

Are Generative AI Agents Effective Personalized Financial Advisors?

Takehiro Takayanagi
takayanagi-takehiro590@g.ecc.u-
tokyo.ac.jp
The University of Tokyo
Tokyo, Japan

Kiyoshi Izumi
izumi@sys.t.u-tokyo.ac.jp
The University of Tokyo
Tokyo, Japan

Javier Sanz-Cruzado
javier.sanz-
cruzadopuig@glasgow.ac.uk
University of Glasgow
Glasgow, United Kingdom

Richard McCreddie
richard.mccreadie@glasgow.ac.uk
University of Glasgow
Glasgow, United Kingdom

Iadh Ounis
iadh.ounis@glasgow.ac.uk
University of Glasgow
Glasgow, United Kingdom

Abstract

Large language model-based agents are becoming increasingly popular as a low-cost mechanism to provide personalized, conversational advice, and have demonstrated impressive capabilities in relatively simple scenarios, such as movie recommendations. But how do these agents perform in complex high-stakes domains, where domain expertise is essential and mistakes carry substantial risk? This paper investigates the effectiveness of LLM-advisors in the finance domain, focusing on three distinct challenges: (1) eliciting user preferences when users themselves may be unsure of their needs, (2) providing personalized guidance for diverse investment preferences, and (3) leveraging advisor personality to build relationships and foster trust. Via a lab-based user study with 64 participants, we show that LLM-advisors often match human advisor performance when eliciting preferences, although they can struggle to resolve conflicting user needs. When providing personalized advice, the LLM was able to positively influence user behavior, but demonstrated clear failure modes. Our results show that accurate preference elicitation is key, otherwise, the LLM-advisor has little impact, or can even direct the investor toward unsuitable assets. More worryingly, users appear insensitive to the quality of advice being given, or worse these can have an inverse relationship. Indeed, users reported a preference for and increased satisfaction as well as emotional trust with LLMs adopting an extroverted persona, even though those agents provided worse advice.

CCS Concepts

• **Information systems** → **Decision support systems**; **Personalization**.

Keywords

large language models, financial advisor, user study, generative AI

ACM Reference Format:

Takehiro Takayanagi, Kiyoshi Izumi, Javier Sanz-Cruzado, Richard McCreddie, and Iadh Ounis. 2025. Are Generative AI Agents Effective Personalized Financial Advisors?. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*, July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3726302.3729897>

1 Introduction

Personalized advice plays a crucial role in our society, particularly in complex and high-stakes domains like healthcare and finance. Advisors and professionals in these fields use their expertise to offer personalized guidance and emotional support to their clients, leveraging people's specific preferences and/or circumstances. However, advisory services are often provided at a high cost, effectively excluding a large portion of the population from this critical advice. In the financial domain, to mitigate this issue, automated decision support systems have been widely studied, with a special focus on investment-related predictions, such as financial asset recommendations [27, 32, 33].

Recent advances in natural language processing and large language models (LLMs) have significantly accelerated the development of conversational agents, presenting the potential to function as personalized assistants for information-seeking and decision-making [40]. These agents can now leverage multi-turn dialogues, enabling dynamic, mixed-initiative interactions where both users and systems can take the lead in conversations [1]. This progression has expanded the application of conversational agents to various tasks, such as recommendation, question answering, and search [10, 24, 31, 40].

The application of these conversational agents for financial decision-making represents a much more complex scenario than others like movie recommendations, because users are not necessarily familiar with the basic terminology and concepts in this space, and mistakes carry a substantial risk that can lead to large monetary losses. While there is a growing interest in building these conversational assistants to provide automated financial advice [18], previous work has mostly targeted agents capable of handling simple inquiries [16, 35, 36]. Compared to these simple systems, helping users navigate financial decisions and market uncertainties poses a much greater challenge. Therefore, it is not



This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

SIGIR '25, Padua, Italy

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1592-1/2025/07

<https://doi.org/10.1145/3726302.3729897>

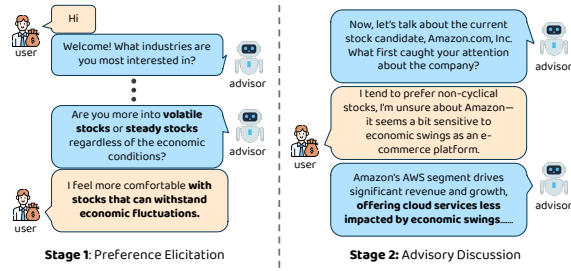


Figure 1: Conceptual illustration of an LLM-advisor with two stages: (1) Preference Elicitation and (2) Advisory Discussion.

yet clear how to develop systems that effectively support complex financial information-seeking and decision-making tasks.

This work aims to close this gap by exploring the effectiveness of LLMs to act as personalized financial advisory agents. In particular, we focus on three problems: (a) eliciting investor preferences through interactive conversations, (b) providing personalized guidance to help users determine whether particular financial assets align with their preferences, and (c) leveraging the personality of the advisor to foster trust on the advisor.

First, the financial literature emphasizes that eliciting user preferences is central to delivering suitable advice [30]. However, it remains unclear whether current conversational technologies, particularly those powered by LLMs, can correctly elicit user preferences in specialized domains where users struggle to articulate their needs. Our work addresses this challenge in the context of financial services.

Second, although personalization is widely regarded as important in the financial decision-support literature [27, 32, 33], its value in a conversational setting remains uncertain. In particular, we explore whether tailoring dialogue around a user’s profile and context improves financial decision-making. Additionally, we also explore how personalization influences user perceptions of the advisor, in terms of aspects like trust and satisfaction.

Finally, in personalized advisory settings within high-stakes domains, the relationship and trust between the client and advisor play a crucial role [18]. Research on conversational agents suggests that agent personality significantly affects users’ perceptions of the system [3, 29]. However, it remains unclear how an advisor’s personality in the financial domain influences both the quality of users’ financial decisions and their overall experience.

To summarize, in this paper, we explore the following questions:

- **RQ1:** Can LLM-advisors effectively elicit user preferences through conversation?
- **RQ2:** Does personalization lead to better investment decisions and a more positive advisor assessment?
- **RQ3:** Do different personality traits affect decision quality and advisor assessment?

To address these questions, we conduct a lab-based user study that explores the effectiveness of LLMs as interactive conversational financial advisors, on which we simulate realistic investment scenarios using investor narratives and stock relevance scores curated by financial experts. Figure 1 illustrates an example conversation with the advisor, divided into two stages: first, the LLM-advisor attempts to capture the investor preferences through conversation; in

the second stage, given an individual asset, the advisor provides information about it to the investor, including how the asset matches (or not) the investor’s preferences. To answer the different questions, we compare different configurations of the LLM-advisor: first, we compare personalized vs. non-personalized advisors, and, then, we compare two personalized advisors with distinct personalities.

