Biased AI Feed Exacerbates Disposition Effect: Evidence from the ETF Market of China

Abstract

Using account-level transaction data on trading platforms with different artificial intelligence (AI) functions, this study explores how AI influences the disposition effect among retail investors in China. Investors on the trading platform with more proactive decision-making AI are found to suffer from greater disposition effect and worse investment performance than investors on the platform without AI function. Investors on the trading platform where AI facilitates instead of replacing decision-making display the same level of disposition effect as investors whose trading platform provides no AI function. The study proposes that AI exacerbates retail investors' disposition effect due to 'corpus bias', which is from the near absence of 'sell' recommendations from analysts in China's capital market. This biased feed leads to the lack of instruction to sell on the platform, enhancing the unwillingness to sell losers and amplifying the disposition effect.

Keywords: disposition effect, artificial intelligence, corpus bias, Chinese ETF market

Acknowledgements

The author gratefully acknowledge the Research Training Program (RTP) scholarship provided by the Australian Government, which supported this research. The data used in this study were provided by an anonymous securities company in China under a confidentiality agreement; the data provider had no role in the design, analysis, or interpretation of the findings. The authors declare no financial relationships or other conflicts of interest relevant to this work. Earlier versions of this paper were presented at the 38th Australasian Finance and Banking Conference (AFBC, UNSW) and the Sydney Banking and Financial Stability Conference (SBFC 2025, University of Sydney); the authors also thank Professor Rangga Handika from Tokyo International University for his detailed and constructive comments as a discussant at SBFC 2025, which led to improvements in the paper. Supervisory guidance was provided by Dr. Grace Lepone and Dr. Di Bu at Macquarie University.

1. Introduction

In an era characterised by the global proliferation of financial technology (FinTech), artificial intelligence (AI) has been widely acclaimed as a potential remedy for the irrationality inherent in human investing (Hasan et al., 2023; Shanmuganathan, 2020). From sophisticated algorithmic trading systems to personalised robot-advisors (robo-advisors), the promise of AI is unequivocal: it aims to liberate investors from behavioural bias by harnessing the dispassionate objectivity and extensive data-processing capabilities of machines, thereby steering investors toward more optimal decision making (Ahmad et al., 2025). However, a troubling paradox is emerging: what if these meticulously designed AI tools do not alleviate the intrinsic forms of investing bias but instead exacerbate them (Eichler & Schwab, 2024; Romeo & Conti, 2025; Zhou et al., 2024)? This study provides an empirical examination of this proposition.

This research question is raised within the unique and explosive context of the exchange-traded funds (ETFs) market in China, in which assets under management (AUM) has just exceeded 6 trillion yuan by the end of 2025, with a participant base comprising over 20 million retail investors. Despite the fact that the size of China's ETF market still lags behind that of mature financial markets (USD 13.5 trillion in the U.S. and USD 3.2 trillion in Europe), with a 20-fold increase in the past eight years, China is undoubtedly one of the world's fastest-growing major markets. This distinctive market structure, the intersection of emerging technological tools and a substantial cohort of non-professional investors, provides an ideal setting in which to assess the real-world behavioural impact of artificial intelligence (AI).

China’s ETF market expansion has coincided with a surge in the popularity of artificial intelligence in securities trading software. While trading platforms were previously relatively uniform in functionality, recent years have seen marked differentiation across providers in terms of news richness, intelligent algorithms, automated investment tools, and AI-driven investment advice. Consequently, service differentiation has accelerated, making real transaction data through the various trading software available, which has enabled this empirical study.

While extensive research has been conducted on retail investor behaviour across various asset classes (Barber & Odean, 2013; Welch, 2022), only a small number of recent studies have addressed retail investing in ETFs (Gempesaw et al., 2024), and the understanding of Chinese retail investors' behaviour in this context remains limited. Studies utilising account-level data are even scarcer.

Through an empirical analysis encompassing over 130,000 real transactions conducted by more than 2,000 investors from a prominent brokerage firm over seven years, this study examines whether—and to what extent—different AI functionalities on various trading platforms affect retail investors' investment bias in China's ETF market. Specifically, the study focuses on the disposition effect—the tendency for investors to sell winning assets too early while holding onto losing assets too long (Shefrin & Statman, 1985; Odean, 1998). Conventional wisdom posits that AI should serve as a natural counterforce to the disposition effect, delivering emotion-free sell signals that assist investors in overcoming their impulses to secure their gains prematurely or to hesitate in cutting their losses. The current study aims to address the following core questions: do AI-powered trading systems of varying sophistication influence the disposition effect among retail investors and, if so, in what manner? What is the subsequent impact on their overall financial performance?

Results from the study provide counter-intuitive evidence: more proactive decision-making AI-embedded trading tools unexpectedly amplify the disposition effect among Chinese ETF investors, leading to demonstrably poorer financial outcomes. A novel mechanism—Corpus Bias—is proposed to explain this finding. We define Corpus Bias as systematic skew in training corpora that leads to asymmetric recommendation distributions. In China's capital market, structural features including distorted analyst incentives and the absence of short-selling mechanisms generate a systematically optimistic financial corpus. When investors hold losing positions, AI models trained on this corpus cannot generate stop-loss suggestions consistent with the training data, instead producing optimistic narratives that provide algorithmic cover for inaction—thereby reinforcing and amplifying the disposition effect.

This study makes three contributions. First, it provides account-level empirical evidence that proactive AI amplifies rather than reduces the disposition effect in an emerging market context, challenging the prevailing techno-optimistic narrative. Second, it introduces the concept of Corpus Bias as a structural mechanism linking AI training data quality to investor behavioural outcomes — a mechanism of universal relevance beyond China. Third, the

findings carry direct implications for FinTech product designers, financial market regulators, and individual investors who depend on AI to navigate market dynamics.

The remainder of this paper is organised as follows: Section 2 reviews the related literature and develops hypothesis; Section 3 describes data and methodology; Section 4 discusses the results and Section 5 concludes.

2. Literature Review and Hypothesis Development

2.1 Research on the Impact of AI Technology on Investor Bias

In recent years (2020–2026), whether AI counteracts or exacerbates investor behavioural bias has become a hot topic in the academic field of finance. Proponents argue that algorithmic objectivity decouples investment decisions from emotional interference: robo-advisors transform subjective risk perceptions into quantitative metrics, countering availability bias and overconfidence (Ahmad et al., 2025; Shanmuganathan, 2020), while deep reinforcement learning models decoupling (Hasan et al., 2023;). Experimental evidence lends some support to this view; Back et al. (2023) find that robo-advisors reduced investors' disposition effect in a controlled setting.

Advancements in AI architecture further enable structural interventions in behavioural loops. Deep reinforcement learning models achieve temporal decoupling, isolating historical experiences from current decisions to mitigate hindsight bias . Real-time monitoring systems intercept emerging forms of bias: algorithmic alerts flag irrational trading impulses (e.g., panic selling during volatility spikes), while personalised financial education modules address cognitive gaps (Gupta, 2025). Crucially, neural networks disrupt the compatibility principle—the human tendency to prioritise belief-consistent information—by independently evaluating new data streams (Hasan et al., 2023).

However, a growing body of evidence suggests that AI can inadvertently intensify behavioural biases through several pathways. First, cognitive offloading triggers responsibility transfer: investors increasingly outsource judgement to algorithms, accepting AI outputs with minimal scrutiny and attributing losses to the AI rather than their own decisions, thereby reinforces self-attribution bias (Rastogi et al., 2022; Shukla & Shukla, 2023). the data on which AI rely inevitably contain subjective components (Eichler &

Schwab, 2024). Second, certain AI designs directly validate inaction. Portfolio optimisation algorithms that prioritise stability may algorithmically justify holding depreciating assets, while anchoring reference points to historical purchase prices, thus cementing mental accounting errors (Shukla and Shukla, 2023; Kahneman and Tversky, 1979). Third, large language models (LLMs) introduce unique amplification risks through financial corpus contamination: training data saturated with market optimism induces overconfidence and sector-specific representativeness errors (Zhou et al., 2024; Eichler & Schwab, 2024). Finally, transparency features can backfire when simplified AI explanations function as accuracy validators, inadvertently strengthening confirmation bias (Romeo & Conti, 2025).

