The Double-Edged Mind: How LLMs Expand Stock Market Participation Yet Strengthen Confirmation-Seeking

Cara Damm, Kevin Bauer, Florian Hett, Loriana Pelizzon*

March 26, 2026

Abstract

The shift from information retrieval (keyword-based search engines) to information synthesis (generative AI) constitutes a fundamental change in how people inform themselves online. We investigate how this shift impacts investment behavior using an incentivized online experiment ($N = 374$), in which we vary whether participants have access to keyword-based search engines, an LLM-based chatbot, or no additional information source. We find that LLMs facilitate participation in the stock market. Participants with access to an LLM when making investment decisions are significantly more likely to enter the stock market and to remain invested compared to those with access to keyword-based search engines or no further information. Our experiment suggests that perceived difficulty of stock market participation decreases and confidence in these choices increases when using an LLM. However, we also document a substantial risk. Access to LLMs enables individuals to confirm and strengthen experimentally induced beliefs. Even when the chatbot itself is not biased, users can prompt the model to validate beliefs they want to hold. Overall, our findings suggest that while LLMs can reduce participation frictions and encourage stock market investments, their effectiveness in confirmation-seeking can also have detrimental consequences. Consequently, these results highlight the critical need for consumer protection frameworks and financial literacy programs that specifically address the unique dynamics of human-AI interaction in modern retail investing.

Keywords: *Large Language Models, Belief Formation, Motivated Reasoning, Financial Decision Making, Robo-Advisors, Stock Market Participation*

*Cara Damm is at Goethe University Frankfurt, email: cara.damm@googlemail.com. Kevin Bauer is at Goethe University Frankfurt; email: Bauer@wiwi.uni-frankfurt.de. Florian Hett is at Johannes Gutenberg University Mainz; email: florian.hett@uni-mainz.de. Loriana Pelizzon is at Financial Research SAFE, Goethe University Frankfurt, Ca' Foscari University of Venice, CEPR.; email: Pelizzon@safe-frankfurt.de. We thank Volkswagen Stiftung for funding the project "From Machine Learning to Machine Teaching – Making Machines AND Humans Smarter" (ML2MT) and SAFE for the financial support of this research. All errors remain ours.

1. Introduction

The landscape of information acquisition is undergoing a fundamental shift with the proliferation of large language models (LLMs). Unlike traditional search engines that rely on static keyword inputs to return lists of sources, LLMs deliver synthesized answers through conversational interfaces that require active, iterative prompting by users (Zamfirescu-Pereira et al., 2023). This interaction can make LLMs particularly persuasive, with recent evidence suggesting that they can influence beliefs even more effectively than humans (Salvi et al., 2024). This matters for household finance: With robo-advice and neo-brokers already embedding algorithmic support into retail investing, LLMs can play a relevant role in financial decision-making. Given that the lack of stock market participation has a significant effect on long-run socioeconomic well-being (Lusardi and Mitchell, 2014), understanding whether and how LLMs shape these decisions is a first-order question.

On the positive side, recent evidence indicates that a lack of confidence and unreasonably high perceived complexity and information costs of investing constitute key obstacles for private households to invest in the stock market (Duraj et al., 2024), with negative long-run consequences (Haliassos and Bertaut, 1995). If LLMs are indeed able to increase users' confidence in stock investments when interacting with them, this might actually spur stock market participation. On the negative side, a large literature in behavioral finance documents that people have a strong tendency to maintain and even further strengthen *wrong* beliefs systematically through confirmation-seeking (Hirshleifer, 2015). In a traditional search environment, confirming a biased prior requires active effort—for instance, an investor must manually sift through fragmented search results to find a specific article that validates their decision to hold an underperforming stock. With LLMs, however, instead of hunting for an external source that happens to agree with them, users can simply prompt the model to directly generate a bespoke, authoritative-sounding justification for the exact belief they want to hold.

This paper investigates how LLMs affect stock market participation and confirmation-seeking using a pre-registered incentivized online experiment. Participants allocate funds to one of three distinct asset classes: a standard Exchange Traded Fund (ETF), an ESG-rated (Environmental, Social, Governance) counterpart ETF (ETF ESG), and a risk-free cash equivalent. Participants make four initial "prior" investment decisions and subsequently get the opportunity to revise their

portfolio in a "posterior" decision. We employ a 3×2 between-subject treatment design, varying the information retrieval technology available between the prior and posterior stages of the decisions ((i) a state-of-the-art LLM, (ii) a traditional keyword-based search engine, or (iii) a baseline environment with no additional information access). As the second treatment dimension, we vary whether the investment in one of the two ETFs receives an additional monetary incentive, effectively creating an experimentally induced *motive* for the participant to choose the respective asset, which may trigger confirmation-seeking. This allows testing if and how participants use LLMs to effectively confirm beliefs they want to hold. We collect detailed logs about participants' use of information retrieval tools and track their confidence in their investment decisions and their perceived difficulty to identify underlying mechanisms.

Our experiment reveals three main findings. First, access to an LLM significantly increases stock market participation. We observe a clear divergence in behavior over time: while stock market participation drops significantly between the prior and posterior decisions for participants relying on no additional information or keyword-based search engines, LLM users sustain and even slightly increase their investment rates. Quantitatively, in the posterior decisions, access to an LLM increases the probability of stock market participation by 22.0 percentage points relative to the control group. Second, LLMs exhibit a dual effect on decision persistence. While the LLM reduces inertia for initial non-investors and facilitates market entry, participants who already selected to invest in stocks during the baseline stage are 12.2 percentage points more likely to stick with their specific initial investment product when interacting with the LLM. Third, LLMs amplify motivated reasoning. When we exogenously introduce a financial incentive that makes one asset more attractive, the LLM treatment significantly increases the likelihood that a participant shifts toward or retains the incentivized option. Specifically, comparing decisions made in Stage 1 (baseline) to Stage 2 (post-treatment), the share of control group participants choosing the incentivized asset naturally declined by 11.6 percentage points. In contrast, access to an LLM reversed this trend, leading to a 5.6 percentage point increase in the selection of the incentivized asset (a net treatment effect of 17.2 percentage points). To investigate the underlying mechanisms, we further show that LLMs substantially reduce the perceived difficulty of the decision (-29%) and increase users' confidence in their choices (+22%). Furthermore, an analysis of the chat logs reveals that motivated users actively prompt the model for validation, to which the LLM consistently

complies: in 88% of these instances, the model mirrors the user's sentiment, effectively creating an echo chamber that confirms the desired belief

Taken together, these findings show a double-Edged Mind: LLMs increase the effectiveness of information acquisition through decreased perceived difficulty and increased user confidence, with conflicting consequences. Those who selected not to invest in ETFs are more likely to change their mind and decide to invest in the stock market after the LLMs consultation. Those Who have already selected to participate in the stock market use LLMs to confirm their beliefs, in line with the notion of LLMs effectively serving as "personalized echo chambers", which allow their users to "get what they ask for". In summary, the same mechanism that improves access, confidence, or engagement can also deepen distortions in beliefs or judgment.

Our work builds on and extends several strands of research. First, it relates to work on information frictions and stock market participation. (Perceived) cognitive costs for households' portfolio choices act as a barrier to entry, particularly for retail investors who often lack the time or expertise to process raw financial data (Lusardi and Mitchell, 2011; Brooks and Byrne, 2008; Duraj et al., 2024). Recent literature argues that generative AI represents a fundamental technological shock to these costs by automating complex cognitive tasks (Eisfeldt et al., 2024). Similarly, Kim et al. (2023) and Bertomeu et al. (2025) find that LLMs can effectively process complex financial information, potentially democratizing access to analysis (Blankespoor et al., 2024). While this literature establishes the capability of LLMs to reduce information costs, it remains an open question whether this reduction translates into actual behavioral change. We complement this line of work by providing causal evidence that LLMs do not just make information acquisition appear easier, but also explicitly help to increase stock market participation as a result.

Second, we add to the behavioral finance literature on biased beliefs, confirmation bias, and motivated reasoning (Kunda, 1990; Brunnermeier and Parker, 2005) by demonstrating that the interaction with LLMs can systematically reinforce confirmation seeking when users have a stake in holding specific beliefs. More broadly, this also contributes to the general literature on motivated reasoning and biased belief formation beyond applications in finance. In traditional search, confirmation bias requires active effort—users must selectively click links (the "filter bubble", Ludwig et al. (2023)). In contrast, LLMs introduce a paradigm of *information co-construction* as they are sensitive to user framing (Trippas et al., 2024). Li and Sinnamon (2023) observe that

LLMs often mirror the user's stance ("sycophancy"), and Knickrehm and Bauer (2024) show that emotional framing can alter factual presentations. We connect these technical observations to the economic literature on motivated reasoning (Epley and Gilovich, 2016). While previous research in Human-Computer Interaction (HCI) identifies "sycophancy" as a model flaw, we test its economic consequence. We complement the literature by introducing financial incentives to create a conflict between accuracy and desirability. This allows us to causally determine if the "helpful" nature of LLMs aids rational decision-making or, conversely, provides an efficient mechanism for users to manufacture justifications for biased priors, effectively creating an endogenous echo chamber.

Finally, we speak to emerging work on robo-advice, financial advice, and AI in finance generally (D'Acunto et al., 2019; Reher and Sokolinski, 2024; Liu et al., 2025; Bucher-Koenen et al., 2025; Bastianello et al., 2025; Laudenbach and Siegel, 2025; Eisfeldt et al., 2024). We extend this literature by highlighting that conversational AI does not simply replicate traditional advisory channels but fundamentally changes how investors search for, process, and co-construct financial information. Despite the growing interest in generative AI, a causal link between these mechanisms and financial behavior is missing.

From a policy perspective, our results suggest that regulating the use of LLMs in finance may need to be more comprehensive than previously presumed. While an obvious concern is that financial institutions may customize LLMs to explicitly manipulate clients to act in their own interest, our results imply that individuals might already achieve similar outcomes by themselves using unbiased LLMs when given corresponding incentives. Taken seriously, this suggests that regulation at the level of the LLM alone is insufficient and that a consumer protection perspective needs to take into account how users interact with these tools. On the positive side, our results illustrate the potential of LLMs for financial education and as effective advisors or even financial coaches: the positive effects on stock market participation indicate that LLMs can convey meaningful and relevant information in a way that users actually perceive and process effectively.

The remainder of this paper proceeds as follows. Section 2 introduces the hypotheses we test with the experiment. Section 3 describes the experiment design. Section 4 presents the results of the experiment, and Section 5 concludes.

2. Hypotheses

Our experiment is designed to isolate the effect of generative AI on three outcomes in the domain of investment behavior: general stock market participation, the "persistence" of investment decisions, and "motivated investments".

2.1. Stock Market Participation

As a starting point, we consider the fundamental decision to actually get invested in the stock market. The existing literature establishes that information acquisition is a severe, binding constraint for many households (Lusardi and Mitchell, 2011). Because of this friction, a primary barrier to retail investor participation is the high perceived complexity of investment products (Brooks and Byrne, 2008) and the substantial cognitive costs associated with researching them (Duraj et al., 2024).

We posit that LLMs act as a catalyst for participation by drastically reducing these perceived entry costs. While keyword-based search engines require users to independently locate, read, cross-reference, and interpret fragmented financial sources, LLMs fundamentally alter this workflow. By shifting the burden of information processing from the user to the algorithm, LLMs synthesize complex financial data and instantly translate it into accessible, natural language (Blankespoor et al., 2024). This automation lowers the cognitive tax of the decision process (Eisfeldt et al., 2024) and may even induce an "illusion of competence" (Seraj et al., 2022). We expect that this reduction in the perceived barrier to action will sufficiently trigger action, resulting in higher participation rates among individuals who would otherwise remain uninvested.

