returns move away from the consensus sentiments on WSB. We, therefore, hypothesise that our consensus and contagion dynamics are negatively related to returns. Finally, increased attention from a new investor group is likely to drive up trading volumes: we, therefore, predict a positive relationship of consensus and contagion with asset trading volumes.

## 5.1 Key variables

**Dependent variables** Section 4.1 proposes a relationship between asset price and hype investor activity. We expand our analysis to consider three variables that summarise stock market activity: i) average returns, ii) variance in returns, and iii) trading volumes. Specifically, we consider the following target variables (where time $t$ is expressed in weeks):

- $\Delta\bar{r}_{j,t}$, the first-difference of stock $j$'s mean daily log-returns between calendar weeks $t$ and $t-1$,
- $\Delta\sigma^2_{j,t}$, the first-difference of the variance in stock $j$'s daily log-returns between weeks $t$ and $t-1$,
- $\Delta v_{j,t} = \Delta V_{j,t}/V_{j,t-1}$, the percent change in stock $j$'s average daily nominal trading volume between weeks $t$ and $t-1$.

Volumes are normalised to compensate for market capitalization heterogeneity among stocks. The relationship between returns and social dynamics hypothetically follows Eq. 16. Volatility and volumes may have a similar relationship, but do not discriminate the direction in which hype investors trade – buying and selling generate the same amount of price volatility and volumes.

**Independent variables** We construct two explanatory variables to evaluate the stock market impact of social dynamics. The first proxies contagion, $\phi_t \Delta A_t / pQ$:

$$\Delta\Omega_{j,t} = \frac{\phi_{j,t-1}}{q_{j,t-1}}(A_{j,t} - A_{j,t-1}), \tag{17}$$

where $A_{j,t}$ is the number of unique authors on WSB whose submissions solely mention stock $j$ in week $t$, $\phi_{j,t-1}$ is the average sentiment to buy ($\phi = +1$) or sell ($\phi = -1$) expressed among submissions in week $t-1$, and $q_{j,t-1} = p_{j,t-1} \times Q_{j,t-1}$ denotes $j$'s average market cap in week $t-1$.

The second proxies consensus formation, $\Delta\phi_t A_t / pQ$:

$$\Delta\chi_{j,t} = \frac{A_{j,t-1}}{q_{j,t-1}}(\phi_{j,t} - \phi_{j,t-1}). \tag{18}$$

Our two independent variables capture distinct components of hype investor demand. Contagion tracks the change in the number of investors interested in stock $j$ as an investment opportunity, fixing the prevailing sentiment. On the other hand, consensus gauges the change in asset demand due to changes in the intensity of hype investor sentiments that determine their investment strategies, keeping their number fixed. In both cases, we divide by the market cap of the relevant stock to control for market depth.

We adjust both independent variables when modelling volatility and volumes:

$$\Delta\Omega^*_{j,t} = \frac{|\phi_{j,t-1}|}{q_{j,t-1}}(A_{j,t} - A_{j,t-1}), \quad \Delta\chi^*_{j,t} = \frac{A_{j,t-1}}{q_{j,t-1}}(|\phi_{j,t}| - |\phi_{j,t-1}|). \tag{19}$$

Absolute values in average weekly sentiments are better suited, since the direction of the strategy, i.e. more bullish or more bearish, is less important as opposed to the absolute size of associated sentiments, which generate new trading activity.

Distinguishing between these two types of asset demand sheds light on the origin of bull runs. Does asset demand stem from information sharing, driving a growing numbers of amateur investors to potentially profitable opportunities, captured by $\Delta\Omega_{j,t}$? Or does coordination among existing users, $\Delta\chi_{j,t}$, who strategically reinforce each other's decision to enter a risky position, drive change in the markets? To answer these questions, we pursue a 2SLS approach.

## 5.2 Empirical Strategy

First, we formulate the linear relationship between our market variables and social dynamic variables of choice. We then propose an estimation strategy, following the findings in Sections 3 and 4.1. The key challenge in measuring the impact of consensus and contagion on stock market variables is the reverse causality with respect to price changes – sentiments have an impact on returns, but the reverse is also true.

**Reduced Form:** We regress changes in weekly log-returns, their variance, and percent change in average daily trading volumes on both measures for contagion and consensus:

$$\Delta\bar{r}_{j,t} = \beta_{\Omega,r}\Delta\Omega_{j,t} + \beta_{\chi,r}\Delta\chi_{j,t} + \eta_{r,t} + \varepsilon_{r,j,t}, \tag{20}$$

$$\Delta\sigma^2_{j,t} = \beta_{\Omega,\sigma}\Delta\Omega^*_{j,t} + \beta_{\chi,\sigma}\Delta\chi^*_{j,t} + \eta_{\sigma,t} + \varepsilon_{\sigma,j,t}, \tag{21}$$

$$\Delta v_{j,t} = \beta_{\Omega,v}\Delta\Omega^*_{j,t} + \beta_{\chi,v}\Delta\chi^*_{j,t} + \eta_{v,t} + \varepsilon_{v,j,t}, \tag{22}$$

where $\beta_\Omega$ and $\beta_\chi$ are coefficients of interest, $\eta_t$ denote time fixed effects, $\varepsilon_{j,t}$ an idiosyncratic error. Subscripts $r, \sigma, v$ serve to differentiate the coefficients, which we estimate separately for each dependent variables.

This Reduced Form setup does not allow us to argue that a causal relationship exists between hype investor activity and stock market activity: the narratives discussed and sentiments expressed at a given point in time are often shaped by real-time news and events.

**First Stage:** We use variation in our contagion and consensus measures, that can be explained by past activity on WSB, to identify our parameters of interest. In doing so, we assume that our target stock market variables are sufficiently uncorrelated between sequential trading weeks.

*Predicting* $\Delta\Omega$ – The number of new authors posting submissions is predicted using the model from Section 3.4:

$$l(a_{j,t}) = \log(a_{j,t}/m_t) = ca_{j,t-1}(1-a_{j,t-1}) + da_{j,t-1} + \beta_1 \bar{r}_{j,t-1} + \beta_2 \sigma^2_{j,t-1} + \eta_{a,t} + \varepsilon_{a,j,t}, \tag{23}$$

where we replace ticker fixed effects, in Eq. 7, with week fixed effects $\eta_{a,t}$. The results are used to predict the author count in the subsequent period, by using the previous week's author count in infrequent tickers, $V_{j,t-1}$:

$$\widehat{A}_{j,t} = V_{j,t-1}\exp\left(\widehat{l(a_{j,t})}\right), \quad \Delta\widehat{\Omega}_{j,t} = \frac{\phi_{j,t-1}}{q_{j,t-1}}\left(\widehat{A}_{j,t} - A_{j,t-1}\right), \quad \Delta\widehat{\Omega}^*_{j,t} = \frac{|\phi_{t-1}|}{q_{j,t-1}}\left(\widehat{A}_{j,t} - A_{j,t-1}\right), \tag{24}$$

where a hat denotes fitted values.

*Predicting* $\Delta\chi$ – In the same vein as in Section 3.2, we predict future sentiment using past stock price behaviour, as well as past sentiments:

$$\Phi^+_{j,t} = \log\left(\frac{\mathrm{P}(\phi_{j,t}=+1)}{\mathrm{P}(\phi_{j,t}=0)}\right) = \lambda^+_r \bar{r}_{j,t-1} + \lambda^+_\sigma \sigma^2_{j,t-1} + \lambda^+_1 \Phi^+_{j,t-1} + \lambda^+_2 \Phi^-_{j,t-1} + \eta^+_t + \varepsilon^+_{j,t}, \tag{25}$$

$$\Phi^-_{j,t} = \log\left(\frac{\mathrm{P}(\phi_{j,t}=+1)}{\mathrm{P}(\phi_{j,t}=0)}\right) = \lambda^-_r \bar{r}_{j,t-1} + \lambda^-_\sigma \sigma^2_{j,t-1} + \lambda^-_1 \Phi^+_{j,t-1} + \lambda^-_2 \Phi^-_{j,t-1} + \eta^-_t + \varepsilon^-_{j,t}, \tag{26}$$

where superscripts differentiate between the average log-odds of a submission in week $t$ expressing bullish (+) versus negative (−) sentiments, over neutral sentiments. Week fixed effects remain in the sentiment models, so that the full estimation strategy rests on within-week variation in all explaining, as well as explained, variables.

The approach outlined above relies on coarser aggregates for sentiments: the probabilities here are not estimated on data for individual submission sentiments, as is the case in Section 3.2. Rather, the probabilities are calculated by averaging the probabilities for *all* submissions in week $t$, discussing ticker $j$, to be bullish ($\mathrm{P}(\phi_{j,t}=+1)$), bearish ($\mathrm{P}(\phi_{j,t}=-1)$), or neutral ($\mathrm{P}(\phi_{j,t}=0)$). Predicted values for our consensus measure follow from the sentiment model predictions:

$$\widehat{\phi}_{j,t} = \frac{\exp\left(\widehat{\Phi}^+_{j,t}\right) - \exp\left(\widehat{\Phi}^-_{j,t}\right)}{1 + \exp\left(\widehat{\Phi}^+_{j,t}\right) + \exp\left(\widehat{\Phi}^-_{j,t}\right)}, \tag{27}$$

$$\Delta\widehat{\chi}_{j,t} = \frac{A_{j,t-1}}{q_{j,t-1}}(\widehat{\phi}_{j,t} - \phi_{j,t-1}), \quad \Delta\widehat{\chi}^*_{j,t} = \frac{A_{j,t-1}}{q_{j,t-1}}(|\widehat{\phi}_{j,t}| - |\phi_{j,t-1}|). \tag{28}$$

In all our estimates, we restrict ourselves to a sub-sample spanning January 2012 to July 2020. As discussed in Section 3.4, this choice serves to limit the amount of missing data in times when activity was relatively sparse.

### 5.3 Results

Table 3: First Stage estimates for consensus and contagion in WSB

| | *Dependent variable:* | | |
|---|---|---|---|
| | $l(a_{j,t})$ | $\Phi^+_{j,t}$ | $\Phi^-_{j,t}$ |
| $a_{j,t-1}(1-a_{j,t-1})$ | 87.05*** (9.77) | | |
| $a_{j,t-1}$ | −52.93*** (7.80) | | |
| $\bar{r}_{j,t-1}$ | 0.84* (0.45) | 0.17 (0.54) | −1.03* (0.57) |
| $\sigma^2_{j,t-1}$ | −1.89** (0.96) | −3.53*** (0.80) | −3.52*** (0.61) |
| $\Phi^+_{j,t-1}$ | | 0.09*** (0.02) | −0.06*** (0.01) |
| $\Phi^-_{j,t-1}$ | | −0.03** (0.01) | 0.17*** (0.01) |
| Week FE | Yes | Yes | Yes |
| Number of obs. | 6,429 | 7,154 | 7,154 |
| Adjusted $R^2$ | 0.32 | −0.004 | 0.02 |
| F Statistic | 793.90 | 18.80 | 55.84 |

*Notes*: column (1) of this table presents OLS estimates for the log ratio of the number of authors discussing stock $j$ in week $t$, over the number discussing stocks mentioned fewer than 31 times. The dependent variable in column (2) is the average log-odds of a given submission in week $t$ on stock $j$ to express bullish over neutral sentiment, and in column (3) – bearish over neutral sentiments. Explanatory variables include: the lag in the share of authors discussing $j$, $a_{j,t-1}$, the interaction with the share of authors not discussing $j$, $a_{j,t-1}(1-a_{j,t-1})$, the average log-return $\bar{r}_{j,t-1}$, and the variance in log-returns $\sigma^2_{j,t-1}$. The logit-transformed sentiments are regressed on the lag of the weekly mean and variance of log-returns, as well as the lag in logit-transformed sentiments. Each specification includes week-specific fixed effects. Accompanying standard errors, displayed in brackets, are clustered at the stock level, and calculated in the manner of MacKinnon & White (1985).
*** Significant at 1% level ** Significant at 5% level * Significant at 10% level

Table 3 helps assess the instruments' strength in predicting contagion and consensus on WSB. In column (1), lags in author shares discussing stock $j$, plus its lagged returns and volatility, are used to predict the share of authors discussing the same stock. This setup is identical to column (4) of Table 2, except in the use of week fixed effects instead of stock fixed effects. The high F-statistic, as well as adjusted $R^2$, confirm that these are good instruments. In columns (2) and (3), we find that lagged weekly mean and variance in returns, combined with lagged sentiments, are significant predictors for the current log-odds in submissions expressing bullish and bearish sentiments.