Two critical knowledge gaps persist in this literature. First, most of existing studies focus on Western markets, leaving the dynamics of retail-dominated markets (e.g., China's ETF ecosystem),where AI-bias interactions may be qualitatively different, largely unexplored. Second, the literature has not systematically distinguished between AI's effects on preference-based versus belief-based biases. Bini et al.(2026) provide a theoretical framework for this distinction: in preference-based tasks, LLM responses become increasingly human-like as models become more advanced, while in belief-based tasks, advanced models frequently generate rational responses. This suggests that perference-based bias, such as disposition effect, may be particularly susceptible to amplification by advanced AI, a proposition this study examines empirically in the context of China's ETF market, where a systematic corpus bias further distorts the information environment.

2.2 Research on the measurement of Disposition Effect

The disposition effect, the tendency to sell winners too early and hold losers too long, is grounded in prospect theory's loss aversion and mental accounting mechanisms (Kahneman & Tversky, 1979; Thaler, 1985). Odean (1998) was the first to systematically document and quantify this bias, proposing a ratio-based measure comparing the proportion of gains realised (PGR) to the proportion of losses realised (PLR). While intuitive and widely cited, this approach requires excluding positions that are neither fully sold nor held throughout the observation period, resulting significant data loss and potential sample selection bias.

Feng and Seasholes (2005) addressed these limitations by introducing survival analysis through the Cox proportional hazards model, treating each position's holding period as the time-to-event variable and gain/loss status as the key time-varying co-variate. This approach

explicitly models the dynamics of the sell decision over time and accommodates censored observations, making it particularly suited to account-level transaction data. Many subsequent studies have widely adopted this framwork (Seru et al., 2010; Frydman & Rangel, 2014; Heimer, 2016; Frydman & Wang, 2020).

More recently, Frydman and Wang (2020) employed an OLS regression framework to study how salience shocks affect the disposition effect, modelling the daily sell decision as a linear function of gain/loss status and interaction terms. This approach offers the key advantage of accommodating additional control variables and interaction effects on the right-hand side of the equation, making it particularly well-suited to examining how a specific factor—in the present study, AI functionality—differentially amplifies or attenuates the disposition effect across investor groups. Following De Winne (2021), who systematically compared ratio-based and econometric approaches and found the latter more appropriate when the research objective is to isolate the marginal effect of a specific explanatory variable, the present study adopts the OLS regression as its primary specification, with the Cox proportional hazards model serving as a robustness check..

2.3 Hypothesis Development

The literature reviewed above establishes that AI's impact on investor behaviour is neither uniformly beneficial nor uniform across bias types. In China's capital market, a structural corpus bias compounds this problem: the near-absence of bearish analyst reports—driven by professional disincentives against negative coverage and the lack of short-selling mechanisms—means that AI models trained on domestic financial corpora are systematically deprived of sell-side logic. When investors hold losing positions, the AI cannot generate stop-loss suggestions consistent with the training data; instead, it produces optimistic narratives that provide algorithmic cover for inaction. For retail investors who rely on AI for decision support, this absence of sell signals is cognitively interpreted as endorsement of holding—transforming corpus bias into an amplifier of the disposition effect.

This mechanism operates through three interlocking channels: loss aversion armed with AI-generated data rationalisation; information overload that triggers cognitive shortcuts and heuristic simplification; and corpus-induced distortion of reference point anchoring. Together, they lead to the following hypothesis:

H1: Retail investors using trading platforms with more proactive decision-making AI exhibit a significantly stronger disposition effect than investors using platforms with basic or tool-based AI functionality.

H2: The amplification of the disposition effect by proactive decision-making AI is associated with worse realized investment performance.

3. Data and Methodology

Building upon the established methods that have examined relationships between disposition effects and influencing factor models, this study re-evaluates these relationships using fresh data from China's rapidly developing ETF market.

This study utilises ETF data from January 2018–December 2024, covering 1,000 ETF products and 2,000 retail investors, supplying over 130,000 transactions. This account-level database comprises investors' profile information as well as their ETF transaction records. This proprietary data set was acquired through a collaboration with a securities brokerage firm in China, under a stringent confidentiality agreement that ensured the protection of both investor anonymity and commercial sensitivity. The current study's data set is very similar to the Chinese data sets used by Feng and Seasholes (2005) and Frydman and Wang (2020), as well as to the US data set used by Odean (1998).

This study's data set comes from a brokerage company which is a provincial state-owned securities company which has operations extending across most provinces in China. While it ranks as a second-tier securities firm on the national level, it holds a leading position within its own province. Its market share for brokerage business consistently fluctuates between 15% and 20% throughout the year, significantly surpassing that of top-tier national brokerage firms. Consequently, the data provided by this firm are expected to be representative of the province's market dynamics.

Based on the assumption of regional homogeneity, this study employs data from investors in a single province as the research sample. This geographic boundary qualification seeks to control the heterogeneous interference of potential confounding factors, such as regional cultural characteristics (e.g., risk preferences, investment traditions) and the market's

institutional environment (e.g., regulatory intensity). Therefore, this methodology enhances the effectiveness of the study's model in identifying causal relationships between core explanatory variables and dependent variables. Although limitations may exist regarding regional extrapolation due to focusing on one specific area, concentrating on this distinct regional market allows this study to more accurately capture the interaction mechanism between trading software choices and types of retail ETF investor bias, which is in line with the "institutionally embedded" research paradigm proposed by Barberis et al. (2018).

In addition, based on the code of the ETF fund being traded, data are obtained regarding the publicly available characteristics of the fund, provided that this information is accessible in at least one of two major financial databases: China Stock Market & Accounting Research (CSMAR) and Wind Economic databases.

3.1 Data and Summary Statistics

The investor profile data include various demographic variables, including gender, age, geographic location and account open dates. The gender distribution among investors in the sample is relatively balanced, with 44% identifying as female and 56% as male. The age structure reflects characteristics typical of middle-aged Chinese individuals, with an average age of 51 years (standard deviation of 11.7 years), ranging from young investors of 24 to senior investors aged 96. All investors in the sample are drawn from long-term clients who opened accounts before 2018.

The transaction file provides data at account-date-ticker level and includes details such as: trade date (year-month-day); ETF ticker symbols; buy/sell indicators; trade volume (value in renminbi [RMB]); trade quantity (number of ETF shares); transaction price (net of transaction costs); discount/premium ratio; and transaction channels, namely, mobile applications or web-based platforms offered by different software companies (marked as A, B, C and Others).

Furthermore, the fund portfolio file includes daily market values of ETF fund holdings for each investor as well as the total size of all assets held within their accounts (including cash, stocks, bonds and other assets).

During the study's observation period from 2018–2024, internet and mobile internet usage in China has been widely prevalent. Retail investors rarely place orders in person at their home

branch of this brokerage company or over the phone. Instead, nearly all retail investors check prices and execute trading orders through computer software or mobile applications during the study's observation period. The choice of transaction channel serves as a key variable within the study's empirical methodology. There are three transaction platforms used by investors in the sample.

Software A is provided directly by the securities firm itself and is exclusively available to its customers. As a second-tier domestic securities company, this company has limited investment in FinTech. While Software B and C are external providers that have captured approximately 80% market share by cooperating with securities firms in China, Software A only offers basic functionality with information derived solely from the perspectives of the securities firm's research team. The display of information and the function menu are traditional and lack proprietary AI functions. Software A is used as the benchmark for the current study.

