Week 12: Client Profiling, Financial Advice & AI

Dr Christian Engels
ce50@st-andrews.ac.uk

FI5614 Behavioural Finance

Department of Finance
University of St Andrews Business School

The Final Lecture – Three Reminders

Module Feedback Questionnaire (MFQ)

- If you have not yet filled it in, please do so today – the link is on Moodle
- Your feedback shapes the module for the 2026–27 cohort
- I will turn around reading-diary feedback faster if the MFQ response rate is high

Final essay – deadline

- The 2022 UK mini-budget and the gilt market crisis through a behavioural lens
- 2,000 words ($\pm$10%, excl. references), worth 30% of the module
- Deadline: 11 May 2026, 12pm (noon) via MMS

Reading diaries

- Marking has started – feedback returned in around three weeks

Readings and Topics

Required readings:

- Ackert & Deaves, Chapter 18 "Debiasing, Education and Client Management"
- "Applications of Client Behavior: A Practitioner's Perspective" in Evensky, Harold, *Financial behavior: players, services, products, and markets*

Topics covered:

1. From architecture to advice – why Week 12 follows Week 11
2. Does traditional financial advice work? The supply-side problem
3. The demand-side problem – who takes advice, and why
4. Client profiling – personality, genes, risk, and the bias blind spot
5. Debiasing – interface, mindfulness, games, and AI prompts
6. Robo-advisors – real gains, real design traps
7. Algorithm aversion, appreciation, and the arrival of LLMs

Outline

Last Week: Architecture. This Week: Intelligence.

Week 11 verdict

- Information and price incentives reach only the active 15%
- Defaults and automatic contributions bypass the decision and work 80× better per £ spent
- For repeated, forecastable decisions – save, contribute, rebalance – architecture is the whole game

But architecture is not enough

- One-off, complex decisions resist defaults: buying a house, choosing an annuity, drawing down wealth
- For these, savers turn to advisors – human and, increasingly, algorithmic
- This week: can advice close the gap that architecture cannot reach?

→ Two questions for the whole lecture: does the advice you receive make you better off, and does the identity of the advisor (person, robot, LLM) change the answer?

Does Traditional Financial Advice Work?

The Central Puzzle: Free, Good Advice Is Refused

5%

accept free, algorithm-generated, unbiased portfolio advice

The offer

8,195 active customers
of a large German broker
(Bhattacharya et al., 2012)

Who follows

Of the 385 who accept,
mean degree of following peaks at
25%

Net effect

Advisees capture
only a third
of the available Sharpe gain

→ Everything that follows sits between two failures: bad advisors, and good clients who will not listen.

Advisor Identity Matters More Than Client Need

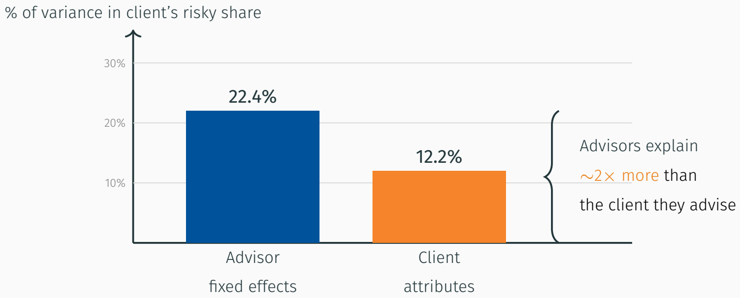

Foerster et al. (2017, *JF*): 10,000 Canadian advisors, 800,000 clients, four dealers. Advice costs 1.5pp/year vs a target-date-fund benchmark.

Linnainmaa et al. (2021, *JF*): advisors' *personal* portfolios mirror their clients' portfolios – they chase returns and hold costly active funds for themselves. The correlation is misguided belief, not conflict of interest.

→ If advisors share their clients' biases, the advisor-client dyad does not cancel errors – it *amplifies* them.