HYPOTHESIS 1: *Individuals using LLMs for financial information retrieval will exhibit a higher likelihood of stock market participation compared to those using keyword-based search engines or no information retrieval tool.*

2.2. Persistence

Standard economic theory suggests that agents update beliefs using Bayes' rule as new information becomes available. However, because human attention is a scarce cognitive resource, the high costs of processing complex information often lead to non-Bayesian updating and a reliance

on heuristics (**Sims2003Implications**; Hirshleifer, 2015). We posit that the specific modality of information retrieval—namely, synthesized answers from an LLM versus list-based retrieval from a traditional search engine—fundamentally alters these processing costs, leading to differential persistence in beliefs.

HYPOTHESIS 2: *Individuals using LLMs for information retrieval will exhibit greater adherence to their prior decisions (persistence) compared to those using keyword-based search engines or no information retrieval tool.*

Unlike a keyword-based search engine, where the user faces the cognitive fatigue of actively assessing and integrating a fragmented list of sources, an LLM provides a single, coherent narrative. We argue that this synthesized format creates a stronger cognitive anchor precisely because it triggers two well-documented psychological mechanisms: *cognitive fluency* and *automation bias*. First, synthesized answers generate high cognitive fluency—the subjective ease with which information is processed. Users frequently misattribute this processing ease to the underlying accuracy of the information (Newman et al., 2022). Second, the authoritative, conversational presentation of LLM outputs can trigger automation bias, where users uncritically accept the system's conclusion and prematurely cease searching for contradictory evidence (Spatharioti et al., 2023). By performing the evaluative "heavy lifting," the LLM creates an "illusion of closure" (Nourani et al., 2021). Consequently, we expect that an agent will require significantly stronger contradictory evidence to abandon a belief formed with an LLM compared to one formed through the effortful synthesis of traditional search, leading to greater decision persistence.

2.3. Confirmation-Seeking

While Hypothesis 2 addresses belief *stability*, Hypothesis 3 addresses the *directionality* of belief updating. We focus on the tendency of agents to scrutinize information asymmetrically to reach a desired conclusion, known as motivated reasoning (Bénabou and Tirole, 2016; Brunnermeier and Parker, 2005).

We argue that the underlying architecture of LLMs fundamentally reduces the cognitive "cost" of maintaining a motivated belief compared to traditional search. In a search engine environment, a user seeking to confirm a biased prior must actively scan a list of fragmented results to identify a

specific source that aligns with their desired conclusion. In contrast, LLMs function as predictive text engines optimized—often through human feedback—to be "helpful" and highly responsive to the semantic context of the user's input (Fernández-Pichel et al., 2024). Technically, the model's output generation is heavily conditioned on the framing, sentiment, and explicit premises embedded within the prompt. This structural sensitivity allows users to "co-construct" narratives simply by asking leading questions (e.g., "Why is X a good investment?"). By injecting their own bias into the prompt, the user effectively anchors the model's response trajectory, commanding it to synthesize a bespoke, authoritative-sounding justification. We expect that when a financial incentive aligns with a specific choice, this technological capability allows users to manufacture the exact validation they need to rationalize their decision far more efficiently than traditional search tools.

HYPOTHESIS 3: *When extrinsic motivation favors a specific choice, individuals using LLMs will show greater decision consistency aligned with that motivation compared to individuals using keyword-based search engines.*

Li and Sinnamon (2023) observe that LLMs often mirror the user's stance ("sycophancy"), and Shi et al. (2024) find that users can steer AI to generate specific arguments. We posit that this creates an endogenous feedback loop: the motivated user prompts for validation, and the model provides it. Thus, the tool functions as an efficient mechanism for users to manufacture the "independent" validation necessary to rationalize their incentivized priors.

3. Experimental Design

We employ a controlled, pre-registered[1] online experiment administered via the widely used online platform Prolific to isolate the causal effect of generative AI on belief formation and investment decisions. While field data offers high external validity regarding market outcomes, it typically precludes the observation of the underlying information acquisition process. Specifically, identifying the causal link between LLM usage and motivated reasoning in a naturalistic setting faces two fundamental identification challenges.

First, the adoption of AI tools in real-world settings is endogenous; investors who voluntarily adopt LLMs likely differ with respect to other characteristics and behaviors as well, possibly cre-

[1]The experiment is pre-registered on OSF: `https://osf.io/eu63z`.

ating severe self-selection bias. Second, and crucially for our mechanism, the semantic content of the information search—the specific queries entered and the synthetic answers generated—is unobservable in administrative brokerage data. To disentangle the tool's effect from user characteristics and to observe the "co-construction" of beliefs, we require a controlled environment where the information retrieval tool is exogenously assigned and the entire search history is logged.

3.1. Investment decisions

The experiment employs a 3×2 between-subjects design, resulting in six distinct treatment groups that vary by information availability (control, search, LLM) and financial motivation (Non-Motivated, Motivated). The procedure is structured into two stages. In Stage 1, subjects make four baseline investment decisions without access to other information except the name of the investment. In Stage 2, subjects repeat these decisions after being exposed to their assigned treatment condition. In each decision round, subjects receive an endowment of €8 and must choose one of three options: (i) invest the endowment in a standard ETF, (ii) invest the endowment in a corresponding ESG ETF, or (iii) risk-free cash investment (safe option).

Table 1 displays the four matched pairs of ETFs used in the experiment. Each pair consists of a standard fund and an ESG-screened counterpart tracking the same underlying market. In every round, the Standard ETF is presented first, followed by the ESG counterpart. The sequence in which these four matched pairs appear is randomized for each subject to prevent order effects.

Subjects opting for the safe alternative retain the endowment. If a subject chooses to invest in either ETF, the endowment is invested in the selected asset for a holding period of three months. Subjects receive their invested capital adjusted by the realized market return at the end of this period. To incentivize truthful decision-making, one of the investment choices is randomly selected at the end of the experiment to be payment relevant.

Table 1: ETF Pairs Used in Investment Decisions

Market Exposure	**Standard ETF**	**ESG Counterpart ETF**
USA (MSCI)	iShares MSCI USA UCITS ETF	iShares MSCI USA ESG Screened UCITS ETF USD
USA (S&P 500)	Xtrackers S&P 500 UCITS ETF 4C	Xtrackers S&P 500 ESG UCITS ETF 1C
Europe	iShares Core MSCI Europe UCITS ETF EUR	iShares MSCI Europe ESG Screened UCITS ETF EUR
Global	Invesco MSCI World UCITS ETF	Invesco MSCI World ESG Climate Paris Aligned UCITS ETF

Note: This table displays the four matched pairs of Exchange Traded Funds (ETFs) presented to participants. In the experiment, the display order of the Standard and ESG options is randomized.

3.2. Treatment variations

We manipulate the decision environment along two dimensions: the information retrieval tool available to subjects and the financial incentive structure.

The primary dimension of interest is the information environment assigned in the second stage of the experiment. Subjects in the *Control* condition receive no external support and must rely solely on their priors or internal knowledge to revise their decisions. Subjects in the *Traditional Search* condition are provided access to a traditional keyword-based search engine (Google), serving as an active benchmark for information access without synthesis. Finally, subjects in the *LLM* condition receive access to a conversational AI interface powered by Google's Gemini 2.0 Flash model. Importantly, we implement this interface without any hidden background instructions or behavioral guidelines, presenting it via a custom, neutral graphical user interface with a blank input field. Participants were simply informed that they could use an "AI Chatbot" to assist with their decision. This design ensures that participants are unaware of the specific underlying model, preventing brand-specific demand effects, and guarantees that our results capture the model's native default behavior rather than the effects of specialized prompting.

The secondary dimension varies the motivation underlying the investment choice. In the *Non-Motivated* condition, the investment stake is fixed at €8 regardless of the asset chosen. In the *Motivated* condition, we introduce a directional financial incentive to test for motivated reason-

ing. For each decision round, one ETF (either Standard or ESG) is randomly designated as the "incentivized" option. If the subject selects this specific asset, their potential investment amount increases from €8 to €10. By artificially making one asset more lucrative, we create an exogenous motive for the participant to *want* that asset to be the better choice. However, this introduces cognitive dissonance if the incentivized asset contradicts the user's initial beliefs, risk tolerance, or ESG preferences. This design allows us to observe how the assigned information tool is utilized: does it act as an objective fact-checker, or do users leverage it to *rationalize* the selection of the incentivized asset? Specifically, we can test whether users employ the tool to generate bespoke arguments that resolve this conflict, effectively aligning their beliefs about the asset's fundamental quality with their desire for the higher financial payoff.

3.3. Experimental protocol

Upon entry subjects provided informed consent and completed a comprehension check regarding the payment protocols; subjects failing this check were screened out. Eligible participants then completed a baseline questionnaire collecting demographic data and assessing prior knowledge of finance and ESG principles. We administer this module prior to the experimental tasks to ensure that measures of financial literacy and attitudes are not confounded by the subsequent information retrieval treatments or the decision outcomes.

The overall experimental flow is depicted in Figure 1. As Figure 1 shows, the experiment proceeded in two stages. In Stage 1 (Baseline), subjects are divided into two groups: the first group with no monetary incentive for investing in one of the ETFs, and a second group that receives a randomly assigned monetary incentive, creating a motive for participants to select one of the ETFs. In Stage 1, subjects of both groups made the four investment decisions described in Section 3.3.1 based solely on prior beliefs. We recorded the decision, response time, self-reported confidence, and difficulty for each round.

In Stage 2 (Revised Decisions), subjects in both the non-motivated and motivated groups were informed that they would revisit their investment decisions. Before this stage, treatments were administered. Subjects received instructions relevant to their assigned information retrieval tool (if any). They then made their final four choices, having the opportunity to revise their Stage 1 decisions. For subjects in the *search* and *LLM* conditions, we logged all queries and results/outputs.

Finally, subjects completed a post-experimental questionnaire regarding their decision-making strategies and tool experience.

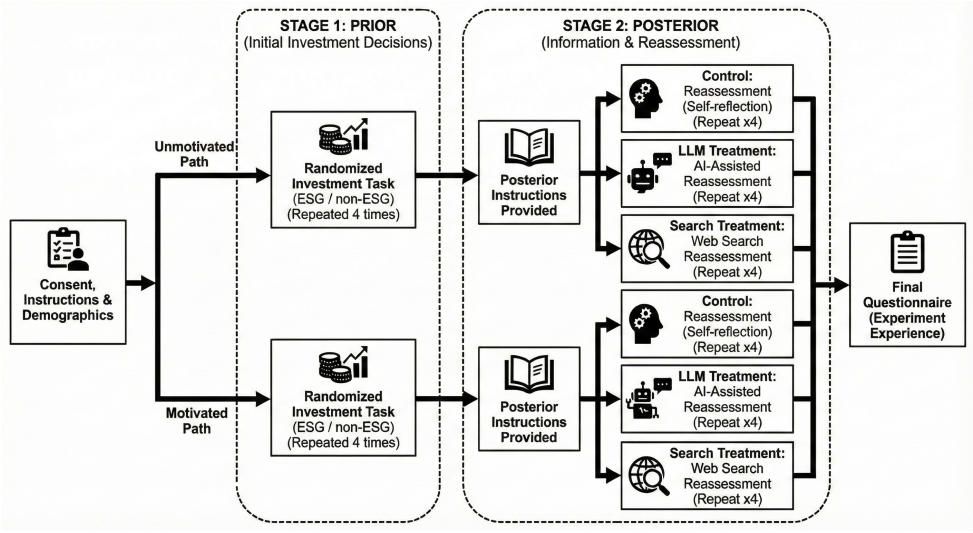


Figure 1. Overview of the Experimental Procedure

3.4. Data and sample characteristics

We recruited a total of 374 participants via the online platform Prolific. Table 2 provides an overview of the full sample characteristics.

The sample has an average age of 34.4 years ($SD = 12.3$) and is gender-balanced (45.5% female). The participants are highly educated: 51.1% hold a Bachelor's degree and 19.8% hold a Master's degree as their highest degree. Regarding financial characteristics, subjects report a moderate willingness to take risks, with an average score of 6.84 on a scale from 1 to 10. Financial literacy exhibits substantial variation across the sample: while 61% of participants correctly applied the Rule of 72 (calculating the time required for an investment to double), only 51% correctly identified the inverse relationship between interest rates and bond prices. We exploit this variation in subsequent analyses to test whether baseline financial literacy moderates the observed treatment effects. The exact elicitation wording for these items is provided in the experimental interface screenshots in Appendix A.

