Personalized Robo-Advising: Enhancing Investment through Client Interaction

Agostino Capponi

Department of Industrial Engineering and Operations Research

Columbia University

ac3827@columbia.edu

(joint work with S. Ólafsson and T. Zariphopoulou)

Mathematical Finance Colloquium

Department of Mathematics

University of Southern California, October 12, 2020

Robo-Advisors

Automated investment platforms providing algorithm-driven investment advice with limited human supervision

- First robo-advisors launched in the wake of the 2008 financial crisis (Betterment, Wealthfront, Personal Capital, . . .)
- Current offerings include:
  - Affordable portfolio management (fees, account size, etc.)
  - Full Automation (portfolio construction and rebalancing)
  - Tax-loss harvesting
- Becoming both day-to-day and long-term money managers

Robo-Advisors

- Currently manage around \$750 billion in the United States (2019)
  - Less than 1.5% of total investable assets

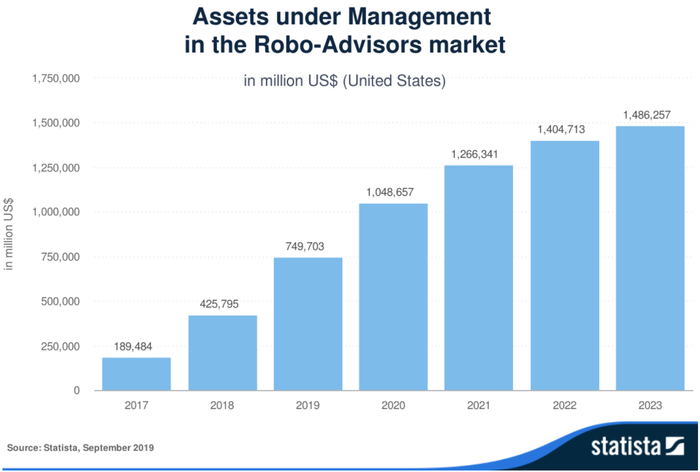

Existing Literature

- D'Acunto et al. [2018] and Rossi and Utkus [2019a,b]
  - Robo-advising vs. traditional financial advising
- Reher and Sokolinski [2020]
  - Robo-advising and financial inclusion
- Beketov et al. [2018] analyze over 200 robo-advisors globally
  - Risk profiling based on online questionnaires
  - Mean-variance portfolio optimization
- Classification of robo-advising systems (D'Acunto and Rossi [2020])
  - Portfolio personalization
  - Client involvement
  - Level of human-advising

Model of Robo-Advising

- We propose the first quantitative model of robo-advising:
  - Investment process accounts for **dynamic** risk preferences and **repeated** interaction between client and robo-advisor
- Our framework is consistent with that of the most prominent stand-alone robo-advising firms
  - High portfolio personalization
  - Low and indirect client involvement
  - Limited or no human-advising

Human-Machine Interaction

Client:

- Dynamic risk aversion process: $(\gamma_n^C)_{n\geq 0}$
- Provides information to the robo-advisor at interaction times

Robo-Advisor:

- Constructs a *model* of the client's risk aversion: $(\gamma_n^R)_{n\geq 0}$
- Differs from $(\gamma_n^C)_{n\geq 0}$ due to imperfect human-machine interaction
  - Changes to client's demographics only observed at interaction times
  - Information communicated by client affected by **behavioral** biases
- Designs an optimal investment strategy

Determinants of Client's Risk Preferences

Client's risk aversion $(\gamma_n^C)_{n \geq 0}$ changes because of:

- Passage of time
  - 1st generation robo-advisors $\approx$ TDFs
- Shocks to demographics
  - Barsky et al. [1997], Guiso and Paiella [2008], ...
- Market returns and economic conditions
  - Fama and French [1989], Cohn et al. [2015], Bucciol and Miniaci [2018], Guiso et al. [2018]

Research Questions

(1) Does stochastic variation in the client's risk aversion produce an intertemporal hedging demand for the risky asset?

(2) Does frequent interaction allow the robo-advisor to implement a strategy closely matching the client's risk profile?

(3) Should the robo-advisor always cater to the client's wishes?