Advice Often Benefits the Bank, Not the Client

Hoechle et al. (2018), Swiss bank

- Trade-level data from a large universal bank
- Advised trades raise bank profit by ≈ CHF 252 per quarter relative to self-directed trades
- The same advised trades produce a 3.1 percentage point annual underperformance on structured products
- The product sold is the product the bank wants to sell

Anagol, Cole & Sarkar (2017), India

- Auditors visit 550 life-insurance agents in Delhi
- 81% recommend a dominated whole-life product; only 5% recommend pure term even when optimal
- Commission on whole-life is 4-6× term
- Disclosure reduces ULIP sales but agents substitute into other opaque products
- Incentives follow commissions, not customer welfare

→ The advisor is not your agent – the advisor is the bank's, and the bank has a sales target.

Misconduct Is Concentrated, Persistent, and Gendered

Statistic	Value	Source
US advisors with a misconduct record	7%	Egan, Matvos & Seru (2019)
Share of misconduct records that are repeats	27%	Egan, Matvos & Seru (2019)
Fired advisors rehired within a year	44%	Egan, Matvos & Seru (2019)
Extra undiversified advice for female auditors	+23pp	Bhattacharya, Illanes & Padi (2024)
Extra cash held by female-advised clients	+5pp	Baeckstrom et al. (2021)
Driver in every case	Male advisor → female client pairing	

- Misconduct clusters in counties with *older, higher-income, less-educated* clients – not random
- Regulators can identify repeat offenders ex ante; instead they get rehired within a year

→ The bad advice is not evenly distributed. It is concentrated in repeat offenders, in older clients, and disproportionately in women.

The "Money Doctors" Story: We Pay for Trust, Not Alpha

Gennaioli, Shleifer & Vishny (2015)

- Active-fund alphas are reliably *negative* after fees
- Yet trillions of dollars pay active-management fees
- Their model: investors are anxious about risk; managers provide confidence, not returns
- Trust is a private good and managers can charge for it
- The "money doctor" analogy: we pay a doctor to recommend medicine we could buy over the counter, because the recommendation lowers anxiety

Why this matters for regulation

- If clients hire advisors for emotional rather than financial reasons, better performance will not drive bad advisors out
- Advice markets can be persistently inefficient even in competition
- Trust-based markets also explain Stolper & Walter's (2018) finding that homophily (same sex, age, background) raises advice-following by up to 8.8pp

→ People do not buy alpha – they buy a relationship that calms them down. That is why trust-based advisors survive bad performance.

The Bright Side: Advice Closes the Numeracy Gap

von Gaudecker (2015), Netherlands

- DNB Household Survey, 1,156 households
- Self-directed investors: each one-point fall in numeracy score → 0.53pp lower risk-adjusted return
- Advised (or default-following) investors: numeracy coefficient falls to zero
- Advice is not a route to above-market returns – it is a route to stopping losses from your own mistakes

So does advice help or hurt?

- Foerster & Linnainmaa: advised clients underperform DIY benchmarks on fees
- Von Gaudecker: advised clients outperform *unadvised* clients of the same low numeracy
- Both true: advice is bad relative to optimal, good relative to untreated
- For the bottom quartile of financial literacy, advice remains welfare-improving

→ Advice is a levelling technology. It hurts the sophisticated relative to DIY and helps the naive relative to doing nothing.

Fiduciary Duty Works – When It Forces Better Advice

Bhattacharya, Illanes & Padi (2025)

- State-level variation in broker fiduciary standards in the US
- Under fiduciary duty, risk-adjusted returns rise by 0.25pp per year
- This is 9% of the mean return in their sample
- The gain comes not from cheaper products – it comes from *better-suited* products being recommended

The regulatory lesson

- Disclosure alone (Anagol, MiFID II) is not enough – customers cannot parse it
- Commission bans shift advisors' mix but not their behaviour
- A standard that makes advisor liable for *unsuitable* advice changes what gets recommended
- UK RDR (2012) and EU MiFID II (2018) move in this direction; US DOL rule repeatedly stalls

→ Supply-side reform is worth pursuing – but only the variety that binds what the advisor is allowed to recommend.

Know Your Client

Profiling Has Moved Beyond Demographics

Three evidence-based lenses

Personality

Big Five + harm avoidance

- Conlin et al. (2015)
- Tauni et al. (2017)
- Oehler et al. (2022)

Predicts participation, trading volume, advice response.

