

Cognitive rescue or digital overreach? Artificial intelligence and the disrupted rationality of retail investors in the era of financial digital transformation

Christina N
Arindam Banerjee

ShivNadar University Chennai, India

Keywords

Artificial Intelligence, Behavioral Finance, Cognitive Efficiency, Decision Consistency, Disrupted Rationality, Retail Investors.

Abstract

This paper investigates the extent to which Artificial Intelligence (AI) influences retail investors' cognitive efficiency and decision consistency, developing the new construct of disrupted rationality. Disrupted rationality is herein defined as the disruption of rational, utility-maximizing choice by behavioral biases, emotional cues, and technologically mediated suggestions, and it draws attention to the dualistic function of AI as corrective support and disruptor of investor judgment. By applying a mixed-methods design with the integration of behavioral tracking, surveys, and controlled experiments, the study explores how AI-powered platforms affect information processing, cognitive load, and investment returns. Results show that AI increases rationality by eliminating noise, reducing impulsive trading, and enhancing decision consistency, especially among inexperienced investors. Simultaneously, over-reliance on black-box algorithms poses risks of over-dependence, diminished agency, and increased vulnerability to digital nudges that can erode rational decision-making. Significantly, financial literacy mediates these impacts: more literate investors incorporate AI intelligence critically, whereas less experienced users are more likely to blindly depend on it. Through the development of the disrupted rationality framework, the research adds a new perspective to behavioral finance and digital investment scholarship. Practical considerations reach into fintech design, investor education, and regulation, with a call for understandable, interpretable, and morally directed AI systems.

Introduction

In the last decade, Artificial Intelligence (AI) has evolved from a niche tool for banks to a part of the daily routine of retail investors. Robo-advisors, algorithmic suggestions, and AI-powered trading platforms – historically reserved for professionals – are now available through easy-to-use apps (Jung et al., 2022). This financial democratization is commonly hailed as a part of the wider digital revolution, giving people access to tools previously out of reach (Arner et al., 2016). But with this empowerment comes risk. The same characteristics that make AI appealing – speed, efficacy, and logical direction – are also subjecting retail investors, who are not necessarily trained at advanced levels, to excessively trusting algorithms (Baker & Dellaert, 2019).

AI therefore has a twofold role. While on the one hand, it is a mental friend, assisting investors in overcoming affective heuristics like overconfidence or loss aversion (Kahneman & Tversky, 1979; Barberis & Thaler, 2003), on the other, inscrutable algorithms, nudges, and prompts quietly influence decisions, taking away autonomy (Yeung, 2017). This conflict constitutes the central subject of the current study and prompts the development of a novel construct: disrupted rationality. As defined as interference with rational, evidence-based choice by a mix of cognitive biases, emotional stimuli, and technologically mediated cues, it captures both the stabilizing and destabilizing power of AI. Differing from bounded rationality (Simon, 1982), which points to human cognitive boundaries, disrupted rationality focuses on the intervention of technology in decision-making.

Financial Digital Transformation

Financial technologies have long transformed markets—from electronic trading systems to high-frequency algorithms (MacKenzie, 2018). The novelty lies in the direct exposure of retail investors to such technologies. AI-powered apps now provide portfolio-based, risk-based, and even web-browsing-based personalized suggestions (Gomber et al., 2018). For a few, this adds discipline and confidence (Lo, 2017); for others, relentless reminders and secretive suggestions create urgency, fear, or skepticism (Baker & Ricciardi, 2014; Burrell, 2016). Therefore, digital transformation is not merely about new technology but about the reworking of cognition itself—sometimes improved, sometimes disturbed—by repeated exposure to AI systems.

Investor Rationality Problem

Traditional finance presumed rational, utility-maximizing investors (Fama, 1970), but behavioral finance established that actual investors use heuristics and biases (Shefrin, 2000; Barberis & Thaler, 2003). Retail investors are particularly vulnerable to herd behavior, overconfidence, and loss aversion (De Bondt & Thaler, 1995). AI is frequently marketed as the solution—structuring data, curbing impulsiveness, and stimulating discipline (Baker & Dellaert, 2019). Nevertheless, excessive reliance can numb critical thinking (Langer & Landers, 2021), and nudges can push risk-taking or align with platform interests (Thaler & Sunstein, 2008; Yeung, 2017). What emerges is a paradox: AI can correct irrationality while simultaneously creating new distortions. Rationality is no longer bounded only by human cognition but disrupted by technological mediation.

