NEOMA BUSINESS SCHOOL

REIMS · ROUEN · PARIS

MS IN MARKETING & DATA ANALYTICS

Academic Year 2024–2025

PROFESSIONAL THESIS

Rethinking Financial Bubbles in the Digital Era: From Irrational Markets to Algorithmic Finance and the Role of AI

Author: Guillaume DUBOIS

Host Company: Microsoft

Position Held during Work-Study Program: Business Manager

Director of the MS MDA: Othman BOUJENA

Professional Thesis Supervisor: Kiyali OUATTARA

Work-Study Tutor (N+1): Baptiste GIRAUDIER

Position: Sales Manager

Defense Date: 12-09-2025

NEOMA
BUSINESS SCHOOL
REIMS · ROUEN · PARIS

MS IN MARKETING & DATA ANALYTICS

Academic Year 2024–2025

DECLARATION OF NON-PLAGIARISM

I, the undersigned, Mr. **Guillaume DUBOIS**, of the current academic year cohort of the Specialized mastere's in marketing & data Analytics (MS/MSc MDA), hereby declare that:

- This professional thesis is an original document and strictly the result of my own personal work,
- the information contained in the professional thesis submitted today under the same program is not plagiarized,
- all bibliographic sources used in the writing of this thesis are fully indicated in the references section at the end of the document.

Date: 25/08/2025
Signature:

NOTICE ON THE USE OF AI

I used AI tools as helpers while writing my thesis. They did not write the content for me. Instead, they helped me organize ideas, tidy up wording, draft small code snippets, and spot weak spots in my arguments. The analysis and conclusions are fully my own.

Tools I used

- ChatGPT (GPT-5): rephrasing, structuring, light drafting/revision, simple Python code, and general feedback.
- Notebook LLM: supporting the literature review by extracting and organizing insights from selected articles.

What AI helped with

- Writing: clearer English, consistent tone (including sections first drafted in French).
- Literature review: summarizing dense papers and linking key ideas.
- Data analysis: quick Python to write coding, classify survey answers and sanity-check my coding.
- Planning & feedback: proposing possible outlines and flagging sections that were dense, repetitive, or unclear.

Examples of prompts (condensed)

Note: these are shortened summaries of the original prompts for transparency.

- "Turn my literature-review conclusion into a synthesis that links Part 1 (psychology/narratives/herding) with Part 2 (algorithms/personalization/hype/predictive analytics/AI paradox); highlight reflexivity, show tech as both amplifier and diagnostic, end with my thesis statement, keep my tone."
- "Revise in my voice: break up dense sections, balance evidence vs. critique, reduce repetition of mechanisms, sharpen the EMH link, and deepen the comparison (why digital bubbles spread faster and are harder to arbitrage)."
- "Create two binary moderators for platform use: algorithm-heavy (Twitter/YouTube/Instagram/TikTok) and community-driven (Reddit/Discord/Telegram). Code 1 if the respondent uses at least one in the category."

Critical reflection

AI was useful but not magic. Prompting took time, and I often had to rewrite outputs to fit my argument and voice. I stayed cautious about potential errors (e.g., fabricated references, generic phrasing) and verified sources and code myself.

Bottom line

AI acted as a supportive assistant for polishing, structure, and checks. The ideas, methods, and judgments in this thesis are entirely mine.

ACKNOWLEDGMENTS

First and foremost, I would like to express my sincere gratitude to Professor Othman Boujena, Director of the MS Marketing & Data Analytics, for his guidance and for fostering an academic environment that made this journey possible.

I am very thankful to my Professional Thesis Supervisor, Kiyali Ouattara, for his valuable insights, and constructive feedback throughout this work. Beyond his direct supervision, I also benefitted from the exchanges and shared experiences of other students under his guidance, which enriched my own reflections and perspective.

I would also like to extend my appreciation to my Work-Study Tutor, Baptiste Giraudier, for his mentorship, encouragement, and availability during my experience at Microsoft. My thanks go as well to Laurent Herbulot, my Microsoft buddy, whose advice and support were invaluable to me both professionally and personally. More broadly, I am grateful to the entire Microsoft LRG Public Sector team and to all those within the organization who contributed to my professional development and supported me throughout this journey.

A special thanks goes to Bella Heejin AHN, whose thoughtful advice, encouragement, and unwavering support have been instrumental throughout the process. Her presence has been a key element in bringing this work to completion.

I would also like to warmly thank my fellow students, Nawal AMELLOUK, Maurad AYACHE, Juliette HONDIER, Lucas GATTERER, and Marius LANGLOIS, for their collaboration, stimulating discussions, and solidarity. More broadly, I am grateful to the entire MS MDA class and staff, whose support and motivation have made this academic journey both a privilege and a source of inspiration.

My gratitude also extends to all the individuals who generously took the time to answer my survey. Your contributions enriched this research and gave it real-world relevance.

SUMMARY NOTE

This thesis examines how financial bubbles evolve in the digital era and the role of algorithms and artificial intelligence (AI) in detecting them. Traditional theories based on rational markets no longer suffice. Today, speculation spreads through viral content, influencers, and algorithm-driven platforms. AI can help identify risks early but also fuels speculation, making its impact ambivalent.

An online quantitative survey (282 responses) gathered investor perceptions. Questions covered investment profiles, psychological biases, digital influence, AI, and speculative behavior. Statistical analysis showed that respondents, mostly young and cautious, viewed algorithms as highly influential. Loss aversion appeared as the main bias. Many admitted hype affects their thinking but claimed to act with restraint, highlighting a gap between intention and action. Digital platforms were the main source of financial information, while trust in AI remained divided. Those with greater trust in AI also demanded stronger accountability from platforms.

The thesis concludes that speculation is shaped not only by human psychology but also by digital system design. It argues that AI should be seen not as a price oracle but as a diagnostic tool for regulators, and stresses the need for transparency, accountability, and governance so that AI supports human judgment without increasing instability.

NOTE DE SYNTHESE

Cette thèse analyse l'évolution des bulles financières à l'ère numérique et le rôle des algorithmes et de l'intelligence artificielle (IA) dans leur détection. Les théories classiques, fondées sur l'idée de marchés rationnels, ne suffisent plus. Aujourd'hui, la spéculation se propage via le contenu viral, les influenceurs et les plateformes algorithmiques. L'IA peut aider à détecter les risques, mais elle alimente aussi la spéculation, rendant son impact ambivalent.

Une enquête quantitative en ligne (282 réponses) a permis de recueillir les perceptions des investisseurs. Les questions portaient sur les profils d'investissement, les biais psychologiques, l'influence du numérique, l'IA et les comportements spéculatifs. L'analyse statistique a révélé que les répondants, majoritairement jeunes et prudents, considèrent les algorithmes comme très influents. L'aversion aux pertes est apparue comme le biais principal. Beaucoup admettent l'impact du « hype », mais disent rester modérés dans leurs actes, montrant un écart entre intentions et comportements. Les plateformes numériques sont leur principale source d'information, tandis que la confiance envers l'IA reste divisée. Ceux qui lui font confiance demandent aussi plus de responsabilité des plateformes.

La thèse conclut que la spéculation est autant le produit de la psychologie humaine que de la conception numérique. Elle propose de voir l'IA non comme un oracle des prix, mais comme un outil diagnostique pour les régulateurs, et insiste sur la transparence, la responsabilité et la gouvernance afin qu'elle soutienne le jugement humain sans accentuer l'instabilité.

TABLE OF CONTENTS

INTRODUCTION

Financial bubbles are the moments when markets lose touch with reality. Indeed, the prices climb too high, excitement builds on itself, and sooner or later the bubble bursts. For years, economists explained these events by assuming markets were for the most part rational and that prices reflected all the available information. However, history showed us that bubbles are not rare accidents as they keep coming back.

Today, the way bubbles form is significantly changing. Social media, online communities, and digital platforms have created new spaces where speculation grows. A tweet, a meme, or a viral video can now move millions of dollars in minutes. Investors do not just follow numbers anymore. They follow stories, influencers, and online trends. Attention itself has become a kind of fuel for speculation.

At the same time, technology is not just part of the background, it actively shapes these markets. Algorithms decide what we see and what we miss. They amplify certain stories, steer collective focus, and create isolated groups were enthusiasm spreads quickly. Artificial intelligence brings an additional dimension. It raises the possibility of detecting risks, anticipating crashes, and even guiding investment decisions. At the same time, it raises new questions. Can these tools really help us identify bubbles early?

This leads to the central question of this thesis: **How are financial bubbles reshaped in the digital era, and to what extent can algorithms and artificial intelligence help in detecting them?** From this overarching problem arises several sub-questions that guide the investigation and structure the analysis.

- What are the limits of traditional financial frameworks in understanding bubbles in the new digital age?
- How are speculative behaviors reshaped by digital communities and new forms of social interaction?
- What can we learn by comparing today's bubbles with historical episodes of speculation?
- How do algorithms and digital infrastructures influence the way market's function?
- To what extent can artificial intelligence and data-driven tools help detect bubbles?
- What challenges and paradoxes arise when technology both drives and seeks to control financial speculation?

To answer this, the thesis unfolds in several stages. It begins with a literature review divided into two parts. Chapter I looks at the shift from classical finance to the digital era, showing how biases, narratives, and online communities give speculation new forms and new intensity. Chapter II focuses on algorithmic finance, examining how platforms, algorithms, and artificial intelligence both fuel speculation and attempt to predict or contain it.

The literature review is then complemented by a quantitative survey, designed to bring evidence on how investors today perceive bubbles, digital platforms, and the role of AI. This makes it possible to connect theory with lived experience.

Building on this, we will present the analysis and interpretation of the survey results, highlight the main insights and linking them back to the theoretical framework. This section also includes recommendations and perspectives, offering ideas for future research, regulation, and practice.

Finally, the conclusion will bring together the key findings. It reflects on how bubbles are reshaped in the digital age, what role algorithms and AI play in their formation, and what this means for the stability and future of finance.

I.FROM CLASSICAL FINANCE TO DIGITAL DISRUPTION: RETHINKING SPECULATIVE BEHAVIOR

A. BEYOND CLASSICAL FINANCE: THE FRAGILITY OF RATIONAL MARKETS

A.1. Limits of Efficiency: Why Rational Expectations Fall Short in Explaining Market Dynamics

The **Efficient Market Hypothesis (EMH)**, as described by Fama (1970), is built on the idea that *"asset prices fully reflect all available information."* It assumes that investors behave rationally, process news without bias, and that markets function smoothly so any mispricing is quickly corrected. This perspective shaped how we design portfolio strategies, develop asset-pricing models, and even set financial regulations. Its strength was that it was elegant, clear, and testable, which is why it became the standard "*null hypothesis*" for understanding market behavior (Fama, 1970).

Yet, from its very beginning, EMH has faced serious empirical challenges. One of the earliest and most influential critiques came from Shiller (1981), who showed that stock prices fluctuate far more than can be explained by changes in dividends, highlighting the problem of *"excess volatility."* In other words, prices seemed to move in ways unrelated to underlying fundamentals, directly challenging EMH's core claim. LeRoy and Porter (1981) reached similar conclusions, further questioning whether markets truly process information efficiently. In response, defenders of EMH (including Fama (1991)) argued that such volatility could be explained by *"time-varying risk premia,"* meaning the extra return investors expect for taking on risk isn't fixed; it changes over time depending on the state of the economy. However, the size and persistence of these deviations made it harder to defend the rationalist view.

The emergence of speculative bubbles poses an especially sharp challenge to the Efficient Market Hypothesis (EMH). A bubble refers to a situation where asset prices rise well above their intrinsic value (the worth justified by fundamentals such as earnings or cash flow), only to eventually collapse. Haykir and Yagli (2022), studying cryptocurrencies during the COVID-19 pandemic, found consistent evidence of bubble dynamics across digital assets, suggesting that this market is structurally prone to non-fundamental forces. Similarly, Cheah and Fry (2015) argued that Bitcoin's *"fundamental value is zero"*, with prices driven largely by speculation rather than by information efficiency.

Historical evidence from the U.S. housing market further reinforces this challenge. Case and Shiller (2003) showed how home prices in the early 2000s became increasingly detached from fundamentals such as rents and household income, warning of speculative excess well before the crash. When the subprime mortgage bubble burst in 2007–2008, it triggered a global financial crisis, demonstrating how narratives of ever-rising house prices, combined with lax credit standards and financial innovation, could sustain a large-scale mispricing. As Gorton (2008) later emphasized, this episode revealed that bubbles are not confined to emerging markets or niche assets but can destabilize the very core of the financial system.

Such findings highlight the difficulty of reconciling recurring, predictable bubbles with a framework based on rational expectations (the assumption that investors use all information optimally) and frictionless arbitrage (the idea that mispricing is swiftly corrected by trading).

Beyond excess volatility and bubbles, a wide range of so-called anomalies systematically contradict EMH predictions. Daniel, Hirshleifer, and Sun (2020) introduced a behavioral factor model to explain patterns such as post-earnings announcement drift (the tendency for stock prices to adjust only gradually to earnings news). They distinguished between short-horizon mispricing (often caused by limited investor attention) and long-horizon mispricing (driven by managers exploiting investor overconfidence). Crucially, these behavioral explanations accounted for anomalies that traditional risk-based models could not, suggesting that systematic psychological biases (such as misjudgment or overreaction) generate return patterns inconsistent with efficiency.

The debate between EMH advocates and critics reveals a deeper theoretical divide. On one side, scholars like Fama and French (1993, 2015) argue that multi-factor risk models (frameworks where returns are explained by exposure to several systematic risks) can account for return variation, interpreting anomalies as compensation for hidden risks. On the other side, behavioral finance researchers argue that these models fall short, since persistent abnormal returns remain unexplained without considering bounded rationality (the idea that investors have cognitive and informational limits). These back-and-forth underscores a fundamental question: are anomalies true evidence of inefficiency, or do they simply reveal gaps in existing models of risk?

Attention also complicates EMH's assumptions. Smales (2021) showed that surges in investor attention (proxied by Google search volumes during the COVID-19 crisis) were linked to higher market volatility, suggesting that attention is neither evenly nor rationally distributed. Barber, Huang, Odean, and Schwarz (2022) demonstrated how Robinhood (a commission-free trading platform popular among retail investors in the United States) and its "Top Mover" feature encouraged herding (concentrated trading in a small set of stocks), often followed by predictable reversals. Such findings indicate that market platforms themselves can shape price distortions, contradicting EMH's premise of unbiased information flows.

Equally problematic is the assumption of homogeneous beliefs (the idea that all investors interpret information in the same way). Giglio, Maggiori, Stroebel, and Utkus (2021) found strong and persistent differences in investor expectations, with portfolio decisions only weakly tied to those beliefs. Factors such as inattention, trading frictions, and lack of confidence prevented rational adjustment. These insights echo earlier work by De Long, Shleifer, Summers, and Waldmann (1990), who argued that *"noise traders"* (investors trading on sentiment rather than fundamentals) can survive and even influence prices, thereby weakening the corrective role that EMH assigns to arbitrage.

Taken together, these strands of evidence expose both the empirical and theoretical limits of EMH. Indeed, volatility is greater than fundamentals can explain, bubbles recur across asset classes, anomalies reflect systematic biases, and differences in attention and beliefs prevent information from being seamlessly reflected in prices. While EMH remains a valuable benchmark for understanding how markets might function under ideal conditions, it fails to capture the real dynamics of speculative markets.

This thesis therefore argues that the limits of EMH reveal the need for a broader framework. The persistence of anomalies, volatility, and speculative excess suggests that rational choice theory, which assumes that investors are fully rational and optimizing, is insufficient on its own to explain market outcomes. The intellectual elegance of EMH is offset by its inability to account for mispricing, collective psychology, and structural frictions. These shortcomings paved the way for behavioral finance, which incorporates insights from psychology to explain systematic departures from rationality. The following section explores behavioral biases such as overconfidence, loss aversion, and emotional contagion that distort investor judgment and destabilize markets.

A.2. Systematic Irrationality: How Behavioral Biases Undermine Market Efficiency

The evidence reviewed so far shows that EMH, while elegant as a benchmark, fails to account for persistent anomalies, volatility, and speculative dynamics. To understand why these inefficiencies arise and endure, it is necessary to look beyond market-level outcomes and examine the psychological mechanisms that drive investor behavior. Behavioral finance highlights several such biases, with overconfidence, loss aversion, emotional contagion, and herding receiving particular attention. These represent distinct but complementary dimensions of decision-making: cognitive (overconfidence), emotional (loss aversion), and social (emotional contagion and herding). Taken together, they explain how systematic departures from rational expectations emerge and destabilize financial markets.

Overconfidence represents one of the most robust behavioral biases. In finance, it refers to the tendency of investors to overestimate the accuracy of their knowledge and predictive ability. This often results in excessive trading and, ultimately, poorer performance. Barber and Odean (2001) provide strong empirical evidence, showing that overconfident investors, particularly men, trade significantly more and earn lower net returns compared to women. Earlier studies by Odean (1999) and Barber and Odean (2000) document similar results, with excessive turnover (trading far more frequently than is optimal) consistently leading to underperformance relative to benchmarks. This challenges EMH's premise of rational, utility-maximizing agents. Overconfidence also extends to corporate decision-making. Daniel, Hirshleifer, and Subrahmanyam (1998) argue that such biases distort aggregate prices, which firms exploit by strategically timing share issuance or repurchases. Building on this, Daniel, Hirshleifer, and Sun (2020) formalize a *"Financing*" factor, showing that managers arbitrage mispricing through financing decisions. Yet even with such strategies, distortions persist, highlighting a tension: whereas EMH predicts arbitrage should swiftly restore efficiency, behavioral evidence suggests that overconfidence can entrench inefficiencies even in the presence of informed actors.

Loss aversion, central to prospect theory, posits that individuals weigh losses more heavily than equivalent gains, systematically shaping investment behavior. A well-documented consequence is the disposition effect, the tendency to sell winners too quickly while holding onto losing positions. Odean (1998) captures this dynamic clearly, noting that investors "*demonstrate a strong preference for realizing winners rather than losers*" a pattern that undermines portfolio returns. This bias becomes particularly consequential in crises. Giglio, Maggiori, Stroebel, and Utkus (2021) study investor behavior during the March 2020 crash, when the COVID-19 pandemic triggered a global equity sell-off that erased more than 30 percent of value in weeks. Despite deteriorating beliefs, they find that "*portfolio adjustments were modest*" showing that even sharp changes in expectations did not translate into proportional shifts in equity exposure. Frictions such as inattention, infrequent trading, and lack of confidence compounded the effects of loss aversion, muting rational rebalancing. In contrast to EMH's prediction that belief revisions should be reflected immediately in portfolio allocations, the evidence reveals asymmetric responses to gains and losses that leave markets exposed to persistent inefficiencies.

Emotional contagion and herding amplify these distortions by transforming individual biases into collective dynamics. Investors often follow one another, whether through information cascades or simple imitation without independent assessment. In cryptocurrency markets, Haykir and Yagli (2022) find that herding intensified during the pandemic but diminished at bubble peaks, suggesting that shifts in risk aversion alter but do not eliminate collective behavior. Vidal-Tomás et al. (2019) and Papadamou et al. (2021) similarly show herding strengthens during downturns, when uncertainty is greatest. In equity markets, Klinge, Ouma, and Hendrikse (2025) portray Tesla as a case in which extreme price movements were driven more by narratives, social media, and retail flows than by fundamentals. Not all research, however, views herding as destabilizing. Gavriilidis (2024) finds that institutional herding may reduce volatility in low-sentiment environments, while Ahn et al. (2024) show that mimicry increases during recessions as investors seek coordination in the face of uncertainty. These perspectives suggest that herding can sometimes stabilize markets. However, such stability is precarious: once sentiment shifts, the same behavior often fuels cascades and amplifies volatility. Even when valuations appear tethered to fundamentals, as in Bruun and Boysøe's (2023) estimate that Tesla was only modestly (around 9 percent) overvalued in early 2023, narratives and trading flows can rapidly push prices far beyond intrinsic values. Moreover, Barber, Huang, Odean, and Schwarz (2021) show how Robinhood's "*Top Mover*" feature channeled attention into a handful of stocks, producing predictable reversals disconnected from fundamentals. These findings underscore that while herding may occasionally appear stabilizing, its prevailing legacy is to weaken corrective arbitrage and magnify market fragility.

These behavioral distortions also interact with attention in ways that amplify their impact. Smales (2021) shows that Google search volumes for terms such as "coronavirus" predicted heightened volatility and lower returns during the pandemic, reflecting fear rather than rational information processing. Liu and Tsyvinski (2018) similarly find that online attention predicts short-term cryptocurrency returns, underscoring the role of speculative sentiment. Such evidence demonstrates that biases are not merely individual traits but are intensified by collective attention cycles, producing systematic and non-random patterns in prices.

The dialogue between EMH proponents and behavioral scholars highlights the stakes of these findings. Defenders such as Fama and French (1993, 2015) maintain that multi-factor risk models account for return variation, framing anomalies as compensation for hidden risks. Their framework remains the mainstream in asset pricing, shaping both theory and practice. In contrast, Daniel, Hirshleifer, and Sun (2020) show that behavioral factors, such as post-earnings announcement drift and financing mispricing, explain anomalies more effectively than risk-based models, directly challenging EMH's sufficiency. The debate centers on whether anomalies are best understood as rational compensation for risk or as systematic mispricing rooted in psychology. Moreover, limits to arbitrage, as highlighted by Shleifer and Vishny (1997), demonstrate why rational investors may fail to correct inefficiencies, leaving room for biases to persist even in sophisticated markets.

Critically, behavioral literature offers both explanatory power and challenges. Its strength lies in linking anomalies to identifiable psychological mechanisms, providing coherent explanations for excess trading, valuation errors, and herding. Yet measurement remains difficult. Survey-based approaches, such as those used by Giglio et al. (2021), may underestimate the true scale of behavioral effects, and disentangling rational risk-taking from bias-driven behavior is inherently complex. Still, the consistency of findings across asset classes, from equities to cryptocurrencies, demonstrates the breadth of behavioral distortions and their capacity to undermine efficient price formation.

In conclusion, overconfidence, loss aversion, and emotional contagion with herding represent systematic departures from rational expectations that weaken the informational efficiency of markets. They amplify volatility, generate persistent inefficiencies, and sustain bubbles, all in inconsistent ways with the assumptions of EMH. While behavioral finance explains much of this through psychological biases, it largely treats them as individual-level errors. This thesis argues that a fuller understanding of market fragility requires moving beyond individual biases to examine how these dispersed tendencies coalesce into collective dynamics. Narrative economics addresses this gap by shifting attention from isolated biases to the broader stories that shape sentiment and coordinate collective behavior (Shiller, 2017).

A.3. From Bias to Story: Narratives as Coordinating Forces in Financial Markets

The conventional understanding of financial markets, rooted in the Efficient Market Hypothesis (EMH), holds that *"asset prices fully reflect all available information"* (Fama, 1970). Yet persistent anomalies, bubbles, and volatility show that fundamentals alone cannot account for market dynamics (Shiller, 1981; LeRoy & Porter, 1981). Behavioral finance has explained much of this by showing how investors deviate from rational expectations through biases such as overconfidence, loss aversion, and emotional contagion (Kahneman & Tversky, 1979; Barber & Odean, 2001). These perspectives highlight how individual errors and social imitation shape markets, but they do not fully explain how dispersed behaviors coalesce into enduring, large-scale movements. Narrative economics addresses this gap by shifting attention from psychological contagion to the broader *"contagious stories that have the potential to change how people make economic decisions"* (Shiller, 2017), which give shared meaning to events, coordinate expectations, and sometimes rival fundamentals in shaping valuations.

Shiller (2017) defines economic narratives as simple, emotionally resonant stories that spread socially and influence economic decisions, regardless of their factual accuracy. These narratives often acquire a viral quality, transmitting rapidly across populations and shaping perceptions of value, risk, and opportunity. Unlike individual errors, narratives operate as coordination mechanisms: they give dispersed investors a shared frame of reference, transform uncertainty into conviction, and move markets in tandem even without direct communication. In this sense, narratives can rival, and at times overwhelm, fundamentals in determining asset prices (Shiller, 2017).

Narratives gain their force by converting diffuse psychological tendencies into collective meaning. For example, ideas such as optimism about technological revolutions or the fear of missing out (FOMO) may begin as individual dispositions, but once embedded in compelling stories, they mobilize large numbers of investors simultaneously (Banerjee, 1992; Shiller, 2017). Biases may thus be seen as the raw material of speculation, while narratives act as the amplifier and organizer that channels them into large-scale market movements. Shiller’s (2017) epidemiological framing highlights how market stories can persist and spread, creating self-reinforcing cycles of enthusiasm or panic.

Recent empirical research strengthens this conceptual account. Taffler, Agarwal, and Obring (2021, 2024) show that emotions embedded in financial media narratives explain a substantial share of market dynamics. Their media-derived emotion variables account for up to 52% of stock market returns and 67% of uncertainty during crisis periods, findings that hold robustly across different episodes. This suggests that collective emotions, transmitted through stories, are not peripheral but central to price formation. Methodologically, their approach also advances the measurement of narratives, moving beyond simple topic models toward emotion-based textual analysis that better captures the affective power emphasized by Shiller (2017). These results underscore how narratives operate not only as cultural scripts but also as quantifiable forces shaping valuation, further challenging the primacy of fundamentals.

This conceptual shift creates tension with rationalist perspectives. From the standpoint of EMH, narratives might be dismissed as “*noise*” or temporary deviations from fundamental value, to be corrected by arbitrage (Fama, 1991). Yet the theory of limits to arbitrage (Shleifer & Vishny, 1997) shows why mispricing can persist: arbitrage is often costly, risky, or poorly synchronized. Within such gaps, narratives can sustain valuations over extended periods. Still, a critical question remains unresolved: are narratives independent drivers of bubbles, or simply by-products of speculation already in motion?

The methodological challenges of narrative economics sharpen this debate. Measurement typically relies on proxies such as textual analysis, Google Trends, or social media data (Smales, 2021; Liu & Tsyvinski, 2018), which capture patterns of attention but not necessarily conviction or trading behavior. A further limitation lies in definitional scope: if any recurring theme is treated as a narrative, the concept risks becoming so expansive that it loses analytical clarity (Shiller, 2017). This expansiveness makes it difficult to distinguish between narratives with genuine economic force and background “*noise*.” Causality adds to another layer of difficulty. Do stories drive prices, or do price movements simply generate stories that rationalize them after the fact? Without clear identification strategies, narrative economics risks slipping into post hoc description rather than causal explanation. Finally, operationalization remains uneven: while Shiller (2017) emphasizes the contagious, viral quality of narratives, empirical studies often reduce them to static text measures, missing the dynamics of diffusion and decay that give the concept its distinctive force. Taken together, these critiques underscore that narrative economics, while intellectually powerful, remains methodologically fragile and should be seen as a complement to, rather than a replacement for, behavioral or rationalist approaches.

In conclusion, narratives represent more than reflections of sentiment; they are organizing forces that transform dispersed psychology into coordinated market action, often rivaling fundamentals in shaping valuations (Shiller, 2017). At the same time, the framework's definitional ambiguity, difficulties of measurement, and unresolved questions of causality mean that its contribution remains contested. Narrative economics shows that financial markets are not only arenas of information aggregation but also spaces of meaning-making and collective storytelling. Yet without sharper methodological tools, its explanatory power risks dilution.

Taken together with the critiques of efficiency in Part A, this analysis points to a broader conclusion: market fragility cannot be explained solely by rational expectations or isolated behavioral biases. Volatility, bubbles, and anomalies (Part A.1) reveal the structural cracks in EMH, while psychological distortions (Part A.2) and narrative contagion demonstrate how collective forces amplify those cracks into systemic dynamics.

This thesis argues that financial theory must move beyond classical rationalist models and embrace a framework that integrates behavioral finance, narrative dynamics, and structural frictions. Only by recognizing markets as both psychological and social constructs can we explain their persistent inefficiencies and recurrent instability.

Building on this foundation, we now turn to digital spaces where investor behavior unfolds. Platforms are not the cause of market inefficiencies in themselves, but they provide environments in which psychological biases, attention cycles, and herding dynamics are amplified and made more visible. What Part A explained as the underlying reasons for market fragility, the next part explores as the social arenas where these forces take shape in practice. If narratives explain why stories matter, the behavior of digital communities shows how those stories spread, gain traction, and transform into collective market movements.

B. DIGITAL COMMUNITIES AND THE INFRASTRUCTURES OF INVESTOR PSYCHOLOGY

B.1. Attention as Currency: How Digital Communities Create Visibility and Shape Investor Focus

In digital finance, attention is one of the scarcest resources. Goldhaber (1997) and Davenport and Beck (2001) argued that the real limit in modern economies is not access to information but the ability to concentrate on it. Financial markets make this especially clear. With thousands of securities available, investors cannot process all signals. What they notice, and what they ignore, strongly influences trading outcomes.

Barber and Odean (2008) demonstrated this in their influential study. They showed that retail investors are disproportionately drawn to "*attention-grabbing*" stocks. These are securities that appear in headlines, show unusual trading volume, or undergo extreme one-day returns. Institutional investors, with more resources and broader mandates, are less affected by this constraint. In digital spaces the gap becomes even wider. Retail investors are confronted with a continuous flow of cues such as memes, hashtags, and viral posts that redirect attention constantly. What once depended on newspapers or television broadcasts now circulates through online communities at high speed.

This makes attention not only scarce but also performative. Investors do not simply receive information passively. They also play a role in producing visibility. A viral meme or a popular Reddit thread does not merely reflect existing interest. It creates attention by drawing others to look and participate. In this sense, attention operates like a currency. The more a stock or token is discussed, the more visible and valuable it appears, regardless of whether its fundamentals have changed.

Empirical research supports the claim that bursts of attention influence markets. Smales (2021) and Liu and Tsyvinski (2018) find that search activity and Twitter mentions predict short-term volatility in equities and cryptocurrencies. Cevik et al. (2022) show that sentiment shapes these effects. Positive attention correlates with temporary gains, while negative attention deepens losses. At the same time, Da, Engelberg, and Gao (2014) caution that these measures capture visibility but not conviction. A surge in search activity may reflect curiosity or amusement rather than an intent to trade.

Theoretical interpretations diverge. Rationalist scholars, following Fama (1970), regard attention-driven trading as temporary noise that arbitrageurs eventually erase. Behavioral finance takes the opposite view. It treats limited attention as a structural bias that filters which information enters the market in the first place (Barber and Odean, 2008). Shleifer and Vishny's (1997) limits-to-arbitrage framework goes further, showing that even obvious mispricings can persist because correcting them is costly, risky, and difficult to time. Finally, sociological and media studies emphasize that attention is socially produced in digital environments. Gillespie (2014) notes that online spaces structure visibility in ways that turn attention into a form of currency. Tufekci (2015) shows that virality is actively cultivated rather than accidental. Shiller's (2017) account of "*narrative economics*" reinforces this point by illustrating how stories harness collective focus. From this perspective, memes, hashtags, and influencer posts do not simply highlight existing interest; they manufacture visibility and recirculate it as a symbolic asset within investor communities.

Treating attention as a currency has critical implications. Like any medium of exchange, it is unevenly distributed and subject to inflationary cycles. A sudden burst of search interest or viral content may not always signal conviction, but once collectively recognized it alters market dynamics. Empirical work confirms this. Smales (2021) and Liu and Tsyvinski (2018) find that attention measures forecast short-term returns and volatility, while Cevik et al. (2022) show that sentiment amplifies both rallies and downturns. In these cases, the "value" of attention is realized in price movement: visibility itself becomes convertible into temporary market power.