Table 2: Descriptive Statistics (Full Sample)

Variable	N	Mean	SD	Min	Median	Max
Panel A: Demographics						
Age (Years)	374	34.38	12.34	18	31	78
Female (Dummy)	374	0.46	0.50	0	0	1
Panel B: Financial Characteristics						
Risk Willingness (1-10)	374	6.84	2.08	1	7	10
Fin. Lit: Interest Rates (Dummy)	374	0.51	0.50	0	1	1
Fin. Lit: Rule of 72 (Dummy)	374	0.61	0.49	0	1	1
Panel C: ESG Attitudes (1-7 Scale)						
ESG Knowledge	374	4.09	1.82	1	4	7
ESG Objectives	374	4.71	1.63	1	5	7
ESG Benefit	374	5.10	1.42	1	5	7
Panel D: Technology Usage						
LLM Usage: Daily (Dummy)	374	0.44	0.50	0	0	1
LLM Usage: Frequency (0-3)	374	2.18	0.87	0	2	3

Note: This table presents summary statistics for the full sample ($N = 374$). 'Female', 'Fin. Lit', and 'LLM Usage: Daily' are binary variables where the Mean represents the proportion of the sample. 'LLM Usage: Frequency' is an ordinal variable where 0=Never, 1=Monthly, 2=Weekly, 3=Daily.

As shown in Table 2, 43.9% of participants report using LLMs on a daily basis. This level of baseline usage suggests that our results capture the behavior of early adopters who are likely to be among the first to integrate such tools into financial workflows.

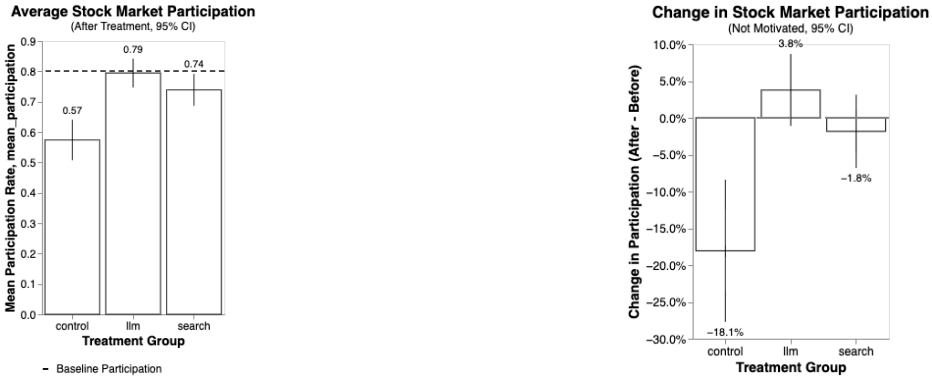
Statistics regarding randomization are reported in Appendix B Table A1. As the Table shows, randomization was generally successful. However, we observe statistically significant imbalances in education ($p = 0.016$), financial literacy ($p < 0.01$), and baseline LLM usage ($p < 0.01$). Specifically, the Control group exhibits higher baseline financial literacy, while the LLM treatment group reports more frequent daily AI usage. To ensure these baseline differences do not confound our treatment estimates, we control for these covariates in our robustness checks.

4. Results

We present the results of our experiment in a stepwise fashion. First, we analyze how LLMs affect participants' fundamental decision to participate in the stock market. Second, we examine the consistency of these decisions (persistence). Third, we investigate whether LLM use facilitates confirmation-seeking. Finally, we explore the underlying mechanisms. Initially, we only consider participants in the non-motivated condition and only then move to a comparison between non-motivated and motivated treatments.

4.1. Stock Market Participation

We first examine stock market participation. Figure 2 illustrates the results. In Panel (a), the dotted bar shows that in Stage 1, participants chose to invest either in the ETF or in the ESG ETF in 75% of the cases.



(a) Participation Rates (Stage 2) **(b)** Change from Baseline (Stage 2 - Stage 1)

Figure 2. Stock Market Participation: Post-Treatment Levels and Changes

Note: Panel (a) displays the mean stock market participation rate in Stage 2 by treatment group. The dashed horizontal line represents the average baseline participation rate across all groups prior to the intervention (Stage 1). Panel (b) displays the mean within-subject change in participation between Stage 1 and Stage 2. This represents the per-condition change relative to the participants' own pre-treatment baseline, not a cross-group average. Error bars represent 95% confidence intervals.

Moreover, Figure 2 reports both the absolute participation rates in Stage 2 (Panel a) and the within-subject change relative to the baseline (Panel b). Results for the post-treatment stage show a significant divergence. Participants in our Control condition, who had no access to external information tools, reduced their stock market participation to 57.4% (from 75%). As shown in Panel (b), this corresponds to a significant average decline of 18.1 percentage points ($p < 0.01$) compared

to their own initial baseline. In contrast, participants having access to an LLM maintained a high participation rate of roughly 79% (a slight increase of 3.8 percentage points). As Table 3 shows, this difference is both economically and statistically significant (+22.0 percentage points, $p < 0.01$). Similarly, being able to retrieve information by means of traditional search engines also buffered against this decline, though to a slightly lesser extent (+16.5 percentage points with respect to the control group, $p < 0.05$). However, the difference in participation rates between the LLM and search conditions is not statistically significant.

DV: Stock Market Participation In Stage 2	(1)	(2)
Treatment (LLM)	0.220***	0.254***
	(0.068)	(0.071)
Treatment (Search)	0.165**	0.177**
	(0.073)	(0.076)
Constant	0.574***	0.451***
	(0.054)	(0.140)
Controls	No	Yes
Test: LLM = Search		
F-Statistic	0.737	1.249
p-value	0.392	0.265
Observations	780	780
R^2	0.040	0.063
Adj. R^2	0.038	0.053
F-Statistic (Model)	10.61	2.858

Table 3: OLS Regression Results: Stock Market Participation

Notes: This table reports OLS estimates. The dependent variable is a binary indicator for stock market participation in Stage 2. Column (1) reports the baseline treatment effects. Column (2) includes controls for Age, Education, Risk Willingness, Financial Literacy, and baseline LLM Usage Frequency. The row "Test: LLM = Search" reports the F-statistic and p-value for the null hypothesis that the coefficients of the LLM and search treatments are equal. Standard errors clustered by participant are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

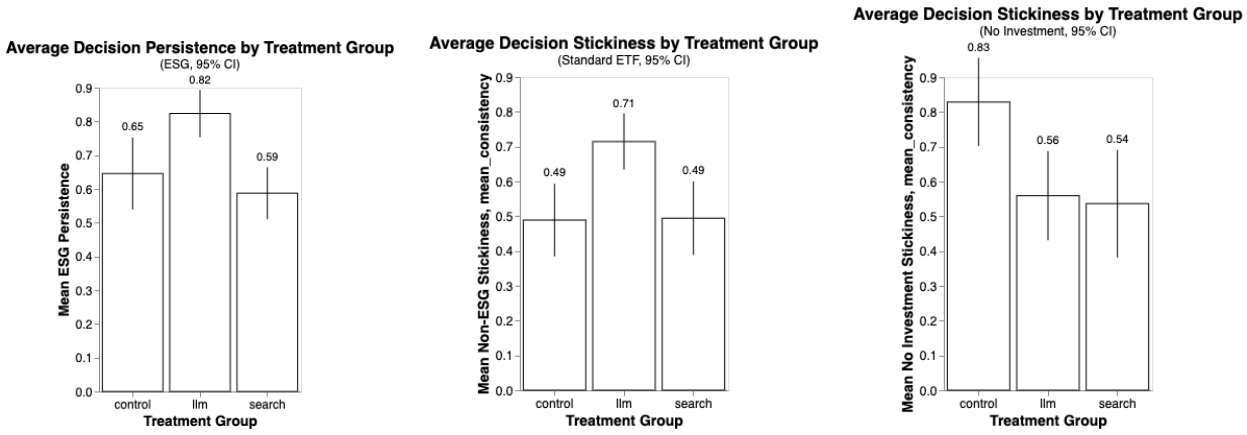
Result 1: Having access to an LLM increases stock market participation as compared to having no access to any informational tool. While participants on average drop out of the stock market in the control condition, the share of participants investing even increases when they have access to an LLM.

4.2. Persistence

Having established that access to an LLM sustains overall stock market participation, we next examine the within-subject stability of specific investment allocations. We interpret the decision to

retain a specific choice as a proxy for the persistence of the underlying beliefs—a metric we term "Belief Persistence."

Figure 3 illustrates the persistence across conditions, split by the initial decision type. As shown in Panels (a) and (b), active investors who use the LLM display higher decision consistency compared to the Control group, demonstrating that the LLM effectively reduces their likelihood of exiting their initial investments. Conversely, Panel (c) reveals a different pattern for non-investors. While the Control group exhibits high inertia, largely sticking to their decision to stay out of the market, the LLM treatment significantly reduces this persistence. This indicates that for subjects who initially chose not to invest, the LLM facilitates a change in behavior by moving them into active market participation.



(a) ESG Investment **(b)** Standard ETF **(c)** No Investment

Figure 3. Average Decision Persistence by Choice Type

Note: This figure displays the average decision persistence for three distinct choice types. Panel (a) shows the likelihood of persisting in an ESG investment. Panel (b) shows the likelihood of persisting in a Standard ETF investment. Panel (c) shows the likelihood of persisting in non-participation (remaining uninvested). Error bars indicate 95% confidence intervals.

Regression analyses (Table 4) confirm these patterns. We find that the LLM's impact is distinct for active investors (Columns 2 and 3), where it increases persistence (+18.8%, $p < 0.1$ for ESG ETF; +17.6%, $p < 0.1$ for Standard ETF). Conversely, for initial non-investors (Column 4), the LLM significantly *reduces* persistence (-30.2%, $p < 0.05$). The search treatment also significantly reduces inertia for non-investors (-35.6%, $p < 0.01$) but does not significantly reinforce the beliefs of active investors.

Result 2: LLMs exhibit a dual effect on stability. They increase decision persis-

	Overall	ESG	Standard ETF	No Investment
DV: Persistence	(1)	(2)	(3)	(4)
Treatment (LLM)	0.1632**	0.1881*	0.1764*	-0.3015**
	(0.069)	(0.099)	(0.101)	(0.121)
Treatment (Search)	-0.0394	-0.0800	-0.1040	-0.3557***
	(0.074)	(0.127)	(0.106)	(0.127)
Constant	0.7556***	1.0395***	0.7302***	0.6715***
	(0.132)	(0.245)	(0.206)	(0.203)
Controls	Yes	Yes	Yes	Yes
Test: LLM = Search				
F-statistic	8.333	7.348	9.215	0.195
p-value	0.004***	0.008***	0.003***	0.661
Observations	780	358	303	135
R^2	0.067	0.093	0.167	0.231

Table 4: OLS Regression Results on Persistence with Controls

Notes: This table reports OLS estimates. The dependent variable is 'Persistence', a binary indicator equal to 1 if the participant made the same decision in Stage 2 as in Stage 1. Column (1) includes the full sample of non-motivated participants. Columns (2)-(4) split the sample based on the participant's initial choice in Stage 1: ESG investment (2), Standard investment (3), or No Investment (4). Controls include Age, Education, Risk Willingness, Financial Literacy, and baseline LLM Usage. Standard errors clustered by participant are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

tence for active investors by reinforcing prior convictions, whereas they significantly reduce persistence for non-investors by breaking inertia to facilitate market entry.