Table 4 presents our main results. Panel A regresses changes in average returns, volatility and nominal trading volumes against *observed* measures for contagion $\Delta\widehat{\Omega}_{j,t}$ and consensus $\Delta\widehat{\chi}_{j,t}$ in WSB data. Panel B in Table 4 presents causal evidence for the impact of consensus formation and contagion among hype investors on stock market variables, using *predicted* user counts and sentiments from the model presented in Table 3. We do not argue that WSB alone affects the markets, but rather that WSB data offers a rich sample of retail investor behaviour. The observed behaviours and dynamics within WSB, in turn, allow us to estimate future stock market activity.

The fitted consensus measure, $\Delta\widehat{\chi}_{j,t}$, displays a strong, negative correlation with changes in average weekly returns in column (1). Sentiments gravitating from 0.5 to 1 on average decrease a stock's weekly return by 8 percentage points (pp), controlling for the previous week's discussion size and the stock's market cap. This estimate is statistically significant at the 1% level when using predicted

Table 4: Market impact of consensus and contagion in WSB

**Panel A**: Reduced Form relationship between WSB and market activity

| | *Dependent variable:* | | |
|---|---|---|---|
| | $\Delta \bar{r}_{j,t}$ | $\Delta \sigma^2_{j,t}$ | $\Delta v_{j,t}$ |
| | (1) | (2) | (3) |
| $\Delta \chi_{j,t}$ | $-0.22^{**}$ (0.09) | | |
| $\Delta \Omega_{j,t}$ | $-0.01$ (0.11) | | |
| $\Delta \chi^*_{j,t}$ | | $-0.19^{***}$ (0.05) | $-16.02^{***}$ (3.93) |
| $\Delta \Omega^*_{j,t}$ | | $0.29^{***}$ (0.06) | $23.09^{***}$ (5.52) |
| Week FE | Yes | Yes | Yes |
| Number of obs. | 7,110 | 7,110 | 7,110 |
| Adjusted $R^2$ | $-0.01$ | 0.02 | 0.16 |
| F-Statistic | 70.31 | 199.99 | 824.79 |

**Panel B**: structural relationship between WSB and market activity

| | | | |
|---|---|---|---|
| $\Delta \widehat{\chi}_{j,t}$ | $-0.16^{***}$ (0.02) | | |
| $\Delta \widehat{\Omega}_{j,t}$ | $-0.001$ (0.003) | | |
| $\Delta \widehat{\chi}^*_{j,t}$ | | $0.02^{***}$ (0.01) | $0.39^{*}$ (0.24) |
| $\Delta \widehat{\Omega}^*_{j,t}$ | | $0.003^{*}$ (0.001) | $0.44^{**}$ (0.18) |
| Week FE | Yes | Yes | Yes |
| Number of obs. | 4,737 | 4,737 | 4,737 |
| Adjusted $R^2$ | $-0.01$ | $-0.05$ | 0.01 |
| F-Statistic | 95.42 | 8.55 | 156.84 |
| J-statistic | 4.255 | 4.51 | 7.003 |

*Notes*: this table presents OLS estimates for stock $j$'s change in average log-return, $\Delta \bar{r}_{j,t}$, change in variance of log-returns, $\Delta \sigma^2_{j,t}$, and percent change in nominal trading volume, $\Delta v_{j,t}$, in week $t$. We filter the sample to stocks mentioned in at least 30 distinct submissions on WSB, and exclude any ETFs. Explanatory variables include a measure for consensus formation, $\Delta \chi_{j,t}$, which tracks the change in average sentiments on WSB, fixing the number of users discussing stock $j$, and contagion, $\Delta \Omega_{j,t}$, which counts the number of new users discussing stock $j$, fixing their sentiment. Both variables are also divided by the market cap of stock $j$. In computing OLS coefficients for volatility and percent changes in trading volumes, both independent variables use the absolute value in sentiments, and are denoted by an asterisk. Each specification includes week-specific fixed effects. Accompanying standard errors, displayed in brackets, are clustered at the stock level, and calculated in the manner of MacKinnon & White (1985). Panel A computes the coefficients using values directly from WSB data, whereas Panel B employs sentiments and stock discussion predicted by past sentiments, stock discussions, as well as returns and return volatility, for which results are in Table 3. The associated J-statistics are recorded at the bottom of Panel B, which are computed by regressing the residuals from the Second Stage on all variables used for predicted $\Delta \widehat{\chi}_{j,t}$ and $\Delta \widehat{\Omega}_{j,t}$.
*** Significant at 1% level ** Significant at 5% level * Significant at 10% level

sentiments. The measure with absolute measures for sentiment, $\Delta \widehat{\chi}^*_{j,t}$, is also a strong and more statistically significant predictor for increases in weekly levels of return variance, in column (2). The coefficient on percent change in trading volumes in column (3) is positive, but only significant at the 10% level. One limitation is that our instrumental variables account for a tiny amount of the variance of weekly average sentiments, as indicated by the $R^2$ in Table 3. This likely contributes additional noise to parameter estimates presented in column (1) of Panel B of Table 4.

Our predicted contagion measure, $\Delta\widehat{\Omega}_{j,t}$, does not have a significant impact on changes in weekly average returns. It's counterpart, with absolute values for sentiment, $\Delta\widehat{\Omega}^{*}_{j,t}$, is a statistically significant predictor for volatility in weekly returns and associated trading volumes. Holding constant a stock's market cap and associated sentiments, one new user's choice to discuss the asset, given their exposure to previous conversations, increases the variance by 0.3pp. The large coefficient on our contagion measure, when predicting volumes, is likely a result of the fact that user numbers are a good proxy for the attention a stock garners among the greater retail investor community.

The results in Panel B of Table 4 stand out next to the Reduced Form estimates for the OLS coefficients from Eqs. 20-22, presented in Panel A. For these estimates, we use the observed, rather than predicted, measures of $\Delta\Omega_{j,t}$, $\Delta\chi_{j,t}$. First, the coefficient on consensus is only statistically significant at the 5% level, indicating that real-time sentiments vary more than those predicted by our consensus formation mechanism. The negative coefficient is surprising, but likely a result of the reaction of other market participants or the fact that asset prices move quickly and in the opposite direction to WSB users. Second, column (2) indicates that sentiments are negatively correlated with changes in volatility. This agrees with our finding that hype investors on WSB exhibit strong risk aversion, and tend to express neutral sentiments when volatility is high. In contrast, our predicted sentiments are positively correlated with volatility, thus supporting the view that consensus formation, orthogonal to market volatility, leads to greater asset price volatility. Looking at column (3), we reach a similar conclusion for trading volumes.

We check the validity of the instruments for consensus and contagion by the J-statistic, reported at the bottom of Table 4. It is calculated as an F-statistic for the hypothesis that the instruments are jointly equal to zero when regressing them against the residuals of the Second Stage. Their low score demonstrates insufficient evidence to reject the hypothesis that the instruments are jointly significant in explaining the residuals for each target stock market variable. This offers statistical support to our identifying assumption that stock market returns and trading volumes are largely uncorrelated from one week to the next.

Clearly, social dynamics among retail investors play an important role in the stock market today. We successfully pin down two channels by which they exert influence: consensus, by which retail investors coordinate a buying or selling strategy as a group, and contagion, by which information on opportunities can be communicated to growing crowds. To the disappointment of this paper's title, we do not pin down instability onto WSB alone. However the forum offers a window into the minds of an expanding number of retail investors, and it is clear that social media plays a key role in their growing clout on the virtual trading floor.

## 6 Conclusion

The importance of narratives and psychology for explaining investor behaviour is a long-standing debate in economics (Shiller 1984, Hirshleifer 2001). In this paper, we contribute to the discussion by documenting the influence of peers on investor opinion formation, using data from WSB as a case study: namely, we observe the existence of i) *consensus* and ii) *contagion* among retail investors. We test for consensus formation by measuring the effect that peer sentiments have on an individual investor's outlook using two different IV approaches. Our results consistently suggest that a causal link exists between the sentiment that an individual investor adopts and that of his peers. User sentiments are, on average, 19% more likely to be bullish rather than bearish, if the odds of peers expressing bullish over bearish sentiments double. Individuals also tend to focus their attention on stocks that others are discussing, underscoring the presence of a contagion mechanism on WSB.

Many impactful studies consider market dynamics, in the presence of rational and irrational investors (Black 1986, Shefrin & Statman 1994, Shleifer & Summers 1990). Drawing upon our empirical findings about peer influence, we model the dynamics of asset returns in the presence of socially-driven hype investors. We find that the stability of the system depends on the degree of consensus formation, and the ability for hype investors to move prices. Depending on the system parameters, sentiments and returns converge, cycle or even diverge. In the regions which are closest to our empirical observations on investor behaviour, the system converges to a stable steady state, with oscillatory dynamics. We thus hypothesise that social dynamics have a transient, but destabilizing impact on markets.

In order to validate our theoretical findings, we directly investigate the impact of social dynamics on assets using WSB data. We do not argue that the social dynamics among WSB users alone drive these results, but rather that the WSB forum allows us to sample broader retail investor sentiments and strategies. We predict 'consensus' and 'contagion' among WSB users using historical forum and market data as our IV. Our predicted 'consensus' measure has a statistically significant correlation with market returns, volatility and trading volumes. Sentiments gravitating from neutral ($\phi = 0$) to bullish ($\phi = +1$) among hype investors on average decreases a stock's weekly returns by 16pp. Though the negative relationship between sentiment and returns is surprising, we explain this by the oscillatory dynamics in our model. This result is also potentially influenced by other institutions trading against retail investors. An absolute change in sentiments by one, for example from neutral ($\phi = 0$) to bearish ($\phi = -1$), increases a stock's volatility by 2pp and trading volumes by 39%. A change in sentiments appears to cause greater uncertainty in an asset's value, echoing the proposed framework of Barlevy & Veronesi (2000). Our predicted contagion metric is strongly predictive of trading volumes – an intuitive relationship since trading activity would likely increase as new investors enter the market.

These findings are of particular importance to how we view the efficiency of financial markets and investor rationality. Our findings favour the study of narratives in financial markets (Shiller 1984, 2005). The important role that peers and asset returns play in formulating investor opinion points to strategic complementarities in information acquisition among investors, underscoring the importance of work by Barlevy & Veronesi (2000) and Hellwig & Veldkamp (2009). Our conclusions complement the earlier experiments of Bursztyn et al. (2014), showing that peer influence plays a role in a broader investment context than that of the authors' experimental setting.

Social media has changed the fabric of society. 4.2 billion people, or 53.6% of the world population, are active social media users, each just a few clicks away from the next popular phenomenon.[9] Now, a growing audience turns to social media for promising stock market gambles. Whether social media has been a boon or a bane to society is a highly contested topic, however, there is little debate over the fact that increased social media usage has generated new and exciting datasets for research. The rich text data from WSB can be used to verify economic behavioural theories: from narrative economics Shiller (1984), to models of 'noise trader' behaviour (Shefrin & Statman 1994, De Bondt & Thaler 1985, Black 1986) and information diffusion among investors (Banerjee et al. 2013, Hellwig & Veldkamp 2009, Barlevy & Veronesi 2000). It also presents new opportunities to leverage techniques in other areas of economics for understanding investor behaviour, such as the well-established peer effects literature (Duflo et al. 2011, Epple & Romano 2011, Sacerdote 2011, Angrist 2014, Bramoullé et al. 2020).