To conclude, attention functions simultaneously as a scarce resource, a behavioral bias, and a currency that organizes participation. Rationalist theories underestimate its force, while behavioral accounts risk overlooking its performative dimension. Treating attention as currency reveals that visibility is not a marginal distortion but a structural driver of digital speculation. It determines which assets enter the spotlight, sets the tempo of trading, and coordinates dispersed investors into shared focus. Crucially, once attention circulates collectively, it invites imitation, endorsement, and reinforcement, transforming the flow of attention into coordinated market action.

This thesis argues that in digital finance, visibility itself has become a market force. This challenges models that rely only on rational expectations or individual biases. Attention is not just temporary noise that disappears through arbitrage; it is a core feature of speculative markets that shapes how value is created and contested. (*The focus here has been on the user inside digital platforms, showing how investors produce, amplify, and circulate attention in ways that reshape market outcomes. The role of platforms themselves, and the infrastructures that condition which stories and signals thrive or fade will be addressed in Part 2.D).*

The next section builds on this claim by examining how attention as currency is “spent” through networked herding, where social cues and shared focus crystallize into synchronized speculation and amplify the fragility identified.

B.2. Networked Herding: How Digital Communities Coordinate Speculation

In digital communities, herding no longer emerges only from observing trades or price charts. It increasingly takes the form of platform-enabled coordination, where investors orient their behavior around visible signals of group sentiment. Features such as upvotes, trending hashtags, memes, and reposted commentary transform scattered individual opinions into synchronized collective momentum (Gillespie, 2014; Cookson et al., 2023). Unlike traditional herding, which unfolded more gradually through price observation, networked herding operates in real time, reducing the costs of coordination and accelerating cascades (Banerjee, 1992; Shiller, 2017).

Classical theories of herding described “*information cascades*” in which investors, uncertain about fundamentals, imitate earlier decisions rather than rely on their own signals (Banerjee, 1992). Behavioral finance added an emotional dimension, showing how biases and sentiments magnify collective tendencies (Shiller, 2017). Digital communities intensify both processes. Emotionally charged content tends to attract greater visibility, and user interactions such as likes, comments, or reposts blur the line between social approval and market signals. As a result, herding online is not simply a spontaneous crowd reaction but a coordinated form of speculation, where social cues provide as much guidance as fundamentals.

This form of networked herding is distinct from other collective phenomena. Emotional contagion and classical cascades explain how imitation spreads through observation, and narrative economics shows how shared stories coordinate expectations (Shiller, 2017). Networked herding combines these elements but adds a crucial dimension: the rapid circulation of visible markers of sentiment that lower the threshold for coordination. When investors see a meme repeated thousands of times or a hashtag trending across platforms, they are not only consuming information but also registering social proof of collective participation. This explains why herding in digital markets escalates faster and more forcefully than in earlier settings.

Empirical studies confirm these effects. Cookson et al. (2023) document that surges in online engagement around meme stocks on Reddit and Twitter coincided with synchronized trading behavior. Long et al. (2022) find that the intensity of online discussions predicted stock returns, demonstrating that digital conversation translated into measurable market outcomes. In cryptocurrency markets, Haykir and Yagli (2022) show that online community clustering intensified during downturns, producing stronger risk-averse cascades. Across these contexts, herding is not merely correlated with volatility but often central to its amplification.

Yet studying these dynamics presents methodological challenges. Counting posts, likes, or search queries can easily mistake curiosity for conviction. Long et al. (2022) emphasize that the cultural content of communication, including memes, slang, humor, and insider references, is crucial for understanding how coordination is sustained. Without this interpretive lens, quantitative measures risk flattening meaningful dynamics into mere background noise.

The implications extend beyond methodology. Digital communities encourage short-termism and hype cycles by privileging content that is emotionally engaging rather than accurate or balanced. They also foster what Sunstein (2001, p. 65) defines as echo chambers, "*environments in which people hear only voices that echo their own*" reinforcing confirmation bias and insulating groups from corrective perspectives. Cookson et al. (2023) show that such self-reinforcing dynamics amplify speculative enthusiasm and mute dissent. Combined with low entry costs and anonymity, these conditions allow participation to accelerate rapidly. As a result, speculative narratives can inflate prices well beyond fundamentals before collapsing just as abruptly.

In conclusion, networked herding illustrates how online communities transform individual biases into collective dynamics. Classical models of imitation explain part of the story, but the distinct feature of digital herding lies in its visibility and cultural reinforcement. Investors do not only follow price movements; they follow memes, influencers, and the visible markers of collective approval that signal belonging. This makes herding in digital finance not simply a by-product of uncertainty but a socially structured form of speculation.

The evidence suggests that online herding is not peripheral but central to contemporary market dynamics. It accelerates cascades, magnifies volatility, and sustains speculative bubbles by embedding coordination into the everyday practices of digital communities. Understanding these dynamics requires treating markets not only as arenas of rational choice but also as social spaces where trust, identity, and symbolic participation shape outcomes.

The next step is to consider the growing role of financial opinion leaders and online influencers. These figures shape both what gains attention and how collective behavior unfolds, making them central to the dynamics of digital speculation. Their influence will be explored in greater detail in the following section.

B.3. Influencers as Opinion Leaders: The cultural authority of financial influencers compared to traditional analysts.

The concept of opinion leadership, defined as the disproportionate influence that certain individuals exert on the behavior of others, is central to understanding information diffusion in markets. Iyengar, Van den Bulte, and Valente (2010) describe opinion leaders as individuals whose positions within networks allow them to act as "*multipliers*" of influence, transmitting ideas more effectively than peers. In financial contexts, this role was historically embodied by credentialed analysts, brokers, and institutions. Their authority rested on professional accreditation, fiduciary duties, and regulatory oversight, which gave their advice an aura of legitimacy and accountability (CFA Institute, 2024).

The rise of financial influencers ("*finfluencers*") represents a fundamental departure from this model. Unlike traditional analysts, their authority is not grounded in institutional credentials but in perceived authenticity, cultural affinity, and the ability to mobilize large online audiences (Hull & Qi, 2024). They communicate in ways that are free, accessible, and visually engaging, often using personal narratives of success or failure to establish credibility. This delivery style fosters parasocial relationships, one-sided yet affectively charged bonds, which enhance trust and identification (CFA Institute, 2024). Younger cohorts, particularly Gen-Z, are drawn to finfluencers not only for their accessibility but also because they distrust established advisors, often viewed as costly, opaque, or conflicted by commission structures. Hull and Qi (2024) further demonstrate that finfluencer popularity is not arbitrary but linked to observable traits such as past performance, shared identity markers, and distinct trading styles. This points to a hybrid authority that blends cultural resonance with selective signals of expertise.

This mode of opinion leadership differs in important ways from both narratives and herding, which were discussed earlier. Narratives operate at the level of ideas: stories that frame events in emotionally compelling ways and spread virally across populations (Shiller, 2017). Herding, by contrast, reflects collective behavioral dynamics, where investors imitate peers or aggregate signals without independent assessment (Banerjee, 1992). Finfluencers synthesize elements of both but add a critical dimension through the personalization of authority. They embody stories through a recognizable figure and channel herding tendencies by acting as focal points for dispersed attention. In this sense, finfluencers transform diffuse psychological and social dynamics into a coherent source of perceived guidance.

The literature diverges sharply on how to interpret this shift. From a rationalist perspective, consistent with the Efficient Market Hypothesis, influencer content is mostly "*noise*" that adds little to information environments and should be arbitraged away (Fama, 1991). Analysts and regulators warn that finfluencers often engage in undisclosed sponsorships, speculative "*tips*," or thinly veiled marketing, practices that risk misleading audiences and encouraging excessive risk-taking (CFA Institute, 2024). From this standpoint, finfluencers erode informational quality rather than enhance it. By contrast, behavioral finance frames finfluencers as systematic opinion leaders whose capacity to capture attention and generate trust makes them central to market sentiment. Herding becomes intensified when individuals substitute influencers' cues for their own judgment, a dynamic amplified by social media echo chambers that reinforce confirmatory biases and sustain collective enthusiasm (Cookson, Engelberg & Mullins, 2023). Evidence from cryptocurrency markets illustrates this tension: Liu and Tsyvinski (2018) show that sentiment-driven attention predicts returns and volatility, while Haykir and Yagli (2022) find that community-driven cascades deepen downturns. In this light, influencers are not incidental noise but catalysts of sentiment cycles that can meaningfully move prices.

Methodologically, research into finfluencers reflects the same fragilities noted in earlier sections. Sentiment analyses of Twitter, Reddit, and YouTube have been used to capture collective mood (Bollen et al., 2011; Long et al., 2022), while Hull and Qi (2024) employ more rigorous causal designs, showing that exogenous variation in influencer exposure affects portfolio allocation. Yet important limitations remain. Engagement metrics such as likes, shares, or follower counts capture reach but not persuasion (Romero et al., 2010). Sponsorship disclosure is inconsistent, blurring the boundary between independent advice and paid promotion (CFA Institute, 2024). Community-specific cultures complicate generalization: humor, memes, and slang often carry meanings invisible to sentiment models but crucial to credibility and uptake (Long et al., 2022). Data access also constrains analysis, since platforms mediate APIs in ways that limit independent scrutiny of influence dynamics.

The broader implications for financial communication are profound. Traditional analysts derive legitimacy from formal expertise and oversight, while finfluencers derive authority from authenticity and relatability. This inversion reshapes how informational authority is constituted: expertise no longer commands attention, but attention itself constructs expertise. While finfluencers may democratize access to financial knowledge by lowering entry barriers, they simultaneously expose investors to hype-driven trust, promotional manipulation, and heightened vulnerability to speculative bubbles. Regulatory frameworks lag behind, struggling to classify or constrain finfluencer activity across jurisdictions (CFA Institute, 2024).

In sum, the rise of finfluencers marks a structural shift in opinion leadership within financial markets. Unlike traditional analysts, whose authority rested on credentials and oversight, influencers derive power from relatability, authenticity, and their ability to mobilize digital communities. This inversion challenges the assumption that expertise and regulation anchor information quality. Instead, attention itself constructs expertise, allowing cultural authority to substitute for formal legitimacy. The result is a mode of opinion leadership that personalizes narratives, concentrates herding tendencies, and exposes investors to both empowerment and new vulnerabilities.

Stepping back, the evidence across Part B points to a broader transformation. Attention now operates as a currency, herding unfolds as a socially coordinated practice within digital communities, and influencers serve as focal points for dispersed psychology. Together, these findings extend the argument developed in Part A. Where Part A concluded that financial theory must move beyond rationalist models by integrating behavioral and narrative dynamics, Part B has shown how these dynamics unfold within digital communities, reshaping investor psychology and amplifying speculative tendencies. This thesis argues that these processes reveal markets as social spaces where visibility, trust, and cultural authority are as consequential as fundamentals in shaping outcomes.

The next step is to examine how these dynamics materialize in practice. The following part, turns to contemporary manifestations of speculative dynamics, focusing on cases such as stocks and cryptocurrencies. These examples demonstrate how the forces of attention, herding, and influencer authority described in Part A and B converge in real-world markets, producing episodes of volatility and fragility that test the limits of classical finance and confirm the need for an expanded theoretical framework.

C. CONTEMPORARY MANIFESTATIONS OF SPECULATIVE DYNAMICS

C.1. Cryptocurrencies as Psychological Laboratories: Bitcoin, Dogecoin, and the Fragility of Bubble Theory

Economists remain divided on how to interpret extreme price swings: some frame them as rational responses to uncertainty, risk premia, or shifting expectations, while others emphasize speculative excess. In the case of Bitcoin, Cheah and Fry (2015) argue bluntly that its "*fundamental value is zero*" since it generates no cash flows and offers little productive use, making price movements largely a product of sentiment. Advocates, by contrast, present Bitcoin as "*digital gold*" where scarcity and adoption provide a novel foundation for value (PwC, 2021). The critical question is whether this scarcity narrative constitutes a genuine anchor or simply another story in the speculative cycle. If gold itself partly derives value from collective belief rather than intrinsic yield, then Bitcoin's analogy may reveal less about fundamentals than about the power of narrative to masquerade as fundamentals (Gerlach et al., 2019).

Bitcoin: recurring bubbles and contested anchors

Bitcoin epitomizes this ambiguity. Its capped supply of 21 million coins underpins the scarcity narrative, yet its price history is dominated by repeated boom-and-bust cycles. Gerlach, Demos, and Sornette (2019) identify "*three main long bubbles and ten additional smaller peaks*" between 2012 and 2018, each followed by massive crashes. Wheatley et al. (2018) likewise document "*near-universal super-exponential growth*" patterns across four major Bitcoin bubbles, with prices accelerating faster than exponential growth models before collapsing. According to them, these dynamics stand in stark opposition to the efficient markets perspective. Such findings undermine the claim that scarcity alone grounds valuation, suggesting instead that Bitcoin's "fundamentals" are fragile anchors, easily overwhelmed by speculative feedback loops.

Dogecoin: the pure narrative coin

Dogecoin (*another cryptocurrency, which in August 2025 ranked eighth in global market capitalization*) makes the critique sharper still. Created in 2013 as parody, it has unlimited supply, minimal technological innovation, and no institutional adoption to speak of. Its valuation has instead been sustained almost entirely by meme culture, retail enthusiasm, and the amplification of celebrity endorsements (Nani, 2022). Oncu (2021) identifies ten distinct bubbles in Dogecoin during 2020–21, many sparked by viral campaigns or Elon Musk's tweets. At its peak in May 2021, Dogecoin's market capitalization reached $88 billion, comparable to General Motors, despite lacking any scarcity mechanism or productive utility. In Shiller's (2017, 2019) terms, this is a "*pure narrative economy*" where contagious stories and community identity alone sustain multi-billion valuations. Dogecoin's rise raises the uncomfortable possibility that speculation untethered from fundamentals is not an aberration but a viable mode of valuation in the digital age.

Cryptocurrencies as a whole also expose the fragility of the existing tools for bubble detection. Econometric tests such as the GSADF procedure (Phillips, Shi & Yu, 2015) can identify explosive dynamics in Bitcoin and Dogecoin, but these methods presuppose a fundamental baseline from which "explosiveness" is measured. In assets like Dogecoin, where no such anchor exists, the category of "bubble" risks becoming tautological: the asset is a bubble because it diverges from fundamentals, yet fundamentals are absent to begin with. Narrative and sentiment-based approaches (Taffler, Agarwal & Obring, 2021, 2024) offer richer texture by tracing cultural dynamics, but as Roos and Reccius (2021) caution, text-mined "topics" often capture curiosity rather than conviction. Thus, crypto markets do not merely display bubbles, they test the very adequacy of the tools designed to study them.

In sum, Bitcoin and Dogecoin dramatize how cryptocurrencies function as "*psychological laboratories*." They reveal that speculative logics can thrive both when quasi-fundamental anchors exist (Bitcoin's scarcity) and when they are absent (Dogecoin's memes). The deeper challenge is conceptual: if assets with no intrinsic cash flows or scarcity can sustain multi-billion valuations, then the very definition of a bubble as "*deviation from fundamentals*" becomes unstable. Cryptocurrencies do not simply illustrate familiar bubble dynamics; they demonstrate that our analytical categories, grounded in rationalist models, are inadequate.

This thesis argues that narratives, attention, and collective belief have become constitutive forces of value in digital markets. Fundamentals still matter, but they no longer hold exclusive authority. Instead, stories and visibility can generate and sustain prices in ways that destabilize the very distinction between fundamentals and speculation. In their most extreme form, as with Dogecoin, cryptocurrencies show that psychology alone can fabricate value at scales once thought impossible. This makes them not only examples of bubbles but decisive evidence that efficiency is fragile and that financial theory must integrate psychological and social dynamics to remain credible.

Seen in this light, cryptocurrencies are not anomalies at the margins of finance but central cases that reveal how markets operate under conditions of narrative-driven speculation. They provide the bridge to equity markets, where meme-stock surges such as GameStop and AMC suggest that the same dynamics of collective belief and cultural authority are reshaping established asset classes.

C.2. Meme Stocks and Beyond: Retail Coordination, Corporate Strategy, and the Limits of the Category

If cryptocurrencies highlight speculation in their purest form, equity markets reveal how similar dynamics reshaped even established companies. The early 2020s saw extraordinary surges in so-called "*meme stocks*" destabilizing long-standing assumptions about efficiency and valuation (Aggarwal et al., 2024). The question was whether these rallies were dysfunctional bubbles detached from fundamentals, or rational responses to uncertainty, new trading technologies, and collective action. The concept of a "*rational bubble*" captures this ambiguity: investors knowingly buy overpriced assets not for dividends but in the expectation of resale at higher prices (Hirano & Toda, 2024). Yet the claim that meme stocks "democratized" finance, having retail overturning Wall Street, has proven misleading. A closer look suggests risks were redistributed downward, while institutional actors ultimately adapted and profited.

GameStop: Narrative and Predatory Trading

GameStop (GME), a declining video game retailer, became the emblematic case. With weak fundamentals and heavy short interest, hedge funds widely anticipated collapse (Aggarwal et al., 2024; Lucchini et al., 2021). Yet in early 2021, retail traders coordinated on zero-commission platforms like Robinhood and on Reddit's WallStreetBets (WSB) (*an online forum*). There, memes, "YOLO" screenshots, and a David-versus-Goliath framing transformed a technical short squeeze into a cultural crusade.

The mechanics were decisive. As prices spiked, hedge funds such as Melvin Capital were forced to cover, losing more than $4 billion (Aggarwal et al., 2024). Scholars called this the first case of "*predatory trading*" by retail investors, where collective enthusiasm destabilized institutional strategies (Lucchini et al., 2021). Yet Shiller's (2017) "*narrative economics*" shows why the frenzy persisted: the David-versus-Goliath story spread contagiously, binding participants to hold even as fundamentals failed to justify valuations.

AMC: Meme Identity and Corporate Strategy

AMC Entertainment (AMC), a debt-laden cinema chain, followed a similar trajectory. Also heavily shorted, AMC became one of WSB's most discussed tickers between 2019 and 2021 (Aloosha et al., 2024). Retail buying triggered a short squeeze, but AMC's distinctiveness lay in management's response. CEO Adam Aron actively leaned into the meme identity: engaging retail investors on Twitter, hosting special screenings, and issuing AMC Preferred Equity Units (APEs) when conventional equity issuance faced resistance (Aggarwal et al., 2024). The firm also exploited inflated prices to raise billions in at-the-market offerings, stabilizing its finances.

The dramatic rallies in GameStop and AMC illustrate how behavioral forces can dominate markets. Behavioral finance highlights "*irrational attention*" (Aloosha et al., 2024), peer and echo chambers that amplified bullish sentiment within Reddit and Robinhood communities (Cookson et al., 2022). Unlike the theoretical accounts discussed earlier, this was a concrete manifestation: closed loops of reinforcement where enthusiasm circulated with little exposure to dissent. Barber et al. (2022) show Robinhood flows closely tracked WSB chatter, reinforcing mispricing through platform design. Rationalists counter that AMC and GameStop converted speculative rallies into fundamental survival (Aggarwal et al., 2024). Yet empowerment rhetoric collapses: shareholder participation fell, many retail traders were left with losses, while hedge funds, market makers, and passive funds adapted and often profited. The David-versus-Goliath story looks inverted: retail bore the volatility while institutions reasserted control.

Tesla: The Breaking Point of the Meme Stock Concept

Tesla (TSLA) complicates the picture. Some commentators grouped it with GameStop and AMC, citing Musk's cult-like following and market-moving tweets. Scholars resist, pointing to profitability from 2023 and leadership in EV markets (Aggarwal et al., 2024; Klinge et al., 2025). Still, behavioral perspectives describe it as "*mementum*," retail enthusiasm propelled by Musk's celebrity and expansive narratives about AI, robotics, and renewable energy (Klinge et al., 2025). Musk's tweets routinely moved prices, showing how charismatic leadership can blur fundamentals and sentiment.

This blurring links Tesla back to the broader issue of opinion leadership introduced in Part B.3. Musk does not operate only as a corporate executive but as a cultural influencer whose persona commands trust, shapes narratives, and mobilizes retail enthusiasm in ways similar to finfluencers on YouTube or Reddit. His role illustrates how influencer-style authority now extends into corporate leadership itself, demonstrating that market outcomes can hinge less on fundamentals than on the symbolic capital of individual figures.

But Tesla also forces a critical question: does its inclusion dilute the very idea of meme stocks? If firms with strong fundamentals qualify simply because narrative and retail enthusiasm play a role, then the category risks incoherence. Either meme dynamics now extend even to industry leaders, or "meme stock" is little more than an ex-post label for price surges that defy easy explanation. Tesla thus becomes a hinge: it destabilizes the meme stock category at the very moment when financial journalism begins to apply it even more broadly.

Beyond Meme Stocks: The 2025 Climate

By mid-2025, financial journalism suggested meme-style speculation had broadened across sectors. Academic work remains scarce, but major outlets consistently framed new surges as "meme" cases:

- Paramount Skydance (Aug. 2025): Barron's reported a 37 percent one-day jump despite no merger developments, attributing gains to retail enthusiasm and chatter.
- Opendoor (July 2025): Reuters and MarketWatch described a 400 percent monthly surge, driven by AI-pivot hype, high short interest, and online coordination.
- Krispy Kreme (July 2025): The Wall Street Journal and Reuters noted a 30 percent spike on nostalgic enthusiasm, before retracing when fundamentals lagged.
- Kohl's (July 2025): Reuters reported shares more than doubled intraday, closing up 40 percent as bullish turnaround narratives spread despite weak operations.

These episodes echo the mechanics of GameStop and AMC: cascades of attention, herding dynamics, and narrative contagion pushing prices far from fundamentals. Yet they also expose a methodological problem. As with crypto, bubble-detection frameworks assume valuations should remain tethered to fundamentals—but even AMC's CEO conceded its stock price was "*not correlated with any fundamental changes*" (Aggarwal et al., 2024). If certain assets are untethered by design, then labeling them as "bubbles" risks circularity: we end up measuring deviations from anchors that no longer function as anchors at all.

Speculative dynamics once confined to fringe assets are now entrenched in mainstream equity markets. GameStop and AMC reveal how retail coordination can destabilize institutions, but also how quickly institutions adapted and profited while retail bore losses. Tesla, meanwhile, destabilizes the meme stock concept itself: either it shows that meme-like forces pervade all markets, or it exposes "meme stock" as analytically incoherent. And the 2025 surges reinforce that such speculation is not episodic but structural.

The critique is therefore twofold. First, the story of retail empowerment proves misleading: rather than democratizing finance, meme stocks shifted risk onto small traders while institutional actors adapted and retained control. Second, the categories we use begin to break down. Crypto already unsettled what we mean by a "bubble," and meme stocks push the challenge further, straining both efficiency theory and the bubble concept itself. If prices can be sustained by memes and narratives rather than fundamentals, then calling them "mispriced" misses the point. The sharper hypothesis is that meme stocks are not just speculative anomalies but indicators that financial theory itself is struggling to capture the logic of digitally mediated speculation. This raises a broader question: how do today's digitally mediated bubbles compare with earlier episodes of financial mania, and what structural continuities or novel features can be discerned?

C.3. Historical vs. Digital Bubbles: Continuities and Ruptures in Speculative Dynamics

Although financial bubbles take different forms across eras, they follow consistent structural patterns. Classic accounts describe them as unfolding in stages of displacement, boom, euphoria, and collapse (Kindleberger & Aliber, 2011; Haykir & Yagli, 2022). Historical episodes such as Tulipmania, the Mississippi Bubble, and the Dot-com boom illustrate these recurring dynamics. Comparative study suggests that while the objects of speculation change, the behavioral underpinnings of excess optimism, herd behavior, and narrative contagion remain remarkably constant.

These stages recur because certain structural continuities persist across eras. Herding, or the tendency to follow others' trades rather than individual judgment, has been repeatedly documented (Pompian, 2017; Haykir & Yagli, 2022). Some studies note that herding may weaken or even reverse during extreme peaks (Chrisostomides, 2022), yet its persistence across markets underscores its explanatory power. Narrative contagion is another constant. Shiller (2017) stresses that emotionally resonant stories, whether about tulips, railroads, or "*new economies*" spread virally and sustain speculative enthusiasm. Media amplification contributes as well. Eighteenth-century pamphlets and today's financial press both fuel feedback loops that reinforce rising prices (Shiller, 2000; Li et al., 2022). Finally, the "*limits to arbitrage*" thesis (DeLong et al., 1990) continues to explain why mispricing endures: rational traders face synchronization and noise-trader risk that prevent them from correcting exuberance.

If these continuities provide the scaffolding, the Dot-com bubble serves as a critical hinge case, showing both how old logic persisted and how new frictions emerged. Ofek and Richardson (2001) show that Internet stock valuations were sustained not simply by optimism but also by institutional frictions. Short-sale constraints prevented pessimists from expressing their views, while heterogeneous beliefs among investors amplified the imbalance. The eventual collapse was triggered by the expiration of IPO lockup agreements, which released a flood of insider shares and introduced new sellers into the market, exposing how structural barriers rather than fundamentals had propped up prices. Brunnermeier and Nagel (2004) further challenge the stabilizing role of rational speculators, demonstrating that hedge funds did not correct mispricing but instead rode the bubble, remaining overweight in technology stocks through the upturn and cutting positions only just before the crash. Narrative amplification was equally central. Cooper, Dimitrov, and Rau (2001) find that firms merely adding ".com" to their names experienced average abnormal returns of 74 percent in the days following announcement, regardless of their actual involvement with the Internet. This was evidence of investor mania in which cultural cues substituted for fundamentals. Rationalist perspectives, however, suggest bubbles need not imply irrationality. Pástor and Veronesi (2005, 2008) develop a general equilibrium model in which technological revolutions produce bubble-like patterns ex post. As investors learn about productivity, expected cash flows rise, but so too does systematic risk, generating high valuations that later collapse. Taken together, these accounts reveal the Dot-com era as a laboratory for the enduring logics of limits to arbitrage, narrative contagion, and herding, dynamics that persist even when speculation is fueled by genuine technological innovation.

If the Dot-com boom complicated older patterns, digital markets go further by introducing qualitatively novel accelerants. Online platforms enable real-time coordination and retail activism at a scale inconceivable in historical bubbles. Reddit communities, for instance, have triggered short squeezes through mass action, creating new forms of bottom-up pressure on institutions (Nani, 2022; CFA Institute, 2024). Algorithmic amplification further intensifies contagion. Search queries, trending topics, and platform metrics reliably predict short-term returns in crypto and meme assets (Liu & Tsyvinski, 2018; Haykir & Yagli, 2022). Influencer-driven trust adds another unprecedented layer, as "finfluencers" cultivate parasocial ties with followers and displace the authority of analysts with a performance of authenticity (CFA Institute, 2024). Meme culture deepens the rupture. Assets like Dogecoin derive value less from scarcity or utility than from their symbolic appeal, embedding finance in cultural play (Nani, 2022).

These features raise a central question: do digital bubbles merely accelerate old dynamics, or do they represent a fundamental break with history? Scholars remain divided. Some argue that cryptocurrencies and meme stocks largely replicate classical bubbles, with GSADF tests showing familiar explosive signatures (Haykir & Yagli, 2022; Corbet et al., 2018a). Others see a qualitative break. Dogecoin's rise as a "*money-meme hybrid*" or the deliberate coordination of retail short squeezes cannot be mapped neatly onto earlier manias without loss of specificity (Nani, 2022; CFA Institute, 2024). Even Shiller's (2017) narrative economics, which treats stories as timeless contagions, is strained by environments where narratives are not just spread but algorithmically amplified and tied to online identity performance.

This debate is mirrored in methodological tensions. Studying historical bubbles usually means piecing together archival evidence, rich in context but often imprecise (Kindleberger & Aliber, 2011). Digital bubbles, by contrast, let researchers analyze millions of Reddit posts or search queries in real time (Long et al., 2022). That precision, however, comes with a catch: online chatter can just as easily reflect passing curiosity as real conviction (Roos & Reccius, 2021). This methodological challenge reflects a deeper problem: traditional tools assume markets stay tied to fundamentals. When even AMC's CEO admitted its stock price was "*not correlated with any fundamental changes*" (Aggarwal et al., 2024), it illustrated how the anchors assumed in historical bubble analysis can vanish entirely. Classical frameworks presuppose a tether to fundamentals, yet cases like AMC shows how modern speculation can detach so fully that "bubble" becomes a label of convenience rather than an analytical category.

Tesla complicates the picture in a different way. Its profitability and technological leadership set it apart from distressed meme stocks, but its valuation still surged on narratives amplified by Musk's persona and online communities (Klinge et al., 2025). This case underscores not simply the ambiguity of the meme-stock label, but the difficulty of applying traditional bubble categories at all. If a market leader can be driven upward by the same narrative and attention dynamics as fragile firms like AMC or GameStop, then the historical distinction between speculative manias and legitimate growth stories becomes increasingly unstable.

In sum, the comparison between historical and digital bubbles reveals both continuity and rupture. The familiar drivers of speculation (herding, contagious stories, and limits to arbitrage) remain as powerful as ever, as the Dot-com case illustrates. What has changed is the infrastructure through which these forces operate. Platforms designed for virality, algorithmic amplification, and influencer authority have compressed the tempo of bubbles from months into hours, embedding speculation into the rhythms of online culture. The promise of democratized finance rings hollow: small traders often bore the steepest losses, while hedge funds, market makers, and even corporations found ways to adapt and profit from inflated valuations.

Taken together, the cases of cryptocurrencies, meme stocks, and digital bubbles show that speculative dynamics are no longer confined to the market's fringes but are entrenched at its core. Retail coordination, meme-driven identities, and charismatic opinion leadership reveal how attention and narratives can generate prices that persist independently of fundamentals. The sharper hypothesis is therefore not simply that digital speculation accelerates old dynamics, but that it destabilizes the very categories used to study them. If assets can be untethered from fundamentals by design, calling them "bubbles" risks circularity, since the very anchors of measurement no longer apply.

Conclusion to Chapter I: From classical finance to digital disruption

Chapter I has shown that financial markets cannot be reduced to rational expectations or efficient information processing. Volatility, bubbles, and anomalies persist because investors are systematically biased, their decisions are shaped by overconfidence, loss aversion, and the stories that coordinate collective belief. In the digital era these tendencies are no longer confined to individual psychology but are amplified within online communities, where attention becomes a form of currency, herding turns dispersed impulses into synchronized cascades, and influencers mobilize trust in ways that rival or surpass traditional analysts. The result is a financial environment where visibility and narrative can generate value independently of fundamentals.

This dynamic is visible in practice. Cryptocurrencies reveal that assets with little or no intrinsic anchor can sustain massive valuations through narratives of scarcity or humor alone, unsettling the very definition of a bubble. Meme stocks show how retail coordination and meme-driven identities can destabilize institutions while simultaneously being absorbed back into corporate strategy. Tesla demonstrates how meme-like speculation is no longer confined to fragile firms but reaches into established leaders, blurring the line between exuberant narratives and legitimate growth stories. These cases confirm that speculation is not an episodic deviation but a structural feature of contemporary markets.

The broader diagnosis is clear. Rationalist frameworks fail, not because they are elegant abstractions, but because they cannot capture the psychological, social, and cultural processes through which markets actually function. Financial theory must therefore move beyond classical models and treat markets as constructs in which attention, narratives, and collective behavior operate alongside fundamentals as enduring drivers of price.