<Insert Figure 1>

Software B specialises in technical analysis and providing software customisation services. It is renowned for its powerful flexible and efficient technical analysis features and formula platform. Its core strength lies in the responsiveness and depth of its K-line, indicator and line drawing tools. It is particularly suitable for users engaged in technical analysis and quantitative investment. By providing powerful tools and platforms, its AI is primarily reflected in the automated execution and optimisation of technical analysis rules. This platform empowers users by empowering them to build and test their own strategies. In this software, AI acts as an implementer, not a consultant.

<Insert Figure 2>

Software C is among the pioneering companies in China to promote mobile application trading software in collaboration with brokerage firms. Its AI platform was introduced in 2017, and by 2018, it had launched a natural language robo-advisor supporting investment queries in Mandarin. In the same year, it partnered with major financial institutions to launch an intelligent investment advisory function that recommends fund portfolios based on customers' risk preferences while utilising asset allocation algorithms. Thus, the current study posits that 2018 marks Software C's entry into AI era. In 2022, Software C released an upgraded large language model built on Transformer architecture that supports complex semantic understanding. The company is widely recognised as a leading adopter of AI and

related technologies, with an ambition to build a comprehensive AI-powered investment service ecosystem.

<Insert Figure 3>

Software C combines conversational and visual approaches which make it extremely user-friendly for retail investors, thus lowering the barrier to entry for AI users. It acts as an AI assistant that is proactive and driven by suggestions. The AI attempts to directly participate in, or even replace, the user's decision-making process, providing clear buy and sell points, hot spot predictions, etc.

<Insert Figure 4>

The transaction data sets and fund portfolio data sets are integrated to construct a holding sample at the investor-ticker-day level. Investor days during which no trades occur are included in the analysis. Purchases and sales are aggregated on a daily basis for each investor-ticker combination. The study documents daily holding status, along with decisions regarding whether to sell each ETF fund per investor, while counting those days on which positions are opened or when ETF funds are sold (including partial sales). For positions for which a purchase price cannot be found in the transaction file—primarily because investors acquired these ETF funds prior to the commencement of the study's sample period—the current study, unlike the work of Frydman and Wang (2020), does not exclude these data. Instead, the current study assumes that these investors purchased the position at the beginning of the study's sample period, with the study using the end-of-day price of that day as the purchase price.

<Insert Table 1>

The descriptive statistics from the study's data sample (Table 1) indicate that investors across the three platforms exhibit homogeneity in terms of demographic characteristics (mean age from 50.5–54.8 years, gender ratio between 0.51 and 0.52). However, their trading behaviour characteristics reveal significant heterogeneity. Specifically:

(1) Software C has the largest sample size (N=1211), demonstrating the highest transaction activity (70 transactions per investor) alongside the lowest single trade volume (12,898 yuan), suggesting that its user base may tend to carry out high-frequency small transactions.

(2) The sample size of Software A users is moderate (N=793), with an average annual transaction frequency per investor that is relatively low (30 transactions). This, combined with a relatively low average transaction amount per trade (22,498 yuan), implies characteristics that are typical of individual investors.

(3) Software B has the smallest sample size (N=441), yet it exhibits a significantly higher average single transaction amount (209,869 yuan) compared to other platforms; Additionally, it boasts a larger average portfolio size (258,765 yuan) and number of years since account-opening year (11 years), with both aligned with profiles that are characteristic of high net worth investors.

Based on these investor profiles, the current study preliminarily predicts that investor characteristics for Software A and C are relatively similar, allowing better assessment of how software functionality influences investment bias in subsequent regression analyses. Conversely, due to Software B's smaller sample size, along with its higher asset scale and single transaction volume, further analysis controls for any estimation bias potentially introduced by these differentiated characteristics.

3.2 Disposition Effect Measurement and Regression

With De Winne (2021) comparing different measurement approaches for the disposition effect, econometric approaches, relying on the estimation of a regression model, are more suitable for the current study than ratio-based approaches. The principle of econometric approaches is to model the decision to sell or to hold a portfolio position and to explain this decision by indicating whether this position is at a gain or at a loss. The main advantage of this approach is that it allows variables and controls to be added on the right-hand side of the model equation. This helps to study the impact of a certain factor on the disposition effect.

The models utilised comprise a simple ordinary least squares (OLS) regression model, a logistic regression model or a hazards model.

The OLS regression is designed to assess the propensity to sell or hold a portfolio position and to elucidate this decision using at least one dummy variable that indicates whether the position is at a gain or at a loss. This approach was employed by Chang et al. (2016) to investigate how the disposition effect is influenced by whether the asset held is a stock or a fund and by Frydman and Wang (2020) to analyse the effects of salience shocks on the

disposition effect. The current study adopts the OLS regression model as its foundational framework as follows:

$$Sale_{i,j,t} = \alpha \quad + \boldsymbol{\beta_1} Gain_{i,j,t} + \boldsymbol{\beta_2} AI_function_i + \boldsymbol{\beta_3} Gain_{i,j,t} * AI_function_i + \gamma Controls \tag{1}$$

To test whether any observation is driven by a specific regression design, another model from the existing literature is used. For the above equation, the variables are described as follows:

$Sale_{i,j,t}$ is equal to 1 if investor *i* sells ETF *j* on day *t*, and 0 otherwise.

$Gain_{i,j,t}$ is a dummy variable equal to 1 if investor *i* has a paper gain on position *j* on day *t*, and 0 otherwise.

AI_function is a categorical variable, with Software A serving as the benchmark. The classification reflects different characteristics of AI sophistication:

AI-function=0, if Software A is used: this represents baseline or minimal AI functionality, typically characterised by rule-based systems or elementary automation features.

AI-function=AI-B=1, if Software B is used: this denotes tool-based AI with this platform empowering users by empowering them to build and test their own strategies. In this way, AI acts as an implementer, not a consultant.

AI-function=AI-C=1, if Software C is used: this signifies decision-making AI, serving as an AI assistant that is both proactive and suggestion-driven. It attempts to directly participate in, or even replace, the user's decision-making process, providing clear buy and sell points, hot spot predictions, etc.

In regression specifications, these categorical variables are typically incorporated using dummy variable encoding. Indicator variables *AI-B* and *AI-C* are for Software B and Software C, respectively, with Software A serving as the reference category.

Categorisation is derived from qualitative assessment of vendor specifications, technical documentation and functional descriptions rather than from quantitative metrics, reflecting that the current state of AI implementation taxonomy is industry practice.

While this categorisation relies on subjective assessment of technical descriptions, this approach is consistent with the emerging literature on AI measurement (Acemoglu et al.,

2022). The ordinal nature of the variables permits the analysis of monotonic relationships between AI sophistication and outcome variables, although the measurement limitations inherent in qualitative technology classification are acknowledged.

The hazards model was initially applied by Feng and Seasholes (2005) to study the disposition effect. It serves as a statistical tool in survival analysis, commonly used within health departments to evaluate mortality occurrences based on various factors such as treatment. The hazard rate, denoted *h(t),* represents a measure of conditional probability. When applied for measuring the disposition effect in the current study, the hazard rate reflects daily probabilities, with these influenced by being in a gain or a loss position. While hazards models share similarities with logistic regression models, they offer significant advantages due to their explicit modelling of time until an event occurs. As for baseline hazard functions, the Cox proportional hazards model (Barber & Odean, 2013) provides great generalisability through its semiparametric specification. Therefore, this hazards model (Feng & Seasholes, 2005; Frydman & Wang, 2020; Heimer, 2016; Seru et al., 2010) was employed as a robustness check to assess the disposition effect. In the current study's analysis framework, the impact of different AI-functions software on disposition effect is estimated utilising the following regression model:

$$h_{i,j,t} = h_0(t)exp(\boldsymbol{\beta_1} Gain_{i,j,t} + \boldsymbol{\beta_2} AI_function_i + \boldsymbol{\beta_3} Gain_{i,j,t} * AI_function_i + \gamma Controls \ ) \quad (2)$$

In this model, each new position of an ETF is treated as a starting point, with the study assuming that this position concludes on the date when investors for the first time sell part, or all, of their holdings. The time-varying co-variables include daily observations indicating whether each position corresponds to a gain or a loss.

