



Working Paper 109-2021 | April 2021

Robo-Advising: Less AI and More XAI?

Augmenting algorithms with humans-in-the-loop

Confidence
must be earned

Amundi

ASSET MANAGEMENT

Robo-Advising: Less AI and More XAI?

Augmenting algorithms with humans-in-the-loop*

Abstract

Milo BIANCHI

Toulouse School of
Economics, TSM, and IUF,
University of Toulouse
Capitole

milo.bianchi@tse-fr.eu

Marie BRIÈRE

Amundi Research
Paris Dauphine University,
Université Libre de Bruxelles

marie.briere@amundi.com

We start by considering some of the key reasons behind the academic and industry interest in robo-advisors. We discuss how robo-advice could potentially address some fundamental problems in investors' decision making as well as in traditional financial advice. We then move on to some of the ongoing issues regarding the future of robo-advice. Firstly, the role Artificial Intelligence (AI) plays, and should play, in robo-advice. Secondly, how far should the personalisation of robo-advice recommendations go. Third, how trust in automated financial advice can be generated and maintained. Fourth, whether robots are perceived as complements or substitutes to human decision-making. Our conclusion outlines some thoughts on what the next generation of robo-advisors might look like. We highlight the importance of recent insights in Explainable AI (XAI) and how new forms of AI applied to financial services would benefit from importing insights from economics and psychology to design effective human/robot interaction.

Keywords: Robo-Advising, Trust, Financial Inclusion, Explainable Artificial Intelligence

JEL classification: G11; G51; G41; G23; D14

Acknowledgement

**This chapter was prepared for the book *Machine Learning in Financial Markets: A Guide to Contemporary Practice* edited by Agostino Capponi and Charles-Albert Lehalle for Cambridge University Press. We thank Pascal Blanqué, Zvi Bodie, Pascal Buisson, Agostino Capponi, Mehrshad Motahari and an anonymous referee for insightful comments. Milo Bianchi acknowledges funding from LTI@Unito, TSE Sustainable Finance Center and ANR (ANR-17-EURE-0010 grant). We have no material interest that relates to the research described in the paper.*

About the authors



Milo Bianchi

Milo Bianchi is Professor of Finance at the Toulouse School of Economics. His current research agenda focuses on fintech and on sustainable finance, with specific focus on individual investors. His work has been published in leading economics and finance journals including *Journal of Finance*, *Review of Economic Studies*, *Journal of Economic Theory*, and *Management Science*. Milo is junior member of the Institut Universitaire de France, director of the FIT-IN Initiative, and member of the Sustainable Finance and the Digital Finance Centers at TSE. Milo has received his PhD from the Stockholm School of Economics and he has held research positions at various institutions including MIT, Paris School of Economics, University College London, and Shanghai University of Finance and Economics.



Marie Brière

Marie Brière is Head of the Investor Research Center at Amundi. She is also an affiliate professor at Paris Dauphine University and an associate researcher at the Centre Emile Bernheim at Solvay Brussels School of Economics & Management. Dr. Brière began her career as a quantitative researcher at the proprietary trading desk at BNP Paribas. She also served as a fixed-income strategist at Crédit Lyonnais Asset Management and as head of fixed income, foreign exchange, and volatility strategy at Crédit Agricole Asset Management. Her scientific articles have been published in international academic journals including the *Financial Analyst Journal*, *Journal of Banking and Finance*, *Journal of International Money and Finance*, *World Development*, etc. She received the Markowitz award for her article with Zvi Bodie on “Sovereign Wealth and Risk Management”, published in the *Journal of Investment Management*. She holds a PhD in economics from the University Paris X and graduated from ENSAE.

1 Introduction

Automated portfolio managers, commonly known as robo-advisors, are attracting a growing interest both in academia and across the investment industry. In this article, we start by reviewing some of the reasons behind this growing interest. We emphasise how robo-advice can be seen in the broader context of the so-called Fintech revolution. We also consider some more specific reasons of interest in automated financial advice, building on the fundamental problems that individual investors face in taking financial decisions, and on the limits often observed in traditional financial advice.

In the second part of the paper, we discuss how robo-advising could potentially address these fundamental problems and highlight robots' main promises. First, promote financial inclusion by reaching under-served investors; second, provide tailored recommendations based on accountable procedures, and, finally, make investors better off. For each of these promises, we review the reasons why some hope can be placed on robots and we take a stand on what the academic literature has shown so far.