The next step is to ask whether new technological infrastructures can succeed where human-centered theory has failed. Part II, '*Algorithmic Finance - From Market Shaping to Predictive Diagnostics and the Paradox of Control'* turns to the role of computation. If speculative dynamics are produced by human psychology and amplified by digital communities, can machine learning, deep learning, and artificial intelligence detect them earlier, predict them more reliably, or even replace human investors as the primary interpreters of financial markets?

II. FROM CLASSICAL FINANCE TO DIGITAL DISRUPTION: RETHINKING SPECULATIVE BEHAVIOR

Part I demonstrated that markets are not purely rational systems but are shaped by biases, narratives, and the dynamics of digital communities. Attention cycles, herding, and influencer authority showed how speculation emerges from psychology and culture as much as from fundamentals. Part II shifts the lens from human behavior to technological infrastructures. The focus is no longer on why investors attend to certain assets, but on how algorithms determine what becomes visible in the first place.

D. ALGORITHMIC INFRASTRUCTURES AS MARKET SHAPERS

D.1. Platform Visibility: How Algorithms Steer Collective Attention

Visibility algorithms work at the collective level: they concentrate attention by steering crowds toward a narrow set of "*hot*" assets, regardless of fundamentals. In financial markets, visibility is never neutral. The question here is not why investors fixate on particular stocks or tokens, but how ranking systems and recommendation engines decide what appears before them at all.

Recommendation engines and ranking systems, the invisible machinery behind trending lists, push notifications, or search results, are often described as filtering tools that *"present items in which the user is likely to be interested in a specific context"* (Zibriczky, 2016). On the surface this sounds neutral. In practice, filtering is neither passive nor unbiased. By deciding what rises to the top of a feed, algorithms grant some assets visibility and consign others to obscurity. In entertainment this may mean wasted time on Netflix; in finance it can redirect billions, distort valuations, and trigger systemic cascades.

Research identifies three structural biases that directly shape financial visibility:

- **Popularity bias**: already-visible assets (such as a surging meme stock or viral token) are recommended more often, creating a *"rich-get-richer"* cycle (Chaney et al., 2017).
- **Position bias**: investors disproportionately click on items placed near the top, regardless of quality
- **Exposure bias**: assets absent from feeds are effectively invisible, even when their fundamentals are strong.

Together these biases create a funnel that privileges spectacle over substance. Sensational and emotionally charged stories generate early engagement, which algorithms interpret as relevance, further amplifying their reach.

This raises a deeper question: are distortions driven primarily by algorithmic design or by human psychology? Some scholars argue that algorithms homogenize behavior by steering diverse investors toward the same narrow set of assets (Chaney et al., 2017). Others see them as mirrors of existing biases, since humans naturally gravitate to novelty and excitement (Tsintzou, 2018). Most evidence points to a recursive cycle: human biases spark early engagement, while algorithms entrench them mechanically. In finance the effects are asymmetric. Retail investors rarely browse thousands of securities; they trade what enters their feeds. Platforms therefore act as bottlenecks of salience, determining not only what investors focus on but also what they ever get the chance to see.

This distinction clarifies the shift from Part I to Part II. Previously, attention was treated as a psychological variable defined by what people search, share, and discuss. Here, visibility is treated as a structural variable defined by what the algorithm surfaces in the first place. The Efficient Market Hypothesis assumes all available information is processed by investors. Yet algorithmic infrastructures redefine *"availability"* itself. Information hidden by ranking systems may exist, but it does not enter prices in practice.

Despite the stakes, financial research lags behind. Most recommender-system studies come from consumer domains such as YouTube or Amazon, where bias produces skewed playlists or product suggestions. Simulations show reinforcement effects but ignore the complexity of capital markets. Audits confirm distortion but rarely separate user choice from algorithm design. Survey-based reviews map the problem broadly but seldom connect it to volatility or systemic risk. The result is under-theorization: we know algorithms bias visibility, but little about how this bias interacts with market fragility.

This matters because visibility is itself a scarce resource. By amplifying hype-driven content, platforms concentrate investor focus on a handful of *"hot"* assets. Informational diversity collapses, dissenting signals are muted, and markets become more prone to speculative surges. At scale, this undermines the Efficient Market Hypothesis, which presumes that markets aggregate dispersed knowledge efficiently. In practice, visibility is not neutral but engineered and traded through clicks, shares, and rankings.

The paradox is that the same systems that destabilize also generate diagnostic signals. Engagement spikes, trending lists, and attention cascades often precede volatility. Algorithms that amplify speculation simultaneously make fragility visible. This dual role of magnifying bubbles while offering clues about their onset sets the stage for the larger question of Part II: whether computational tools can not only amplify but also anticipate speculative dynamics.

The central point is that amplification in finance is structurally engineered rather than incidental. Recommendation engines do not simply reflect curiosity; they manufacture salience by embedding popularity, position, and exposure biases into the architecture of visibility. In doing so, they transform fragmented attention into synchronized speculation before any trade occurs. This helps explain why some memes or *"hot stocks"* rise to prominence while others, equally relevant, remain unseen.

What Part I described as investor-driven attention and herding is therefore extended here: algorithms are not neutral filters but active organizers of visibility, concentrating focus and accelerating speculative cascades. The unresolved question is whether these distortions originate primarily in algorithmic design or whether they simply entrench the psychological tendencies already documented by behavioral finance. This thesis argues that the two cannot be separated. Human biases may provide the spark, but algorithms institutionalize those signals, hard-wiring them into the structure of visibility itself. The result is a market environment where efficiency falters not only because of cognitive limits but also because the informational field is engineered in systematically skewed ways.

Amplification demonstrates how algorithms shape speculation at the collective level, while personalization represents the next step, fracturing that collective field into individualized streams. Instead of one shared set of signals, investors receive tailored feeds that reinforce prior beliefs, producing informational silos and parallel financial realities. While this section has shown how amplification synchronizes attention, the next examines how personalization fragments it. This shift exposes not only the fragility of digitally mediated markets but also the difficulty of treating algorithms as predictive tools rather than engines of bias.

D.2. Personalization Algorithms: Fragmenting Attention into Informational Silos

Personalization algorithms work at the individual level: they fragment attention by tailoring feeds to prior beliefs, dividing investors into insulated informational silos. In Part I, echo chambers were explained primarily as social phenomena that arose from imitation, herding, and narrative contagion within shared environments. Here the mechanism is different. Echo chambers are not organic outcomes of human interaction but engineered directly by personalization. Investors do not simply cluster with like-minded peers; they are guided into algorithmically curated worlds where exposure to dissenting signals is systematically reduced.

This distinction matters. Classical accounts of speculation assumed a broadly shared informational environment, even if investors disagreed on interpretation. In dot-com stocks or meme rallies, optimism and pessimism competed within the same marketplace of signals. Personalization undermines this premise. Instead of one collective mania, attention fractures into micro-clusters: bullish tribes immersed in optimism,

bearish tribes steeped in gloom. These parallel informational worlds sustain localized "*micro-manias*" that would not exist without algorithmic curation.
Why personalization leads to segmentation is contested. Some emphasize algorithmic design. Chaney et al. (2017) describe "*algorithmic confounding*," where recommender systems reinforce the very preferences they helped create. This shows that these loops intensify over time, making polarizing content more entrenched. From this perspective, segmentation is a structural outcome of algorithms designed to maximize engagement. Others stress psychology. Tsintzou (2018) argues that personalization mirrors confirmation bias and selective exposure. Investors, like news readers, prefer content that affirms what they already believe. Platforms intensify this through homophily, clustering like-minded users around the same influencers or hashtags. In reality, both dynamics interact: personalization reflects bias but also amplifies it mechanically.

Practically, personalization repurposes the same structural biases discussed in the visibility section (popularity, position, and exposure) but applies them unevenly across individuals. In the visibility case, these biases homogenize investor focus, pushing everyone toward the same narrow set of "*hot*" assets. In personalization, the effect is the opposite: the very same mechanisms are distributed differently across users, fragmenting attention rather than synchronizing it. Two investors scrolling the same platform may inhabit radically different informational universes: one immersed in bullish optimism, the other fed a steady stream of bearish gloom. Both environments skew toward emotionally charged content, but the outcome is not a collective cascade. Instead, it is the division of the market into parallel, insulated silos.

The consequences are significant. First, confirmation bias deepens bullish investors are shown more bullish stories, bearish investors more bearish ones. Corrective signals are filtered out. Second, trading has become polarized. Conflicts between retail enthusiasm and institutional short positions, as in meme stocks, reveal the risks when opposing bubbles collide. Third, information diversity shrinks. When markets fragment into silos, they struggle to converge on fundamentals. Cookson, Engelberg, and Mullins (2021) link such siloing to lower returns, highlighting material costs.

Research methods shed light on these dynamics but also reveal important limits. Simulation models such as Chaney et al. (2017) show how feedback loops make user behavior increasingly uniform, but they simplify away the complexity of real markets. Platform-level studies, for example Flaxman, Goel, and Rao (2013), demonstrate large-scale ideological segregation online, yet they cannot fully disentangle algorithmic effects from user self-selection. Finance-specific work, such as Cookson, Engelberg, and Mullins (2021), directly links personalization to trading outcomes and lower returns, but these studies face problems of representativeness and the opacity of proprietary algorithms. As a result, much of what we know about financial personalization remains indirect, and its full scope is still underexplored.

In synthesis, personalization illustrates how speculation no longer unfolds in one shared marketplace of ideas but within fragmented informational silos. By tailoring feeds to prior beliefs, it deepens confirmation bias, sustains micro-manias, and creates self-reinforcing echo chambers that investors cannot easily exit. These algorithmically engineered divisions transform herding from a collective cascade into a patchwork of insulated feedback loops. The risk is not only that these silos polarize markets but that, when they collide, they produce sudden volatility and systemic fragility.

Placed alongside amplification, a clearer pattern emerges. Where visibility synchronized attention by pushing investors toward the same “hot” assets, personalization divides them into separate informational worlds. Together these dynamics reveal the dual face of digital speculation: mass cascades on one side, tribalized micro-markets on the other. Fundamentals still anchor valuation in many parts of the market, but in these speculative arenas their role is increasingly mediated by attention cycles and fragmented silos. Rather than replacing fundamentals outright, amplification and personalization overlay them with new dynamics that can distort, delay, or temporarily override price formation. This thesis argues that it is precisely this interaction between traditional anchors of value and digitally mediated speculation that exposes the limits of existing financial theories.

Yet this account has limits. We already know from Part I that speculation depends on stories, signals, and shared meanings that investors generate and circulate. What remains less understood is how, in digital settings, the production of this content is increasingly automated and scaled. Personalization explains how investors are sorted into fragmented silos, but it does not explain what fills those silos or why some messages endure while others fade. To address this, the next section shifts focus from the distribution of information to its production, examining how generative and automated systems industrialize financial communication and further shape the environments in which speculation takes place.

D.3. Generative and Automated Narratives: The Next Frontier in Market Narratives

We saw how algorithms shape financial attention by filtering and distributing content. Visibility algorithms synchronize attention around the same sensationalist assets, while personalization fragments attention into micro-segments, producing bullish and bearish tribes. Both cases assumed a human starting point: stories written by analysts, posts shared on forums, memes circulating through networks. Generative AI shifts the terrain. The issue is no longer which stories rise to the top or who gets to see which version. It is which stories exist in the first place. Large Language Models (LLMs), bots, and neural networks are no longer just intermediaries of financial information. They are authors of it. This shift from curation to generation industrializes financial storytelling, flooding markets with synthetic narratives at a scale and speed that human discourse cannot match.

LLMs such as GPT-style architectures can synthesize market news, analyst commentary, and social media chatter into coherent narratives in seconds (Lo & Ross, 2024). Newer systems go beyond synthesis: Kim, Muhn, and Nikolaev (2024) demonstrate how models can extract “*narrative signals*” from numerical data, identifying storylines latent in earnings calls, stock reports, or even price dynamics. These capacities collapse the cost of producing plausible-sounding financial content, enabling a flood of synthetic commentary with little human involvement. In effect, machines now generate the very stories that investors encounter, not just determine which ones they see.

This industrialization has two major implications. First, it decouples narrative production from human limits. Unlike journalists or analysts, generative systems do not verify, reflect, or tire. They can endlessly output bullish or bearish commentary, sustaining hype cycles long after organic interest would have faded. Second, it blurs the distinction between credible analysis and manufactured noise. Human readers struggle to separate fact-based insight from AI-generated hype, particularly when the latter imitates the tropes and rhetorical styles of professional finance (Lo & Ross, 2024). Generative systems, in this sense, are not merely amplifiers of sentiment; they are producers of synthetic consensus, capable of fabricating the appearance of widespread belief.

Scholars interpret this shift in different ways. One camp emphasizes risk. Lo and Ross (2024) argue that poorly aligned models can hallucinate, misrepresent, or produce coherent but false narratives, which in financial contexts may distort prices or trigger herd reactions on flawed premises. Another camp sees continuity. Cookson, Engelberg, and Mullins (2021) document how human investors already gravitate toward confirmatory and emotionally charged stories. From this perspective, AI does not introduce a new distortion so much as accelerate long-standing tendencies, acting as an industrial multiplier of confirmation bias and narrative contagion.

This thesis takes a stronger view: automated narratives are not merely an acceleration but a structural novelty. Earlier, algorithms shaped which stories gained attention (visibility) or who saw which versions of them (personalization). Generative systems now determine which stories exist in the first place. This is not just narrative contagion in Shiller's (2017) sense, where stories spread virally among humans. It is the industrialization of story supply itself, with machines generating new waves of hype designed to sustain sentiment and speculation.

Concrete manifestations of this are already visible. Bots can produce streams of bullish memes, fake "*breaking news*" or emotionally charged forum posts, keeping hype alive with little human participation. Unlike traditional amplification, where engagement gradually fades, generative content ensures that echo chambers never run empty. Bulls encounter endless bullish "*analysis*"; bears are fed a steady diet of collapse warnings. This entrenches polarization and helps sustain micro-bubbles, even in the absence of fresh fundamental news.

Research on automated narratives remains concentrated in computer science rather than finance. Existing work measures tone or simulates bot cascades but rarely connects outputs to capital allocation. Lo and Ross (2024) emphasize hallucination risk, while Kim, Muhn, and Nikolaev (2024) stress potential in extracting narrative signals from data. Both perspectives leave open a central gap: we lack tools that track not only the production of synthetic stories but their deployment as market strategies. In entertainment or consumer platforms, measuring tone suffices. In finance, where information itself is a price signal, the stakes are far higher.

The implications are serious. First, generative systems can manufacture synthetic consensus, flooding markets with coordinated bullish or bearish narratives and locking investors into polarized silos. Second, they erode trust in financial information ecosystems. If investors cannot reliably distinguish credible insight from machine-written hype, then the informational signals markets rely on to coordinate expectations become unstable.

In conclusion, generative and automated narratives represent an important new stage in speculative dynamics. While still emerging, these tools already demonstrate the capacity to generate synthetic stories at scale, lowering the cost of financial commentary and sustaining cycles of hype. The risks are clear: even modest volumes of machine-written analysis can fabricate the appearance of consensus and blur the boundary between credible insight and manufactured noise.

Placed alongside amplification and personalization, this development completes a structural shift. Platforms do not only determine which assets gain visibility or how investors are divided into informational silos; they now influence what kinds of narratives circulate in the first place. Taken together, these dynamics show that markets are increasingly coordinated not just by human psychology and collective storytelling, but by computational infrastructures that amplify, fragment, and generate the signals through which investors interpret prices.

This thesis argues that such a transformation requires rethinking the foundations of financial theory. If attention and narratives can be partly engineered by machines, then markets can no longer be understood only through efficiency theory or behavioral finance. The challenge is double: algorithms destabilize markets by amplifying and producing speculative signals, yet they also create new kinds of data that might be harnessed for prediction. This opens the next section on predictive analytics and algorithmic foresight, which asks whether tools such as machine learning and AI can help identify and even anticipate bubbles, and introduces the question whether such systems could one day rival or replace human investors in managing speculative risk.

E. PREDICTIVE ANALYTICS AND ALGORITHMIC FORESIGHT

E.1 Natural Language Processing for Market Sentiment: Quantifying Narrative Intensity and Hype Cycles

This section shifts focus to methodology: the computational tools researchers use to measure market psychology. Natural Language Processing (NLP), broadly defined as *“the use of computational techniques to analyze and represent human language data”* (Kaminski & Gloor, 2014), has become one of the most common ways of converting online chatter into numerical *“sentiment scores.”* The appeal is obvious: if bubbles are fueled by contagious narratives, then quantifying narrative intensity across platforms could, in theory, provide early warning signals of speculative overheating. Yet the empirical record tells a different story. Most sentiment indices either lag behind price movements or confuse hype with conviction, raising serious doubts about their credibility as forecasting devices.

The central weakness is reductionism. Standard NLP frameworks collapse noisy, ironic, or meme-driven talk into a crude *“positive versus negative”* scale. This simplification strips away the cultural cues that make financial discourse intelligible in the first place. Kaminski and Gloor (2014) find that Twitter-based sentiment measures usually trail market returns, implying that what looks like investor *“mood”* is often just a delayed echo of price action. Ciganovic and D’Amario (2024), examining Reddit sentiment during cryptocurrency surges, likewise find no consistent predictive power even when volatility is extreme. In such cases, *“sentiment”* is less a window onto crowd psychology than a mirror of price dynamics dressed up as foresight.

A deeper problem is that sentiment is not a universal measure of mood but a platform-specific artifact. Each digital arena embeds its own biases:

- **Twitter**: sentiment tracks intraday volatility but rarely predicts it. Retweet cascades inflate apparent consensus without signaling genuine commitment.
- **Reddit**: upvotes and long discussions can produce weak predictive effects, yet what rises to the top reflects engagement bias rather than information quality.
- **TikTok**: mood is encoded in sound, visuals, and emoji, making text-based sentiment measures misleading. Almeida and Gonçalves (2023) stress that viral content often signals entertainment, not trading intent.
- **Discord**: real-time chat reveals coordination more clearly, but its semi-private, fragmented servers make generalization impossible.

These examples reinforce a broader critique: sentiment indices are not neutral instruments. They are artifacts shaped by platform norms, community cultures, and algorithmic filters. Roos and Reccius (2021) argue that topic models often capture curiosity or novelty seeking rather than conviction, raising the question of whether such tools measure the forces that actually drive capital flows. Put differently, what appears as *“market mood”* is often nothing more than the structure of visibility imposed by platforms themselves.

This methodological fragility carries consequences. By compressing messy, culturally embedded discourse into tidy numerical scores, researchers risk overstating the predictive value of sentiment indices. Almeida and Gonçalves (2023) warn that sentiment-based trading can amplify bubbles, as both humans and algorithms herd around misleading cues. Cookson, Engelberg, and Mullins (2021) demonstrate that echo chamber dynamics already distort what content is visible; sentiment measures built on such biased samples reflect reinforcement loops rather than market-wide psychology. Even the strongest positive findings, such as Liu and Tsyvinski’s (2018) evidence that Google search intensity forecasts short-term crypto returns, highlight how fleeting these effects are. They flag bursts of attention, not enduring trends.

In conclusion, NLP-based sentiment analysis is too fragile to function as a forecasting device. By collapsing complex, culturally embedded discourse into crude polarity scores, these methods misinterpret hype for conviction and often lag behind the very price movements they claim to predict. At best, sentiment indices serve as barometers of speculative intensity, flagging when narratives themselves begin to drive pricing. They do not, however, provide a stable foundation for anticipating when bubbles will form or collapse.

This limitation underscores the broader claim of this thesis: speculative dynamics cannot be understood through discourse analysis alone. If language stripped of context cannot consistently anticipate crises, then predictive tools must shift from what investors say to how markets behave. This is the rationale for moving beyond sentiment indices toward structural diagnostics that identify fragility directly in patterns of prices, volatility, and liquidity. The next section therefore turns to machine learning and early warning models, which promise to move past linear thresholds and toward nonlinear diagnostics capable of detecting speculative risk before it erupts.

E.2. Machine Learning in Early Warning Systems: From thresholds to nonlinear diagnostics of market fragility

Early-warning systems (EWS) represent a methodological departure from language-based sentiment analysis. Instead of reducing discourse to numerical scores, they focus directly on the structure of financial markets. Bussiere and Fratzscher (2006, p. 107) define EWS as "*statistical or rule-based models that aim to predict the probability of a crisis before it materializes*" Their purpose is diagnostic: to detect anomalies in prices, volatility, or liquidity early enough to provide regulators and investors with warning signals. The challenge, however, is that the history of EWS reveals a recurring tension. Models that are transparent tend to oversimplify, while models that capture complexity often do so at the cost of interpretability. This trade-off between clarity and predictive power sets the stage for the integration of machine learning into early-warning design.

The earliest EWS relied on fixed thresholds: crisis warnings were triggered when macro-financial variables such as reserves, exchange rates, or credit growth breached predefined levels. Kaminsky and Reinhart (1999) describe this "*signal approach*" as useful for broad classification, yet prone to false alarms and missed crises. Reinhart and Rogoff (2011) likewise stress that while such indicators helped catalogue centuries of financial folly, they offered little in terms of real-time guidance. Regression-based logit and probit models (Frankel & Rose, 1996; Eichengreen & Rose, 1998; Bussiere & Fratzscher, 2006) refined the approach by assigning statistical probabilities to crisis events. Yet they too assumed static, linear relationships in markets that are inherently nonlinear. Wang et al. (2019) conclude that these econometric models "systematically underestimated the timing and severity of crashes."

To address such limits, volatility and regime-switching models introduced a dynamic perspective. As mentioned earlier, the Johansen–Ledoit–Sornette (JLS) model conceptualized bubbles as "*super-exponential accelerations*" in prices driven by herding, producing an "*alarm index"* that anticipated some crashes and rebounds (Yan & Sornette, 2011). Its theoretical grounding in behavioral dynamics was innovative, but its sensitivity to parameter calibration generated frequent false positives. Similarly, SWARCH models (Hamilton & Susmel, 1994) classified markets into high- and low-volatility states, improving adaptability by allowing thresholds to evolve endogenously (Wang et al., 2019). These methods represented progress over static thresholds, but they remained econometric in spirit: dependent on restrictive distributional assumptions and vulnerable to misspecification.

Machine learning (ML) was introduced as a more radical alternative. As Nag and Mitra (1999) argued in an early contribution, neural networks can capture "*relationships not easily specified in econometric form*" Supervised algorithms such as Support Vector Machines and Random Forests have since been trained to classify bubble versus non-bubble regimes more flexibly than probit regressions (Wang et al., 2019). Deep learning architectures, particularly Long Short-Term Memory (LSTM) networks, incorporate sequential dependencies that are critical for volatility and crash detection. Fischer and Krauss (2018) show that LSTMs "significantly outperform linear benchmarks" in predicting stock returns, while Wang et al. (2019) demonstrate that combining SWARCH with an LSTM achieved 96.6% accuracy in forecasting daily crashes with an average lead time of 2.4 days. These results suggest that ML can extend, rather than merely repackage, the predictive frontier.

Yet fragility persists. Samitas et al. (2020) emphasize that training data for crises are inherently scarce, making rare-event prediction vulnerable to overfitting. The black-box nature of deep models undermines interpretability, limiting their practical utility for regulators who require transparent and causal signals. Explainable AI methods such as Shapley values (Lundberg & Lee, 2017) provide post-hoc attributions of variable influence, but as Purnell et al. (2024) caution, such tools "*risk mistaking correlation for causality*." Reflexivity deepens the problem: once widely adopted, predictive models risk producing "*prediction theatre*" where apparent accuracy reflects ex post volatility rather than genuine foresight.

The integration of ML into early warning systems therefore illustrates both progress and constraint. Classical models such as JLS and SWARCH remain valuable for their theoretical grounding and interpretability, while ML contributes adaptability and predictive sharpness. Yet neither resolves the underlying paradox: models that aspire to forecast fragility are themselves fragile. The reflexive nature of markets ensures that the very act of prediction reshapes the conditions it seeks to anticipate. Forecasting through price and volatility models may improve technical precision, but it cannot escape the problem of data scarcity and opacity. The limits of forecasting through price and volatility models have encouraged a parallel turn. Rather than modeling markets directly, big data approaches examine the digital traces of investor behavior as potential barometers of systemic stress.

In sum, machine learning expands the toolkit of early-warning systems by moving beyond fixed thresholds and linear econometric models. Neural networks and deep learning architectures capture nonlinear dependencies that older models overlooked, and they often outperform classical benchmarks in short-term crash detection. Yet their strengths expose their limits. Scarce crisis data make rare-event prediction vulnerable to overfitting, while black-box opacity undermines their credibility for regulators who need transparent signals. Reflexivity compounds these challenges: once adopted, predictive models risk reshaping the very dynamics they seek to anticipate. Machine learning therefore enhances diagnostic precision but does not escape the fragility that defines the phenomena it measures.

Placed alongside E.1, the parallel becomes clear. NLP-based sentiment indices falter because they reduce culturally embedded discourse into oversimplified polarity scales. ML-based volatility models stumble because they mistake technical sharpness for foresight in environments defined by reflexivity and data scarcity. Both illustrate a broader thesis claim: computational tools extend the frontier of prediction, but they cannot resolve the structural instability of speculation or fully displace the interpretive role of human judgment.

These limits set the stage for the next frontier. If discourse analysis struggles with cultural nuance and price models collapse under reflexivity, perhaps the most promising signals lie in the aggregate traces of digital activity itself. The following section turns to Big Data as Systemic Risk Barometers, asking whether large-scale behavioral signals can provide regulators with early indicators of stress that neither sentiment scores nor volatility models can reliably deliver.

E.3 Big Data as Systemic Risk Barometers: Search, flows, and engagement as institutional diagnostics

Big data approaches pursue foresight along a different path. Rather than modeling financial variables directly, they analyze the digital residues of investor behavior. Lazer et al. (2009) define big data as the *"digital residues of human behavior that can be collected and analyzed at scale."* In finance, these residues include search queries, clicks, transaction flows, and social media engagement, which are increasingly treated as high-frequency barometers of systemic stress. The promise is that such traces can reveal speculative overheating before fundamentals adjust. Yet this institutionalization of behavioral signals raises its own set of methodological and reflexive challenges.

The canonical example is the Google Search Volume Index (GSVI). Da, Engelberg, and Gao (2011) transformed casual online curiosity into a formal finance variable, showing that spikes in search activity preceded short-term price increases and reversals. Their contribution was methodological: legitimizing digital search as a measurable proxy for attention. Yet limits were visible from the outset. Their dataset was confined to Russell 3000 stocks, and they cautioned that GSVI may reflect curiosity rather than trading intent. Subsequent studies exposed further weaknesses. Lai, Chang, Hu, and Chou (2022) found that positive shocks in search intensity pushed prices upward in retail-driven markets, but negative shocks failed to produce symmetrical declines due to constraints on short-selling. A meta-review of 56 studies by Ayala, Gonzálvez-Gallego, and Arteaga-Sánchez (2024) concludes that GSVI consistently predicts volatility and trading activity but has *"unstable and contradictory"* relationships with returns. At the macro level, Puhr and Müllner (2021) caution that Google Trends can complement conventional indicators of globalization or credit cycles but cannot replace them.

Social media data broadens the terrain but intensifies the problem of instability. Mao, Counts, and Bollen (2011) reported that Twitter *"bullishness"* preceded both Google queries and stock returns, suggesting discourse could serve as a leading indicator of attention. Yet their effects decayed within days. Platform-specific work shows even sharper divergences. Baklanova (2025) finds that a Reddit sentiment index correlates with Bitcoin volatility but functions recursively as both signal and effect. Chang and Peng (2022) demonstrate that TikTok enthusiasm reliably forecasts subsequent declines, implying virality is more a symptom of exhaustion than of conviction. What appears predictive on one platform often dissolves into noise or even reverses direction on another.

Narrative-driven approaches seek to embed these digital traces in systemic theory. Nyman, Kapadia, and Tuckett (2021) employ algorithmic text analysis of financial news to track emotional language shifts, arguing that rising *“approach”* terms foreshadowed pre-crisis exuberance in 2007–08. By grounding signals in psychological theory, they move beyond raw sentiment indices. Yet as Ayala (2024) notes, hindsight bias remains powerful: detecting exuberance retrospectively is easier than providing actionable foresight.

The institutionalization of big data has reinforced these tensions. Regulators and risk managers are drawn to its promise of immediacy. Wang, Zong, and Ma (2019) demonstrate how integrating regime-switching econometrics with ML can sharpen predictive precision, while Purnell, Etemadi, and Kamp (2024) present an explainable ML framework that flagged instability before the Silicon Valley Bank collapse. These contributions illustrate the appeal of embedding digital signals into systemic monitoring. Yet critics remain cautious. Loughran and McDonald (2011) show that generic sentiment dictionaries systematically misclassify financial terminology, undermining text-based indices. Kaminski and Gloor (2014) warn against *“mistaking correlation for causation in markets that approximate random walks.”* Even Purnell et al. (2024) concede that their framework identifies associations rather than causal mechanisms.

Taken together, big data turns the digital residues of investor behavior into potential barometers of systemic stress. Search volumes, flows, and engagement metrics promise immediacy and granularity, yet their instability and reflexivity undermine their reliability. The Google Search Volume Index, for instance, transformed curiosity into a measurable proxy for attention (Da, Engelberg & Gao, 2011), but meta-reviews show “*unstable and contradictory*” effects on returns (Ayala, Gonzálvez-Gallego & Arteaga-Sánchez, 2024). Twitter bullishness once appeared predictive (Mao, Counts & Bollen, 2011), only to decay within days. Reddit indices correlate with Bitcoin volatility but act as both cause and effect (Baklanova, 2025). TikTok enthusiasm even forecasts decline (Chang & Peng, 2022). What looks predictive on one platform often reverses on another. Reflexivity compounds the problem: as Soros (1987) argued, markets are not neutral mirrors but self-fulfilling systems, and once digital signals are integrated into trading or risk management they risk amplifying the very dynamics they were meant to anticipate. Barber and Odean’s (2008) *price pressure hypothesis* already warned that attention-driven buying inflates valuations; Nyman et al. (2021) confirm that narrative exuberance fed by data-driven monitoring becomes part of the instability. Big data, in short, provides probabilistic diagnosis rather than deterministic foresight: it illuminates fragility but never escapes it.

Placed alongside the previous sections, a broader picture emerges. NLP sentiment analysis reduced discourse into misleading polarity scores, stripping away cultural context. Machine learning introduced adaptability and nonlinear diagnostics, yet it struggled with interpretability, the scarcity of crisis data, and the reflexive nature of markets. Big data institutionalized behavioral traces, but in doing so it risked turning prediction itself into a new source of systemic instability. Taken together, these approaches highlight both the promise and the limits of computational foresight: each extends predictive reach, yet each generates vulnerabilities of its own.

The thesis position is that computational intelligence, while powerful, cannot by itself replace human investors or stabilize markets through predictive precision. Its tools have shown real capacity to extend foresight, but because markets are reflexive, predictive instruments never remain external observers: they feed back into the system they measure, blurring the line between diagnosis and participation. This does not render prediction futile, but it does mean its role is not corrective in the strict sense. Instead of resolving fragility, computational foresight reconfigures it, shifting risks into new forms that themselves require oversight and interpretation.

