

Adaptive Heuristics in Stock Selection (Extended Abstract)

Byoung-Hyoun Hwang*

NTU - Singapore

Don Noh

HKUST

Sean Seunghun Shin

KAIST

This draft: March 15, 2026

Abstract

Once investors decide to invest in individual stocks, they face a difficult inference problem. Predicting return distributions across thousands of securities requires processing large amounts of noisy information, and investors often rely on simple heuristics to guide their decisions. While such shortcuts are commonly viewed as behavioral biases that lead investors away from optimal choices, we examine whether heuristics in stock selection may instead represent adaptive responses to this complex environment.

To study this question, we combine a structured survey with adaptive AI-led interviews that elicit how investors select stocks in their own words. We interview 1,540 investors across ten countries, including many with substantial investable wealth. We then apply large language models (LLMs) to the interview transcripts and let the data reveal which heuristics investors invoke and how their use varies across countries.

We find that investors differ in the objectives they prioritize, such as high average returns, homerun potential, safety, or non-financial considerations such as ESG. Moreover, the heuristics investors invoke vary substantially across countries. Strikingly, the heuristics used most frequently in a given country tend to be precisely those that are most effective in that local market.

Taken together, our findings suggest that investors systematically align decision rules with their objectives and adapt these rules to the market environments in which they operate. Heuristics in stock selection therefore do not merely reflect behavioral biases but appear to represent adaptive responses to a difficult inference problem.

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

Keywords: Heuristics, Investor Decision Making, Field Interview, Large Language Models, Artificial Intelligence

*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, Lauren Cohen, Zhi Da, Steve Dimmock, David Hirshleifer, Lawrence Jin, Abhiroop Mukherjee, Mikael Paaso, Johannes Schwab and seminar participants at KAIST, NTU - Singapore, 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. This research is supported by the Ministry of Education, Singapore.

1. INTRODUCTION

Once investors decide to invest in individual stocks, they face a daunting task. Analyzing a stock and predicting its future return distribution is difficult, and there are thousands of securities to choose from. In practice, investors therefore rely on heuristics. These mental shortcuts replace the original, complex question: “which stock has the future return distribution that best matches my preferences?,” with questions that are easier to answer, such as “which stock offers a high and steady dividend yield?” or “which stock has performed well recently?”

A commonly held view is that while such shortcuts simplify decision-making, they move investors far away from the correct answer to the original question and therefore constitute a form of behavioral bias that investors should try to overcome. In the popular press and much financial advice, heuristics are often portrayed as simplistic rules of thumb that lead investors astray (e.g., Malkiel, 1973; Bogle, 1999; Ellis, 2002).

In this paper, we revisit this conventional view and check whether heuristics are all that bad. Conditional on active stock selection, we ask whether investors systematically align their heuristics with the investment objectives they care most about and whether, subsequently, these decision rules actually deliver the characteristics investors seek in their respective market environments.

Our central premise is that heuristics in equity selection may not merely reflect cognitive limitations but, instead, represent adaptive responses to a difficult inference problem. Forecasting expected returns, higher moments, and risks across thousands of securities requires estimating a large number of parameters from limited and noisy data. In such environments, full optimization is computationally demanding and estimation is fragile. Simple decision rules may therefore perform well, and, sometimes, even outperforming full optimization, because they reduce estimation error and direct attention toward relatively stable regularities in the data. Such a perspective is consistent with recent arguments in psychology that heuristics can be adaptive: individuals learn over time which rules work in a given environment, and simple decision rules can outperform more complex ones when the underlying inference problem is high-dimensional and noisy (Gigerenzer and Gaissmaier, 2011).

The ideal empirical design for evaluating this proposition would involve studying investors across multiple environments, systematically documenting which heuristics they invoke and then testing whether these heuristics deliver in that particular environment. Our empirical approach is motivated by this ideal. We first conduct a structured survey with 1,540 investors across ten countries, including many with substantial

investable wealth,¹ and then immediately follow with an open-ended AI-driven interview. In the survey, we ask, among others, how important the following four objectives are to them personally: average returns, homerun potential, safety and non-financial considerations. In the interview, we first ask respondents to explain, in their own words, why they purchased specific stocks in the past. Our AI agent then poses ten to fifteen adaptive follow-up questions designed to clarify reasoning, uncover latent considerations, and elicit concrete examples. We apply large language models (LLMs) on the resulting interview transcripts and systematically map how signals, beliefs, and preferences translate into stock-purchase decisions.