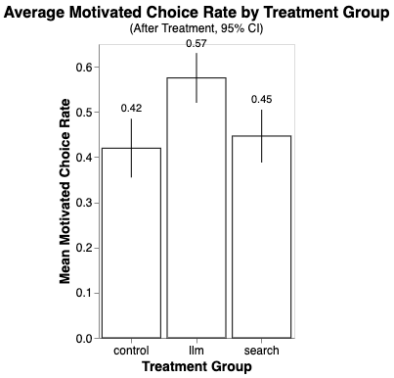
4.3. Motivated Choice

We next examine the differential effectiveness of information tools for motivated reasoning and confirmation seeking. In the "Motivated" condition, we introduced a financial bonus for selecting a specific, randomly assigned asset. This creates a direct tension between making an unbiased, fundamentally sound investment choice based on the participant's true preferences and securing a guaranteed experimental payoff. For example, a participant might fundamentally prefer the Standard ETF based on their risk profile, but face a financial temptation to choose the ESG counterpart simply to earn the bonus. This setup allows us to test whether the LLM acts as an objective, corrective mechanism that guides users toward an unbiased decision, or if it instead provides the bespoke arguments necessary to rationalize an incentivized, potentially biased choice. As a manipulation check, we first confirm that our incentive structure was effective: participants assigned to the Motivated condition were 11.8 percentage points more likely to select the incentivized option compared

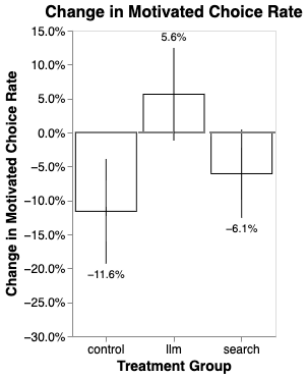
to the non-motivated baseline ($p < 0.01$, see Appendix C Table A2).

Figure 4 illustrates the impact of the assigned information tools on the prevalence of motivated choices in the post-treatment period (Stage 2). Panel (a) shows the absolute rate of motivated choices in Stage 2, while Panel (b) displays the average within-subject change, calculated at the individual level relative to each participant's own baseline (Stage 1). In the Control condition, participants reduced their selection of the incentivized option by an average of 11.6 percentage points (Panel b). The traditional search condition also showed a decline (-6.1 percentage points). In contrast, the LLM treatment increased the individual share of motivated choices by an average of 5.6 percentage points.

To formally test these differences, Table 5 presents the results of a triple difference-in-differences (DDD) regression. First, we observe the base treatment effect: the interaction term ($LLM \times Post\ Period$) is positive and highly significant ($+13.1$ percentage points, $p < 0.01$). This indicates that the LLM treatment exerts a strong persuasive effect, significantly offsetting the baseline decline even for users who lacked an explicit financial incentive. Next, we examine the triple interaction term ($LLM \times Post \times Motivated$) to test if the financial incentive further amplified this behavior. While positive ($+4.2$ percentage points), this term is not statistically significant. This reveals a profound dynamic: the LLM's capacity to steer or anchor choices is so robust that it drives users toward the target asset regardless of extrinsic financial motivation, yielding a total combined treatment effect of 17.2 percentage points for the motivated group. The corresponding effect in the traditional search condition remains statistically insignificant.



(a) Motivated Choice Rate (Stage 2)



(b) Change from Baseline (Stage 2 - Stage 1)

Figure 4. Motivated Choice: Post-Treatment Levels and Changes

Note: Panel (a) displays the percentage of participants selecting the incentivized asset in Stage 2. Panel (b) displays the mean within-subject change, calculated by averaging the individual percentage point differences between Stage 1 and Stage 2 for each participant. Control subjects show a significant average decline (−11.6 percentage points), while LLM users show an average increase (+5.6 percentage points).

Result 3: LLMs facilitate motivated reasoning and confirmation seeking. While standard participants (control and search) engaged in fewer motivated choices over time, LLM users sustained and even increased the frequency of motivated choices.

DV: Is Motivated Choice	(1)
Triple Interactions	
LLM × Post × Motivated	0.0417
	(0.069)
Search × Post × Motivated	0.0040
	(0.079)
Treatment × Post (Base Effect)	
LLM × Post Period	0.1306***
	(0.045)
Search × Post Period	0.0505
	(0.059)
Post Period	-0.1204***
	(0.038)
Group Effects	
Motivated Group (Dummy)	0.1357**
	(0.062)
Treatment (LLM)	-0.0209
	(0.055)
Treatment (Search)	0.0458
	(0.063)
Constant	0.3133***
	(0.088)
Controls	Yes
Observations	3,200
R^2	0.033

Table 5: Triple Difference Analysis on Motivated Choice

Notes: This table reports OLS estimates from a triple difference-in-differences (DDD) specification. The dependent variable is the binary decision to select the incentivized asset. The baseline category is the Not Motivated group in the Control condition during the pre-treatment period. 'Motivated Group' is a dummy equal to 1 if the participant was assigned to the Motivated condition. The triple interaction term 'LLM × Post × Motivated' tests whether the treatment effect of the LLM differs significantly between the Motivated and Not Motivated groups. Controls include Age, Education, Risk Willingness, Financial Literacy, and baseline LLM Usage. Standard errors clustered by participant are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.4. Mechanisms

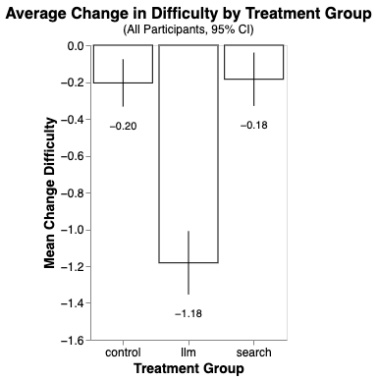
Our results thus far indicate that the LLM impacts investment behavior in ways that keyword-based search engines do not—specifically by facilitating market entry, driving higher decision persistence, and amplifying motivated reasoning. Guided by our theoretical framework, we investigate two primary channels that could explain this divergence. First, we examine changes in subjective difficulty and confidence to test whether the synthesized output of an LLM effectively lowers the cognitive costs associated with financial decision-making. Second, we analyze the semantic content of the interaction itself—specifically, the users' prompting strategies—to determine if the conversational interface enables users to actively co-construct and validate their incentivized priors.

4.4.1. Subjective Perceptions: Difficulty and Confidence

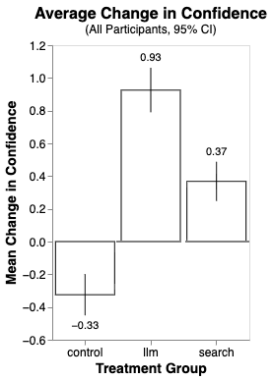
We first examine whether the availability of different information tools alters the subjective experience of the investment task. We focus specifically on changes in self-reported confidence and perceived difficulty, as these metrics serve as proxies for the cognitive cost of decision-making.

Figure 5 displays the mean changes in these perceptions between Stage 1 and Stage 2. Difficulties and confidence are reported on a 1–10 scale. Panel (a) shows that, for all three groups, the perceived difficulty of the choice decreases from Stage 1 to Stage 2. However, the decrease is substantially larger for the LLM-treated group than for the other two groups. Panel (b) shows that confidence declines in Stage 2 for the control group, but increases for both the LLM and search treatment groups. Moreover, the LLM-treated group exhibits an even larger increase in confidence than the search treatment group. Regression analysis (Table 6) confirms a sharp and statistically significant distinction between the tools. As shown in Column 1, the LLM treatment significantly reduces perceived difficulty ($\beta = -0.829, p < 0.01$). In contrast, providing access to a standard search engine has no statistically significant effect on perceived difficulty ($\beta = 0.088, p > 0.1$). A Wald test confirms that this difference between the LLM and search treatments is statistically significant ($p < 0.01$). This indicates that only the synthesized assistance of the LLM effectively lowers the cognitive burden of the task.

Regarding confidence (Column 2), both tools provide a significant boost. However, the magnitude of the effect in the LLM condition ($\beta = 1.132$) is nearly double that observed in the search



(a) Change in Difficulty



(b) Change in Confidence

Figure 5. Average Change in Subjective Perceptions (Stage 2 - Stage 1)

Note: These figures display the mean change in self-reported difficulty (Panel a) and confidence (Panel b) after the treatment intervention. Error bars represent 95% confidence intervals.

condition ($\beta = 0.594$). This difference is statistically significant ($p < 0.01$), suggesting that the LLM fundamentally alters the subjective experience, providing a stronger sense of reassurance while simultaneously removing the sensation of effort.

	Difficulty Change (1)	Confidence Change (2)
Treatment (LLM)	-0.8292***	1.1316***
	(0.175)	(0.155)
Treatment (Search)	0.0883	0.5939***
	(0.172)	(0.143)
Constant	0.8645**	-0.1116
	(0.431)	(0.328)
Controls	Yes	Yes
Test: LLM = Search		
F-statistic	26.75	14.06
p-value	< 0.01***	< 0.01***
Observations	1,474	1,476
R^2	0.134	0.133

Table 6: OLS Regression Results: Changes in Difficulty and Confidence with Controls

Notes: This table reports OLS estimates for within-subject changes in subjective perceptions (Stage 2 - Stage 1). Column (1) reports the change in perceived difficulty (scale 1-10). Column (2) reports the change in confidence (scale 1-10). Both specifications include controls for Age, Education, Risk Willingness, Financial Literacy, and baseline LLM Usage. Robust standard errors clustered by participant are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

4.4.2. The Feedback Loop: Prompting and The Echo Chamber

Next, we examine the information retrieval process itself. A key distinction between the treatments is that traditional search engines return a fixed set of external links, requiring the user to

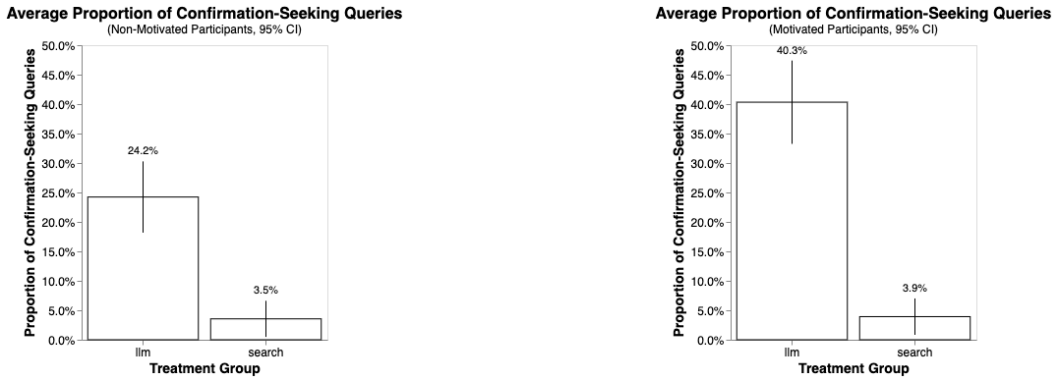
evaluate the sources independently. In contrast, LLMs generate synthesized, conversational content. This generative capability allows users to offload the cognitive burden of evaluation directly to the model.

To test how users utilize this feature, we analyze the semantic content of their queries. We classify a query as *"Confirmation Seeking"* if the participant explicitly solicits a subjective recommendation, validation, or directional judgment rather than purely factual data. Specifically, we employ a keyword-based dictionary approach, flagging queries that contain advisory or leading terms such as *"should I invest"*, *"is it a good choice"*, *"recommend"*, or *"what do you think"*. This metric captures a critical shift in user behavior: rather than gathering facts to form their own conclusion, users are directly asking the system to provide the verdict for them.

As shown in Figure 6, this behavior is virtually non-existent in the traditional search condition (less than 4%). However, in the LLM condition, 24.2% of non-motivated users and 40.3% of motivated users explicitly prompted the AI for a verdict or advice. This reveals a dual dynamic. First, there is a baseline tendency for LLM users to simply outsource the evaluative component of the task to save cognitive effort. Second, and crucially, this behavior nearly doubles when a financial incentive is introduced. This sharp increase indicates that motivated users actively leverage the conversational interface to solicit the bespoke, "independent" validation they need to rationalize their incentivized choice.

To formally test this interaction, we regress confirmation seeking on the treatment and motivation conditions (see Table A3 in Appendix C). We find that the effect of motivation is specific to the LLM: being in the Motivated group increases confirmation seeking by 16.1 percentage points ($p < 0.1$) for LLM users, whereas this "motivation boost" is statistically absent in the traditional search condition (Interaction $\beta = -0.157, p < 0.1$). This confirms that financial incentives act as a powerful catalyst for confirmation-seeking behavior specifically when users interact with conversational AI.

