How Investors Pick Stocks:
Global Evidence from 1,540 AI-Driven Field Interviews

Byoung-Hyoun Hwang* Don Noh Sean Seunghun Shin

NTU Singapore HKUST KAIST

This draft: December 2, 2025

Abstract

We conduct AI-driven field interviews with 1,540 actual investors, including many "millionaires," across ten countries to examine how individuals pick stocks. Textual analysis of the transcripts uncovers thirteen recurrent mechanisms that together form an empirically grounded taxonomy of actual investor behavior. Several frequently invoked mechanisms are only partially reflected in current mainstream asset-pricing theories. We also document substantial heterogeneity both across and within investors: while some emphasize certain mechanisms, others rely on markedly different ones, and many draw simultaneously on multiple approaches. Our evidence points to the need for theoretical refinements that accommodate heterogeneous agents with distinct preferences and belief systems.

JEL Classification: G11, G12, G14, G40, G50

Keywords: Field Interview, Textual Analysis, Large Language Models, Artificial Intelligence, Risk, Behavioral Finance, Social Finance

*Hwang is with NTU - Singapore. Noh is with HKUST. Shin is with KAIST. E-mail: bh.hwang@ntu.edu.sg, noh@ust.hk and sean.shin@kaist.ac.kr.

The authors would like to thank Briana Chang, Hui Chen, Darwin Choi, Zhi Da, Steve Dimmock, David Hirshleifer, Lawrence Jin, Abhiroop Mukherjee, Mikael Paaso, Johannes Schwab and seminar participants at KAIST, the University of Hong Kong, University of Washington and audiences at the 2025 European Summer Symposium in Financial Markets for their valuable suggestions and comments. The authors also thank Jay Choi for excellent research assistance.

1. Introduction

The finance literature has developed a wide range of theories of investor and stock market behavior.[1] The variety of models raises the natural question of which framework, if any, most accurately describes the real world. One empirical strategy to tackle this question is to test which models' predictions are most consistent with observed patterns in the stock market. The difficulty with this approach, however, is that many theories overlap in their predictions (Liu, Peng, Xiong, and Xiong, 2022).

While many theories yield similar predictions, the assumed investor behavior that gives rise to these predictions differs markedly. As an alternative strategy, one could therefore test which model's assumption is most consistent with how investors actually behave.

In this paper, we follow this approach. We interview actual investors, including those with substantial assets under management (AUM), and inquire what led them to invest in particular stocks in the past. Field interviews are a uniquely powerful method for uncovering how individuals think and feel about complex tasks. Unlike other methods that researchers have begun using to describe actual investors' decision-making processes, such as surveys (Choi and Robertson, 2020; Bender, Choi, Dyson, and Robertson, 2022) and textual analysis of analysts' and bloggers' buy recommendations (Chen, Hwang, and Peng, 2025), field interviews are dynamic and adaptive, allowing the interviewer to probe for greater detail and context (Lamont, 2009; Kaplan, 2008). The interviewer can ask clarifying follow-ups and solicit concrete examples. These adaptive clarifying questions can transform short, generic, and ambiguous statements into rich, specific, and interpretable information.

A brief example, drawn from an actual interview transcript in our sample, illustrates the value of this dynamic and adaptive structure. When asked how they choose between stocks, a respondent initially replied, "A combination of factors including stability and dividends." Non-interactive methods would stop at this generic statement, leaving unclear what "stability" means to this particular investor. In contrast, our interviewer followed up with "Thanks for sharing that. When you mention stability, what does that mean to you in the context of picking a stock?" After probing further, we learn that "stability" to this investor means "will the firm exist in five years," which the respondent judges by the firm's "relevance to the world," proxied by "the products it produces, their usefulness," and everyday adoption ("Do people I know use their products such as Apple/iPhone?"). In the end, we learn that this investor selects stocks they believe are safe in the sense

[1] We describe these theories and the corresponding papers in Section 3.1.

that they think the underlying company is likely to endure, which they infer from product-market strength. This interactive dialogue thus surfaces a latent consideration that a one-shot response would have missed and might otherwise have conflated with entirely different, more textbook-oriented notions, such as a stock contributing stability to the investor's portfolio because of its low covariance with the overall market.

Because well-conducted interviews are open-ended, they also avoid common concerns associated with surveys, particularly those related to how questions and answer choices are worded and framed (McCracken and McCracken, 1988; Schuman and Presser, 1996; Kvale and Brinkmann, 2009; Patton, 2014). The open structure also allows respondents to raise motivations and nuances that researchers may not have anticipated and therefore could not have possibly incorporated into their survey design.

While we rely on field interviews, our approach incorporates a key innovation. Recent studies note that artificial intelligence (AI) agents driven by Generative Pre-trained Transformers (GPT) can conduct interviews of quality comparable to those of human experts (Chopra and Haaland, 2023; Geiecke and Jaravel, 2024). Building on this evidence, we conduct field interviews, but replace human interviewers with AI agents.[2]

AI-led interviews offer several advantages. First, they hold interviewer quality constant. Second, they remove unwanted variation due to personal chemistry (or the lack thereof) between the interviewer and the interviewees. Most importantly, AI agents are cheaper than human interviewers and enhance scalability. Traditional human-led field interviews are typically limited to small samples because interviewers must be trained individually and dispatched to conduct conversations. For example, Dichev, Graham, Harvey, and Rajgopal (2016) draw inferences from in-person interviews with twelve chief financial officers and two standard setters. In contrast, AI-led interviews can be initiated simply by sharing a link to a chatroom. Once a human enters the chatroom, the AI agent commences the conversation. At least in principle, once one AI agent has been trained, researchers can therefore conduct thousands of AI-driven interviews simultaneously and in multiple languages.

We leverage these advantages and conduct AI-led interviews on a uniquely large and diverse pool of investors, made possible through a collaboration with CoreData Research. CoreData Research is a market research firm that conducts investor surveys for major financial institutions. Our pool of interviewees comprises 1,540 actual investors from ten different countries: 280 investors from the United States (US) and 140 investors each from Australia, Canada, France, Germany, India, Japan, South Korea, Singapore, and the

[2]GPT and other large language models (LLMs) are a type of AI. Throughout the paper, we use the terms "GPT," "LLM," and "AI" interchangeably.

United Kingdom (UK). Within each country, we stratify respondents by wealth: 29% have investable assets of USD 1,000,000 or more, 29% have between USD 500,000 and USD 999,999, and 42% have between USD 1,000 and USD 499,999. We also stratify by age, education and gender, allowing for a relatively clean comparison across the ten countries.

We develop an AI-based interactive web app. Upon answering some background information questions, investors are directed to our app (https://hns-ai-interview.streamlit.app/) and asked in their preferred language what led them to purchase particular stocks in the past (over other stocks they ended up not buying). The AI agent then poses 10–15 follow-up questions to probe and clarify their motivations. We instruct the agent to adhere to best practices from field interviewing, including openness, non-directedness, and empathy, while still pressing for specificity and coherence. The average respondent spends twenty-three minutes conversing with our AI agent and answering our questions.

After the interviews, we use GPT to conduct a textual analysis of the transcripts. For each transcript, our analysis produces a causal map that traces the investor's stated reasoning from signals, beliefs and preferences through to the final buy decision. These maps are then compared across investors to identify recurrent mechanisms that consistently appear in practice. Our process yields a total of thirteen recurrent mechanisms that together capture almost the full range of preferences and belief formation processes of the investors in our sample. The thirteen mechanisms emerge entirely from the transcripts in a bottom-up fashion, without the imposition of any external theoretical priors.

These are the mechanisms, along with descriptions and the fractions of transcripts in which a given mechanism is a "meaningful contributor:"

1. *Fundamental Strength* [40.5%]: Investors buy stocks when audited financial information indicate robust current performance, which, to investors invoking this mechanism, also signals high expected future returns.
2. *Growth/Innovation* [37.6%]: Investors purchase stocks when they see substantial growth potential tied to innovative products.
3. *Familiarity/Brand Affect* [28.4%]: Brand familiarity and positive product experiences generate intuitive comfort, which, to investors invoking this mechanism, prompts purchase decisions.
4. *Blue-Chip Comfort* [28.1%]: Investors favor large, stable firms whose established operating history and low perceived risk create a sense of safety.

5. *Authority-Follow* [18.6%]: Investors delegate stock selection to trusted experts or platforms, relying on source credibility rather than independent analysis.
6. *Momentum* [17.9%]: Investors buy stocks with strong, recent returns and high trading volume as these stocks' prices are expected to continue to drift upwards.
7. *Confluence* [14.6%]: Investors act only when multiple independent signals from different sources give a buy recommendation.
8. *Dividends* [14.0%]: Investors emphasize dividends over capital gains and purchase stocks whose dividend yields and payout stability satisfy their income requirements.
9. *Social-Copy* [13.6%]: Investors follow stock picks of trusted peers or family members, relying on relational trust rather than analytical verification.
10. *Valuation/Mispricing* [13.6%]: Investors buy when valuation metrics indicate that a stock trades below its intrinsic value.
11. *Buy-the-Dip* [12.5%]: Investors interpret temporary price declines as buying opportunities as they believe in mean reversion.
12. *ESG/Values* [10.3%]: Investors restrict purchases to firms that meet their ethical or ESG standards.
13. *Technical Analysis* [6.8%]: Investors buy when technical patterns generate predefined entry signals independent of fundamental considerations.

Only 0.8% of interview transcripts remain unexplained, in the sense that none of the above mechanisms are a meaningful contributor to the purchasing decision.

Our study makes three broad observations and contributions to the finance literature. Our first contribution is the identification and description of thirteen recurrent mechanisms that investors invoke when selecting stocks. Derived inductively from interviews with 1,540 investors across ten countries, including many "millionaires," these mechanisms offer a comprehensive and empirically grounded taxonomy of actual investor behavior. Our documentation complements the literature on investor behavior and household finance, which has traditionally inferred investor motivations from observed trading decisions, survey responses and textual analysis of investor buy recommendations (e.g., Barber and Odean, 2000, 2001; Bender, Choi, Dyson, and Robertson, 2022; Chen, Hwang, and Peng, 2025).

Recent work has begun to explore whether LLM-based agents can form beliefs similar to those of actual investors (e.g., Bhagwat, Cookson, Dim, and Niessner, 2025; Bybee, 2023). Relatedly, researchers have started to simulate financial markets in which multiple AI agents trade against one another (e.g., Henning, Ojha, Spoon, Han, and Camerer, 2025; Lopez-Lira, 2025). To conduct such analyses, researchers must endow LLMs with distinct investor personas. Our evidence on which mechanisms investors invoke (and, importantly, also to what extent) should provide useful guidance to this rapidly growing literature, improving both the realism of LLM-based measures of investor beliefs and the fidelity of LLM-based market simulations.

Several of the recurrent mechanisms that we identify align with existing theoretical frameworks. For instance, *Momentum* aligns with the "extrapolation framework," which assumes that investors (over)extrapolate recent trends into the future (e.g., Barberis and Shleifer, 2003; Barberis, Shleifer, and Vishny, 1998; Barberis, Greenwood, Jin, and Shleifer, 2015, 2018); *Social-Copy* aligns with the "social finance framework," which emphasizes the role of social interactions in shaping investment choices (e.g., Brown, Ivković, Smith, and Weisbenner, 2008; Hirshleifer, 2020; Huang, Hwang, and Lou, 2021; Chen and Hwang, 2022; Han, Hirshleifer, and Walden, 2022; Gelman, Hirshleifer, Levi, and Reiter Gavish, 2024); and *Blue-Chip Comfort* aligns broadly with the "risk framework," which assumes that investors value stability and willingly accept modest expected returns if an asset helps them achieve such stability.

Important caveats, however, also emerge. Although *Momentum* is related to the extrapolation framework, a defining feature of our *Momentum* mechanism is that investors believe recent price trends will continue only when accompanied with substantial trading volume. This thinking is in line with Lee and Swaminathan (2000), who find that the momentum effect is substantially stronger when trading volume is high, but it is currently absent from existing extrapolation models. Similarly, while the "risk framework" defines a "safe" stock as one with low covariance with adverse states of the world, investors in our sample invoking the *Blue-Chip Comfort* mechanism do not allude to concepts tied to covariance. Instead, safety is measured by whether the company is likely to survive in the future or by the stock's stand-alone volatility. Neither of these two safety definitions appear in standard formulations of the risk framework.

We also find that a non-negligible share of transcripts prominently feature mechanisms that do not align, or, align only weakly, with mainstream asset pricing theory, such as *Confluence* or *Dividends*. For instance, traditional asset-pricing theory focuses on total returns and treats the two sources of total returns, capital gains and dividends, as perfectly substitutable, building on the dividend irrelevance result of Miller and Modigliani (1961). In contrast, investors invoking the *Dividends* mechanism consider dividends separately and only purchase stocks whose dividend yield meets their minimum income requirement. Our results point to behavioral dimensions that are absent in current leading theories, thereby offering guidance for future theoretical refinement and extension.

Our second broader observation and contribution to the literature is the documentation of substantial heterogeneity in the mechanisms investors invoke, both within and across individuals. Nearly all investors in our sample report drawing from multiple mechanisms when deciding which stocks to purchase, and these mechanisms often point in different directions. For example, some investors report relying on both the *Growth/Innovation* mechanism, which emphasizes speculation on upside potential, and the *Blue-Chip Comfort* mechanism, which prioritizes safety over returns. Heterogeneity also exists across investors. The modal investor in our sample simultaneously invokes two or three mechanisms. When we consider all possible investor pairs, we find that when the focal investor invokes two mechanisms (three mechanisms), the fraction of times that the other investor invokes *none* of the mechanisms that the focal investor relies on is 56% (41%).

The substantial heterogeneity within and across investors points to the need for "super-models." One possibility is to construct models in which heterogeneous groups of investors, each characterized by distinct preferences and belief-formation processes, interact with one another. A second possibility is to develop models in which individual investors simultaneously exhibit multiple behavioral tendencies, such as holding lottery-type preferences while also displaying extrapolative belief tendencies. Yet another, third possibility is that investors invoke different mechanisms for different segments of their portfolios. Finally, a fourth possibility could be a "super-model" in which investors shift between mechanisms over time, depending on perceived market conditions or on a mechanism's perceived performance.

We are aware of two papers that follow the last approach. Hong, Stein, and Yu (2007) model the implications of investors' sudden "paradigm shifts" between models based on their perceived usefulness. Barberis, Greenwood, Jin, and Shleifer (2018) model a world where investors randomly "waver" between extrapolation (as in our *Momentum* mechanism) and comparing market prices to perceived fundamental values (as in our *Valuation/Mispricing* mechanism). Our results suggest that extending such frameworks and developing models that capture other forms of within- and across-investor heterogeneity may have the greatest chance of capturing the complexity of real-world behavior and produce predictions consistent with observed market outcomes. To help guide such research, we describe in the main body of the paper the mechanism pairs that most frequently co-occur in investors' transcripts.