In the third part of the article, we address what we believe are fundamental ongoing issues in the future of robo-advice. Firstly, we discuss what role Artificial Intelligence (AI) plays, and should play, in robo-advice. We take a rather general textbook definition of AI, intended broadly speaking as any attempt to design robo-advisors as "intelligent agents" that collect information on the investor and on the market and propose strategies aiming at maximizing the chances to fulfill the investor's goals. We stress the constraints that may limit which forms of AI can be placed into robo-advice, in terms both of the regulatory challenges and of the conceptual advances of portfolio theory. We also emphasise how the quest for simplicity and the ease of explanation in recommendations could make some forms of AI undesirable even if feasible. Secondly, we discuss how far we should go into the personalisation of robo-recommendations, highlighting the trade-off between aiming to bring a portfolio closer to a specific individual's needs and the risks related to possible measurement errors of relevant individual characteristics (say, risk aversion) and to the sensitivity of algorithms to parameter uncertainty. Third, we discuss how robo-advice can shed light on the broader issues of human/robot interactions and on the mechanics of trust in automated financial services. We review the arguments of algorithm aversion, and the possible ways to reduce it, and how these can be applied in the context of automated financial advice. Finally, we discuss some evidence of whether robots are perceived as complements or substitutes to human decision-making.

We conclude with our thoughts on what the next generation of robo-advisors may look. Rather than continuing the trend of using more data, more complex models and more automated interactions, we define an alternative path that builds on the key premise of robo-advice in terms of increased accountability and financial inclusion, and on the key challenge of developing trust in financial technology. We highlight the importance of recent insights on XAI (Explainable Artificial Intelligence), centered around building algorithms in which the underlying model or the model's predictions can be explained to the users. We also stress how new forms of AI applied to financial services can benefit from importing insights from social sciences, such as economics and psychology.

This review does not aim to be exhaustive. Rather, it should be seen as complementary to existing reviews (such as D’Acunto and Rossi, 2020) and to the other chapters in this book.

2 Why So Popular?

Robo-advisors use automated procedures, ranging from relatively simple algorithms that use limited information on the client to complex systems built around big data, with the purpose of recommending how to allocate funds across different types of assets. First, a client profiling technique is used to assess an investor’s characteristics (risk aversion, financial knowledge, investment horizon...) and goals. Second, an investment universe is defined and, third, a portfolio is proposed by taking into account the investment goals and desired risk level. As documented in Beketov, Lehmann and Wittke (2018), in most cases, the optimal portfolio builds on modern portfolio theory, dating back to Markowitz (1952). In addition to recommending an initial allocation of funds, algorithms can be designed to continuously monitor portfolios and detect deviations from the targeted profile. Whenever deviations are identified, the client is alerted and/or the portfolio is automatically rebalanced. The portfolio can also be automatically rebalanced to reduce risk as time goes by or when the investor changes their risk tolerance or investment goals. Some robots also propose to implement “tax harvesting” techniques: selling assets that experience a loss and using the proceeds to buy assets with similar risk, to decrease capital gains and taxable income without affecting the portfolio’s exposure to risk. Apart from the portfolio allocation, the robot can display statistics of interest to the client, such as the expected annual return and volatility, often by using historic performance and Monte Carlo simulations of the possible future outcomes of the portfolio allocation.

The market is growing rapidly. Most practitioners estimate that the global market is currently around \$1 trillion (Statista), as compared to \$100bn in 2016 (S&P Global Market Intelligence, Backend Benchmarking, Aite Group - see Buisson (2019)). Worldwide assets under management in the robo-advice segment worldwide are projected to reach between \$1.7trn and \$4.6trn in 2022 (Statista, BI Intelligence). The number of users is expected to amount to 436m by 2024 (Statista 2020). This growth is driven by the entry of large incumbents in the digital service arena (for example, JPMorgan and Goldman Sachs announced the launch of a digital wealth management services in 2020) and the migration of assets managed by large financial institutions to their robo-advice operations, which amounts to 8% of their AUM and to one-quarter of the assets in accounts with less than \$1m. At the same time, clients have increased their demand for digital investment tools, and particularly for low-cost portfolio management and adjacent services such as financial planning. The United States remains, by far, the leading market for robo-advising (with more than 200 robo-advisors registered), but the number of robo-advisors is growing rapidly in Europe (more than 70), and also in Asia, driven by an emerging middle class and high technological connectivity (Abraham, Schmukler and Tessada (2019)).¹

We refer to Grealish and Kolm (2021) for more details on the functioning of robo-advisors and on recent market trends. We next highlight a few reasons which may be motivating such a rapid market growth and increased interest in academic and policy circles.