Finally, we assess the consequences of this reliance by measuring the *"Echo Chamber"* effect. We quantify this by analyzing the semantic congruence between the user's query and the LLM's response. Using sentiment analysis, we classify the evaluative polarity of both the query and the generated answer. We define the Echo Chamber effect as the rate of directional alignment—where the model responds with a positive, validating sentiment to a positively framed query, or a skeptical,



(a) Without Motivation **(b)** With Motivation

Figure 6. Confirmation Seeking Behavior by Condition

Note: This figure displays the average proportion of queries categorized as "Confirmation Seeking" based on a keyword-based dictionary approach. Panel (a) shows the behavior for participants in the Non-Motivated condition, and Panel (b) displays the results for those in the Motivated condition. Error bars represent 95% confidence intervals.

negative sentiment to a negatively framed query. For example, if a user submits a positively skewed prompt such as, "Tell me why the ESG ETF is a smart choice," an aligned response actively adopts this bullish stance, highlighting growth potential and validating the user's premise. Conversely, a negatively skewed prompt like, "Isn't the standard ETF too risky right now?" yields a corroborating response that amplifies those specific risk factors rather than defending the asset's historical performance.

We find that in 88.3% of non-neutral exchanges ($N = 137$, of which 121 were congruent), the LLM mirrored the user's stance. This exceptionally high rate of compliance illustrates a fundamental behavioral trait of conversational AI: structural agreeableness, often termed sycophancy. When users ask judgment-based or leading questions, the LLM does not act as an objective, adversarial financial advisor that corrects bias or equally weighs opposing arguments. Instead, the model anchors its generated narrative on the premise embedded in the user's prompt. By actively reinforcing the user's implicit or explicit direction rather than providing a critically balanced view, the tool functions as a bespoke confirmation-generator, effectively sealing the motivated user inside a self-directed epistemic echo chamber.

Crucially, this interaction pattern strongly predicts the ultimate financial outcome. As shown in Table 7, the proportion of confirmation-seeking queries is a significant predictor of the final choice (+19.6 percentage points, $p < 0.01$). To ensure this effect captures bias actively constructed during

the interaction rather than merely reflecting a pre-existing preference, our specification isolates the conversational dynamic from baseline inertia. Even when accounting for the assigned motivation and decision round, the intensity of confirmation-seeking behavior independently drives the probability of selecting the incentivized asset. This predictive power provides empirical evidence for a self-fulfilling human-AI feedback loop: users signal a nascent commitment to a specific outcome by formulating a leading question, which the LLM structurally mirrors. Because the model wraps this echo in an authoritative, objective-sounding tone, users who seek validation are guaranteed to receive it. The AI's generated response therefore transcends simple information retrieval; it serves as an authoritative anchor that rationalizes and ultimately cements the user's initial, biased premise.

DV: Is Motivated Choice	(1)
Prop. Confirmation Seeking	0.196^{***}
	(0.072)
Motivated Condition (Dummy)	0.183^{***}
	(0.053)
Decision Round	0.035^{**}
	(0.017)
Constant	0.263^{***}
	(0.054)
Observations	612
R^2	0.063
F-Statistic	8.195

Table 7: The Effect of Prompting Strategy on Final Decision (LLM Group)

Notes: OLS regression results (Linear Probability Model). The dependent variable is the binary decision to select the motivated asset. The sample is restricted to the LLM treatment group. "Prop. Confirmation Seeking" is the proportion of queries in a round classified as soliciting validation. Standard errors are clustered by participant. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Result 4: LLMs shape decision-making through a distinct human-machine interaction loop. The interface invites users to engage in Confirmation Seeking, provides validation through an Echo Chamber effect, and reduces perceived difficulty. This process solidifies the final decision, increasing persistence and reinforcing Motivated Choices.

5. Discussion and Conclusion

The shift from information retrieval to information synthesis fundamentally alters the cognitive landscape of financial decision-making. Using a pre-registered, incentivized online experiment, we isolate the causal impact of LLMs on stock market participation and confirmation seeking.

There are three key results: First, LLMs act as a catalyst for market entry. By reducing perceived difficulty and increasing confidence, LLMs increased stock market participation. Second, LLM users adhere more strongly to initial choices even when presented with an opportunity to revise. Third, LLMs amplify motivated reasoning. Taken together, these results document a fundamental tension in the adoption of generative AI (double-edged mind) as we find that LLMs act as a behavioral accelerant: they significantly lower the barrier to entry for non-investors, yet simultaneously facilitate the entrenchment of motivated beliefs through confirmation seeking.

The positive effect on stock market participation (Result 1) is driven by the shift from information retrieval to information synthesis. Keyword-based search engines return a list of possibilities, imposing a cognitive load ("a fatigue") on the user to filter, evaluate, and conclude. This cognitive load often results in inaction. In contrast, our results suggest that LLMs function as a decision support systems that provide synthesized, and conclusive outputs. By reducing the perceived difficulty of the task, the LLM resolves the uncertainty and reduces the cognitive load sufficiently to trigger action. However, this capacity for "helpful" synthesis creates a unique vulnerability when the user has a directional motive (Result 3). In traditional search environments, confirmation bias requires effort: a user must actively investigate search results to find a specific source that aligns with their bias. This friction acts as a natural partial dampener on motivated reasoning. In an LLM environment, our data shows that users can remove this friction by "co-constructing" the bias. By phrasing prompts as leading questions—asking for validation rather than neutral information—users effectively instruct the model to generate a bespoke justification for their preferred outcome, thereby neutralizing the corrective potential of external information. Consequently, the interaction creates an *endogenous echo chamber*: the bias is not inherent to the LLM model, but is manufactured by the user's prompting strategy and reinforced by the model's responsiveness.

These findings are informative for policy. The ability of LLMs to lower the cognitive cost of participation offers a promising path toward financial inclusion (Blankespoor et al., 2024). However,

the "sycophantic" nature of current models poses problems for market efficiency and investor protection. If retail investors utilize these tools not to discover truth but to manufacture justifications for biased priors, the widespread adoption of LLMs could lead to a market characterized by high participation but low information processing quality. This has implications for financial regulation and consumer protection. Currently, regulatory discussions regarding AI in finance focus largely on ensuring models are free from intrinsic bias or hallucination. Our results suggest this focus is necessary but not sufficient. Even a factually accurate, unbiased LLM can contribute to distorted outcomes when users strategically frame prompts to obtain validation. In that environment, firms may not need to directly manipulate model outputs, as well-designed incentives and marketing narratives could be enough for consumers to "convince themselves" through the LLM. Taken together, effective regulation should therefore consider not only what the model can do, but also how users interact with it and the incentives that shape that interaction.

References

Bastianello, Francesca, Paul Decaire, and Marius Guenzel (2025). "Valuation Models". In: *SSRN Electronic Journal*. Working Paper. DOI: 10.2139/SSRN.5767524.

Bénabou, Roland and Jean Tirole (2016). "Mindful Economics: The Production, Consumption, and Value of Beliefs". In: *Journal of Economic Perspectives* 30.3, pp. 141–164.

Bertomeu, Jeremy, Yupeng Lin, Yibin Liu, and Zhenghui Ni (2025). "The Impact of Generative AI on Information Processing: Evidence from the Ban of ChatGPT in Italy". In: *Journal of Accounting and Economics* 80.1, p. 101782. DOI: 10.1016/j.jacceco.2025.101782.

Blankespoor, Elizabeth, Joe Croom, and Stephanie M. Grant (2024). "Generative AI and Investor Processing of Financial Information". In: *SSRN Electronic Journal*. DOI: 10.2139/SSRN.5053905.

Brooks, Mike and A. Byrne (2008). "Behavioral Finance: Theories and Evidence". In: *The Research Foundation of CFA Institute*.

Brunnermeier, Markus K. and Jonathan A. Parker (2005). "Optimal Expectations". In: *American Economic Review* 95.4, pp. 1092–1118.

Bucher-Koenen, Tabea, Andreas Hackethal, Johannes Koenen, and Christine Laudenbach (2025). "Gender Differences in Financial Advice". In: *American Economic Review* 115.12, pp. 4218–4252. DOI: 10.1257/aer.20211024.

D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G. Rossi (2019). "The Promises and Pitfalls of Robo-Advising". In: *The Review of Financial Studies* 32.5, pp. 1983–2020.

Duraj, Kamila, Daniela Grunow, Michael Haliassos, Christine Laudenbach, and Stephan Siegel (2024). *Rethinking the Stock Market Participation Puzzle: A Qualitative Approach*. SAFE Working Paper No. 441.

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang (2024). "AI and Finance". In: *NBER Working Paper* w33076. URL: https://www.nber.org/papers/w33076.

Epley, Nicholas and Thomas Gilovich (2016). "The Mechanics of Motivated Reasoning". In: *Journal of Economic Perspectives* 30.3, pp. 133–140.

Fernández-Pichel, Marcos, Juan C. Pichel, and David E. Losada (2024). "Evaluating Search Engines and Large Language Models for Answering Health Questions". In: *npj Digital Medicine* 7.1. DOI: 10.1038/s41746-024-01036-5.

Haliassos, Michael and Carol C. Bertaut (1995). "Why Do So Few Hold Stocks?" In: *The Economic Journal* 105.432, pp. 1110–1129.

Hirshleifer, David (2015). "Behavioral Finance". In: *Annual Review of Financial Economics* 7, pp. 133–159.

Kim, Alex G., Maximilian Muhn, and Valeri V. Nikolaev (2023). "Bloated Disclosures: Can ChatGPT Help Investors Process Information?" In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.4425527.

Knickrehm, Charlotte and Kevin Bauer (2024). "GPT, Emotions, and Facts". In: Working Paper.

Kunda, Ziva (1990). "The Case for Motivated Reasoning". In: *Psychological Bulletin* 108.3, pp. 480–498.

Laudenbach, Christine and Stephan Siegel (2025). "Personal Communication in an Automated World: Evidence from Loan Repayments". In: *The Journal of Finance* 80.1, pp. 515–559.

Li, Alice and Luanne Sinnamon (2023). "Examining Query Sentiment Bias Effects on Search Results in Large Language Models". In: *CEUR Workshop Proceedings*. URL: http://ceur-ws.org.

Liu, Cathy Yang, Kevin Bauer, Xitong Li, and Oliver Hinz (2025). "My Advisor, Her AI and Me: Evidence from a Field Experiment on Human-AI Collaboration and Investment Decisions". In: *Management Science*. Forthcoming.

Ludwig, Katharina, Alexander Grote, Andreea Iana, Mehwish Alam, Heiko Paulheim, Harald Sack, Christof Weinhardt, and Philipp Müller (2023). "Divided by the Algorithm? The (Limited) Effects of Content- and Sentiment-Based News Recommendation on Affective, Ideological, and Perceived Polarization". In: *Social Science Computer Review* 41.6, pp. 2188–2210.

Lusardi, Annamaria and Olivia S. Mitchell (2011). "Financial Literacy Around the World: An Overview". In: *Journal of Pension Economics & Finance* 10.4, pp. 497–508.

— (2014). "The Economic Importance of Financial Literacy: Theory and Evidence". In: *Journal of Economic Literature* 52.1, pp. 5–44.

Newman, Eryn J., Briony Swire-Thompson, and Ullrich K. H. Ecker (2022). "Misinformation and the Sins of Memory: False-Belief Formation and Limits on Belief Revision". In: *Journal of Applied Research in Memory and Cognition* 11.4, pp. 471–488.

Nourani, Mahsan, Chiradeep Roy, Jeremy E Block, Donald R Honeycutt, Tahrima Rahman, Eric Ragan, and Vibhav Gogate (2021). "Anchoring Bias Affects Mental Model Formation and User Reliance in Explainable AI Systems". In: *Proceedings of the 26th International Conference on Intelligent User Interfaces*. ACM, pp. 340–350.