Risk attitude

FinaMetrica-style tests

- Hartnett et al. (2019)
- 233k completed tests
- The industry default

But noisy for 20% of respondents – a full quintile of mis-classification.

Biases

Kahneman–Riepe checklist

- West et al. (2012)
- Barnea, Cronqvist & Siegel (2010)
- Cronqvist & Siegel (2014)

Biases are 25–46% heritable and covered by a blind spot.

The old Vanguard typology (planners vs avoiders, Marconi & Utkus 2002) is still used in practice – MBTI too – but neither has predictive validity for investing.

→ The advisor's job is not to label the client – it is to pick an action whose success depends on how the client will actually behave.

Personality Predicts Both Participation and Response to Advice

Conlin et al. (2015), Finland

- 2,197 Finns, linked tax records and Cloninger's Temperament & Character Inventory
- Harm avoidance reduces stock-market participation by 1.7pp – over 10% of the baseline participation rate
- Among high-SES individuals, extraversion raises participation by 8.8pp
- Standard demographics (age, income, wealth, education) do not absorb these effects

Tauni et al. (2017), Pakistan

- 541 Pakistani equity investors, Big Five inventory
- Baseline: one extra SD of neuroticism raises trading by 0.93 trades per month
- With a financial advisor: the effect jumps to 2.21 trades per month
- Advice amplifies, rather than dampens, trait-driven trading
- Word-of-mouth channels act in the opposite direction

→ The same piece of advice yields different behaviour in different personalities – client profiling is not cosmetic, it is outcome-determining.

Nature and Nurture – Biases Are Partly Heritable

Twin studies, Sweden

Barnea, Cronqvist & Siegel (2010) exploit identical vs fraternal twins to decompose variation in household portfolio choice:

- 29% of stock-market participation is heritable
- 28% of the equity share in financial wealth
- 37% of portfolio volatility
- Shared family environment is close to zero for adults – peer and life experience dominate

Even the biases are heritable

Cronqvist & Siegel (2014), same data, decompose specific behavioural biases:

- Under-diversification: 25% genetic
- Home bias: 46% genetic
- Disposition effect: 29% genetic
- Finance-sector work experience moderates the genetic effect; general education does not

→ A third of the biases you will inherit cannot be educated away – but targeted experience can dampen them. Client profiling has to live with this.

Risk Tolerance Questionnaires Are Noisier Than We Admit

Hartnett et al. (2019)

- 233,000 FinaMetrica risk-tolerance tests over 12 years
- For 20% of respondents, within-test inter-item response variability is very high – the same test gives unstable scores
- A one-SD rise in IRV predicts a 3.5 point change in absolute FRT score – the width of a risk-tolerance quintile
- These are the clients an advisor would least want to misclassify

West, Meserve & Stanovich (2012): the bias blind spot

- Harvard undergraduates rate themselves as *less biased than the average person* on 14 classic biases
- Cognitive sophistication (SAT, Need-for-Cognition) is *positively* correlated with the blind spot
- Smart people know more about biases in the abstract – and judge themselves more immune in consequence
- The "I am aware, therefore I am debiased" fallacy

→ The client who most needs profiling is least likely to admit it. Numerical profiling is noisy. Self-profiling is biased upward.

Debiasing – What Works, What Does Not

Four Routes to Debiasing – and Their Ceilings

Interface design

Hide the anchor, change the default

Frydman & Rangel (2014)

State of mind

Mindfulness, cognitive distance

Hafenbrack et al. (2014)

Training

Serious games, repeated practice

Sellier et al. (2019)

External systems

Rules, checklists, algorithms

Robos, prompt-debiased LLMs

- Wilson, Centerbar & Brekke (2002): a bias is only removed if the person has awareness, motivation, magnitude knowledge, and ability – most interventions fail at step 1
- Education (Week 11) touches *knowledge* and decays; the four routes above touch *behaviour*

→ If you cannot change the investor, change the environment in which the investor decides.