Evaluating economic theory is of great importance of its own right, however, we must also remember that the financial markets do not exist in isolation: investor decisions have broader implications on the cost of capital. A growing body of research considers how social media impacts society, with polarization, the spread of fake news and other societal challenges being some of its documented consequences (Tucker et al. 2018). In light of our findings, a relevant question may be how these consequences may impact capital allocation. Does the government have a right to track and penalise the spread of misinformation about asset prices on social media, despite its mandate to defend free speech? Do we need to think carefully about the power of social media personalities, given their potential to destabilise the financial system?[10] Perhaps it is an important time to re-evaluate regulatory structures within the financial system, which currently closely monitor financial institutions and large players, but overlook smaller investors.

With the first publicly acclaimed victory of Main Street over Wall Street, in the form of the GameStop short squeeze, it is unlikely that socially-driven asset volatility will simply disappear. In fact we observe the opposite: WSB grew from approximately 1.8 million users at the start of January 2021 to over eleven million users in February 2022.[11] The safety of the retail investor is emerging as a prominent concern. The excitement of potential gains has attracted a greater number of individuals to online forums and new trading platforms, with many offering incentives to lure in the unsophisticated trader. Institutional investors have also become keenly aware of the ties between retail investor social dynamics and the markets, with profits to be made from influencing online investor discussions. At this junction, it is perhaps more important than ever to consider our findings in the broader economic context, and ask what insights economic theory can offer to ensure financial stability and prosperity at a time when social media is rapidly changing the investment landscape.

# Notes

[1] https://www.alexa.com/topsites

[2] https://redditblog.com/2019/12/04/reddits-2019-year-in-review/

[3] https://subredditstats.com/r/wallstreetbets

[4] https://andriymulyar.com/blog/how-a-subreddit-made-millions-from-covid19

[5] https://www.reddit.com/r/wallstreetbets/

[6] https://pushshift.io/

[7] Since the means of $\epsilon_i^+$ and $\epsilon_i^-$ are zero when we fit the model to data, the location parameters actually take a negative value. As seen in Proposition 3, this constant would cancel out when subtracting the respective probabilities for bullish or bearish sentiments, so we may just set it to zero without altering the final outcome.

[8] https://andriymulyar.com/blog/how-a-subreddit-made-millions-from-covid19

[9] https://datareportal.com/reports/digital-2021-global-overview-report

[10] https://www.theguardian.com/technology/2022/jan/14/dogecoin-value-soars-after-elon-musk-says-it-will-be-accep

[11] https://subredditstats.com/r/wallstreetbets

## References


Anderson, C. A., Lepper, M. R. & Ross, L. (1980), 'Perseverance of social theories: The role of explanation in the persistence of discredited information.', *Journal of personality and social psychology* **39**(6), 1037.

Angrist, J. D. (2014), 'The perils of peer effects', *Labour Economics* **30**, 98–108.

Araci, D. (2019), 'Finbert: Financial sentiment analysis with pre-trained language models', *arXiv preprint arXiv:1908.10063* .

Aral, S., Muchnik, L. & Sundararajan, A. (2009), 'Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks', *Proceedings of the National Academy of Sciences* **106**(51), 21544–21549.

Aral, S. & Nicolaides, C. (2017), 'Exercise contagion in a global social network', *Nature communications* **8**(1), 1–8.

Avery, C. & Zemsky, P. (1998), 'Multidimensional uncertainty and herd behavior in financial markets', *American economic review* pp. 724–748.

Bailey, M., Cao, R., Kuchler, T. & Stroebel, J. (2016), Social networks and housing markets, Technical report, National Bureau of Economic Research.

Banerjee, A., Chandrasekhar, A. G., Duflo, E. & Jackson, M. O. (2013), 'The diffusion of microfinance', *Science* **341**(6144).

Banerjee, A. V. (1993), 'The economics of rumours', *The Review of Economic Studies* **60**(2), 309–327.

Barberis, N. C. (2013), 'Thirty years of prospect theory in economics: A review and assessment', *Journal of Economic Perspectives* **27**(1), 173–96.

Barlevy, G. & Veronesi, P. (2000), 'Information acquisition in financial markets', *The Review of Economic Studies* **67**(1), 79–90.

Bifulco, R., Fletcher, J. M. & Ross, S. L. (2011), 'The effect of classmate characteristics on post-secondary outcomes: Evidence from the add health', *American Economic Journal: Economic Policy* **3**(1), 25–53.

Black, F. (1986), 'Noise', *The journal of finance* **41**(3), 528–543.

Blume, L. E., Brock, W. A., Durlauf, S. N. & Ioannides, Y. M. (2011), Identification of social interactions, *in* 'Handbook of social economics', Vol. 1, Elsevier, pp. 853–964.

Bollen, J., Mao, H. & Zeng, X. (2011), 'Twitter mood predicts the stock market', *Journal of Computational Science* **2**(1), 1 – 8.

Bordalo, P., Gennaioli, N. & Shleifer, A. (2012), 'Salience theory of choice under risk', *The Quarterly journal of economics* **127**(3), 1243–1285.

Bordalo, P., Gennaioli, N. & Shleifer, A. (2019), 'Diagnostic expectations and stock returns', *The Journal of Finance* **74**(6), 2839–2874.

Bordalo, P., Gennaioli, N. & Shleifer, A. (2020), 'Memory, Attention, and Choice*', *The Quarterly Journal of Economics* **135**(3), 1399–1442.

Bouchaud, J.-P. (2013), 'Crises and collective socio-economic phenomena: simple models and challenges', *Journal of Statistical Physics* **151**(3), 567–606.

Bramoullé, Y., Djebbari, H. & Fortin, B. (2020), 'Peer effects in networks: A survey', *Annual Review of Economics* **12**, 603–629.

Breza, E. & Chandrasekhar, A. G. (2019), 'Social networks, reputation, and commitment: evidence from a savings monitors experiment', *Econometrica* **87**(1), 175–216.

Brock, W. A. & Durlauf, S. N. (2001), 'Discrete choice with social interactions', *The Review of Economic Studies* **68**(2), 235–260.

Bulow, J. I., Geanakoplos, J. D. & Klemperer, P. D. (1985), 'Multimarket oligopoly: Strategic substitutes and complements', *Journal of Political economy* **93**(3), 488–511.

Bursztyn, L., Ederer, F., Ferman, B. & Yuchtman, N. (2014), 'Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions.', *Econometrica* **82**(4), 1273.

Chohan, U. (2021), 'Counter-hegemonic finance: The Gamestop short squeeze', *SSRN Electronic Journal* .

Cont, R. & Bouchaud, J.-P. (2000), 'Herd behavior and aggregate fluctuations in financial markets', *Macroeconomic Dynamics* **4**(2), 170–196.

De Bondt, W. F. & Thaler, R. (1985), 'Does the stock market overreact?', *The Journal of finance* **40**(3), 793–805.

De Long, J. B., Shleifer, A., Summers, L. H. & Waldmann, R. J. (1990), 'Noise trader risk in financial markets', *Journal of political Economy* **98**(4), 703–738.

Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K. (2018), 'Bert: Pre-training of deep bidirectional transformers for language understanding', *arXiv preprint arXiv:1810.04805* .

Duflo, E., Dupas, P. & Kremer, M. (2011), 'Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in Kenya', *American Economic Review* **101**(5), 1739–74.

Epple, D. & Romano, R. E. (2011), Peer effects in education: A survey of the theory and evidence, *in* 'Handbook of social economics', Vol. 1, Elsevier, pp. 1053–1163.

Garcia, D. (2013), 'Sentiment during recessions', *The Journal of Finance* **68**(3), 1267–1300.

Gennaioli, N. & Shleifer, A. (2010), 'What comes to mind', *The Quarterly journal of economics* **125**(4), 1399–1433.

Gennaioli, N. & Shleifer, A. (2018), *A crisis of beliefs*, Princeton University Press.

Gentzkow, M., Kelly, B. & Taddy, M. (2019), 'Text as data', *Journal of Economic Literature* **57**(3), 535–74.

Grossman, S. J. & Stiglitz, J. E. (1980), 'On the impossibility of informationally efficient markets', *The American economic review* **70**(3), 393–408.

Hellwig, C. & Veldkamp, L. (2009), 'Knowing what others know: Coordination motives in information acquisition', *The Review of Economic Studies* **76**(1), 223–251.

Hirshleifer, D. (2001), 'Investor psychology and asset pricing', *The journal of Finance* **56**(4), 1533–1597.

Hirshleifer, D. (2020), 'Presidential address: Social transmission bias in economics and finance', *The Journal of Finance* **75**(4), 1779–1831.

Hommes, C. (2013), *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*, Cambridge University Press.

Hommes, C. (2021), 'Behavioral and experimental macroeconomics and policy analysis: A complex systems approach', *Journal of Economic Literature* **59**(1), 149–219.

Kahneman, D. & Tversky, A. (1972), 'Subjective probability: A judgment of representativeness', *Cognitive psychology* **3**(3), 430–454.

Kahneman, D. & Tversky, A. (1973), 'On the psychology of prediction.', *Psychological review* **80**(4), 237.

Kahneman, D. & Tversky, A. (1979), 'Prospect theory: An analysis of decision under risk', *Econometrica* **47**(2), 263–291.

Kirman, A. (1993), 'Ants, rationality, and recruitment', *The Quarterly Journal of Economics* **108**(1), 137–156.

Kumar, A. & Lee, C. M. (2006), 'Retail investor sentiment and return comovements', *The Journal of Finance* **61**(5), 2451–2486.

Lahno, A. M. & Serra-Garcia, M. (2015), 'Peer effects in risk taking: Envy or conformity?', *Journal of Risk and Uncertainty* **50**(1), 73–95.

Lux, T. (1998), 'The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions', *Journal of Economic Behavior & Organization* **33**(2), 143–165.

MacKinnon, J. G. & White, H. (1985), 'Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties', *Journal of Econometrics* **29**(3), 305–325.

Patacchini, E. & Zenou, Y. (2016), 'Social networks and parental behavior in the intergenerational transmission of religion', *Quantitative Economics* **7**(3), 969–995.

Pool, V. K., Stoffman, N. & Yonker, S. E. (2015), 'The people in your neighborhood: Social interactions and mutual fund portfolios', *The Journal of Finance* **70**(6), 2679–2732.

Ross, L., Lepper, M. R. & Hubbard, M. (1975), 'Perseverance in self-perception and social perception: biased attributional processes in the debriefing paradigm.', *Journal of personality and social psychology* **32**(5), 880.

Sabherwal, S., Sarkar, S. K. & Zhang, Y. (2011), 'Do internet stock message boards influence trading? evidence from heavily discussed stocks with no fundamental news', *Journal of Business Finance & Accounting* **38**(9-10), 1209–1237.

Sacerdote, B. (2011), Peer effects in education: How might they work, how big are they and how much do we know thus far?, *in* 'Handbook of the Economics of Education', Vol. 3, Elsevier, pp. 249–277.

Securities & Commission, E. (1963), 'Report of special study of securities markets of the securities and exchange commission', *House Document* **95**(4).

Shefrin, H. & Statman, M. (1994), 'Behavioral capital asset pricing theory', *Journal of financial and quantitative analysis* **29**(3), 323–349.

Shiller, R. J. (1984), Stock prices and social dynamics, Technical Report 2, The Brookings Institution.

Shiller, R. J. (2005), *Irrational Exuberance: (Second Edition)*, Princeton University Press.

Shiller, R. J. (2014), 'Speculative asset prices', *American Economic Review* **104**(6), 1486–1517.

Shiller, R. J. (2017), 'Narrative economics', *American Economic Review* **107**(4), 967–1004.

Shleifer, A. & Summers, L. H. (1990), 'The noise trader approach to finance', *Journal of Economic perspectives* **4**(2), 19–33.

Tetlock, P. C. (2007), 'Giving content to investor sentiment: The role of media in the stock market', *The Journal of finance* **62**(3), 1139–1168.

Tucker, J. A., Guess, A., Barberá, P., Vaccari, C., Siegel, A., Sanovich, S., Stukal, D. & Nyhan, B. (2018), 'Social media, political polarization, and political disinformation: A review of the scientific literature', *Political polarization, and political disinformation: a review of the scientific literature (March 19, 2018)* .