A brief example illustrates the value of this approach. When asked how they choose between stocks, one 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 and how the investor goes about finding such “stable” stocks. In contrast, our AI agent follows up and we learn that, for this investor, “stability” means “will the firm exist in five years” and that the respondent predicts such stability based on a firm’s product popularity (“Do people I know use their products such as Apple/iPhone?”). In the end, we learn that in addition to looking at dividends, this investor uses the heuristic *Familiarity / Brand Affect* to select stocks that deliver on the objective she cares about. 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 predicting a stock’s covariance with the overall market.

In the end, our procedure yields a comprehensive and empirically grounded taxonomy of thirteen stock-selection mechanisms, including *Dividends*, *Familiarity/Brand Affect*, *Fundamental Strength*, and *Growth/Innovation*, just to name a few. These mechanisms emerge inductively from the data in a bottom-up manner, without imposing theoretical priors. We observe substantial heterogeneity across countries. For example, *Fundamental Strength* is the most frequent heuristic among Indian investors, appearing prominently in more than 60% of transcripts, whereas in Germany this heuristic appears prominently in only about 35% of transcripts. Similarly, *Growth/Innovation* appears prominently in more than 60% of transcripts in South Korea but in fewer than 30% in Canada.

Strikingly, we find that the heuristics investors rely on most frequently in a given environment tend to be

¹Our participant pool includes 280 investors from the United States (US) and 140 investors each from Australia, Canada, France, Germany, India, Japan, South Korea, Singapore, and the 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 relatively clean comparisons across countries.

precisely those that are most effective in that environment. For instance, when sorting stocks by profitability, we find that high-profitability firms exhibit much higher future average returns in India (where *Fundamental Strength* is most commonly invoked) than in Germany (where it is much less frequently invoked). Similarly, when sorting stocks by R&D scaled by firm size, we find that high-R&D firms exhibit unusually high future average returns in South Korea but not in Canada. This pattern aligns closely with *Growth/Innovation* appearing prominently in more than 60% of transcripts in South Korea but in fewer than 30% in Canada.

Taken together, our findings suggest that heuristics in stock selection are not merely behavioral biases. Conditional on engaging in active equity selection, investors appear to adapt their decision rules to their objectives and to the market environments in which they operate, and these rules are effective at selecting stocks with the desired characteristics in their particular country.

Our study contributes to three large literature streams. First, our paper contributes to the literature on heuristics and decision making. In much of the behavioral economics and psychology literature, heuristics are often discussed in connection with systematic judgment errors and deviations from optimal statistical reasoning (e.g., Tversky and Kahneman, 1974; Kahneman, 2003; Barberis and Thaler, 2003a). In parallel, the popular discourse surrounding investing frequently portrays simple decision rules as misguided shortcuts that lead investors away from sound financial decisions (e.g., Malkiel, 1973; Bogle, 1999; Ellis, 2002). In contrast, a growing body of work in psychology emphasizes that heuristics can be adaptive responses to complex and uncertain environments. Rather than representing flawed approximations of optimal decision rules, simple heuristics may perform well precisely because they ignore part of the available information, thereby reducing estimation error and focusing attention on relatively robust signals (Gigerenzer and Gaissmaier, 2011). Our paper brings this perspective to the context of equity selection. Using AI-driven interviews with investors across ten countries, we document the heuristics investors invoke when selecting stocks and evaluate whether these rules deliver the characteristics investors seek in their respective market environments. We show that the heuristics most commonly used in a given country tend to be precisely those that are most effective in that environment. These findings suggest that heuristics in stock selection may reflect adaptive decision rules rather than merely behavioral mistakes.

Second, our paper relates to the literature on retail investors' active stock selection. Much of this literature relies on brokerage or administrative trading data to document which types of stocks investors purchase and how these choices relate to demographic characteristics and subsequent performance. Our study complements this work by providing direct evidence on the reasoning underlying why investors purchase particular stocks.