2 Related Work

2.1 Personalization and Preference Elicitation

Information systems, especially those focused on search and recommendation benefit from personalization [14]. Specifically, personalization techniques play a crucial role in enhancing user experience [17, 22, 41]. Interactive approaches, such as conversational preference elicitation represent the frontier of personalization. This problem has received growing attention, as advances in generative AI now provide a functional mechanism to collect user preferences dynamically in a free-form manner [40]. This interactive approach can capture more diverse and targeted insights than static approaches like questionnaires [6, 10, 23, 24, 31]. Indeed, recent studies have proposed various methods for effective conversational preference elicitation [31, 42], as well as user studies on the perceived quality of this process in domains such as e-commerce, movies, fashion, books, travel, and restaurant recommendations [2, 7, 15, 23, 31, 45].

However, we argue that for some important domains, trying to directly collect preferences is insufficient. An implicit assumption of these studies is that if directly asked, the user will be able to accurately express their preferences. It is reasonable to expect that this assumption would hold for scenarios like movie recommendation; we can ask a user “do you like horror movies?” and expect a useful response. On the other hand, this will not hold for complex tasks, where the user lacks the knowledge to form an accurate response [10, 39]. For instance, in an investment context if we asked “do you prefer ETFs or Bonds?”, it is not clear that an inexperienced user would be able to produce a meaningful answer. In these cases, an ideal agent needs to fill the gaps in the user knowledge through conversation, as well as infer the user preferences across multiple (often uncertain) user responses. But how effective are generative AI agents at this complex task? This paper aims to answer that question for the domain of financial advisory; a particularly challenging domain given its technical nature and high risks if done poorly.

2.2 Financial advisory

In the financial domain, advisors help individuals manage their personal finances by offering guidance on investments and assisting with decision-making [34]. While financial advisors can be beneficial, their services often come at a high cost, making them unaffordable for many people. To mitigate this issue, automated (non-conversational) financial decision support systems such as financial recommender systems have been widely studied [44]. The majority of research in this area has been focused on how to find profitable assets (i.e. those that will make money if we invest in them). These works assume a simplified user-model, where an investor is only concerned with maximizing return-on-investment over a fixed period of time [27, 32, 33]. These studies frame financial advisory as a ranking problem, where the goal is to rank financial

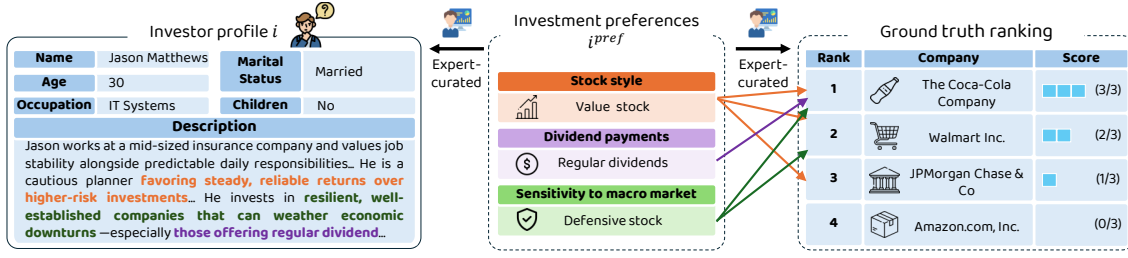


Figure 2: Example of an investor profile, investment preferences, and ground truth ranking. Dashed line components are used for evaluation (and therefore, they are not shown to the user/LLM).

assets for a user over a specified time period. However, a recent study suggests that a large part of the value offered by human financial advisors stems from their ability to personalize investment guidance to clients' specific needs, build relationships, and foster trust [13], rather than simply presenting suitable assets.

Reflecting on these findings, the development of conversational financial advisors has drawn increasing attention, as it enables a dynamic understanding of users' needs, personalized guidance, and the potential to build trustworthy relationships [9, 16, 34, 43]. In particular, the conversational agents' personality has gained attention as a factor that can help build relationships with clients and foster trust [18], especially given the successes of conversational agents using the Big Five personality model [20] to enhance the end-user experience [4, 30]. Although conversational agents show potential in finance, how to configure them to match the value of human advisors remains unclear. Therefore, we conduct a user study to examine how personalizing investment guidance and the advisor's personality shape users' financial decision-making effectiveness and overall user experience.

3 Methodology

In this paper we aim to determine to what extent current generative language models can act as an effective financial advisor. Indeed, given the need to personalize for the user, emotional implications, the technical nature of the information-seeking task, and high impact if failed, we argue that this is an excellent test case for the limits of generative large language models. To structure our evaluation, we divide our study into two phases, as illustrated in Figure 1, where we evaluate the success of both:

- (1) **Preference Elicitation:** During this stage, we have the LLM-advisor hold a natural language conversation with a human, where it is directed to collect information regarding the person's investment preferences. The human in this interaction is pretending to have preferences from a given investor profile.
- (2) **Advisory Discussion:** During the advisory discussion, the LLM-advisor again has a natural language conversation with the human (acting on an investor profile), where the human collects information about whether a company is a suitable investment for them. This is repeated for multiple companies per investor profile.

We provide preparatory information and discuss each stage in more detail below:

3.1 Investor Profiles

To fairly evaluate the ability of any LLM-advisor, we need to have them interact with human users with real needs. Given the open-ended nature of free-form conversations, it is desirable to repeat each experiment with different people such that we can observe variances in conversation paths, as those variances may influence task success. However, to enable repeatability, we need to hold the investor needs constant across repetitions. Hence, we define three archetypal investor profiles $i \in I$ based on input from a financial expert, where our human participants are given one to follow when conversing with the LLM-advisor:

- **Investor 1: Growth-Oriented Healthcare Enthusiast:** Prefers healthcare innovations, values high-growth opportunities, and takes measured risks.
- **Investor 2: Conservative Income Seeker:** Seeks stable returns, invests in well-established companies, values regular dividend payouts.
- **Investor 3: Risk-taking Value Investor:** Targets undervalued companies with strong long-term potential, tolerates short-term volatility, and invests in cyclical sectors.

For each of these investor profiles, we select three key investment preferences, chosen from well-known investment characteristics such as industry sector, stock style, consistency in dividend payments, and sensitivity to global market changes [8]. We denote the set of investor preferences as i^{pref} . In our experiments, we simulate a realistic elicitation scenario where the advisor collects the preferences from the participants. Therefore, we do not straightforwardly provide the preferences to the participants. Instead, we present them as text narratives of between 150 to 200 words. A financial expert was consulted to confirm the quality and reliability of these narratives. An example narrative representing Investor 2 is illustrated in Figure 2, where we highlight the sentences referring to specific investor preferences.

3.2 Stage 1: Preference Elicitation

The goal of stage 1 of our study is to determine to what extent an LLM-advisor can effectively collect a user's investment preferences through conversation. Formally, given a participant of the user study u and an investor profile i , during the elicitation stage, the LLM-advisor aims to obtain an approximated set of preferences, denoted i_u^{LLM} , that matches the investor preferences (i^{pref}). To achieve this, the generative model produces a series of questions that participants answer by interpreting the investor narrative.

Responses to those questions, denoted as R_i^u , are used by the LLM-advisor to generate the user profile i_u^{LLM} . Success is then measured by manually evaluating the overlap between i^{pref} and i_u^{LLM} .