Our third principal contribution is the systematic comparison of investor behavior across countries. We observe substantial cross-country heterogeneity, even though the demographic composition of participants is relatively homogeneous across countries. For instance, *Fundamental Strength* is the most frequent mechanism among Indian investors (60.0%) but only the fourth most frequent among US investors (30.0%; Δ = 30.0%). Likewise, in Japan, *Familiarity/Brand Affect* ranks first, appearing prominently in nearly half of the transcripts (43.6%), whereas in Singapore it ranks only sixth (19.3%; Δ = 24.3%). When we group countries based on their reliance on the thirteen mechanisms, three clusters emerge endogenously from the data. Investors in India resemble those in Singapore and South Korea. Investors in Australia, Canada, France, Germany, the UK, and the US form a largely similar cluster. Investors in Japan constitute a distinct group of their own. These observations should help the literature gauge to what degree results based on data from one country may extend to other countries.

Our study closely relates to Chopra and Haaland (2023) and Geiecke and Jaravel (2024). Chopra and Haaland (2023) compare the effectiveness of open-ended survey questions with GPT-driven interviews by asking users of the online survey platform Prolific why they choose not to participate in the stock market. Chopra and Haaland (2023) show that GPT-driven interviews elicit substantially longer and more nuanced responses than open-ended survey questions.

Geiecke and Jaravel (2024) similarly assess the value of GPT-driven interviews as a tool for social science research. They conduct interviews on three topics: meaning in life, political preferences, and educational or occupational choices. The authors find that AI-led interviews generate greater engagement and a richer understanding of motivations than open-ended surveys. Notably, they have trained sociologists evaluate the interview transcripts. The sociologists conclude that across all topics, the quality of AI-led interviews is comparable to that of interviews conducted by human experts.

Together, Chopra and Haaland (2023) and Geiecke and Jaravel (2024) provide compelling evidence that GPT-driven interviews produce accurate and insightful accounts of respondents' behavior. Our study builds on their methodological innovation and extends the literature by conducting AI-led interviews with a large and diverse pool of 1,540 actual investors across ten countries, 29% (58%) of whom have investable assets exceeding USD 1 million (USD 500,000). Our subject pool allows for a comprehensive and global perspective on actual investors' decision-making. Crucially, we apply textual analysis to the interview transcripts to identify recurrent mechanisms. We then map them to established theoretical frameworks in finance. In doing so, we extend the methodological innovation of Chopra and Haaland (2023) and Geiecke and Jaravel (2024) into direct theoretical contributions to core debates about investor behavior and financial markets.

The remainder of our paper proceeds as follows. Section 2 describes our subject pool, our interview design and our method to process the interview transcripts. Section 3 presents the corresponding results. In

Section 4, we discuss advantages and important limitations of our AI-led interviews compared with other empirical strategies to study investor decision-making, and we allude to possible directions for future research. Section 5 concludes.

2. Sample and Methodology

2.1 Subject Pool

To ensure that our findings convey information about actual investor motivations, it is essential that we conduct interviews with real investors. It is also crucial that our subject pool includes participants with varying levels of investable assets, particularly those with substantial AUM. If our sample were restricted to individuals with small AUM, which are typically the only type accessible through online survey platforms such as Prolific (https://www.prolific.com), it would be unclear whether our findings generalize to the broader investor population and carry meaningful asset pricing implications.

To meet these criteria, we collaborate with CoreData Research, a market research firm that conducts investor surveys for major financial institutions. CoreData Research maintains a large global panel of retail and institutional investors who have agreed, and cleared compliance, to receive occasional survey requests in exchange for monetary compensation. These surveys are typically commissioned by large financial institutions and include questions about market outlook and investment trends. For our study, participants were informed that they would be taking part in an academic study and participate in an AI-driven interview about how they pick stocks. We designed the AI-driven interview and paid CoreData Research a lump sum; CoreData Research handled participant outreach, participant compensation, and the delivery of the completed outputs to us.

Our final sample includes 280 investors from the US and 140 investors each from Australia, Canada, France, Germany, India, Japan, South Korea, Singapore, and the UK. Except for the US, each country includes 40 investors with USD 1,000,000 or more in investable assets, 40 investors with investable assets between USD 500,000 and USD 999,999, and 60 investors with investable assets between USD 1,000 and USD 499,999. The corresponding numbers for the US are 80, 80 and 120. In total, we interview 1,540 investors during the months of July and August 2025.

Table 1, which presents descriptive statistics for the investors in our sample, shows that our pool of interviewees exhibit wide variation in gender, age, education, investable assets and self-reported investment knowledge and risk attitude.

Table 2 reports the distribution of investors by country across key demographic variables (gender, age, education and investable assets). As per our instructions to CoreData Research, the investor composition is largely homogeneous across countries. The main exceptions are that the Japanese investors in our sample are predominantly male, and the Indian investors are younger and almost all hold a university degree.

2.2 Interview Design

Our interview consists of three stages: (i) a background information stage, (ii) the main AI-driven interview stage, and (iii) a summary and evaluation stage.

1. **First Stage – Background Information:** CoreData Research sends investors a brief closed-end (multiple-choice) survey to gather or confirm basic background information, including country of residence, age, gender, educational background, personal investments (asset allocation and invested dollar amounts), frequency of investment style changes, sources of income, perceived investment knowledge, and perceived risk appetite. The full questionnaire is shown in Online Appendix A.

   If a respondent indicates that they currently have money allocated to individual stocks, they are directed to our AI-driven interactive web app: https://hns-ai-interview.streamlit.app/.

2. **Second Stage – AI-driven Interview:** Once participants open our app, they engage in a one-on-one interview with an AI agent. The AI agent is driven by GPT-4.1, which, at the time we developed the AI-led interview, was the most recent GPT model. We describe our instructions to the AI agent in Online Appendix B. These instructions follow best research practices and are designed to mimic the training we would provide to a human interviewer.

   While each interview is unique in both content and line of questioning, reflecting its open-ended and adaptive nature, it adheres to a consistent structure:

   (2.a) Upon entering our platform, participants are asked in five languages (English, French, German, Japanese, or Korean) to provide two pieces of information: (i) their preferred language (English, French, German, Japanese, or Korean) and (ii) their unique survey ID, which they receive after completing the initial background survey (see figure below). The remainder of the interview is conducted in the selected preferred language.

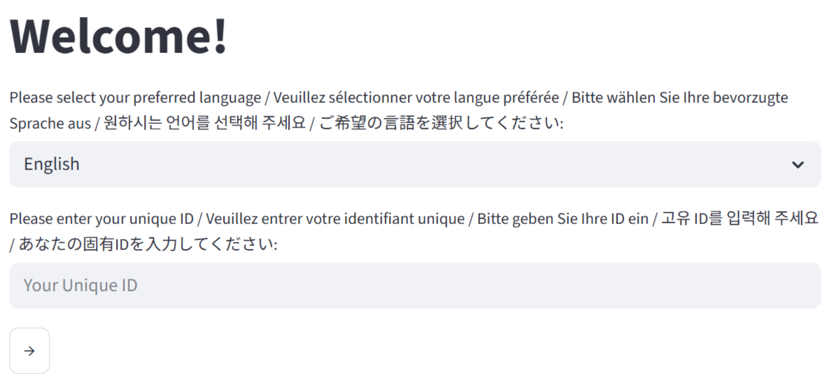

(2.b) The respondent subsequently receives a greeting and brief instructions, and the AI agent initiates the conversation with the opening question (see figure below).

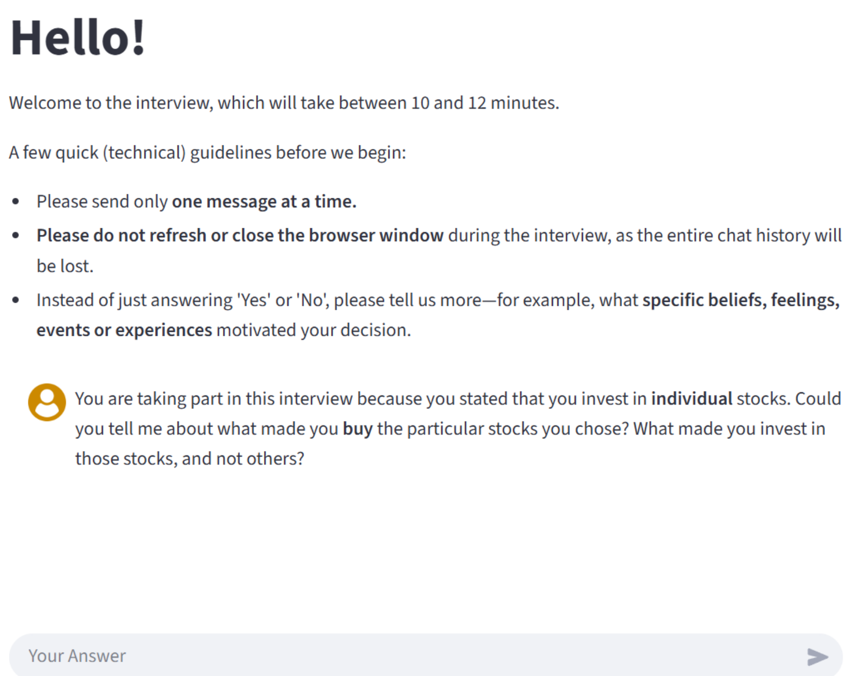

(2.c) The interview proceeds with 10 to 15 follow-up questions. These questions adapt to the conversation and aim to probe deeply, yet non-directively, into the interviewees' reasons for purchasing specific stocks.

After completing the follow-up questions, the AI agent asks:

> Is there anything else that plays an important role when you decide which stocks to buy? Anything we haven't touched on yet?

If the interviewee indicates that there is something else, the AI agent asks two to three further probing questions on the untouched subjects.

(2.d) After the interviewee confirms that all relevant areas have been covered, the AI agent asks:

> Just one quick follow-up question: After you **sell** a stock, do you continue to track its performance and how its **price evolves**?
> 1 (Never)
> 2 (Sometimes)
> 3 (Always)
> Please only reply with the associated number.

The interview then moves to the third and final stage.

3. **Third Stage – Summary and Evaluation.**

(3.a) In the first part of this stage, the AI agent summarizes the key points from the interview and invites the respondent to evaluate its accuracy:

> How well does the above summarize how you approach investing in individual stocks?
> 1 (It's a poor summary)
> 2 (It's a fair summary)
> 3 (It's a good summary)
> 4 (It's a very good summary)
> Please only reply with the associated number.

(3.b) The AI agent then asks the respondent to (i) rate the overall quality of the interview on a four-point scale; (ii) select whether they would have preferred a human interviewer (over an AI agent); and (iii) indicate whether they would have preferred an open-ended survey question (over an AI-led conversation). We display the exact questions and answer choices that our participants see in Online Appendix B.

(3.c) The AI agent closes the session with a brief thank-you note.

On average, respondents spent approximately four minutes completing the background information survey and nineteen minutes in the AI-driven interview, summary, and evaluation stages. The interview duration is broadly similar across demographic groups, but male, older, and higher-wealth investors spent an additional one to three minutes on average (Table OA1). Because all participants consented ex ante to completing the full protocol, the dropout rate is zero.

Online Appendix Figure OA1 reports the responses to questions (3a) and (3b). Ninety-four percent rate the summary as either "very good" (62%) or "good" (32%), indicating that the AI agent captures their investment approach with a high degree of fidelity. Similarly, ninety-two percent rate the quality of the overall interview as either "very good" (60%) or "good" (32%). Only 11% express a preference for a human interviewer over an AI interviewer and only 15% would have preferred a free-text response ("open-ended survey question") despite the significant amount of time they spent on the AI interview. Table OA2 shows that the preference for an AI interviewer is broadly similar across demographic groups, though younger investors and respondents in Asian countries exhibit a somewhat stronger preference for an AI interviewer. Taken together, these quality metrics indicate that the AI-led interviews yield content that respondents regard as accurate and engaging, with the overwhelming majority viewing the AI-based format favorably.

2.3 Interview Transcript Processing / Causal Map Construction

After receiving the 1,540 interview transcripts, we analyze them using a three-step procedure driven by GPT-5, which, at the time of our analysis, was the most recent GPT model available (see figure below):

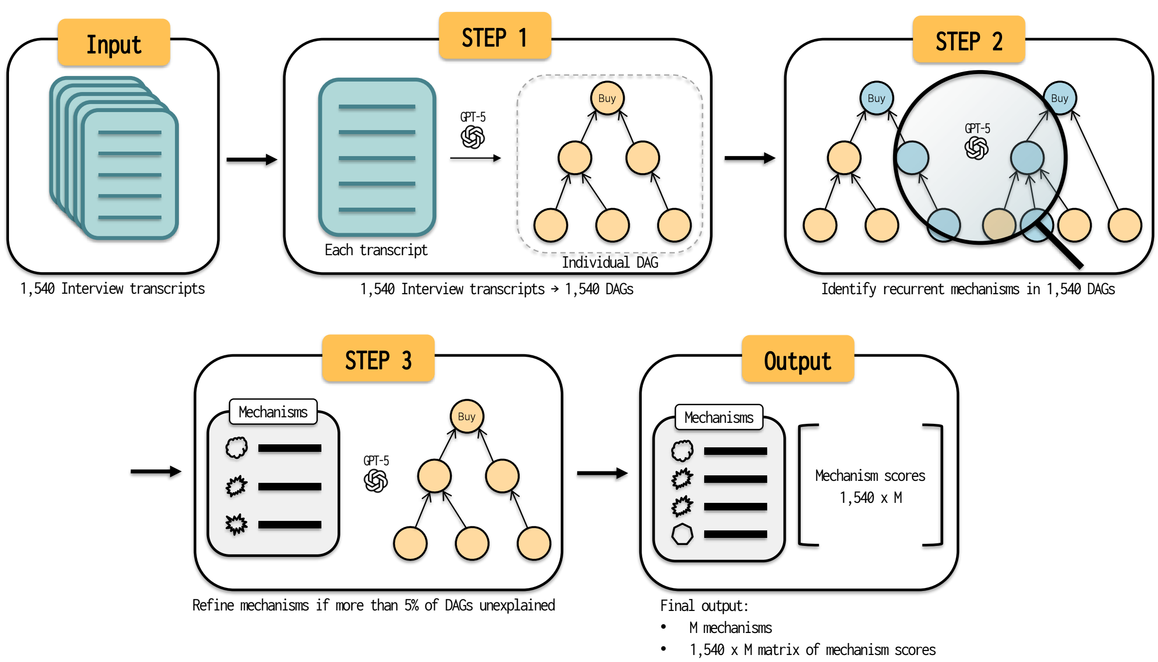

Step 1: We convert each interview transcript into an "argument map," specifically, a directed acyclic graph (DAG) that outlines the investor's key reasoning.

Step 2: We then examine all 1,540 maps and identify recurring decision patterns, which we term as "mechanisms."

Step 3: Finally, we review all 1,540 maps again and determine whether more than 5% are not captured by the current set of mechanisms. If so, we refine or expand the set of mechanisms.

Output: Our process yields a final set of thirteen mechanism descriptions and, for each interview transcript, thirteen scores indicating how closely the investor's reasoning aligns with each of the identified thirteen mechanisms.

The ensuing subsections provide more information on each of the three steps and the output. We emphasize intuition whenever possible. For technical details, we refer interested readers to Online Appendix C.