This progression leads directly to the final part of the thesis. Since algorithms are themselves woven into the reflexive dynamics of markets, the question is not simply whether they can replace human investors, but how they might contribute to market diagnosis without reinforcing the very fragilities they aim to detect. Their capacity to extend foresight is evident, yet it is inseparable from the risks they introduce once absorbed into practice. The challenge, then, is to identify the ethical, institutional, and regulatory conditions under which such tools can add diagnostic value while minimizing destabilization.

F. THE PARADOX OF AI IN FINANCE - ENGINE AND ANTIDOTE

F.1 AI as Market Diagnostic Tool: Between Foresight and Reflexivity

The very infrastructures that have destabilized contemporary markets are now being enlisted to stabilize them. This section looks at that paradox in two ways. First, it asks whether artificial intelligence, used as a diagnostic tool, really provides foresight or whether it simply repackages existing volatility in technical language. Second, it examines the ethical and regulatory issues that follow from the growing use of models that are powerful, opaque, and widely shared.

The optimistic case for AI starts with performance. Several studies report striking results from machine-learning early-warning systems. Wang, Zong and Ma (2019), for example, present a hybrid LSTM-SWARCH model that they claim can forecast crash regimes in Chinese equity markets with test accuracies above 96% and an average lead of two trading days. Ke, Kelly and Xiu (2019) likewise show that text-based models can extract incremental information from financial language, improving return forecasts beyond what dictionary methods achieve. Taken at face value, such findings suggest a qualitative shift: by processing high-dimensional data streams, from intraday prices to corporate narratives, AI seems able to detect weak signals that humans or traditional econometrics would miss.

Yet doubts about such claims have a longer history. Kaminski and Gloor (2014) showed that emotional bursts on social media often follow volatility rather than precede it, while Engelberg and Parsons (2009) demonstrated that media coverage tends to react to market movements rather than anticipate them. These earlier insights raise an uncomfortable question for more recent work: if the inputs to predictive models are themselves endogenous to prices, then reported accuracies may reflect retrospective fit rather than genuine foresight. Sophisticated pattern recognition can appear predictive while merely tracking the surface of volatility. The term *prediction theatre* captures this problem: numbers, probabilities, and lead times can look scientific without truly anticipating turning points.

Even when this issue is handled carefully, opacity remains a second problem. Deep learning models such as LSTMs can generate confident warnings while offering no clear explanation of how they were produced. Industry responses rely on *explainability*. Shapley-value methods, for instance, claim to identify which inputs influenced a prediction (Lundberg and Lee, 2017; Purnell, Etemadi and Kamp, 2024). Yet Buckman, Joseph and Robertson (2021) argue that much of this is post hoc storytelling: plausible-sounding rationalizations that satisfy readers but cannot be tested as causal accounts. For risk managers and supervisors this is not a theoretical problem but a practical one. If the pathways from inputs to outputs remain inaccessible, responsibility cannot be meaningfully assigned when warnings prove wrong.

A third difficulty is reflexivity. Markets adapt to the tools that monitor them. Harras and Sornette (2007) show how convergence on similar trading rules can create instability. Regulators make the same point in policy terms: IOSCO (2025) warns that widespread reliance on the same AI systems could increase systemic risk by synchronizing attention and action. Predictive dashboards may begin as warning devices but once they are widely used, they risk becoming coordination devices, aligning behavior across institutions and accelerating the very herding they were meant to prevent.

These critiques must be set against evidence that automation can help in calmer conditions. Hendershott and Riordan (2013) find that algorithmic trading improves price efficiency, while Fischer and Krauss (2018) show that LSTM models outperform linear benchmarks on certain prediction tasks. The challenge is reconciling these benefits with the problems of endogeneity, opacity, and reflexivity. A useful way to frame the issue is by context. In routine periods, fast automated responses can tighten spreads and speed up the incorporation of straightforward information. In periods of stress, the same speed and shared reliance on common models can amplify correlations, drain liquidity, and spread error. The contrast between *AI and humans* is therefore misleading. A better description is that AI institutionalizes human heuristics - herding, confirmation seeking, sentiment chasing - at machine speed and scale. If the data are saturated with these biases, the models will not rise above them; they will reproduce and accelerate them.

The implications are both methodological and practical. To be credible, claims of predictive power must be tested against hard timestamps and fixed designs that rule out look-ahead bias. They must be robust across different training windows and feature sets, and they must rely on drivers that are not themselves direct reflections of price action. Most importantly, they should be judged not only on statistical fit but on their behavioral consequences. A model that improves classification metrics but encourages synchronous de-risking across institutions may worsen systemic fragility. Few studies examine these second-order effects. By contrast, the most promising uses of AI seem to lie not in point forecasts of returns but in regime diagnostics: identifying shifts in volatility structures, thinning liquidity, concentrated positioning, or the dominance of certain narratives. These are conditions regulators can act upon with circuit breakers, margin adjustments, or supervisory guidance, without creating the false precision and coordination risks that come with deterministic price forecasts.

The position of this thesis is therefore cautious but not dismissive. AI can play a role in stabilizing markets, but only if its function is carefully recast. Predictive dashboards should not be treated as oracles of price direction but as monitors of state: sensors for liquidity stress, fragility in market microstructure, or the clustering of attention. Their use must be subject to auditable evidence and strict evaluation. Without such a reorientation, predictive systems risk becoming amplifiers disguised as antidotes, producing confident warnings that guide participants onto the same bridge just as it begins to sway.

This conditional promise makes the ethical and regulatory questions unavoidable. Transparency, responsibility, and investor protection are not secondary to technical performance but central to whether these systems can operate safely at scale.

F.2 Ethical and Regulatory Dimensions of AI in Finance

If the technical case for AI is mixed, the governance challenge is even sharper. The key issue is how to design institutions so that powerful but opaque tools can operate with accountability and investor protection. This is not a purely technical matter. Without safeguards, predictive dashboards risk amplifying the same speculative bubbles they are meant to contain. Bubbles thrive on opacity, weak accountability, and the fast spread of persuasive narratives. For that reason, governance must be understood as a form of bubble prevention. Current debates center on three axes: transparency, responsibility, and investor protection, with a risk of concentration running through all three.

Transparency is the first axis. Regulators have drawn a clear line. IOSCO (2021, 2025) states that black-box systems (models that generate outputs without showing how decisions are made) are unacceptable in areas where compliance or prudential judgments are at stake. This is directly relevant to bubbles. In periods of speculative excess, investors are already inclined to defer to seemingly authoritative signals. If a model produces confident predictions that cannot be interrogated, it risks reinforcing herd behavior at exactly the wrong moment. Industry responses often rely on post hoc interpretability. Purnell, Etemadi and Kamp (2024), for example, use ensembles with Shapley-value analysis to claim they can reveal which factors drive predictions. Buckman, Joseph and Robertson (2021) are more skeptical. They argue that these explanations are cosmetic, offering stories that sound plausible but cannot be tested. That kind of surface-level clarity can feed bubble dynamics by turning complex outputs into easy narratives. Regulators also recognize that transparency can backfire. IOSCO (2025) warns that too much disclosure might allow manipulation by revealing how models work. The paradox is clear: transparency is essential for legitimacy, but if poorly designed, it can undermine trust and amplify speculative surges. Workable regimes will need to be layered. These could include glass-box requirements for safety-critical uses such as margining, independent audits of data lineage when trade secrecy blocks disclosure, regular reporting of sensitivity and stability, and enforceable incident logs that keep responsibility alive across staff turnover. None of these measures solves opacity completely, but they can prevent models from becoming unchallengeable *oracles* that fuel bubbles.

Responsibility is the second axis. Here, the gap between rhetoric and practice is stark. Platforms often claim neutrality, presenting themselves as passive conduits of information. Evidence suggests the opposite. Studies show that engagement-driven systems boost low-credibility financial content once it attracts enough interaction. Cookson, Engelberg and Mullins (2023) find similar distortions in retail trading linked to engagement amplification. These are not accidents; they are design outcomes. And they matter for bubbles. If the very systems that deliver analysis and signals amplify hype, they are also inflating speculative surges. Regulators have responded by insisting that firms retain *ultimate responsibility* for AI-driven decisions (IOSCO, 2025). But this principle assumes a straight chain of accountability. In practice, errors travel through layers: a platform's ranking algorithm, a broker's execution logic, a third-party vendor's model. Without clear allocation of duties beforehand and reliable logs afterward, responsibility disappears. In this vacuum, bubbles thrive, because no one actor is held to account. Effective regimes will need both prospective allocation clauses that set duties across the stack and retrospective logs that allow failures to be traced across proprietary boundaries.

Investor protection is the third axis, and its stakes are most visible during bubbles. On the positive side, AI can improve fraud detection, enhance disclosures, and alert retail traders when conditions are unusual. On the negative side, the same technologies can generate synthetic narratives that look like genuine analysis, automate persuasion at scale, and speed up sentiment cycles that expose small investors to volatility. These risks peak in bubble phases. The Financial Stability Board (2024) captures this tension, praising AI's monitoring power while warning about bias and concentration. IOSCO (2025) stresses that dependence on common models could itself amplify speculative vulnerabilities. Scholars add that institutions often roll out powerful tools faster than they put safeguards in place, widening gaps between institutional players and retail investors (Buckman, Joseph and Robertson, 2021). In every bubble, this asymmetry shows: institutions exit first, leaving retail traders most exposed. Remedies have been proposed, from human-in-the-loop oversight to continuous audits and standardized disclosures. Yet these often slide into governance theatre, procedures that create an appearance of accountability while leaving core problems untouched. More promising are design-level interventions: trading interfaces that flag the difference between attention-driven and fundamental signals, provenance mechanisms that watermark machine-generated analysis, and dashboards that prevent mono-signal dependence by showing alternative perspectives. These are modest but they act directly where bubbles form.

Running across all three axes is the problem of concentration risk. When many institutions rely on the same vendors, model families, or pre-trained foundations, errors stop being isolated and become systemic. IOSCO (2025) calls this a single-point-of-failure problem. The comparison with ratings agencies and Value-at-Risk models is instructive. Before the 2008 crisis, widespread reliance on similar risk metrics created a false sense of security. When the bubble burst, uniformity made the collapse faster and deeper. The same danger applies to AI. If too many firms optimize against the same signals, correlations will rise in calm markets, and fragility will surface suddenly under stress. Structural solutions are needed. These include diversity mandates to ensure heterogeneity in models and data, stress tests that simulate correlated error and publish herding indices, and switch-off protocols for entire signal classes when monoculture risk appears. Such measures go further than disclosure, but they are more likely to prevent predictive dashboards from becoming bubble accelerators.

The analysis of transparency, responsibility, and investor protection shows that governance is not an afterthought but the central hinge of AI's financial role. Without layered transparency regimes, accountability mechanisms that travel across institutional stacks, and protections that address asymmetries between institutional and retail actors, predictive dashboards risk reinforcing the very bubbles they aim to contain. The lesson is that ethical and regulatory design cannot simply follow technical innovation; it must anticipate how opacity, amplification, and concentration risk will interact with speculative dynamics.

Taken together, Sections F.1 and F.2 frame AI not as a neutral forecasting tool but as a market institution. Its predictive power is inseparable from its systemic effects. As a diagnostic tool, AI can illuminate fragility by identifying volatility regimes, liquidity stress, or narrative concentration. Yet once embedded in practice, those same tools shape behavior, redistribute responsibility, and create new channels of systemic risk. Foresight, opacity, and reflexivity therefore converge with transparency, responsibility, and investor protection: technical and ethical questions are two sides of the same coin.

The thesis position is cautious but constructive. AI should not be dismissed as inherently destabilizing, nor celebrated as a new oracle of market foresight. Instead, it should be redefined as a conditional diagnostic infrastructure: one that can support stability only when coupled with robust governance. Predictive dashboards must be judged not by their statistical accuracy alone but by their behavioral and systemic consequences. Used narrowly as sensors of state, under regimes of transparency and accountability, they can help regulators act without fueling new cascades. Used as unregulated predictors of price, they risk becoming amplifiers disguised as safeguards. The final claim is therefore that AI's contribution to finance lies not in replacing human judgment but in augmenting it - provided institutions embedding ethical design and regulatory oversight at the very core of technological adoption.

Conclusion to Part II

Part II has shown that digital infrastructures do not simply mediate speculation, they constitute it. Algorithms of visibility, personalization, and generation actively engineer the informational environment of markets. They determine which assets rise to prominence, how investors are clustered into silos, and which narratives gain traction at scale. These mechanisms transform speculation from a byproduct of human psychology into a structural outcome of computational design. Amplification synchronizes attention around "hot" assets, personalization fragments it into parallel informational worlds, and generative systems industrialize narrative production. Speculation is therefore no longer solely human-driven but increasingly algorithmically organized.

Predictive analytics build on this infrastructure by presenting themselves as tools of foresight. Natural language processing quantifies sentiment, machine learning models classify crash regimes, and big data traces are turned into barometers of systemic stress. Each technique expands diagnostic capacity, but each carries structural weaknesses: NLP reduces cultural nuance to polarity scores, machine learning often trades interpretability for complexity, and big data risks amplifying reflexivity by transforming diagnostic signals into trading cues. Instead of delivering stable foresight, these methods fluctuate between technical precision and what amounts to prediction theatre.

These tensions converge in the paradox of technology. The same computational systems that destabilize markets are increasingly deployed as solutions to that instability. AI can function as a market diagnostic tool if understood as a sensor of systemic states rather than an oracle of price direction. Yet without transparency, clear accountability, and robust protection for retail investors, predictive dashboards risk serving as amplifiers disguised as safeguards. Part II therefore concludes that the promise of algorithmic finance is conditional: its effectiveness depends less on technical accuracy than on the ethical, institutional, and regulatory frameworks in which it is embedded.

Conclusion to Parts I and II

Taken together, Parts I and II trace the transformation of speculation across two dimensions, from the psychological to the infrastructural, and from the human to the algorithmic.

Part I demonstrated that speculation cannot be explained by rationalist frameworks of efficient markets. Investor biases, contagious narratives, and digitally networked communities produce volatility and bubbles as structural features of contemporary finance. Attention acts as currency, herding turns dispersed impulses into cascades, and influencers who increasingly rival analysts as coordinators of belief. Markets must therefore be understood as cultural and social systems as much as informational ones.

Part II extended this analysis into the algorithmic domain. Here, infrastructures of visibility, personalization, and generation do not merely reflect human biases but institutionalize them in code. Platforms synchronize attention, fragment it into informational silos, or flood it with synthetic narratives, while predictive systems attempt to turn these same dynamics into early warning signals. Yet the reflexivity of markets ensures that these tools are never external observers: once adopted, they feed back into the very processes they aim to monitor. Technology thus emerges simultaneously as amplifier and diagnostic, destabilizer and safeguard.

The broader synthesis is clear. Speculative dynamics today cannot be reduced to either human psychology or computational precision alone. They are generated at the intersection of cognitive biases, cultural narratives, digital communities, and algorithmic infrastructures. Stability and instability alike emerge from this entanglement. Rationalist theory falters because it cannot account for these recursive, technologically mediated processes; algorithmic prediction falters when it forgets that markets adapt to their own diagnostic tools. To understand contemporary speculation, finance must be reconceptualized as a reflexive system where human and machine, psychology and infrastructure, amplification and diagnosis are inseparably intertwined.

This leads to the central question of this thesis: **How are financial bubbles reshaped in the digital era, and to what extent can algorithms and artificial intelligence help in detecting them?**

From this overarching problem, the hypotheses can be grouped into three thematic domains that guide the analysis.

1. Investor Profiles and Psychology

The first domain concerns the individual dimension of speculation. H1 suggests that investment status does not influence perceptions of algorithmic exposure and personalized news feed reinforcement. H2 further proposes that investors may reject rationalist assumptions by relying on biases and narratives instead of market efficiency. Together, these hypotheses highlight the psychological and behavioral underpinnings that shape how investors engage with markets in the digital age.

2. Collective Dynamics in Digital Communities

The second domain addresses the social amplification of speculation. H3 emphasizes that digital communities transform individual cognitive biases into collective speculative behavior, while H4 argues that cryptocurrencies and meme stocks exemplify how online narratives and hype outweigh traditional fundamentals. These hypotheses underscore the role of digital communities in reshaping collective behavior and driving new forms of market volatility.

3. Technological Infrastructures and Ambivalence

The third domain focuses on the technological layer of financial speculation. H5 proposes that algorithms curate exposure and shape beliefs about markets, while H6 highlights that investors and potential investors trust AI-based financial tools as credible sources of foresight. Together, these hypotheses point to the ambivalent role of technology, both as a driver of speculative dynamics and as a potential means of detecting and regulating them.

In conclusion, these three domains crystallize the central dynamics of digital-era bubbles. To address the hypotheses, I now turn to the next part of this thesis, where methodological choices and empirical strategies will be developed.

PART II: FIELD STUDY

CHAPTER 1: METHODOLOGICAL DESIGN

1.1 METHODOLOGICAL CHOICES

To build on the literature review, this thesis continues with a quantitative survey that grounds the theoretical insights in real-world data and captures how investors and potential investors perceive bubbles, digital platforms, and the role of AI. This step directly connects back to the central research question: *How are financial bubbles reshaped in the digital era, and to what extent can algorithms and artificial intelligence help in detecting them?*

For this reason, I chose a quantitative rather than a qualitative approach. My aim was not only to explore how bubbles form in the digital era, but also to measure how widespread certain behaviors and perceptions are across different groups. A qualitative approach, such as interviews or case studies, would have provided rich detail and personal insights, but it would not have allowed me to compare groups, test specific hypotheses, or identify broader patterns with the same level of reliability. Since bubbles are collective phenomena, what matters is not only individual stories but also how often and how strongly these behaviors appear across a population.

Choosing a quantitative design made it possible to translate abstract ideas into something measurable. Instead of leaving concepts such as *"investors rely on narratives"* or *"algorithms influence visibility"* at the level of theory, I was able to turn them into concrete survey questions. This approach allowed me to collect responses that could be analyzed statistically. It made it possible to spot recurring patterns, test relationships between variables, and check whether the mechanisms described in the literature also appear in the behavior and perceptions of investors and potential investors.

Just as importantly, this choice helps address a gap in existing research. Much of the current work on bubbles, algorithms, and digital speculation is descriptive, focusing on particular cases, or conceptual, aiming to build theoretical frameworks. What is often missing is quantitative evidence that links these ideas to observable data. By using a survey-based method, this thesis contributes that missing layer, grounding theoretical claims in measurable results and showing how they play out in practice.

1.2 SAMPLING

The target population was broadly defined as individual investors and people interested in investing, particularly those active on digital platforms and exposed to speculative environments. Since bubbles and hype cycles are collective phenomena, the aim was to reach a diverse sample that could reflect different profiles, levels of experience, and risk appetites. The survey was distributed primarily through social media channels, specifically Instagram, LinkedIn, Reddit, and Discord. This choice was consistent with the subject of the research, since these platforms play a central role in contemporary speculation. Instagram and Reddit are widely used by retail investors for sharing memes, following hype cycles, and discussing cryptocurrencies. Discord hosts numerous investing and trading communities where coordination and real-time discussion take place. LinkedIn, by contrast, reaches a more professional audience, including finance practitioners and institutional investors.

As a result, the sample was naturally tilted toward investors or people interested in investing, reflecting the communities present on those platforms. However, a small number of respondents indicated that they were not interested in investing at all. This was not an issue, since those few responses were removed during data cleaning and are therefore not referenced in the analysis. Even if they had been included, they could still have provided a useful point of contrast by highlighting how perceptions of speculation differ between active investors and outsiders.

This approach does carry the risk of overrepresenting younger and digitally active respondents. However, these are also the groups most actively involved in meme stocks, cryptocurrencies, and algorithm-driven trading, making the sample particularly relevant for the goals of this study. The objective was to obtain at least 100 valid responses for statistical analysis. In practice, 282 responses were collected, exceeding expectations and offering a richer dataset. Within this, 177 participants completed the entire survey, while others left some questions unanswered. This distinction was taken into account during analysis.

1.3 DATA COLLECTION INSTRUMENTS

Data was collected through an anonymous online survey created with Google Forms. This platform was chosen because it is accessible, easy to use, and allows responses to be exported directly for analysis. To encourage honesty and reduce bias, participation was entirely voluntary, and no personal information was collected.

The questionnaire was divided into six thematic sections that reflected the focus of the research, followed by a final demographic section:

1. **Investment Profile and Intentions**
2. **Psychological Biases**
3. **Influence of Digital Platforms**
4. **Algorithms and Artificial Intelligence**
5. **Speculative Behaviors and Scenarios**
6. **General Information (age, gender, education)**

The survey contained 21 questions in total. This is slightly higher than the average for short exploration surveys, but it was necessary because of the complexity of the topic. Speculation in the digital age is influenced by psychology, social interactions, and technology all at once. Capturing these dynamics properly requires more than just a few items. Each section of the literature review was translated into targeted questions, making it possible to test the six hypotheses of the study in a structured way.

Most questions used a Likert-scale format. This design was chosen because Likert scales are easy for respondents to understand and complete, while also producing data that allows attitudes and perceptions to be measured consistently across individuals. This format makes it possible to compare responses, identify patterns, and conduct statistical tests with reliability. At the same time, Likert scales have certain limitations.

They can encourage neutral or middle-ground answers, they sometimes fail to capture the intensity of strong opinions, and they may oversimplify complex attitudes by reducing them to fixed points on a scale. These limits were kept in mind during the design and interpretation of the survey.

Although no formal pilot study was conducted, the items were adapted from well-established theoretical constructions discussed in the literature review. Their clarity and simplicity reduced the risk of misinterpretation, which increased the validity of the instrument.

The survey length struck a balance between detail and feasibility. On average, participants could complete it in about five to seven minutes, which was short enough to prevent fatigue but still comprehensive enough to cover the main research dimensions. Before launching it more widely, I first ran a small beta test by asking people I know to complete the survey and share feedback on whether it felt too long, repetitive, or unclear. Their input confirmed that the length was appropriate and engaging

1.4 METHOD OF ANALYSIS

Once I collected the responses, I first exported them from Google Forms to Google Sheets for storage and basic handling. I then imported the dataset into Python (via Google Colab) to carry out the cleaning and statistical analysis. At this stage, I distinguished between two datasets. The first was made up of the 177 respondents who completed the entire survey. I used this version for multivariable analysis, since complete answers were necessary to compare variables reliably. The second was the full set of 282 responses, which I used for single-variable or descriptive analysis. As a result of the cleaning process, some people who had indicated no interest in investing were also excluded, but this was simply because they had not completed all questions.

For the analysis, I combined both descriptive and inferential techniques. On one level, I used descriptive statistics to summarize overall trends and provide a clear picture of how respondents perceive bubbles, digital platforms, and AI. On another level, I applied different statistical tests and models to compare groups, explore associations, and identify the factors that shape attitudes and behaviors. I also looked for broader patterns in the data by grouping respondents into profiles that reflected distinct ways of approaching investment.

Overall, this mix of methods offered both breadth and depth, allowing me to move from surface-level observations to more detailed insights. By structuring the analysis in this way, I ensured that the results were not only robust but also closely tied back to the theoretical perspectives developed in the literature review.

With this methodological foundation established, we can now turn to *Chapter 2: Analysis and Interpretation of Results*, where the empirical findings are presented and examined in detail.

CHAPTER 2: ANALYSIS AND INTERPRETATION OF RESULTS

2.1. WHO ARE THE RESPONDENTS? DEMOGRAPHICS AND INVESTOR PROFILES

As a first step, I examined the demographics and investor profiles of respondents, since understanding who participated in the survey provides essential context for interpreting the results that follow.

For the descriptive analysis, I worked with the full set of 282 responses and consistently compared them with the 177 fully completed surveys. Both samples displayed the same tendencies, which gave me confidence in the reliability of the findings.

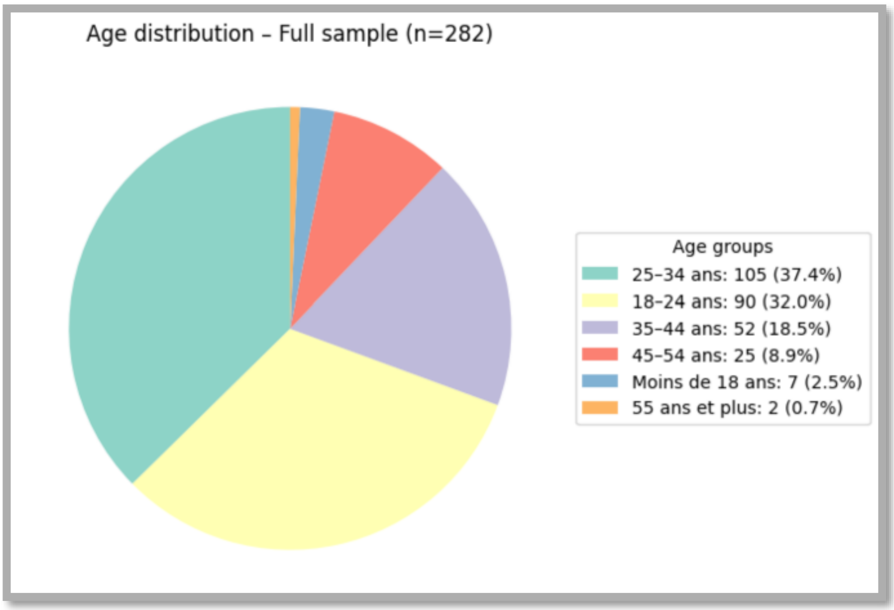

To begin, I looked at the "Age distribution (n=282) Pie Chart" chart, which shows a clear skew toward younger groups, with nearly 70 percent of respondents between 18 and 34 years old. Older investors are underrepresented, which is not surprising given that the survey was distributed via Instagram, Reddit, Discord, and LinkedIn. While these bias limit representativeness of the general investor population, it is also highly relevant. Younger, digitally active groups are precisely those most engaged with cryptocurrencies, meme stocks, and algorithm-driven speculation.

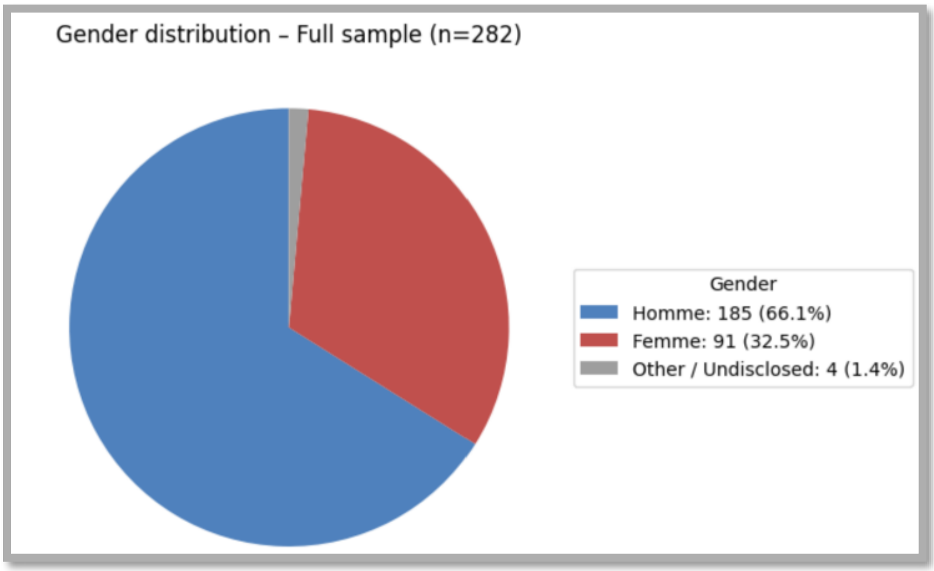

Next, I examined the "Gender distribution Pie Chart (n=282)" chart, which reveals another imbalance: two-thirds of respondents are men, one-third are women, and only a very small fraction identified as “other” or preferred not to disclose. We can assume that this skew reflects broader trends in retail trading and crypto communities, where male participation has historically been dominant.

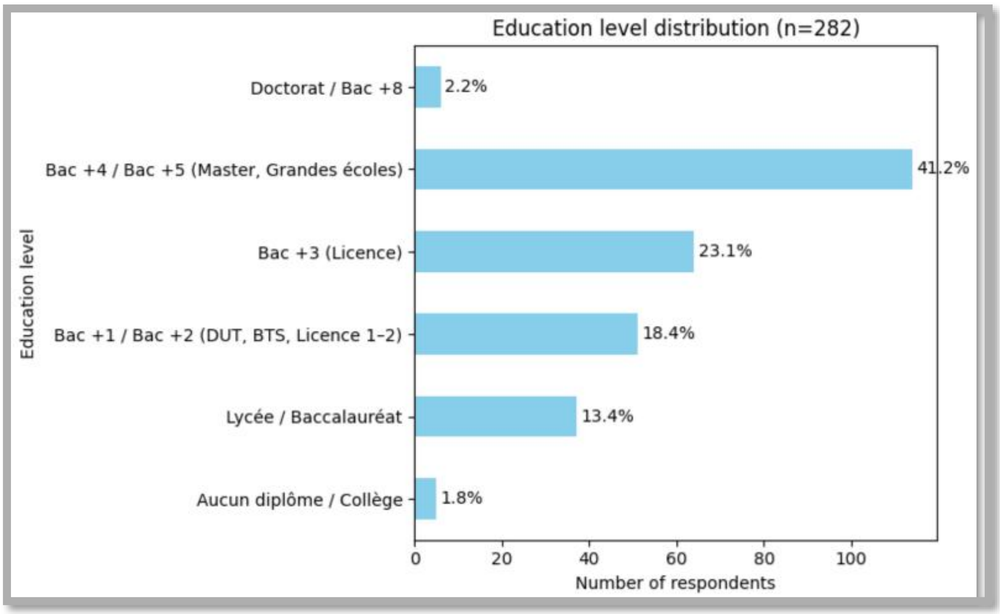

Then, I turned to the “Education level distribution Bar Chart (n=282)” chart, which shows that the sample leans heavily toward the well-educated. Over 40 percent hold a master’s degree (or equivalent), and nearly a quarter hold a Bachelor’s. Only a handful reported having no diploma or only secondary schooling. We can make the assumption that this profile challenges the stereotype of uninformed retail traders rushing blindly into speculation and instead shows that highly educated individuals are also drawn to these environments.