For user elicitation, we adopted a System-Ask-User-Respond (SAUR) paradigm [42]. During the conversation, the advisor proactively inquires about the user's preferences given a set of target preferences (e.g., industry type, acceptable risk). After the human participant responds to a question, the LLM-advisor checks whether the collected preferences cover all of the target preferences. If the advisor is confident that they do, it ends the conversation and prompts the user to proceed to the next stage; otherwise, it continues asking follow-up questions in a loop.

3.3 Stage 2: Advisory Discussion

Stage 2 of our study investigates to what extent an LLM-advisor can provide the same benefits as a real human advisor when exploring investment options. Note that the goal here is not to have the LLM-advisor promote any one asset, but rather to provide accurate and meaningful information such that the human can find the best investment opportunity for them. To this end, we structure our experiment such that the human (acting on an investor profile) has one conversation with the LLM-advisor for each of a set of assets being considered.¹ After all assets are presented to the participant, a stock ranking is generated by sorting the stocks by the participant rating in descending order.

Importantly, as we know the investor profile i^{pref} for each conversation about an asset a , we can objectively determine whether a is a good investment given i^{pref} , forming a ground truth against which we can compare to the rating provided by our human participant after their conversation with the LLM-advisor. For each asset a , a financial expert produced a score between 0 and 3 by manually checking whether a satisfied each of the three investment criteria contained in i^{pref} . A ground-truth ranking was produced by sorting the assets by the expert scores. We show an example of the ranking construction in Figure 2. During evaluation, the closer the participant ranking is to the ranking produced by expert judgments, the better the LLM-advisor performed.

Baseline Prompt: As we are working with an LLM-advisor and the nature of financial information-seeking is time-sensitive, we need to provide any information that might change over time to the LLM within the prompt. As such, for each asset a , we pre-prepared a standard asset descriptor block after consulting with a financial expert, containing:

- **Stock Prices:** We collect monthly stock prices from 2023 using Yahoo! Finance.²
- **Business Summary:** We gather each company's business overview from Yahoo! Finance.
- **Recent Performance and Key Financial Indicators (e.g., EPS):** We obtain earnings conference call transcripts³ from Seeking Alpha for the last quarter of 2023.

¹These were manually selected, however in a production environment these might be produced by an asset recommendation system.

²The scenario for the financial advising of our user study is set to December 30, 2023. By basing our experiment at the end of 2023, we avoid the problem of data contamination [25].

³Earnings conference calls, hosted by publicly traded companies, discuss key aspects of their earnings reports and future goals with financial analysts and investors, thus covering critical financial indicators and recent performance insights [21]. These

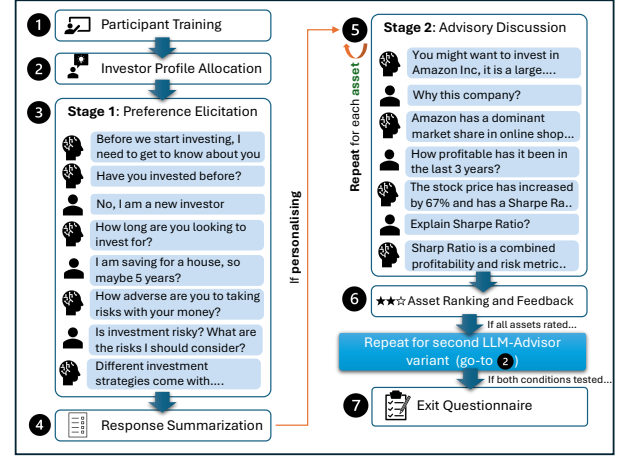


Figure 3: User study structure.

The advisor using this prompt acts as our baseline for the advisory discussion study. We augment this baseline with additional context and instructions to form two additional experimental scenarios, discussed below:

+Personalization: As discussed earlier, one of the core roles of the financial advisor is to personalize to the individual customer, based on their financial situation, needs, and preferences. To enable the LLM-advisor to personalize for the user, we integrate the generated profile from the preference elicitation (Stage 1) i_u^{LLM} into the prompt. We represent each preference as a series of short sentences.

+Personality: In Section 2.2 we discussed how human financial advisors provide emotional support as well as financial advice. While it is unlikely that an LLM-advisor could do this as well as a human (it lacks both emotional intelligence and non-conversational clues to the customer's mental state [38]), it might be possible to provide a better end-user experience by directing the LLM-advisor to adopt a *personality*. As noted in Section 2 it is possible to do this via prompt engineering, such as instructing the LLM to take on the traits of one or more of the Big-Five personality types [20].

As we are performing a user study with humans, it would be impractical to exhaustively test every combination of personality types, hence as an initial investigation we experiment with two distinct personality profiles [29]:

- **Extroverted:** High in extroversion, agreeableness, and openness; low in conscientiousness and neuroticism.
- **Conscientious:** Low in extroversion, agreeableness, and openness; high in conscientiousness and neuroticism.

We adopted the prompting method from Jiang et al. (2024) to assign a Big Five personality trait to the LLM agent [12], choosing it for its simplicity and effectiveness among various proposed approaches for embedding personality in LLMs (including both prompting and fine-tuning) [11, 12, 28]. To ensure a high standard of professionalism and accurate representation of the intended personality, we consulted financial professionals to review the texts generated by LLMs adopting both personas.

transcripts cover significant financial indicators and provide explanations of recent performance.

3.4 Experimental Design

In our experiment, we conducted two studies: a personalization study (for RQ2) and an advisor persona study (for RQ3). In the personalization study, participants compared a non-personalized (Baseline) advisor with a personalized (+Personalized) version. In the advisor persona study, they compared different LLM-advisor personality types (+Extroverted vs. +Conscientious). Participants are randomly assigned to one of these two studies.

Figure 3 shows the structure of our user study for a single participant, comprising seven steps:

- (1) **Participant Training:** Participants are given a general overview of the user study and given instructions on their expected roles during preference elicitation, advisory discussions, asset ranking, and advisor assessment.
- (2) **Investor Profile Allocation:** The user u is randomly allocated one of the investor profiles (See Section 3.1) that they will follow. Each profile is assigned to 42 participants.
- (3) **Preference Elicitation (Stage 1):** The participant interacts with the LLM-advisor as if they were a new investor. The conversation ends once the LLM-advisor determines that they know enough about the investor to personalize for them. The median time spent on preference elicitation was 5 minutes and 11 seconds.
- (4) **Response Summarization:** Given the aggregator of user responses R_i^u , we instruct an LLM to generate an investor profile i_u^{LLM} . For each investor preference in i^{pref} , if there is any relevant information in the responses R_i^u , that information is included in i_u^{LLM} . Otherwise, i_u^{LLM} indicates that no relevant information is available for that specific preference.
- (5) **Advisory Discussion (Stage 2):** To simplify the conversation flow we have the participant hold separate conversations with the LLM-advisor for each asset they might invest in. The LLM-advisor is provided with context about the current asset (see Section 3.3), and depending on the experimental scenario, optionally personalization information (step 4 output) and/or a target personality context statement. Each conversation continues until the user is satisfied that they have enough information to rate the asset. The order in which the assets are discussed is randomly assigned to avoid position bias.
- (6) **Asset Ranking and Feedback:** Participants rank all the stocks (four in total) discussed in the advisory session according to their desire to invest in each. They also assess the advisor they interacted with using a 7-point Likert scale for the items listed in Table 1 (see Section 4).