2.3.1 Step 1: Constructing "Argument Maps"

We convert each interview transcript into an investor-level structural causal model (SCM) represented by a DAG. Since Pearl (1995, 2009), SCMs and DAGs have been widely employed in sociology, political science, epidemiology, and computer science. Their principal advantage lies in their ability to encapsulate complex statements and arguments within a compact representation, thereby providing a clear and intuitive summary of underlying causal relationships and sharpening the GPT model's ability to identify recurring mechanisms.[3]

DAGs can be viewed as a form of an argument map, also known as an argument diagram. In traditional argument mapping, writers create a visual diagram to break down and analyze the structure of an argument. The figure below is an example of a simple, traditional argument map.

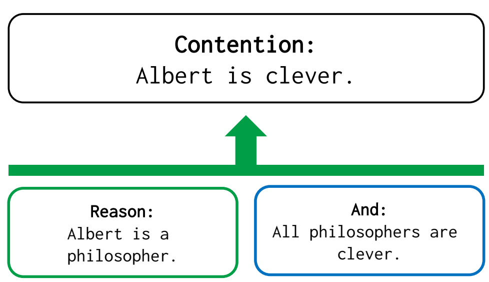

Argument maps use boxes and directed lines to represent the logical connections between a conclusion and its supporting reasons ("premises"). This approach helps the writer organize their arguments and identify potential logical gaps. Once the argument map is complete, the writer then translates the map into prose. In

[3]We hereafter refer to "SCMs and DAGs" simply as "DAGs."

our setting, we reverse this process: rather than moving from map to text, we start with text and deduce its underlying map.

Each DAG consists of three components: (i) nodes, (ii) directed edges, and (iii) a node dictionary. All are generated from the interview corpus in an entirely bottom-up, data-driven manner.

(i) Nodes represent factors that influence an investor's stock-purchase decision, either directly or through intermediate nodes. Typical nodes include external signals (e.g., price movements, peers' advice), beliefs or expectations (e.g., return expectations), affect (e.g., confidence, fear), preferences (e.g., ESG orientation), and constraints (e.g., budget limits); nodes are standardized across all 1,540 DAGs.

(ii) Directed edges encode causal relationships between two nodes. For example, (X, Y) denotes that node X causally influences node Y; a set of edges $\{(X, Z), (Z, Y)\}$ indicates that X influences Y through the intermediary node Z.

(iii) The node dictionary defines each node in one short sentence.

DAGs are well suited to the interactive nature of our AI-led interviews as each adaptive follow-up question can become a new node or edge, explicitly linking additional information into the investor's evolving causal chain. To illustrate, consider Panel (a) in Figure 1.

[Figure 1 here]

An investor states that they purchase a stock because it has "next-gen products." The AI interviewer then asks what "next-gen products" signal to the investor, and the investor responds that such products indicate substantial possible growth. Such conversation would be encoded through:

- "node set" $V = \{\text{InnovativeProduct}, \text{Growth/Innovation}, \text{BuyStock}\}$
- "edge set" $E = \{(\text{InnovativeProduct}, \text{Growth/Innovation}), (\text{Growth/Innovation}, \text{BuyStock})\}$
- a "node dictionary" that defines all three nodes.

Panel (b) in Figure 1 depicts a slightly more elaborate DAG. The investor selects stocks by filtering on ESG criteria. Before purchasing, the investor assesses firms' ESG standing using website content and donation disclosures and monitors news reports to avoid firms with potential ethical conflicts. The DAG for this investor would comprise:

- "node set" V = {DonationInfo, WebsiteContent, NewsReport, ValueAlignment, ESGConflictAversion, FilterByESG, BuyStock}
- "edge set" E = {(DonationInfo, ValueAlignment), (WebsiteContent, ValueAlignment), (NewsReport, ESGConflictAversion), (ValueAlignment, FilterByESG), (ESGConflictAversion, FilterByESG), (FilterByESG, BuyStock)}
- a "node dictionary" that defines all seven nodes.

2.3.2 Step 2: Data-Driven Identification of Recurring Mechanisms Across Investors

In the second step, we analyze all 1,540 investor DAGs (nodes, directed edges, and the node dictionary) to identify recurrent mechanisms, that is, families of pathways that share a common decision-making process leading to stock purchases.

Our objective is to let the data fully reveal a set of mechanisms that, taken together, explain how investors select stocks. To this end, we instruct the GPT model to analyze the DAGs for recurring path patterns. Using the GPT model allows us to uncover patterns in the underlying logic and decision-making process leading to stock purchases, beyond mere overlaps in the same nodes or edges. Each identified recurrent mechanism comes with four output-fields: (i) a text statement of the core logic; (ii) canonical pathways, that is, typical routes in the DAGs from initial cues to stock purchases; (iii) diagnostic markers, that is, exemplar edges that instantiate the mechanism pathways; and (iv) causal roots, that is, recurrent starting nodes of the pathways.

For illustration, a recurrent mechanism may take the following form:

(i) **Text statement:** Investors buy when sustained price and volume trends reinforce beliefs in continued upward movement and predictable gains.

(ii) **Canonical pathways:** PriceUpTrend / HighVolume → UptrendPersist → BuyStock

(iii) **Diagnostic markers:**

{(PriceUpTrend, UptrendPersist), (HighVolume, UptrendPersist), (UptrendPersist, BuyStock)}

(iv) **Causal roots:** {PriceUpTrend, HighVolume}

The key benefit of identifying recurrent mechanisms through DAGs is that they help us capture not merely *what* investors do, but also *why* they do so.

2.3.3 Step 3: Validation and Refinement

In Step 3, we return to the original 1,540 interview transcripts to validate how well the identified recurrent mechanisms explain investors' behavior and to potentially refine the mechanism set to account for any important residual behaviors. Validation and refinement have been shown to improve the quality of LLM outputs (e.g., Madaan, Tandon, Gupta, Hallinan, Gao, Wiegreffe, Alon, Dziri, Prabhumoye, Yang, et al., 2023).

The specific actions are as follows:

(i) *Validation:* For each investor DAG, we assign a four-point residual score (1–4) and extract the corresponding residual patterns. A residual score of four indicates that major causal paths to the stock-purchase decision are not well explained by any of the mechanisms identified thus far; a score of one indicates that there is no residual, meaning that all key paths are effectively covered by the current mechanisms.[4] The extracted residual patterns capture the portion of the decision process left unexplained by any of the current mechanisms.

(ii) *Refinement:* If the fraction of DAGs with a residual score of greater than three exceeds 5%, we refine the mechanism set by lightly generalizing the existing mechanism descriptions and, at most, adding one new mechanism derived from the most prominent residual patterns. The refinement is data driven and automated by the GPT model.

If the fraction of DAGs with a residual score of greater than three is less than 5%, the mechanism set is deemed final and the iteration stops.

2.3.4 Final Outputs

Once the final set of recurrent mechanisms is established, we assign a four-point score (1–4) to each mechanism for every investor DAG. The mechanism score reflects how well a mechanism governs a particular investor's stock-selection decision. Each mechanism is scored independently; multiple mechanisms may therefore receive high scores for a single investor if they jointly explain the behavior. The scoring scale is summarized in the table below.

[4] We evaluate the DAG against the recurrent mechanisms' output fields (see Step 2). If clarification is required beyond the DAG, the model refers back to the original transcript for additional context.

Mechanism Score	Interpretation
4	“Primary driver” (most consequential)
3	“Meaningful contributor” (consequential but non-dominant)
2	“Fragmentary presence” (peripheral)
1	“Absent” (or contradictory)

For robustness, we repeat the scoring fifteen times and take the median score across the fifteen runs as the final mechanism score.[5]

The final output comprises thirteen distinct recurrent mechanisms (and their associated output fields) and a 1,540 × 13 matrix of final mechanism scores. Only twelve investors (0.8%) in our sample remain unexplained, in the sense that none of the thirteen mechanisms are “meaningful contributors” in these investors’ transcripts (i.e., receive a score of ≥ 3).

3. Results: Recurrent Mechanisms in Investor Interview Transcripts

Our analysis identifies thirteen recurrent mechanisms that offer a comprehensive and empirically grounded taxonomy of actual investor behavior. The mechanisms are: (i) *Fundamental Strength*, (ii) *Growth/Innovation*, (iii) *Familiarity/Brand Affect*, (iv) *Blue-Chip Comfort*, (v) *Authority-Follow*, (vi) *Momentum*, (vii) *Confluence*, (viii) *Dividends*, (ix) *Social-Copy*, (x) *Valuation/Mispricing*, (xi) *Buy-the-Dip*, (xii) *ESG/Values*, and (xiii) *Technical Analysis*.

Table 3 reports descriptions of each of the thirteen mechanisms. These thirteen mechanism descriptions are generated by GPT and not written by us.

[Table 3 here]

Figure 2 displays, for each mechanism, the proportion of transcripts in which the mechanism features prominently, based on the scoring procedure described in Section 2.3.4. The dark-gray bars represent the share of transcripts where a particular mechanism is deemed the “primary driver” of the investor’s decision. The light-gray bars denote the share of transcripts where a particular mechanism is a “meaningful contributor.”

[Figure 2 here]

[5] It turns out that the scores are highly consistent across runs, indicating strong procedural stability.

The fraction of transcripts in which a particular mechanism is deemed the "primary driver" never exceeds 16%. As we discuss later, this reflects the fact that almost all investors invoke multiple mechanisms when deciding which stocks to purchase rather than relying on a single dominant one. Throughout the remainder of the paper, we therefore emphasize the share of transcripts in which a mechanism is a "meaningful contributor." Using this broader notion of prominence, the *Fundamental Strength* mechanism appears in 40.5% of transcripts, and the *Growth/Innovation* mechanism appears in 37.6%.

Other mechanisms that feature prominently relate to investors' preference for safety and aversion to information uncertainty. *Familiarity/Brand Affect* is a "meaningful contributor" in 28.4% of transcripts, *Blue-Chip Comfort* in 28.1%, and *Confluence* in 14.6%.

Mechanisms grounded in price history also appear frequently. The fractions for *Momentum*, *Buy-the-Dip* and *Technical Analysis* are 17.9%, 12.5%, and 6.8%, respectively.

Finally, mechanisms associated with delegation (*Authority-Follow*) and social interactions (*Social-Copy*) feature prominently in 18.6% and 13.6% of transcripts, respectively. Dividends (*Dividends*) are a key focus in 14.0% of transcripts. The "standard" stock valuation approach (*Valuation/Mispricing*) is a "meaningful contributor" in 13.6% of transcripts, while the value-based mechanism *ESG/Values* appears as a "meaningful contributor" in 10.3% of transcripts.

In the next subsection (3.1), we examine how the thirteen mechanisms relate to existing theoretical frameworks in finance. In the subsections that follow (3.2–3.4), we discuss how the prevalence of these mechanisms varies within and across investors, and across countries and investor demographics.

3.1 Recurrent Mechanisms and Leading Theoretical Frameworks

There are likely more than one hundred theoretical models of investor behavior and financial markets. It is infeasible and certainly well beyond the scope of this paper to assess, for each existing model, whether our mechanisms align with the model or not. Instead, we highlight a couple of broader observations concerning the relationship between our recurrent mechanisms and leading theoretical frameworks.

3.1.1 Where Mechanisms Align with, and Extend Beyond, Existing "Mainstream Asset-Pricing Theories"

The first broader observation is that while many of the recurrent mechanisms reported by actual investors align with key assumptions in mainstream theoretical frameworks, there remain aspects of investors' thought processes that current models do not incorporate.

Momentum* and *Fundamental Strength For example, one of the most prominent classes of behavioral theories are rooted in extrapolation (e.g., Barberis and Shleifer, 2003; Barberis, Shleifer, and Vishny, 1998; Barberis, Greenwood, Jin, and Shleifer, 2015, 2018). These models assume that investors observing an increase in a stock's price or cash flows believe that such growth will persist. The tendency to extrapolate could be rooted in heuristics, such as representativeness or recency effects, which make recent historical data appear more informative than they truly are.

The *Momentum* mechanism description in Table 3 indicates that actual investors do, in fact, extrapolate past stock returns ("Investors buy when sustained price and volume trends reinforce beliefs in continued upward movement and predictable gains."). It is noteworthy, however, that investors in our sample expect recent patterns to persist only when past price movements are accompanied by high trading volume. This reasoning is consistent with the empirical finding of Lee and Swaminathan (2000) that the momentum effect is substantially stronger when trading volume is high. To the best of our knowledge, this conditioning on trading volume is not yet a central feature of theoretical models of momentum, including those that assume investor belief extrapolation.

The *Fundamental Strength* mechanism description in Table 3 suggests that actual investors also extrapolate cash flows ("Investors purchase stocks when audited financial information—such as revenues, profitability, and balance-sheet strength—indicates robust current performance and attractive expected returns."). Notably, investors invoke the *Fundamental Strength* mechanism far more frequently (40.5%) than the *Momentum* mechanism (17.9%). The substantially greater prevalence implies that cash-flow extrapolation may be even more central to real-world decision-making than return extrapolation.

The *Fundamental Strength* mechanism is consistent not only with behavioral frameworks, but also rational theories. In particular, investors' equation of high profitability with high future stock returns aligns with investment-based asset-pricing models in which high profitability implies high discount rates in equilibrium. For example, Hou, Xue, and Zhang (2015) show analytically that high profitability relative to investment implies higher required returns in equilibrium, while Novy-Marx (2013) documents that, in the cross-section, gross profitability predicts returns as strongly as the book-to-market ratio.

Growth/Innovation Another prominent class of behavioral models assumes that investors exhibit Cumulative Prospect Theory (CPT) preferences. A key element of CPT is that individuals place large "decision weights" on low-probability events, a phenomenon known as "probability weighting." These decision weights reflect a psychological tendency to overweight rare but salient possibilities. As a result of these decision weights, CPT investors willingly pay a premium for assets that they believe could deliver extreme positive outcomes and become the next Amazon or Google, even if the likelihood of such outcomes materializing is very low (Barberis and Huang, 2008; Barberis, Jin, and Wang, 2021).

The second most frequently invoked mechanism, *Growth/Innovation* mechanism, is broadly consistent with models rooted in CPT preferences and indicates that many investors indeed behave as posited by that framework ("Investors buy when they perceive a strong forward-looking growth runway, typically linked to product innovation, technological advantage, or structural demand expansion. The emphasis is on long-term upside potential rather than current fundamentals [...]").

A notable feature here is that investors in our sample do not gauge upside potential by conditioning on historical stock returns; instead, they evaluate upside by considering the product market. Chen, Hwang, and Peng (2025) show that the stocks that analysts and bloggers perceive to have the greatest upside potential are not the ones that, historically, have the highest positive return skewness. This discrepancy can lead to misleading conclusions when using historical return skewness to test CPT's prediction that stocks with high perceived skewness earn low average returns. Our finding that investors turn to the product market to gauge a stock's upside potential offers one possible explanation for the disconnect between perceived upside potential and historical return skewness.

As with the *Fundamental Strength* mechanism, the *Growth/Innovation* mechanism is not only consistent with a behavioral framework but also aligns with rational theories of investor behavior. For example, Gârleanu, Kogan, and Panageas (2012) argue that investors who are concerned about being replaced by new technologies in the workplace seek out growth stocks to hedge against such "displacement risk."