The study collects data about the trading software employed by each investor and about every transaction, classifying them according to their AI function. The hazard rate, $h_{i,j,t}$, represents investor *i*'s probability of selling position *j* on day *t* conditional upon not having sold prior to day *t*, while $h_0(t)$ denotes the baseline hazard. Notably, the Cox proportional hazards model does not impose any specific structure on this baseline hazard function. Other variables are as defined previously. The coefficient, β_1 measures the disposition effect; the hazard ratio, e^{β_1}, captures the ratio between the hazard rate of selling winners and the hazard rate of selling losers. Furthermore, the coefficient of interaction term β_3 reflects how AI functions influence disposition effects. $Sale_{i,j,t}Gain_{i,j,t}$

3.3 Control Variables and Individual Fixed Affect

Consistent with the empirical literature on salience and investor behaviour (Frydman & Wang, 2020), the current study's specification incorporates a comprehensive set of control variables to mitigate omitted variable bias. The three investor groups that use different trading software, as previously described in Table 1, show certain differences in trading frequency, average single transaction amount and portfolio size. These factors have been suggested in the literature as potentially influencing the disposition effect. Barber and Odean (2000) showed that frequent trading exacerbates forms of behavioural bias. Ben-David and Hirshleifer (2012), in examining transaction frequency as well as transaction size, found the disposition effect to be prevalent across both large and small transactions. To account for investor-level heterogeneity that may jointly influence disposition effect outcomes and to ensure that this study's estimates capture the net effect of AI sophistication, the following control variables are added and defined as follows:

Log-price is the logarithm of investor *i*'s historical average purchase price for fund *n*.

Log-asset is the logarithm of investor *i*'s total assets.

Trade-freq is the total number of transactions by investor *i* in all ETFs divided by the total transaction duration.

Trade-avg-value is the total transaction amount of investor *i* in all ETFs divided by the total number of transactions in fund *n*.

Position-value is investor *i*'s daily trading value in all ETFs divided by the total assets on that day.

Return_plus=1 when return>0; otherwise, it is 0. Return is investor *i*'s book return on fund *n*=(market value of fund *n*/historical average purchase cost of fund *n*)-1.

Mkt-ret is the daily return of the CSI 300 index.

The comprehensive nature of the current study's control variables minimises the risk that omitted factors drive the results, although the inherent limitations in establishing causality with observational data are acknowledged. Although the regression controls observable investor characteristics, it may omit time-invariant factors which are unobservable or difficult to measure. To further mitigate this heterogeneity concern, the study introduces individual

fixed effect. This approach helps to eliminate the influence of permanent differences between individuals and provides estimates that are closer to causal identification.

3.4 Performance Comparison

Beyond illustrating investors' forms of bias related to disposition effects, this study also evaluates performance differences between investors who utilise different trading software platforms. Three returns are computed: (1) buy-and-hold returns; (2) realised returns; and (3) holding portfolio returns. The respective returns are annualised in accordance with the holding period for easier comparison.

The buy-and-hold return quantifies total returns accrued by investors from purchasing ETFs until the conclusion of the study's observation period—regardless of any sales made during that time frame—thus reflecting outcomes from an entire passive investment strategy focused solely on buying and holding assets. The buy-and-hold return represents the opportunity cost of active trading by isolating an investor's fundamental Excess Buy-and-Hold Return. It answers the question: "was the investor's initial idea a good one?" A high buy-and-hold return means a well-chosen portfolio of ETFs that performed well overtime. A low buy-and-hold return suggests the initial ETF selection was poor, regardless of subsequent trading.

The realised return accounts for actual earnings/losses achieved by investors through all buying and selling activities, along with cash flows throughout these periods. It thus provides insight into real-world performance stemming from an active trading strategy. More importantly, the gap between buy-and-hold returns and realised returns quantifies the 'behavioural tax' paid by an investor for his/her bias.

The holding portfolio return is the daily rate of return on an investor's portfolio compared to his/her historical holding costs, averaged and annualised over the observation period. This return is the one felt every time an investor logs in. The holding portfolio return captures the investor's daily psychological journey and its influence on his/her perception of risk and future decisions. It reflects the constant stream of paper gains and losses that shape an investor's sentiment and willingness to engage with the market.

By analysing all three returns, this study paints a comprehensive picture including the quality of the initial decision (buy-and-hold return), the behavioural penalty paid (the gap between

the buy-and-hold return and the realised return), and the daily psychological experience that drives these behaviours (holding portfolio return).

4. Results and Discussion

4.1 Main Findings

This study estimates the OLS regression in (1) using linear regression, with the results reported in Table 2 Column 1. The coefficient associated with the *Gain* dummy variable is significantly positive, indicating that the benchmark investor (using basic Software A) already demonstrates the disposition effect. Being in a gain position increases his/her daily probability of selling by 1.06 percentage points, which is highly significant (p=0). This confirms a baseline irrationality.

The positive and significant coefficient of the interaction term Gain*AI-C implies that using the most advanced Software C amplifies the disposition effect. A Software C user's disposition effect is the baseline plus the interaction (1.06+0.42=1.48). This means his/her bias is significantly stronger ($p<0.005$) than that of the benchmark user. The coefficient of the interaction term Gain*AI-B is positive but statistically insignificant (p=0.28). Based on this model, this implies that using Software B has no discernible impact on the disposition effect compared to the benchmark. The control variables perform largely as expected by economic theory. For instance, the current study finds that age is negatively associated with the probability of selling, while log-asset, trade-freq and mkt-ret are all positively and significantly associated with selling, consistent with prior literature.

In summary, the results of the OLS regression model reveal a preliminary, although somewhat puzzling, correlation: using Software C is significantly positively associated with a stronger disposition effect.

<Insert Table 2>

The previous regression results answered the question: do investors who use advanced AI software have a stronger bias? This correlation can be caused by the impact of AI tools, or it

can be interpreted as a selection effect. This could mean that these investors, who actively choose to use the most complex and 'professional' Software C from the beginning, are themselves a special group that is more aggressive, more confident and more likely to produce a disposition effect. Individual fixed effect is added to further mitigate this heterogeneity concern. The model with individual fixed effect isolates the pure processing effects of the software by comparing each investor only to his/her own changes over time. This eliminates any stable, pre-existing differences between users.

Table 2 Column 2 reports the estimation results of an OLS regression in (1) with individual fixed effect. Overall, the previous conclusion, that the software with more proactive decision-making AI enhances investors' disposition effect, is still supported, but the significance level has decreased slightly. The coefficient of *Gain* is 1.16, $p<0.001$. This shows that, on average, when an investor's position is in a gain, their probability of selling increases by 1.16 percentage points, slightly higher than in the situation without fixed effect. The coefficient of *Gain*AI-C* is 0.28, $p=0.088$. In a fixed-effects model, this positive and marginally significant coefficient implies a causal treatment effect. It suggests that when an investor starts using Software C, it causes his/her disposition effect to increase by 0.28 percentage points (relative to his/her own baseline behaviour on Software A). This is a counter-intuitive finding. It suggests that the 'advanced AI' software is not making investors more rational; instead, it may be amplifying their various forms of bias. The coefficient of *Gain*AI-B* is positive but statistically insignificant, meaning that investors using Software B do not suffer from the disposition effect more or less than the baseline investors using Software A.