Interface Design: Hide the Anchor, Kill the Disposition Effect

Frydman & Rangel (2014), lab

Participants trade three virtual stocks under two interfaces:

- **High salience:** purchase price shown alongside current price
- **Low salience:** only current price shown
- Same information, same task, same stakes

High-salience participants show a textbook 6.8% disposition effect.
Low-salience participants show -9.0% – the bias reverses.

Hafenbrack et al. (2014), meditation

- 15-minute recorded mindfulness induction vs a mind-wandering control
- Classic sunk-cost vignette: do you see the movie you hate?
- Control condition resistance rate: 44%
- Mindfulness condition: 78%
- Mechanism: reduced temporal focus → reduced negative affect → cleaner decision

→ Remove the reference point from the screen; the bias dissolves. Remove the reference point from the mind; the bias weakens almost as much.

Training That Actually Transfers – and a Warning About AI

Sellier, Scopelliti & Morewedge (2019)

- 290 consultants in a one-off 80-minute serious game targeting confirmation bias
- Biased choice on a transfer task falls from 72% → 59%
- Training effect still visible 43–52 days later
- Gamified, personalised feedback beats didactic lectures
- First evidence of durable, field-transferable debiasing training

Winder, Hildebrand & Hartmann (2025)

- Ask ChatGPT, Copilot and Gemini for a $10,000 portfolio under 27 persona conditions
- Average LLM portfolio: 93% US equities, heavy top-3 concentration, 51% actively managed
- Benchmark (VT ETF) holds 59% US, broadly diversified
- *Broad* debiasing prompts ("avoid lack of diversification, cluster risks, active management") work; narrow prompts do not

→ Debiasing the *human* takes 80 minutes of training. Debiasing the *LLM* takes one prompt – if you know which one to write.

Enter the Robots

Robo-Advisors Deliver Material Welfare Gains

Rossi & Utkus (2024), *JFE*

- 11,000 self-directed Vanguard investors adopt the Personal Advisor Services hybrid robo
- Sharpe ratio rises from 0.33 → 0.38 (+16%)
- Within-investor log-Sharpe gain: +10.9%
- Idiosyncratic volatility falls 17%
- Welfare gain: 0.77% of lifetime consumption under CARA utility
- 70% of the gain comes from diversification, 30% from the age/wealth glide path

Reher & Sokolinski (2024), Wealthfront

- Wealthfront cuts its account minimum from $5,000 to $500 in July 2015
- Middle-class (Q2–Q3) participation rises +107% (+16pp)
- No change in Q1 (incomplete democratisation); no decline in Q4–Q5
- Middle-class Sharpe jumps from 0.45 to 0.75
- Welfare gain: 0.77% of lifetime consumption – and 1.68% for the 56–65-year-olds

→ Access and diversification are the whole story. Robo-advisors do for portfolios what auto-enrolment did for pensions: they bypass the client's decision and give them a sensible default.

Who Gains From a Robo? Not Everyone Equally

D'Acunto, Prabhala & Rossi (2019), India

- Large Indian brokerage rolls out a behavioural robo to a subset of clients
- Clients who held 1–2 stocks beforehand: portfolio variance falls by half; number of holdings rises to 3–4
- Disposition effect falls 30%; trend-chasing falls 1.2pp; rank effect falls 26%
- Clients who already held 11+ stocks: volatility rises after adoption – robo destroys their prior diversification

Capponi, Oláfsson & Zariphopoulou (2022)

- Life-cycle model of a myopic-loss-averse client interacting with a robo
- Optimal rebalancing frequency: 2–6 months
- Continuous rebalancing reduces welfare – it triggers behavioural reviews of every drawdown
- Wealthfront's practice of limiting client-initiated changes to one per month is consistent with theory
- Less is more – even for rational reasons

→ The robo is not a universal upgrade. It helps the underdiversified, and leaves (or hurts) those who already knew what they were doing.