Blue-Chip Comfort The classic, traditional risk-based framework assumes that investors are Bayesian and exhibit Expected Utility Theory preferences. In this framework, investors value assets that provide high payoffs when marginal utility is high, in other words, assets that perform well during bad states of the world (e.g., Lucas, 1978; Breeden, 1979; Barro, 2006; Bansal and Yaron, 2004). Because such assets offer valuable hedging services, investors are willing to pay higher prices for them and accept lower expected returns.

The *Blue-Chip Comfort* mechanism description in Table 3 is broadly consistent with the traditional risk-based framework ("Investors prefer large, stable firms that are perceived as safe and low risk.") At the same time, a key discrepancy emerges. In the traditional framework, cross-sectional variation in expected returns is determined entirely by the *covariance* of asset payoffs with adverse states of the world. Yet none of the investors in our sample, including those invoking the *Blue-Chip Comfort* mechanism, describe safety in terms of covariance. In the *Blue-Chip Comfort* mechanism, "Signals of operating history, enduring consumer demand, and low volatility create a sense of security that can outweigh concerns about modest expected returns" (Table 3). In other words, investors operationalize safety through perceptions of company longevity via strengths in financials and product-market positioning. If there is a statistic that matters to investors in the *Blue-Chip Comfort* mechanism, it is not the second-order *mixed* moment of returns with state variables, but rather the second-order *central* moment of the company's own stock return.

Our finding mirrors Chinco, Hartzmark, and Sussman (2022), who survey U.S. investors to test the textbook prediction that the correlation between stock returns and consumption growth materially influences portfolio choice. They find that "only 11% reported thinking about consumption-growth correlations in a manner consistent with textbook theory" (p. 2186). Similarly, Chen, Hwang, and Peng (2025) show that although safety-related terms are common in analyst reports, the term "covariance" appears only 123 times across 1.72 million reports containing approximately 1.88 billion words.

3.1.2 Where Mechanisms Largely Fall Outside Existing "Mainstream Asset-Pricing Theories"

Our second broader observation regarding the relationship between our recurrent mechanisms and leading theoretical frameworks is that several mechanisms observed in the data do not fit neatly into any established models and instead reveal dimensions of investor behavior that current theories capture only partially.

Confluence The *Confluence* mechanism description in Table 3 shows a desire for confirmation and coherence across multiple information cues ("Investors require confirmation from multiple independent sources before acting. Expert opinions, online consensus, and other signals must align; conflicting information leads to inaction even when some indicators are favorable.").

This notion has been broached by the psychology and sociology literature. For instance, Harkins and Petty (1987) show that persuasive messages are more effective when coming from multiple independent sources versus a single source. Centola and Macy (2007) introduce the concept of "complex contagion," proposing that many behaviors spread through social networks only after confirmatory signals from multiple distinct contacts. Centola (2010) finds empirical support for this mechanism in the context of health behavior adoption. While the notion of *Confluence* has been explored in psychology and sociology, to the best of our knowledge, it has yet to be incorporated into mainstream asset pricing.

Dividends In the *Dividends* mechanism, "Investors purchase only when dividend yield and payout stability meet their income thresholds." This notion diverges from traditional asset-pricing theory, which focuses on total returns and treats the two sources of total returns, capital gains and dividends, as perfectly substitutable, building on the dividend irrelevance result of Miller and Modigliani (1961).

At the same time, individual investors' focus and preference for dividends is not new to the literature. Shefrin and Statman (1984) propose that investors maintain separate mental accounts for dividends and capital gains and use dividends as a self-control device to smooth consumption. Graham and Kumar (2006) find that older and low-income retail investors exhibit a stronger preference for dividend-paying stocks, which the authors attribute to self-control motives as in Shefrin and Statman (1984), but also tax incentives. More recently, Hartzmark and Solomon (2019) document that investors trade as if dividends are disconnected from price change.

The *Dividends* mechanism in our data is broadly consistent with the arguments of Shefrin and Statman (1984), Graham and Kumar (2006) and Hartzmark and Solomon (2019). A meaningful subset of investors in our sample impose minimum-dividend-yield and payout-stability thresholds and, beyond maintaining separate mental accounts, treat these as binding constraints that effectively limit their investment universe.

Our finding calls for future models that do not center on total returns but instead explicitly account for the fact that, in many investors' minds, capital gains and dividends are distinct. Incorporating such dividend-specific perceptions and preferences may yield richer theories that better align with real-world investor behavior and observed market outcomes.

3.2 Variation in Mechanisms Within Investors

The results presented thus far represent aggregate patterns. In this subsection, we show that there is substantial variation *within* investors as most investors simultaneously invoke multiple, and sometimes contrasting, mechanisms when making investment decisions.

[Figure 3 here]

Figure 3 shows the distribution of the number of mechanisms that are "meaningful contributors" to an investors' decision-making. 86.8% of investors in our sample rely on multiple mechanisms, whereas only 12.5% rely on a single one. The modal investor simultaneously employs either two (34.7%) or three (35.6%) mechanisms.

[Figure 4 here]

Figure 4 reports the ten most frequent mechanism pairs among the 1,347 investors invoking two or more mechanisms. The most common pair, (*Fundamental Strength*, *Growth/Innovation*), accounts for 21.9% of multi-mechanism investors. The next two most common pairs are (*Blue-Chip Comfort*, *Fundamental Strength*) and (*Familiarity/Brand Affect*, *Growth/Innovation*) at 11.5% and 11.2%, respectively.

Several of the frequently co-occurring mechanisms are rooted in similar higher-level construct. For instance, one frequently occurring pair is *Authority-Follow* and *Blue-Chip Comfort*, indicating that some investors feel uncertain and wary of the investment process. These investors tend to favor expert guidance and investments in well-known, stable firms that provide psychological comfort and a sense of safety.

Other co-occurring mechanisms do not share the same higher-level root. Consider the most common pair, (*Fundamental Strength*, *Growth/Innovation*). *Fundamental Strength* aligns with extrapolation of cash flows, while *Growth/Innovation* aligns with cumulative prospect theory preferences, suggesting that some investors exhibit both non-traditional preferences and irrational beliefs rooted in over-extrapolation.

Some mechanism combinations appear even theoretically inconsistent and may reflect underlying cognitive tension. The co-occurrence of *Blue-Chip Comfort* and *Growth/Innovation*, for example, mixes risk-averse and risk-seeking orientations: the former prioritizes stability and capital preservation, while the latter emphasizes speculation and potential upside. Such combinations may indicate that investors simultaneously pursue competing goals for different parts of their portfolios: seeking excitement and returns from innovation for some of their holdings while retaining the psychological reassurance of safety or familiarity in others. Alternatively, investors may shift over time between *Blue-Chip Comfort* and *Growth/Innovation*, depending on perceived market conditions or on a mechanism's perceived performance.

Taken together, our results suggest that investors generally draw on multiple behavioral mechanisms rather than adhering to a single coherent mechanism. A key implication for future research is the importance of incorporating such multi-dimensional thinking into theoretical models. There are several possibilities future work could do so. One is to develop models in which individual investors simultaneously exhibit multiple behavioral tendencies, such as exhibiting lottery-type preferences while also forming extrapolative beliefs. A second is to allow investors to invoke different mechanisms for different segments of their portfolios. A third is to allow investors to shift between mechanisms over time, depending on perceived market conditions or perceived performance. Irrespective of which avenue future work adopts, our documentation of the mechanisms that most frequently co-occur, thereby narrowing the set of plausible combinations, should make the modeling of multi-dimensional behavior less daunting.

3.3 Variation in Mechanisms Across Investors

The previous subsection documents substantial variation *within* investors. While several mechanisms are invoked relatively frequently, there is also considerable disagreement and substantial variation in the mechanisms invoked *across* investors.

To illustrate this variation, consider the 1,185,030 possible investor pairs that can be formed from our 1,540 respondents:

$$1{,}185{,}030 \;=\; \frac{1{,}540 \times 1{,}539}{2}.$$

As noted earlier, the modal investor invokes either two (34.7%) or three (35.6%) mechanisms. Consider an extreme form of disagreement where two investors invoke entirely non-overlapping sets of mechanisms: for example, one investor invokes mechanisms 1, 4, and 7, while the other invokes mechanisms 3, 5, and 10. Figure 5 shows that such extreme cases of across-investor-disagreement are far from rare.

[Figure 5 here]

Figure 5 reports the fraction of investor pairs in which none of the mechanisms invoked by one investor is invoked by the other; put differently, the fraction of pairs in which the two investors have zero mechanisms in common. We report these fractions separately for pairs in which one of the investors invokes exactly one mechanism, two mechanisms, three mechanisms, or more than three mechanisms. When one investor invokes only one mechanism, the fraction of pairings in which the other investor relies on entirely different mechanisms is 77%. When the focal investor invokes two mechanisms (three mechanisms), the corresponding fractions are 56% (41%). Even when an investor invokes four or more mechanisms, the fraction of times that the other investor invokes none of these four or more mechanisms is 29%.

Taken together, these results point to substantial variation in mechanisms invoked not only within investors but also across investors. A key implication for future research is the importance of incorporating such heterogeneity into theoretical models by allowing for different investors to rely on different mechanisms and have them interact with one another.

3.4 Variation in Mechanisms Across Countries and Investor Demographics

Figure 6 reports the prevalence of the thirteen mechanisms separately for each of the ten countries in our sample. Each radar chart arranges the mechanisms radially, with the distance from the center representing the proportion of investor transcripts tied to a given mechanism. As before, the dark-gray shading indicates the share of transcripts where a particular mechanism is deemed the "primary driver" to the investor's decision. The light-gray shading represents the share of transcripts where a particular mechanism is deemed a "meaningful contributor."

[Figure 6 here]

The radar charts reveal pronounced cross-country heterogeneity in the prevalence of mechanisms. Visual inspection shows that the overall shapes of the charts differ markedly across countries. For instance, India exhibits a pronounced spike in the *Fundamental Strength* mechanism, reflecting the exceptionally high prevalence of this mechanism among Indian investors (60.0%), whereas the US displays a visible dent in the same mechanism, consistent with its relatively low prevalence (30.0%). Japan shows a notable spike in the *Familiarity/Brand Affect* mechanism, reflecting the fact that this mechanism is a "meaningful contributor" in nearly half of the Japanese transcripts (43.6%). In contrast, countries such as Singapore, where this mechanism plays a less prominent role (19.3%), exhibit no such spike. While these comparisons are based on mean differences, our sample is stratified to ensure broadly comparable demographic compositions across countries in terms of age, gender, education, and wealth.

To formalize these visual patterns, we conduct a hierarchical clustering analysis based on country-level mechanism prevalence. For each country, we construct a vector of length $N = 13$, with the proportion of transcripts in which the mechanism features prominently, based on a score of at least three. These vectors are standardized and clustered using complete linkage with Euclidean distance as the similarity metric. This approach aligns naturally with the visual notion of proximity in the radar charts.

The algorithm begins with ten individual countries and iteratively merges the two most similar clusters until all observations are combined.[6]

[6] At each step, complete linkage minimizes the maximum pairwise distance between elements of the candidate clusters, thereby ensuring compactness in the resulting groupings.

[Figure 7 here]

Figure 7 Panel (a) displays the resulting dendrogram. The vertical axis measures Euclidean distance between clusters, with greater height indicating greater dissimilarity. Merges at lower heights reflect more similar countries: for example, Australia and Canada merge early (distance ≈ 2.7), whereas Japan remains distinct until the final merge (distance ≈ 8.1). Applying a linkage distance threshold of 6.0 yields three clusters. The first includes South Korea, India, and Singapore, characterized by a high prevalence of fundamentals and growth-oriented mechanisms. The second consists solely of Japan, distinguished by elevated prevalence of *Dividends* and *Familiarity/Brand Affect*. The third encompasses Australia, Canada, France, Germany, the UK, and the US, which exhibit relatively balanced mechanism profiles.

The dendrogram results suggest that investors in South Korea, India, and Singapore share broadly similar behavioral tendencies that differ from those observed in the other seven countries. Likewise, investors in Australia, Canada, France, Germany, the UK, and the US display comparable behavioral profiles. Japanese investors appear distinct from all other groups.

To assess whether these clusters can be traced to differences in culture, we conduct a parallel clustering analysis using Hofstede's six cultural dimensions: power distance, individualism, masculinity, uncertainty avoidance, long-term orientation, and indulgence (Hofstede, Hofstede, and Minkov, 2010). Panel (b) of Figure 7 displays the resulting dendrogram based on standardized Hofstede scores.

The culture-based clustering in Panel (b) yields groupings that broadly align with the mechanism-based clusters in Panel (a). India, Singapore, and South Korea cluster together, reflecting relatively high power distance (77, 74, and 60, respectively) and low individualism (24, 43, and 58). Japan remains culturally distinct, with exceptionally high uncertainty avoidance (92) and long-term orientation (100). Most Western countries—Australia, Canada, Germany, the UK, and the US—cluster together, characterized by high individualism (73, 72, 79, 76, and 60) and relatively low power distance (38, 39, 35, 35, and 40). The exception is France, which joins the India–Singapore–South Korea cluster in the culture-based dendrogram due to its relatively high power distance (68) and uncertainty avoidance (86).

Taken together, the mechanism- and culture-based clusters broadly align, suggesting that cross-country differences in how investors draw on mechanisms may have cultural roots. This interpretation is consistent with prior work, such as Chui, Titman, and Wei (2010), who find evidence that certain behavioral biases vary systematically with cultural characteristics.

A key implication of our finding for future research is that insights drawn from investor behavior in a particular country, such as the US, and the theories developed from those insights may generalize to other markets within the same cluster, but not necessarily to markets in different clusters. Our findings thus provide guidance on when it may be worthwhile to revisit a research question previously examined using data from one country with data from another country.

We next examine variation across investor demographics. Investors are partitioned into two groups along each of six dimensions: gender, age, education, wealth, self-reported investment knowledge, and self-reported risk attitude. For each subsample, we repeat the same procedure applied to the individual countries, thereby generating a total of twelve radar charts.

[Figure 8 here]

Figure 8 reveals some demographic variation. Female investors exhibit higher prevalence of *Authority-Follow*, *Blue-Chip Comfort*, and *Social-Copy*, whereas male investors show higher prevalence of *Fundamental Strength*, *Growth/Innovation*, and *Valuation/Mispricing*. These patterns are consistent with prior evidence that women tend to display greater risk aversion (Barber and Odean, 2001). High-wealth investors show somewhat higher prevalence of *Authority-Follow* and lower prevalence of *Social-Copy*. Beyond these patterns though, differences across age, education, knowledge, and risk attitude are limited.

In robustness tests, we estimate logistic regressions with investor demographics and country fixed effects. This approach mitigates two potential limitations of the radar chart comparisons. First, logistic regressions account for differences in the unconditional prevalence of mechanisms (e.g., *Fundamental Strength* appears in 41% of transcripts, whereas *Technical Analysis* appears in only 7%). They do so by analyzing odds ratios rather than raw proportions. Second, by including both demographic indicators and country fixed effects within a single regression specification, we can isolate the influence of each characteristic while holding others constant. As shown in Online Appendix Figure OA2, the results reinforce our key take-away from this analysis: There are strong cross-country differences. Demographic effects are relatively modest.