Vasileiou, E., Bartzou, E. & Tzanakis, P. (2021), 'Explaining Gamestop short squeeze using intraday data and Google searches.', *Available at SSRN 3805630* .

Wang, Y.-H., Keswani, A. & Taylor, S. J. (2006), 'The relationships between sentiment, returns and volatility', *International Journal of Forecasting* **22**(1), 109–123.

Yan, X., Guo, J., Lan, Y. & Cheng, X. (2013), A biterm topic model for short texts, *in* 'Proceedings of the 22nd International Conference on World Wide Web', WWW '13, Association for Computing Machinery, New York, NY, USA, p. 1445–1456.

Zenou, Y. (2016), 'Key players', *The Oxford Handbook of the Economics of Networks* **11**.

# A Data appendix

## A.1 Tickers Mentioned on WSB

Conventionally, submissions or comments that mention a ticker will spell it using uppercase letters, or following a dollar sign. However, a challenge is that not all uppercase words are valid tickers.

We first match words in WSB submissions to assets by identifying any succession of two to five capital letters. Subsequently, we used a pre-determined list of tickers from CRSP to check whether a match is indeed present in the available financial data. Some abbreviations or capitalised words which are not valid tickers might still show up, such as 'USD' (*ProShares Ultra Semiconductors*), 'CEO' (*CNOOC Limited*), and 'ALL' (*The Allstate Corporation*). Single characters also appear, such as 'A'

Table 5: Most Frequent Ticker Mentions

| Ticker | Name | Comments | Submissions | Sum |
|---|---|---|---|---|
| SPY | S&P 500 Index | 291,279 | 9,408 | 300,687 |
| AMD | Advanced Micro Devices, Inc. | 124,685 | 5,721 | 130,406 |
| TSLA | Tesla, Inc. | 124,222 | 5,910 | 130,132 |
| MU | Micron Technology, Inc. | 86,611 | 3,941 | 90,552 |
| AAPL | Apple Inc. | 48,345 | 1,880 | 50,225 |
| AMZN | Amazon.com, Inc. | 44,426 | 1,534 | 45,960 |
| MSFT | Microsoft Corporation | 41,152 | 1,799 | 42,951 |
| SNAP | Snap Inc. | 40,766 | 2,043 | 42,809 |
| NVDA | NVIDIA Corporation | 38,012 | 1,556 | 39,568 |
| SPCE | Virgin Galactic Holdings, Inc. | 30,758 | 1,640 | 32,398 |
| FB | Facebook, Inc. | 26,143 | 1,446 | 27,589 |
| DIS | The Walt Disney Company | 25,611 | 1,088 | 26,699 |
| BYND | Beyond Meat, Inc. | 23,299 | 906 | 24,205 |
| NFLX | Netflix, Inc. | 20,800 | 936 | 21,736 |
| JNUG | Direxion Daily Jr Gld Mnrs Bull 3X ETF | 15,761 | 1,095 | 16,856 |
| GE | General Electric Company | 15,730 | 929 | 16,659 |
| RAD | Rite Aid Corporation | 14,781 | 839 | 15,620 |
| SQ | Square, Inc. | 14,003 | 824 | 14,827 |
| ATVI | Activision Blizzard, Inc. | 13,076 | 674 | 13,750 |
| USO | United States Oil | 12,949 | 667 | 13,616 |

*Notes*: this table lists the 20 most mentioned assets on WSB, observed by submissions which uniquely mention the related ticker. 'Comments' is the number of comments posted on these submissions, 'Submissions' counts submissions, and 'Total' is the sum of the two. The name of the asset corresponding to the identified ticker is retrieved from *Yahoo Finance*.

(*Agilent Technologies, Inc.*). We manually created a list of such tickers, and removed matches featured in WSB submissions, to build a preliminary list of candidate ticker mentions. We refined a second list of candidates by checking whether a collection of one to five letters, lower or uppercase, is preceded by a dollar sign. Any mentions of '$CEO' or '$a' count as the tickers 'CEO' and 'A', respectively. These extracts are, again, checked against the list of available tickers.

A small fraction of the 4,650 tickers we extract dominate the discourse on WSB. 90% of tickers are mentioned fewer than 31 times, and more than 60% are mentioned fewer than five times. The frequency distribution of tail of ticker mentions demonstrates this point, for which Figure 11 displays a QQ-plot. We arbitrarily selected tickers with the number of mentions in the top $10^{\text{th}}$ percentile. Even though threshold of mentions for this top decile is 30 submissions, the most popular, SPY, features in almost 8,000 submissions. The orange crosses in Figure 11 locate the empirical densities, on a log scale, which are plotted against the theoretical quantiles of an exponential distribution on the x-axis. Under the assumption that ticker mentions are heavy-tailed (similarly to vocabulary distributions), the logarithm of the mentions follows an exponential distribution, with the intercept at the threshold, and the slope equal to the inverse of the tail index. Indeed, the linear fit in Figure 11 is close to perfect, supporting the assumption that the popularity of assets in WSB is heavy-tailed, with an estimated tail exponent of approximately 1.03. In what follows, we used submissions for which

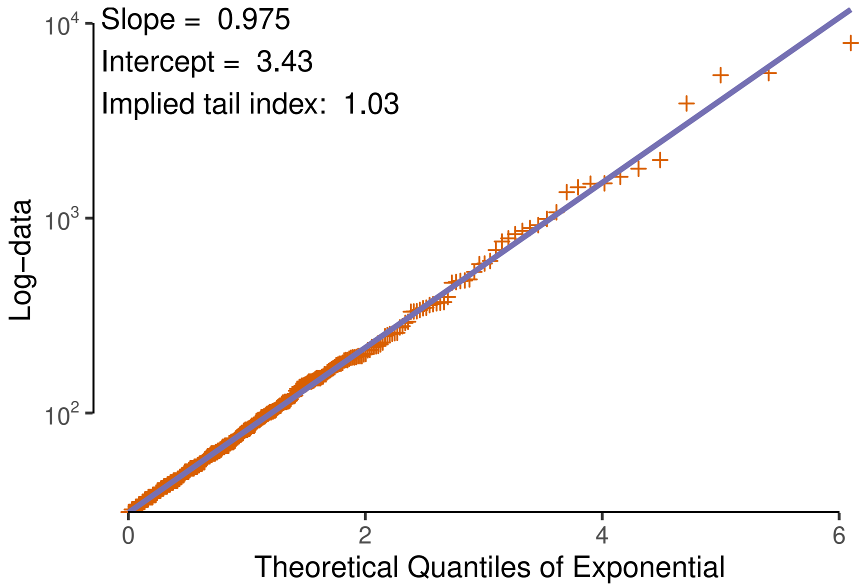


Figure 11: **QQ Plot of the Tail in Ticker Mentions on WSB**; the number of submissions for each ticker (on a log-scale) is plotted against the theoretical quantiles of an exponential distribution. Quantiles are calculated as $q(i) = -\log(1 - i/(N+1)$, where $N$ is the number of observations, and $i$ the order of the statistic, from 1 to $N$. The linear fit suggests that the data follows a Pareto distribution, with the tail index equal to the inverse of the slope. The threshold for a ticker to be part of the 'tail' is 31 mentions; note the intercept, at $\exp(3.43) \approx 31$.

we identified a single ticker, unless otherwise specified, forming a dataset of 103,205 submissions with unique ticker mentions by our cutoff date.

## A.2 Sentiment Modeling in WSB Posts

| True Label \ Predicted Label | - | **0** | + |
|---|---|---|---|
| - | 64% | 28% | 7% |
| **0** | 6% | 77% | 17% |
| + | 6% | 27% | 67% |

Table 6: **Fine-tuned FinBERT Confusion Matrix:** We use 10% of our hand-labeled data to test the performance of FinBERT on out-of-sample sentiment prediction. The results highlight the model's ability to predict sentiment with reasonably high accuracy.

In order to thoroughly understand the social dynamics of asset discussions, it is not sufficient to simply identify what assets are being discussed; it is important to understand what is being said about them. Our goal, with regards to the text data in WSB, is to gauge whether discussions on certain assets express an expectation for their future price to rise, the 'bullish' case, to fall, the 'bearish' case, or to remain unpredictable, the 'neutral' case.

A series of studies link sentiment, measured through diverse approaches, to stock market performance (Garcia 2013, Tetlock 2007, Bollen et al. 2011). Gentzkow et al. (2019) offer a thorough review. Many of these works use lexicon approaches, whereby documents are scored based on the prevalence of words associated with a certain sentiment. Recently, machine learning has offered alternative, powerful tools, such as Google's Bidirectional Encoder Representations from Transformers (BERT) algorithm (Devlin et al. 2018). The BERT algorithm trains a final layer of nodes in a neural network from a pre-trained classifier on labelled data. The classifier itself is a neural net, pre-trained by Google on a corpus of Wikipedia entries to i) predict the probability distribution of words appearing in a given sentence (Masked Language Modeling), and ii) predict the relationship between sentences (Next Sentence Prediction). BERT has further been modified through pre-training on a financial text corpus to produce FinBERT (Araci 2019).

Among other alternatives, we pursued a supervised-learning approach to identify the sentiment expressed about an asset within a WSB submission. This required a training dataset, for which we manually labelled 4,932 random submissions with unique ticker mentions as either 'bullish', 'bearish' or 'neutral' with respect to the authors' expressed expectations for the future price. We used the FinBERT algorithm for labeling. Work not shown here implements an alternative regression-based approach as a robustness check, but FinBERT performs better out-of-sample.

We trained FinBERT on 75% of the labelled data, and used the remaining 25% for validation and the test set. Table 6 plots the out-of-sample confusion matrix. For the out-of-sample test, we train FinBERT on 75% of the available data and use 15% for validation; we then compute what the algorithm predicts for the remaining 10% of data. We achieve 70% accuracy on the test set. This is better than a LASSO regression's accuracy, which was implemented separately and is not cover here.

### A.3 Sentiment variable creation

We begin with the output of our sentiment classifier, detailed in Appendix A.2. It assigns three probability scores to each submission about a ticker: the probability of a submission being bullish, $P(\phi = +1)$, bearish, $P(\phi = -1)$, neutral, $P(\phi = 0)$. The probabilities sum to one. At the time $t$ when an author $i$ posts about asset $j$, we use the probability scores above to calculate a continuous sentiment score between $(-\infty, \infty)$:

$$\Phi_{i,j,t} = \frac{1}{2}\log\left(\frac{\mathrm{P}(\phi_{i,j,t} = +1)}{\mathrm{P}(\phi_{i,j,t} = -1)}\right).$$

Submissions labeled as bullish ($\mathrm{P}(\phi = +1) = 1$), or bearish ($\mathrm{P}(\phi = +1) = 1$), are set to $\mathrm{P}(\phi = +1) = 0.98$, or $\mathrm{P}(\phi = -1) = 0.98$, to retrieve a finite value for the log-odds.

### A.4 Market variables

We include a set of market return and volatility control variables. The data source for these variables are the daily stock files issued by the Center for Research in Security Prices (CRSP), accessed through Wharton Research Data Services.

**Market variables in Section 3.2** The following market variables serve as controls.

$r_{j,t}$: the log return for asset $j$ on trading day $t$. From CRSP, we calculate it using their 'RET' variable: $r_{j,t} = \log(RET_{j,t} - 1)$, which automatically corrects the percentage change in closing prices for share splits and dividend distributions.

$\bar{r}_{j,t}$: the average log returns for asset $j$ in the five days prior to $t$ (the log return on day $t$ is not included). A minimum of three daily log-return observations is required, otherwise the observation is set as missing.

$\sigma^2_{j,t}$: the variance of log returns for asset $j$ in the five days prior to $t$ (the log return on day $t$ is not included). A minimum of three daily log-return observations is required, otherwise the observation is set as missing.