For example, Graham and Kumar (2006) show that, within the retail investor population, older investors more frequently hold dividend-paying stocks. One interpretation is that older investors explicitly invoke a *Dividends* heuristic to select dividend-paying stocks. Another possibility is that they seek to reduce portfolio volatility and therefore select stocks with lower market beta, which happen to be dividend-paying firms. Even if one assumes that investors deliberately choose dividend-paying stocks, administrative data do not reveal *why* they do so. Do investors primarily believe that dividend-paying stocks offer higher total returns? Or do they view these stocks as safer investments that generate less volatile outcomes?

Knowing the answer to this question—the “why”—is crucial for evaluating investors’ skill in active stock selection. Prior research typically assumes a common objective, namely the maximization of risk-adjusted returns, and assesses performance accordingly. Based on this metric, early evidence suggests that, on average, retail investors underperform after costs (Barber and Odean, 2000, 2001). More recent studies paint a more nuanced picture. Retail trades positively predict returns in certain contexts (Boehmer, Jones, and Zhang, 2008; Kaniel, Saar, and Titman, 2008; Kelley and Tetlock, 2013), and some investors appear to improve with experience (Seru, Shumway, and Stoffman, 2010).

Our findings suggest that investors do not share a single objective. Some prioritize high average returns with limited regard for risk, others seek homerun potential or lottery-like payoffs, and still others value safety. If investors differ in their objectives, evaluating all investors using the same risk-adjusted return metric may obscure whether their stock selections align with what they intend to achieve. An investor who deliberately pursues homerun potential may appear to underperform on a Sharpe ratio basis, even if the heuristic they employ successfully delivers the characteristic they value.

Consequently, the finding of low average risk-adjusted returns in the academic literature does not necessarily imply that retail investors lack skill in active stock selection. Rather, it may reflect a mismatch between the researcher’s benchmark and the investor’s objective. Our results suggest that future research should explicitly account for heterogeneity in investor objectives when assessing the effectiveness of retail investors’ stock selections.²

Third, our paper contributes to the expanding literature that uses non-traditional empirical methods to study investor reasoning. Recent work relies on closed-ended surveys to elicit investors’ beliefs and preferences (e.g., Choi and Robertson, 2020; Bender, Choi, Dyson, and Robertson, 2022; Liu, Peng, Xiong, and Xiong,

²One may argue that some investors pursue suboptimal objectives, for example preferring homerun potential despite lower expected returns. Whether investors should maximize a particular objective is a normative question that is distinct from whether they successfully select stocks that deliver their intended outcome. Our analysis focuses on the latter.

2022). While surveys generate standardized and easily comparable responses, they are inherently top-down and therefore limited to considerations anticipated by the researcher. Other studies conduct textual analyses of pre-existing text, such as analyst reports or investment blogs (e.g., Ke, 2024; Bastianello, H. Décaire, and Guenzel, 2024; Chen, Hwang, and Peng, 2025). These approaches benefit from scale and long historical coverage but rely on static text that does not permit clarification or adaptive probing.

Our study complements these approaches by employing AI-driven field interviews that are dynamic, adaptive, and scalable. Unlike closed-ended surveys, our method is bottom-up: respondents drive the content of the conversation, allowing previously unrecognized mechanisms to emerge. Unlike textual analyses of pre-existing documents, our interactive structure enables clarification of ambiguous statements. By combining the depth and non-leading nature of interviews with the scalability of AI, we provide a systematic and empirically grounded account of how investors select stocks. More broadly, our evidence illustrates how AI-led interviews can serve as a powerful tool for studying economic decision-making at scale.

Relatedly, our paper contributes to the growing literature on AI and finance. Recent work explores whether LLM-based agents can form beliefs similar to those of actual investors (e.g., Bhagwat, Cookson, Dim, and Niessner, 2025; Bybee, 2023) and whether simulated markets populated by AI agents replicate empirical regularities (e.g., Henning, Ojha, Spoon, Han, and Camerer, 2025; Lopez-Lira, 2025). Such research requires realistic investor personas. Our evidence on which mechanisms investors invoke, to what degree, and in what combinations provides guidance for constructing such agents, improving both the realism of LLM-based belief measures and the fidelity of AI-based market simulations.