4. Comparison with other Empirical Methods to Study Investor Decision Making, Limitations, Considerations and Future Research Directions

Before concluding, we discuss the relative merits and limitations of AI-driven interviews compared to other methods to examine investor decision-making. We also outline potential directions for future research.

4.1 Investor Surveys

Investor surveys have become a popular tool to better understand investor decision-making. Choi and Robertson (2020) survey US households and ask what factors broached by the academic literature determine their equity allocations. Bender, Choi, Dyson, and Robertson (2022) focus on US individuals with at least $1 million in investable assets and use a survey to examine how well leading academic theories explain the beliefs and choices of "millionaires." Liu, Peng, Xiong, and Xiong (2022) survey Chinese retail investors and link their responses to trading records from the Shenzhen Stock Exchange to assess the motives behind trading activity and compare competing behavioral theories.

All of the above studies rely on closed-ended surveys. The comparative advantage of interviews is that they are dynamic and adaptive. An interviewer can probe for additional detail and context (Lamont, 2009; Kaplan, 2008) and solicit concrete examples. These adaptive clarifying questions can potentially generate richer information.

Well-conducted interviews are also inherently "bottom-up" and non-leading: the respondent, rather than the researcher, drives the conversation and provides nearly all of the verbal input. This structure mitigates common concerns associated with surveys, especially those involving the wording and framing of questions and answer choices (McCracken and McCracken, 1988; Schuman and Presser, 1996; Kvale and Brinkmann, 2009; Patton, 2014). It can also surface aspects of investor behavior that the academic literature has not yet fully recognized and, therefore, could not have incorporated as predefined options in a multiple-choice setting.

At the same time, closed-ended survey questions are generally more "top-down" and rooted in the theoretical frameworks of interest to the survey designer. This feature can produce responses that are more standardized, comparable, and useful for adjudicating among competing theories.

A separate literature uses open-ended (open-text-field) surveys to elicit respondents' reasoning (e.g., Almås, Cappelen, Sørensen, and Tungodden, 2024; Stantcheva, 2021). These essay-type survey questions are neither dynamic nor adaptive. It is therefore unsurprising that Chopra and Haaland (2023) and Geiecke and Jaravel (2024) find that AI-led interviews generate a substantially richer understanding of individuals' motivations than open-ended surveys. Open-ended online surveys also face a growing challenge: nowadays, many respondents can use AI writing tools such as ChatGPT, Claude, or Gemini to quickly create answers and copy and paste them into the online survey. Respondents have far less incentive to use AI tools during interactive interviews, which involve shorter responses to a series of questions.

Future research could combine the interview approach with closed-end survey questions. In the first step, researchers could conduct bottom-up field interviews to identify the key considerations that are on investors' minds. Once these considerations have been uncovered, the second step could involve a top-down, theory-driven survey designed to more precisely understand investor behavior and the underlying drivers of that behavior.

For example, one limitation of our study is that we do not know whether the investors who invoke multiple mechanisms (i) do so simultaneously for all of their stock selections, (ii) rely primarily on one mechanism for some of their holdings, while invoking another for their other positions, or (iii) switch between mechanisms over time. Now that it is clear that nearly all investors in our sample invoke multiple mechanisms, one could construct a follow-up survey listing all of the here-identified thirteen mechanisms, ask investors to select the mechanisms they rely on, and then query them explicitly on how they combine them: option (i), (ii), or (iii). One could also embed a brief, targeted follow-up interview to probe why investors follow a particular strategy and what is leading them to choose specific combinations of mechanisms.

Another important limitation of our paper is that our mechanism descriptions are broad and do not (yet) allow for clean differentiation between theories. For instance, as alluded to in Section 3.1, the *Growth/Innovation* mechanism is consistent both with behavioral theories grounded in CPT and with rational theories based on displacement risk.

One could again follow up with a survey. If participants identify themselves as invoking the *Growth/Innovation* mechanism, researchers could administer a theory-driven survey, as in Choi and Robertson (2020) and Bender, Choi, Dyson, and Robertson (2022), to pin down the underlying drivers of this mechanism and determine whether they align more closely with one theoretical framework than another. Such a hybrid approach would leverage the strengths of both interviews and surveys and could substantially deepen our understanding of investor decision-making.

4.2 Textual Analysis on Already Existing Text

Another emerging approach in the literature for understanding investor behavior and beliefs is textual analysis on already existing text. Ke (2024) and Bastianello, H. Décaire, and Guenzel (2024) apply LLMs to analyst reports to examine which topics influence analysts' earnings forecasts and price targets. They find that analysts are primarily concerned about sales, costs, margins, corporate management, financial conditions, and macroeconomic trends, and that these topics differentially relate to their forecasts. Décaire, Sosyura, and Wittry (2024) apply a similar methodology to study how analysts determine discount rates when valuing future cash flows.

More relevant to our study, Chen, Hwang, and Peng (2025) examine analysts' and bloggers' buy recommendations for a special subset of stocks, namely those that reside in the short leg of anomalies. Prior literature shows that these stocks trade at comparatively high prices and earn low future returns. Chen, Hwang, and Peng (2025) conduct a textual analysis of analyst reports and investment blogs, which recommend that investors buy these "short-leg securities," to examine what beliefs and preferences draw investors to this particular set of stocks despite their well documented poor average performance.

The key strength of textual analysis on already existing text lies in the abundance of written expressions that are publicly accessible. For instance, while our study parses through 1,540 interview transcripts, Chen, Hwang, and Peng (2025) analyze roughly 1.7 million analyst reports and 140,000 blogs written over a fifteen-year period. The long historical record of written expressions also allows future research to explore the evolution of investor beliefs and preferences over time and to shed light on what investors were thinking at specific historical junctures, an analysis that AI-driven interviews cannot perform.

Conversely, the use of field interviews allows researchers to conduct follow-up probing and clarify ambiguous or unclear statements. This interactive nature allows for cleaner inferences than static textual analysis. In the end, it appears to us that both empirical strategies have their comparative advantages and can be useful tools for future research.

4.3 Human-to-Human Interviews

To the best of our knowledge, the only finance paper conducting human-to-human interviews with investors is Duraj, Grunow, Haliassos, Laudenbach, and Siegel (2024). Without mentioning investing upfront, Duraj, Grunow, Haliassos, Laudenbach, and Siegel (2024) elicit open-ended reflections on "money," allowing respondents to bring up investing on their own. The authors code the interview transcripts using a detailed scheme and conclude that a widespread misunderstanding, namely that successful equity investing requires frequent monitoring, significantly inflates perceived entry and participation costs and lowers stock market participation.

Human-to-human interviews excel at probing respondents' thought processes and uncovering previously unrecognized factors. However, they are costly to conduct and difficult to scale. These challenges can at least be partially addressed by AI-to-human interviews.

4.4 Variations of AI-to-Human Interviews

Our participants already spent an average of 23 minutes in our study, and there are natural limits to how deeply we can probe without exceeding our research budget or exhausting participants' attention spans. Future research could attempt to make the interviews verbal rather than written. A verbal format would allow for more frequent back-and-forth within a fixed time window and could yield richer, more nuanced information. A verbal format could have mitigated some of the limitations noted above. For example, a more frequent back-and-forth could have allowed us to probe more deeply and examine whether investors who invoke multiple mechanisms (i) rely on them simultaneously for all of their stock selections, (ii) use different mechanisms for different holdings, or (iii) switch between mechanisms over time.

Technologically, a switch from written to verbal is not difficult. But it also raises new complications. For example, when interviewing a female investor, should the interviewer's voice be female or male? Should it sound young or old? What should the general tone of the interviewer be? Future research must take care to avoid introducing excessive noise, or, worse, systematic bias, into the data-generating process.

Finally, rather than conducting AI-to-human interviews, future research could also take the next step and explore AI-to-AI interviews. Such an approach would be even less costly and could feasibly support much larger samples and much longer, detailed interviews. At the same time, it remains unclear whether the responses of AI agents to an AI interviewer would resemble the responses of human investors to an AI interviewer.

4.5 Other Limitations, Considerations and Future Research Directions

Our interview is designed to shed light on why investors choose to *purchase* certain stocks rather than others. A further limitation of our study is that we do not examine the factors that investors consider when deciding which stocks to *sell*. We elected not to probe selling decisions beyond the brief question reported in (2.d) of Section 2 ("After you sell a stock, do you continue to track its performance and how its price evolves?") because, again, doing so would have lengthened the interview substantially and exceeded the time allocated to us by CoreData Research under our contractual budget. We were also concerned that a longer interview might reduce respondent patience and engagement. Future work could examine investors' selling decisions.

Responses to question (2.d) indicate that a large majority of investors (87%) continue to at least sometimes track prices after selling, with about 31% reporting that they always do so (Online Appendix Figure OA3). This pattern points to another promising avenue for future research: examining how post-sale behavior and the post-sale performance of previously held stocks shape subsequent investment decisions. The fact that most investors at least occassionally continue to monitor prices even after selling suggests that their experiences are informed not only by realized returns, but also by hypothetical returns had they not sold.

Our paper suffers from at least one more shortcoming. The *Social-Copy* mechanism and, to some extent, the *Authority-Follow* mechanism are consistent with the social finance framework. Yet these mechanisms likely understate the full importance of social interactions. Many of the mechanisms we document did not arise in a vacuum; rather, they were likely transmitted and reinforced through social channels, much in the spirit of the narrative dynamics emphasized by Shiller (2017).

A parallel observation applies to the literature on experience effects (Malmendier and Nagel, 2011, 2016; Malmendier, Nagel, and Yan, 2021). None of our mechanisms explicitly reference past investment experiences, yet investors' histories may plausibly shape whether they adopt or reject particular mechanisms. For example, Japanese investors in our sample rarely invoke the *Momentum* mechanism. One possible explanation is that momentum effects are relatively weak in Japan, so investors' personal histories may reinforce the belief that momentum-based strategies are ineffective.

These observations point to a broader limitation of our study. While we are able to document *which* mechanisms investors invoke, time and resource constraints prevent us from probing further and uncovering *why* particular mechanisms are adopted (or, not adopted). Understanding these origins represents an important direction for future research as they can help us further understand which theoretical frameworks are the most relevant in explaining the behaviors of investors and financial markets.

5. Conclusion

The question of what drives the behavior of investors and ultimately that of financial markets lies at the heart of finance and has been the focus of extensive research (e.g., Fama and French, 1992; Barberis and Thaler, 2003). Our paper revisits this classic question with a new method. We conduct AI-driven field interviews with 1,540 actual investors from ten countries and a wide range of investor demographics.

Our paper contributes to the literature by documenting, comprehensively and systematically, what actually motivates investors in their decisions. Our results show that investor stock selection is organized around thirteen recurrent mechanisms. Several of the most frequently invoked mechanisms are only partially captured by current mainstream asset-pricing theories, pointing to the need for theoretical refinement and extension. Our findings also suggest that future work should accommodate the substantial heterogeneity both within and across investors, for example, by allowing interactions among groups with distinct preferences and belief-formation processes. Finally, we provide a systematic comparison of investor behavior across countries.

More broadly, our study advocates AI-driven field interviews as a complementary empirical tool to study investor behavior. With continued advances in AI, this approach offers valuable opportunities to gain deeper insights into investor decision-making and the dynamics of financial markets. Our paper outlines the methods that future research can use not only to conduct such interviews but also to analyze the resulting transcripts.

References

Almås, Ingvild, Alexander W Cappelen, Erik Ø Sørensen, and Bertil Tungodden, 2024, Attitudes to inequality: Preferences and beliefs, *Oxford Open Economics* 3, i64–i79.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *The Journal of Finance* 59, 1481–1509.

Barber, Brad M., and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *The Journal of Finance* 55, 773–806.

——— , 2001, Boys will be boys: Gender, overconfidence, and common stock investment, *The Quarterly Journal of Economics* 116, 261–292.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2015, X-CAPM: An extrapolative capital asset pricing model, *Journal of Financial Economics* 115, 1–24.

——— , 2018, Extrapolation and bubbles, *Journal of Financial Economics* 129, 203–227.

Barberis, Nicholas, and Ming Huang, 2008, Stocks as lotteries: The implications of probability weighting for security prices, *American Economic Review* 98, 2066–2100.

Barberis, Nicholas, Lawrence J. Jin, and Baolian Wang, 2021, Prospect Theory and Stock Market Anomalies, *The Journal of Finance* 76, 2639–2687.

Barberis, Nicholas, and Andrei Shleifer, 2003, Style investing, *Journal of Financial Economics* 68, 161–199.

——— , and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–343.

Barberis, Nicholas, and Richard Thaler, 2003, A survey of behavioral finance, in *Handbook of the Economics of Finance*, vol. 1 of *Financial Markets and Asset Pricing* . pp. 1053–1128 (Elsevier).

Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *The Quarterly Journal of Economics* 121, 823–866.

Bastianello, Francesca, Paul H. Décaire, and Marius Guenzel, 2024, Mental Models and Financial Forecasts, *SSRN Working Paper*.

Bender, Svetlana, James J. Choi, Danielle Dyson, and Adriana Z. Robertson, 2022, Millionaires speak: What drives their personal investment decisions?, *Journal of Financial Economics* 146, 305–330.

Bhagwat, Vineet, J Anthony Cookson, Chukwuma Dim, and Marina Niessner, 2025, The market's mirror: Revealing investor disagreement with LLMs, *FEB-RN Research Paper*.

Breeden, Douglas T., 1979, An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–296.

Brown, Jeffrey R., Zoran Ivković, Paul A. Smith, and Scott Weisbenner, 2008, Neighbors Matter: Causal Community Effects and Stock Market Participation, *The Journal of Finance* 63, 1509–1531.

Bybee, J Leland, 2023, The ghost in the machine: Generating beliefs with large language models, *arXiv preprint arXiv:2305.02823* pp. 349–389.

Centola, Damon, 2010, The spread of behavior in an online social network experiment, *Science* 329, 1194–1197.

——— , and Michael Macy, 2007, Complex contagions and the weakness of long ties, *American Journal of Sociology* 113, 702–734.

Chen, Hailiang, and Byoung-Hyoun Hwang, 2022, Listening in on investors' thoughts and conversations, *Journal of Financial Economics* 145, 426–444.

——— , and Zhuozhen Peng, 2025, Why Do Investors Like Short-leg Securities? Evidence from a Textual Analysis of Buy Recommendations, *The Review of Financial Studies* p. Advance online publication.

Chinco, Alex, Samuel M. Hartzmark, and Abigail B. Sussman, 2022, A New Test of Risk Factor Relevance, *The Journal of Finance* 77, 2183–2238.

Choi, James J., and Adriana Z. Robertson, 2020, What Matters to Individual Investors? Evidence from the Horse's Mouth, *The Journal of Finance* 75, 1965–2020.

Chopra, Felix, and Ingar Haaland, 2023, Conducting Qualitative Interviews with AI, *SSRN Working Paper*.

Chui, Andy CW, Sheridan Titman, and KC John Wei, 2010, Individualism and momentum around the world, *The Journal of Finance* 65, 361–392.