To enable more effective pair-wise comparison of LLM-advisor variants, we have each participant test *two variants per study*. If the user has only tested one variant at this point, then they repeat the user study (starting at step 2) with the second variant. The order in which participants experience each variant is randomly assigned.

- (7) **Exit Questionnaire:** Once a pair of LLM-advisor variants have been tested, the user fills in an exit questionnaire that is designed to ask the overall experience in the user study.

Table 1: Operational definitions used in the advisor assessment questionnaire for all response dimensions.

Response Dimension	Operational Definition
Perceived Personalization [14]	The advisor understands my needs.
Emotional Trust [14]	I feel content about relying on this advisor for my decisions.
Trust in Competence [14]	The advisor has good knowledge of the stock.
Intention to Use [14]	I am willing to use this advisor as an aid to help with my decision about which stock to purchase.
Perceived Usefulness [22]	The advisor gave me good suggestions.
Overall Satisfaction [22]	Overall, I am satisfied with the advisor.
Information Provision [37]	The advisor provides the financial knowledge needed.

In our experiments, we use Llama-3.1 8B as the background model for all our LLM-advisor variants.⁴

3.5 Participants

We recruited 64 participants from the authors’ affiliated university for our study: 32 participants for the personalization study and 32 participants for the advisor persona study, utilizing the university’s online platform and blackboard for recruitment. Participants were required to be fluent in English, over 18 years old, and have an interest in finance and investment, mirroring the target demographic of our system’s users. After excluding invalid data, 29 participants remained in the personalization study and 31 in the advisor persona study. We conducted a power analysis using the Wilcoxon signed-rank test for matched pairs, with the experimental conditions as the independent variable and users’ response to the advisor assessment questionnaire as the dependent variable [26]. The analysis determined that 29 participants are needed to observe a statistically significant effect on user-perceived quality. Our recruitment criteria and compensation (£10/hour) for approximately one hour of participation were approved by our organization’s ethical board.

4 Evaluation Metrics and Statistics

In this section we discuss how we quantify effectiveness for the preference elicitation and advisory discussion stages, respectively, in addition to summarizing dataset statistics for each.

4.1 Preference Elicitation Metrics (Stage 1)

To evaluate the quality of the first preference elicitation stage, we want to measure how well the LLM-advisor has captured the investor preferences as defined in the investor profile i (see Section 3.1). Each investor profile $i \in I$ defines key features of the investor, such as preferring high-growth stocks, or favoring regular payouts, denoted i^{pref} . We have three investor profiles ($|I|=3$), with 10 (n) participants performing elicitation on i_u^{LLM} for each profile and each LLM variant, i.e. there are 120 elicitation attempts in total, with 30 attempts per LLM-advisor variant. Following the notation in Section 3, i_u^{LLM} in this case denotes a similar list of features to i^{pref} that LLM-advisor learned about the investor during conversation with a participant u , which we derive from a manual analysis of the elicitation output (i.e. what is produced by response summarization). Intuitively, the closer the features produced from

⁴Further details about the LLM configuration, investor narratives, relevant scores, prompts and scripts for data analysis can be accessed at the following repository: https://github.com/TTsamurai/LLMAdvisor_supplementary

Table 2: General statistics of the collected conversation data.

Participants	60
Time Period	2024/10/24 ~ 2024/11/7
Total Turns	10,008
Stage 1: Preference Elicitation	
Total Turns	1,788
Number of Sessions	120
Avg. Turns/Session	15.8
Avg. User Words/Turn	9.8
Stage 2: Advisory Discussion	
Total Turns	8,220
Number of Sessions	480
Avg. Turns/Session	18.2
Avg. User Words/Turn	13.0

any elicitation attempt i_u^{LLM} is to i^{pref} , the better the LLM-advisor is performing. To this end, we report elicitation accuracy for each investor profile, calculated as:

$$\text{ElicitationAccuracy}(i) = \frac{1}{n} \sum_{j=1}^n \frac{|i_j^{LLM} \cap i^{pref}|}{|i^{pref}|} \quad (1)$$

Human Advisor: To provide a point of comparison, we also conduct a preference elicitation with a financial expert using the same prompt and instructions as the LLM. This allows us to evaluate how close LLMs are to a paid human advisor undertaking the same task. More specifically, for each investor profile, three participants engaged with this expert, who then produced a set of preferences i_u^{Expert} , which can be used instead of i_u^{LLM} in Equation 1.

4.2 Advisory Effectiveness Metrics (Stage 2)

Ranking correlation (Spearman’s Rho): In the second stage, we evaluate how well the LLM-advisor can support an investor to select financial assets that are suitable for them to invest in. Recall from Figure 3 that after a participant finishes discussing all assets with the LLM-advisor, they rank those assets $a \in A_i$ based on the likelihood they will invest in each, i.e. each participant u acting on a profile i we have an asset ranking $R(A_i, i_u)$. As illustrated in Figure 2, each investor profile i was derived from a ground truth set of investor preferences i^{pref} , which an expert used to create a ground truth ranking $R(A_i, i^{pref})$, i.e. the “correct” ranking of assets. Intuitively the closer the $R(A_i, i_u)$ is to $R(A_i, i^{pref})$, the better the advisor is performing, as the participant was better able to distinguish suitable assets vs. unsuitable ones. Hence, to evaluate the effectiveness of the advisory task, we report the mean ranking correlation (Spearman’s Rho) between $R(A_i, i_u)$ and $R(A_i, i^{pref})$ across participants u for each LLM-advisor.

Advisor Assessment Questionnaire: Lastly, we also gather qualitative data from each participant via a questionnaire. In particular, after ranking assets each participant, reports how they feel the LLM-advisor performed in terms of 7 dimensions, listed in Table 1, such as perceived usefulness, trust, and user satisfaction. We use this data later to evaluate how sensitive the user is to differences in the LLM-advisor.

4.3 Dataset Statistics

Table 2 summarizes the statistics of the data collected during the two stages of our user study. Each conversation that a participant had with an LLM-advisor in either stage 1 or 2 is referred to as

Table 3: Stage 1 - Comparison of Elicitation Accuracy of an expert vs. different LLM-advisors for each investor profile. The best advisor is highlighted in bold. Arrows denote percentage increases (↑) or decreases (↓) compared to the expert.

Investor Profile	Expert	LLM-Advisors			
		LLM	+Extr.	+Cons.	Average
Growth-Oriented	0.78	0.76	0.80	0.79	0.78 ^{→0.0%}
Conservative-Income	0.89	0.82	0.75	0.87	0.82 ^{↓7.8%}
Risk-Taking	0.89	0.48	0.60	0.55	0.53 ^{↓40.5%}
Average	0.85	0.69	0.70	0.73	0.70 ^{↓17.6%}

a session, e.g. during Stage 1, there were 3 investor profiles * 10 participants * 4 LLM-advisors, resulting in 120 sessions. Stage 2 has 4x the number of sessions, as there are four assets associated with each profile (A_i) to discuss with the LLM-advisor.