The Humanisation Trap – Trust Can Cost You

Hildebrand & Bergner (2021)

- Same robo-advisor, two user interfaces: conversational (chat-style, anthropomorphised) vs static (form-style)
- Conversational design raises affective trust by +1.10 (7-point scale) and benevolence by +0.63
- But acceptance of an unsuitable recommended portfolio rises from 40% to 73%
- Disclaimers do not undo the effect

Hodge, Mennecke & Sprong (2021)

- Labelling an algorithm with a human name helps reliance on *simple* tasks (+9pp) but hurts it on *complex* tasks (—15pp)
- For human advisors, credentialing helps in both cases (+19pp)
- Robot anthropomorphism is a double-edged sword: it can attract and mislead in equal measure

→ The same design choices that build trust in a fintech app can lead clients to accept portfolios they should refuse. UX is policy.

AI & LLMs – The Next Advisor

Algorithm Appreciation, Not Aversion

Weight-of-advice on *identical* numerical advice

Logg, Minson & Moore (2019)

0.45	**0.30**
Algorithm source	**Human** source
"An algorithm combined previous ratings..."	"Another person rated this photo..."

- Laypeople *appreciate* algorithms; JDM researchers predicted the opposite
- **Dietvorst et al. (2015)**: after the algorithm makes *any* visible mistake, model-use drops to 26% – the same human errors are forgiven

→ The public is more pro-algorithm than academics assume. The challenge is not adoption – it is retention after the first visible failure.

When Is Algorithmic Advice Trusted? Task Type Matters

Castelo, Bos & Lehmann (2019)

- Seven studies across domains
- Trust in algorithm is high for objective tasks: route planning (82%), tax calculation, radiology
- Trust is low for subjective tasks: choosing a joke (30%), relationship matching, hiring
- Financial portfolio construction sits near the objective end – LLMs and robos should inherit trust
- Re-framing a subjective task in quantitative terms restores trust

Implication for advice design

- Portray portfolio choice as diversification and tax efficiency (objective) rather than life goals (subjective)
- This is why the robo-advisor onboarding experience emphasises fees, correlations, and historical returns – not dreams and ambitions
- The Hildebrand trap then reappears: re-framing increases uptake, but does it increase *good* uptake?

→ Algorithms are trusted for the parts of finance that look like engineering. The framing of the advice is half the product.

GPT Is Already Financially Literate

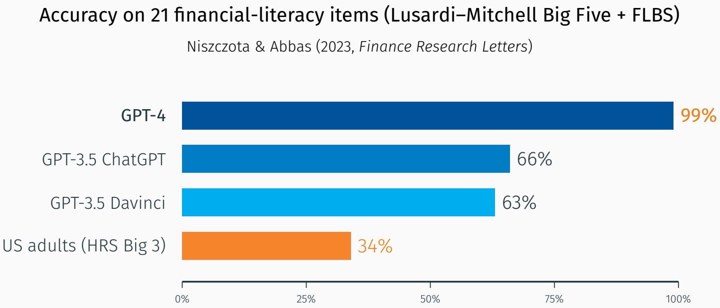

- **And people listen to it:** on a $33k savings problem, mean weight-of-advice on GPT = 0.65 – more than double Logg's 0.30 for humans
- Low-literacy participants' WOA on GPT rises to 0.74 – they lean *more* on the machine, not less

→ GPT-4 is literate enough to pass your exam. The next policy question is not whether it can advise, but whether it should, and on what regulatory basis.

But LLM Portfolios Echo the Biases of the Humans Who Trained Them

Winder et al. (2025) again

- Ask for a 30-year-old's $10k portfolio, across 3 LLMs
- LLM average: 93% US equities, vs 59% in the benchmark
- Top-3 concentration: +28pp above benchmark
- Active management share: +51pp
- Explicit language style: highly assertive, low-disclaimer
- The LLM inherits the US-home-bias and active-picking habits of the finance writing it was trained on

What to do about it

- A broad debiasing prompt ("avoid cluster risks and active management") materially improves the portfolio
- A narrow prompt ("no management fees") does not
- An ESG prompt reduces US concentration but introduces sector concentration ("green" cluster risk)
- Prompting is now a regulatory surface – who gets to set the default?

→ The machine has absorbed the same biases as the books; debiasing the client is being replaced by debiasing the prompt.