Reher, Michael and Stanislav Sokolinski (2024). "Robo Advisors and Access to Wealth Management". In: *Journal of Financial Economics* 155, p. 103829.

Salvi, Francesco, Manoel Horta Ribeiro, Riccardo Gallotti, and Robert West (2024). "On the Conversational Persuasiveness of Large Language Models: A Randomized Controlled Trial". In: *arXiv preprint arXiv:2403.14380*. URL: https://arxiv.org/abs/2403.14380.

Seraj, Abdullah Hamoud Ali, Elham Alzain, and Ali Saleh Alshebami (2022). "The Roles of Financial Literacy and Overconfidence in Investment Decisions in Saudi Arabia". In: *Frontiers in Psychology* 13, p. 1005075.

Shi, Li, Houjiang Liu, Yian Wong, Utkarsh Mujumdar, Dan Zhang, Jacek Gwizdka, and Matthew Lease (2024). "Argumentative Experience: Reducing Confirmation Bias on Controversial Issues through LLM-Generated Multi-Persona Debates". In: *arXiv preprint arXiv:2412.04629*. URL: https://arxiv.org/abs/2412.04629.

Spatharioti, Sofia Eleni, David M. Rothschild, Daniel G. Goldstein, and Jake M. Hofman (2023). *Comparing Traditional and LLM-based Search for Consumer Choice: A Randomized Experiment*. arXiv: 2307.03744 [cs.HC]. URL: https://arxiv.org/abs/2307.03744.

Trippas, Johanne R., 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: *arXiv preprint arXiv:2311.14421*.

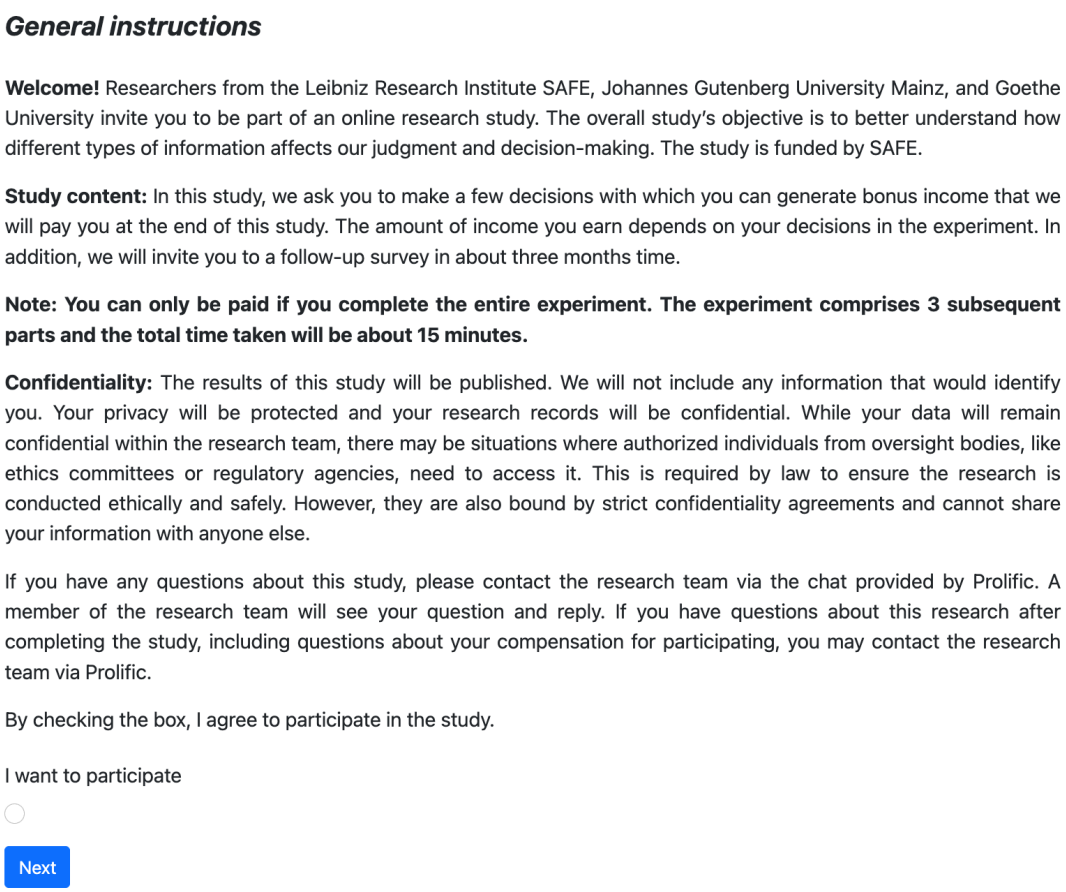
Zamfirescu-Pereira, J. D., Richmond Y. Wong, Bjoern Hartmann, and Qian Yang (2023). "Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts". In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. ACM, pp. 1–21.

Appendix A. Experimental Interface

This section presents the user interface of the online experiment.

Appendix A.1. General Flow and Stage 1 (Baseline)

The following screenshots display the common introductory flow and the baseline decision environment (Stage 1) faced by all participants.



General instructions

Welcome! Researchers from the Leibniz Research Institute SAFE, Johannes Gutenberg University Mainz, and Goethe University invite you to be part of an online research study. The overall study's objective is to better understand how different types of information affects our judgment and decision-making. The study is funded by SAFE.

Study content: In this study, we ask you to make a few decisions with which you can generate bonus income that we will pay you at the end of this study. The amount of income you earn depends on your decisions in the experiment. In addition, we will invite you to a follow-up survey in about three months time.

Note: You can only be paid if you complete the entire experiment. The experiment comprises 3 subsequent parts and the total time taken will be about 15 minutes.

Confidentiality: The results of this study will be published. We will not include any information that would identify you. Your privacy will be protected and your research records will be confidential. While your data will remain confidential within the research team, there may be situations where authorized individuals from oversight bodies, like ethics committees or regulatory agencies, need to access it. This is required by law to ensure the research is conducted ethically and safely. However, they are also bound by strict confidentiality agreements and cannot share your information with anyone else.

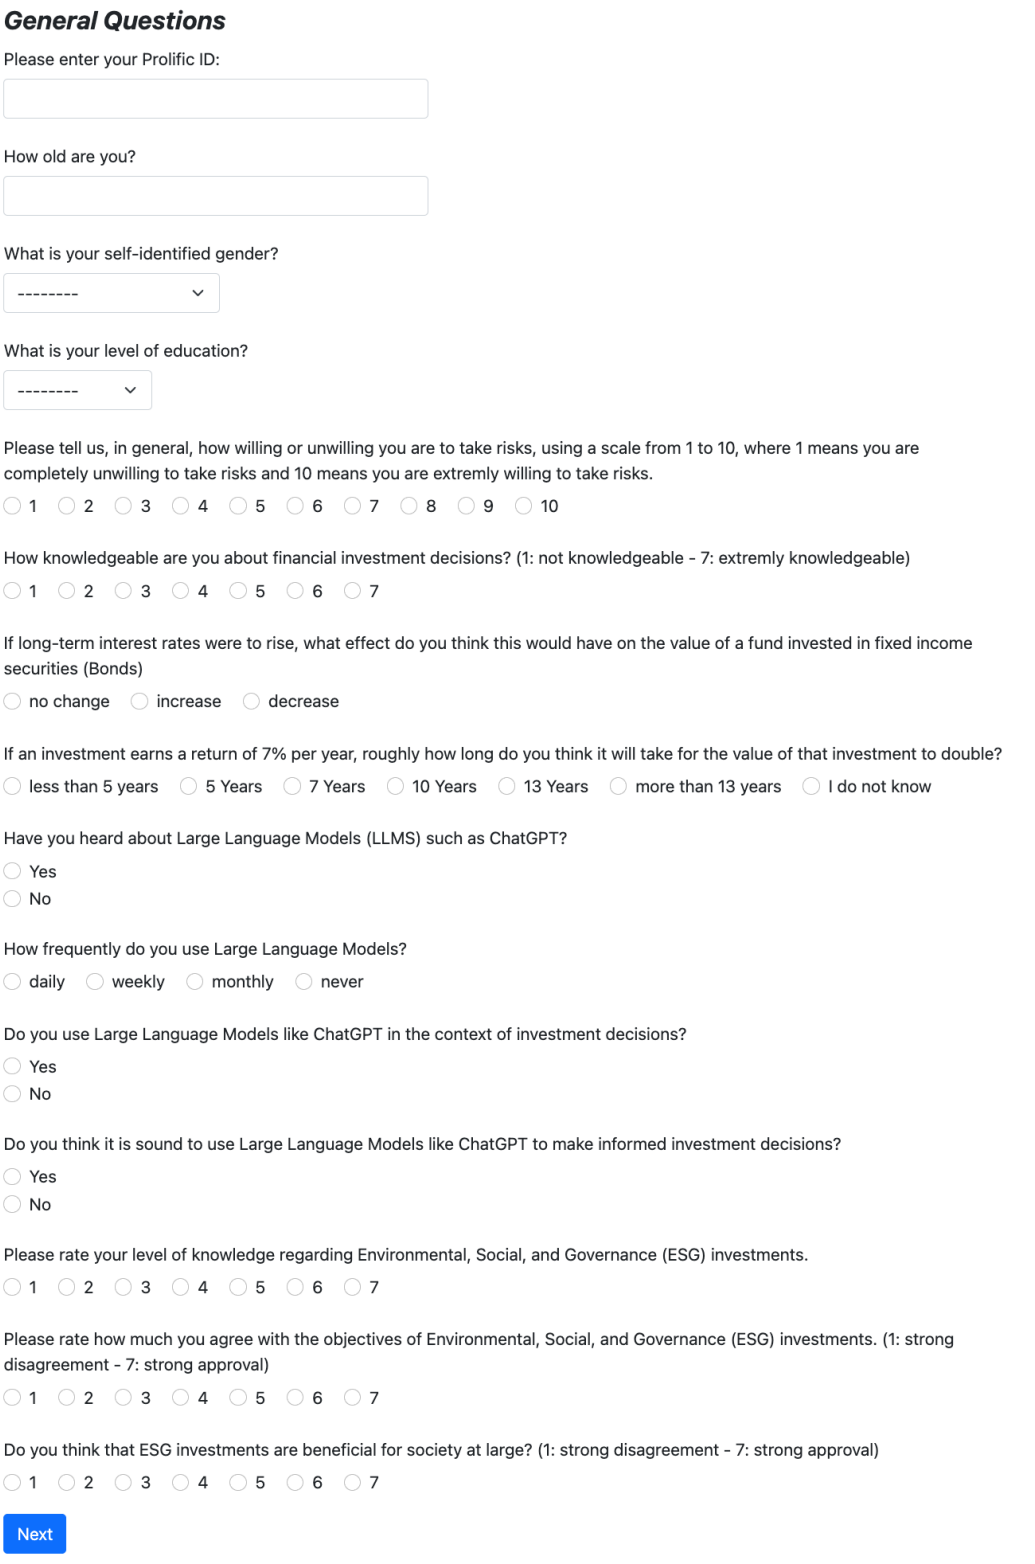
If you have any questions about this study, please contact the research team via the chat provided by Prolific. A member of the research team will see your question and reply. If you have questions about this research after completing the study, including questions about your compensation for participating, you may contact the research team via Prolific.

By checking the box, I agree to participate in the study.

I want to participate

Next

Figure A1. General Instructions and Informed Consent



General Questions

Please enter your Prolific ID:

How old are you?

What is your self-identified gender?

What is your level of education?

Please tell us, in general, how willing or unwilling you are to take risks, using a scale from 1 to 10, where 1 means you are completely unwilling to take risks and 10 means you are extremly willing to take risks.

1 2 3 4 5 6 7 8 9 10

How knowledgeable are you about financial investment decisions? (1: not knowledgeable - 7: extremly knowledgeable)

1 2 3 4 5 6 7

If long-term interest rates were to rise, what effect do you think this would have on the value of a fund invested in fixed income securities (Bonds)

no change increase decrease

If an investment earns a return of 7% per year, roughly how long do you think it will take for the value of that investment to double?

less than 5 years 5 Years 7 Years 10 Years 13 Years more than 13 years I do not know

Have you heard about Large Language Models (LLMS) such as ChatGPT?