From Table 2 we observe that in contrast to other conversational tasks [35, 36], financial information-seeking appears to require more extended interactions. On average, preference elicitation involves 15 turns per session with 9.8 words per turn, whereas advisory discussions involve 18 turns per session with 13.0 words per turn, highlighting the overall complexity of the task.

5 Results

In this work, we explore how to design conversational financial advisors that enhance both decision-making and positive experience. To achieve this, our user study is guided by 3 core research questions.

- **RQ1:** Can LLM-advisors effectively elicit user preferences through conversation?
- **RQ2:** Does personalization lead to better decisions and more positive advisor assessment?
- **RQ3:** Do different personality traits affect decision quality and advisor assessment?

5.1 RQ1: Elicitation accuracy

We begin by examining how effective the LLM-advisors are at identifying investment preferences during conversations in Stage 1. Elicitation Accuracy is the primary metric, where we contrast the mean accuracy across 10 sessions in comparison to a human expert tackling the same task (see Section 4.1). Table 3 reports elicitation accuracy for each LLM-advisor and the Human Expert across investment profiles. Arrows denote percentage increases (↑) or decreases (↓) of the LLM-advisor compared to the expert.

To set expectations, we first consider the performance of the expert in the first column in Table 3, as we might expect, the expert maintains consistently high performance across all profiles, averaging 85% accuracy (random accuracy is 50%). This forms an expectation of the performance ceiling for the task.

Next, we compare the expert performance to each LLM-advisor. From the perspective of preference elicitation, there are three LLM-advisor configurations, those that use only the Baseline Prompt (denoted LLM) from the personalization study, and those that include a defined personality (either extroverted, +Extr., or conscientious, +Cons.) from the advisor persona study.⁵ From Table 3, we observe

⁵Note we cannot have a personalized variant here, as the personalization evidence is derived from this stage.

that the LLM-advisor’s performance is generally strong for growth-oriented, and conservative-income investors (with accuracy around 80%) on average, which is similar to the human advisor. However, for the risk-taking investor profile, the LLM-advisor’s elicitation accuracy was substantially lower (~40.5%).

From a manual failure analysis, we observed the following trends that contribute to the performance gap with the human advisor, particularly for the risk-taking profile. First, it is notable that elicitation failures can originate from the investor (participant) rather than the LLM. Recall that one of the aspects that makes finance more challenging than domains like movie recommendation is that the “user” is inexpert, and so may give incorrect information during the conversation. Indeed, we observed cases where the participant confused concepts such as the difference between a growth and a value stock, as well as cyclical/non-cyclical assets. On the other side, preference hallucination is a core issue for the LLM-advisor. The LLM is a probabilistic token generator conditioned on the baseline prompt and prior conversation, and as a result, in some scenarios, the contextual content can override a statement by the investor. This type of error is more likely when the investor is unsure in their responses or when they provide contradictory statements. For instance, an investor expressing an interest in the consumer discretionary sector while simultaneously opting for non-cyclical stocks, despite consumer discretionary being inherently cyclical.

To answer RQ1, our results demonstrate that LLM-advisor’s are able to elicit preferences from a user via conversation and that for 2/3’s of the user profiles tested, elicitation accuracy was consistently equivalent or close to that of an expert human advisor. However, we observed a clear failure mode when testing the risk-taking profile, where misunderstandings by the investors and hallucinations within the LLM compound to result in accuracy that is close to random. Overall, we consider this a promising result, as the majority of the time it is effective, and the failure mode observed might be rectified by better context crafting and the addition of contradiction detection; both directions for future research.

5.2 RQ2: Effectiveness of personalization

Having shown that automatic preference elicitation is possible, we now examine stage 2 of our study, namely the advisory discussions. Given the inherently personalized nature of financial advice, we expect that the customer preferences obtained during stage 1 will be key to enabling LLM-advisors to provide effective investment advice. Hence, in this section, we compare the performance of an LLM-advisor using only the Baseline Prompt to one that includes the preferences obtained during stage 1 (+Personalized). However, as we observed that preference elicitation is not always successful, we also examine what effect elicitation performance has on the LLM-advisor.

5.2.1 Non-personalized Decision-making Effectiveness: We initially establish how effective the LLM-advisor is without any information regarding the investor. LLM-advisor effectiveness is measured based on how well the investor was able to rank the assets discussed by suitability to them. The primary metric is average Spearman’s Rho correlation between the investor ranking and the ground truth ranking (see Section 4.2), reported in Table 4 row 1. As we expect,

Table 4: Investor decision-making effectiveness, expressed as the Spearman’s Rho correlation between the investor’s asset ranking and the expert asset ranking (higher is better). [†] indicates statistical improvements (Welch’s t-test with $p < 0.05$) over the not personalized baseline, while [§] indicates significant differences between cases with successful and unsuccessful preference elicitations.

Advisor Config		Investor vs. Expert (Spearman’s Rho)		
Personalization	Personality	All	Preference Elicitation	
			Successful	Unsuccessful
Baseline	None	0.110	–	–
+Personalized	None	0.310	0.481 ^{†§}	-0.228
+Personalized	+Extroverted	0.122	0.243 [§]	-0.286
+Personalized	+Conscientious	0.26	0.365	-0.025

baseline advisory performance is low, with only a very weak positive correlation to the ground truth ranking of 0.11. This indicates that without further evidence, the LLM is not able to meaningfully guide the investor.

5.2.2 Personalized Decision-making Effectiveness: Having established our baseline, we now examine the impact that adding the investor preferences collected during stage 1 has, comparing Table 4 row 1 (baseline) to row 2 (personalized). As we anticipated, personalization is beneficial, with investor decision-making effectiveness increasing from 0.11 to 0.31 (average Spearman’s Rho correlation to the expert ranking). However, this correlation is still weak, illustrating that while discussing assets with the LLM-advisor is better than no help at all, our participants are still struggling to evaluate the suitability of financial assets.

This correlation is an average over all the participants in the user study, regardless of how effective their preference elicitation was in stage 1. Hence, we might ask whether the low correlation is due to the LLM-advisor being confused by poor preference elicitation data. To explore this, Table 4 also reports investor decision-making effectiveness stratified based on whether stage 1 was successful (column 4) or not (column 5).⁶ As expected, we see a statistically significant increase in investor decision-making effectiveness when preference elicitation was successful when compared to non-personalized sessions (0.481 vs. 0.110). More concerningly, we also see the LLM-advisor has a strong negative influence on the investors’ decision-making capability if preference elicitation fails, as illustrated by the negative correlations with the expert in column 5. This result highlights both that effective preference elicitation is crucial, but also that the LLM-advisor can easily influence the investor into making poor decisions, as the human is heavily reliant on the agent to navigate the relatively unfamiliar financial information space.