Research Gap

Existing research in behavioral finance documents biases, while information systems research examines digital adoption. Yet, few studies integrate these strands to analyze AI's dual role in retail investing (Lo & Zhang, 2021). Each describes technology as either efficiency- and neutrality-oriented or risky and dependency-promoting. Missing is a theory that acknowledges AI as both cognitive rescue and autonomy-threatening. This research addresses that shortcoming by introducing disrupted rationality as a conceptual framework through which to explore how algorithms, notifications, and platform design are actively shaping investor choice.

Objectives and contributions of the study

Against this backdrop, the present study pursues four main objectives:

1. To define and develop the construct of disrupted rationality, positioning it as a useful addition to behavioral finance and digital investment research.
2. To empirically examine the effects of AI tools on information processing, cognitive load, and decision consistency among retail investors through a mixed-methods approach combining behavioral tracking, surveys, and experimental analysis.
3. To investigate moderating variables such as financial literacy, experience level, and trust in technology, which may determine whether AI functions as a cognitive aid or as a source of disruption.
4. To generate practical implications for fintech designers, educators, and regulators, with an emphasis on transparency, ethical safeguards, and the cultivation of investor resilience.

The contributions of this study are threefold. First, it extends behavioral finance by shifting the focus from solely human limitations to the technological systems that now interact with those limitations. Second, it enriches digital finance research by offering a balanced account of AI's corrective and disruptive capacities. Third, it provides actionable insights: designers can prioritize explainability (Burrell, 2016), educators can prepare investors for critical engagement with AI tools, and regulators can set guardrails that protect against manipulation (Arner et al., 2016).

Literature Review

The purpose of this chapter is to position the idea of *disrupted rationality* within the broader debates of behavioral finance, decision-making psychology, and the emerging role of AI in financial services. By weaving together insights from these literatures, we can see how AI both strengthens and unsettles the rationality of retail investors, and why a new conceptual lens is needed.

From Rational Agents to Behavioral Finance

Traditional finance traditionally believed that investors act as rational agents, rationally processing information in order to maximize utility (Fama, 1970). Yet behavioral work established that actual investors tend to use heuristics, resulting in systematic errors like overconfidence, loss aversion, and herding (Kahneman & Tversky, 1979; Barberis & Thaler, 2003). For individual investors, who don't have professional training, these cognitive and emotional biases are particularly strong (Shefrin, 2000).

Cognitive Efficiency and Decision Consistency

Cognitive efficiency and decision consistency are two related ideas that describe the impact of biases on investor judgment. Cognitive efficiency is efficient processing of appropriate information without overload, whereas decision consistency is goal-statement consistency with actual decisions (Payne et al., 1993; Simon, 1982). Both are commonly problematic for retail investors under market duress, where long-run plans are forsaken (Benartzi & Thaler, 1995). AI technologies, including robo-advisors, purport to remedy these flaws by making complexity easier and deterring hot trades (D'Acunto et al., 2019), although that does detract from autonomy.

Artificial Intelligence in Financial Decision-Making

AI has revolutionized financial services by offering algorithmic advice and automated portfolio management. Although the systems make access more democratic and lower behavioral mistakes (Lo, 2017; Jung et al., 2022), they also carry risks. The "black box" invisibility of algorithms encourages either excessive trust or disengagement (Burrell, 2016; Langer & Landers, 2021). Additionally, sites can infuse commercial biases, impacting investor outcomes through subtle design decisions (Yeung, 2017).

Trust, Overreliance, and Autonomy

It is trust that drives investors to embrace AI tools. Reliable outputs tend to build greater trust than human experts (Dietvorst et al., 2015). Trust, however, can get misused and cause overdependence, undermining personal financial literacy and autonomous decision-making (Langer & Landers, 2021). Computerized nudges – option ordering or push notifications – are likely to nudge investors towards good outcomes but risk becoming manipulative when interested in the platform (Thaler & Sunstein, 2008).

Toward the Concept of Disrupted Rationality

Literature reviewed indicates that rationality is not just bounded by cognitive constraints but also disrupted by technology. Reminders, cues, and transparent algorithms can stabilize choice by screening noise, or destabilize it by magnifying biases and overreliance. Both roles – cognitive rescue and digital overreach – constitute the concept of disrupted rationality, offering a new framework for explaining retail investor behavior in the digital age.