Model of Robo-Advising

Model components:

- Market model for available investment securities
- **Dynamic** model for client's risk aversion
- Mechanism to decide the human-machine interaction schedule
- Optimal adaptive investment criterion

Market Dynamics

- Risky asset $(S_n)_{n\geq 0}$ and a risk-free asset $(B_n)_{n\geq 0}$,

$$S_{n+1} = (1 + Z_{n+1}(Y_n))S_n$$
$$B_{n+1} = (1 + r(Y_n))B_n$$

- Price dynamics modulated by a Markov regime switching model of economic conditions $(Y_n)_{n\geq 0}$ (Hamilton [1989])
- Given $Y_n = y$, the risk-free rate is $r(y) \geq 0$, and the risky asset's return has mean $\mu(y) > r(y)$, and variance $\sigma^2(y) > 0$
- Probability space $(\Omega, \mathcal{F}, \mathbb{P})$ also supports a sequence $(\epsilon_n)_{n\geq 1}$ of random variables, independent of $(Y_n)_{n\geq 0}$ and $(Z_n)_{n\geq 1}$
- Filtration $(\mathcal{F}_n)_{n\geq 0}$ defined by $\mathcal{F}_n = \sigma(Y_{(n)}, Z_{(n)}, \epsilon_{(n)})$

Market Dynamics

- Self-financing investment strategy $\pi = (\pi_n)_{n \geq 0}$
  - π_n is the amount of wealth invested in the risky asset at time n
- Under π, the client's wealth process $(X_n^\pi)_{n \geq 0}$ satisfies

$$X_{n+1}^\pi = R_{n+1} X_t^\pi + \widetilde{Z}_{n+1} \pi_n$$

where $R_{n+1} := 1 + r(Y_n)$ and $\widetilde{Z}_{n+1} := Z_{n+1}(Y_n) - r(Y_n)$

Client's Risk Aversion Process

- Captures stylized features of retail investors' risk profiles:

$$\gamma_n^C := \gamma_n^C(Y_{(n)}, Z_{(n)}, \epsilon_{(n)}) = e^{\eta_n} \gamma_n^{id} \gamma_n^Y$$

- The first component $(e^{\eta_n})_{n \geq 0}$ is of the form

$$e^{\eta_n} = e^{-\alpha(T-n)}$$

and captures age-related increase in risk aversion

- The second component $(\gamma_n^{id})_{n \geq 0}$ is of the form

$$\gamma_n^{id} = \gamma_{n-1}^{id} e^{\epsilon_n}, \qquad \epsilon_n = \begin{cases} \mathcal{N}(0, \sigma_\epsilon^2), & \text{w.p. } p_\epsilon \\ 0, & \text{w.p. } 1 - p_\epsilon \end{cases}$$

and captures idiosyncratic shocks to the client's risk aversion

Client's Risk Aversion Process

- The third component is a state-dependent coefficient $\gamma_n^Y = \bar{\gamma}(Y_n)$, which is increasing in the market Sharpe ratio:

$$\frac{\mu(y_i) - r(y_i)}{\sigma(y_i)} \geq \frac{\mu(y_j) - r(y_j)}{\sigma(y_j)} \implies \bar{\gamma}(y_i) \geq \bar{\gamma}(y_j)$$

- Risk aversion and market Sharpe ratio (λ) are **countercyclical**
  - Higher at business cycle troughs than at peaks
  - Lettau and Ludvigson [2010], Campbell and Cochrane [1999]:

$$Y \downarrow \implies \gamma^C \uparrow \implies \lambda \uparrow \implies S \downarrow$$

- "Unadvised" client is inclined to reduce market exposure when the market Sharpe ratio is high

Human-Machine Interaction

- The interaction schedule $(T_k)_{k\geq 0}$ is an increasing sequence of stopping times with respect to the filtration $(\mathcal{F}_n)_{n\geq 0}$
- Interaction can be triggered by any combination of client-specific events, economic state changes, and market events
- Define $(\tau_n)_{n\geq 0}$ where $\tau_n := \sup\{T_k : T_k \leq n\}$ is the most recent interaction time occurring prior to or at time n
- **Deterministic schedule:** The sequence $(T_k)_{k\geq 0}$ is given by

$$T_k = k\phi, \quad k \geq 0$$

where $\phi \geq 1$ is the time between consecutive interactions

Human-Machine Interaction

- At an interaction time n, the risk aversion value communicated by the client is

$$\xi_n = \gamma_n^C \gamma_n^Z$$

where γ_n^C is the client's risk aversion and, for $\beta \geq 0$,

$$\gamma_n^Z = e^{-\beta \left(\frac{1}{\phi} \sum_{k=n-\phi}^{n-1} (Z_{k+1} - \mu_{k+1}) \right)}$$

captures client's behavioral biases (e.g., trend-chasing)