5.2.3 Participant Assessment of the Advisor: So far we have demonstrated that there is a large difference between a non-personalized LLM-advisor and a personalized one, in terms of how they can alter the decision-making of the investor/participant. But can the participant tell the differences between them?

Table 5 reports the aggregation of the qualitative data we collected from each participant after they finished interacting with each LLM-advisor in terms of 7 dimensions, where we start by

⁶We define that an elicitation session is successful if more than 50% of the investor’s preferences were correctly captured

Table 5: Average participant users’ response to advisor assessment questionnaire under different advisor conditions. Columns labeled with advisor condition (Baseline, +Pers., +Cons., +Extr.) contain a 7-point Likert scale (higher is better). “p” column contains Wilcoxon signed-rank test p -values for (RQ2) Baseline vs. +Personalized (Pers.), and (RQ3) +Conscientious (Cons.) vs. +Extroverted (Extr), for both the full data (All) and the subset where the elicitation accuracy is above 0.5. “Successful Elicitation” refers to the subset where elicitation accuracy was ≥ 0.5 . For RQ2, this subset consists of pairs for which +Pers elicitation is successful, while for RQ3, it consists of pairs for which both +Extr and +Cons elicitation are successful. Boldface indicates significant effects with † for $p < 0.1$ and ‡ for $p < 0.05$.

Response Dimension	(RQ2) Baseline vs. +Personalized						(RQ3) +Conscientious vs. +Extroverted					
	All			Successful Elicitation			All			Successful Elicitation		
	Baseline	+Pers.	p	Baseline	+Pers.	p	+Cons.	+Extr.	p	+Cons.	+Extr.	p
Perceived Personalization	5.759	5.724	0.838	5.762	5.905	0.751	5.500	5.500	0.663	5.588	5.706	0.941
Emotional Trust	5.103	5.241	0.446	5.143	5.333	0.537	5.038	5.154	0.600	4.706	5.235	0.034[‡]
Trust in Competence	5.690	5.690	0.817	5.810	5.857	0.782	5.962	6.077	0.538	6.000	6.000	1.000
Intention to Use	5.310	5.483	0.505	5.429	5.714	0.166	4.885	5.462	0.005[‡]	4.941	5.588	0.013[‡]
Perceived Usefulness	5.241	5.517	0.183	5.381	5.810	0.194	5.423	5.538	0.425	5.176	5.118	0.968
Overall Satisfaction	5.345	5.690	0.116	5.429	5.810	0.098[†]	5.269	5.577	0.179	5.118	5.529	0.244
Information Provision	5.517	5.966	0.026[‡]	5.714	6.143	0.053[†]	5.692	5.654	0.953	5.588	5.765	0.490

focusing on the RQ2-All columns, i.e. comparing the baseline and personalized variants. The important observation to note here is that the participant preference scores for both variants are statistically indistinguishable, except under the quality of information provision criteria. This means that our participants cannot tell if the LLM-advisor is personalizing to them, and trust the worse agent just as much as the better one. Furthermore, if we consider the best case scenario where the preference elicitation was successful (RQ2 Successful Elicitation columns) we observe the same pattern, even though the difference between the baseline and the personalized variants in terms of the effect it has on the participant decision-making is more pronounced. This underlines one of the core risks of using LLM-advisors in the financial domain; since our users are inherently inexperienced they lack the fundamental skills to judge to what extent the LLM is providing good advice, meaning that there is no safety net if the LLM makes a mistake.

To answer RQ2, our results show that a personalized LLM-advisor is able to provide useful financial advice when it has accurate information regarding the preferences of the investor. This is demonstrated by better decision-making capability by participants using the personalized advisor in comparison to the non-personalized one. However, we also identified two important challenges to adoption. First, the impact the LLM-advisor has is strongly tied to the quality of the preference elicitation data provided, where poor preference elicitation will cause the agent to actively direct the investor to the wrong assets. Second, while the participants were positive regarding the LLM-advisors across all questionnaire criteria, they were not able to consistently tell the difference between good and bad advisors; leading to an increased risk of humans acting on bad advice.

5.3 RQ3: Effectiveness of personalities

Once we have confirmed the utility of personalization for LLM-advisors, we now study the effect that the personality of the advisor has on users’ financial information-seeking. As previous studies have shown [29], chatbot personality can affect the way humans interact with the chatbot, and therefore affect the effectiveness and perception of LLM-advisors. To understand whether personality affects LLM financial advisors, we compare two personalized LLM-advisors on which we have injected a pre-defined personality: an

extroverted personality and a conscientious personality.⁷ While we could consider the personalized LLM-advisor discussed in Section 5.2 as a third distinct personality (the base LLM personality of the LLM), we shall not compare it with our personality-injected models, because different sets of participants were used in the personalization study and the advisor-persona study.

5.3.1 Decision-making Effectiveness: We first examine the impact of adding personality to the advisors on the decision-making process, by measuring the capacity of the participants to correctly rank the assets (as previously done in Section 5.2). As a primary metric, we again use the average Spearman’s Rho correlation between the investor ranking and the ground truth ranking reported in Table 4 rows 3 (extroverted advisor) and row 4 (conscientious advisor).

We first observe the results for the full set of participants in the user study. Interestingly, we observe a difference between the two advisors, with the conscientious LLM-advisor providing better guidance than the extroverted one (0.26 vs. 0.122). This observation is consistent when we restrict our analysis to those cases where the preference elicitation is successful. While, expectedly, the effectiveness of both advisors improves when the elicitation is successful (0.243 vs. 0.122 in the case of the extroverted advisor and 0.365 vs. 0.26 in the case of the conscientious one), the conscientious advisor has an advantage over the extroverted one (0.365 vs. 0.26).

These results highlight that providing different personalities to an LLM-advisor can notably impact the capacity of the advisor to provide useful information to the investors.

5.3.2 Participant Assessment of the Advisor: We have observed so far that the use of different personalities affects the user decision-making process. But how do these personalities affect the perception that users have of the LLM-advisor? We observe this in Table 5, in terms of the seven dimensions captured during the advisor assessment questionnaire.

We first look at the RQ3-All columns, comparing the two personalities. Notably, for the majority of the dimensions, users barely distinguish between both systems. The only answer where we observe a statistically significant difference is the intention to use the

⁷Refer to Section 3.3 for a full description of each personality.

system in the future. Surprisingly, despite providing worse guidance to the investor, participants expressed a higher interest in using the extroverted advisor than the conscientious one. When we limit our study to those participants who experienced a successful preference elicitation in both advisor variants, this issue is stressed, as those users also develop a significantly greater emotional trust with the extroverted advisor.

These observations are worrisome, as they reveal that the personality of a financial advisor cannot only affect the quality of the advice but also lead the investors to trust more on those systems providing worse advice.