**Matching submission timings to trade timings** If a post occurs before 16:00:00 EST on day $t$, we match it with the log-return on the same day $t$. If a post occurs after 16:00:00 EST on a given day, we match it with market data for the next trading day, $t+1$. This is done to capture the fact that many news announcements occur after hours and someone posting after the market close may be exposed to these after-hour moves. Instance in which submissions are made on weekends, or holidays, are matched to the next possible trading day. For example, a submission made at 5pm on Friday is paired to the observed log return for the following Monday.

**Market variables in Sections 3.4, 5** The following market variables serve as independent variables.

$\bar{r}_{j,t-1}$: the average log returns for asset $j$ in calendar week $t-1$.

$\sigma^2_{j,t-1}$: the variance of log returns for asset $j$ in calendar week $t-1$.

Both variables are constructed from the same daily log returns panel as those in Section 3.2, described earlier in this appendix.

# B Details on testing for consensus and contagion

## B.1 Target independent variable

We extend the utility framework to suit our empirical strategy. Under the assumption that $u_{i,t}$ is drawn from a standard type-I Extreme Value Distribution, we model the log-odds of expressed investor sentiments $\phi_i$ by a standard multivariate logistic function,

$$\log\left(\frac{P(\phi_{i,t}=+1)}{P(\phi_{i,t}=0)}\right) = g(b_{i,t}) + f(\bar{\phi}_{-i,(t-1,t)}) - \theta\sigma^2_{i,t} + u^+_{i,t}, \tag{29}$$

$$\log\left(\frac{P(\phi_{i,t}=-1)}{P(\phi_{i,t}=0)}\right) = -g(b_{i,t}) - f(\bar{\phi}_{-i,(t-1,t)}) - \theta\sigma^2_{i,t} + u^-_{i,t}, \tag{30}$$

where $t$ denotes time, and $(t-1,t)$ an interval preceding $t$. The goal of this paper, in light of Proposition 1, is to test empirically whether $f(\cdot)$ is increasing. To that end, we aggregate bullish and bearish sentiments into one continuous variable, $\Phi_{i,t}$, leaving out the neutral benchmark:

$$\Phi_{i,t} = \frac{1}{2}\log\left(\frac{P(\phi_{i,t}=+1)}{P(\phi_{i,t}=-1)}\right) = g(b_{i,t}) + f(\bar{\phi}_{-i,(t-1,t)}) + \frac{u^+_{i,t}-u^-_{i,t}}{2}. \tag{31}$$

In the main body, the error term is expressed as $\epsilon_{i,t}$. Under the assumption that $u^+_{i,t}$ and $u^-_{i,t}$ are independent identically distribution, $u^+_{i,t}-u^-_{i,t}$ will follow a logistic distribution with finite variance. The intensity of neutral sentiments are an interesting manifestation of uncertainty in WSB, but ultimately not insightful in our approach to measuring the degree of social contagion in $f(\cdot)$.

## B.2 Consensus formation

This section details the variables used in Section 3, and provides additional results in support of the reported findings.

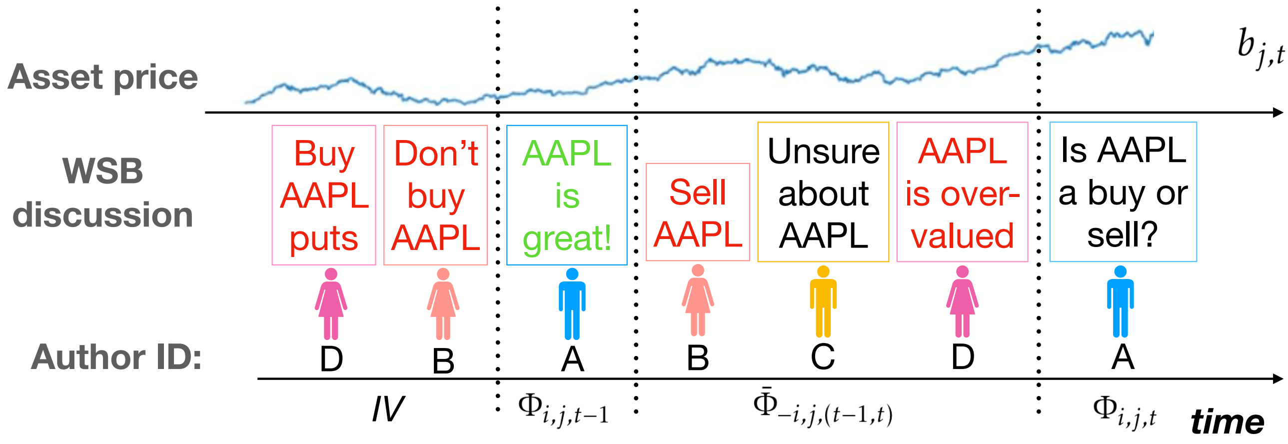


Figure 12: **Data on Investor Influence**; We consider how investor **A** is influenced by asset market movements and the sentiment of others. We use the past sentiments of author **A**'s peers (authors **B** and **D**) as an instrumental variable to predict the future sentiment of peers.

### B.2.1 Frequent Posters Approach – Further Results and Details on Estimation Strategy

In the *Frequent Posters* approach, we attempt to find the influence that peers have on an individual who posts multiple times about the same asset. We estimate this through the coefficient on $\bar{\Phi}_{-i,j,(t-1,t)}$ in the following equation:

$$\Phi_{i,j,t} = \kappa\bar{\Phi}_{-i,j,(t-1,t)} + X_{i,j}\beta_{i,j} + \epsilon_{i,j}.$$

The set of controls, $X_{i,j}$, include the author's own previous sentiment, market movements and ticker fixed effects.

Table 7, column (1), presents the full result of estimating peer effects using a regression approach. We control for the asset's returns $(r_{j,t}, \bar{r}_{j,t})$, volatility $(\sigma^2_{j,t})$ and the author's own previous sentiment $(\Phi_{i,j,t-1})$.

One challenge to accurately estimating the coefficient on $\bar{\Phi}_{-i,j,(t-1,t)}$ is the potential for both the author and his peers to experience an exogenous shock, such as a news announcement, between an author's posts. This would incorrectly be attributed to peer effects in a naive estimation of the model. For this reason, we employ an IV approach: we approximate the sentiments of peers using the peer's previously expressed sentiment about the asset.

Figure 12 illustrates our approach. We attempt to estimate the influence of peers **B**, **C**, **D** on author **A**. We observe that peers **B**, **D** post about the same asset prior to **A**'s initial post. We, therefore, estimate **B** and **D**'s view about the asset using their historical view (highlighted as *IV* in Figure 12). Any exogenous shocks that occur between *(t-1,t)* would not affect the historic views of **B**, **D**. This method, therefore, allows us to eliminate confounding variables. Market moves, labeled $b_{j,t}$ and ticker fixed effects are also included as controls.

The full estimates from our second stage regression are presented in Table 7. Both the reduced form estimates and the complete second stage support Corollary 1.1 and Proposition 2.

### B.2.2 Network Approach – Further Results

In this section, we present the results from our full second stage and reduced form regressions, from our network analysis.

Table 7: Peer Influence: Frequent Posters – Full Regression Estimates

| | | *Dependent Variable:* $\Phi_{i,j,t}$ | | |
|---|---|---|---|---|
| | | Reduced Form | Full Second Stage | Random Peers |
| | | (1) | (2) | (3) |
| *Independent Variables* | $\Phi_{i,j,t-1}$ | 0.13 (0.01) *** | 0.13 (0.01) *** | 0.15 (0.01) *** |
| | $\bar{\Phi}_{-i,j,(t-1,t)}$ | 0.06 (0.02) *** | 0.19 (0.05) *** | 0.01 (0.02) |
| | $r_{j,t}$ | 0.94 (0.17) *** | 0.95 (0.17) *** | 0.89 (0.14) *** |
| | $\bar{r}_{j,t}$ | 0.88 (0.50) * | 0.87 (0.50) * | 0.82 (0.44) * |
| | $\sigma^2_{j,t}$ | 0.23 (1.22) | 0.24 (1.21) | -0.38 (0.61) |
| | Ticker Fixed Effects | Yes | Yes | Yes |
| | No. Observations: | 11,129 | 11,122 | 14,391 |
| | $R^2$: | 0.08 | 0.08 | 0.11 |
| | $R^2_{adj}$: | 0.06 | 0.06 | 0.08 |

*Notes*: The dependent variable is individual investor sentiment about an asset, scaled continuously between $(-\infty, \infty)$, is estimated by the individual's previously expressed sentiment about the same asset ($\Phi_{i,j,t-1}$) and a set of market control variables ($r_{j,t}, \bar{r}_{j,t}, \sigma^2_{j,t}$), using OLS. The sentiment of peers ($\bar{\Phi}_{-i,j,(t-1,t)}$) is estimated in several ways. In Column (1), we use observed, average sentiment of peers between an author's two posts. In Column (2), we estimate the sentiment of peers using an IV. In Column (3), we select a random cohort to estimate peer sentiment. Robust standard errors, clustered at the ticker level, are presented in parentheses. Observations with incomplete market data are dropped.
*** Significant at 1% level ** Significant at 5% level * Significant at 10% level

Table 8 presents the reduced form for our network regression, as well as the full second stage estimates. Both the reduced form estimates and the complete second stage support Corollary 1.1 and Proposition 2.

### B.2.3 Evidence of Identification Strategy

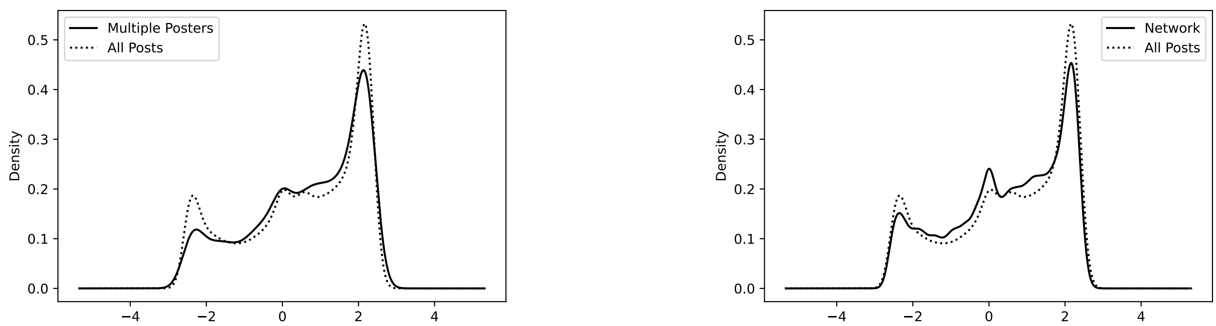


(a) Frequent Posters Sentiment PDF

(b) Commenters Sentiment PDF

Figure 13: **Density Plot of Sentiments Expressed on WSB**; We present the density plot of the sentiments expressed by users on WSB who post multiple, labeled as *Multiple Posters*, those who comment on others' posts, labeled as *Network*, and that of all submissions, labeled as *All Posts*.