- The **bias** γ_n^Z is based on recent stock market performance:
  - Market outperforming the client's expectations: $\gamma_n^Z < 1$
  - Market underperforming the client's expectations: $\gamma_n^Z > 1$

Human-Machine Interaction

- Magnitude of the bias γ_n^Z is increasing in the sensitivity β but decreasing in the time between interactions ϕ:

$$\mathbb{E}_n[\gamma_n^Z] \approx e^{\frac{1}{2}\frac{\beta^2\sigma^2(Y_n)}{\phi}}$$

- Sensitivity β assumed known to the robo-advisor at the outset
  - Behavioral biases related to financial literacy, experience, and cognitive abilities (e.g., Oechssler et al. [1997], Seasholes and Feng [2005])

Robo-Advisor's Model of Client

- Robo-advisor's filtration $(\mathcal{F}_n^R)_{n\geq 0}$ is generated by

$$D_n := (Y_{(n)}, Z_{(n)}, \tau_{(n)}, \xi_{(n)})$$

which consists of market information $(Y_{(n)}, Z_{(n)})$ and **information from client-interaction** $(\tau_{(n)}, \xi_{(n)})$

- Robo-advisor's model of the client's risk aversion $(\gamma_n^R)_{n\geq 0}$ is a process adapted to the robo-advisor filtration:

$$\gamma_n^R := \gamma_n^R(D_n)$$

- The filtration $(\mathcal{F}_n^R)_{n\geq 0}$ grows with the frequency of interaction

Robo-Advisor's Model of Client

- The robo-advisor's model $(\gamma_n^R)_{n\geq 0}$ is given by

$$\gamma_n^R = \mathbb{E}_n[\gamma_n^C]\gamma_{\tau_n}^Z = \xi_{\tau_n}\, e^{\eta_n - \eta_{\tau_n}} \frac{\gamma_n^Y}{\gamma_{\tau_n}^Y}$$

  where γ_n^C is the client's risk aversion and $\gamma_{\tau_n}^Z$ is the bias realized at the previous time of interaction

- Updated in real time based on the passage of time, realized market returns, and changes in economic conditions

Adaptive Mean-Variance Criterion

- **Partial Equilibrium framework:** For a fixed horizon $T \geq 1$, the robo-advisor maximizes a **mean-variance** objective

$$J_n(x, d, \pi) := \mathbb{E}_{n,x,d}\left[\frac{X_T^\pi - X_n}{X_n}\right] - \frac{\gamma_n^R}{2} Var_{n,x,d}\left[\frac{X_T^\pi - X_n}{X_n}\right]$$

  where π is a self-financing strategy
- Initial condition fixes the robo-advisor's information set:
  - $\mathbb{P}_{n,x,d}(\cdot) := \mathbb{P}(\cdot | X_n = x, D_n = d)$
- Sequence of objective functions $(J_n)_{0 \leq n < T}$
  - Robo-advisor's model of client's risk preferences $\gamma_n^R > 0$ adapts to market, economic, and client-communicated information
- **Time-inconsistent stochastic control problem**

Adaptive Mean-Variance Criterion

- Multi-period mean-variance optimization is time-inconsistent
  - Tower property of conditional expectations does not apply to $(\mathbb{E}_{n,x,d}[X_T])^2 \implies$ Bellman optimality principle does not hold
  - Stochastic risk-return coefficient is an additional source of time-inconsistency
- Multitude of future decision-making "selves" that may not act in the best interest of previous "selves"
  - If π^* is the control law that maximizes J_n, then π^* restricted to $\{n+1, n+2, \ldots, T\}$ is suboptimal for the objective function J_{n+1}

Adaptive Mean-Variance Criterion

- Myopic investment
  - At each time n, apply the control π_n^* where π^* is the control law that maximizes J_n (reoptimize at each time)
  - Ignores time-inconsistency
- Pre-committed investment
  - At time $n = 0$, find the control law π^* that maximizes J_0 and commit to this strategy throughout the investment horizon
  - Ties the hands of future "selves"
- **Equilibrium control**
  - At each time n, apply the control π_n^* that maximizes J_n, *given* that future "selves" will act in their own best interest
  - Non-cooperative game with one player for each time n
  - Subgame-perfect equilibrium