Synthesis

Five Themes, One Spine

Theme	Key mechanism	Signature finding
Advice supply	Misguided beliefs, commissions, misconduct	7% of US advisors have a record; 22% of risky-share variance is advisor-driven
Advice demand	Trust, homophily, inattention	5% take up free unbiased advice; advisees follow $<$25% of it
Client profiling	Personality, genes, biases	25–46% of biases are heritable; neuroticism $\times$ advice doubles trading
Robo-advisors	Diversification, glide path, access	Sharpe +16%; welfare +0.77% of lifetime consumption
AI & LLMs	Appreciation, task framing, prompt design	GPT-4 scores 99% on Big Five; weight-of-advice 0.65; but 93% US tilt

$\longrightarrow$ Each row solves a problem the previous row created, and introduces a new one. There is no endpoint – every generation of advice technology must be debiased afresh.

The Lecture in One Sentence – and One for You

Policy takeaway

Advice markets fail on both supply and demand;
algorithms close the supply gap –
the demand gap is now the frontier.

Personal takeaway

- Distrust advice that sells *you* a product – demand advice that sells *itself*
- If you use a robo or LLM advisor, write a broad debiasing prompt and check the US-bias
- Your own bias blind spot is largest where you feel most competent – put external rules between you and decisions you care most about
- Rebalance quarterly, not constantly – the Capponi ceiling applies to you too

→ The architecture lesson from Week 11 generalises: when you cannot trust the decision-maker – whether the decision-maker is an advisor, an algorithm, or yourself – design the system so the decision does not depend on trust.

Thanks for the Semester

The module in one thread

- **Weeks 1–5:** how classical finance sees markets, and where the cracks first appeared
- **Weeks 6–8:** the anomalies, the biases, and why prices drift away from fundamentals
- **Weeks 9–10:** managers, firms, and the macro scale of behavioural finance
- **Weeks 11–12:** the household – and what we do about it

Good luck with the essay, and the rest of your studies.

References i

- Ackert, L. & Deaves, R. (2010). *Behavioral Finance: Psychology, Decision-Making, and Markets*, Chapter 18.
- Anagol, S., Cole, S. & Sarkar, S. (2017). Understanding the advice of commissions-motivated agents: evidence from the Indian life insurance market. *Review of Economics and Statistics*, 99(1), 1–15.
- Baeckström, Y., Silvester, J. & Pownall, R. A. J. (2021). Millionaire investors: financial advisors, attribution theory and gender differences. *European Journal of Finance*, 27(11), 1–21.
- Barnea, A., Cronqvist, H. & Siegel, S. (2010). Nature or nurture: what determines investor behavior? *Journal of Financial Economics*, 98(3), 583–604.
- Bhattacharya, U., Hackethal, A., Kaesler, S., Loos, B. & Meyer, S. (2012). Is unbiased financial advice to retail investors sufficient? *Review of Financial Studies*, 25(4), 975–1032.
- Bhattacharya, V., Illanes, G. & Padi, M. (2024). Fiduciary duty and the market for financial advice. *Journal of Finance*, forthcoming.
- Bhattacharya, V., Illanes, G. & Padi, M. (2024). Gender and the market for financial advice. *Journal of Finance*, forthcoming.
- Capponi, A., Oláfsson, S. & Zariphopoulou, T. (2022). Personalized robo-advising: enhancing investment through client interaction. *Management Science*, 68(4), 2485–2512.
- Castelo, N., Bos, M. W. & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research*, 56(5), 809–825.