5.3.3 Differences in language: To further understand how personalities affect financial advisory, we analyze the differences in the linguistic patterns provided by extroverted and conscientious advisors. Analyzing participants' reported overall experience from the exit questionnaires in the advisor persona study, over 20% (7 of 31) described the extroverted advisor as clear, assertive, and cheerful while perceiving the conscientious advisor as straightforward, analytical, yet less confident.⁸ Therefore, to quantify the linguistic differences in the advisors, we conduct a financial sentiment analysis of the utterances generated by each advisor. For each utterance, we count the occurrences of positive, negative, and uncertain words from the Loughran and McDonald Financial Sentiment Dictionary [19]. We normalize these counts by the length of the sentences and average the results across all dialogues.

Figure 4 shows the results, showing the extroverted sentiment scores in blue, and the conscientious scores in orange. For the three sentiment dimensions, differences between advisors are statistically significant (Welch's t-test with $p < 0.01$). Figure 4 shows that extroverted advisors tend to use more positive language in their interactions, while conscientious advisors prefer negative and uncertain tones. Through manual analysis of the conversation, we observe that this results in the extroverted advisor focusing on the positive aspects of investments while overlooking serious drawbacks, whereas the conscientious advisor provides a more balanced view of the assets. Because of this, participants guided by conscientious advisors may make more well-informed financial decisions. Meanwhile, the positivity of the extroverted advisor seems more appreciated by the users, which is reflected in higher advisor assessment scores from the post-discussion questionnaire.

To answer RQ3, our results show that different personalities of a personalized LLM-advisor can affect the utility of the provided advice. This is demonstrated by the better decisions of the study participants when using an advisor with a conscientious personality than when using an advisor with an extroverted personality. Moreover, the personality of the advisor affects the perception of humans towards the system, and it has the risk of leading investors to further trust those systems that provide worse advice.

6 Conclusion

In this paper, we have conducted a lab-based user study to examine how effective large language models are as financial advisors. We focus on three core challenges: preference elicitation, investment personalization, and advisor personality.

⁸Participants were unaware of the specific personas during the study.




Figure 4: Average sentiment scores by advisor personality. Error bars indicate the standard deviation.

First, our analysis shows that LLMs are effective tools for preference elicitation through conversation. In a majority of cases, they are capable of obtaining investor's preferences with an accuracy close to or equivalent to that of an expert human advisor. However, there are some clear failure cases, as LLMs are vulnerable to contradictory statements and hallucinations, which, in the case of complex investor profiles, can decrease the accuracy of the elicitation to random levels. Although LLMs are promising for elicitation, in a complex domain like finance, investors do not always fully understand their own preferences (or they have difficulties expressing them). Therefore, future work should explore the development of LLM-advisors capable of resolving conflicting user needs.

Second, personalizing LLMs to provide investment advice can improve the decisions made by the investors, but only when the personalized LLM-advisor receives accurate information about the investor's preferences. If the preference elicitation is not successful, the agent actively directs the investors to the wrong assets on which to invest. This underscores how crucial a good preference elicitation is for providing useful financial advice.

Finally, our results suggest that investors are not necessarily aware of what constitutes good financial advice, and therefore, are vulnerable to acting on bad advice provided by LLMs. In the comparison between a non-personalized and a personalized LLM-advisor, although the personalized system led to better decisions, participants were unable to distinguish between the systems. More worryingly, when comparing two personalized advisors with extroverted and conscientious personalities, we observed that, even though the extroverted advisor provided lower-quality advice, participants trusted this advisor more than the conscientious one.

Our findings highlight that, while personalized LLM-advisors represent a promising research direction, their use in high-stakes domains like finance is not free of risks: due to the limitations of LLMs at capturing complex investment preferences, and the difficulty of investors to discern whether the advice they receive truly serves their interests, LLMs have a notable risk to drive investors to bad financial assets (leading not only to a low satisfaction but also to potentially large monetary losses). However, these drawbacks open interesting research directions not only from a system perspective, but also from a human-centered approach: automated advisory development where we do not just focus on improving the quality of automated systems to guide investors, but also on how the investors will adopt, trust and interact with these AI agents [5].

Acknowledgments

This work was supported by Daiwa Securities Group Inc.