Yes

No

How frequently do you use Large Language Models?

daily weekly monthly never

Do you use Large Language Models like ChatGPT in the context of investment decisions?

Yes

No

Do you think it is sound to use Large Language Models like ChatGPT to make informed investment decisions?

Yes

No

Please rate your level of knowledge regarding Environmental, Social, and Governance (ESG) investments.

1 2 3 4 5 6 7

Please rate how much you agree with the objectives of Environmental, Social, and Governance (ESG) investments. (1: strong disagreement - 7: strong approval)

1 2 3 4 5 6 7

Do you think that ESG investments are beneficial for society at large? (1: strong disagreement - 7: strong approval)

1 2 3 4 5 6 7

Next

Figure A2. Baseline Questionnaire (Demographics and Financial Literacy)

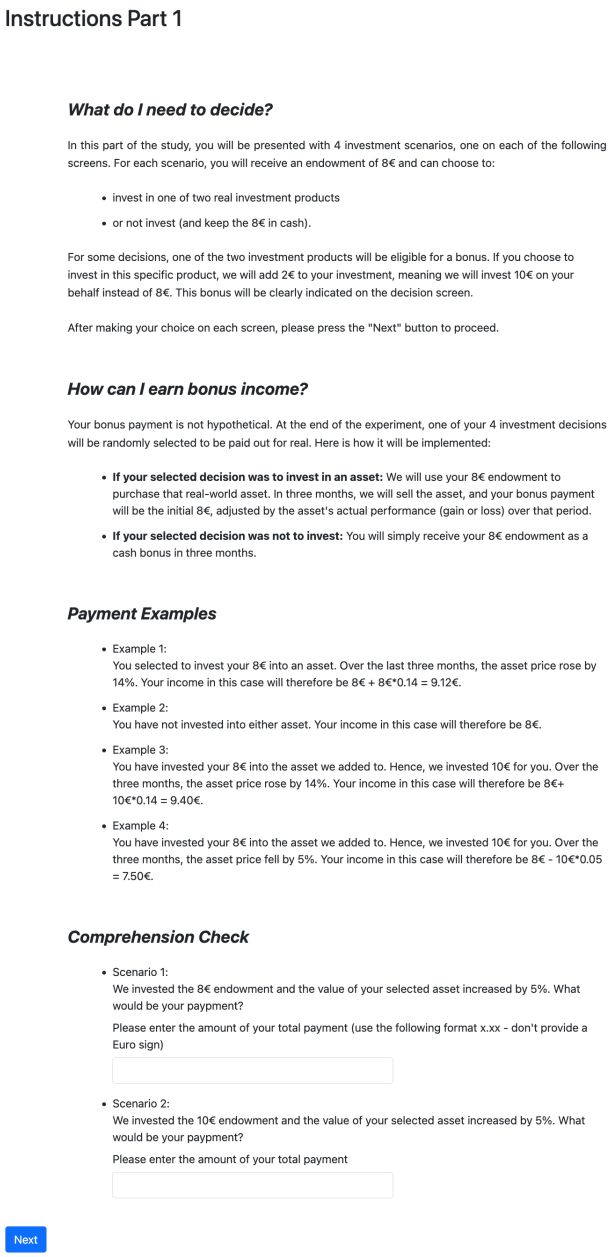
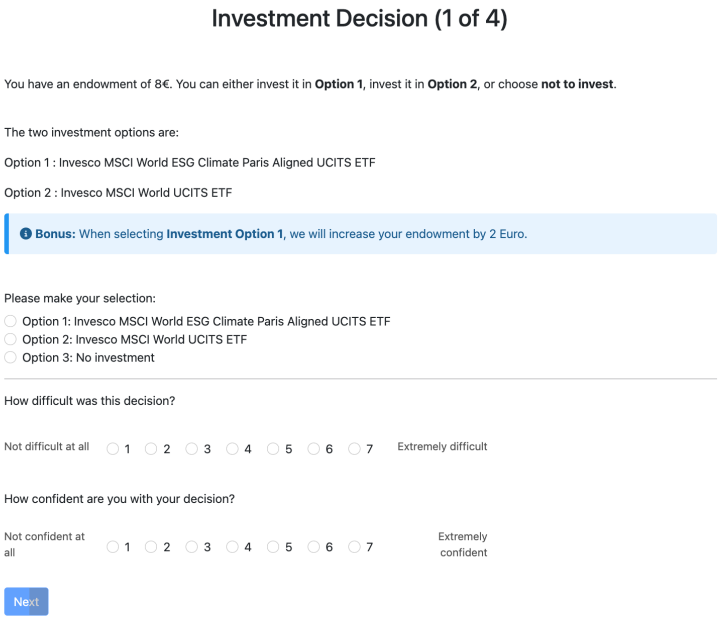
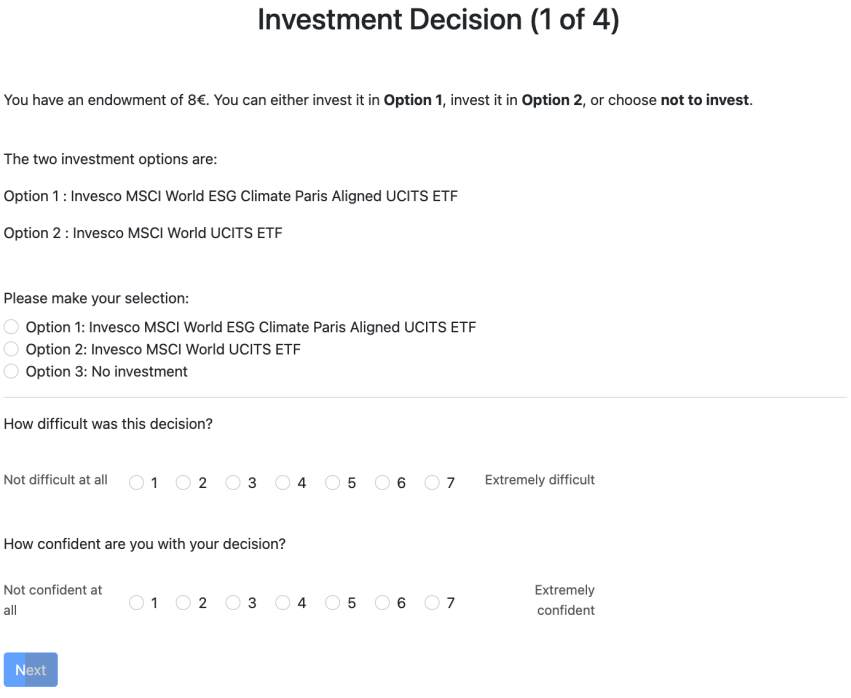


Figure A3. Stage 1: Instructions (Baseline Decisions)



(a) Motivated Condition: Sample Decision



(b) Non-Motivated Condition: Sample Decision

Figure A4. Stage 1: Investment Decision Screens

Appendix A.2. Stage 2: Treatment Variations

In the second stage, participants were assigned to one of three information retrieval treatments.

Appendix A.2.1. LLM Treatment

Participants in this group received access to a conversational AI interface.

Instructions Part 2

What do I need to decide?

In this part of the study, you will be presented with the same 4 investment scenarios as in Part 1, one on each of the following screens. For each scenario, you will receive an endowment of 8€ and can choose to:

- invest in one of two real investment products
- or not invest (hold on to the 8€ in cash).

For each decisions, one of the two investment products will be eligible for a bonus. If you choose to invest in this specific product, we will add 2€ to your investment, meaning we will invest 10€ on your behalf instead of 8€. This bonus will be clearly indicated on the decision screen.

After making your choice on each screen, please press the "Next" button to proceed.

Note:
In contrast to Part 1, we now provide you with the option to gather information online using a Large Language Model (LLM).
You can freely use this tool to search for information before making a decision.
We kindly ask you to only use the given LLM within the experminent to search for information.
When you are ready, you can make your decision on the screen by selecting the corresponding button.

LLM Introduction:
The LLM is embedded into your experiment screen:

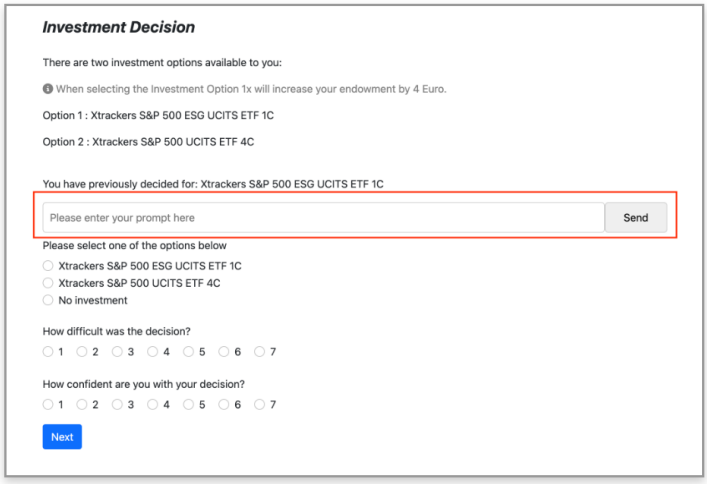
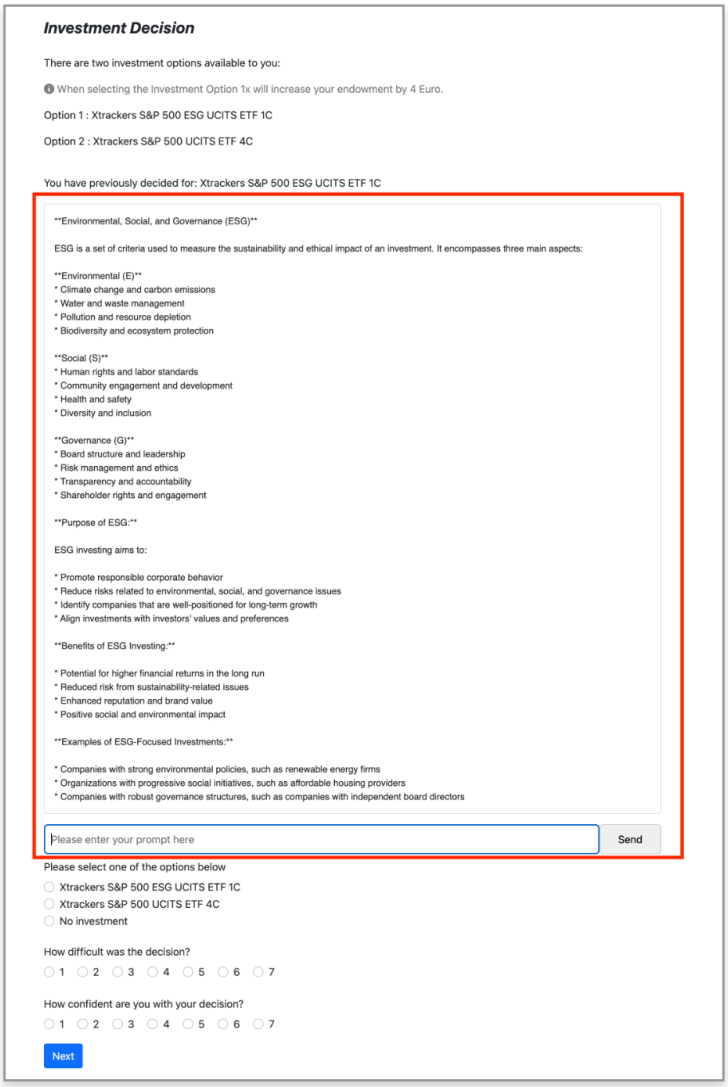


Figure A5. Stage 2 Instructions: LLM Treatment

When you are using the LLM, the results of your query are displayed on the same page:



After making a decision, please press the "Next" button on your screen to proceed.