Adaptive Mean-Variance Criterion

- Look for an optimal control π^* which is time-consistent:

$$\sup_{\pi \in A^*_{n+1}} J_n(\pi) = J_n(\pi^*), \quad \text{for all} \;\; 0 \leq n < T$$

  where $A^*_{n+1} = \{\pi : \pi_{n+1:T} = \pi^*_{n+1:T}\}$ is the set of control laws that coincide with π^* *after* time n

- Any candidate optimal control $(\pi_n)_{n\geq 0}$ is of the form

$$\pi_n = \pi_n(x, d) \in \mathcal{F}^R_n$$

  such that $\mathbb{E}_0\Big[\sum_{n=0}^{T-1} \pi_n^2\Big] < \infty$

Optimal Investment

Theorem

The optimal proportion of wealth invested in the risky asset at time n is

$$\pi_n^* = \frac{1}{\gamma_n^R} \frac{\mathbb{E}_n[\widetilde{Z}_{n+1} r_{n+1}^{\pi^*}]}{Var_n[\widetilde{Z}_{n+1} r_{n+1}^{\pi^*}]} - R_{n+1} \frac{Cov_n(r_{n+1}^{\pi^*}, Z_{n+1} r_{n+1}^{\pi^*})}{Var_n[\widetilde{Z}_{n+1} r_{n+1}^{\pi^*}]}$$

where $\widetilde{Z}_{n+1}$ is the excess market return at $n+1$, and

$$r_{n+1}^{\pi^*} = \frac{X_T^{\pi^*}}{X_{n+1}}$$

is the value of one dollar invested in the optimal portfolio between time $n+1$ and the terminal date T.

The allocation π_n^* depends on both the **current** *risk-return tradeoff γ_n^R its* **future dynamics** *through the conditional expectations.*

Structure of Optimal Strategy

$$\pi_n^* = \frac{1}{\gamma_n^R} \frac{\mathbb{E}_n[\widetilde{Z}_{n+1} r_{n+1}^{\pi^*}]}{Var_n[\widetilde{Z}_{n+1} r_{n+1}^{\pi^*}]} - R_{n+1} \frac{Cov_n(r_{n+1}^{\pi^*}, Z_{n+1} r_{n+1}^{\pi^*})}{Var_n[\widetilde{Z}_{n+1} r_{n+1}^{\pi^*}]}$$

- $\frac{1}{\gamma_n^R} \frac{\mathbb{E}_n[Z_{n+1} r_{n+1}^{\pi^*}]}{Var_n[Z_{n+1} r_{n+1}^{\pi^*}]}$: standard single-period Markowitz strategy, but also accounting for the future portfolio return $r_{n+1}^{\pi^*}$
- $\frac{Cov_n(r_{n+1}^{\pi^*}, Z_{n+1} r_{n+1}^{\pi^*})}{Var_n[Z_{n+1} r_{n+1}^{\pi^*}]}$: **intertemporal** hedging component
  - Incorporates the effect of market returns and economic conditions on the client's risk aversion

$$\widetilde{Z}_{n+1} \uparrow \implies \gamma_{n+1}^R \downarrow \implies \pi_{n+1}^* \uparrow \implies r_{n+1}^{\pi^*} \uparrow \implies Cov_n(\dots) > 0$$

Portfolio Personalization

- Relative difference between the robo-advisor's model and the client's risk aversion process:

$$\mathcal{R}(\phi, \beta) := \mathbb{E}\left[\frac{1}{T}\sum_{n=0}^{T-1}\left|\frac{\gamma_n^R - \gamma_n^C}{\gamma_n^C}\right|\right]$$

  - $\phi \geq 1$: time between consecutive interaction times
  - $\beta \geq 0$: strength of client's behavioral bias

- Robo-advisor faces a tradeoff between **information acquisition rate** and **accuracy of acquired information**

Proposition

1. *There exists a unique value of ϕ that minimizes $\mathcal{R}(\phi, \beta)$*
2. *Optimal value of ϕ is increasing in β*