In the individual fixed-effect model, all the time-invariant control variables, such as *age* and *gender*, are absorbed. The remaining control variables perform largely as expected by economic theory. For instance, *log-asset* and *mkt-ret* are both positively and significantly associated with selling, consistent with prior literature.

Comparing regression results before and after adding individual fixed effects opens the black box and splits the AI-enhanced disposition effect into two parts: selection effect or treatment effect. In selection effect, the type of investor who is attracted to Software C in the first place is already more biased than the average investor. In treatment effect, the software causally amplifies this pre-existing bias, most likely by corpus bias, loss aversion or information overload. Corpus bias is a key factor, especially in the Chinese capital market. Due to the lack of short-selling mechanisms and skewed analyst incentives, the corpus feed to AI is severely

lacking in negative or bearish information, leading to an asymmetric optimism bias. This bias is inherited and amplified by AI, ultimately strengthening the disposition effect.

The study's regression analysis does not find significant evidence that Software B has affect disposition effect like Software C does. One possible reason is that the type of AI function in Software B is significantly different from that in Software C. Unlike Software C, Software B is more about providing tools that facilitate self-guided analysis and decision making than proactively offering investment recommendations . As such, Software B is less affected by corpus bias as much as Software C is.

We acknowledge that the relatively low R-squared in our regression suggests that our model does not capture the full spectrum of factors influencing an investor's decision to sell. In the context of financial markets, trading decisions, particularly sells, are inherently complex and often driven by idiosyncratic, unobservable liquidity shocks or personal circumstances that are not reflected in standard datasets. For instance, an investor may be compelled to liquidate a position due to an urgent and unforeseen need for cash — a variable that is empirically difficult to observe and quantify. Given this inherent complexity, the primary objective of our regression is not to build a comprehensive predictive model for selling behavior. Rather, our aim is hypothesis testing: we seek to isolate and identify whether our variable of interest (AI adoption) has a statistically significant marginal effect on the sell decision after controlling for observable characteristics. The low R-squared does not undermine this objective; it simply reflects the high level of noise typical of individual-level financial data. Our results provide clear evidence that, despite this noise, investors with more proactive decision-making AI exhibit systematically different selling patterns, thereby confirming our core hypothesis.

4.2 Robustness Check

Following Heimer (2016) and Seru et al. (2010), this study employs a Cox proportional hazards model to estimate in （2） the change in probability when using different AI software platforms, as discussed in Chapter 4. Table 3 provides estimation results in (2) from the Cox proportional hazards model, with the current study's results remaining robust to this alternative method. The key output of a Cox model is the hazard ratio (*HR*) which is column 'exp(coef)' in Table 3. When $HR>1$, selling is more likely to happen; when $HR<1$, selling is less likely to happen. The *HR* of *Gain* is 1.35 ($p<0.005$), with this being the baseline disposition effect for the benchmark group (Software A). The *HR* is highly significant and

35% greater than 1. This means that, for an investor who uses Software A, a position in a gain state is 35% more likely to be sold at any given moment than a position in a loss state. This confirms a strong presence of the disposition effect: investors are irrationally eager to sell winners. The *HR* of *Gain*AI-C* is 1.08 (p<0.005), thus being statistically significant and greater than 1. This coefficient multiplies the baseline hazard ratio to compute the total hazard for Software C users: 1.35*1.08=1.458. This means that Software C amplifies the disposition effect, making investors an additional 8% more likely to sell their winners than the benchmark group. Their total bias is a 45.8% increased likelihood of selling a winner. The *HR* of *Gain*AI-B* is 1.05 (p=0.12). It is higher than 1, but statistically insignificant,potentially due to the reduced corpus bias compared to Software C.

When the study adds individual fixed effect to the Cox proportional hazards model, the results are similar to those in the OLS regression model when individual fixed effect is added. The HR of Gain is 1.32 (p<0.005) which is the baseline disposition effect, after accounting for all stable investor differences. The bias remains strong and highly significant: on average, an investor is 32% more likely to sell a winning position than a losing one. The HR of Gain*AI-C is 1.04, p=0.08, with the effect marginally significant (at the p<0.10 level), while the coefficient of the interaction term with AI-B is still statistically insignificant. It cannot be concluded that using Software B has any causal effect on investors' bias.

By comparing the Cox proportional hazards model results before and after adding individual fixed effects, this study can quantitatively outline the two components of the AI-enhanced bias as discussed earlier. Before adding individual fixed effects, the use of Software C makes investors an additional 8% more likely to sell their winners than the benchmark group. After adding individual fixed effects, the increment narrows to 4%. This means that the software itself makes the user about 4% more biased, with the remaining 4% of the effect is due to selection. The type of investor attracted to Software C in the first place is already about 4% more biased than the average investor.

<Insert Table 3>

4.3 Worse Performance of Proactive Decision-making AI Software

Does the bias amplification on the platform hurt investors' financial performances? The answer is clearly yes. Table 4 presents the key metrics:

1) Annualized Buy-and-Hold Return (BHR), which measures the total return an investor would have earned by purchasing an ETF at the initial transaction price and holding it until the end of the observation period, irrespective of any intermediate trading activity. To benchmark this ability against the market, we also report Excess Buy-and-Hold Return (BHR − Index) , defined as the difference between the buy-and-hold return and the return of the CSI 300 ETF index over the same period. This metric reflects the investor's security picking ability—the quality of the initial security selection. A positive (negative) value indicates that the investor's chosen ETFs outperformed (underperformed) the market, absent subsequent trading intervention.

2) Annualized Realized Return, which captures the actual returns investors achieve through their sequence of buy and sell decisions, incorporating all trading gains, losses, and transaction costs. This metric reflects the net outcome of active trading. To assess whether investors successfully time the market or add value through trading, we compute Excess Realized Return (RR − Index) , the difference between realized return and the return of the CSI 300 ETF index. A positive (negative) value indicates that the investor's trading activities generated returns above (below) a simple passive investment in the market index. We define Behavioural Tax as the difference between Buy-and-Hold Return and Realized Return (BHR-RR). This metric quantifies the value destroyed by trading.

3) Annualized Holding Portfolio Return, which is the daily rate of return on an investor's portfolio compared to his/her historical holding costs. This is the return that an investor considers every time he/she logs in, capturing the investor's daily psychological journey and influencing his/her sentiment and willingness to engage with the market. Based on the daily portfolio return and its volatility, we reported Sharpe Ratio, which measures risk-adjusted performance by quantifying the excess return per unit of total risk. It is calculated as the mean of daily excess returns (daily portfolio return minus the risk-free rate) divided by the standard deviation of those excess returns. This metric allows us to assess whether higher raw returns, when observed, merely reflect compensation for bearing higher volatility, or whether they represent genuine skill in generating return per unit of risk. A Sharpe ratio of zero indicates returns equivalent to the risk-free rate; negative values suggest that the investor's returns do not adequately compensate for the risks undertaken.

Table 4 presents a nuanced picture of performance across the three investor groups. Software C users, on average, exhibit the highest Excess Buy-and-Hold Return (0.68%, $p<0.01$) among

all groups. This finding suggests that proactive decision-making AI, with its superior information-processing capabilities, effectively identifies promising ETFs and generates sound buy recommendations. Investors relying on Software C appear to benefit from this enhanced selection ability at the initial purchase stage. However, this advantage does not translate into superior realized outcomes. Despite their strong security-picking ability, investors from Software C achieve the lowest realized return (-10%), both in absolute terms and after adjusting for market movements(RR-index).This striking divergence indicates that active trading decisions—likely influenced by the platform's AI functionality—substantially erode the potential gains embedded in their chosen portfolios.