References

- [1] James E. Allen, Curry I. Guinn, and Eric Horvitz. 1999. Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications* 14, 5 (1999), 14–23.
- [2] Ashay Argal, Siddharth Gupta, Ajay Modi, Pratik Pandey, Simon Shim, and Chang Choo. 2018. Intelligent travel chatbot for predictive recommendation in echo platform. In *2018 IEEE 8th Annual Computing and Communication Workshop and Conference (CCWC 2018)*. IEEE, 176–183.
- [3] Wanling Cai, Yucheng Jin, and Li Chen. 2022. Impacts of personal characteristics on user trust in conversational recommender systems. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI 2022)*. Article 489, 14 pages.
- [4] Gary Charness, Uri Gneezy, and Alex Imas. 2013. Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior & Organization* 87 (2013), 43–51.
- [5] Erin K. Chiou and John D. Lee. 2023. Trusting automation: Designing for responsibility and resilience. *Human factors* 65, 1 (2023), 137–165.
- [6] Konstantina Christakopoulou, Filip Radlinski, and Katja Hofmann. 2016. Towards conversational recommender systems. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining (KDD 2016)*. 815–824.
- [7] Berardina De Carolis, Marco de Gemmis, Pasquale Lops, and Giuseppe Palestra. 2017. Recognizing users feedback from non-verbal communicative acts in conversational recommender systems. *Pattern Recognition Letters* 99 (2017), 87–95.
- [8] Eugene F Fama and Kenneth R French. 1998. Value versus growth: The international evidence. *The journal of finance* 53, 6 (1998), 1975–1999.
- [9] Christian Hildebrand and Anouk Bergner. 2021. Conversational robo advisors as surrogates of trust: onboarding experience, firm perception, and consumer financial decision making. *Journal of the Academy of Marketing Science* 49, 4 (2021), 659–676.
- [10] Dietmar Jannach, Ahtsham Manzoor, Wanling Cai, and Li Chen. 2021. A survey on conversational recommender systems. *Comput. Surveys* 54, 5 (2021), 1–36.
- [11] Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2024. Evaluating and inducing personality in pre-trained language models. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS 2023)*.
- [12] Hang Jiang, Xiajie Zhang, Xubo Cao, Cynthia Breazeal, Deb Roy, and Jad Kabbara. 2024. PersonalLLM: Investigating the Ability of Large Language Models to Express Personality Traits. In *Findings of the Association for Computational Linguistics: NAACL 2024*. 3605–3627.
- [13] Francis M. Kinniry Jr., Colleen M. Jaconetti., Michael A. DiJoseph., Yan Zilbering., Donald G. Bennyhoff., and Georgina Yarrowood. 2020. *Putting a value on your value: Quantifying Vanguard Adviser's Alpha in the UK*. Technical Report. The Vanguard Group, Valley Forge, Pennsylvania, USA.
- [14] Sherrie Y.X. Komiak and Izak Benbasat. 2006. The effects of personalization and familiarity on trust and adoption of recommendation agents. *MIS quarterly* (2006), 941–960.
- [15] Ivica Kostrić, Krisztian Balog, and Filip Radlinski. 2021. Soliciting user preferences in conversational recommender systems via usage-related questions. In *Proceedings of the 15th ACM Conference on Recommender Systems*. 724–729.
- [16] Kausik Lakkaraju, Sara E. Jones, Sai Krishna Revanth Vuruma, Vishal Pallagani, Bharath C. Muppasani, and Biplav Srivastava. 2023. LLMs for Financial Advice: A Fairness and Efficacy Study in Personal Decision Making. In *Proceedings of the 4th ACM Conference on AI in Finance (ICAIF 2023)*. 100–107.
- [17] Cong Li. 2016. When does web-based personalization really work? The distinction between actual personalization and perceived personalization. *Computers in human behavior* 54 (2016), 25–33.
- [18] Andrew W. Lo and Jillian Ross. 2024. Can ChatGPT Plan Your Retirement?: Generative AI and Financial Advice. *Harvard Data Science Review* (2024). Issue Special Issue 5.
- [19] Tim Loughran and Bill McDonald. 2011. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of finance* 66, 1 (2011), 35–65.
- [20] Robert R. McCrae and Oliver P. John. 1992. An introduction to the five-factor model and its applications. *Journal of personality* 60 2 (1992), 175–215.
- [21] Sourav Medya, Mohammad Rasoolinejad, Yang Yang, and Brian Uzzi. 2022. An Exploratory Study of Stock Price Movements from Earnings Calls. In *Companion Proceedings of the Web Conference 2022 (WWW 2022)*. Association for Computing Machinery, 20–31.
- [22] Pearl Pu, Li Chen, and Rong Hu. 2011. A user-centric evaluation framework for recommender systems. In *Proceedings of the 5th ACM conference on Recommender Systems (RecSys 2011)*. 157–164.
- [23] Filip Radlinski, Krisztian Balog, Bill Byrne, and Karthik Krishnamoorthi. 2019. Coached conversational preference elicitation: A case study in understanding movie preferences. In *Proceedings of the 20th Annual SIGDial Meeting on Discourse and Dialogue (SIGDIAL 2019)*. 353–360.
- [24] Filip Radlinski and Nick Craswell. 2017. A theoretical framework for conversational search. In *Proceedings of the 2nd Conference on Human Information Interaction and Retrieval (CHIIR 2017)*. 117–126.
- [25] Oscar Sainz, Jon Campos, Iker García-Ferrero, Julen Etxaniz, Oier Lopez de Lacalle, and Eneko Agirre. 2023. NLP Evaluation in trouble: On the Need to Measure LLM Data Contamination for each Benchmark. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, 10776–10787.
- [26] Tetsuya Sakai. 2018. Laboratory experiments in information retrieval. *The information retrieval series* 40 (2018), 4.
- [27] Javier Sanz-Cruzado, Edward Richards, and Richard McCreadie. 2024. FAR-AI: A Modular Platform for Investment Recommendation in the Financial Domain. In *Proceedings of the 46th European Conference on Information Retrieval (ECIR 2024)*, Part V. Springer-Verlag, 267–271.
- [28] Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. Character-LLM: A Trainable Agent for Role-Playing. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP 2023)*. Association for Computational Linguistics, 13153–13187.
- [29] Tuva Lunde Smestad and Frode Volden. 2019. Chatbot personalities matters: improving the user experience of chatbot interfaces. In *5th International Conference Internet Science: (INSCI 2018)*. Springer, 170–181.
- [30] David J Streich. 2023. Risk preference elicitation and financial advice taking. *Journal of Behavioral Finance* 24, 3 (2023), 259–275.
- [31] Yueming Sun and Yi Zhang. 2018. Conversational recommender system. In *Proceedings of the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2018)*. 235–244.
- [32] Takehiro Takayanagi, Chung-Chi Chen, and Kiyoshi Izumi. 2023. Personalized Dynamic Recommender System for Investors. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2023)*. Association for Computing Machinery, 2246–2250.
- [33] Takehiro Takayanagi, Kiyoshi Izumi, Atsuo Kato, Naoyuki Tsunedomi, and Yukina Abe. 2023. Personalized Stock Recommendation with Investors' Attention and Contextual Information. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2023)*. Association for Computing Machinery, 3339–3343.
- [34] Takehiro Takayanagi, Masahiro Suzuki, Kiyoshi Izumi, Javier Sanz-Cruzado, Richard McCreadie, and Iadh Ounis. 2025. FinPersona: An LLM-Driven Conversational Agent for Personalized Financial Advising. In *Proceedings of the 47th European Conference on Information Retrieval (ECIR 2025)*, Part V. Springer-Verlag, 13–18.
- [35] Johanne R. Trippas, Sara Fahad Dawood Al Lawati, Joel Mackenzie, and Luke Gallagher. 2024. What do Users Really Ask Large Language Models? An Initial Log Analysis of Google Bard Interactions in the Wild. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2024)*. 2703–2707.
- [36] Johanne R. Trippas, Luke Gallagher, and Joel Mackenzie. 2024. Re-evaluating the Command-and-Control Paradigm in Conversational Search Interactions. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM 2024)*. Association for Computing Machinery, 2260–2270.
- [37] Patchara Vanichvasin. 2021. Chatbot Development as a Digital Learning Tool to Increase Students' Research Knowledge. *International Education Studies* 14, 2 (2021), 44–53.
- [38] Xuena Wang, Xueting Li, Zi Yin, Yue Wu, and Jia Liu. 2023. Emotional intelligence of large language models. *Journal of Pacific Rim Psychology* 17 (2023), 18344909231213958.
- [39] Pontus Wärnestål. 2005. User evaluation of a conversational recommender system. In *Proceedings of the 4th Workshop on Knowledge and Reasoning in Practical Dialogue Systems*.
- [40] Hamed Zamani, Johanne R Trippas, Jeff Dalton, Filip Radlinski, et al. 2023. Conversational information seeking. *Foundations and Trends® in Information Retrieval* 17, 3-4 (2023), 244–456.
- [41] Markus Zanker, Laurens Rook, and Dietmar Jannach. 2019. Measuring the impact of online personalisation: Past, present and future. *International Journal of Human-Computer Studies* 131 (2019), 160–168.
- [42] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W Bruce Croft. 2018. Towards conversational search and recommendation: System ask, user respond. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM 2018)*. 177–186.
- [43] Huaqin Zhao, Zhengliang Liu, Zihao Wu, Yiwei Li, Tianze Yang, Peng Shu, Shaochen Xu, Haixing Dai, Lin Zhao, Gengchen Mai, et al. 2024. Revolutionizing Finance with LLMs: An Overview of Applications and Insights. *arXiv preprint arXiv:2401.11641* (2024).
- [44] Dávid Zibriczy. 2016. Recommender systems meet finance: a literature review. In *Proceedings of the 2nd International Workshop on Personalization & Recommender Systems in Financial Services (FinRec 2016)*. 1–10.
- [45] Liv Zieffeld, Daan Di Scala, and Anita HM Cremers. 2025. The effect of preference elicitation methods on the user experience in conversational recommender systems. *Computer Speech & Language* 89 (2025), 101696.