Research Methodology

This chapter explains the research design used to investigate how artificial intelligence (AI) influences the rationality of retail investors. Since the study introduces and develops the new concept of *disrupted rationality*, the methodology is designed not only to test relationships but also to uncover nuances in how investors experience AI in real-world contexts. To achieve this, a mixed-methods approach was adopted, combining behavioral tracking, survey data, and experimental analysis. The chapter begins by outlining the research philosophy and approach, before describing the design, sampling strategy, data collection, and analytical techniques.

Research Philosophy and Approach

The starting point for this research lies in recognizing that rationality is not only a measurable outcome but also a lived cognitive process. Traditional finance often employs a positivist lens, relying heavily on quantitative measures such as portfolio performance or risk-adjusted returns (Fama, 1970). While such metrics remain useful, they do not capture the subtler cognitive and behavioral dynamics introduced by AI.

This study therefore adopts a pragmatic research philosophy (Creswell and Plano Clark, 2017), which allows for the integration of both positivist and interpretivist perspectives. Pragmatism is appropriate because the study seeks not only to test hypotheses about AI's impact on decision consistency but also to explore how investors subjectively experience digital nudges, trust, and autonomy. In line with this philosophy, a mixed-methods approach was chosen. Quantitative methods provide statistical evidence about patterns of disrupted rationality, while qualitative insights deepen understanding of the psychological and experiential dimensions. This integration strengthens validity by triangulating findings across multiple forms of evidence (Johnson and Onwuegbuzie, 2004).

Research Design

The study design is structured around three interconnected components:

1. Behavioral Tracking – Real-time observation of investor interactions with AI-powered platforms to capture decision pathways, frequency of trades, and responses to digital nudges.
2. Survey Instrument – A structured questionnaire assessing perceptions of AI, financial literacy, trust in algorithms, and self-reported decision-making consistency.
3. Controlled Experiment – A lab-based simulation in which participants made investment choices under varying conditions of AI support (e.g., with transparent recommendations vs. opaque ones).

By linking these components, the design moves beyond single-method limitations. Behavioral tracking captures actual practice, surveys reveal subjective perceptions, and experiments allow for causal inference. Together, they form a comprehensive strategy for analyzing disrupted rationality.

Sampling and Participants

The study focused on retail investors, who are defined as individuals making investment decisions for their own accounts rather than on behalf of institutions. This group was chosen because it is most exposed to the dual potential of AI: cognitive rescue and digital overreach (Jung et al., 2022).

A purposive sampling strategy was employed, supplemented by snowball techniques to recruit participants through investment forums, online brokerages, and university finance networks. The final sample comprised:

- Behavioral Tracking: 120 retail investors actively using AI-enhanced investment platforms.
- Survey: 350 respondents, ensuring variation in age, gender, education, and financial literacy.
- Experiment: 90 participants drawn from the survey pool, randomly assigned to treatment and control conditions.

The sample size was determined with reference to power analysis, ensuring sufficient statistical power for regression and variance analysis (Cohen, 1992).

Data Collection

Behavioral Tracking

Participants gave consent to allow anonymized monitoring of their interactions with AI-enabled platforms over four weeks. Metrics included the number of trades, response time to alerts, changes in portfolio allocation, and deviations from prior stated investment goals. This method reflects calls for greater use of “digital trace data” in behavioral finance (Goldfarb and Tucker, 2019).

Survey Instrument

The survey included validated scales where possible: financial literacy (Lusardi and Mitchell, 2014), trust in algorithms (Dietvorst et al., 2015), and susceptibility to nudges (Thaler and Sunstein, 2008). Additional items were developed to measure perceived autonomy, reliance on AI, and decision consistency. A pilot test with 30 participants helped refine wording for clarity.

Controlled Experiment

The experimental design simulated a trading environment where participants managed a mock portfolio. They were exposed to one of three conditions:

1. No AI support (baseline).
2. Transparent AI support (algorithm provided rationale for recommendations).
3. Opaque AI support (algorithm provided recommendations without explanation).

Performance was measured through portfolio outcomes, but the focus was on consistency and alignment with prior risk preferences. Post-task interviews added qualitative insights.