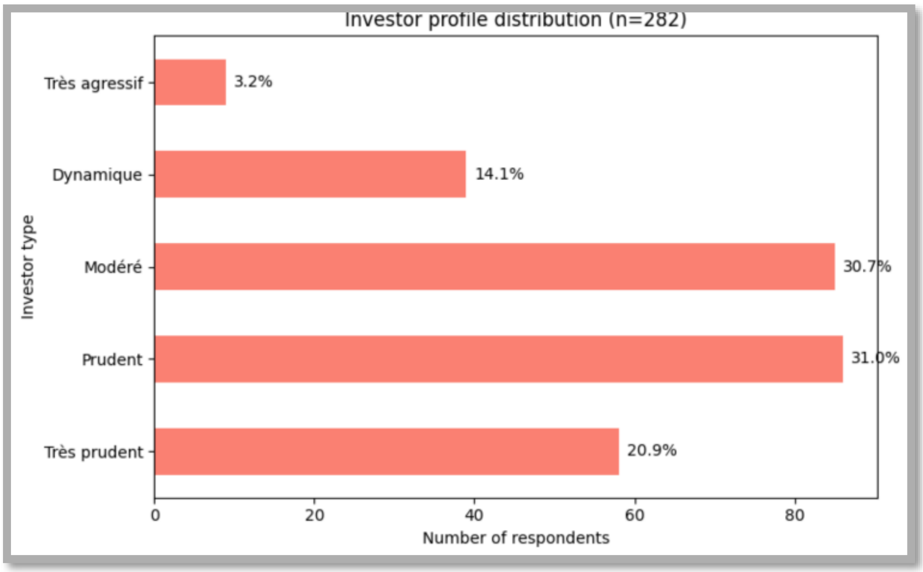

After that, I examined the “Investor profile distribution Bar Chart (n=282)” chart, which sheds light on self-perceived risk appetite. Most respondents described themselves as cautious: 31 percent called themselves “prudent,” another 31 percent “moderate,” and 21 percent “very prudent.” By contrast, only 14 percent identified as “dynamic” and just 3 percent as “very aggressive.” It is important to stress that these are self-assessments. We can make the assumption that people often

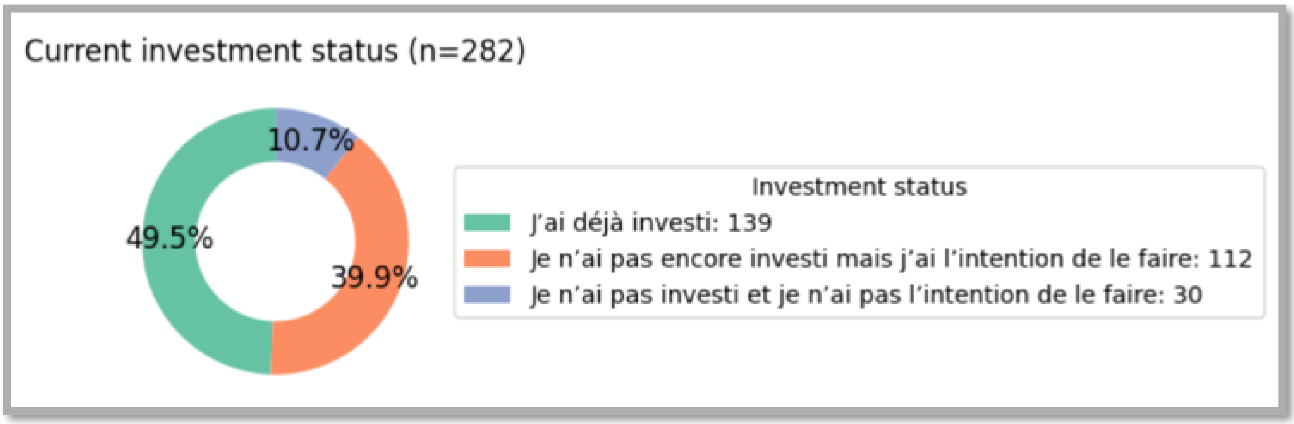

see themselves as careful even when their actions suggest otherwise, especially in volatile and hype-driven environments. This could suggest that speculative waves are not fueled only by self-declared risk takers but also by larger groups who perceive themselves as cautious yet still end up participating in collective enthusiasm.

Finally, I looked at the “Current investment status Donut Chart, (n=282)” chart, which shows that about half of respondents had already invested, 40 percent intended to invest, and only 11 percent said they had no plans to do so. We can assume that this mix could highlight how speculation are fueled not only by active market participants but also by those preparing to enter, reminding us that bubbles thrive as much on expectations as in actual trades.

Altogether, the descriptive analysis shows that the sample is young, male-dominated, highly educated, and self-identified as cautious. This profile does not represent the full diversity of the general investor population, and the reliance on social media platforms for distribution likely contributed to these imbalances. Rather than treating these characteristics as causal findings, it is important to acknowledge them as limitations that shape how the results should be interpreted. At the same time, the concentration of digitally active and investment-curious respondents aligns with the focus of this thesis, since these groups are precisely those most exposed to speculative hype, online narratives, and algorithmic visibility.

In light of the literature review, these characteristics are significant. Earlier studies have stressed that bubbles in the digital era cannot be explained only by irrational gamblers or uninformed traders. Instead, speculation often involves educated individuals who consider themselves cautious yet operate in environments where narratives and social contagion amplify risk-taking. The fact that a substantial share of respondents intends to invest but have not yet done so further illustrates how expectations and anticipation, rather than only active participation, can fuel speculative dynamics.

Taken together, these descriptive findings highlight both the opportunities and the limits of the dataset. While the sample cannot claim to represent the general investor population, it captures groups that are central to understanding how digital speculation unfolds: young, digitally active individuals, many of whom are either already investors or preparing to enter financial markets. Their self-perceived prudence, combined with high levels of education, might echo the literature's insight that bubbles are not simply the product of reckless or uninformed actors but of collective dynamics in which even cautious participants can be drawn into speculative cycles. At the same time, it could also suggest that fundamentals and rational considerations still play a role, since many respondents see themselves as careful and deliberate investors. In a theoretical context, this aligns with the idea that investors tend to respect fundamentals and approach markets with prudence, even if their behavior may at times diverge under the influence of narratives and hype. This makes investor status and psychology a logical starting point for hypothesis testing

2.2 INVESTOR PROFILES AND PSYCHOLOGY

The literature review emphasized how algorithms and personalized news feeds shape financial decision-making by curating visibility, amplifying narratives, and reinforcing prior beliefs. This builds on prior research showing that bubbles are not confined to irrational or inexperienced participants but rather unfold as collective phenomena that affect diverse groups of investors simultaneously. Narratives and algorithmically curated content reach across categories, meaning that both experienced and novice investors can be drawn into similar dynamics. To test this empirically, I formulated *H1: Investment status doesn't influence perceptions of algorithmic exposure and personalized news feed reinforcement*.

To test H1, I examined whether investment status influences perceptions of algorithmic exposure and personalized news feed reinforcement. I first calculated the mean Likert scores for each investment status group Q1* (active investors, potential investors, and non-investors) then for (Algorithms influence the financial content I am exposed to) Q.12** and then for (Personalized news feeds reinforce my investment opinions) Q13. *** These averages were then compared visually and statistically.

**Q1. Current investment status ->I have already invested / I have not yet invested but intend to do so / I have not invested and do not intend to do so*

***Q12. Algorithms: Algorithms (TikTok, YouTube, Twitter, etc.) strongly influence the financial content I am exposed to.** -> Strongly disagree / Disagree / Neutral / Agree / Strongly agree

****Q13. News feeds: Personalized news feeds reinforce (or would reinforce) my investment opinions (or future investment opinions).** -> Strongly disagree / Disagree / Neutral / Agree / Strongly agree

Investment status (Q1)	Mean – Algorithms influence (Q12)	Mean – Personalized feeds (Q13)
I have not yet invested but intend to	**3.75**	**3.39**
I have not invested and do not intend to	3.57	3
I have already invested	3.57	3.05

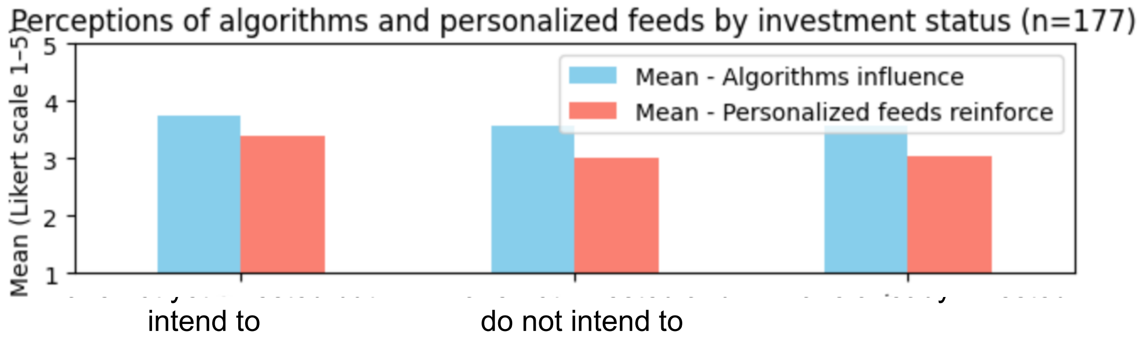

The results are presented in the bar chart “Perceptions of algorithms and personalized feeds by investment status (n=177).” The averages suggest broadly similar perceptions across groups. Potential investors reported slightly higher scores (3.75 for algorithmic influence and 3.39 for personalized feeds) compared with active investors (3.57 and 3.05) and non-investors (3.57 and 3.00). While the gap is small, it is consistent with the idea that individuals preparing to enter the market may be more attuned to the role of algorithms and feeds in shaping their financial outlook.

To formally test this observation, I conducted one-way ANOVA tests. The null hypothesis (H0) stated that investment status has no effect on perceptions of algorithmic influence or personalized feed reinforcement, while the alternative hypothesis (H1) stated that investment status does affect these perceptions. For Q12 (algorithms), the ANOVA yielded $F(2,174) = 0.48$, $p = 0.62$, and for Q13 (personalized feeds), $F(2,174) = 1.89$, $p = 0.15$. Both p-values are above the conventional 0.05 threshold, which means I cannot reject the null hypothesis. Although potential investors scored slightly higher in the descriptive results, the differences are too small to be considered statistically significant.

In sum, these findings indicate that perceptions of algorithmic and feed influence are widespread and relatively independent of whether one is already investing, intends to invest, or has no plans to do so. This echoes the literature’s claim that algorithmic effects are a structural feature of digital speculation, shaping exposure across the board rather than targeting only specific categories of investors.

Since investment status did not show significant differences in how respondents perceive algorithmic influence, I next turned to another dimension of investor psychology: self-declared risk orientation. The literature review highlighted that bubbles are not only driven by participation status but also by psychological factors such as risk perception, narrative susceptibility, and behavioral biases. Exploring risk profiles therefore provides a deeper way to test whether individuals' self-conceptions as cautious or aggressive investors shape how they recognize algorithmic and feed influences.

This was tested using Q3 (Which type of investor best describes you), Q12 (Algorithms influence the financial content I am exposed to), and Q13 (Personalized news feeds reinforce my investment opinions). I first calculated the mean Likert scores for each risk group and then presented the results in the bar chart *"Perceptions of algorithms and personalized feeds by risk profile (n=177)."*

Risk profile (Q3)	Mean – Algorithms influence (Q12)	Mean – Personalized feeds (Q13)
Very cautious	**4.14**	**3.41**
Cautious	3.82	3.29
Moderate	3.45	3.11
Dynamic	3.47	3.07
Very aggressive	**2.62**	**2.62**

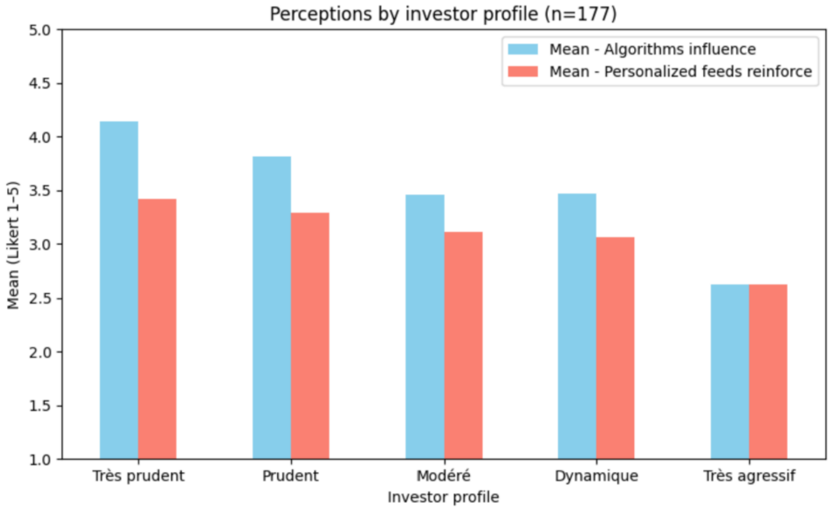

The descriptive results show clearer differences than those observed for investment status. Very cautious investors reported the strongest agreement that algorithms influence them (mean = 4.14) and that personalized feeds reinforce their opinions (mean = 3.41). Cautious investors followed with slightly lower averages (3.82 and 3.29). From there, the values declined gradually: moderate (3.45 and 3.11), dynamic (3.47 and 3.07), and very aggressive investors, who stood out with the lowest scores on both items (2.62). This pattern suggests that perceptions of algorithmic influence vary systematically with self-identified risk orientation. Cautious investors see themselves as more affected by algorithms and feeds, while aggressive investors consider themselves largely unaffected. At the same time, these are self-assessments, and it is possible that aggressive investors underreport such influence even if their behavior is shaped by it in practice.

To test whether these differences were statistically significant, I conducted one-way ANOVA tests. For algorithmic influence (Q12), the ANOVA yielded $F(4,172) = 3.72$, $p = 0.006$, which is below the 0.05 threshold. This indicates that risk profile systematically shapes how respondents perceive the role of algorithms. Looking back at the descriptive results, very cautious investors reported the highest perceived influence, while very aggressive investors reported the lowest, and the ANOVA confirms that these differences are not due to chance.

For personalized feeds (Q13), the ANOVA result was $F(4,172) = 1.01$, $p = 0.40$, above the significance threshold. Although the means vary across risk groups, the differences are not statistically reliable. Interestingly, respondents generally acknowledge that algorithms shape the content they see, but they do not strongly perceive that this reinforcement translates into actual investment decisions.

In sum, the results offer limited support for H1. Investment status itself does not significantly shape perceptions of algorithmic exposure or personalized feed reinforcement, which supports the literature's claim that bubbles and algorithmic influence are collective phenomena cutting across investor categories. However, for algorithmic influence (Q12), the ANOVA yielded *$F(4,172) = 3.72$, $p = 0.006$, below the 0.05 threshold*, indicating that risk profile plays some role in how respondents perceive algorithmic impact. Very cautious investors reported the highest perceived influence, while very aggressive investors reported the lowest. This was the only significant effect observed, so it should be interpreted cautiously, but it suggests that while algorithmic effects are broadly universal, perceptions of their influence may still vary with psychological self-identification as cautious or aggressive.

Having considered how investment status and risk profiles relate to perceptions of influence, I now move to the next part of the analysis, which looks more directly at investor behavior.

The literature review highlighted how behavioral finance challenges the efficient market hypothesis by showing that investors are often guided by cognitive biases and narratives rather than purely rational calculation. Based on this, I formulated the hypothesis:

H2 investors may reject rationalist assumptions by relying on biases and narratives instead of market efficiency.

To test this, I examined three survey questions designed to capture these behavioral dimensions. Q4 measured overconfidence ("I believe I can (or could) beat the market through my own decisions"), Q5 measured loss aversion ("A loss of €100 would hurt me more than a gain of €100 would make me happy"), and Q7 assessed whether respondents value narratives over fundamentals ("Narratives and stories matter more than fundamentals in explaining asset prices").

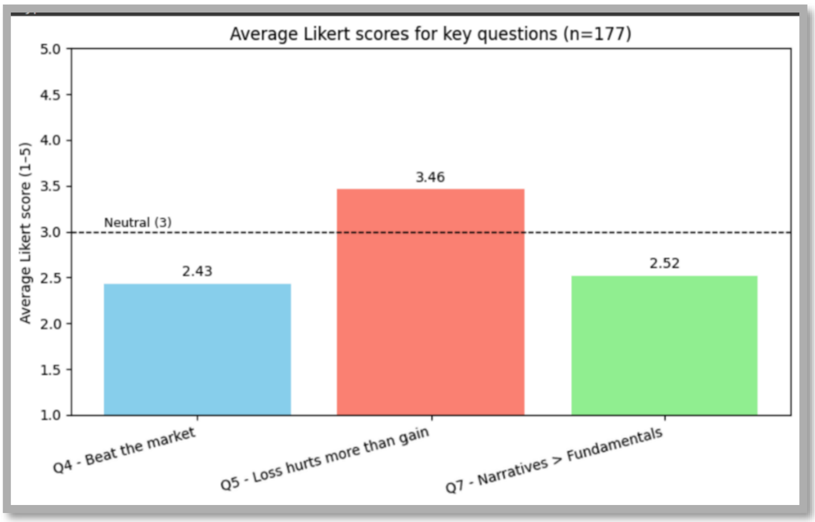

Fig. Bar chart Average Likert score

The bar chart *"Average Likert scores for key questions (n=177)"* highlights three important aspects of investor psychology. The most striking finding is on Q5, where the mean score of 3.46 is clearly above the neutral midpoint, confirming the presence of loss aversion. Respondents acknowledge that losses weigh more heavily than equivalent gains, which is consistent with one of the most robust biases identified in behavioral finance.By contrast, the average scores for Q4 (2.43) and Q7 (2.52) both fall below the neutral point. This suggests that respondents generally do not see themselves as able to consistently beat the market, nor do they claim to prioritize narratives over fundamentals when explaining asset prices.

Analytically, this points to a nuanced pattern: the main bias at play is emotional sensitivity to losses, not exaggerated self-belief or reliance on narratives. This is especially notable given that the literature review highlighted the growing role of narratives in shaping bubbles (as emphasized by authors such as Shiller). The fact that respondents did not strongly endorse narratives over fundamentals suggests either a reluctance to admit narrative influence, or that narratives operate more subtly than investors consciously recognize.

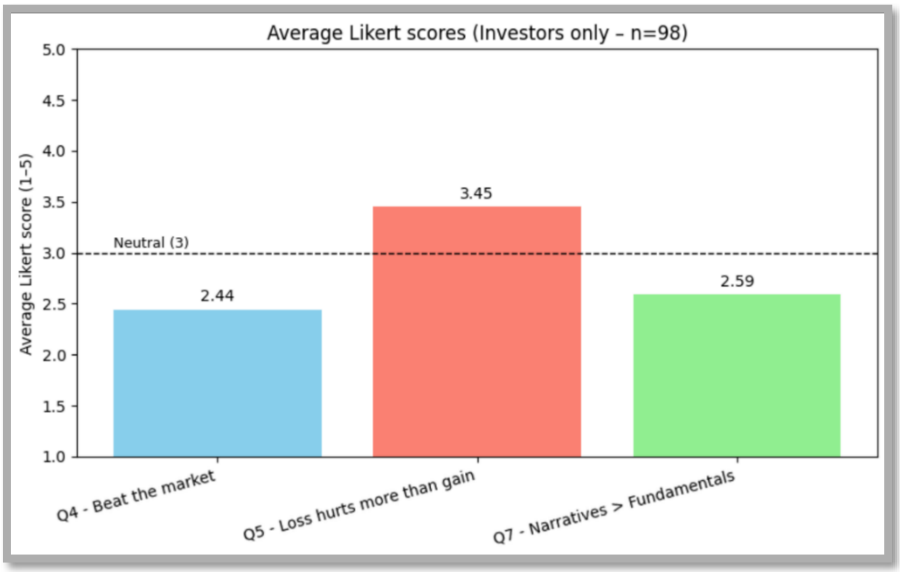

Fig. Bar chart Average Likert score

The bar chart “Average Likert scores (Investors only – n=98)” shows that the patterns observed in the full sample are also present among actual investors. Once again, the mean score for loss aversion (Q5 = 3.45) lies clearly above the neutral midpoint, confirming that investors feel losses more intensely than equivalent gains. In contrast, the averages for overconfidence (Q4 = 2.44) and narrative weighting (Q7 = 2.59) remain below 3, indicating that most investors do not believe they can outperform the market and continue to give more weight to fundamentals than to stories.

This consistency across the subgroup of investors and the broader sample strengthens confidence in the results. It suggests that loss aversion is the dominant behavioral bias across both groups, while overconfidence and narrative reliance appear less pronounced.

In this sense, the findings align with H1 and with the literature review, since they show that investors do not fully adhere to rationalist assumptions and are shaped by psychological biases. At the same time, there is a divergence: while my results highlight loss aversion as the central bias, the literature review placed greater emphasis on the role of narratives in driving bubbles. One possible explanation is that in a survey setting, respondents may recognize loss aversion more readily because it feels rational to admit that losses hurt more than gains. By contrast, acknowledging the power of narratives may be harder, since from the inside they are not always perceived as irrational forces but simply as part of normal decision-making. This remains a hypothesis, but it suggests that narratives might exert their influence in ways that investors themselves do not consciously recognize or admit when reflecting on their own behavior.

Taken together, my analysis of investor profiles and psychology shows that speculative dynamics reflect broader collective patterns rather than simple distinctions between experienced and inexperienced investors. Biases such as loss aversion are widely present, and algorithmic influence is perceived across categories, while risk orientation adds a small nuance: cautious investors tend to recognize these influences more than aggressive ones. These findings confirm the perspective developed in the literature review that bubbles are collective phenomena shaped by universal tendencies, even if investors do not always acknowledge the influence of narratives. I therefore understand investor psychology not as a matter of good versus bad decision-making but as a shared behavioral landscape in which biases, self-perceptions, and collective dynamics interact.

Having established this, I now move to the next part of the analysis, where I look at Collective Dynamics in Digital Communities.

2.3 COLLECTIVE DYNAMICS IN DIGITAL COMMUNITIES.

The literature review highlighted how speculation in the digital era is no longer explained only by rational expectations or isolated biases. Instead, attention functions as a scarce but powerful resource, herding takes the form of visible and accelerated cascades through online platforms, and influencers increasingly act as focal points of trust and coordination. Together, these mechanisms suggest that digital communities transform dispersed psychological tendencies into collective speculative patterns.

With H3, I shift my focus from identifying individual biases to examining how these biases might be amplified, coordinated, and expressed at the collective level. Whereas earlier analyses showed that investors are shaped by systematic tendencies such as loss aversion, here I investigate whether digital infrastructures convert these private dispositions into visible speculative behavior. In doing so, I am not only asking whether biases exist, but whether they are translated into market dynamics through the influence of online communities.

Building on this framework, *H3 suggests: Digital communities transform individual cognitive biases into collective speculative behavior.*

Survey evidence gives me a way to examine this question in practice. I started by looking at how respondents said they are influenced by others: quite a few admitted being swayed by majority behavior (41.8%, Q6), some pointed to viral trends like memes or hashtags (28.8%, Q9), and a smaller group mentioned following financial influencers over traditional analysts (13.0%, Q10). These answers already fit the picture from the literature review, where visibility, imitation, and symbolic authority shape how attention moves online.

But I wanted to go a step further. Instead of stopping at what people say influences them, I compared their general intentions with how they reacted when faced with a more concrete scenario. This, I believe, gives a better test of H3 because it shows whether the hype that spreads so easily in digital communities actually translates into speculative behavior.

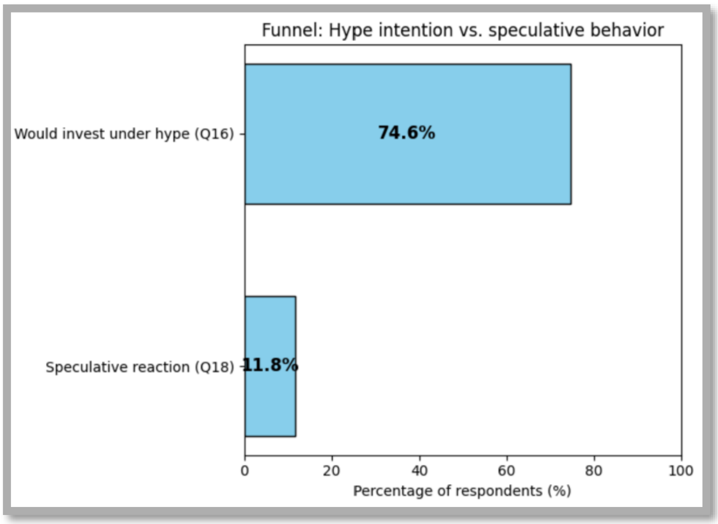

Fig: Hype intention vs. speculative behavior

The chart titled *“Funnel: Hype intention vs. speculative behavior”* makes the gap between words and actions very clear. Almost three quarters of respondents told me they would invest under hype in principle (74.6%, Q16). Yet when I placed them in a more concrete situation with a sudden 200 percent price surge, only 11.8 percent reacted in a speculative way (Q18).

For me, this sharp drop is revealing. It shows that digital communities succeed in building strong speculative narratives, but when it comes to actual decisions many investors hold back. In other words, hype fuels the conversation, yet it does not always translate into behavior at the same scale.

To prove this further I ran chi-square tests. The patterns went in the expected direction. Respondents who felt more influenced by the majority were more likely to invest because of hype, 81 percent versus 70 percent (Q6 × Q16). Those influenced by viral trends also showed a slightly higher tendency to react speculatively to the 200 percent surge scenario (Q9 × Q18). Yet neither result was statistically significant. For example, the majority–hype link produced $\chi^2(1, N=177) = 2.28$, $p = .13$. This suggests that while the descriptive data point toward community influence, I cannot confirm the effect with statistical confidence in my sample.

That is why I see H3 as only partly supported. Compared with H2, where I already found evidence that social influence shapes speculative tendencies, the results here suggest that these effects may be more fragile and context dependent than the narratives alone would imply. Linking back to the literature review, I would say that digital communities amplify and coordinate biases into collective stories, but the actual step from discourse to trading behavior remains uneven. Narratives set the stage, yet the leap into speculation is not guaranteed. Which naturally raises the next question: are digital-era bubbles such as cryptocurrencies and meme stocks ultimately driven more by hype, narratives, and online communities than by traditional fundamentals?

In my review of existing research, I saw that cryptocurrencies, NFTs, and meme stocks are often portrayed as speculative assets whose value depends less on fundamentals and more on collective stories and online attention. Some accounts treat them as “pure hype,” while others argue that investors still rely on conventional reasoning even when engaging with these novel asset classes. This raised the question of whether investors show a stronger reliance on hype when dealing with digital assets compared to more traditional ones.

H4: Cryptocurrencies and meme stocks exemplify how online narratives and hype outweigh traditional fundamentals.

To test this, I divided respondents into two behavioral categories. I considered them hype-driven if they said they would invest under hype (Q16) or chose speculative options in the scenario question such as buying quickly or following influencers (Q18). I considered them fundamentals-driven if they chose more cautious responses in Q18, such as waiting to check fundamentals or refusing to invest due to bubble risk. I then grouped respondents by the type of assets they reported holding (Q2).

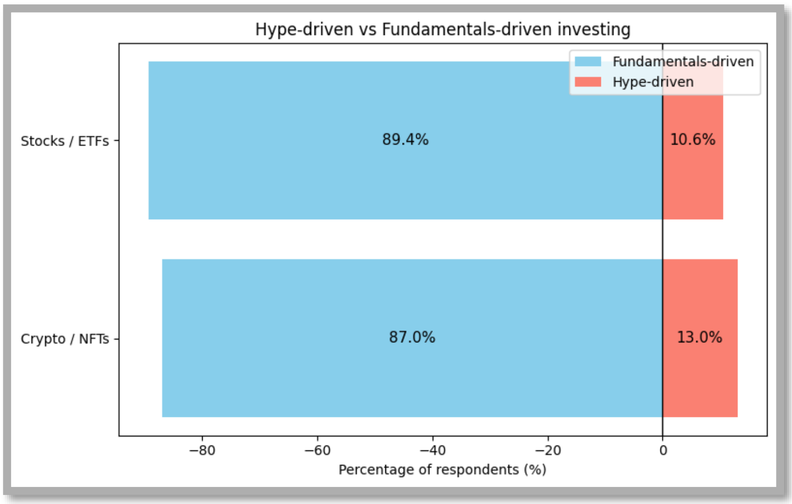

Fig : Hype-driven vs Fundamentals-driven investing

The results are illustrated in the chart titled *“Hype-driven vs Fundamentals-driven investing.”* The chart compares the proportion of fundamentals-driven and hype-driven investors across two categories. Among those who invested in stocks or ETFs, 89.4 percent were classified as fundamentals-driven and 10.6 percent as hype-driven. Among those who invested in cryptocurrencies or NFTs, 87.0 percent were fundamentals-driven and 13.0 percent were hype-driven. The pattern is similar across both groups, although crypto and NFT investors show a slightly higher share of hype-driven behavior.

Because many respondents reported holding a mix of assets, I also looked only at those who invested exclusively in one category. In the crypto and NFT only group (N = 37), 10.8 percent were hype-driven, compared with 4.4 percent in the stocks and ETFs only group (N = 45). While fundamentals remain dominant in both groups, crypto investors in my sample showed more than double the rate of hype-driven behavior compared to traditional investors. The difference is modest and not statistically significant, but it still points in the expected direction.

To complement the descriptive and bivariate analyses, I conducted an exploratory cluster analysis to examine whether distinct investor profiles could be identified based on community influence and speculative orientation. I included five variables: herding (Q6), influence of viral trends (Q9), reliance on financial influencers (Q10), hype investing (Q16), and speculative reaction to the 200 percent scenario (Q18). All variables were standardized prior to clustering. Model selection criteria (inertia, silhouette, Calinski–Harabasz, Davies–Bouldin) indicated that a four-cluster solution provided the best balance of separation and interpretability. Although the silhouette scores were modest (≈0.31), this solution offered a clearer picture than the three-cluster model, particularly by distinguishing between investors who were exposed to community narratives but remained cautious in their behavior, and those who combined community influence with speculative reactions.

The four clusters can be summarized as follows:

- **Cautious fundamentals (N ≈ 42):** low herding and viral influence, not hype investors, and not speculative in scenarios.
- **Pragmatic hype investors (N ≈ 63):** acknowledge hype in principle but refrain from speculative scenario reactions.
- **Hardcore community speculators (N ≈ 21):** high herding and viral influence, strong hype orientation, and speculative in the scenario.
- **Narrative-driven investors (N ≈ 51):** very high scores on herding, viral trends, and influencers, nearly all hype investors, but ultimately cautious in their scenario behavior.