Turning to risk-adjusted measures, the Sharpe Ratios are negative across all three platforms, indicating that investors' returns do not adequately compensate for the risks they undertake. Notably, investors from Software C exhibit the most negative Sharpe Ratio (-0.54, $p<0.01$), significantly worse than the benchmark group. This pattern suggests that the higher realized losses of Software C users are not merely a reflection of lower risk-taking. On the contrary, they achieve the worst risk-adjusted performance despite their elevated risk exposure. This suggests that proactive decision-making AI not only fails to improve risk-return trade-offs but may actively degrade them.

The Behavioural Tax, defined as the difference between buy-and-hold return and realized return (BHR-RR), quantifies the value destroyed by trading. Investors using proactive decision-making AI driven Software C pay a substantial behavioural tax of 11%, significantly higher than the other groups. This observation aligns with our earlier finding that Software C users suffer from the most pronounced disposition effect. By selling winning positions too early while holding losers too long, they forfeit the higher return potential embedded in their superior initial selections. The behavioural tax thus captures the tangible financial cost of AI-amplified bias—a cost that is most severe for investors relying on proactive decision-making AI.

The holding portfolio returns provide further insight into the mechanisms underlying the performance divergence documented above. As shown in Table 4, all three investor groups experience negative holding portfolio returns throughout most of the observation period, a phenomenon that reflects the broader experience of retail investors in China's capital market during that period. This means that, for the majority of that time, investors would log into their accounts to find their portfolios trading at a loss relative to historical purchase costs. It is

within the context that the role of AI—particularly the proactive decision-making AI embedded in Software C—becomes critically important. As argued earlier, the Chinese capital market is characterized by a systematic corpus bias. Due to institutional constraints such as the lack of short-selling mechanisms and analyst incentives that discourage bearish reports, the training data for financial AI models are severely depleted of negative or sell-side information. Consequently, AI assistants trained on such corpora rarely generate sell recommendations. The interaction between this corpus bias and investor natural prefence against selling losers creates a powerful behavioural reinforcement loop. When investors open their accounts to find their portfolios in paper loss, the AI—lacking the linguistic and analytical frameworks to articulate sell logic—offers no guidance to exit losing positions. For retail investors who rely on AI for decision support, this absence of sell signals is cognitively interpreted as an endorsement of inaction. The AI's silence becomes, in effect, a recommendation to hold. Thus, proactive decision-making AI, trained on biased corpora, transforms the passive experience of holding losses into an algorithmically reinforced inertia, amplifying the disposition effect and results in quantifiable financial cost.

This mechanism explains why Software C users, despite their superior security-picking ability, end up with the lowest realized returns, the most negative Sharpe ratios, and the largest behavioural tax. The very technology designed to enhance decision-making becomes a catalyst instead for behavioural bias, resulting in a measurable financial cost on its users.

<Insert Table 4>

A group comparison is conducted on the returns generated by different software platforms.. The results indicate that only the Realized Return and Excess Realized Return are significantly different between Software C and Software A ($p<0.01$). No other comparisons—whether for BHR, BHR-index or Sharpe ratios—reach conventional levels of statistical significance. This pattern of results suggests that the detrimental effect of proactive decision-making AI is most pronounced in the realized performance dimension—the actual outcomes of trading activity—rather than in security selection or risk-adjusted metrics. This directly shifts the focus of the discussion to the realm of behavioural finance. This result confirms and is completely consistent with the core hypothesis: that 'proactive decision-making AI affects investor performance by amplifying different forms of behavioural bias. The significant differences in realized returns are precisely the manifestation of the amplification of these forms of behavioural bias' (e.g., disposition effect) in returns. Software C, with more

proactive decision-making AI, did not impair users' security selection ability. Instead, it acted as a behavioural amplifier, magnifying the disposition effect embedded in human nature and systematically eroding their realized returns. This provides a more specific and feasible direction for regulation, algorithm design and investor education than 'improving security selection ability'.

On the other hand, the fact that only the C-A comparison for realized return is significant indicates that the effect is specific to the more proactive decision-making AI platform relative to the baseline, rather than a general phenomenon across all AI-enhanced platforms.

<Insert Table 5>

5. Conclusion

This study investigates the paradoxical relationship between the increasing sophistication of artificial intelligence (AI) in financial technology (FinTech) and the enduring irrationality of retail investors. The central research question is whether AI-powered trading systems, designed to promote rational decision making, mitigate or amplify ingrained forms of behavioural bias, such as the disposition effect. By analysing a unique, account-level data set of over 130,000 ETF trades in the rapidly growing Chinese market, this study provides compelling evidence for its counter-intuitive and concerning conclusion. The study finds that more proactive decision-making AI is associated with a significant amplification of the disposition effect, leading to demonstrably poorer financial outcomes for investors. Corpus bias is a key factor leading to the AI-amplified disposition effect in China's capital markets.

The finding presented in this study challenges the prevailing techno-optimistic narrative that AI is a panacea for human bias in financial markets. Rather, it serves as a behavioural amplifier, scaling the impact of investors' innate psychological tendencies without altering their fundamental direction. In the Chinese ETF market, where AI training corpora suffer from systematic bias toward optimistic narratives, this amplification effect manifests as a pronounced exacerbation of the disposition effect. Software C users, equipped with the most proactive decision-making AI, did not become more rational investors; they became more efficient versions of their biased selves—quicker to sell winners, slower to cut losses, and ultimately poorer in realized returns.

The implications of this study are significant and wide-ranging, extending to FinTech designers, regulators and investors themselves. For the FinTech industry, these findings serve as a critical call to action. The value of AI lies not in its technical sophistication alone, but in the alignment between its design and the psychological realities of its users. When that alignment fails, technology does not correct human nature—it magnifies its flaws. Decision-making AI enhances user stickiness and trading frequency by crafting a 'pleasurable' trading experience, for example, through hot-topic pushes and one-click copy trading. In the short run, this proves advantageous for the company's revenue, specifically its commissions. Nevertheless, from a long-term perspective, it undermines users' long-term wealth accumulation and erodes trust. This requires a departure from merely catering to various forms of user preference and focusing on the alignment between its design and the psychological realities of its users.

For regulators, this study highlights the urgent need for a new framework of algorithmic governance. As AI becomes more influential in shaping financial decisions, supervision must evolve to consider the behavioural side-effects of these complex systems. Disclosures about a platform's potential to amplify various forms of bias may become as crucial as traditional risk warnings. Regulators need to acknowledge that one key cause of forms of AI bias lies within the data used to train these systems. Promotion of the legitimacy of 'bearish thinking' is a pressing need. At the institutional level, regulators should actively encourage and support sound bearish research and reporting. Fostering a diverse range of voices in the market is essential for maintaining a healthy and efficient financial ecosystem. To this end, regulations should be designed to safeguard those who issue objective negative evaluations from unjust pressure. This not only protects the freedom of information dissemination but also serves as the bedrock for improving the overall market information environment. By ensuring a balanced flow of information, regulators can help prevent market over-optimism and potential bubbles.

For investors, this study reveals that utilising AI software is not merely about enjoying a service; rather, it is about forging a 'cooperative relationship'. In the realm of investment, cross-validation, critical thinking and an international perspective are indispensable. When it comes to decision-making AI, investors should view it as a "barometer of mainstream market optimism" rather than an infallible investment guide. The recommendations provided by such

AI merely reflect what most in the market might be talking about, yet this does not necessarily equate to objective truth.

It is important to acknowledge the limitations of this study. The findings are based on transaction data and therefore establish a strong correlation. Although efforts were made to get a closer causal identification by adding control variables and individual fixed effects, the inherent limitations in establishing causality with empirical approaches are acknowledged. The unique characteristics of the Chinese ETF market, such as T+1 settlement, daily price fluctuation caps, and the prohibition of naked short selling, imposes trading limitations that are not present in most developed capital markets. Consequently, the generalizability of our results to other, less-regulated market contexts should be approached with caution. Further research can aim to disentangle the causal mechanisms and conduct comparative studies in more mature markets. With the growing prevalence of AI among investors, future research could draw on the Adaptive Market Hypothesis to explore how AI-driven behaviour shapes market ecology. Moving beyond individual-level analysis, subsequent studies should employ formal modeling and simulation-based approaches to examine the systemic implications of AI integration.