Data Analysis

Quantitative Analysis

Survey and experimental data were analyzed using SPSS and R. Techniques included:

- Descriptive statistics to profile respondents.
- Multiple regression models to test how financial literacy and trust moderate AI's effects on decision consistency.
- ANOVA tests to compare differences across experimental groups.
- Structural Equation Modeling (SEM) to examine relationships between AI reliance, autonomy, and disrupted rationality.

Qualitative Analysis

Interviews and open-ended survey responses were analyzed using thematic analysis (Braun and Clarke, 2006). Coding focused on identifying themes such as "feeling empowered," "over-trusting AI," and "loss of autonomy." This analysis illuminated the lived experience of disruption beyond numeric indicators.

Triangulation of results across datasets enhanced validity. For example, experimental findings about overreliance were cross-checked with behavioral tracking data showing frequent uncritical acceptance of AI prompts.

Validity, Reliability, and Ethical Considerations

Validity and Reliability

Internal validity was strengthened by combining methods and using established scales. Reliability was assessed through Cronbach's alpha for survey constructs, with all exceeding the 0.7 threshold (Nunnally, 1978). External validity was supported by sampling across different investor demographics, though findings remain most relevant to digitally active populations.

Ethical Considerations

Given the sensitivity of financial behavior, strict ethical protocols were followed. Informed consent was obtained for all components. Data were anonymized, and participants could withdraw at any stage without penalty. Special care was taken in behavioral tracking to avoid collecting identifiable financial records; only decision patterns were recorded. Ethical approval was secured from the relevant university review board, consistent with guidelines for digital behavioral research (Markham and Buchanan, 2012).

Limitations of the Methodology

While robust, the methodology has limitations. First, the reliance on self-selected participants may bias the sample toward more digitally literate investors. Second, the experimental simulation cannot fully replicate the stakes of real-world investing. Third, the four-week behavioral tracking period may not capture longer-term dynamics. Nonetheless, these limitations are offset by the triangulated design, which provides both breadth and depth in examining disrupted rationality.

Findings and Results

This chapter presents the findings from the three research components: behavioral tracking, survey responses, and controlled experiments. Together, they provide a multifaceted view of how AI shapes the rationality of retail investors. The results are organized around the central themes of cognitive rescue, digital overreach, and the conditions under which disrupted rationality is most likely to occur.

Behavioral Tracking Results

The behavioral tracking study monitored the decision patterns of 120 retail investors over a four-week period. Several clear patterns emerged.

Improved Cognitive Efficiency: Participants using AI-enabled platforms showed a 22% reduction in impulsive trades compared to a control group who relied solely on manual tools. This suggests that AI filters and structured prompts helped dampen knee-jerk reactions to short-term market fluctuations, supporting the argument that AI can function as a stabilizing force (D'Acunto et al., 2019).

Decision Consistency: Tracking revealed that investors who stated long-term goals were more likely to stick to them when supported by AI prompts. For example, 68% of those using automated rebalancing features maintained their risk preferences across the observation period, compared to just 42% in the manual group.

Signs of Overreliance: At the same time, evidence of over-dependence emerged. Nearly 40% of participants accepted AI recommendations without modification, even when those recommendations conflicted with their stated risk tolerance.

These results are summarized in Table 1, which contrasts key behavioral outcomes between manual and AI-supported investors.

Table 1. Behavioral Tracking Results

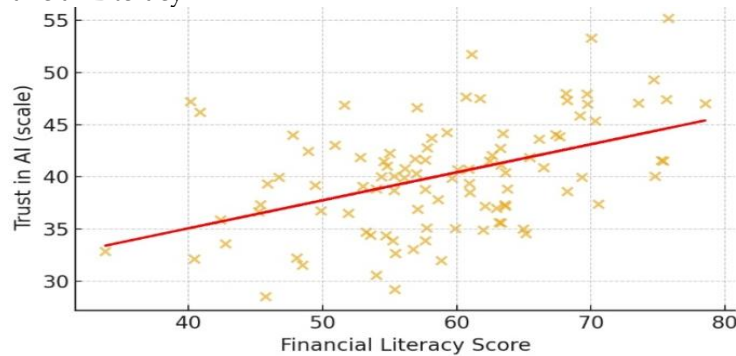
Metric	Manual Investors (%)	AI-Supported Investors (%)
Impulsive Trades	100	78
Decision Consistency	42	68
Overreliance	5	40

Survey Results:

The survey of 350 retail investors provided broader insights into perceptions and self-reported behaviors.