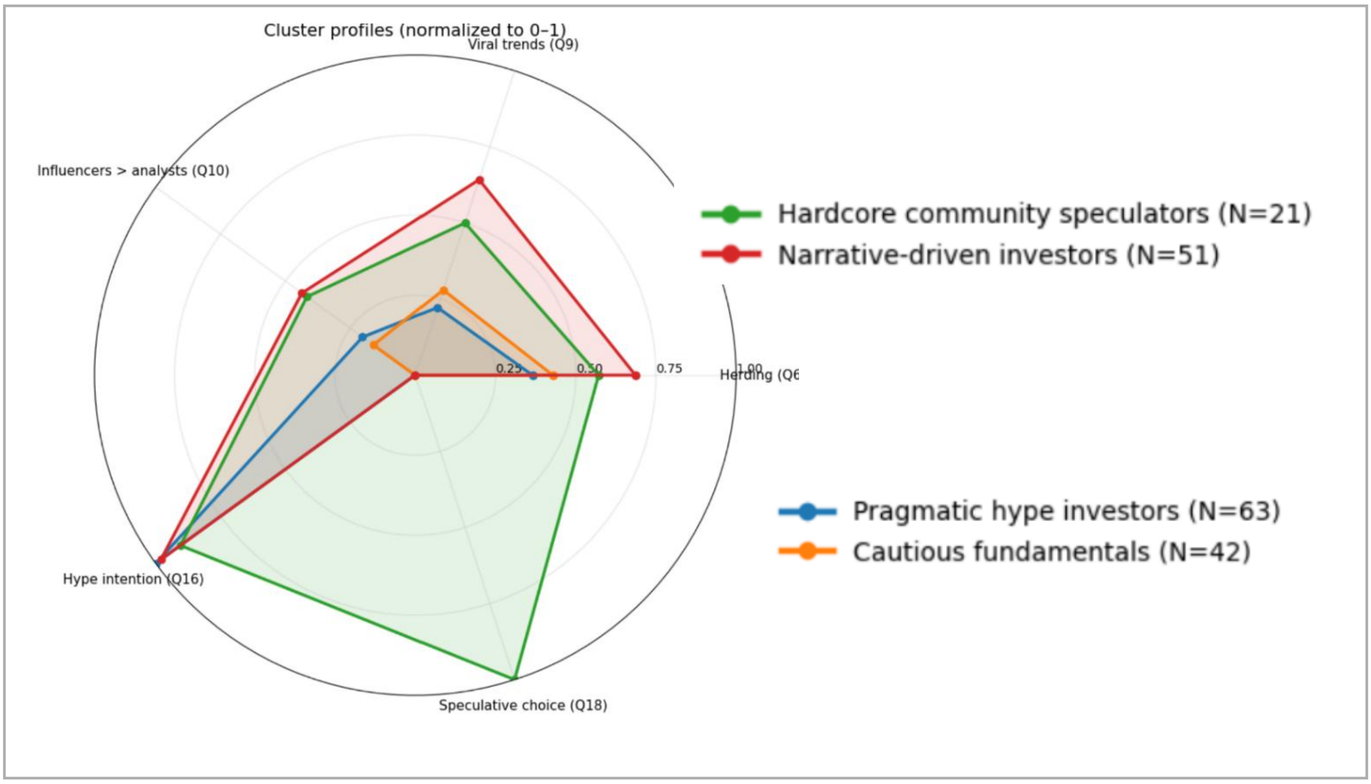

Radar figure

The radar figure provides a visual summary of these profiles. The chart plots mean scores for each cluster across the five variables. As shown, cautious fundamentals remain close to the center across all axes, reflecting low exposure to hype and speculation. Pragmatic hype investors extend outward only on hype intention but remain restrained on speculative choice. Hardcore community speculators stand out with consistently high values across all dimensions, while narrative-driven investors display strong community influence and hype intention but little speculative action.

Taken together, these results suggest that while most respondents remain fundamentally oriented, a visible minority of around 10 to 15 percent are clearly shaped by digital community dynamics and speculative tendencies. The cluster analysis adds nuance to the descriptive findings: susceptibility to hype does not translate uniformly into speculative action but manifests in multiple forms, ranging from pragmatic acknowledgment to extreme speculation. It is important to stress, however, that the silhouette scores were modest and cluster sizes uneven, with one group comprising only 21 respondents. These patterns are therefore exploratory rather than conclusive. Nonetheless, they reinforce the overall finding that digital communities give rise to heterogeneous investor orientations and highlight the presence of a minority cluster of community-driven speculators.

For me, the key takeaway is that digital assets do not simply replace fundamentals with hype. Most respondents continue to frame their decisions in terms of fundamentals, even when trading in markets where the very meaning of fundamentals is contested. This links back to what I noted in my literature review: digital speculation does not just accelerate traditional dynamics; it also unsettles the categories we use to interpret them. If assets can be created in ways that are structurally detached from conventional anchors, then the boundary between hype and fundamentals becomes increasingly blurred.

I therefore see H4 as only partly supported. Crypto and NFT investors show a slightly higher tendency toward hype, but fundamentals still dominate across the board. In line with the literature, I interpret this as evidence that digital communities shape speculative behavior not by erasing fundamentals altogether, but by reshaping the way investors understand the balance between narratives and value.

Taken all together, these findings suggest that collective dynamics in digital communities are less about replacing fundamentals with hype and more about reshaping the landscape in which investors interpret both. Narratives, memes, and influencers create powerful frames for how markets are imagined, but most investors still temper these stories with fundamental reasoning. What emerges is not a wholesale shift toward speculation but a hybrid space where fundamentals and hype coexist, and where investors position themselves differently depending on their susceptibility to community influence.

This perspective brings me to the next part of the analysis. If digital communities provide the narratives and social cues that shape speculative tendencies, algorithms determine which narratives gain visibility in the first place. In section 2.4, I turn to the role of algorithmic curation in structuring exposure and shaping beliefs about markets.

2.4. TECHNOLOGICAL INFRASTRUCTURES AND AMBIVALENCE

In the literature review, I explained how attention functions as a form of currency, how networked herding coordinates collective behavior, and how influencers transform cultural authority into market signals. Each of these dynamics is mediated by platform infrastructures such as ranking systems, recommendation engines, trending lists, and personalized feeds. Building on this logic, I test the following hypothesis:

H5. Algorithms curate exposure and shape beliefs about markets.

If algorithms determine what I see, then (i) the sources from which I obtain financial news should be disproportionately shaped by algorithmic platforms, and (ii) within those platforms, greater exposure should have a stronger association with herding-oriented beliefs and behaviors.I started by looking at the information diet of respondents. To do this, I summarized their multi-select answers about where they follow finance and investing, and I grouped the open “Other” responses into themes. This made it possible to see whether people’s exposure is shaped more by algorithmic platforms, community-driven spaces, or traditional media. To test H5 more directly, I ran moderation models. The key outcome was self-reported herding (Q6), predicted by how much financial content people said they were exposed to (Q8, measured on a 1–5 scale). The moderator was whether or not respondents used a given platform type, which let me compare Users and Non-users. For clarity, I divided platforms into two groups: algorithm-heavy (X/Twitter, YouTube, Instagram, TikTok) and community-driven (Reddit, Discord). At first I tried a four-way split that also included traditional/professional outlets and “other learning” sources, but because nearly everyone used algorithm-first apps, that approach added noise without producing new insight. I therefore kept to the simpler two-category comparison. For each group, I estimated OLS models with an interaction between exposure and user status, and I plotted the slopes separately for Users and Non-users. The interactions did not reach conventional levels of statistical significance, and the models explained only a small share of variance. Even so, I report and visualize them, since the slope patterns still give useful signals for evaluating H5.

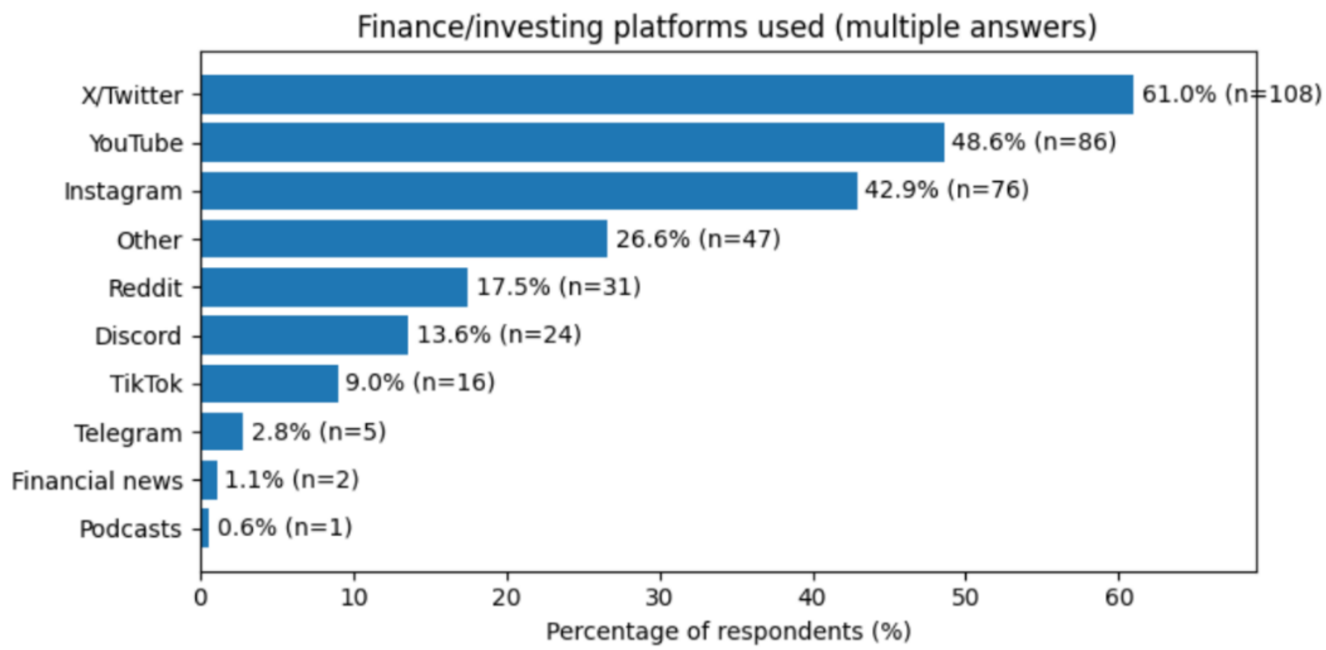

Fig 3.1 Most common platforms for finance/investing updates

From the bar chart, respondents overwhelmingly cite algorithm-ranked social platforms: X/Twitter 61.0% (n=108), YouTube 48.6% (n=86), Instagram 42.9% (n=76). Smaller shares mention Reddit 17.5% (n=31) and Discord 13.6% (n=24), while TikTok 9.0% (n=16) is present but niche. “Other 26.6% (n=47)” appears sizable, but traditional outlets barely register: Financial news 1.1% (n=2) and Podcasts 0.6% (n=1). This pattern already hints that curated feeds dominate how people *enter* the finance conversation.

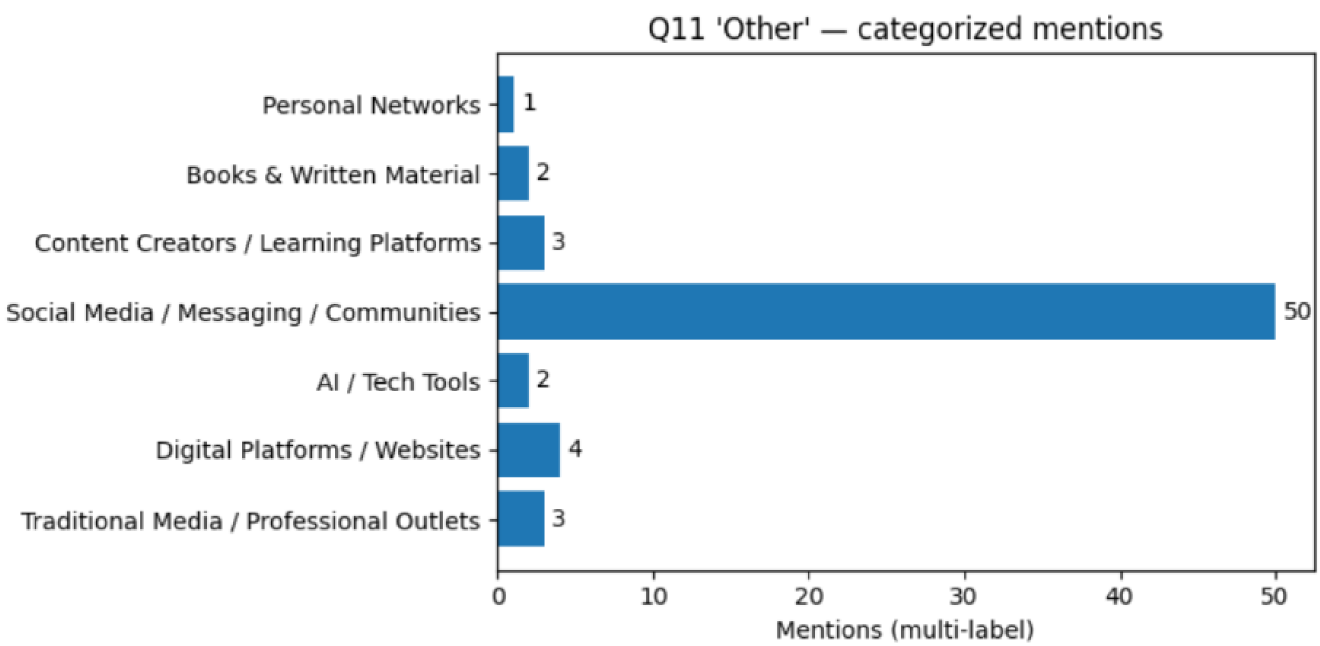

Fig 3.2 Categorization of open-ended ‘primary source

When I categorized “Other” sources, the skew toward “Social Media / Messaging / Communities” (n=50) dwarfs all else. The remaining categories are small: Digital Platforms/Websites (n=4), Traditional Media/Professional (n=3), Content Creators/Learning (n=3), AI/Tech Tools (n=2), Books & Written (n=2), Personal Networks (n=1). Even outside the preset options, respondents still point to social/algorithmic environments. Together with Figure 3.1, this supports the *exposure is infrastructured* premise behind H5.

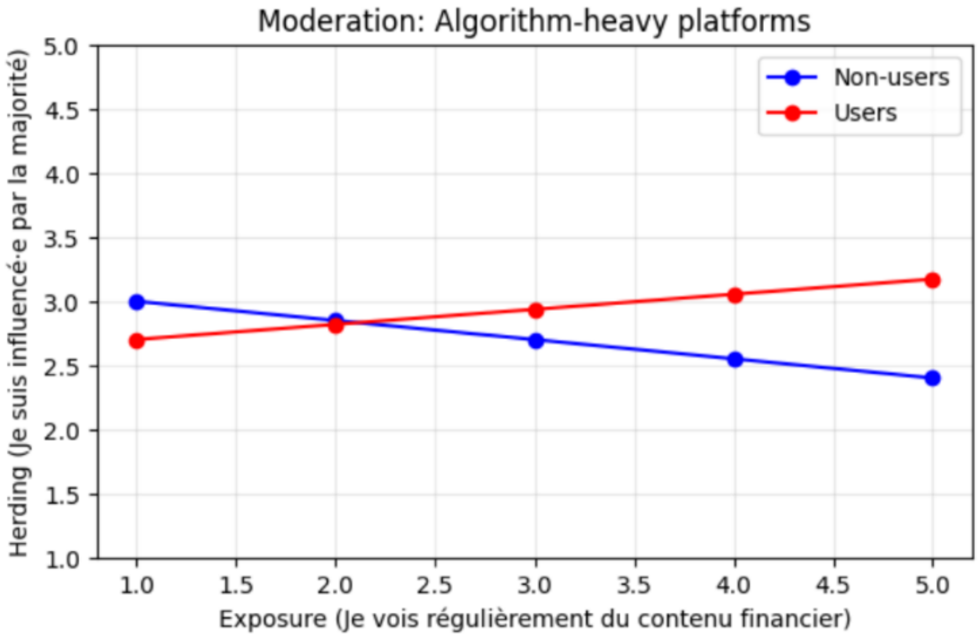

Fig 3.3Caption: “Exposure (Q8) vs. Herding (Q6), split by Users vs. Non-users of algorithm-heavy platforms.”

On algorithm-heavy platforms, the Users line (red) slopes upward: more exposure links to *more* herding. The Non-users line (blue) slopes downward: more exposure links to *less* herding. The divergence suggests that when algorithms curate the feed, additional exposure may nudge beliefs toward perceived majority views—exactly the amplification mechanism theorized in the review.

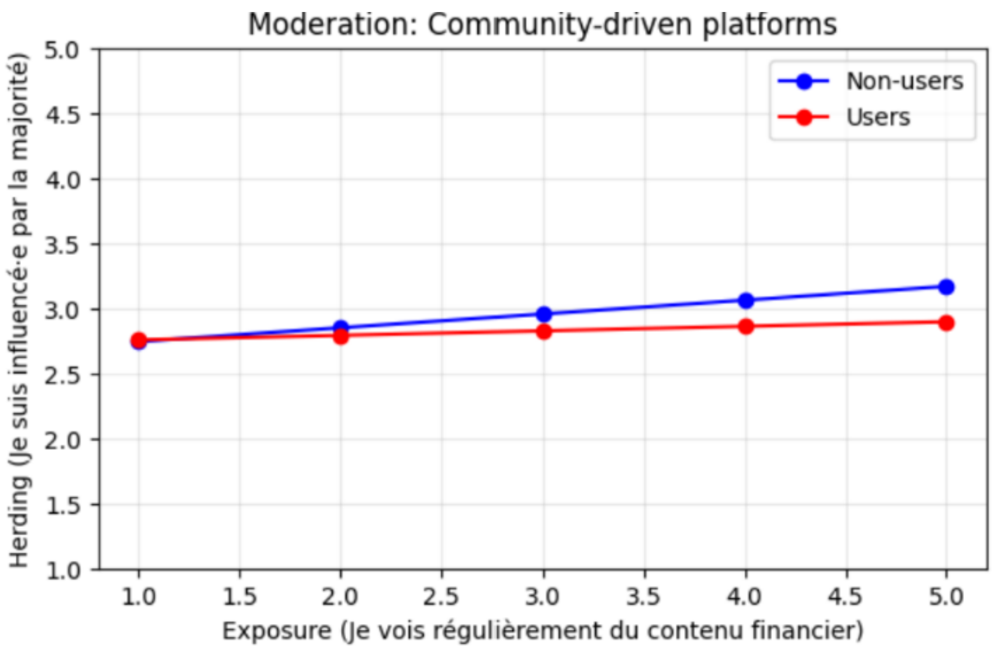

Fig 3.4: “Exposure (Q8) vs. Herding (Q6), split by Users vs. Non-users of community-driven platforms.”

For Reddit/Discord, both lines are shallow and almost parallel: exposure relates weakly and similarly to herding for Users and Non-users. That implies community spaces might *add* conversations without strongly *reshaping;* how exposure maps into conformity.

The descriptive patterns align with the literature. Most respondents operate within algorithmic feeds (Figures 3.1–3.2), and it is only in those environments that I see a direction consistent with H5: exposure relates more strongly to herding among Users (Figure 3.3). This echoes prior work on attention as currency, networked herding, and algorithmic amplification. Ranking systems and recommendation engines do more than display content; they shape what appears popular or consensual, nudging beliefs toward the perceived majority.

The statistical evidence, however, is weaker. The interaction terms were not significant, and the models offered only modest fit, which limits any strong causal claim. There are three likely reasons. First, because nearly everyone uses algorithmic platforms, the distinction between Users and Non-users is too blunt to capture meaningful variation. Second, the measures are noisy: single-item exposure and herding scales, together with self-reported platform use, inevitably reduce precision. While questions on algorithm perceptions (Q12–Q13) are closer to the construct, they were not embedded in the moderation models. Third, timing matters. Algorithmic amplification may be most powerful during volatile, narrative-driven episodes, whereas a cross-sectional snapshot outside such shocks leaves little variation to exploit.

If I were to tighten the test, I would replace binary user flags with measures of usage intensity and feed personalization, modeling the interaction between exposure and algorithmic personalization directly, whether through SEM or hierarchical OLS. Adding controls for demographics and trading practices would help reduce omitted-variable bias. Re-estimating the models around events such as earnings announcements, viral posts, or volatility spikes would also create conditions where algorithmic amplification is more likely to leave a trace. Ideally, survey data could be paired with behavioral traces such as time-on-feed or network structure to move from perceptions toward mechanisms.

Taken together, the figures are consistent with the spirit of H5: algorithms curate exposure and subtly steer beliefs. At the same time, the dataset cannot isolate these effects statistically. This tension is itself revealing. Algorithmic influence appears pervasive enough to shape descriptive patterns, yet too diffuse and my measures too coarse to identify cleanly through simple interaction models. This conclusion fits with the broader literature: platforms are not neutral channels but shaping environments, and detecting their influence requires research designs that match the moments when algorithmic amplification is most active.With this in view, I now shift from what algorithms do to how investors relate to them

The literature review also pointed to a different kind of ambivalence: not only do platforms filter and amplify information, they increasingly act as *advisors* through AI-driven systems. Research on robo-advisors, predictive analytics, and algorithmic decision aids shows that while investors value the efficiency and consistency of such tools, they also question their transparency, credibility, and potential to destabilize markets. This tension sets the stage for H6, which asks : *Investors and potential investors trust AI-based financial tools as credible sources of foresight.*

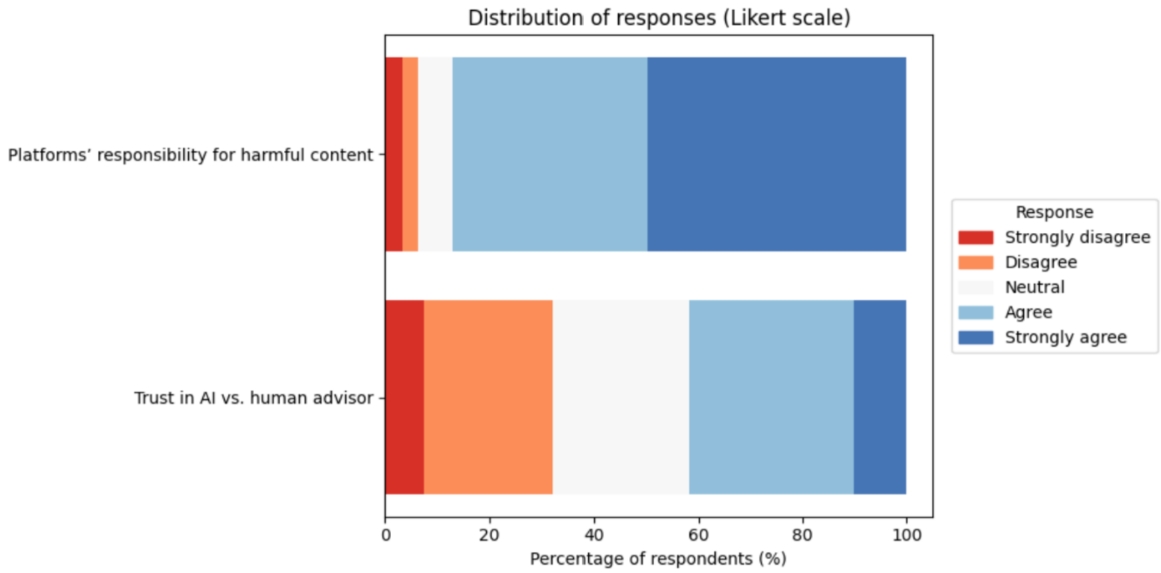

Figure 4.1 "Distribution of responses (Likert scale)".

To begin, I looked at the descriptive distributions of the two central items: *trust in AI compared to a human advisor* and *platforms' responsibility for harmful content*. These are shown in Figure 4.1 "Distribution of responses (Likert scale)".

The figure makes the contrast very clear. On Trust in AI vs. human advisor, many respondents chose the Neutral category, while the remainder were split fairly evenly between agreement and disagreement. This divided pattern indicates that AI is still perceived as an uncertain or emerging option: some respondents see efficiency and objectivity as its main advantages, while others remain cautious about its opacity and potential risks.

By contrast, the Platforms' responsibility item shows far less hesitation. In Figure 4.1, the majority of respondents selected Agree or Strongly agree, with only a handful disagreeing. This suggests that investors overwhelmingly view platforms not as passive pipes but as active agents who shape the financial information environment and should therefore be held accountable.

Taken together, the figure suggests a paradox also identified in the literature: while personal trust in AI remains unsettled, expectations of platform responsibility are much more firmly established. Yet descriptive results alone cannot show whether these two attitudes are connected or independent of each other. To test that, I next turned to a Chi² test of independence.

After looking at the descriptive patterns, I wanted to know whether the two attitudes are actually connected. To do this, I compared trust in AI as an advisor with expectations of platform responsibility. The results are shown in *Figure 4.2 "Observed frequencies".*

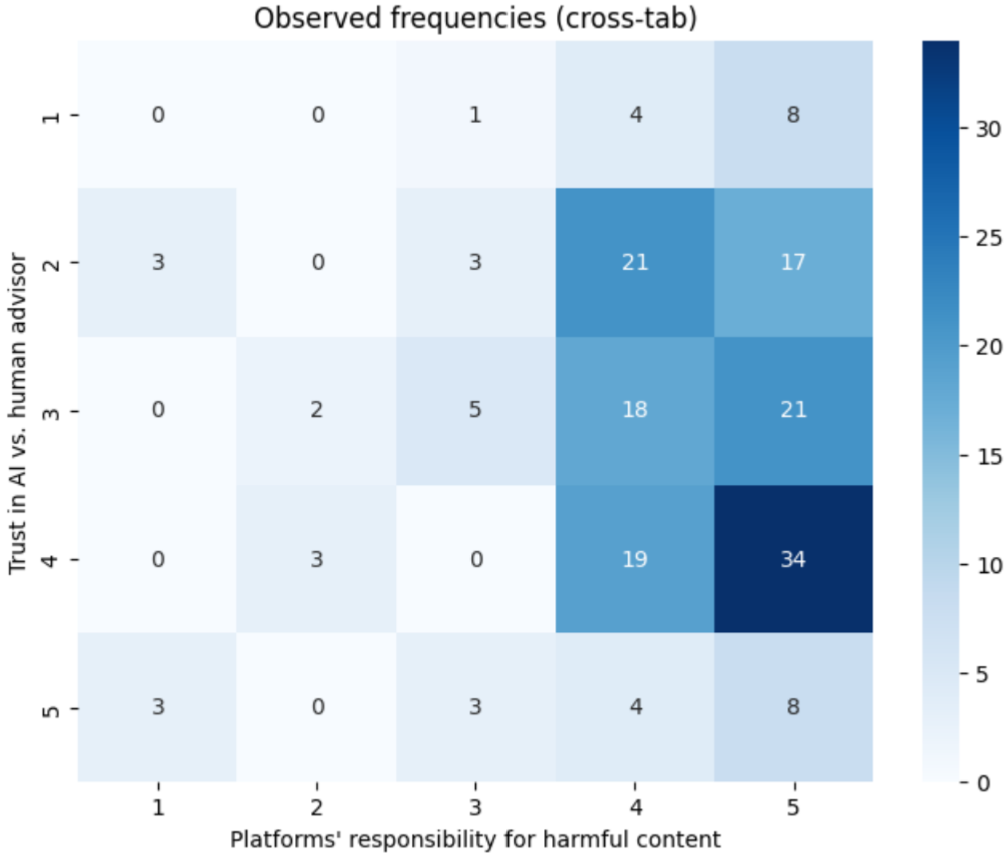

Figure 4.2 "Observed frequencies".

When I look at the heatmap, I can see that the darkest areas are in the upper right. This means that the largest number of respondents both trusted AI more and also agreed that platforms should be responsible. For instance, there are 34 respondents in the cell for Trust = 4 and Responsibility = 5, and 21 in the cell for Trust = 3 and Responsibility = 5. By contrast, the lower-left part of the table is almost empty, which shows that very few respondents both distrusted AI and rejected responsibility.

To test this more formally, I ran a Chi² test of independence. The result gave me a Chi² statistic of 31.9 with 16 degrees of freedom and a p-value of 0.0103. Since the p-value is below 0.05, I can reject the idea that the two attitudes are independent. In other words, trust in AI and expectations of responsibility tend to move together. People who are more open to trusting AI are also more likely to demand strong accountability from platforms. On the other hand, people who are skeptical of AI often show weaker or uncertain views about platform responsibility.

For me, this extends what I saw in the descriptive results. It shows that the ambivalence around AI and the consensus around responsibility are not two separate stories. They are connected in a way that makes sense: some investors hold a coherent position where confidence in AI as an advisor comes together with strong expectations that platforms should act responsibly in shaping financial information.

This finding also connects back to the literature I reviewed earlier. Scholars have pointed out that algorithms create a paradox for investors. On one hand, they raise doubts because of their opacity and possible destabilizing effects. On the other hand, these same systems are expected to provide safeguards and stability. My Chi² analysis confirms this paradox in the survey data. Individual trust and collective demands for governance are intertwined, which shows that investors think about AI not in isolation but as part of the broader system that shapes financial markets.

Now that I have evidence that trust in AI and expectations of responsibility move together, I want to know who is most likely to hold these views. The Chi² test shows association, not prediction or group differences. To move from “are they linked” to “who holds them,” I turn to logistic regression.

After the descriptives and the Chi² test showed that trust in AI and expectations of responsibility move together, I estimated two logistic models to see who holds these views. I first tried the survey’s full category detail, but sparse cells created non-convergence and unstable coefficients. I therefore simplified the predictors so the models would run cleanly: education to low vs high, risk profile to prudent, moderate, aggressive, age to under 25, 25–44, 45+, and investor status to already invested, intends, non-investor. With these adjustments the models converged and produced interpretable results (Table 1, N = 177).

Predictor	Trust in AI (Model A)	95% CI	p-value	Responsibility (Model B)	95% CI
Investor status (ref = Intends)					
Already invested	1.97	0.96 – 4.06	0.065	0.97	0.34 – 2.72
Non-investor	0.75	0.12 – 4.82	0.766	0.62	0.06 – 6.41
Risk profile (ref = Aggressive)					
Moderate	0.78	0.32 – 1.92	0.584	1.03	0.30 – 3.46
Prudent	0.37*	0.15 – 0.92	0.032	1.6	0.45 – 5.63
Age (ref = <25)					
25–44	0.7	0.34 – 1.44	0.327	0.42	0.14 – 1.25
45+	0.67	0.21 – 2.09	0.487	1.52	0.16 – 14.21
Education (ref = High)					
Low	3.72**	1.56 – 8.86	0.003	1.01	0.30 – 3.37

Model A: Trust in AI.

Trust is unevenly distributed. Respondents who had already invested were almost twice as likely to trust AI compared to those who intend to invest (OR = 1.97, $p = .065$, marginal). Prudent investors were less likely to trust AI than aggressive investors (OR = 0.37, $p = .032$), which points to risk orientation as a key line of differentiation. Lower education was associated with higher trust (OR = 3.72, $p = .003$). Age did not matter once other factors were included.

Model B: Responsibility.

Support for platform responsibility did not vary systematically across groups. None of the predictors were significant and the odds ratios sat close to one. In other words, the expectation that platforms should curb harmful content is broadly shared.

Putting this together, I read H6 as follows. Trust in AI is socially differentiated, shaped by experience in markets, risk psychology, and education. Responsibility is universal, a cross-cutting norm that spans those differences. This maps directly onto the literature I reviewed: investors relate to algorithms with ambivalence at the individual level while converging on demands for governance at the system level. The same infrastructures that can feel opaque or destabilizing are also the ones investors expect to provide safeguards. In short, H6 shows the paradox of algorithmic finance in action: infrastructures oscillate between fueling speculative dynamics and providing mechanisms of control.

I showed that platforms play a central role in shaping the information diet of investors. As Figures 3.1 and 3.2 illustrate, most respondents rely on algorithmic feeds for financial news, which supports the idea from the literature that attention, herding, and influencer power are routed through ranking and recommendation systems.

I also found signs of amplification, but only at a descriptive level. In Figure 3.3, exposure is more strongly linked to herding among users of algorithm-heavy platforms, while Figure 3.4 shows little difference in community spaces. The interactions were not statistically significant and the models had modest fit, which I interpret as evidence that algorithmic influence is pervasive but difficult to capture cleanly with my measures and cross-sectional design.

Turning to AI, the results revealed ambivalence. As Figure 4.1 shows, investors are divided on whether they trust AI-based tools, but they are much more united in expecting platforms to take responsibility for harmful content. Figure 4.2 and the Chi² test ($\chi^2 = 31.9$, $df = 16$, $p = 0.0103$) confirmed that higher trust in AI often goes hand in hand with stronger support for accountability.

The logistic regressions refined this picture. Trust in AI is socially differentiated: it is somewhat higher among those already invested, lower among prudent investors, and significantly higher among respondents with lower formal education. By contrast, support for platform responsibility is broadly shared and does not vary by status, risk profile, age, or education.