- Clifford, C. P., Ellis, J. A. & Gerken, W. C. (2023). Racial concordance in the market for financial advice. *Review of Corporate Finance Studies*, 12(4), 758–805.
- Conlin, A., Kyröläinen, P., Kaakinen, M., Järvelin, M.-R., Perttunen, J. & Svento, R. (2015). Personality traits and stock market participation. *Journal of Empirical Finance*, 33, 34–50.
- Cronqvist, H. & Siegel, S. (2014). The genetics of investment biases. *Journal of Financial Economics*, 113(2), 215–234.
- D'Acunto, F., Prabhala, N. & Rossi, A. G. (2019). The promises and pitfalls of robo-advising. *Review of Financial Studies*, 32(5), 1983–2020.
- Deaves, R. (2024). *Household Finance*, Chapter 11. Oxford University Press.
- Dietvorst, B. J., Simmons, J. P. & Massey, C. (2015). Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126.
- Egan, M., Matvos, G. & Seru, A. (2019). The market for financial adviser misconduct. *Journal of Political Economy*, 127(1), 233–295.
- Egan, M., Matvos, G. & Seru, A. (2022). When Harry fired Sally: the double standard in punishing misconduct. *Journal of Political Economy*, 130(5), 1184–1248.
- Foerster, S., Linnainmaa, J. T., Melzer, B. T. & Previtero, A. (2017). Retail financial advice: does one size fit all? *Journal of Finance*, 72(4), 1441–1482.

- 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.
- Gennaioli, N., Shleifer, A. & Vishny, R. (2015). Money doctors. *Journal of Finance*, 70(1), 91–114.
- Hafenbrack, A. C., Kinias, Z. & Barsade, S. G. (2014). Debiasing the mind through meditation: mindfulness and the sunk-cost bias. *Psychological Science*, 25(2), 369–376.
- Hartnett, N., Gerrans, P., Faff, R. & Smith, T. (2019). Inter-item response variability in financial risk-tolerance assessment. *Financial Analysts Journal*, 75(2), 67–85.
- Hildebrand, C. & Bergner, A. (2021). Conversational robo advisors as surrogates of trust: onboarding experience, firm perception, and consumer financial decision making. *Journal of the Academy of Marketing Science*, 49, 659–676.
- Hodge, F. D., Mennecke, B. E. & Sprong, N. (2021). The effect of humanizing robo-advisors on investor judgments. *Contemporary Accounting Research*, 38(1), 770–792.
- Hoechle, D., Ruenzi, S., Schaub, N. & Schmid, M. (2018). Financial advice and bank profits. *Review of Financial Studies*, 31(11), 4447–4492.
- Linnainmaa, J. T., Melzer, B. T. & Previtero, A. (2021). The misguided beliefs of financial advisors. *Journal of Finance*, 76(2), 587–621.
- Logg, J. M., Minson, J. A. & Moore, D. A. (2019). Algorithm appreciation: people prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103.

- Niszczota, P. & Abbas, S. (2023). GPT has become financially literate. *Finance Research Letters*, 58, 104333.
- Oehler, A., Horn, M. & Wendt, S. (2022). Investor characteristics and their impact on the decision to use a robo-advisor. *Journal of Financial Services Research*, 62, 91–125.
- Reher, M. & Sokolinski, S. (2024). Robo-advisors and access to wealth management. *Journal of Financial Economics*, 155, 103829.
- Rossi, A. G. & Utkus, S. P. (2024). The diversification and welfare effects of robo-advising. *Journal of Financial Economics*, 157, 103869.
- Sellier, A.-L., Scopelliti, I. & Morewedge, C. K. (2019). Debiasing training improves decision making in the field. *Psychological Science*, 30(9), 1371–1379.
- Stolper, O. & Walter, A. (2019). Birds of a feather: the homophily effect in financial advice. *Review of Financial Studies*, 32(11), 4255–4291.
- Tauni, M. Z., Rao, Z.-U.-R., Fang, H.-X. & Gao, M. (2017). Does investor personality moderate the relationship between information sources and trading behavior? *International Journal of Bank Marketing*, 35(5), 857–877.
- von Gaudecker, H.-M. (2015). How does household portfolio diversification vary with financial literacy and financial advice? *Journal of Finance*, 70(2), 489–507.
- West, R. F., Meserve, R. J. & Stanovich, K. E. (2012). Cognitive sophistication does not attenuate the bias blind spot. *Journal of Personality and Social Psychology*, 103(3), 506–519.

- Winder, P., Hildebrand, C. & Hartmann, J. (2025). Biased echoes: large language models reinforce investment biases and increase portfolio risks of private investors. *PLoS ONE*, 20(6), e0325459.