A potential concern with the approach in Section 3.2.1 is whether the sentiments expressed by

Table 8: Peer Influence: Quantifying Peer Influence – Network Regressions

| | | *Dependent Variable* – $\Phi_{i,j,t}$ | | |
|---|---|---|---|---|
| | | Reduced Form | Full Second Stage | Random Network |
| | | (1) | (2) | (3) |
| *Independent Variables* | $\phi^{-1}_{i,j,t-1}$ | -0.34 (0.04) *** | -0.34 (0.04) *** | -0.35 (0.04) *** |
| | $\phi^{0}_{i,j,t-1}$ | 0.07 (0.03) ** | 0.07 (0.03) ** | 0.07 (0.04) ** |
| | $\phi^{+1}_{i,j,t-1}$ | 0.24 (0.04) *** | 0.24 (0.04) *** | 0.25 (0.04) *** |
| | $\bar{\Phi}_{-i,j,t-1}$ | 0.05 (0.01) *** | 0.31 (0.07) *** | 0.01 (0.01) |
| | $r_{j,t}$ | 0.84 (0.12) *** | 0.84 (0.12) *** | 0.85 (0.12) *** |
| | $\bar{r}_{j,t}$ | 0.90 (0.42) ** | 0.90 (0.42) ** | 0.95 (0.43) ** |
| | $\sigma^2_{j,t}$ | -0.01 (0.51) | -0.01 (0.51) | 0.00 (0.53) |
| | Ticker Fixed Effects | X | X | X |
| | No. Observations: | 24,963 | 24,963 | 25,284 |
| | $R^2$: | 0.09 | 0.06 | 0.09 |
| | $R^2_{adj}$: | 0.06 | 0.09 | 0.06 |

*Notes*: The dependent variable is individual investor sentiment about an asset, scaled continuously between $(-\infty,\infty)$. We estimate it using the individual's previously expressed sentiment about the same asset ($\phi_{i,j,t-1}$) as a categorical variable, with the author not having posted previously ($\phi^{NA}_{i,j,t-1}$) as the baseline. We control for a set of market control variables ($r_{j,t}, \bar{r}_{j,t}, \sigma^2_{j,t}$). The sentiment of posts that the author commented on previously ($\bar{\Phi}_{-i,j,t-1}$) is estimated several ways. In column (1), we present the estimate using the sentiment of posts the author previously commented on. In column (2), we use an IV to predict the sentiment of posts the author comments on. In column (3), we randomly rewire the network, connecting the author to a random set of posts about the same ticker. Robust standard errors, clustered at the ticker level, are presented in parentheses. Observations with incomplete market data are dropped.

*** Significant at 1% level ** Significant at 5% level * Significant at 10% level

individuals who post multiple times or are part of the commenters network follow the same distribution as all submissions on the forum. Figure 13a presents the distribution of sentiments for the second or later post of an author about a ticker, versus that of all submissions across all tickers. Figure 13b presents the distribution of sentiments for those who comment on other's posts, versus that of all submissions across all tickers. Figure 13 provides evidence that the distributions are very similar, which supports the hypothesis that our analysis offers insight into how all individuals on WSB form opinions.

A second concern is whether we effectively control for unobserved ticker characteristics. Similarly to Patacchini & Zenou (2016), we run 'placebo tests', where we replace the composition of an authors peers (who post between an author's posts about a ticker or posts that an author comments on) with a random cohort of people who post on WSB about the same ticker. The random cohort is chosen as follows. We observe how many peers an individual author has (how many other authors post about the same ticker mentioned in the author's posts between an author's two post or how many posts an author comments on). We then select a random sample of the same number of individuals, without replacement, who post *before* the author's original post (if fewer individuals post before, we select all of those individuals), or through a random network rewiring. The results are presented in Tables 7 and 8, column (3). We observe that all the coefficients remain close to their original values, except for the peer effect, which becomes insignificant. This lends credibility to our peer identification strategy and shows that unobserved factors that influence within ticker variation are not confounding our estimates.

# C Details on modelling market impact

## C.1 Aggregating sentiments

**Aggregating sentiments** In this section, we describe how investors reach a buy/sell decision about an asset on aggregate. First, we consider that they form a sentiment $\phi_i$ about an asset based on the signals we consider in Section 3. Their choice to buy / sell the asset is modelled as a quantal response function. The precise steps are outlined here.

In Section 3, we model the log-odds of investor sentiment $\phi_i$ by a standard multivariate logistic function. This stems from a utility ranking between adopting a positive ($\phi_i = +1$) or negative ($\phi_i = -1$) position on the asset, as a function of observed returns and peer sentiment:

$$U(\phi_{i,t+1} = +1) = \epsilon_i^+ + \alpha\phi_t + \beta r_t - \gamma r_t^2, \quad U(\phi_{i,t+1} = -1) = \epsilon_i^- - \alpha\phi_t - \beta r_t - \gamma r_t^2.$$

An assumption on the distribution of the idiosyncratic terms $\epsilon_i^+$ and $\epsilon_i^-$, which we state in Assumption 3, allows us to derive an expression for the probability of bullish versus bearish sentiment. We write the probability of bullish sentiment as

$$\begin{aligned}
\mathrm{P}[U(+1) > U(-1)] &= \mathrm{P}(\epsilon_i^+ + \alpha\phi_t + \beta r_t - \gamma r_t^2 > \epsilon_i^- - \alpha\phi_t - \beta r_t - \gamma r_t^2), && (32)\\
&= \mathrm{P}(\epsilon_i^+ + 2\alpha\phi_t + 2\beta r_t > \epsilon_i^-) && (33)\\
&= \int_{-\infty}^{+\infty} \exp\left[-\exp\left(-\frac{\epsilon_i^+ + 2\alpha\phi_t + 2\beta r_t}{2\lambda}\right)\right] \times \frac{1}{2\lambda}\exp\left[-\frac{\epsilon_i^+}{2\lambda} - \exp\left(-\frac{\epsilon_i^+}{2\lambda}\right)\right] d\epsilon_i^+, && (34)\\
&= \frac{\exp\left(\frac{\alpha\phi_t + \beta r_t}{\lambda}\right)}{\exp\left(\frac{\alpha\phi_t + \beta r_t}{\lambda}\right) + 1}, && (35)
\end{aligned}$$

where step 34 follows from the fact that we are integrating the cumulative density of $\epsilon_i^-$, with scale $2\lambda$, over the support of $\epsilon_i^+$, which also has scale $2\lambda$. The final simplification comes from the fact that $U(\phi_{i,t+1} = +1)$ and $U(\phi_{i,t+1} = -1)$ are symmetric in parameters $\alpha, \beta$ and $\gamma$. If we want to keep them distinct, we arrive at the quantal response function for the probability of positive or negative sentiment:

$$\phi_{i,t+1}^D = \begin{cases} +1, & \text{with probability} \frac{\exp\left[(\alpha\phi_t + \beta r_t - \gamma r_t^2)/2\lambda\right]}{\exp\left[(\alpha\phi_t + \beta r_t - \gamma r_t^2)/2\lambda\right] + \exp\left[(-\alpha\phi_t - \beta r_t - \gamma r_t^2)/2\lambda\right]}, \\ -1, & \text{with probability } \frac{\exp\left[(-\alpha\phi_t - \beta r_t - \gamma r_t^2)/2\lambda\right]}{\exp\left[(\alpha\phi_t + \beta r_t - \gamma r_t^2)/2\lambda\right] + \exp\left[(-\alpha\phi_t - \beta r_t - \gamma r_t^2)/2\lambda\right]}. \end{cases}$$

Therefore, we can model the aggregate buying intensity by summing $\phi_{i,t+1}^D$ across the $N$ hype investors, which yields the well-known hyperbolic tangent function:

$$\phi_{t+1} = \tanh\left[(\beta r_t + \alpha\phi_t)/\lambda\right].$$

**$\epsilon_i$ type-I Extreme Value (EV) distribution** We investigate empirically where our claim that $\epsilon_i$'s, individual investor's errors/priors, follow a type-I EV distribution. In order to study this, we look at the residuals from the following two regressions, which we run independently:

$$\log\left(\frac{P(\phi_{i,t+1} = +1)}{P(\phi_{i,t+1} = 0)}\right) = \epsilon_i^+ + \alpha\phi_t + X_{i,t}\beta, \quad \log\left(\frac{P(\phi_{i,t+1} = -1)}{P(\phi_{i,t+1} = 0)}\right) = \epsilon_i^- - \alpha\phi_t + X_{i,t}\beta. \tag{36}$$

Consistently with the *Frequent Posters* approach, the set of controls $X_{i,t}$ includes the author's own previous sentiment, market movements and ticker fixed effects. We observe a reasonable fit of the type-I EV distribution in Figure 14, although the observations appear to cluster around certain levels. We attribute this to the sentiment model, FinBERT. The output of the NLP model provides the probabilities for the sentiment of a post being positive, negative or neutral (with all probabilities summing to one). However, as a model trained on a classification task, it is more likely to highly weight a single category, rather than give an equal probability of a statement falling into all three categories. The three humps are, therefore, tied to the way in which probabilities are assigned by our NLP model, and are also clearly visible in the distribution of sentiments in Figures 13a and 13b.

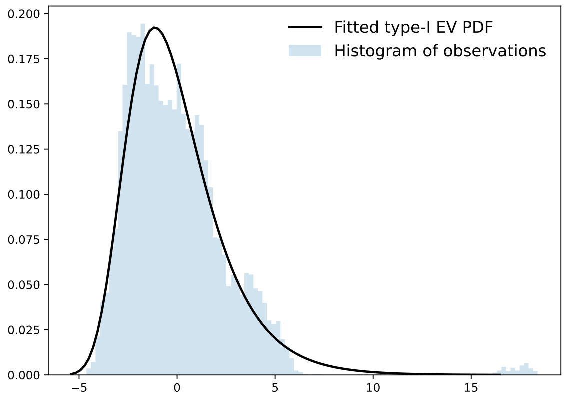


(a) Distribution of $\epsilon_i^-$ with location parameter $-1.13$ and scale parameter 1.91

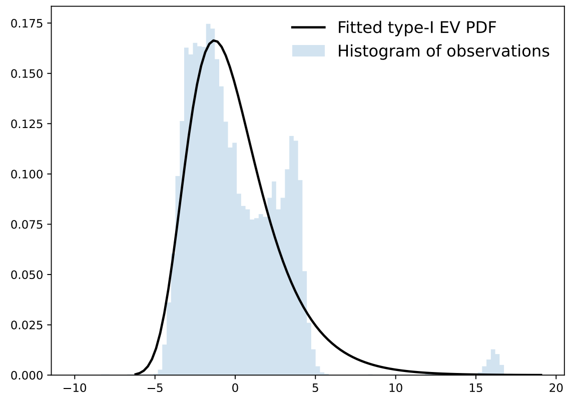


(b) Distribution of $\epsilon_i^+$ with location parameter $-1.30$ and scale parameter 2.21

Figure 14: **Distribution of $\epsilon_i^+$ and $\epsilon_i^-$**; We plot the observed distribution of residuals from the regressions presented in Eq. 36, which we use to estimate author's $\epsilon_i$. We overlay our histogram of observations with a fitted type-I error distribution Probability Density Function (PDF).

### C.2 Stability Analysis

In this section of the appendix, we discuss the methodology behind Section 4.2. We precisely outline the steps taken to find our system steady states, as well as show the expected system behaviour around the steady states.

As mentioned in Section 4.2, $r_t$ must be zero at the steady state. This means that steady states of the dynamic system in Eqs. 14-15 are those for which $\phi_t$ is a solution to Eq. 14, which is the hyperbolic tangent function displayed in Figure 6. This means that zero is a unique steady state when $\alpha/\lambda < 1$, and two further steady states emerge when $\alpha/\lambda > 1$. Those two additional steady states are solved numerically in all simulation exercises, using a solving algorithm.

The behaviour of our system depends not only on the existence of steady states, but also on the types of steady states that we observe in different stability regions. In a discrete time system, the type of stability around a steady state is dependent on the eigenvalues of the Jacobian matrix at the steady state. Table 9 shows some types of common behaviours around a single steady state, as they relate to the Jacobian eigenvalues. When a system has multiple eigenvalues, different behaviours can emerge. For example, if we consider two eigenvalues $x_1$ and $x_2$ and $|x_1| > 1$ while $x_2 < 1$, we would observe a saddle. Various resources exist to review system steady states, largely originating from the study of physical systems, with Hommes (2013) providing a relevant review for the economic setting.

Table 9: System behaviour around steady states and Jacobian Eigenvalues

| | | |
|---|---|---|
| $\lvert x_i \rvert < 1$ | $\mathbb{I}(x_i) = 0$ | stable node / sink |
| $\lvert x_i \rvert < 1$ | $\mathbb{I}(x_i) \neq 0$ | stable focus / spiral sink |
| $\lvert x_i \rvert > 1$ | $\mathbb{I}(x_i) = 0$ | unstable node / source |
| $\lvert x_i \rvert > 1$ | $\mathbb{I}(x_i) \neq 0$ | unstable focus / spiral source |

*Notes*: The behaviour of our system around our steady states is dependent on the eigenvalues of the Jacobian at the steady states. Here we describe some of the most common eigenvalue combinations and behaviours. For a more comprehensive review, see Hommes (2013).