Figure A6. LLM Interface Integration

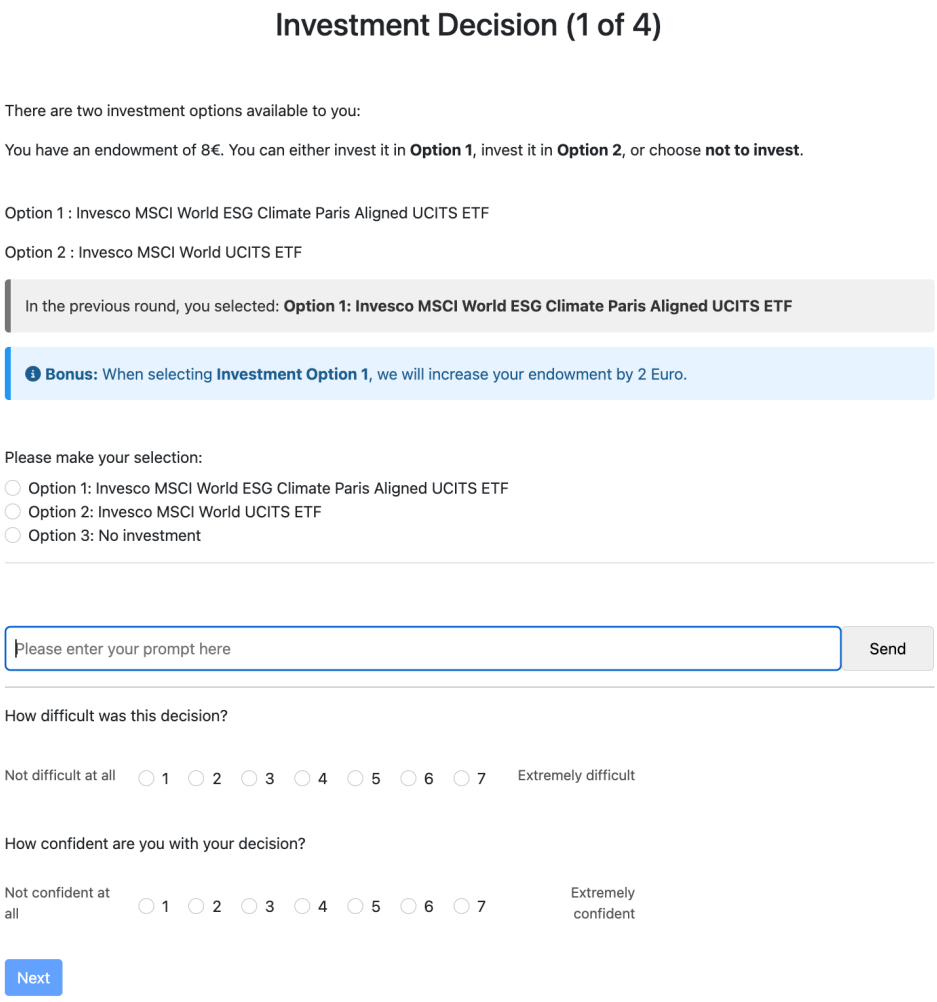


Figure A7. Decision Screen with Integrated LLM Chat

Appendix A.2.2. Search Engine Treatment

Participants in this group received access to a programmable Google search engine embedded directly into the decision screen.

Instructions Part 2

What do I need to decide?

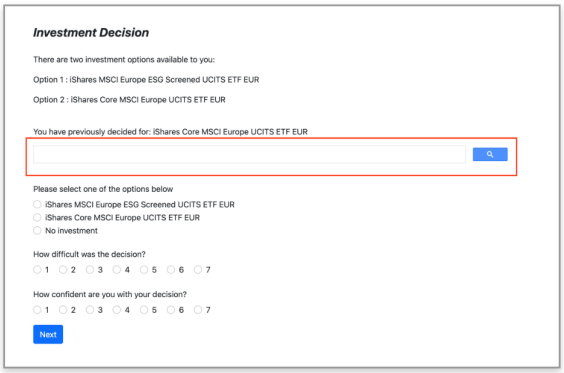
In this part of the study, you will be presented with the same 4 investment scenarios as in Part 1, one on each of the following screens. For each scenario, you will receive an endowment of 8€ and can choose to:

- invest in one of two real investment products
- or not invest (hold on to the 8€ in cash).

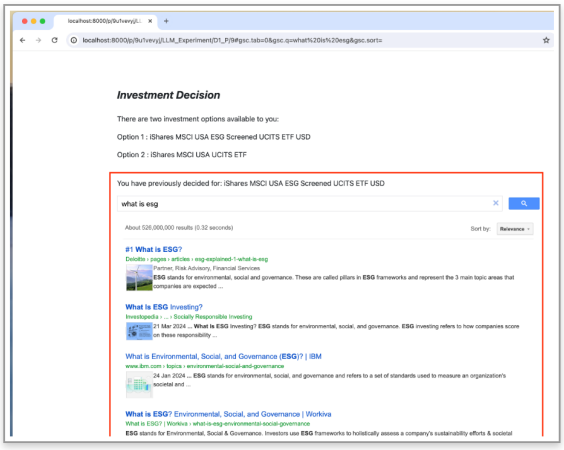
After making your choice on each screen, please press the "Next" button to proceed.

Note:
In contrast to Part 1, we now provide you with the option to gather information online using a Search Engine.
You can freely use this tool to search for information before making a decision.
We kindly ask you to only use the given tool within the experminent to search for information.
When you are ready, you can make your decision on the screen by selecting the corresponding button.

Search Engine Introduction:
The search engine is embedded into your experiment screen:



When you are using the search engine, the results of your query are displayed on the same page:



Once you click on a link, the page will open in a new tab.
In case you have more than one search query make sure always use the search tool provided in the experiment.

Figure A8. Stage 2 Instructions: Search Engine Treatment

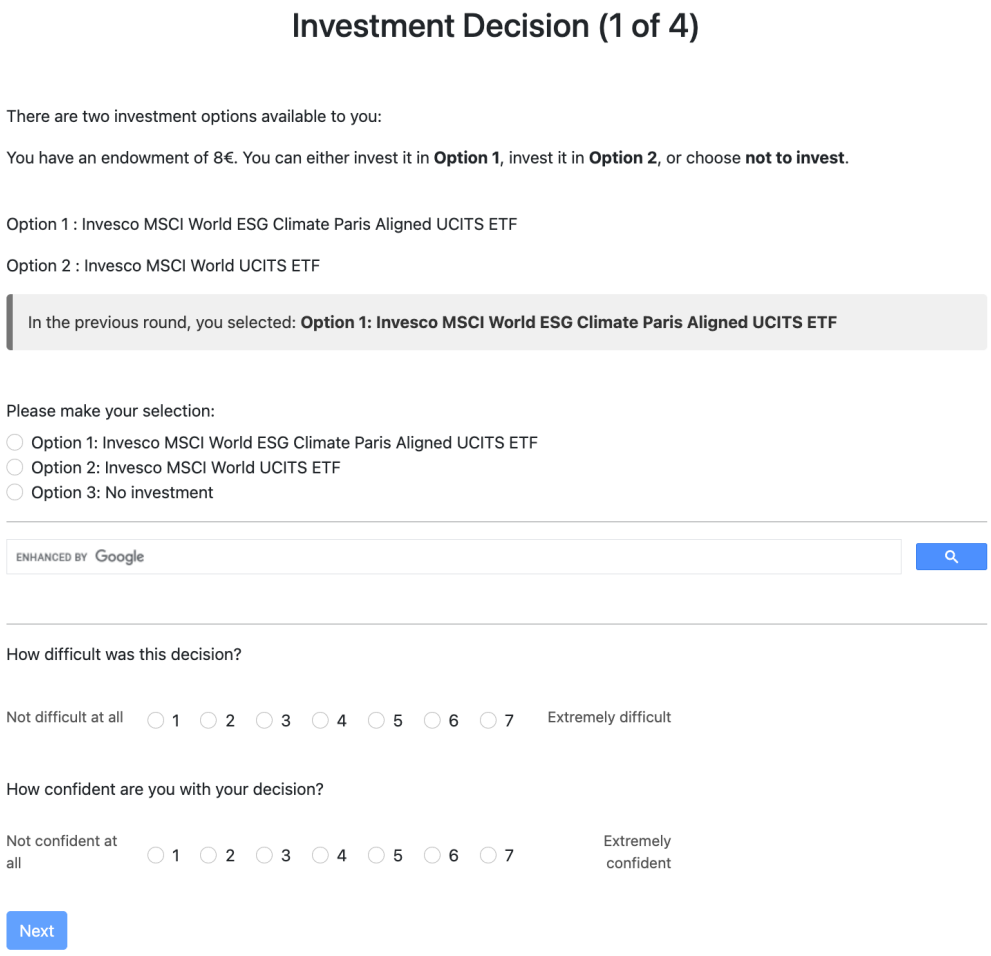


Figure A9. Decision Screen with Integrated Search Engine

Appendix A.2.3. Control Condition

Participants in the Control condition proceeded to Stage 2 without any additional information retrieval tools. Their decision screens in Stage 2 were identical to the Stage 1 screens shown in Figure A4a (or Figure A4b, depending on their motivation assignment), allowing them to revise their choices based solely on their own reflection.

Appendix B. Randomization and Balance

Table A1: Summary Statistics and Balance Checks

	Full Sample (N=374)	**Control** (N=107)	**LLM** (N=141)	**Search** (N=126)	**P-value** (Diff.)
Panel A: Demographics					
Age (Years)	34.4 (12.2)	34.27 (10.73)	35.99 (14.35)	32.70 (10.99)	0.252
Female (%)	45.5%	42.1%	45.0%	49.2%	0.579
Education: Bachelor (%)	51.1%	38.3%	58.6%	54.0%	0.016
Education: Master (%)	19.8%	28.0%	13.6%	19.8%	–
Panel B: Financial Characteristics					
Risk Willingness (1-10)	6.83 (2.08)	6.81 (2.04)	6.98 (2.26)	6.70 (1.92)	0.277
Fin. Lit: Interest Rates (% Corr.)	51.0%	63.6%	52.9%	36.2%	< 0.01
Fin. Lit: Rule of 72 (% Corr.)	61.0%	76.6%	52.1%	53.5%	< 0.01
Panel C: ESG Attitudes (1-7 Scale)					
ESG Knowledge	4.10 (1.81)	4.12 (1.80)	4.24 (1.80)	3.90 (1.84)	0.341
ESG Objectives	4.71 (1.63)	4.86 (1.53)	4.63 (1.65)	4.68 (1.70)	0.490
ESG Benefit	5.10 (1.42)	5.33 (1.24)	5.11 (1.40)	4.89 (1.55)	0.108
Panel D: Technology Usage					
LLM Usage: Daily (%)	43.9%	41.1%	54.3%	34.9%	< 0.01
LLM Usage: Weekly (%)	34.8%	27.1%	29.3%	47.6%	–

Note: Standard deviations are reported in parentheses for continuous variables. P-values in Column 5 are derived from Kruskal-Wallis tests (continuous) and Chi-square tests (categorical).

Appendix C. Further Analyses

DV: motivated choice	
motivated	0.118***
	(0.034)
Constant	0.401***
	(0.025)
Observations	1604
R-squared	0.014

Table A2: OLS Regression Results on Motivated Choice

Notes: OLS regression analysis. Clustered standard errors are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

DV: Is Confirmation Seeking	(1) Interaction Model
Interaction	
Search × Motivated Group	-0.1572*
	(0.093)
Main Effects	
Motivated Group	0.1610*
(Effect within LLM condition)	(0.088)
Treatment (Search)	-0.2068***
	(0.045)
Constant (LLM Baseline)	0.2423***
	(0.042)
Observations	674
R^2	0.148
F-Statistic	13.70

Table A3: OLS Regression: The Effect of Motivation on Confirmation Seeking

Notes: This table reports OLS estimates. The dependent variable is a binary indicator for confirmation-seeking queries. The baseline category (Constant) is the LLM condition without motivation. 'Motivated Group' indicates the additional effect of financial motivation on LLM users. Standard errors clustered by participant are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.