¹Robo-advisors are already present in China, India, Japan, Singapore, Thailand and Vietnam.

2.1 Fintech Revolution

Part of the interest in robo-advising comes from the broader trend of applying new technologies and novel sources of data in the financial domain, a phenomenon often dubbed as fintech. The word has played a central role in many academic and policy debates in the past few years. Enthusiasts about fintech talk about a revolution that promises to disrupt and reshape the financial service industry.²

Buchanan (2019) discusses the global growth of the AI industry and its application to the finance industry. She documents an impressive growth in AI startups during the past five years, driven by the advances in computing power, leading to a decline in the cost of processing and storing data, while at the same time by the availability of data of increased size and scope. Similarly, AI related patent publications (denoted by the AI keyword) in the US have grown from around 50 in 2013 to around 120 in 2017. In China, such growth has been even more dramatic, with around 120 patents in 2013 rising to 640 patents in 2018. Buchanan (2019) also discusses the broad range of ways in which AI is changing the financial services industry, not only in terms of robo-advising but also for fraud detection and compliance, chatbots, and algo-trading.

In academic circles, the increased attention can be seen for example from the exponential growth in finance academic studies primarily centered around AI. Bartram, Branke and Motahari (2020) analyze the number of AI-related keywords in the title, abstract, or listed keywords of all working papers posted in the Financial Economics Network (FEN) between 1996 and 2018.³ In 1996, no working paper with any AI-related keyword was uploaded, in 2018 the number of posted papers including such keywords were 410, accounting for 3% of all papers posted that year.

Robo-advisors promise to apply new technologies and procedures to improve financial decision making, as we discuss below, and as such they can be seen as a piece of the broader fintech revolution, just like digital currencies that promise to redefine the role of traditional money and platform lending which promises to redefine the role of traditional access to credit.

2.2 Fundamental Problems with Investors

From a more specific perspective, one key interest in robo-advising is that it is now commonly understood that many investors could substantially improve their financial decisions. In the past decades, the literature has documented various ways in which investors' decisions may deviate, sometimes in a fundamental way, from the standard premises of a fully rational economic agent, who knows the entire set of possible alternatives, the associated outcomes in a probabilistic sense, and can correctly match all the information in order to maximize life-time utility.

Restricting to the investment domain, which has so far been the typical focus of robo-advisors, investors have been found to display low participation (Mankiw and Zeldes, 1991), underdiversification (Grinblatt, Keloharju and Linnainmaa, 2011; Goetzmann and Kumar, 2008; Bianchi and Tallon, 2019), default bias (Benartzi and Thaler, 2007), portfolio inertia (Agnew, Balduzzi and Sunden, 2003; Biliias, Georgarakos and Haliassos, 2010),

²See e.g. The Economist (2015) on The Fintech Revolution or The World Economic Forum (2017) on Beyond FinTech: A pragmatic assessment of disruptive potential in financial service.

³Keywords included artificial intelligence, machine learning, cluster analysis, genetic algorithm or evolutionary algorithm, lasso, natural language processing, neural network or deep learning, random forest or decision tree, and support vector machine.

excessive trading (Odean, 1999), trend chasing (Greenwood and Nagel, 2009), and a poor understanding of matching mechanism (Choi, Laibson and Madrian, 2009). Many of those investment behaviours are associated with a poor understanding of basic financial principles (Lusardi, Michaud and Mitchell, 2017; Lusardi and Mitchell, 2014; Bianchi, 2018).

Several surveys provide a comprehensive list of biases and associated trading mistakes (see e.g. Guiso and Sodini, 2013, Barber and Odean, 2013, Beshears, Choi, Laibson and Madrian (2018)). For the purpose of this article, two points are worth stressing. First, these mistakes are not small; on the contrary, their welfare implications can be substantial (Campbell, 2006, Campbell, Jackson, Madrian and Tufano (2011)). Second, they do not cancel out in equilibrium; rather, they have important effects on the functioning of financial markets and on broader macroeconomic issues such as wealth inequality (Vissing-Jorgensen, 2004; Lusardi, Michaud and Mitchell, 2017; Bach, Calvet and Sodini, 2020; Fagereng, Guiso, Malacrino and Pistaferri, 2020).