The Jacobian matrices at the different steady states are

$$J(0,0) = \begin{bmatrix} \alpha & \beta \\ C(\alpha-1) & C\beta \end{bmatrix}, \quad J(\phi,0) = \begin{bmatrix} \alpha \mathrm{sech}^2(\alpha\phi) & \beta \mathrm{sech}^2(\alpha\phi) \\ C\left(\alpha \mathrm{sech}^2(\alpha\phi) - 1\right) & C\beta \mathrm{sech}^2(\alpha\phi) \end{bmatrix}, \tag{37}$$

where $\phi \in \{\phi^+, \phi^-\}$. The corresponding eigenvalues for steady state $(0,0)$ are

$$x_1 = \frac{C\beta + \alpha + \sqrt{(C\beta+\alpha)^2 - 4C\beta}}{2}, \quad x_2 = \frac{C\beta + \alpha - \sqrt{(C\beta+\alpha)^2 - 4C\beta}}{2}, \tag{38}$$

and those for steady states $(\phi,0)$ are

$$x_1 = \frac{\mathrm{sech}^2(\alpha\phi)(C\beta+\alpha) + \sqrt{\mathrm{sech}^4(\alpha\phi)(C\beta+\alpha)^2 - 4C\beta \mathrm{sech}^2(\alpha\phi)}}{2}, \tag{39}$$

$$x_2 = \frac{\mathrm{sech}^2(\alpha\phi)(C\beta+\alpha) - \sqrt{\mathrm{sech}^4(\alpha\phi)(C\beta+\alpha)^2 - 4C\beta \mathrm{sech}^2(\alpha\phi)}}{2}. \tag{40}$$

These expressions trace out different regions of stability as a function of $C\beta$ and $\alpha$ in Figure 7.

Figure 7 is drawn by solving numerically for the points $C\beta$ for which the eigenvalues in Eq. 38, on the left, and Eqs. 39-40, on the right, switch between the different stability regions. Specifically, we solve for when the complex part of the eigenvalues are equal to zero or not equal to zero, and are within / outside the unit circle. These points are computed by taking a range of values for $\alpha$, then applying a solving algorithm. These solutions are then used to distinguish the different regions of stability in Figure 7.

Figures 15 and 16 show the phase diagrams in the regions surround our steady states, for the different parameter regions discussed in Section 4.2. We observe how the system changes around the steady states, depending on our choices of parameters, whether the system appears to move towards stability, and how it moves towards the steady state (directly or through cycling around it).

**Stability at** $(0,0)$ Steady state $(0,0)$ displays five different types of behaviours as a function of parameters $C\beta$ and $\alpha$. These spaces are represented in Figure 7a for relevant parameter values. In region (A), both eigenvalues are real and greater than one. This means that $(0,0)$ is an unstable node. In region (B), both eigenvalues are complex, and their modulus is greater than one; in this region, $(0,0)$ is an unstable focus. In regions (C) and (D), both modulii are less than one, indicating that $(0,0)$ is stable. The eigenvalues in (C), unlike those in (D), are complex, meaning that $(0,0)$ is a focus instead of a node in that region. In the final region (E), both eigenvalues are real, but one is greater than one, meaning that $(0,0)$ is a saddle.

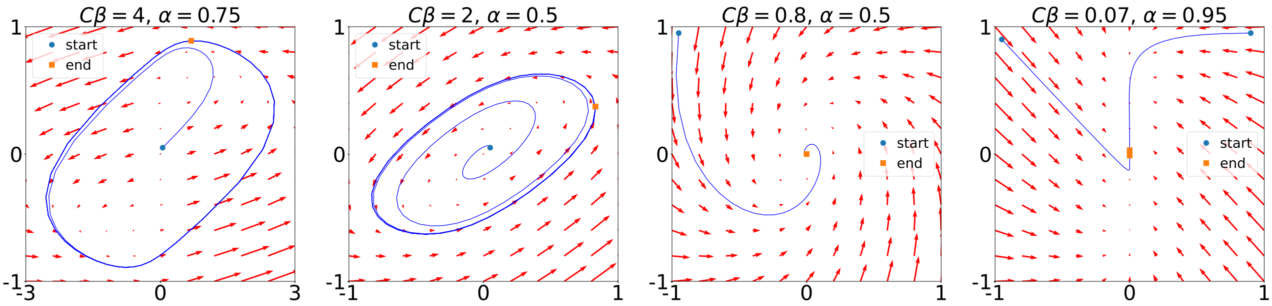


(a) Region (A) (b) Region (B) (c) Region (C) (d) Region (D)

Figure 15: **Stability Regions from Figure 7a**; we plot the phase portrait for our dynamic system showing qualitatively which direction our system will move from time $t$ to time $t+1$ for a given starting point. Returns $r$ are plotted on the x-axis and aggregate sentiments $\phi$ are plotted on the y-axis. In all of these regions, our system has one steady state. The trajectories for the system for different starting points are simulated for 200 time steps using a continuous time, Ordinary Differential Equation approximation and displayed as blue lines.

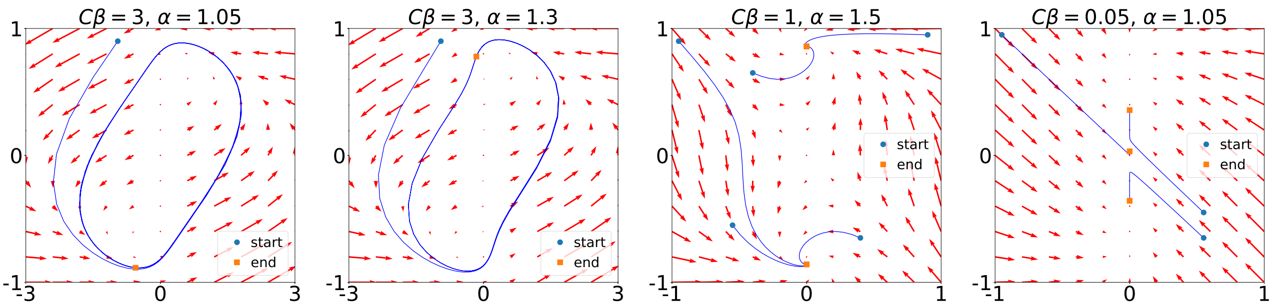


(a) Region (F) (b) Region (G) (c) Region (H) (d) Region (I)

Figure 16: **Stability Regions from Figure 7b**; we plot the phase portrait for our dynamic system showing qualitatively which direction our system will move from time $t$ to time $t+1$ for a given starting point. Returns $r$ are plotted on the x-axis and aggregate sentiments $\phi$ are plotted on the y-axis. In all of these regions, our system has three steady states. The trajectories for the system for different starting points are simulated for 200 time steps using a continuous time, Ordinary Differential Equation approximation and displayed as blue lines.

A change in capacity $C$ can therefore be significant, in two qualitative ways. Assuming that consensus formation is relatively weak, such that $\alpha < 1$, an exogenous increase in capacity from a state initially close to zero introduces oscillations in the path to the stable node at $(0,0)$. The dynamics change in observed dynamics between regions (D) and region (C) can be attributed to hype investors overreacting to large returns, thus overshooting the steady state in region (D). The ensuing change in sentiments fails to cause returns to diverge form their steady state in perpetuity, and investors correct their positions in the opposite direction in the subsequent period. A more consequential bifurcation happens when the change in capacity is large enough, so that $C\beta > 1$ and the parameters move from region (C) to region (B). This type of bifurcation is known as a 'Hopf bifurcation', as two complex eigenvalues cross the unit circle (Hommes 2013).

**Stability at $(\phi^+, 0)$ and $(\phi^-, 0)$** When consensus formation is relatively high ($\alpha > 1$), two new steady states emerge; one with positive sentiment $(\phi^+, 0)$ and one with negative sentiment $(\phi^-, 0)$. Both display identical behaviours in terms of stability, as a function of parameters $C\beta$ and $\alpha$. These are separately visible in Figure 7b, since they only exist when $\alpha > 1$. In regions (F) and (G), the modulii of both eigenvalues are greater than one, but are real only in region (F). In contrast, both modulii are less than one in regions (H) and (I), and real only in region (I). This means that both steady states are stable focuses in region (H), and stable nodes in region (I).

Generally, the stronger consensus formation is in the dynamic system, the larger the combined trend following component and investor capacity has to be in order to destabilise these two steady states. Functionally, the historical signal is too weak to overcome the strong consensus effect, and hype investors hardly stray from their equilibrium sentiment. In this sense, social dynamics in the form of consensus formation may actually serve to stabilise markets.

### C.3 Bifurcation diagram with respect to consensus

We produce two additional bifurcation diagrams for each state variable, varying $\alpha$. Figure 7 suggests two interesting values for $C$ – one at 0.3, and a second at 2.5. We assign parameter $\beta$ a value of one. We then initialise sentiment $\phi_0$ to be zero, and randomly draw one hundred values for returns $r_0$ from a normal distribution with mean zero and variance 0.25. For a set of values of $\alpha$, ranging from zero to two, we iterate forward $\phi_t$ and $r_t$ one thousand times. The final values for $\phi_t$ and $r_t$ are then plotted as a function of $\alpha$.

Those plots are presented in Figure 17. Figure 17a illustrates the final values for $\phi_t$ when $C$ is set to 2.5, and Figure 17b does the same for $r_t$. For low values of $\alpha$, sentiments, and concurrently returns, display quasi-periodic dynamics (in the clouded regions), interrupted by stable cycles (for example when $\alpha$ is around 0.8). Larger values of $\alpha$ eventually dominate the effect of hype investors' large capacity, and the system settles at it steady states, with either positive or negative sentiment. This in stark contrast to a scenario where capacity is too small, as in Figures 17c and 17d where $C$ is set to 0.3. The meaningful change in the behaviour of the system is that sentiments settle at one value, either positive or negative, when $\alpha$ is greater than one – already present in the one-dimensional hyperbolic tangent map from Figure 5. It is interesting to note that there is slight residual dispersion in returns in Figure 17d (note the scale of the axis) when $\alpha$ reaches that threshold, as one steady state loses its stability, and two new ones barely emerge.

## D Topic model and narratives

### D.1 Topic model

Does WSB reflect new information for the larger market to trade on, or social activity that drives perceived changes in value, regardless of fundamentals? A topic model offers a simple method to evaluate the content of WSB discussions. Figure 18 presents our preferred topic model, namely the Biterm Topic Model (BTM), which is optimal for smaller bodies of text (Yan et al. 2013). Submissions from April 2012 to February 2021 give a time series of almost 100 months. A random sample of submissions is drawn for the months of January and February 2021 in order to prevent these two months, with a high number of submissions, from skewing the topic model results.