Décaire, Paul, Denis Sosyura, and Michael D. Wittry, 2024, Resolving Estimation Ambiguity, *SSRN Working Paper*.

Dichev, Ilia, John Graham, Campbell R Harvey, and Shiva Rajgopal, 2016, The misrepresentation of earnings, *Financial Analysts Journal* 72, 22–35.

Duraj, Kamila, Daniela Grunow, Michael Haliassos, Christine Laudenbach, and Stephan Siegel, 2024, Rethinking the Stock Market Participation Puzzle: A Qualitative Approach, *SSRN Working Paper*.

Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *The Journal of Finance* 47, 427–465.

Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas, 2012, Displacement risk and asset returns, *Journal of Financial Economics* 105, 491–510.

Geiecke, Friedrich, and Xavier Jaravel, 2024, Conversations at Scale: Robust AI-led Interviews with a Simple Open-Source Platform, *SSRN Working Paper*.

Gelman, Michael, David Hirshleifer, Yaron Levi, and Liron Reiter Gavish, 2024, Social Interaction Intensity and Investor Behavior, *SSRN Working Paper*.

Graham, John R, and Alok Kumar, 2006, Do dividend clienteles exist? evidence on dividend preferences of retail investors, *The Journal of Finance* 61, 1305–1336.

Han, Bing, David Hirshleifer, and Johan Walden, 2022, Social transmission bias and investor behavior, *Journal of Financial and Quantitative Analysis* 57, 390–412.

Harkins, Stephen G, and Richard E Petty, 1987, Information utility and the multiple source effect., *Journal of Personality and Social Psychology* 52, 260.

Hartzmark, Samuel M, and David H Solomon, 2019, The dividend disconnect, *The Journal of Finance* 74, 2153–2199.

Henning, Thomas, Siddhartha M. Ojha, Ross Spoon, Jiatong Han, and Colin F. Camerer, 2025, Llm agents do not replicate human market traders: Evidence from experimental finance, *arXiv Working Paper*.

Hirshleifer, David, 2020, Presidential address: Social transmission bias in economics and finance, *The Journal of Finance* 75, 1779–1831.

Hofstede, G., G.J. Hofstede, and M. Minkov, 2010, *Cultures and Organizations: Software of the Mind, Third Edition* (McGraw Hill LLC).

Hong, Harrison, Jeremy C Stein, and Jialin Yu, 2007, Simple forecasts and paradigm shifts, *The Journal of Finance* 62, 1207–1242.

Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting anomalies: An investment approach, *The Review of Financial Studies* 28, 650–705.

Huang, Shiyang, Byoung-Hyoun Hwang, and Dong Lou, 2021, The rate of communication, *Journal of Financial Economics* 141, 533–550.

Kaplan, Sarah, 2008, Framing contests: Strategy making under uncertainty, *Organization Science* 19, 729–752.

Ke, Shikun, 2024, Analysts' Belief Formation in Their Own Words, *SSRN Working Paper*.

Kvale, Steinar, and Svend Brinkmann, 2009, *Interviews: Learning the craft of qualitative research interviewing* (sage).

Lamont, Michèle, 2009, *The dignity of working men: Morality and the boundaries of race, class, and immigration* (Harvard University Press).

Lee, Charles M.C., and Bhaskaran Swaminathan, 2000, Price momentum and trading volume, *The Journal of Finance* 55, 2017–2069.

Liu, Hongqi, Cameron Peng, Wei A. Xiong, and Wei Xiong, 2022, Taming the bias zoo, *Journal of Financial Economics* 143, 716–741.

Lopez-Lira, Alejandro, 2025, Can Large Language Models Trade? Testing Financial Theories with LLM Agents in Market Simulations, *arXiv Working Paper*.

Lucas, Robert E., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.

Madaan, Aman, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al., 2023, Self-refine: Iterative refinement with self-feedback, *Advances in Neural Information Processing Systems* 36, 46534–46594.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: Do macroeconomic experiences affect risk taking?, *The quarterly journal of economics* 126, 373–416.

——— , 2016, Learning from inflation experiences, *The Quarterly Journal of Economics* 131, 53–87.

——— , and Zhen Yan, 2021, The making of hawks and doves, *Journal of Monetary Economics* 117, 19–42.

McCracken, Grant David, and Grant McCracken, 1988, *The long interview* (Sage).

Miller, Merton H, and Franco Modigliani, 1961, Dividend policy, growth, and the valuation of shares, *the Journal of Business* 34, 411–433.

Novy-Marx, Robert, 2013, The other side of value: The gross profitability premium, *Journal of Financial Economics* 108, 1–28.

Patton, Michael Quinn, 2014, *Qualitative research & evaluation methods: Integrating theory and practice* (Sage publications).

Pearl, Judea, 1995, *Probabilistic reasoning in intelligent systems: networks of plausible inference* (Elsevier).

———, 2009, *Causality* (Cambridge university press).

Schuman, Howard, and Stanley Presser, 1996, *Questions and answers in attitude surveys: Experiments on question form, wording, and context* (Sage).

Shefrin, Hersh M, and Meir Statman, 1984, Explaining investor preference for cash dividends, *Journal of Financial Economics* 13, 253–282.

Shiller, Robert J., 2017, Narrative economics, *American Economic Review* 107, 967–1004.

Stantcheva, Stefanie, 2021, Understanding Tax Policy: How do People Reason?, *The Quarterly Journal of Economics* 136, 2309–2369.

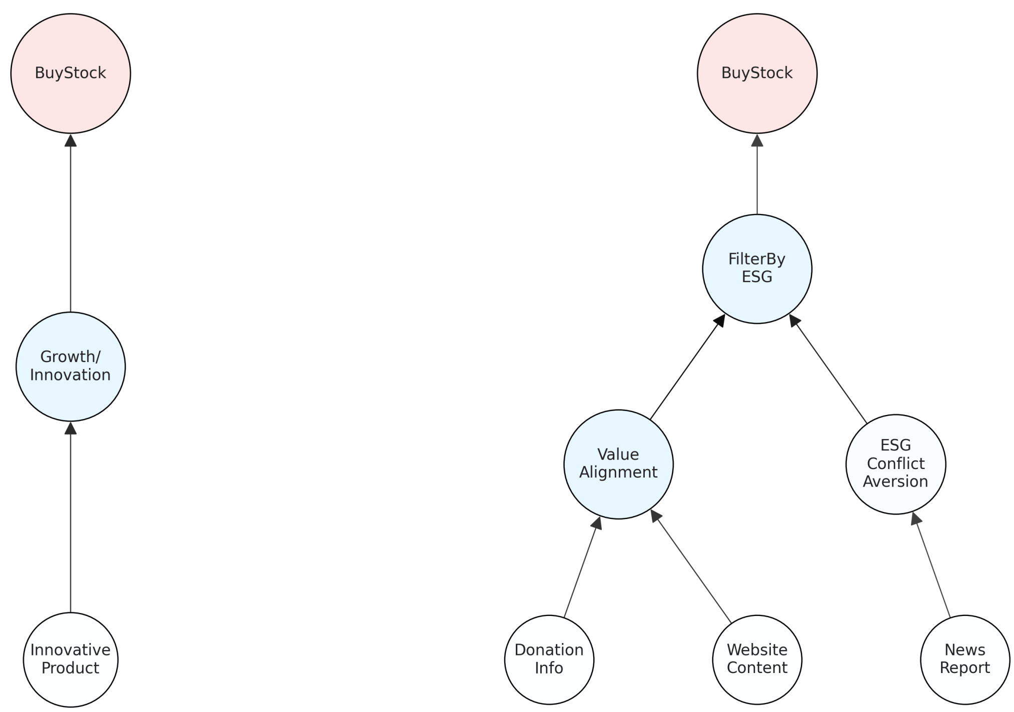

(a) Example 1

(b) Example 2

Figure 1: DAG Examples

This figure illustrates two example investor-level DAGs. Panel (a) depicts an investor whose stock purchases are driven by firms' innovative products, and Panel (b) depicts an investor who buys stocks by filtering on ESG criteria.

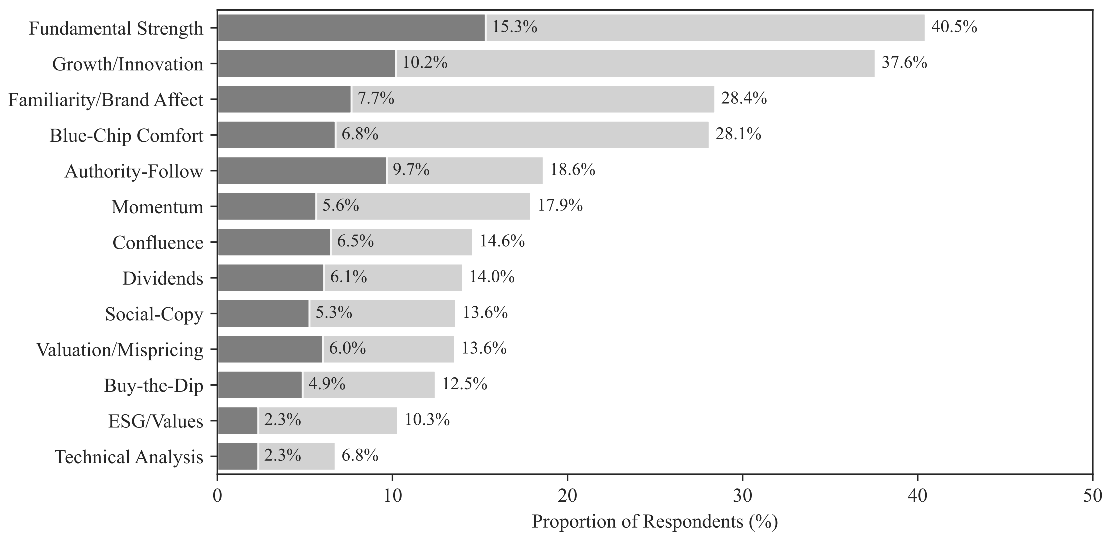

Figure 2: Proportion of Respondents Matched to Each Mechanism.

This figure displays the proportion of investor transcripts invoking each of the identified thirteen recurrent mechanism. For each investor, we check whether a given mechanism is a “meaningful contributor” or a “primary driver” to the investor’s stock-purchase decisions (as detailed in Section 2.3.3). The figure plots the fraction of investors for which a mechanism is a “meaningful contributor” in light-gray shading and the fraction for which it is a “primary driver” in dark-gray shading.

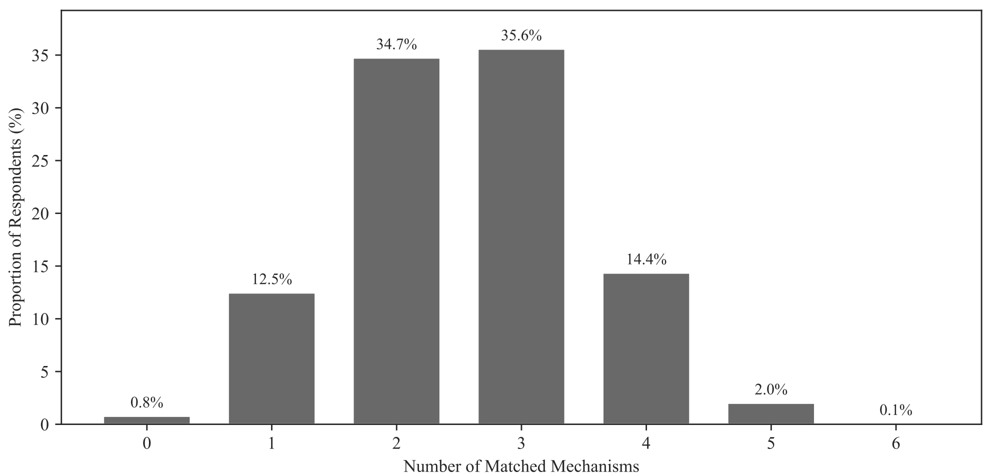

Figure 3: Number of Mechanisms Simultaneously Invoked by Same Investor.

This figure displays the distribution of the number of mechanisms investors in our sample simultaneously invoke. For each investor, we list and count the number of the mechanisms that are "meaningful contributors" to the investor's stock-purchase decisions (as detailed in Section 2.3.3). We then plot the distribution of the counts.

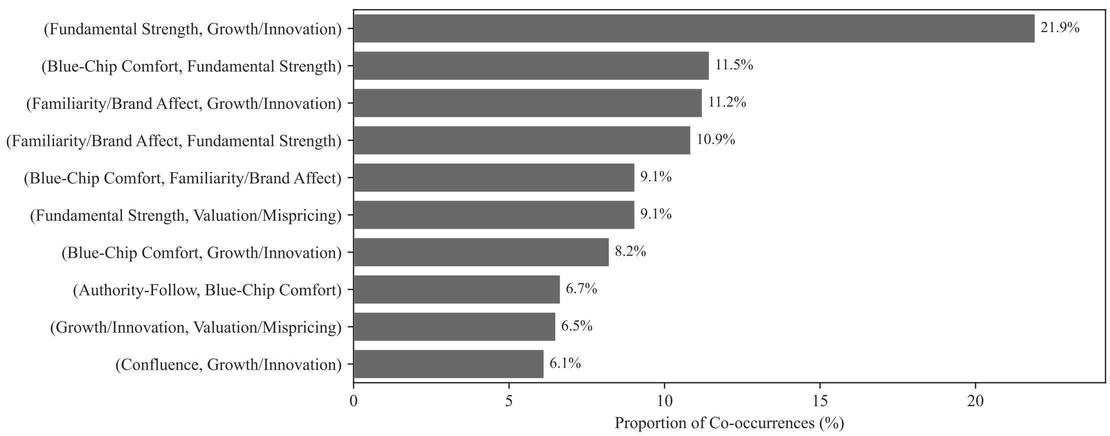

Figure 4: Mechanism Pairs Most Frequently Invoked by Same Investor.

This figure displays the ten most frequent mechanism pairs. Given that there are thirteen mechanisms, there are 78 possible mechanism pairs. For each pair, we count the number of investors who simultaneously invoke both mechanisms in their interview transcripts and divide this by the total number of investors who invoke multiple mechanisms. We then list, in descending order, the ten most frequently co-occurring mechanisms.

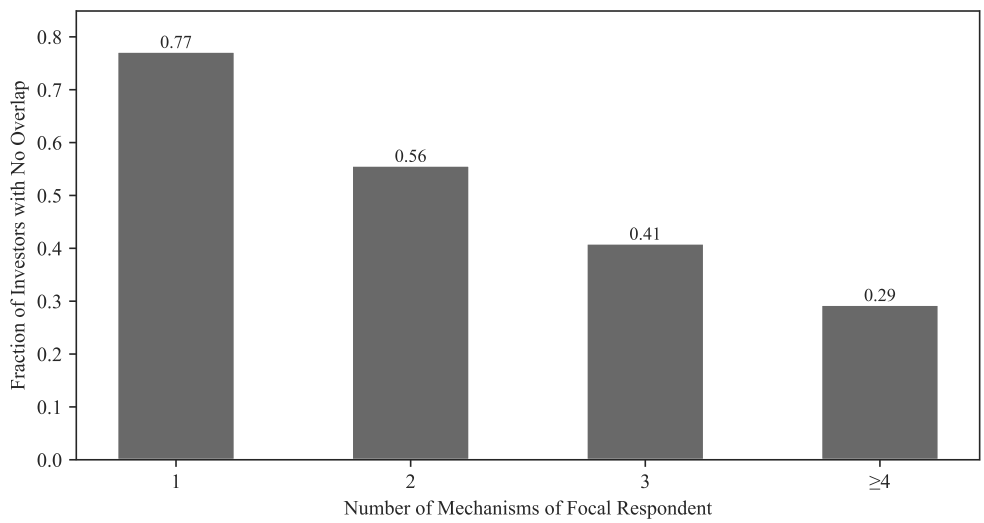

Figure 5: Fraction of Investors with Non-Overlapping Mechanisms.