Optimal Interaction Frequency

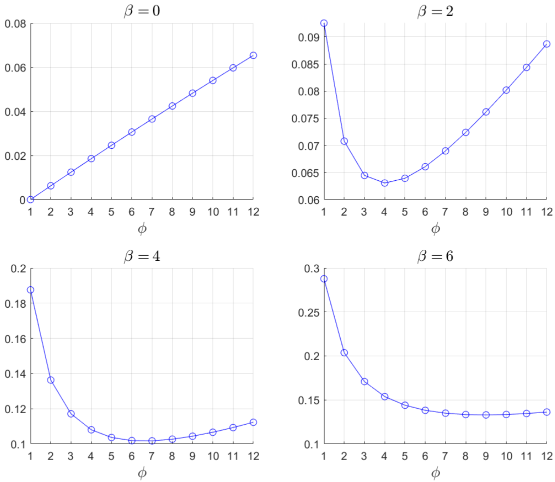

- Magnitude of behavioral bias increases with the interaction frequency
  - Myopic loss aversion (Benartzi and Thaler [1995])

Economic Transitions and Sharpe Ratios

- Economy with two states $\mathcal{Y} = \{1, 2\}$ corresponding to economic expansions and contractions, and such that

$$\frac{\mu(1) - r(1)}{\sigma(1)} < \frac{\mu(2) - r(2)}{\sigma(2)}$$

- Consider an investment strategy of the form

$$\pi_n^*(y) = \begin{cases} \bar{\pi}, & y = 1 \\ \bar{\pi}(1 + \delta), & y = 2 \end{cases}$$

- $\delta < 0$: "unadvised" client shifting wealth away from risky asset when the market Sharpe ratio is high
- $\delta > 0$: "robo-advised" client doing the opposite

Catering or Going Against Client's Wishes?

Proposition

The optimal portfolio's Sharpe ratio $s^{\pi^*}(\delta)$ is increasing and concave around $\delta = 0$:

$$\frac{\partial s^{\pi^*}(\delta)}{\partial \delta} > 0, \qquad \frac{\partial^2 s^{\pi^*}(\delta)}{\partial \delta^2} < 0$$

- Compared to a buy-and-hold strategy:
  - Sharpe ratio lower if client is left unassisted ($\delta < 0$)
  - Sharpe ratio higher with robo-advising ($\delta > 0$)
- The drop in Sharpe ratio **is greater** when tilting away from the risky asset compared to the gain when tilting towards the risky asset

Economic Transitions and Portfolio Returns

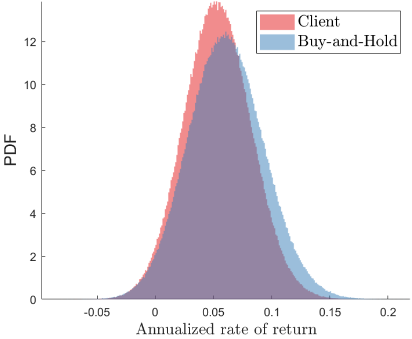

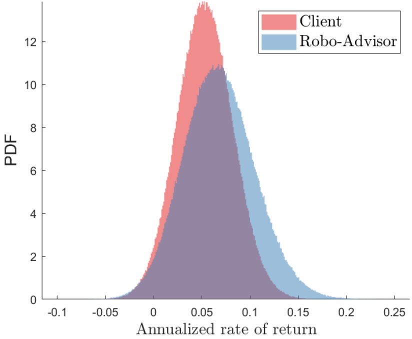

Left: Client ($\delta < 0$) vs. Buy-and-Hold ($\delta = 0$).

Right: Client ($\delta < 0$) vs. Robo-Advisor ($\delta > 0$).

Future Developments

- Should the robo-advisor's recommendations depend on the client's account **monitoring frequency**?
  - Effect of short-term losses outweighs the effect of short-term gains $\rightarrow$ gradual increase in risk aversion and shift to safer investments
  - Gneezy and Potters [1997]: causal relationship between frequency at which returns are evaluated and willingness to accept risk
  - Myopic loss aversion exacerbated by technology?
- Threshold-based rebalancing (Beketov et al. [2018])
  - Portfolio weights adjusted only if there are sufficiently large changes in market prices and client characteristics

Future Developments

- Strategically determine interaction times to avoid inflated risk aversion and ensure good investment performance
    - Take into account investment performance since the previous interaction time - reference point of client
    - Restrict interaction during bad economic conditions to keep client invested in stock market
- Strategically communicate investment performance to clients
    - Emphasize projected long-term wealth distribution
    - Display the performance of a well diversified portfolio rather than the performance of individual asset classes

Future Developments

- Robo-advisor's recommendations accounting for uncertainty in elicitation of risk preferences
  - Noisy and biased risk aversion coefficients
  - Disentangling risk aversion from client expectations
  - Accuracy of elicitation methods depend on financial literacy and numerical skills (Dave et al. [2010])
- Clients may not be easily contacted or willing to respond
  - Investor attentiveness (e.g. Abel et al. [2007, 2013] and Gargano and Rossi [2018])

Thank you!