Overall, these results mirror the literature review. Platforms are not neutral pipes: they curate exposure and nudge beliefs, while also carrying a governance expectation. Investors remain cautious about relying on AI as an advisor, yet they want the same infrastructures to safeguard the financial information environment.

I recognize the limits of my data. Binary user flags, single-item scales, and a snapshot taken outside periods of volatility likely muted the effects I could detect. A stronger design would measure usage intensity and feed personalization directly, integrate perceptions of algorithmic influence (Q12–Q13), and capture behavior during moments of market stress or through digital trace data.

Taken together, my findings highlight the paradox of technological infrastructures in finance. They fuel speculative dynamics by amplifying exposure and social cues, while at the same time being expected to provide mechanisms of control. This ambivalence is the central lesson of my analysis.

RECOMMENDATIONS AND PERSPECTIVES

RECOMMENDATIONS

The methodological choices adopted in this thesis provided valuable insights, but they also highlighted several limitations that suggest concrete recommendations for future research. The reliance on a quantitative survey successfully offered measurable data regarding how investors perceive financial bubbles in the digital age, yet the responses also revealed gaps that should be addressed to strengthen validity and depth. A first limitation comes from the design of the questionnaire itself. Most questions were structured in the form of Likert scales, which enabled comparability and statistical treatment but restricted the richness of the answers. For example, when participants were asked about their trust in artificial intelligence as a diagnostic tool for financial markets, the majority selected “neutral” or “somewhat agree,” which left little room to understand why they felt this way. Some respondents may have been cautious due to a lack of familiarity with AI, while others may have recognized both benefits and risks but were unable to express this dual perspective within a closed scale. To overcome this limitation, future surveys should incorporate open-ended questions or scenario-based items that allow respondents to explain their reasoning. For instance, presenting a short vignette where an AI system successfully predicts a market downturn, or conversely fails to detect a bubble, could invite participants to articulate their confidence or skepticism in more detail. This type of instrument would reduce the ambiguity of neutral responses and provide a more nuanced understanding of investor attitudes.

The sampling strategy also raises important methodological considerations. In this study, a significant portion of the respondents were young, digitally engaged individuals, many of whom reported using platforms such as Reddit, Twitter, or Discord to inform their investment decisions. While this profile is consistent with the phenomenon of meme stocks and digital speculation, it does not reflect the full diversity of the investing population. Professional traders, institutional investors, and older cohorts were largely absent from the dataset. This creates a risk of overgeneralization, since the perceptions of digitally native retail investors may differ substantially from those of more experienced or conservative actors. For example, when asked whether online communities exert a strong influence on market movements, younger respondents tended to strongly agree, whereas in professional contexts this influence might be perceived as temporary noise. To strengthen external validity, future research should adopt a stratified sampling method, deliberately including different groups by age, profession, and region. Such comparisons would make it possible to assess whether enthusiasm for digital platforms and AI tools is specific to certain demographics or whether it represents a broader trend across the financial ecosystem.

The timing of data collection also limited the scope of this study. The survey was conducted at a single moment, which means it captured a snapshot of investor perceptions but not their evolution over time. This is particularly problematic in the context of speculative bubbles, which are inherently dynamic and subject to rapid shifts. For example, during the field study, several respondents referred to the volatility of cryptocurrencies, noting that their confidence in assets such as Bitcoin could fluctuate dramatically within weeks. However, the survey design did not allow these changes to be systematically captured. A longitudinal design, with repeated surveys at different points in time, especially during market events such as sudden crashes or viral rallies, would provide a much clearer picture of how digital narratives and algorithmic signals influence behavior dynamically. This would also allow researchers to distinguish between stable attitudes (e.g., general distrust of speculative assets) and event-driven reactions (e.g., temporary enthusiasm during a meme stock surge).

Another area for improvement concerns the analytical tools used. In this thesis, descriptive statistics and correlations offered useful insights into the relationship between variables, such as the link between age and reliance on digital platforms. For instance, the results showed that younger respondents were significantly more likely to trust influencers or online communities as a source of financial information, while older respondents relied more on traditional media. While these patterns are informative, more advanced statistical methods would provide deeper insights. Regression models could help determine whether age remains a significant factor once other variables, such as trading experience or exposure to AI tools, are controlled for. Structural equation modeling could even test complex relationships, such as whether digital engagement mediates the link between attention to narratives and susceptibility to bubbles. By adopting such techniques, future research could move beyond descriptive associations and begin to explain causal mechanisms.

The study also revealed the limitations of relying exclusively on self-reported perceptions. For example, when asked whether they had participated in speculative episodes such as meme stock rallies or cryptocurrency surges, several respondents reported “no” while at the same time acknowledging frequent trading activity on platforms like Robinhood or Binance. This discrepancy suggests a degree of underreporting or reinterpretation, possibly due to bias in social desirability. Participants may hesitate to admit being influenced by hype, even if their trading records suggest otherwise. To address this, future research should combine survey data with behavioral evidence, such as digital trace analysis or anonymized trading data. For instance, sentiment analysis of Reddit posts could be cross-referenced with survey responses to evaluate whether declared attitudes align with actual participation in online discussions. This triangulation would reduce self-reporting bias and create a more robust foundation for conclusions.

Experimental designs represent another promising avenue. In the present study, participants were asked to report their general level of trust in algorithms or influencers, but they were not placed in a simulated decision-making situation. As a result, the findings may reflect abstract opinions rather than real reactions. A complementary approach would be to embed experiments directly into the survey. For example, respondents could be shown two price charts, one accompanied by a viral social media post and the other not, and asked to choose whether they would invest. Observing the differences in choice would provide direct evidence of the effect of digital narratives on decision-making. Similarly, exposing participants to AI-generated trading recommendations could reveal whether they follow automated advice even when it contradicts their initial judgment. Such experiments would significantly enhance the explanatory power of future studies.

Cultural context is another area that deserves more attention. Most respondents in this thesis came from similar socio-cultural backgrounds, which limited the ability to assess whether attitudes toward speculation and digital platforms vary across countries. Yet financial behavior is deeply shaped by cultural factors such as trust in institutions, risk tolerance, and attitudes toward technology. For example, investors in highly regulated markets may view AI-driven trading tools with more skepticism than those in less regulated contexts. Translating and adapting the questionnaire for multiple regions would therefore broaden the generalizability of findings and highlight potential cultural differences in the perception of bubbles and digital finance.

Finally, ethical considerations must be part of any methodological reflection. The anonymity of participants was preserved in this study, but future research incorporating digital traces or trading data would need to establish clear safeguards for privacy and consent. Furthermore, as AI increasingly becomes both the subject and the tool of research, transparency about how algorithms are used in data analysis is essential. Respondents should know how their information is processed, and researchers must ensure that AI-generated insights are critically assessed rather than accepted at face value. This is especially important given that the very topic of this research highlights the risks of overreliance on opaque algorithms.

In summary, the field study demonstrated both the value and the limitations of a survey-based design. While it provided clear evidence that investors are influenced by digital platforms and narratives, the methodology constrained the depth, diversity, and dynamism of the findings. Future research should therefore expand the design along several dimensions: more open and scenario-based questions, broader and more representative samples, longitudinal and experimental approaches, triangulation with behavioral data, and careful attention to cultural and ethical dimensions. Only through such methodological pluralism can future studies capture the full complexity of speculative behavior in the digital era, where financial decisions are no longer driven solely by fundamentals but also by stories, communities, and algorithms.

PERSPECTIVE

Looking ahead, this thesis also invites several perspectives regarding methodological innovation. The current design proved effective for capturing perceptions of digital speculation, yet the rapid transformation of financial markets suggests that future research must adopt more flexible and adaptive methods. Speculation today is not static phenomenon but one that evolves at the intersection of technology, culture, and investor psychology. As such, methodologies must be designed to evolve in parallel, capable of tracking new practices as they emerge.

One important perspective is the integration of real-time data collection. While surveys offer valuable snapshots, speculative dynamics unfold in fast and unpredictable cycles, often triggered by sudden bursts of attention on social media. To capture these shifts, future research could rely on tools that gather data continuously. For example, scraping digital platforms in real time to monitor trading discussions, memes, or influencer content could be paired with investor feedback to map how narratives evolve hour by hour. Combining this type of live observation with survey responses would create a richer and more temporally sensitive methodology, offering a clearer link between discourse, sentiment, and actual trading outcomes.

Another promising perspective concerns the use of experimental and simulation-based methods. As digital speculation increasingly resembles a complex interaction between humans and algorithms, future research should replicate these environments in controlled settings. Simulations could, for instance, expose participants to algorithmic trading recommendations, varying the framing or reliability of the advice, and record how individuals adapt their decisions. Similarly, virtual communities could be recreated to observe how herding behavior emerges when participants are exposed to visible markers of popularity such as likes, trending tags, or comment counts. These controlled experiments would provide insights into causal mechanisms that cannot be fully disentangled in observational surveys.

A further methodological perspective lies in multi-modal triangulation. Speculative phenomena are shaped not only by self-reported attitudes but also by emotions, discourses, and behavioral traces. Future research could therefore combine quantitative instruments with qualitative approaches such as in-depth interviews, digital ethnography, or discourse analysis. For example, interviewing a small number of retail traders about their experiences during a meme stock rally could complement survey data by uncovering motivation that are not easily measurable. Ethnographic immersion in online investor communities might also reveal cultural dynamics, humor, and symbolic practices that traditional questionnaires fail to capture. These qualitative perspectives would enrich the empirical texture and ensure that the complexity of digitally mediated speculation is not reduced to mere numbers.

The use of advanced computational methods also opens important methodological perspectives. Sentiment analysis, natural language processing, and machine learning models now allow researchers to process massive datasets of social media posts or trading records. Future studies could employ these techniques not only to analyze discourse but also to develop predictive indicators of speculative behavior. For instance, shifts in narrative intensity or emotional tone on platforms such as Reddit or Twitter could be modeled to anticipate sudden surges in trading activity. By integrating these computational tools with survey-based approaches, future research could develop hybrid methodologies that bridge human perception with algorithmic signals.

Another perspective relates to cross-cultural and comparative research designs. The findings of this thesis were grounded in a particular context, but speculation is a global phenomenon that unfolds differently depending on regulatory regimes, cultural norms, and levels of technological adoption. A future study could, for example, compare investor attitudes in the United States, Europe, and Asia, highlighting how local narratives and institutional frameworks shape digital speculation in distinct ways. Such comparative perspectives would strengthen the external validity of findings and help identify whether the patterns observed here are universal or context-dependent.

In addition, longitudinal perspectives remain crucial. While cross-sectional surveys provide useful baselines, they cannot fully capture the cyclical and recursive nature of bubbles. Long-term studies that follow the same participants across several years, or that observe investor sentiment across multiple speculative episodes, would offer a more dynamic understanding of how perceptions evolve. This is especially relevant in the digital era, where platforms, influencers, and algorithms change at a rapid pace. A longitudinal design would also make it possible to distinguish between temporary hype cycles and structural shifts in the way speculation is perceived and enacted.

Finally, methodological perspectives must remain attentive to ethical considerations. The incorporation of digital trace data, algorithmic tools, and experimental designs raises new challenges around privacy, consent, and transparency. Future researchers must ensure that the tools used to study speculation do not replicate the same opacity or manipulation that they are meant to investigate. Clear ethical frameworks, explicit communication with participants, and transparency about the role of AI in data analysis will be essential in maintaining the integrity of research in this field.

In conclusion, the perspectives outlined here suggest that the methodology used in this thesis should be seen as a foundation rather than an endpoint. Future research can build upon it by incorporating real-time data, experimental simulations, multi-modal triangulation, advanced computational techniques, cross-cultural comparisons, longitudinal designs, and robust ethical safeguards. Together, these perspectives point toward a methodological toolkit capable of capturing the complexity of speculation in the digital era, where financial decisions are no longer isolated acts of rational calculation but part of a dynamic interplay between narratives, communities, and algorithms.

CONCLUSION

I began this thesis with a simple but unsettling question. I wanted to know how bubbles are being reshaped by digital infrastructures and whether AI can help us detect them before they implode. Along the way, I had to accept that markets today do not only process information. They manufacture it, curate it, and circulate it through platforms that reward speed, emotion, and visibility. Attention behaves like a currency, communities act as amplifiers, and influence often looks more like culture than expertise. I do not treat these as anecdotal distortions. I treat them as structure.

I argued that the old choice between rational markets and behavioral anomalies is no longer useful on its own. What I described across the chapters is a reflexive system where psychology, culture, and computation feed each other. Platforms push stories into the spotlight, fragment audiences into silos, and now even generate the stories that will be believed tomorrow. In other words, the supply of narrative itself has become computational. I do not see this as a faster version of the past. I see it as a qualitative shift.

This is why I built the paradox of AI at the center of my analysis. AI extends foresight, yet it also changes the thing it claims to observe. Once predictive dashboards are embedded in practice they stop being external instruments. They become part of market reflexes. I therefore treat AI less as a neutral forecaster and more as a market institution whose signals can coordinate behavior for better or worse. I call out the risk of prediction theatre, where numbers look confident without delivering causal grip on turning points. I also insist that endogeneity and opacity are not side notes. They are the main story.

I confronted these claims with data, knowing the limits of a single survey. What I saw is consistent with the literature and with my own argument. Investors are split on trusting AI as an advisor, yet they converge on holding platforms responsible for the quality and safety of the information environment. Statistically, trust varies with experience and education, while the demand for responsibility does not. This is the paradox in miniature. Individually, people hesitate. Systemically, they ask for governance. I read that alignment as an opening for institutional design rather than as a contradiction to be explained away.

I will not overclaim on my empirical design. I know that a cross sectional snapshot outside high volatility phases mutes the very mechanisms I care about. I also know that binary usage flags are crude, that Likert scales compress nuance, and that self reporting often hides the gap between what investors say and what they actually do. If I were to redo the study I would embed event time, measure feed personalization and intensity, pair surveys with digital traces, and stage realistic decisions inside the questionnaire. The point is not that my results are wrong. The point is that speculation is dynamic, recursive, and cultural, and methods must meet it on that moving ground.

Looking forward, I do not think the right question is whether AI will replace human judgment. I think the right question is how to define a diagnostic role that helps without deepening fragility. I come out of this work with a conditional answer. AI should be used as a sensor of state rather than an oracle of price. It can map regime shifts in volatility, thinning liquidity, narrative concentration, and correlated positioning. That information is actionable for circuit breakers, margin policy, and supervisory guidance without pretending to know tomorrow's price. This redefinition demands evidence that survives hard timestamps and out of sample tests, and it demands evaluation of behavioral consequences rather than only statistical fit.

Because markets adapt to their instruments, I argue for governance that addresses concentration, opacity, and synchronized error. I favor layered transparency that distinguishes input data, model logic, and deployment context. I want accountability that travels with the signal as it moves through vendors, platforms, and institutions. I also want structural safeguards against monoculture, including diversity mandates in data and models, stress tests that simulate correlated error, and switch off protocols for whole classes of signals when crowding appears. I am not persuaded that disclosure alone will be enough. I am persuaded that institutional design can make a difference if it aims at structure rather than cosmetics.

This leaves me with a clear position. AI is neither a savior nor a villain. It is a conditional diagnostic infrastructure. Used under regimes of transparency and responsibility, it can illuminate fragility without becoming a new source of it. Used as an unregulated predictor of price, it risks turning warning lights into traffic signals that send everyone onto the same bridge at the same time. My conclusion is therefore cautious but constructive. I do not want to ban prediction. I want to domesticate it.

I also need to acknowledge something important. Writing this as a master's thesis was challenging because I was tackling a highly actual issue. The landscape is evolving while I write, and new findings are being made as we speak. That tension between studying a moving target and wanting to frame it analytically was both frustrating and motivating. I accept that my work will soon be overtaken by new data, but I also claim that my framework gives tools to interpret those updates when they arrive.

Finally, I return to why I wrote this in the first place. I wanted to understand bubbles in the world I actually live in. I now see that the line between engine and antidote runs straight through the same infrastructures. The platforms that synchronize attention are the places where detection must occur. The models that promise foresight can stabilize or destabilize depending on how we embed them. My survey confirmed that investors already sense this ambivalence and already expect responsibility. I take that seriously. I end by insisting on an agenda that joins method and governance. Real time, event based, multi modal research to see speculation as it forms. Institutional designs that slow monoculture before it spreads. An explicit norm that judges predictive systems by their systemic footprint, not only by their backtests. If I keep that standard in view, I can say that AI can help. Not by replacing judgment, but by giving it better instruments and better guardrails.

This is my final claim. Bubbles in the digital era are not accidents at the edge of rational markets. They are produced by the entanglement of human psychology with computational infrastructure. I cannot unwind that entanglement. I can decide what role I want AI to play inside it. I choose a role that is narrow, transparent, accountable, and plural. I choose sensors over oracles. I choose design over faith. And I accept that in reflexive systems, the humility to measure without pretending to control may be the most radical form of foresight we can achieve.

BIBLIOGRAPHY

Aggarwal, R., Lucchini, L., Aloosha, A., et al. (2024). Meme stocks and market dynamics. [Working paper].

Ahn, K., Cong, L., Jang, H., & Kim, D. S. (2024). Business cycle and herding behavior in stock returns: Theory and evidence. Financial Innovation, 10(1), 6.

Almeida, J., & Gonçalves, T. C. (2023). A systematic literature review of investor behavior in the cryptocurrency markets. Journal of Behavioral and Experimental Finance, 37, 100785.

Aloosha, A., et al. (2024). AMC and retail coordination in meme stock trading. [Working paper].

Ayala, M. J., Gonzálvez-Gallego, N., & Arteaga-Sánchez, R. (2024). Google search volume index and investor attention in stock market: A systematic review. Financial Innovation, 10(1), 70.

Baklanova, V. (2025). The relationships between RedditSI and BTC exchange characteristics: Do Reddit users still control the market? Eurasian Economic Review, 15(1), 285–306.

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

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

Barber, B. M., & Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. Quarterly Journal of Economics, 116(1), 261–292.

Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. Review of Financial Studies, 21(2), 785–818.

Barber, B. M., Huang, X., Odean, T., & Schwarz, C. (2022). Attention-induced trading and returns: Evidence from Robinhood users. Journal of Finance, 77(6), 3141–3190.

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

Brunnermeier, M. K., & Nagel, S. (2004). Hedge funds and the technology bubble. Journal of Finance, 59(5), 2013–2040.

Buckmann, M., Joseph, A., & Robertson, H. (2021). Opening the black box: Machine learning interpretability and inference tools with an application to economic forecasting. In Data Science for Economics and Finance: Methodologies and Applications (pp. 43–63). Springer.

Bussière, M., & Fratzscher, M. (2006). Towards a new early warning system of financial crises. Journal of International Money and Finance, 25(6), 953–973.

Case, K. E., & Shiller, R. J. (2003). Is there a bubble in the housing market? Brookings Papers on Economic Activity, 2003(2), 299–362.

Cevik, S., Dibooglu, S., & Kutan, A. M. (2022). Investor sentiment and asset return dynamics. International Review of Economics & Finance, 77, 272–285.

Cheah, E.-T., & Fry, J. (2015). Speculative bubbles in Bitcoin markets? An empirical investigation into the fundamental value of Bitcoin. Economics Letters, 130, 32–36.

Chrisostomides, L. D. (2022). Herding in cryptocurrencies: CSSD and CSAD approaches. In FINIZ 2022 – Conference Proceedings (pp. 52–56). Singidunum University.

Ciganovic, M., & D'Amario, F. (2024). Forecasting cryptocurrencies log-returns: A LASSO-VAR and sentiment approach. Applied Economics, 56(58), 8112–8138.

Cookson, J. A., Engelberg, J., & Mullins, W. (2023). Echo chambers: Social media and retail investors. Review of Financial Studies, 36(2), 450–500.

Cooper, M. J., Dimitrov, O., & Rau, P. R. (2001). A rose.com by any other name. Journal of Finance, 56(6), 2371–2388.

Da, Z., Engelberg, J., & Gao, P. (2011). In search of attention. Journal of Finance, 66(5), 1461–1499.

Da, Z., Engelberg, J., & Gao, P. (2015). The Sum of All FEARS: Investor sentiment and asset prices. Review of Financial Studies, 28(1), 1–32.

Daniel, K., Hirshleifer, D., & Subrahmanyam, A. (1998). Investor psychology and security market under- and overreactions. Journal of Finance, 53(6), 1839–1885.

Daniel, K., Hirshleifer, D., & Sun, L. (2020). Financing-based asset pricing. Journal of

Financial Economics, 135(1), 122–149.
Davenport, T. H., & Beck, J. C. (2001). The attention economy: Understanding the new currency of business. Harvard Business School Press.
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.
Eichengreen, B., Rose, A. K., & Wyplosz, C. (1998). Contagious currency crises. NBER Working Paper No. 5681.
Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. Journal of Finance, 25(2), 383–417.
Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. Journal of Financial Economics, 33(1), 3–56.

Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. Journal of Financial Economics, 116(1), 1–22.
Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270(2), 654–669.
Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. Public Opinion Quarterly, 80(S1), 298–320.
Frankel, J. A., & Rose, A. K. (1996). Currency crashes in emerging markets: An empirical treatment. Journal of International Economics, 41(3–4), 351–366.
Gerlach, J.-C., Demos, G., & Sornette, D. (2019). Dissection of Bitcoin's multiscale bubble history from January 2012 to February 2018. Royal Society Open Science, 6(7), 180643.
Giglio, S., Maggiori, M., Stroebel, J., & Utkus, S. (2021). Five facts about beliefs and portfolios. American Economic Review, 111(5), 1481–1522.
Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), Media Technologies: Essays on Communication, Materiality, and Society (pp. 167–194). MIT Press.
Gorton, G. (2008). The panic of 2007. In Maintaining Stability in a Changing Financial System (pp. 131–262). Federal Reserve Bank of Kansas City.
Goldhaber, M. H. (1997). The attention economy and the Net. First Monday, 2(4).
Hamilton, J. D., & Susmel, R. (1994). Autoregressive conditional heteroskedasticity and changes in regime. Journal of Econometrics, 64(1–2), 307–333.
Harras, G., & Sornette, D. (2011). How to grow a bubble: A model of myopic adapting agents. Journal of Economic Behavior & Organization, 80(1), 137–152.
Haykir, S., & Yagli, I. (2022). Herding behavior in cryptocurrency markets during the COVID-19 pandemic. Finance Research Letters, 44, 102–140.
Hendershott, T., & Riordan, R. (2013). Algorithmic trading and the market for liquidity. Journal of Financial and Quantitative Analysis, 48(4), 1001–1024.
Hirano, T., & Toda, A. A. (2024). Rational bubbles: A review. Journal of Economic Perspectives, 38(2), 203–226.
Hull, I., & Qi, Y. (2024). The impact of finfluencers on retail investment. SSRN Working Paper.
IOSCO. (2025). Digital engagement practices (DEPs). Final report.

Iyengar, R., Van den Bulte, C., & Valente, T. W. (2010). Opinion leadership and social contagion in new product diffusion. Marketing Science, 30(2), 195–212.
Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2), 263–291.
Kaminski, J., & Gloor, P. (2014). Nowcasting the Bitcoin market with Twitter signals. arXiv preprint.
Kaminsky, G. L., & Reinhart, C. M. (1999). The twin crises: The causes of banking and balance-of-payments problems. American Economic Review, 89(3), 473–500.
Ke, Z. T., Kelly, B. T., & Xiu, D. (2019). Predicting returns with text data. NBER Working Paper No. 26186.
Kindleberger, C. P., & Aliber, R. Z. (2011). Manias, Panics, and Crashes: A History of Financial Crises (6e éd.). Palgrave Macmillan.
Kim, J., Muhn, M., & Nikolaev, V. (2024). Narrative extraction with generative AI: Risks and opportunities. Journal of Accounting Research, 62(3), 545–578.

Klinge, M., et al. (2025). Tesla, narratives, and valuation in the digital era. [Working paper].
Lai, H.-H., Chang, T.-P., Hu, C.-H., & Chou, P.-C. (2022). Can Google Search Volume Index predict the returns and trading volumes of stocks in a retail investor-dominant market? Cogent Economics & Finance, 10(1), 2014640.
Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D., & Van Alstyne, M. (2009). Computational social science. Science, 323(5915), 721–723.
LeRoy, S. F., & Porter, R. D. (1981). The present-value relation: Tests based on implied variance bounds. Econometrica, 49(3), 555–574.
Liu, Y., & Tsyvinski, A. (2018). Risks and returns of cryptocurrency. Review of Financial Studies, 34(6), 2689–2727.
Long, H., Zhang, S., & Zhao, R. (2022). Social media and stock returns: Evidence from meme stocks. Journal of Banking & Finance, 137, 106472.
Lo, A., & Ross, D. (2024). Hallucination risks in generative AI models. Journal of Financial Data Science, 6(1), 45–63.
Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. Journal of Finance, 66(1), 35–65.
Lucchini, L., et al. (2021). Retail predatory trading: Evidence from GameStop. [Working paper].

Mao, H., Counts, S., & Bollen, J. (2011). Predicting financial markets: Comparing survey, news, Twitter, and search engine data. arXiv preprint.
Nag, A. K., & Mitra, A. (1999). Neural networks and early warning indicators of currency crisis. Reserve Bank of India Occasional Papers, 20(2), 183–222.
Nani, F. (2022). Dogecoin: A case study in speculative narratives. Journal of Behavioral Finance, 23(3), 345–359.
Nyman, R., Kapadia, S., & Tuckett, D. (2021). News and narratives in financial systems: Exploiting big data for systemic risk assessment. Journal of Economic Dynamics and Control, 127, 104119.
Odean, T. (1998). Are investors reluctant to realize their losses? Journal of Finance, 53(5), 1775–1798.
Odean, T. (1999). Do investors trade too much? American Economic Review, 89(5), 1279–1298.
Ofek, E., & Richardson, M. (2001). Dotcom mania: The rise and fall of internet stock prices. Journal of Finance, 58(3), 1113–1137.
Oncu, T. S. (2021). Dogecoin's rise: A bubble or a joke? Finance Research Letters, 43, 101977.
Papadamou, S., Kyriazis, N. A., & Tzeremes, N. G. (2021). Herding behavior and cryptocurrency market dynamics. Research in International Business and Finance, 58, 101492.
Papadamou, S., et al. (2019). Herding behavior and asset prices. Economic Modelling, 82, 44–53.
Pástor, Ľ., & Veronesi, P. (2005). Rational IPO waves. Journal of Finance, 60(4), 1713–1757.
Pástor, Ľ., & Veronesi, P. (2009). Technological revolutions and stock prices. American Economic Review, 99(4), 1451–1483.
Phillips, P. C. B., Shi, S., & Yu, J. (2015). Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. International Economic Review, 56(4), 1043–1078.
Pompian, M. M. (2017). Behavioral finance and investor types: Managing behavior to make better investment decisions. Wiley.
Puhr, H., & Müllner, J. (2021). Let me Google that for you: Capturing globalization using Google Trends. SSRN Working Paper.

Purnell, D., Jr., Etemadi, A., & Kamp, J. (2024). Developing an early warning system for financial networks: An explainable machine learning approach. Entropy, 26(9), 796.
PwC. (2021). Time for trust: How blockchain will transform business and the economy. PwC Global Report.
Roos, M., & Reccius, C. (2021). Narrative mining in financial markets. Journal of Economic Behavior & Organization, 188, 954–972.

Romero, D. M., Galuba, W., Asur, S., & Huberman, B. A. (2011). Influence and passivity in social media. In Proceedings of the 20th International Conference on World Wide Web (pp. 113–122). ACM.
Semenova, V., & Winkler, J. (2021). Social contagion and asset prices: Reddit's self-organised bull runs. INET Oxford Working Paper No. 2021-04.
Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? American Economic Review, 71(3), 421–436.
Shiller, R. J. (2017). Narrative economics. Princeton University Press.
Shiller, R. J. (2019). Narrative economics: How stories go viral and drive major economic events. Princeton University Press.
Shleifer, A., & Vishny, R. W. (1997). The limits of arbitrage. Journal of Finance, 52(1), 35–55.
Smales, L. A. (2021). Investor attention and financial markets. International Review of Financial Analysis, 76, 101784.
Soros, G. (1987). The alchemy of finance: Reading the mind of the market. Simon & Schuster.
Sunstein, C. R. (2001). Republic.com. Princeton University Press.
Taffler, R. J., Agarwal, V., & Obring, A. (2021). Narrative and sentiment analysis in financial markets. Journal of Behavioral Finance. [Et note 2024 connexe si applicable].
Tsintzou, V., Pitoura, E., & Tsaparas, P. (2018). Bias disparity in recommendation systems. arXiv preprint.
Vidal-Tomás, D., Ibáñez, A. M., & Farinós, J. E. (2019). Herding in the cryptocurrency market: CSSD and CSAD approaches. Finance Research Letters, 30, 181–186.
Wheatley, S., Sornette, D., Huber, T., Reppen, M., & Gantner, R. N. (2018). Are Bitcoin bubbles predictable? Combining a generalized Metcalfe's law and the LPPLS model. Swiss Finance Institute Research Paper 18-22.

Yan, W., & Sornette, D. (2011). Effect of positive feedbacks on external shocks in social systems. Physical Review E, 84(4), 046107.

APPENDICES

Survey :

Survey Instrument

<u>Section 1 of 7: Investing in the Digital Era: Psychology, Social Media, and Artificial Intelligence</u>

This questionnaire is part of a doctoral research project on the formation of speculative bubbles in the digital age.
Its objective is to analyze the role of behavioral biases, social media, digital marketing, as well as algorithmic technologies and artificial intelligence (AI) in current or future investment decisions.

- *Duration: approximately 5–7 minutes*
- *All responses are anonymous and confidential*
- *There are no right or wrong answers — only your perceptions matter*

Thank you for your participation!

<u>Section 2 of 7: Investment Profile and Intentions</u> - *Your current habits or future plans regarding investment*

1. Current investment status:

- *I have already invested*
- *I have not yet invested but intend to do so*
- *I have not invested and do not intend to do so*

2. Types of assets invested in or considered (multiple answers possible):

- *Stocks*
- *Cryptocurrencies*
- *NFT(s)*
- *ETFs / Index funds*
- *Bonds*
- *Options / Derivatives*
- *Other (please specify)*

3. Which type of investor best describes you (or would best describe you if you invested)?

- *Very cautious*
- *Cautious*
- *Moderate*
- *Dynamic*
- *Very aggressive*

<u>Section 3 of 7: Psychological Biases</u> - *How you perceive gains, losses, and the influence of other investors*

4. I believe I can (or could) beat the market through my own decisions.

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

5. A loss of €100 would hurt me more than a gain of €100 would make me happy.

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

6. I am (or would be) influenced by what the majority of other investors do.

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

7. Narratives and stories matter more than fundamentals in explaining asset prices.
(By "narratives and stories," we mean discourses, slogans, or collective explanations circulating in media and on social networks, e.g. "this stock will revolutionize its sector," "this crypto is the future of finance").
(By "fundamentals," we mean actual economic data such as company earnings, growth, or book value).