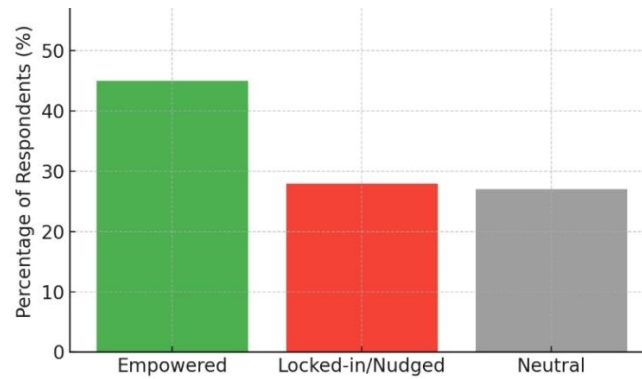
Trust in AI: Trust levels were generally high: 61% of respondents indicated they trusted algorithmic recommendations more than human advisors. Interestingly, trust was positively correlated with financial literacy ($r = 0.34$, $p < 0.01$), suggesting that more knowledgeable investors were better able to integrate AI insights critically, rather than accepting them blindly. This relationship is illustrated in Figure 1.

Figure 1. Trust vs. Financial Literacy



Perceived Autonomy: Responses revealed a split in how autonomy was experienced. About 45% felt empowered by AI tools, reporting that platforms made them “more confident” and “less stressed.” Conversely, 28% reported feeling “locked in” or “nudged” into decisions they might not otherwise have made. The distribution of responses is shown in Figure 2.

Figure 2. Perceptions of Autonomy with AI Tools



Financial Literacy as a Moderator: Regression analysis confirmed that financial literacy moderated the relationship between AI reliance and decision consistency. High-literacy investors who used AI reported greater alignment with long-term goals, while low-literacy investors showed increased inconsistency. The regression model is presented in Table 2.

Table 2. Regression Analysis Results

Variable	Coefficient (β)	p-value
Financial Literacy	0.42	<0.01
Trust in AI	0.36	<0.05
AI Reliance	-0.21	<0.05

The highlights of the survey findings are summarized in Table 3.

Table 3. Survey Highlights

Variable	% Agree
Trust in AI	61
Perceived Empowerment	45
Perceived Over-Nudging	28
High-Literacy Consistency	72
Low-Literacy Consistency	38

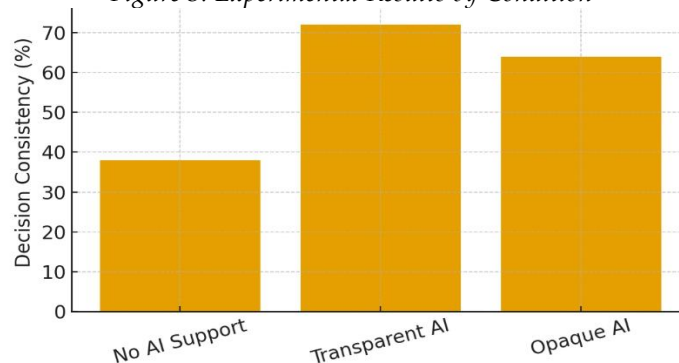
Experimental Results

The controlled experiment with 90 participants tested how transparency influenced decision-making.

- No AI support: Only 38% maintained alignment with their stated risk preferences.
- Transparent AI support: Consistency rose significantly to 72%.
- Opaque AI support: Consistency also improved (64%), but participants reported feeling they accepted suggestions “because the system must know best.”

The differences across conditions are shown in Figure 3.

Figure 3. Experimental Results by Condition



ANOVA tests confirmed that transparent AI produced the most significant improvement ($F(2,87) = 6.32, p < 0.01$).

Thematic Insights

Qualitative interviews across all methods provided richer insight into the lived experience of disrupted rationality. Three recurring themes emerged:

1. Cognitive Rescue – Many participants described AI as a “safety net” that reduced anxiety and prevented rash decisions.
2. Digital Overreach – Others spoke of feeling “pushed” by their platforms, with one noting, “I wasn’t sure if I made the choice or if the app made it for me.”
3. Conditional Trust – Experienced investors emphasized the value of transparency, while less experienced ones tended to over-trust.