References

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1), S293-S340.

Ahmad, U., Van Keulen, M., Briassouli, A., & Saad, M. (2025). Cognitive biases, robo advisor and investment decision psychology: An investor's perspective from New York stock exchange. *Acta Psychologica*, 256, 105048.

Back, C., Morana, S., & Spann, M. (2023). When do robo-advisors make us better investors? The impact of social design elements on investor behavior. *Journal of behavioural and Experimental Economics*, 103, 101984.

Barber, B. M., Huang, X., & Odean T. (2016). Which factors matter to investors? Evidence from mutual fund flows. *Review of Financial Studies*, 29(10), 2600-2642.

Barber, B. M., Lin, S., & Odean, T. (2023). Resolving a paradox: Retail trades positively predict returns but are not profitable. *Working paper*, University of California, Davis, San Francisco State University and University of California, Berkeley.

Barber, B. M., & Odean, T. (2000). Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The Journal of Finance*, 55(2), 773-806.

Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *The Review of Financial Studies,* 21(2), 785-818.

Barber, B. M., & Odean, T. (2013). The behavior of individual investors. Chapter 22 in *Handbook of the economics of finance* (2B, pp. 1533-1570). Elsevier.

Barberis, N., Greenwood, R., Jin, L., & Shleifer, A. (2018). Extrapolation and bubbles. *Journal of Financial Economics*, 129(2), 203-227.

Ben-David, I., Franzoni, F., Kim, B., & Moussawi, R. (2023). Competition for attention in the ETF space. *The Review of Financial Studies*, 36(3), 987-1042.

Bini, P., Cong, L. W., Huang, X., & Jin, L. J. (2026). Behavioral economics of AI: LLM biases and corrections. *NBER Working Paper* 34745.

Chang, T. Y., Solomon, D. H., & Westerfield, M. M. (2016). Looking for someone to blame: Delegation, cognitive dissonance, and the disposition effect. *The Journal of Finance,* 71(1), 267-302.

De Winne, R. (2021). Measuring the disposition effect. *Journal of behavioural and Experimental Finance,* 29, 100468.

Eichler, K. S., & Schwab, E. (2024). Evaluating robo-advisors through behavioural finance: A critical review of technology potential, rationality, and investor expectations. *Frontiers in behavioural Economics*, 3.

Feng, L., & Seasholes, M. S. (2005). Do investor sophistication and trading experience eliminate behavioural biases in financial markets? *Review of Finance*, 9(3), 305-351.

Frydman, C., & Rangel, A. (2014). Debiasing the disposition effect by reducing the saliency of information about a stock's purchase price. *Journal of Economic Behavior & Organization*, 107, 541-552.

Frydman, C., & Wang, B. (2020). The impact of salience on investor behavior: Evidence from a natural experiment. *The Journal of Finance*, 75(1), 229-276.

Gempesaw, D., Henry, J. J., & Xiao, H. (2024). Retail ETF investing. *European Financial Management*, 30(4), 2305-2342.

Hasan, Z., Vaz, D., Athota, V. S., Désiré, S. S. M., & Pereira, V. (2023). Can artificial intelligence (AI) manage behavioural biases among financial planners? *Journal of Global Information Management*, 31(2), 1-18.

Heimer, R. Z. (2016). Peer pressure: Social interaction and the disposition effect. *The Review of Financial Studies*, 29(11), 3177-3209.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-292.

Noelle-Neumann, E. (1974). The spiral of silence: A theory of public opinion. *Journal of Communication*, 24(2), 43-51.

Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of Finance*, 53(5), 1775-1798.

Rastogi, C., Zhang, Y., Wei, D., Varshney, K. R., Dhurandhar, A., & Tomsett, R. (2022). Deciding fast and slow: The role of cognitive biases in AI-assisted decision-making. *Proceedings of the ACM on Human–Computer Interaction*, 6(CSCW1), 1-22.

Romeo, G., & Conti, D. (2025). Exploring automation bias in human–AI collaboration: A review and implications for explainable AI. *AI & Society.*

Seru, A., Shumway, T., & Stoffman, N. (2010). Learning by trading. *The Review of Financial Studies*, 23(2), 705-739.

Shanmuganathan, M. (2020). Behavioural finance in an era of artificial intelligence: Longitudinal case study of robo-advisors in investment decisions. *Journal of behavioural and Experimental Finance*, 27, 100297.

Shukla, S., & Shukla, S. K. (2023). Impact of automation bias and status quo bias on capital market investment decisions of Indian investors: An explanatory research. *Financial Markets, Institutions and Risks*, 1-11.

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing Science*, 4(3), 199-214.

Vaarmets, T., Liivamägi, K., & Talpsepp, T. (2019). How does learning and education help to overcome the disposition effect? *Review of Finance*, 23(4), 801-830.

Welch, I. (2022). The wisdom of the Robinhood crowd. *The Journal of Finance*, 77(3), 1489-1527.

Zhou, Y., Ni, Y., Gan, Y., Yin, Z., Liu, X., Zhang, J., Liu, S., Qiu, X., Ye, G., & Chai, H. (2024). Are LLMs rational investors? A study on detecting and reducing the financial bias in LLMs. *arXiv*: 2402.12713. arXiv.



Figure 1: Screenshot of menu items and information layout of Software A

Figure 1 shows the screenshot of Software A, which provides menu items and information layout. This is a typical template of trading software in China. The top menu bar included options such as Buy, Sell, Watchlist, F10 (in-depth Information), Cycle, Stock Selection, Panorama, Individual Stock, Sectors, Capital Markets, Research Reports, STAR market (Science and Technology Innovation Board), Dragon and Tiger List (a list in China's stock market that publishes trading information of the most active stocks daily, including top buyers, sellers and turnover rates), New-listed Stock (stocks that have just been issued and listed, available for investors to subscribe to and trade on the Stock Exchange).



Figure 2: Screenshot of investor's coding file embedded in Software B

Figure 2 shows the screenshot of Software B which provides tools for investors entering a file path or URL. Investors set conditions and rules to select daily top active sectors in their coding files, and then upload the files into Software B. The platform integrates index and stock movements, updating and outputting the results in real time.



Figure 3: Screenshot of robot-advisor in Software C

Figure 3 shows the screenshot of robot-advisor in Software C, which supports communicating with AI using natural language about stock selection. Even if investors do not ask questions, it will present a series of hot questions, guiding them to click and answer.



Figure 4: Screenshot of conditional fund selection in Software C

Figure 4 shows the screenshot of conditional fund selection in Software C, which supports fund selection using a single natural language sentence and lists numerous options related to fund characteristics, allowing investors to select and then output their fund selection results.

Table 1 Characteristics of investors on three platforms

	Software A		Software B		Software C	
Number of investors	793		441		1211	
Number of trades	23946		18873		84807	
Average no. of trades per investor	30		43		70	
Total trade value (million Rmb)	539		3961		1094	
Average value per trade (Rmb)	22498		209869		12898	
	Mean	**Median**	**Mean**	**Median**	**Mean**	**Median**
Portfolio size (RMB)	105781	16786	258765	22177	81200	11188
Number of ETFs	6	3	3	3	5	2
Age	51.2	51.0	54.8	54.0	50.5	50.0
Gender	0.52	0.98	0.51	0.86	0.51	0.91
Account open years	10.1	8.3	11.0	10.8	10.8	10.2

This table provides descriptive statistics on the basic characteristics and trading behaviour of the entire sample of investors used in this study. It is divided into two parts: Panel A presents the trading characteristics of investors across the three platforms, including number of investors, trading frequency and volume per trade. Panel B presents individual characteristics of the three investor groups, including age, gender, account opening time and portfolio size. This provides preliminary evidence for the subsequent exploration of heterogeneity across different user groups.