This figure displays the fraction of investor pairs for which there is zero overlap in the mechanisms the two investors invoke. We consider all 1,185,030 possible pairs formed from our 1,540 investors and compute the fraction of pairs for which none of the mechanisms invoked by one investor is invoked by the other investor; put differently, the fraction of pairs where the two investors have zero mechanisms in common. We report these fractions separately for pairs in which one of the investors invokes exactly one mechanism, two mechanisms, three mechanisms, or more than three mechanisms.

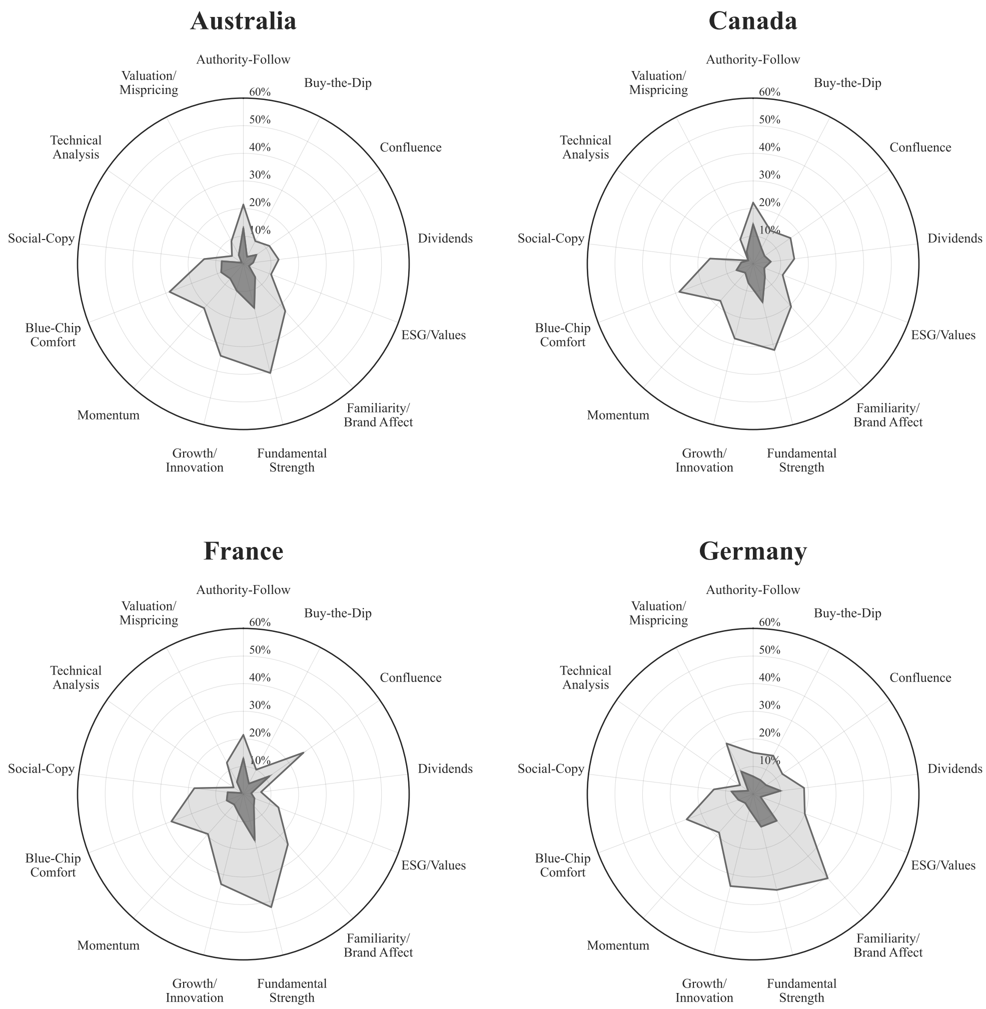

Figure 6: Mechanism Prevalence by Country.

This figure displays the proportion of investor transcripts invoking each of the thirteen identified recurrent mechanisms by country. Each panel plots the thirteen mechanisms arranged radially, with the distance from the center representing the proportion of investors to whom each mechanism is assigned. Light-gray shading indicates assignment as a "meaningful contributor," and dark-gray shading indicates assignment as a "primary driver" (as detailed in Section 2.3.3). Concentric circles mark 10-percentage-point intervals from 0% to 60%.

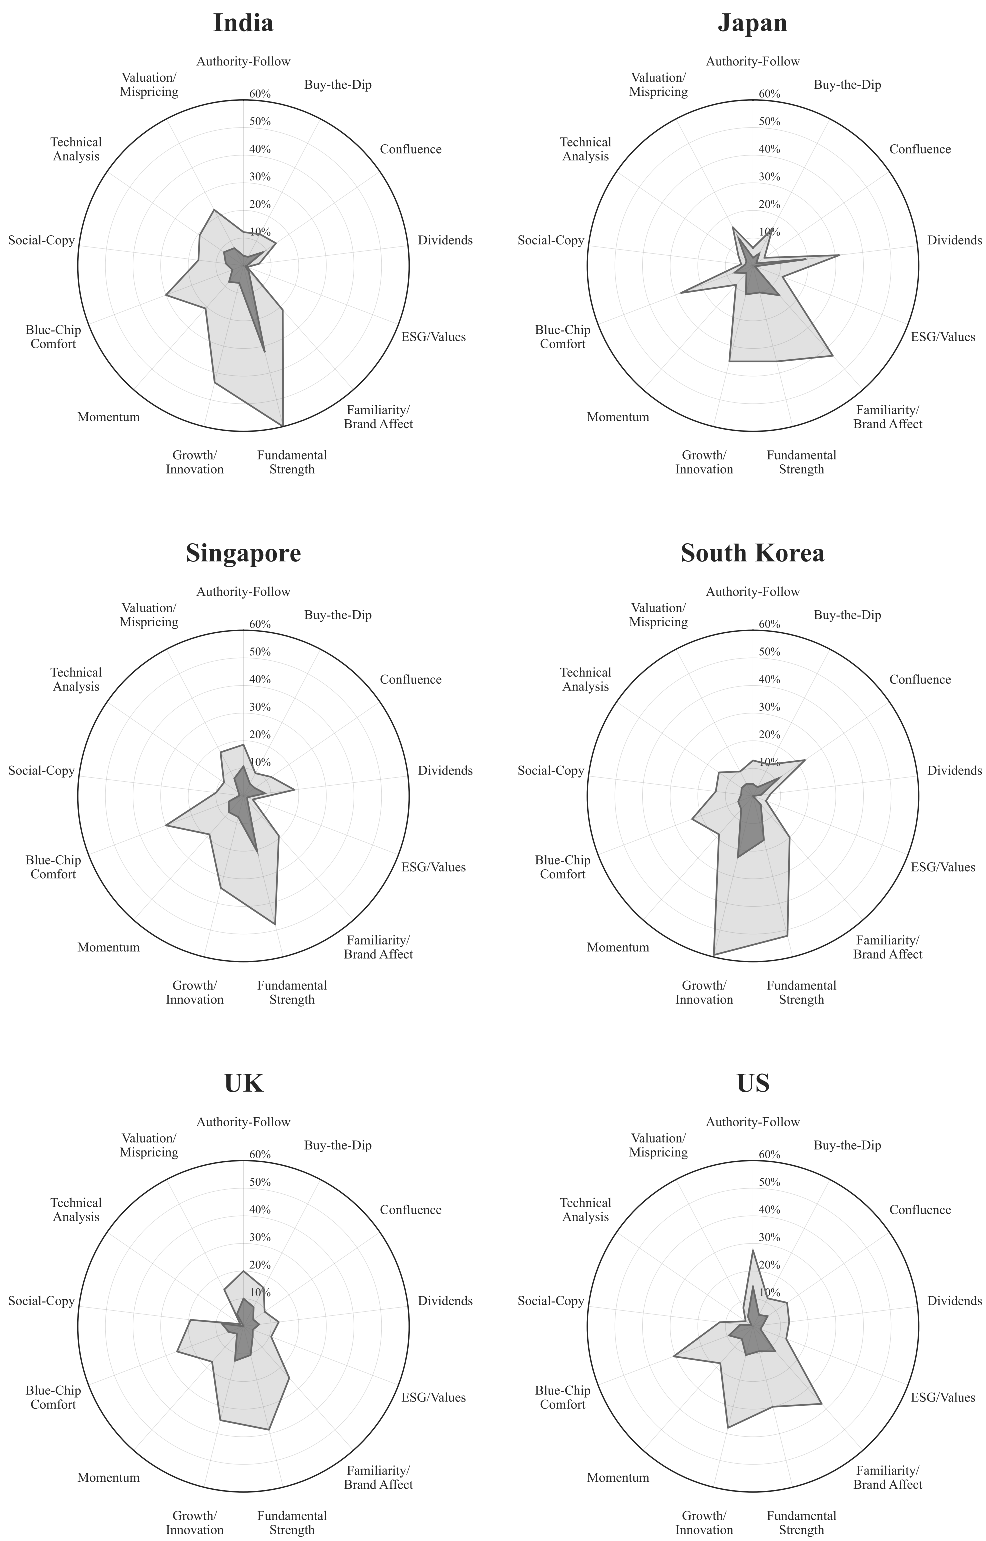

Figure 6: Mechanism Prevalence by Country (continued).

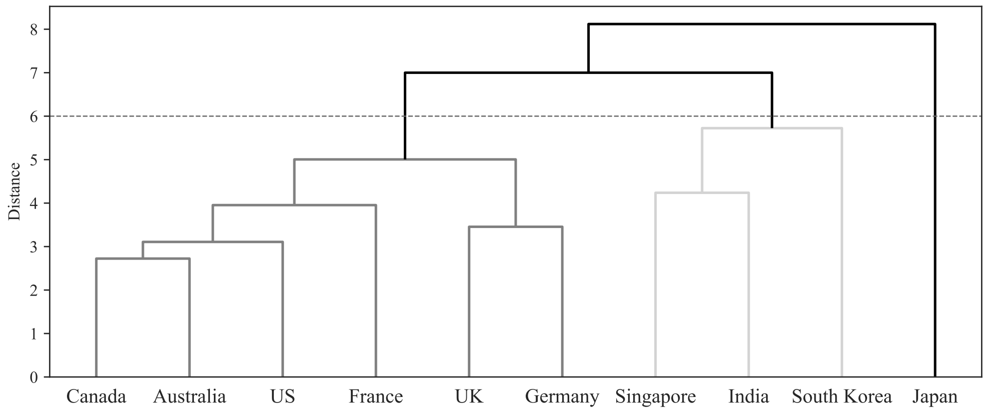

(a) By Mechanism Prevalence

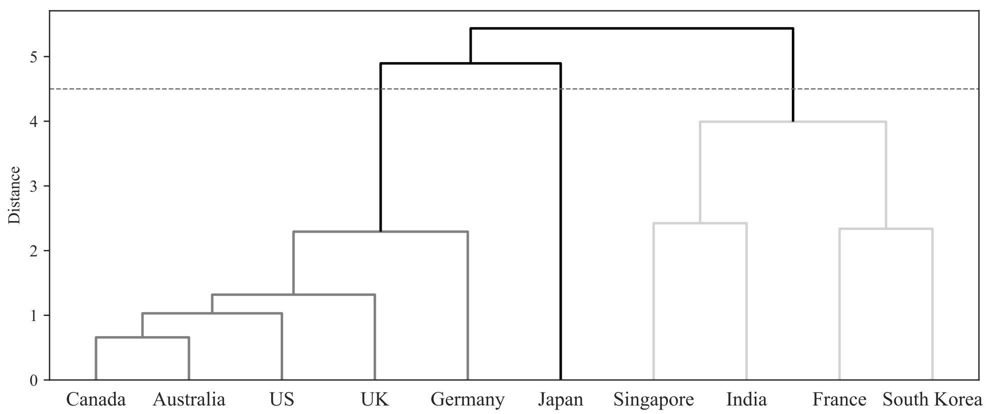

(b) By Culture Scores

Figure 7: Hierarchical Clustering Dendrograms.

This figure displays hierarchical clustering dendrograms illustrating country similarities based on two different measures. Panel (a) clusters countries based on standardized mechanism-prevalence patterns across our sample. For each country, we compute the proportions of investors for whom a particular mechanism is a "meaningful contributor" (as detailed in Section 2.3.3). We standardize these proportions and apply agglomerative clustering. Panel (b) clusters the same countries using Hofstede's culture scores. In both panels, the vertical axis measures the Euclidean distance between clusters, with greater height indicating greater dissimilarity between the grouped countries.