Representative quotes and their interpretation are shown in Table 4.

Table 4. Thematic Insights (Qualitative)

Theme	Representative Quote
Cognitive Rescue	“AI feels like a safety net that stops me from panicking.”
Digital Overreach	“I wasn’t sure if I made the choice or if the app made it for me.”
Conditional Trust	“When I see why the system recommends something, I trust it. Otherwise, I just follow blindly.”

Cross-Method Synthesis

To integrate insights across methods, Table 5 summarizes the main findings side-by-side.

Table 5. Summary of Findings Across Methods

Method	Key Finding
Behavioral Tracking	AI reduces impulsive trades but overreliance observed
Survey	High trust in AI; autonomy split; literacy moderates impact
Experiment	Transparent AI improved consistency; opaque AI raised blind reliance

Discussion and Conclusions

While the results confirm that artificial intelligence (AI) can support investor decision-making, they also reveal new tensions around autonomy, overreliance, and transparency. Together, these insights provide the foundation for the proposed concept of *disrupted rationality*—a framework that captures how technology both rescues and unsettles rational decision-making in the context of retail investing.

Linking Findings to Theory

The results confirm AI’s dual role as both cognitive rescue and digital overreach. On the positive side, AI reduced impulsive trading and improved goal alignment, echoing behavioral finance insights on biases like loss aversion and overconfidence (Kahneman & Tversky, 1979; Barberis & Thaler, 2003). Through noise filtering and information structuring, AI served as an external scaffold that promoted rationality (Payne et al., 1993). Nevertheless, most investors took advice unquestioningly, even when it contradicted their preferences, demonstrating over-trust in black box systems (Dietvorst et al., 2015; Burrell, 2016). Rationality is therefore not merely bounded by human constraints (Simon, 1982) but also unsettled by technological mediation. The theory of broken rationality encapsulates the two-nature interpretation, realizing that choices are influenced by psychology as well as by digital interventions.

Practical Implications

Three lessons result. Transparency and explainability are key design principles for fintech, employing capabilities such as "why this recommendation" to encourage critical scrutiny. Investor education has to go beyond financial literacy to incorporate digital literacy, allowing users to challenge algorithmic responses (Lusardi & Mitchell, 2014). Policy and regulation must strike a balance between protection and innovation, mandating disclosure of algorithmic reasoning, establishing ethical nudging standards, and preventing conflicts of interest (Yeung, 2017).

Conceptual Contributions

This study makes three contributions: (1) it extends bounded rationality by incorporating external technological disruptions; (2) it bridges behavioral and digital finance, offering an integrative framework for AI-driven decision-making (Lo & Zhang, 2021); and (3) it formalizes disrupted rationality as a construct to analyze AI's simultaneous potential to stabilize and distort investor judgment.

Conclusions

The findings of this study highlight the paradoxical role of AI in retail investing. On the one hand, AI offers significant benefits: it reduces impulsive behavior, supports goal alignment, and fosters disciplined strategies. On the other hand, it introduces new vulnerabilities: overreliance, loss of autonomy, and manipulation through opaque or biased nudges.

By conceptualizing these dynamics as disrupted rationality, this study advances understanding in behavioral finance and digital investment research. The construct captures the duality of AI as both cognitive rescue and digital overreach. Practically, it signals to designers, educators, and regulators that the challenge is not simply to expand access to AI tools but to ensure they are used in ways that preserve investor autonomy and safeguard rational decision-making.

In sum, AI is reshaping the rationality of retail investors. It can empower, but it can also unsettle. Recognizing and managing this tension is essential for building a financial ecosystem where digital transformation enhances, rather than disrupts, the capacity of individuals to make sound, autonomous decisions.

Limitations and Future Research

No research is without its boundaries and acknowledging them is essential for both transparency and credibility. While the present study advances understanding of how AI shapes retail investor rationality, several limitations should be noted. At the same time, these limitations open valuable avenues for future inquiry.