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

***Section 4 of 7: Influence of Digital Platforms* -** *How social media, influencers, and online content can affect your decisions*

8. I regularly see financial content on social media.

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

9. Viral trends (memes, hashtags, buzz) can influence my investment decisions.

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

10. I give more weight to financial influencers' advice than to traditional analysts.

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

11. Which platforms do you mainly use to stay informed about finance or investing? (multiple answers possible)

- *Reddit*
- *Twitter / X*
- *TikTok*
- *YouTube*
- *Discord*
- *Instagram*
- *None*
- *Other (please specify)*

***Section 5 of 7: Algorithms and Artificial Intelligence* -** *Your perception of the role of algorithms and AI in financial information and investment decisions*

12. Algorithms (TikTok, YouTube, Twitter, etc.) strongly influence the financial content I am exposed to.

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

13. Personalized news feeds reinforce my investment opinions (or would reinforce them if I invested).

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

14. I would have as much or more trust in an AI-based investment tool as in a human advisor.

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

15. Platforms should be responsible for limiting the spread of misleading or excessively risky financial content.

- *Strongly disagree / Disagree / Neutral / Agree / Strongly agree*

<u>Section 6 of 7: Speculative Behaviors and Scenarios</u> - *Your reactions to market situations marked by "buzz" or rapid price increases*

16. Have you ever invested (or would you invest) primarily due to online hype?

- *Yes, I have already done so*
- *Not yet, but I could do so*
- *No, never*

17. If yes (or if you would be willing to), what would be your main reason?

- *Community enthusiasm / fear of missing out (FOMO)*
- *Possibility of quick profit*
- *Belief in the long-term narrative*
- *Direct influence of an influencer*
- *Other (please specify)*

18. Scenario: a stock or cryptocurrency rises +200% in one week and goes viral. What would be your most likely reaction?

- *Buy quickly*
- *Wait and check fundamentals*
- *Not invest (bubble risk)*
- *Follow the advice of influencers*

<u>Section 7 of 7 — General Information</u> - *Some basic questions to better understand your profile*

19. Age:

- *Under 18*
- *18–24*
- *25–34*
- *35–44*
- *45–54*
- *55 and above*

20. Gender:

- *Male*
- *Female*
- *Prefer not to answer*
- *Other (please specify)*

21. Highest level of education completed or currently in progress:

- *No diploma / Middle school*
- *High school diploma (Baccalaureate or equivalent)*
- *1–2 years post-secondary (Associate degree, Licence 1–2, DUT, BTS)*
- *3 years post-secondary (Bachelor's degree)*
- *4–5 years post-secondary (Master's degree, Grandes Écoles)*
- *Doctorate / PhD*
- *Other (please specify)*

Link to survey answer https://docs.google.com/spreadsheets/d/1sbv6hzBokLT-ylV6Ybvf5liKZ7Re-DnHKqfwobE5q7c/edit?usp=sharing

Screenshots of the charts created with Python, representing the data collected from the form above

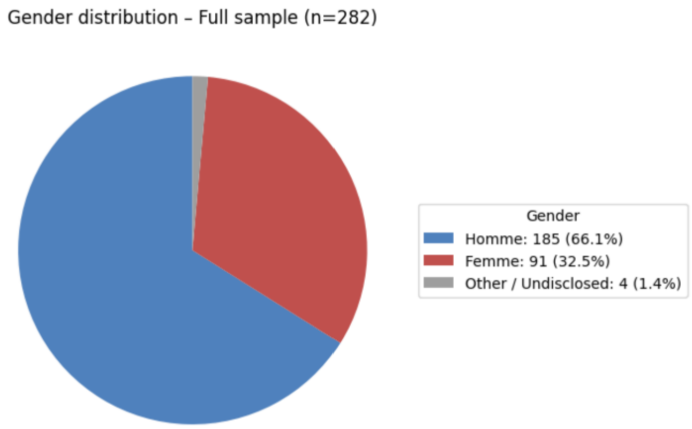

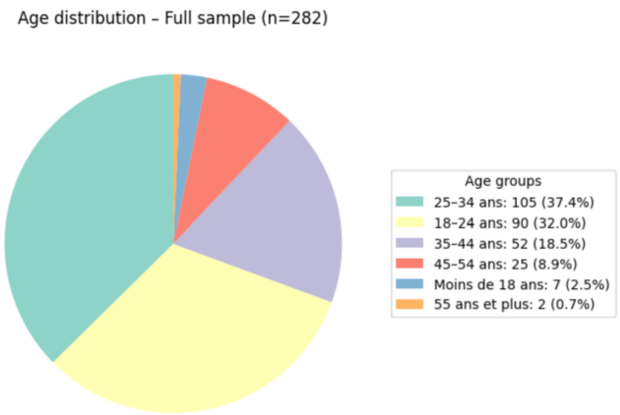

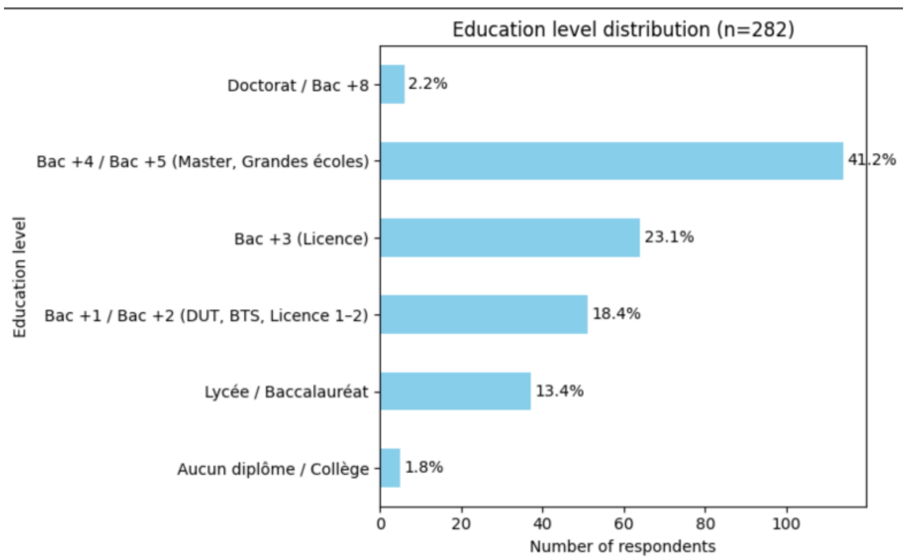

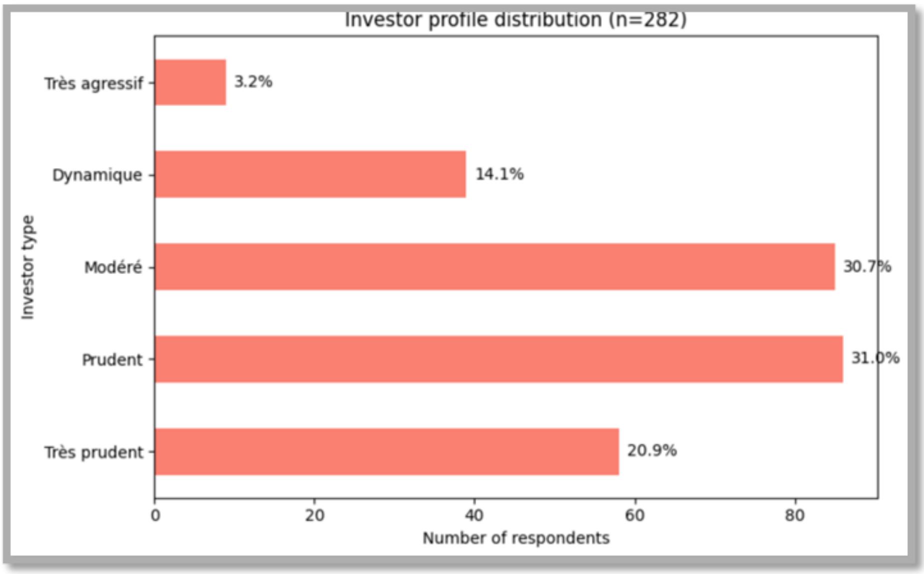
Investor profile distribution (n=282)
Très agressif
Dynamique
Modéré
Prudent
Très prudent
3.2%
14.1%
30.7%
31.0%
20.9%
Investor type
Number of respondents


Current investment status (n=282)
10.7%
49.5%
39.9%
Investment status
J'ai déjà investi: 139
Je n'ai pas encore investi mais j'ai l'intention de le faire: 112
Je n'ai pas investi et je n'ai pas l'intention de le faire: 30

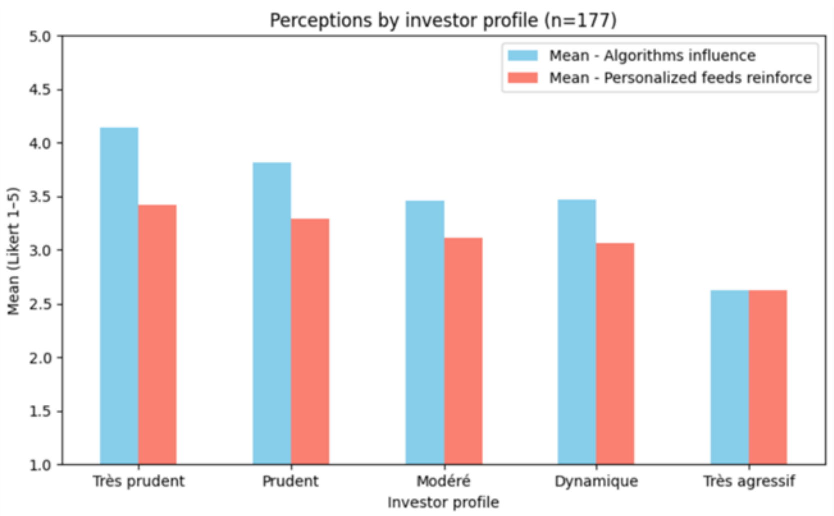
Perceptions by investor profile (n=177)
Mean - Algorithms influence
Mean - Personalized feeds reinforce
Mean (Likert 1–5)
Très prudent
Prudent
Modéré
Dynamique
Très agressif
Investor profile

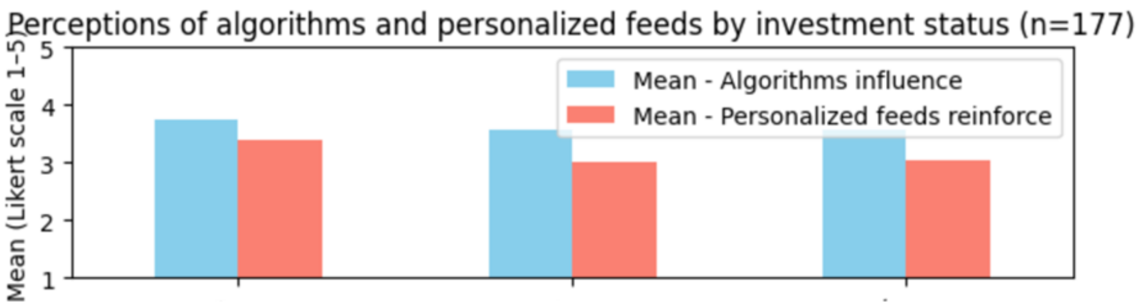
Perceptions of algorithms and personalized feeds by investment status (n=177)
Mean - Algorithms influence
Mean - Personalized feeds reinforce
Mean (Likert scale 1–5)

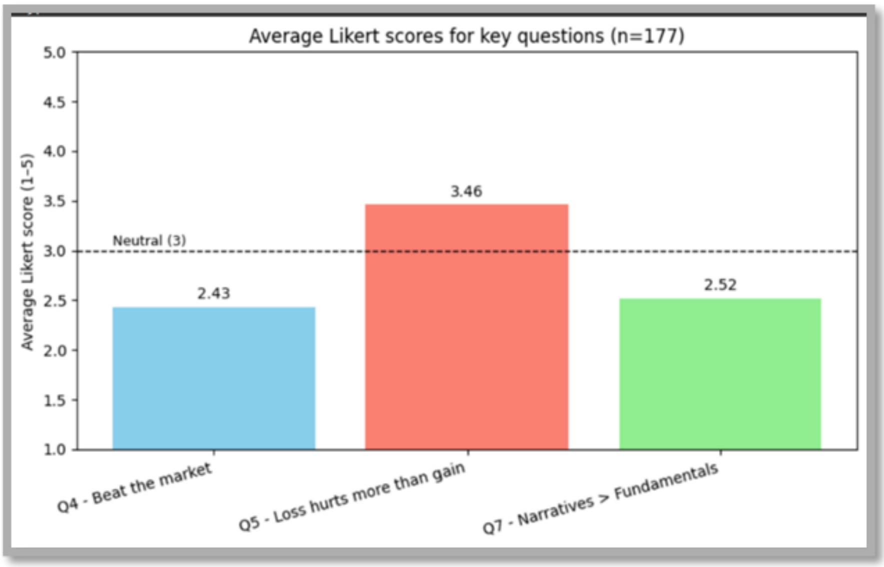
Average Likert scores for key questions (n=177)
Average Likert score (1–5)
5.0
4.5
4.0
3.5
3.0
2.5
2.0
1.5
1.0
Neutral (3)
2.43
3.46
2.52
Q4 - Beat the market
Q5 - Loss hurts more than gain
Q7 - Narratives > Fundamentals

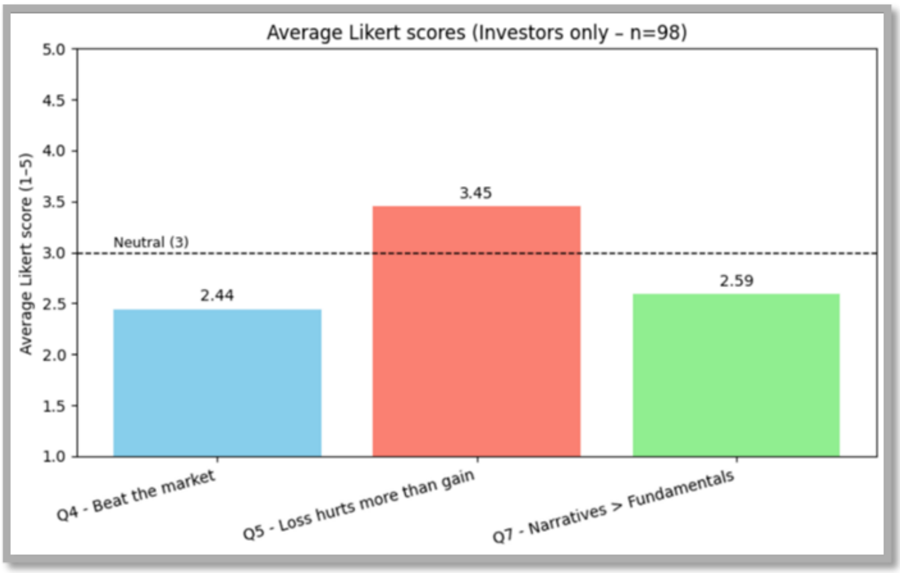
Average Likert scores (Investors only – n=98)
Average Likert score (1–5)
5.0
4.5
4.0
3.5
3.0
2.5
2.0
1.5
1.0
Neutral (3)
2.44
3.45
2.59
Q4 - Beat the market
Q5 - Loss hurts more than gain
Q7 - Narratives > Fundamentals

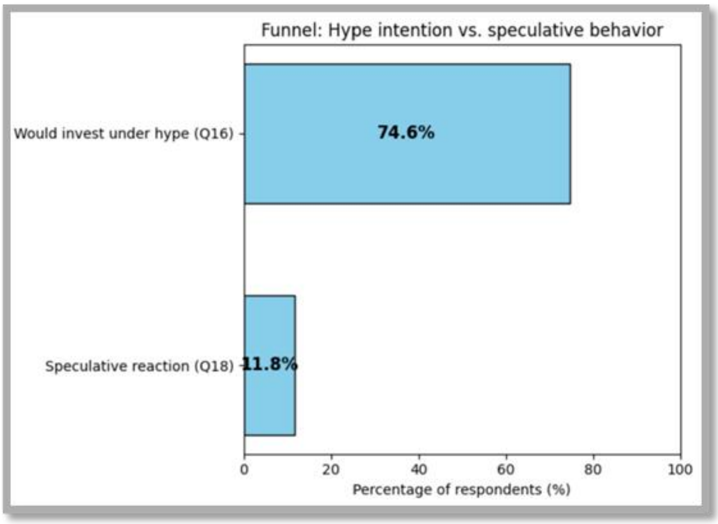
Funnel: Hype intention vs. speculative behavior
Would invest under hype (Q16)
74.6%
Speculative reaction (Q18)
11.8%
0
20
40
60
80
100
Percentage of respondents (%)

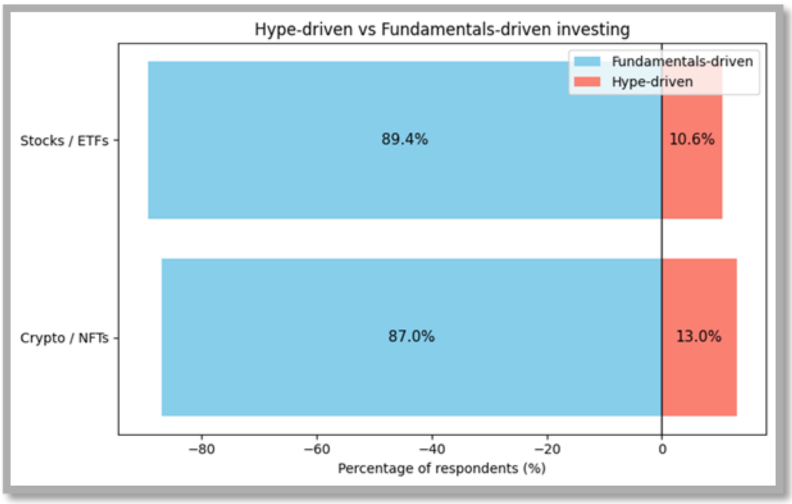
Hype-driven vs Fundamentals-driven investing
Fundamentals-driven
Hype-driven
Stocks / ETFs
89.4%
10.6%
Crypto / NFTs
87.0%
13.0%
−80
−60
−40
−20
0
Percentage of respondents (%)


Cluster profiles (normalized to 0–1)
Viral trends (Q9)
Influencers > analysts (Q10)
Hype intention (Q16)
Speculative choice (Q18)
0.25
0.50
0.75
1.00

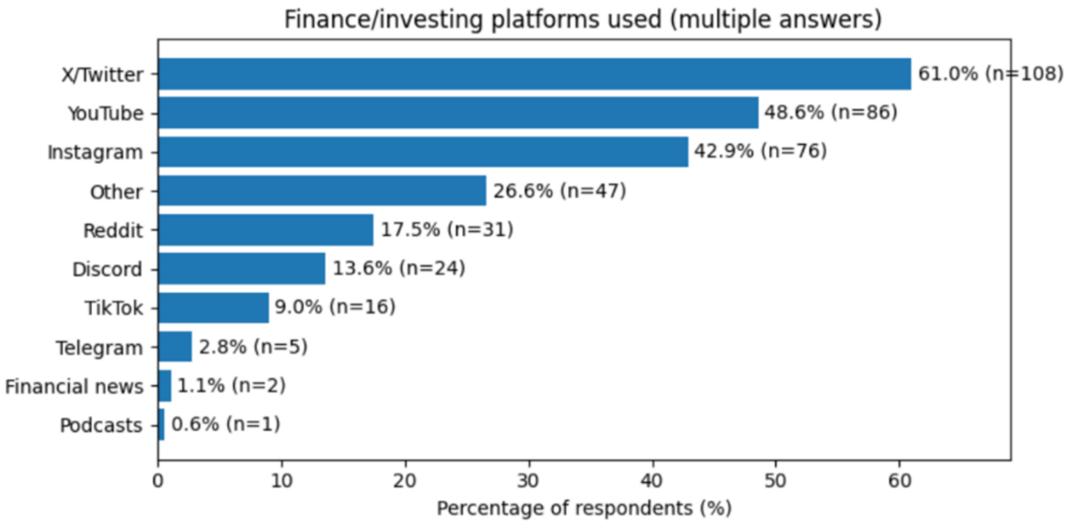
Finance/investing platforms used (multiple answers)
X/Twitter
61.0% (n=108)
YouTube
48.6% (n=86)
Instagram
42.9% (n=76)
Other
26.6% (n=47)
Reddit
17.5% (n=31)
Discord
13.6% (n=24)
TikTok
9.0% (n=16)
Telegram
2.8% (n=5)
Financial news
1.1% (n=2)
Podcasts
0.6% (n=1)
0
10
20
30
40
50
60
Percentage of respondents (%)

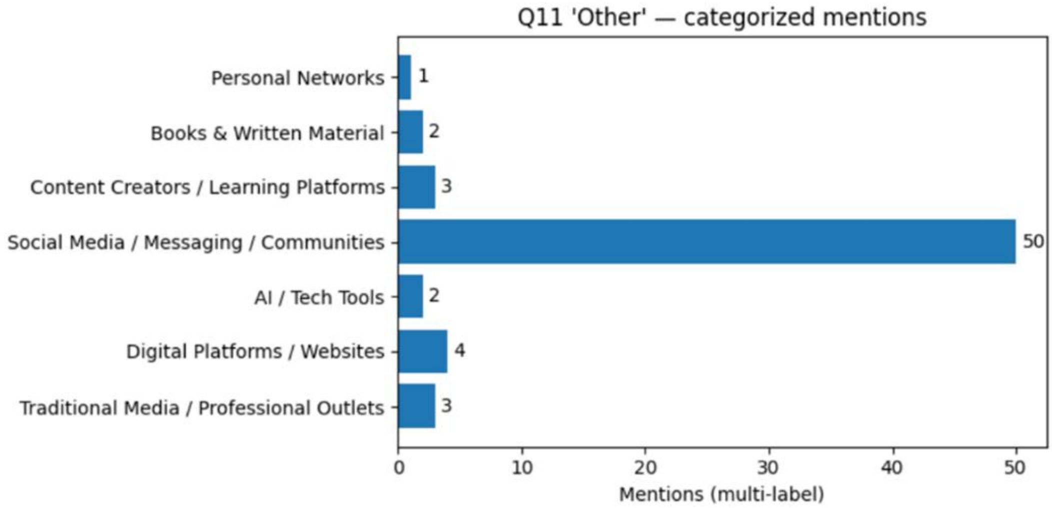
Q11 'Other' — categorized mentions
Personal Networks
Books & Written Material
Content Creators / Learning Platforms
Social Media / Messaging / Communities
AI / Tech Tools
Digital Platforms / Websites
Traditional Media / Professional Outlets
1
2
3
50
2
4
3
0
10
20
30
40
50
Mentions (multi-label)

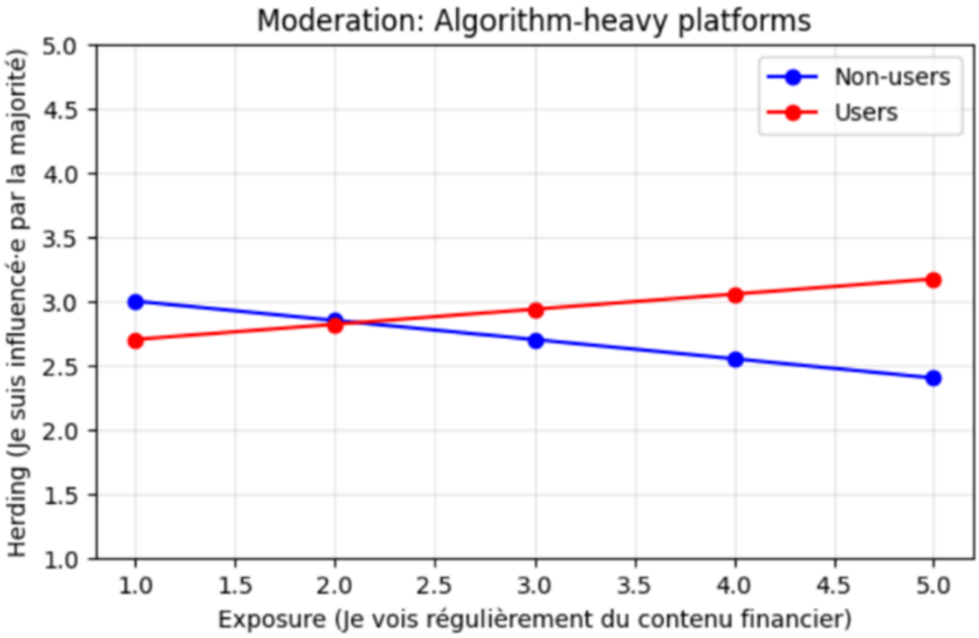
Moderation: Algorithm-heavy platforms
Non-users
Users
Herding (Je suis influencé·e par la majorité)
Exposure (Je vois régulièrement du contenu financier)
1.0
1.5
2.0
2.5
3.0
3.5
4.0
4.5
5.0

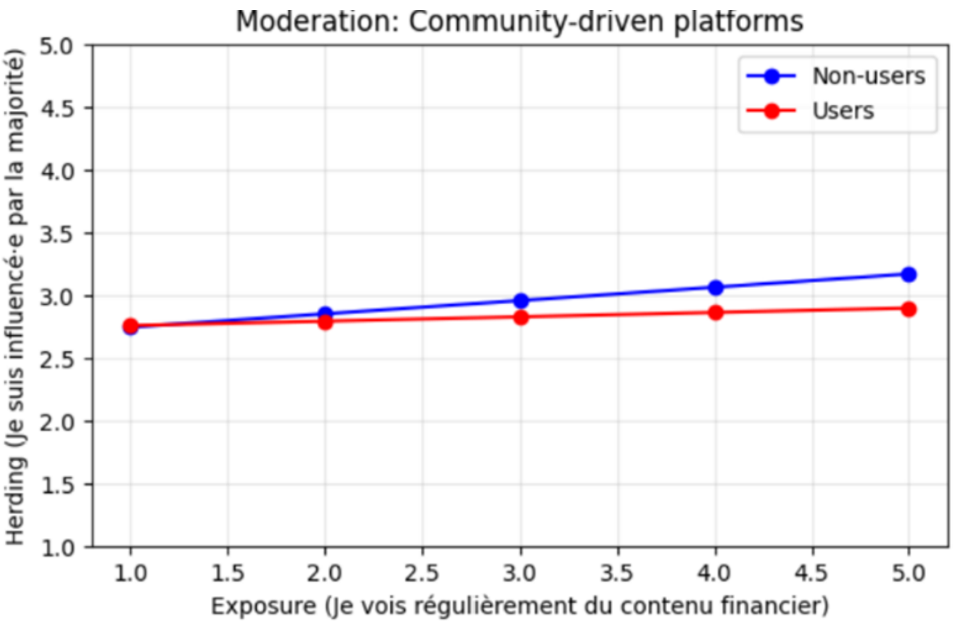
Moderation: Community-driven platforms
Non-users
Users
Herding (Je suis influencé·e par la majorité)
Exposure (Je vois régulièrement du contenu financier)
1.0
1.5
2.0
2.5
3.0
3.5
4.0
4.5
5.0

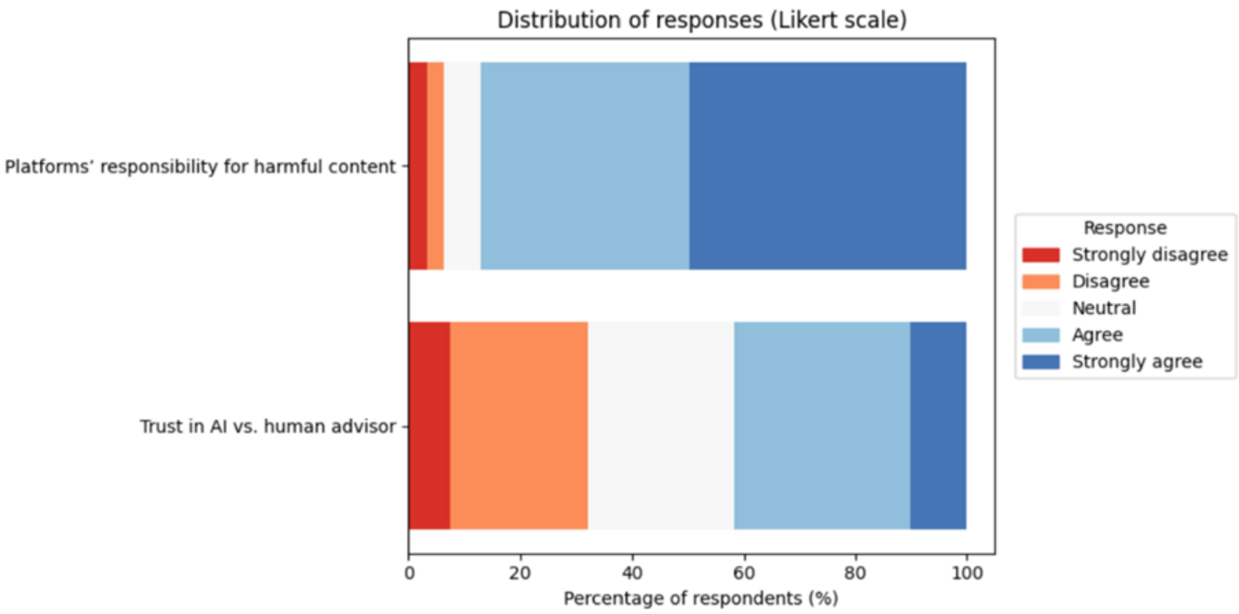
Distribution of responses (Likert scale)
Platforms’ responsibility for harmful content
Trust in AI vs. human advisor
Response
Strongly disagree
Disagree
Neutral
Agree
Strongly agree
0
20
40
60
80
100
Percentage of respondents (%)

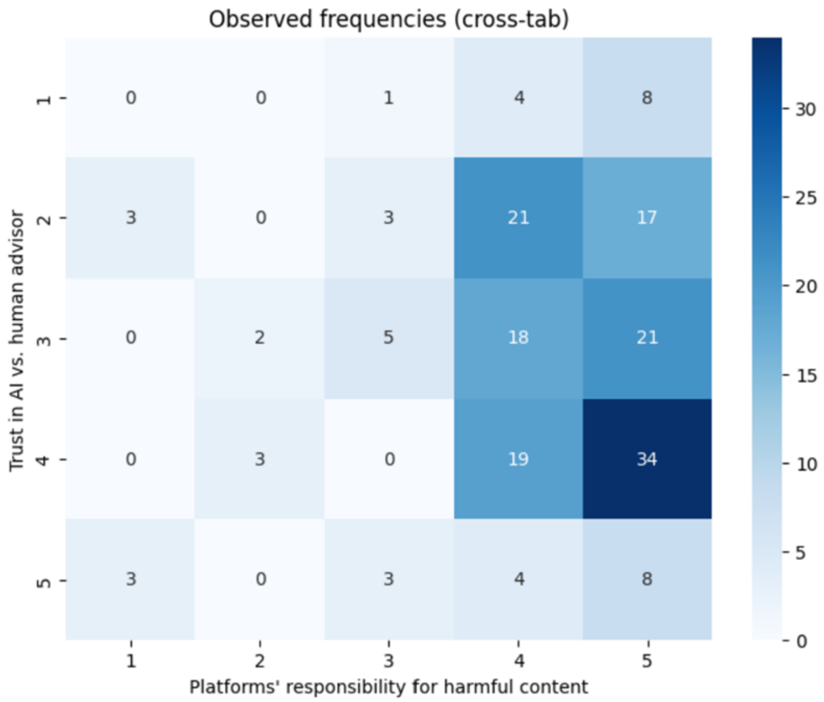
Observed frequencies (cross-tab)
Trust in AI vs. human advisor
Platforms' responsibility for harmful content
1
2
3
4
5
0 0 1 4 8
3 0 3 21 17
0 2 5 18 21
0 3 0 19 34
3 0 3 4 8
0
5
10
15
20
25
30