Figure 18 presents a stacked plot of the monthly submission count of a selected subset of discussions, normalised by the total across the selected topics. It begins in 2015 when the forum gained a consistent user base. On one hand, some topics persist in the overall discussion: people consistently ask for advice about trading accounts and anonymously share details of how their trading is affecting

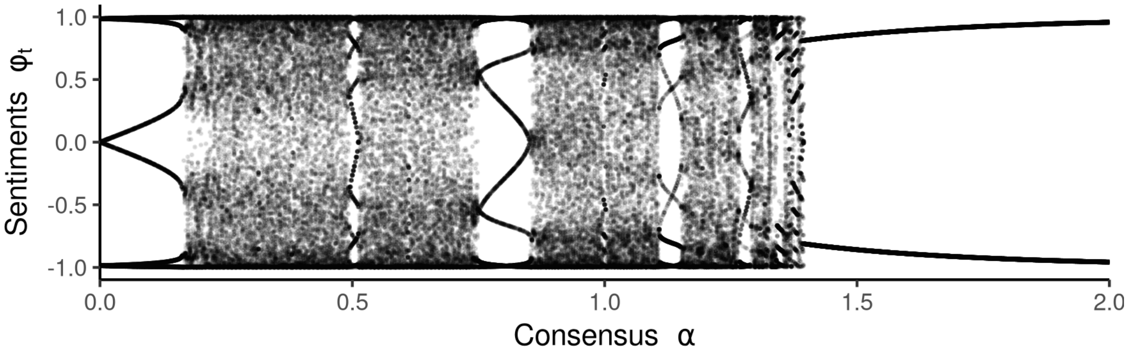


(a) $C = 2.5$

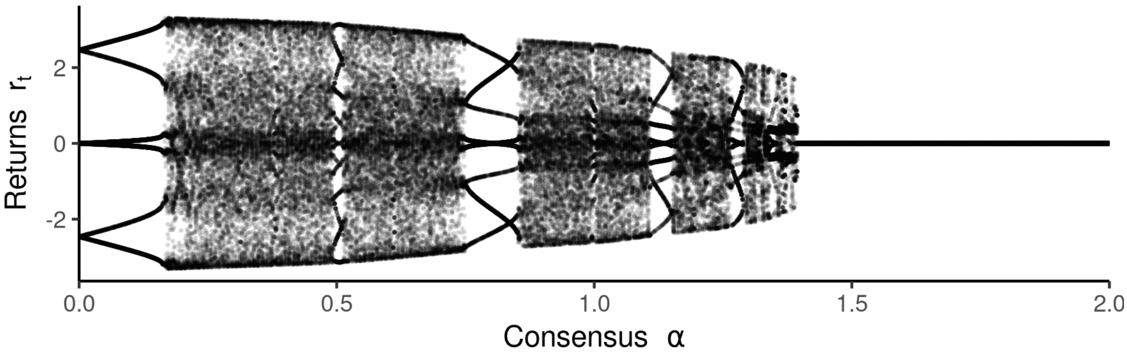


(b) $C = 2.5$

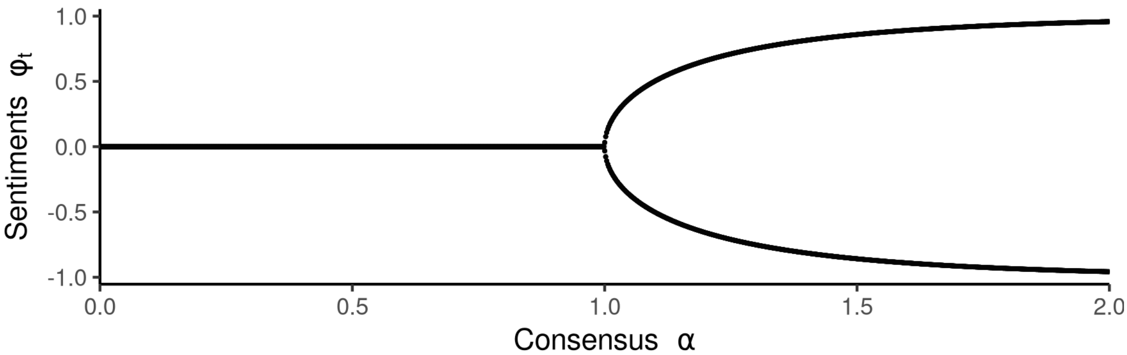


(c) $C = 0.3$

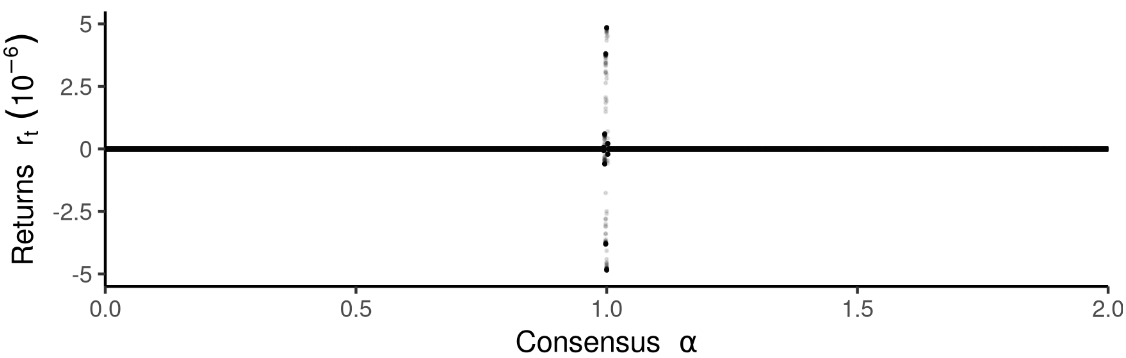


(d) $C = 0.3$

Figure 17: **Bifurcation diagrams with respect to consensus parameter** $\alpha$; the final values from one thousand iterations of the dynamic system in Eqs. 14-15, are plotted, for two values of $C$ (2.5 and 0.3) and a range of values for $\alpha$ between zero and two. Remaining parameters are constant for $\beta = 1$ and $\lambda = 1$. One hundred initial values for $r$ are drawn from a normal distribution with mean zero and variance 0.25, with $\phi$ initiated at zero.

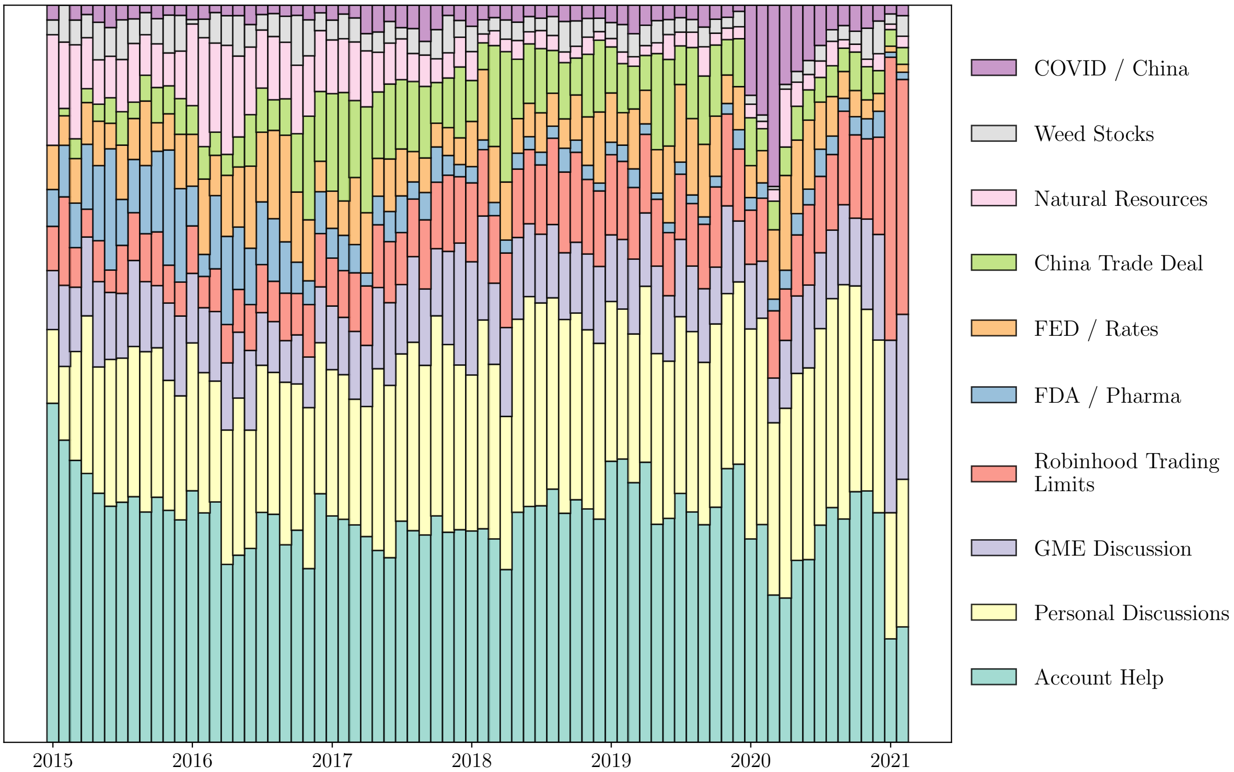


Figure 18: **Temporal trends in topics**; the stacked count, normalized to 100% at each time period, showing the prevalence of a select subset of topics discussed on WallStreetBets.

their personal lives. On the other hand, topics concerned with larger economic trends wax and wane over the observation period. Two examples of this are the uptick in submissions discussing GME and Roinhood account trading limits, coinciding with the GME short squeeze, and the COVID-19 topic, which is negligible until January 2020, but gains prominence in the subsequent months. Pharmaceutical company and natural resource discussions, on the other hand, seem to loose popularity. A full list of topics with their respective keywords is presented below.

| Topic Title | Top Words | Topic Prevalence (%) |
|---|---|---|
| Robinhood Trading Limits | robinhood, gme, account, nkla, margin, order, app, limit, broker, orders | 1.8 |
| International Trade | expected, yr, china, usd, europe, japan, pmi, manufacturing, korea, data | 0.8 |
| Retail Sales + Amazon | sales, amazon, home, stores, business, online, companies, food, store, retail | 3.0 |
| Top Stock Picks / Positions | tsla, news, sold, aapl, weeks, holding, hold, amd, months, dip | 11.6 |
| Other | comments, daily, best, moves, spy, weekend, fo, fn, fm, fp | 1.1 |
| Other | mentions, vote, log, wsbvotebot, submission, posts, check, reverse, mention, great | 0.2 |
| Electric Cars | tsla, energy, car, ev, cars, nio, electric, battery, elon, space | 2.3 |
| Revenues, Earnings, Ratings | revs, beats, tgt, line, eps, neutral, downgraded, initiated, fy, reports | 0.9 |
| FDA / Pharma | drug, fda, phase, patients, vaccine, trial, treatment, results, clinical, covid | 2.9 |
| Revenues, Earnings, Ratings | revenue, million, growth, quarter, share, sales, billion, net, expected, eps | 4.0 |
| China Trade Deal | trump, china, said, president, deal, bill, house, election, chinese, news | 3.1 |
| Social Media Stocks | fb, game, snap, aapl, disney, games, video, google, users, netflix | 2.7 |
| GME Discussion | companies, gme, investors, world, years, believe, hedge, actually, value, funds | 7.3 |
| Financial News | data, information, news, financial, report, based, find, sec, research, investors | 4.3 |
| Personal Discussions | life, wife, ass, little, said, spy, getting, old, red, went | 7.6 |
| Weed Stocks | million, share, capital, ceo, cannabis, ipo, public, management, merger, billion | 2.1 |

Table 10: **Topics Extracted from BTM Model (1)**

| Topic Title | Top Words | Topic Prevalence (%) |
| --- | --- | --- |
| Software Tech Stocks | amd, aapl, intc, companies, data, cloud, software, technology, services, tech | 3.4 |
| Earnings Release | release, estimates, consensus, share, revenue, move, open, beat, average, interest | 1.0 |
| Natural Resources | oil, gold, prices, silver, gas, futures, crude, production, demand, companies | 1.8 |
| FED / Rates | fed, rates, rate, economy, markets, economic, said, interest, growth, inflation | 6.0 |
| Other Tech Stocks | tsla, pltr, elon, musk, mods, ban, gme, retards, gains, autists | 2.0 |
| COVID / China | virus, covid, cases, coronavirus, china, world, weeks, corona, states, news | 4.8 |
| Other | spy, bear, gay, text, bears, bull, gang, msft, words, stonks | 1.7 |
| Debt / Loans | debt, cash, pay, credit, loans, loan, million, interest, billion, bank | 3.4 |
| Other | usd, bln, exp, revenue, newswires, eps, co, share, symbol, live | 1.4 |
| Other | spy, chart, close, index, month, performance, major, past, futures, sectors | 1.3 |
| Account Help | account, help, investing, best, start, robinhood, advice, work, years, please | 7.4 |
| Other | spy, volume, chart, support, bullish, low, trend, resistance, bearish, line | 3.7 |
| Other | amet, calendar, releases, wed, link, thurs, tues, fri, analyst, close | 0.5 |
| Options / Risk | option, spy, profit, strike, spread, risk, loss, value, selling, position | 6.3 |

Table 11: **Topics Extracted from BTM Model** (2)