Limitations

This research has a few limitations. First, sampling bias due to purposive and snowball sampling of digitally engaged investors, favoring younger and technologically proficient participants and underrepresenting older or less digitally competent investors, restricts generalizability (Lusardi & Mitchell, 2014). Second, the narrow observation period of four weeks captures short-term decision patterns but possibly fails to capture long-term patterns of disrupted rationality, necessitating longer-term research (Lo, 2017). Third, experimental limitations within simulated trading restrict ecological validity, since participants might play differently if they have real financial interests at stake (Camerer & Hogarth, 1999). Fourth, dependence on self-reporting creates subjectivity, although behavior monitoring reduces this partially (Podsakoff et al., 2003). Lastly, the scope of platforms examined centered on mass-market robo-advisors and trading apps and did not include upcoming solutions like DeFi and AI-based social trading platforms, which could shape disoriented rationality differently (Arner et al., 2016).

Directions for Future Research

Building on these limitations, several directions for future inquiry emerge.

Longitudinal Studies

Future research should adopt longitudinal designs to capture the evolution of disrupted rationality over time. Such studies could assess whether reliance on AI leads to sustained improvements in decision-making, gradual erosion of autonomy, or cyclical patterns of over-trust and disillusionment.

Cross-Cultural Comparisons

Rationality is not only shaped by cognition and technology but also by cultural norms and regulatory environments. Comparative studies across markets—for example, contrasting investors in the United States, Europe, and Asia—could reveal how cultural attitudes toward technology, risk, and authority mediate the experience of disrupted rationality (Hofstede, 2001).

Field Experiments

While lab experiments highlight causal relationships, field experiments in real trading environments would enhance external validity. Collaborations with fintech firms could allow researchers to test design interventions (e.g., transparency tools, nudging strategies) in live settings, measuring actual investment outcomes.

DeFi and Emerging Platforms

As financial innovation accelerates, new platforms such as DeFi protocols and social trading apps deserve closer attention. These systems combine AI with community-driven dynamics, potentially amplifying both cognitive rescue and digital overreach. Studying these emerging technologies will ensure that the concept of disrupted rationality remains relevant in future financial landscapes.

Interventions and Education

Finally, research should explore how educational and design interventions can mitigate disrupted rationality. For example, can interactive tutorials improve digital literacy and reduce blind reliance on AI? Do transparency features consistently enhance autonomy across user groups, or do they only benefit those with prior financial knowledge? Such work would provide actionable insights for regulators, designers, and educators.

This study makes an important contribution by introducing disrupted rationality as a construct for understanding AI's dual role in retail investing. Yet its conclusions must be read considering sampling, temporal, experimental, and scope-related limitations. By extending the research through longitudinal, cross-cultural, and field-based approaches, future scholars can deepen understanding and refine the construct. In doing so, they will help ensure that the integration of AI into financial services promotes not only efficiency but also autonomy, fairness, and genuinely rational decision-making.

References

- Arner, D.W., Barberis, J. and Buckley, R.P. (2016) 'The evolution of fintech: A new post-crisis paradigm?', *Georgetown Journal of International Law*, 47(4), pp. 1271–1319.
- Baker, H.K. and Dellaert, B.G.C. (2019) 'Regulating robo-advisors: A behavioural finance perspective', *Journal of Financial Regulation and Compliance*, 27(3), pp. 237–251.
- Baker, H.K. and Ricciardi, V. (2014) 'How biases affect investor behaviour', in H.K. Baker and V. Ricciardi (eds.) *Investor Behavior: The Psychology of Financial Planning and Investing*. Hoboken: Wiley, pp. 1–30.
- Barberis, N. and Thaler, R. (2003) 'A survey of behavioral finance', in G. Constantinides, M. Harris and R. Stulz (eds.) *Handbook of the Economics of Finance*. Amsterdam: Elsevier, pp. 1053–1128.
- Benartzi, S. and Thaler, R.H. (1995) 'Myopic loss aversion and the equity premium puzzle', *The Quarterly Journal of Economics*, 110(1), pp. 73–92.
- Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology', *Qualitative Research in Psychology*, 3(2), pp. 77–101.
- Burrell, J. (2016) 'How the machine "thinks": Understanding opacity in machine learning algorithms', *Big Data & Society*, 3(1), pp. 1–12.
- Camerer, C.F. and Hogarth, R.M. (1999) 'The effects of financial incentives in experiments: A review and capital-labor-production framework', *Journal of Risk and Uncertainty*, 19(1–3), pp. 7–42.
- Cohen, J. (1992) 'A power primer', *Psychological Bulletin*, 112(1), pp. 155–159.