Table 2 Analysis of impact of AI function on disposition effect: OLS Regression

	OLS Regression without individual fixed effect	OLS Regression with individual fixed effect
Gain	1.06*** *(0.000)*	1.16*** *(0.000)*
Gain*AI-B	0.31 *(0.280)*	0.27 *(0.282)*
Gain*AI-C	0.42*** *(0.005)*	0.28* *(0.088)*
AI-B	0.42* *(0.090)*	0.44 *(0.152)*
AI-C	0.11 *(0.341)*	-0.17 *(0.674)*
log price	0.03 *(0.792)*	-0.08 *(0.553)*
log asset	0.13*** *(0.000)*	0.34*** *(0.000)*
trade_freq	23.47*** *(0.000)*	
trade_avg_value	-0.00 *(0.792)*	
position_value	0.00** *(0.029)*	0.00** *(0.037)*
return_plus	-3.48*** *(0.000)*	-3.65*** *(0.000)*
age	-0.01*** *(0.000)*	
gender	0.02 *(0.833)*	
mkt_ret	16.28*** *(0.000)*	16.36*** *(0.000)*
R-squared	0.0637	0.0038
No. Observations	2137735	2137735

*This table presents the results of an OLS regression in which the dependent variable is a dummy equal to 1 if an ETF is sold (fully or partially) on a given day, and 0 otherwise. Column 1 presents the coefficients of baseline OLS regression without fixed effect, while Column 2 presents the coefficients of OLS regression with individual fixed effect. The core independent variables are profit status (Gain) and its interaction term with a software type category variable (Software B, Software C). Gain is a dummy variable equal to 1 if the paper return since purchase is positive, and 0 otherwise. Log-price is the logarithm of investor i's historical average purchase price for fund n. Log-asset is the logarithm of investor i's total assets. Trade-freq is the total number of transactions by investor i in all ETFs divided by the total transaction duration. Trade-avg-value is the total transaction amount of investor i in all ETFs divided by the total number of transactions on fund n. Position-value is investor i's daily trading value in all ETFs divided by the total assets on that day. Return_plus=1, when return>0; otherwise, it is 0, while return is investor i's book return on fund n=(market value of fund n/historical average purchase cost of fund n)-1. Mkt-ret is the daily return of the CSI 300 index. Those time-invariant control variables are absorbed by individual fixed effect. All coefficients are multiplied by 100. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively. The numbers in parentheses under coefficients represent p-values.*

Table 3 Robustness check using Cox proportional hazards model

	Cox Model without individual fixed effect	Cox Model with individual fixed effect
Gain	1.35*** *(<0.005)*	1.32*** *(<0.005)*
Gain*AI-B	1.05 *(0.12)*	1.04 *(0.22)*
Gain*AI-C	1.08*** *(<0.005)*	1.04* *(0.08)*
AI-B	1.17*** *(<0.005)*	1.05 *(0.12)*
AI-C	1.25*** *(<0.005)*	1.00 *(0.99)*
log price	1.04*** *(<0.005)*	1.01* *(0.06)*
log asset	1.04*** *(<0.005)*	1.04*** *(<0.005)*
trade_freq	1.43*** *(<0.005)*	
trade_avg_value	1.00*** *(<0.005)*	
position_value	1.00*** *(<0.005)*	1.00** *(0.01)*
return_plus	0.40*** *(<0.005)*	0.59*** *(<0.005)*
age	1.00*** *(<0.005)*	
gender	1.00 *(0.87)*	
mkt_ret	1668.88*** *(<0.005)*	2224.15*** *(<0.005)*
Cox & Snell R-squared	0.003	0.231
Nagelkerke R-squared	0.008	0.635
No. Subjects/Periods	1715913	1715913

*This table presents the results of the Cox proportional hazards model in which the dependent variable is the 'hazard' or instantaneous risk of the event occurring (selling an ETF). Column 1 presents the hazard ratio (HR) which is 'exp(coef)' of baseline Cox model without fixed effect, while Column 2 presents the hazard ratio (HR) of Cox model with individual fixed effect. The core independent variables are profit status (Gain) and its interaction term with a software type category variable (Software B, Software C). Gain is a dummy variable equal to 1 if the paper return since purchase is positive, and 0 otherwise. Log-price is the logarithm of investor i's historical average purchase price for fund n. Log-asset is the logarithm of investor i's total assets. Trade-freq is the total number of transactions by investor i in all ETFs divided by the total transaction duration. Trade-avg-value is the total transaction amount of investor i in all ETFs divided by the total number of transactions on fund n. Position-value is investor i's daily trading value in all ETFs divided by the total assets on that day. Return_plus=1, when return>0; otherwise, it is 0, while return is investor i's book return on fund n=(market value of fund n/historical average purchase cost of fund n)-1. Mkt-ret is the daily return of the CSI 300 index. Some time-invariant control variables are absorbed by individual fixed effect. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.*

Table 4 Comparison of investment performance by trading software user group

	Software A	Software B	Software C
Buy-and-Hold Return (BHR)	0.90%	1.14%	0.96%
Realized Return (RR)	45.45%	31.67%	-10.0%
Behavioural Tax (BHR-RR)	-45%	-31%	11%
Holding Portfolio Return (HPR)	-0.71%	-2.43%	-1.56%
Excess Buy-and-Hold Return (BHR-index)	0.29% (0.211)	0.66%* (<0.1)	0.68%*** (<0.01)
Excess Realized Return (RR-index)	45.42% (0.292)	31.58%* (<0.1)	-9.93% (0.292)
Sharpe Ratio	-0.25*** (<0.01)	-0.50*** (<0.01)	-0.54*** (<0.01)

*This table presents a comparative analysis of the investment performance across different investor cohorts identified in this study. The key metrics are reported: Annualized Buy-and-Hold Returns, which reflect the performance of the underlying assets; Annualized Realized Returns, which capture the actual outcomes of investors' trading activities; Behavioural Tax, defined as the difference between Buy-and-Hold Return and Realized Return (BHR-RR), which captures the trading cost; Annualized Holding Portfolio Returns, which are the daily rate of return on an investor's portfolio compared to his/her historical holding costs; Excess Buy-and-Hold Return (BHR−Index), defined as the difference between the buy-and-hold return and the return of the CSI 300 ETF index over the same period; Excess Realized Return((RR−Index), the difference between the realized return and the CSI 300 ETF index return; and Sharpe Ratio, which measures risk-adjusted performance by quantifying the excess return per unit of total risk. For all benchmark-adjusted measures—Excess Buy-and-Hold Return (BHR-Index), Excess Realized Return (RR-Index), and Sharpe Ratio—we test whether the mean values are significantly different from zero using one-sample t-tests. Statistical significance is indicated by asterisks in the tables: ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.*

Table 5 Comparison of difference analysis of investment performance

	Software C - Software A	Software B - Software A
Buy-and-Hold Return (BHR)	0.06% (0.85)	0.24% (0.63)
Excess Buy-and-Hold Return (BHR-index)	0.39% (0.18)	0.37% (0.40)
Realized Return (RR)	-55.45%*** (<0.01)	-13.78% (0.60)
Excess Realized Return (RR-index)	-55.35%*** (<0.01)	-13.84% (0.60)
Holding Portfolio Return (HPR)	-0.85% (0.34)	-3.14% (0.15)
Sharpe Ratio	-0.29 (0.20)	-0.25 (0.24)

*This table presents the statistical significance of the difference in means between the groups, tested by a standard t-test, with p-values reported accordingly. Statistical significance is indicated by asterisks in the tables: ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. The results show that only the Realized Return and Excess Realized Return have significant differences between the software C and software A. The other returns and ratios are not significantly different from each other.*