Bibliography

A. Abel, J. Eberly, and S. Panageas. Optimal Inattention to the Stock Market. *American Economic Review*, 2, 244–249, 2007.

A. Abel, J. Eberly, and S. Panageas. Optimal Inattention to the Stock Market with Information Costs and Transactions Costs. *Econometrica*, 81(4), 1455–1481, 2013.

F. D'Acunto, N. Prabhala, and A. Rossi. The Promises and Pitfalls of Robo-Advising. *Review of Financial Studies*, 32(5), 1983–2020, 2019.

F. D'Acunto and A. Rossi. Robo-Advising. *Palgrave Macmillan Handbook of Technological Finance*, Forthcoming. Editors: R. Rau, R. Wardrop, and L. Zingales

R. Barsky, F. Juster, M. Kimball, and M. Shapiro. Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study. *Quarterly Journal of Economics*, 112(2), 537–579, 1997.

M. Beketov, K. Lehmann, and M. Wittke. Robo Advisors: Quantitative Methods Inside the Robots. *Journal of Asset Management*, 19, 363–370, 2018.

S. Benartzi and R. Thaler. Myopic Loss Aversion and the Equity Premium Puzzle. *Quarterly Journal of Economics*, 110(1), 73–92, 1995.

Bibliography

A. Bucciol and R. Miniaci. Financial Risk Propensity, Business Cycles and Perceived Risk Exposure. *Oxford Bulleting of Economics and Statistics*, 80(1), 160–183, 2018.

J. Campbell and J. Cochrane. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *J Political Economy*, 107(2), 205–251, 1999.

A. Capponi, S. Ólafsson, and T. Zariphopoulou. Personalized Robo-Advising: Enhancing Investment through Client Interaction. *Management Science*. Revise and Resubmit, 2020.

A. Cohn, J. Engelmann, E. Fehr, and M. Maréchal. Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals. *American Economic Review*, 105(2), 860-–885, 2015.

C. Dave, C. Eckel, C. Johnson, and C. Rojas. Eliciting Risk Preferences: When is Simple Better? *Journal of Risk and Uncertainty*, 41, 219–243, 2010.

Bibliography

E. Fama and K. French. Business Conditions and Expected Returns on Stocks and Bonds. *Journal of Financial Economics*, 25(1), 23–49, 1989.

A. Gargano and A. Rossi. Does It Pay to Pay Attention? *Review of Financial Studies*, 31(12), 4595–4649, 2018.

U. Gneezy and J. Potters. An Experiment on Risk Taking and Evaluation Periods. *Quarterly Journal of Economics*, 631–645, 1997.

L. Guiso and M. Paiella. Risk Aversion, Wealth, and Background Risk. *Journal of the European Economic Association*, 6(6), 1109–1150, 2008.

L. Guiso, P. Sapienza, and L. Zingales. Time Varying Risk Aversion. *Journal of Financial Economics*, 128(3), 403–421, 2018.

J. Hamilton. A new Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57, 357–384, 1989.

M. Lettau and S. Ludvigson. Measuring and Modeling Variation in the Risk-Return Tradeoff. *Handbook of Financial Econometrics*, 1, 617–690, 2010.

Bibliography

J. Oechssler, A. Roider, and P. Schmitz. Cognitive Abilities and Behavioral Biases. *Journal of Economic Behavior & Organization*, 72(1), 147–152, 1997.

M. Reher and S. Sokolinski. Does FinTech Democratize Investing? *Working Paper*, 2020.

A. Rossi and S. Utkus. Who benefits from Robo-Advising? Evidence from Machine Learning. *Working Paper*, George Washington University, 2019.

A. Rossi and S. Utkus. The Needs and Wants in Financial Advise: Humans versus Robo-Advising. *Working Paper*, George Washington University, 2019.

C. Sahm. How Much Does Risk Tolerance Change? *Quarterly Journal of Finance*, 2(4), 1–38, 2012.

M. Seasholes and L. Feng. Do Investor Sophistication and Trading Experience Eliminate Behavioral Biases in Financial Markets? *Review of Finance*, 9, 305–351, 2005.