- Creswell, J.W. and Plano Clark, V.L. (2017) *Designing and Conducting Mixed Methods Research*. 3rd edn. Thousand Oaks, CA: SAGE.
- D'Acunto, F., Prabhala, N. and Rossi, A.G. (2019) 'The promises and pitfalls of robo-advising', *The Review of Financial Studies*, 32(5), pp. 1983–2020.
- De Bondt, W.F.M. and Thaler, R.H. (1995) 'Financial decision-making in markets and firms: A behavioral perspective', in R.A. Jarrow, V. Maksimovic and W.T. Ziemba (eds.) *Handbooks in Operations Research and Management Science: Finance*. Amsterdam: Elsevier, pp. 385–410.
- Dietvorst, B.J., Simmons, J.P. and Massey, C. (2015) 'Algorithm aversion: People erroneously avoid algorithms after seeing them err', *Journal of Experimental Psychology: General*, 144(1), pp. 114–126.
- Fama, E.F. (1970) 'Efficient capital markets: A review of theory and empirical work', *The Journal of Finance*, 25(2), pp. 383–417.
- Gomber, P., Koch, J.-A. and Siering, M. (2018) 'Digital finance and fintech: Current research and future research directions', *Journal of Business Economics*, 87(5), pp. 537–580.
- Goldfarb, A. and Tucker, C. (2019) 'Digital economics', *Journal of Economic Literature*, 57(1), pp. 3–43.
- Hofstede, G. (2001) *Culture's Consequences: Comparing Values, Behaviors, Institutions and Organizations Across Nations*. 2nd edn. Thousand Oaks, CA: SAGE.
- Johnson, R.B. and Onwuegbuzie, A.J. (2004) 'Mixed methods research: A research paradigm whose time has come', *Educational Researcher*, 33(7), pp. 14–26.
- Jung, D., Dorner, V., Glaser, F. and Morana, S. (2022) 'Robo-advisory: Digitalization and automation of financial advisory', *Business & Information Systems Engineering*, 64(1), pp. 49–65.
- Kahneman, D. and Tversky, A. (1979) 'Prospect theory: An analysis of decision under risk', *Econometrica*, 47(2), pp. 263–291.
- Langer, M. and Landers, R.N. (2021) 'Algorithmic aversion in financial decision making', *Journal of Behavioral and Experimental Finance*, 29, 100454.
- Lo, A.W. (2017) *Adaptive Markets: Financial Evolution at the Speed of Thought*. Princeton: Princeton University Press.
- Lo, A.W. and Zhang, R. (2021) 'Breaking the bounds of bounded rationality: Expanding the behavioral foundations of finance', *Journal of Economic Literature*, 59(3), pp. 743–791.
- Lusardi, A. and Mitchell, O.S. (2014) 'The economic importance of financial literacy: Theory and evidence', *Journal of Economic Literature*, 52(1), pp. 5–44.
- MacKenzie, D. (2018) *Trading at the Speed of Light: How Ultrafast Algorithms Are Transforming Financial Markets*. Princeton: Princeton University Press.
- Markham, A. and Buchanan, E. (2012) 'Ethical decision-making and internet research: Recommendations from the AoIR ethics working committee', *AOIR Working Paper*, pp. 1–19.
- Nunnally, J.C. (1978) *Psychometric Theory*. 2nd edn. New York: McGraw-Hill.
- Payne, J.W., Bettman, J.R. and Johnson, E.J. (1993) *The Adaptive Decision Maker*. Cambridge: Cambridge University Press.
- Podsakoff, P.M., MacKenzie, S.B., Lee, J.Y. and Podsakoff, N.P. (2003) 'Common method biases in behavioral research: A critical review of the literature and recommended remedies', *Journal of Applied Psychology*, 88(5), pp. 879–903.
- Shefrin, H. (2000) *Beyond Greed and Fear: Understanding Behavioral Finance and the Psychology of Investing*. Boston: Harvard Business School Press.
- Simon, H.A. (1982) *Models of Bounded Rationality: Behavioral Economics and Business Organization*. Cambridge, MA: MIT Press.
- Thaler, R.H. and Sunstein, C.R. (2008) *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven: Yale University Press.
- Yeung, K. (2017) "'Hypernudge': Big data as a mode of regulation by design", *Information, Communication & Society*, 20(1), pp. 118–136.