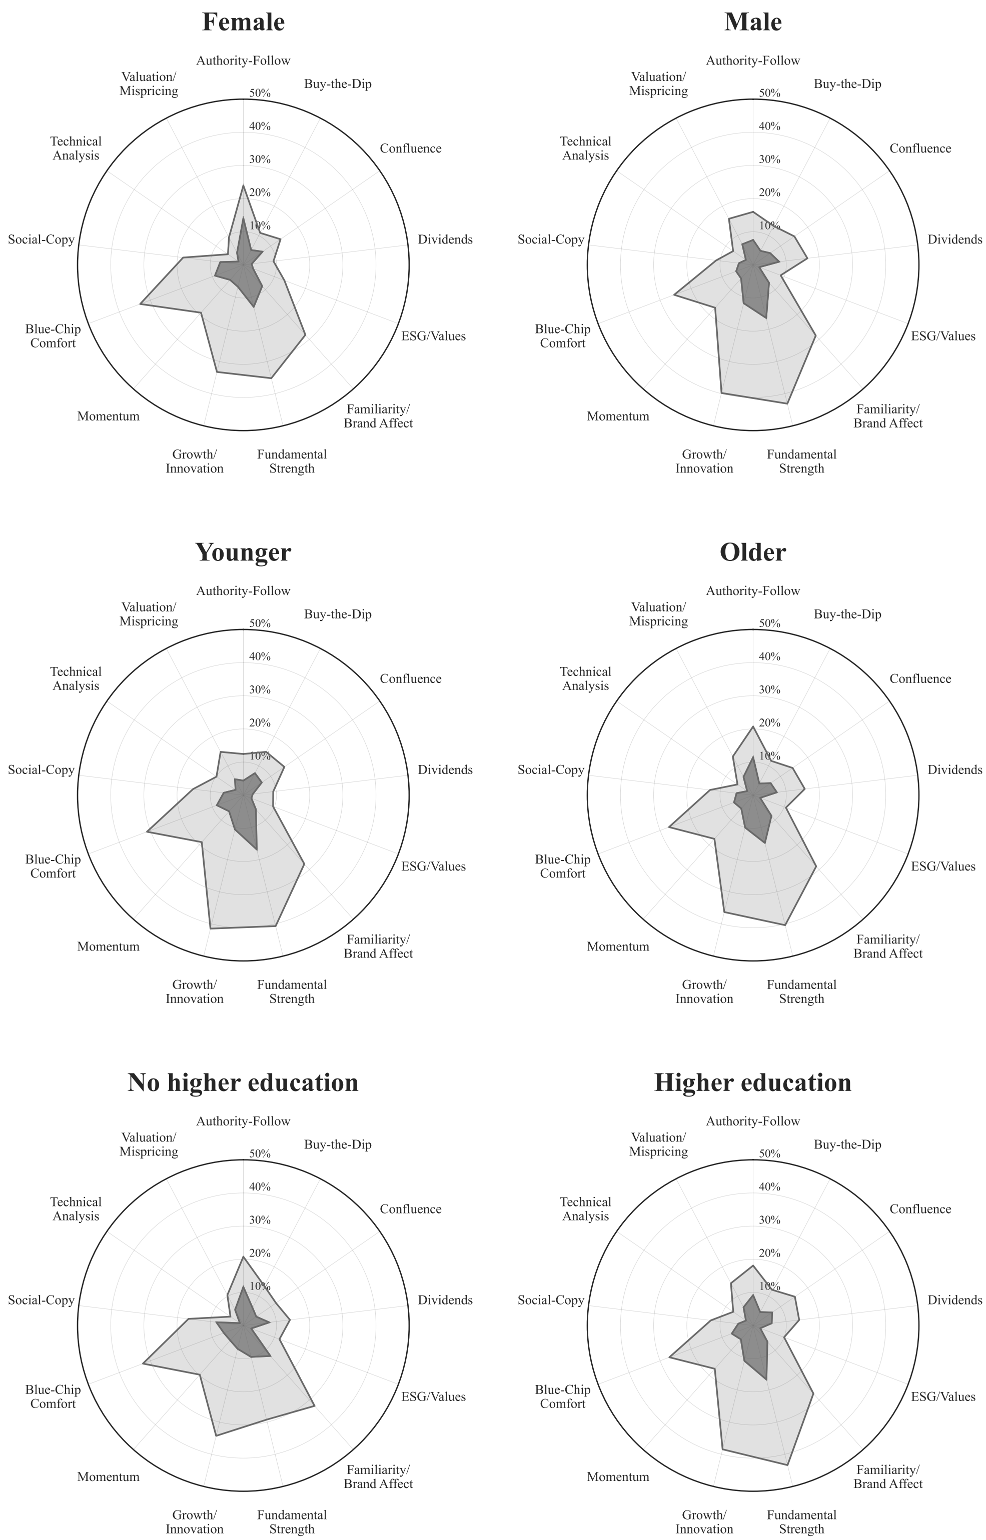

Figure 8: Mechanism Prevalence by Investor Background.

This figure displays the proportion of investor transcripts invoking each of the thirteen identified recurrent mechanisms by investor group. Each panel plots the thirteen mechanisms arranged radially, with the distance from the center representing the proportion of investors to whom each mechanism is assigned. Light-gray shading indicates assignment as a "meaningful contributor," and dark-gray shading indicates assignment as a "primary driver" (as detailed in Section 2.3.3). Concentric circles mark 10-percentage-point intervals from 0% to 60%.

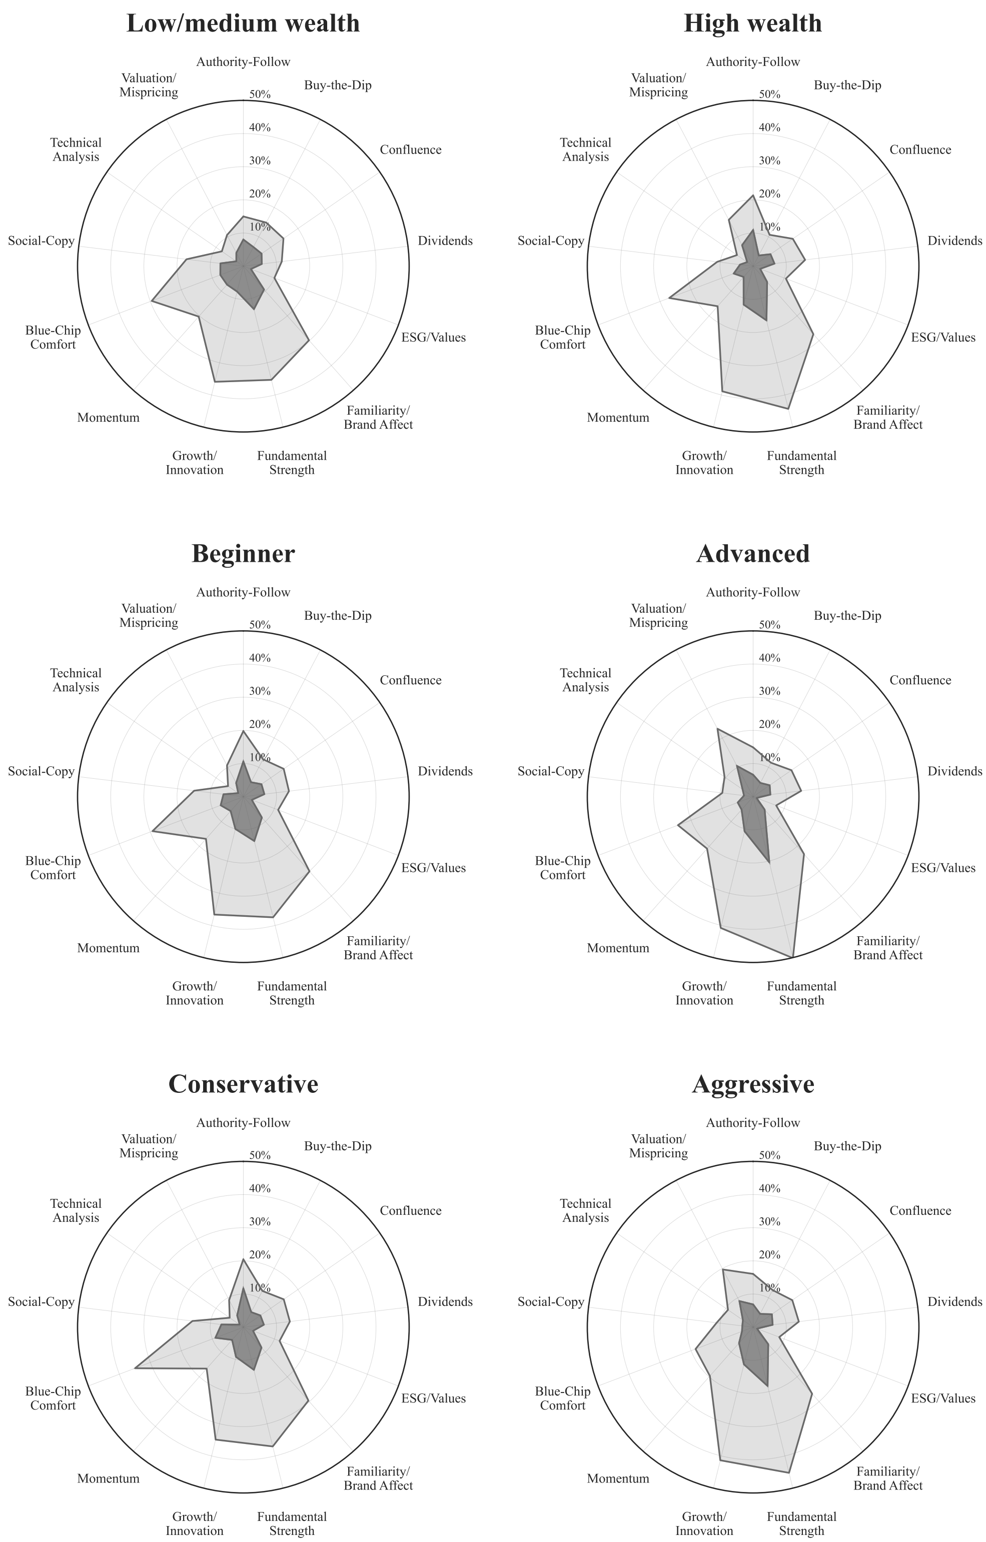

Figure 8: Mechanism Prevalence by Investor Background (continued).

Table 1: Sample Demographics.

This table reports descriptive statistics for all investors in our sample as well as by investor groups. Each row corresponds to an investor group. Columns (1) and (2) report the numbers and proportions of investors in each group. Columns (3) through (6) show the proportions of investors within each group who report owning individual stocks, mutual funds or exchange-traded funds, derivatives, and cryptocurrencies, respectively.

	N (1)	% Sample (2)	% Own Stocks (3)	% Own Funds (4)	% Own Deriv. (5)	% Own Crypto (6)
Overall	1,540	100.0	100.0	61.6	14.7	38.6
Gender						
Female	514	33.4	100.0	58.4	11.7	29.0
Male	1,025	66.6	100.0	63.2	16.3	43.4
Other	1	0.1	100.0	0.0	0.0	0.0
Age						
21-24	64	4.2	100.0	57.8	17.2	59.4
25-34	322	20.9	100.0	66.8	22.0	54.0
35-44	404	26.2	100.0	65.3	22.0	46.8
45-54	312	20.3	100.0	60.9	11.2	38.5
55-64	213	13.8	100.0	56.3	5.6	22.5
65+	225	14.6	100.0	54.2	4.0	11.1
Education						
No degree	24	1.6	100.0	50.0	16.7	8.3
School diploma	268	17.4	100.0	49.3	7.8	36.9
Bachelor's	757	49.2	100.0	61.3	13.1	38.3
Master's	397	25.8	100.0	70.3	21.2	42.1
Doctorate	68	4.4	100.0	70.6	25.0	39.7
Other	26	1.7	100.0	50.0	7.7	34.6
Investable assets						
1K-10K	142	9.2	100.0	40.8	9.9	38.0
10K-25K	123	8.0	100.0	44.7	7.3	36.6
25K-50K	123	8.0	100.0	56.1	7.3	42.3
50K-100K	100	6.5	100.0	52.0	14.0	38.0
100K-250K	105	6.8	100.0	53.3	14.3	39.0
250K-500K	67	4.4	100.0	62.7	14.9	43.3
500K-1M	440	28.6	100.0	62.7	15.5	40.5
1M+	440	28.6	100.0	77.3	20.0	35.7
Self-reported invest. know.						
Complete beginner	70	4.5	100.0	35.7	4.3	25.7
Rudimentary	345	22.4	100.0	45.5	5.8	25.5
Intermediate	762	49.5	100.0	62.5	12.5	37.5
Advanced	310	20.1	100.0	80.0	26.8	54.5
Expert	53	3.4	100.0	79.2	49.1	62.3
Self-reported risk attitude						
Conservative	87	5.6	100.0	40.2	9.2	23.0
Average	812	52.7	100.0	54.2	8.7	29.7
Aggressive	546	35.5	100.0	71.8	21.2	49.6
Very aggressive	95	6.2	100.0	85.3	33.7	65.3

Table 2: Sample Demographics by Country.

This table reports the proportions of respondents who are female, aged 35 or above, hold at least a bachelor's degree, or own investable assets of at least USD 500,000.

	No. Respondents (1)	% Female (2)	% Age 35+ (3)	% Bachelor's+ (4)	% High Assets (5)
Australia	140	35.7	66.4	75.7	60.7
Canada	140	42.9	78.6	75.7	61.4
France	140	30.0	75.7	84.3	61.4
Germany	140	22.1	67.1	67.9	60.7
India	140	30.0	50.7	97.1	62.9
Japan	140	14.3	85.7	87.1	60.7
Singapore	140	33.6	75.7	81.4	65.0
South Korea	140	28.6	81.4	87.1	62.9
UK	140	28.6	75.7	75.7	59.3
US	280	50.7	83.6	70.4	60.7
Overall	1,540	33.4	74.9	79.4	61.5

Table 3: Recurrent Mechanisms.

This table reports the thirteen recurrent mechanisms identified from the investor interview transcripts. For each mechanism, we present the GPT-generated title, its prevalence in the transcripts, defined as the fraction of transcripts in which the mechanism is a "meaningful contributor," and the GPT-generated mechanism description (all as detailed in Section 2.3).

Mechanism	Prevalence	Description
Fundamental Strength	40.5%	Investors purchase stocks when audited financial information—such as revenues, profitability, and balance-sheet strength—indicates robust current performance and attractive expected returns. This mechanism is historical data driven, grounded in verifiable accounting data.
Growth/Innovation	37.6%	Investors buy when they perceive a strong forward-looking growth runway, typically linked to product innovation, technological advantage, or structural demand expansion. The emphasis is on long-term upside potential rather than current fundamentals, and purchase decisions often precede earnings confirmation.
Familiarity/Brand Affect	28.4%	Investors act on personal familiarity, brand trust, or positive product experiences. These cues generate affective comfort and intuitive confidence, leading to rapid decisions that largely bypass detailed analysis. Diminished familiarity sharply reduces willingness to buy.
Blue-Chip Comfort	28.1%	Investors prefer large, stable firms that are perceived as safe and low risk. Signals of operating history, enduring consumer demand, and low volatility create a sense of security that can outweigh concerns about modest expected returns.
Authority-Follow	18.6%	Investors delegate stock selection to trusted professionals or platforms, such as advisors, brokers, financial experts, or investment apps. Source credibility substitutes for independent analysis, and perceived deterioration in that credibility substantially weakens buy intentions.
Momentum	17.9%	Investors buy when sustained price and volume trends reinforce beliefs in continued upward movement and predictable gains.
Confluence	14.6%	Investors require confirmation from multiple independent sources before acting. Expert opinions, online consensus, and other signals must align; conflicting information leads to inaction even when some indicators are favorable.
Dividends	14.0%	Investors purchase only when dividend yield and payout stability meet their income thresholds. A perceived decline in dividend safety sharply reduces buying even when yields rise.
Social-Copy	13.6%	Investors follow the stock choices of trusted peers or family members. Confidence derives from social trust and relational bonds rather than analytical verification, and repeated buys often trace to consistent social sources.
Valuation/Mispricing	13.6%	Investors buy when valuation metrics—such as low price-to-book ratios, discounted-cash-flow estimates, or NAV signals—indicate that the stock trades below intrinsic value. The mechanism requires a clear valuation gap rather than strong fundamentals alone.
Buy-the-Dip	12.5%	Investors interpret temporary price declines as opportunities to capitalize on expected mean reversion. Entry often occurs after evidence of stabilization or rebound, and the belief that the drop is transient is central.
ESG/Values	10.3%	Investors restrict purchases to firms aligned with their ethical or ESG standards. Values-based filters act as mandatory gates that can override otherwise attractive fundamentals, momentum, or valuation signals.
Technical Analysis	6.8%	Investors buy when specific technical patterns—such as breakouts, moving-average crossovers, or support bounces—generate predefined entry signals. Chart-based rules drive timing independent of broader narratives.