Motivated by this evidence, it is clear that improving financial decision making can be seen as a major goal of financial innovation, and part of the interest in robo-advising lies in its promise to help investors in this dimension.

2.3 Fundamental Problems with Advisors

A natural response to investors' poor financial decision-making is to delegate the task to professional experts, who have the time and skill to serve investors' best interest. The argument relies on a few important assumptions, which may sometimes be difficult to meet in practice. First, advisors are required to be able to recognise and adapt their strategies to match their clients' preferences and needs. This is far from obvious and recent evidence suggests that advisors may themselves have misguided beliefs. Foerster, Linnainmaa, Melzer and Previtro (2017) analyse the trading and portfolio decisions of around 10,000 financial advisors and 800,000 clients in four Canadian financial institutions. They show that clients' observable characteristics (risk tolerance, age, income, wealth, occupation, financial knowledge) jointly explain only 12% of the cross-sectional variation in clients' risk exposure. This is remarkably low, especially compared to the effect of being served by a given advisor, which explains 22% of the variation in clients' risk exposure. In terms of incremental explanatory power, adding advisor effects to a model in which investors' risk exposure is explained by their observable characteristics improves the adjusted R^2 from 12% to 30%. This evidence suggests that, in some cases, financial recommendations are closer to "one size fits all" than being fully tailored to clients' specific preferences and needs. Furthermore, Linnainmaa, Melzer and Previtro (2020) show that some advisors, when trading with their own money, display very similar trading biases to their clients: they prefer active management, they chase returns and they are not well diversified. Similar flaws have been noticed for investment advice provided to participants in self-directed retirement plans (Bodie, 2003).

A second key aspect is that advisors need to have the incentives to act in clients' best interests, rather than pursuing their own goals. Again, recent evidence suggests this need not be the case. Mullainathan, Noeth and Schoar (2012) conducted a study by training auditors, posing as customers of various financial advisors, and (randomly) asking them to represent different investment strategies and biases. They show that advisors display a significant bias towards active management, they initially support clients' requests but their final recommendations are orthogonal to clients' stated preferences. At the end,

advisors failed to correct clients' biases and even made clients worse off. Similarly, Foà, Gambacorta, Guiso and Mistrulli (2019) document banks' strategic behaviours in their offer of mortgage contracts. A more extensive review of advisors' conflicted advice is provided in Beshears et al. (2018).

A third key aspect is that, even removing the previous concerns, financial advising is costly, and a significant part of the cost has a fixed component (say, the advisor's time). This implies that financial advice may not be accessible to investors with lower levels of wealth, who may in fact be those who need it most.

3 Promises

3.1 Accountable Procedures and Tailored recommendations

Robo-advisors' services offer accountable procedure to allocate an individual's portfolio across various asset classes and different types of funds, depending on individual characteristics. Two stages of the process are crucial to this : (1) client profiling, and (2) asset allocation. While tailored recommendations are offered to clients, there is considerable heterogeneity in the recommended allocations for the same investor's profile, and the exact algorithms used by robo-advisors are typically not transparent.

Client profiling

Robo-advisors typically use an online questionnaire to assess investors' financial situation, characteristics and investment goals. This questionnaire is a regulatory requirement under SEC guidelines in the US (SEC, 2006; SEC, 2019). A "suitability assessment" is also mandatory under MiFID (Markets in Financial Instruments Directive) regulation in Europe.⁴

Individual characteristics, such as age, marital status, net worth, investment horizon and risk tolerance are used to assess the investor's situation. Interestingly, a large variety of questions can be used to estimate one particular characteristic. For example, if you consider risk tolerance, most of the robo-profilers use subjective measures of risk aversion based on a self-assessment. Some robo-profilers use risk capacity metrics (measuring the ability to bear losses), estimated from portfolio loss constraints, financial obligations or expenses, balance sheet information, etc. In Europe, under MiFID II, advisors also assess the clients' "experience and knowledge" to understand the risks involved in the investment product or service offered.⁵ Robo-advisors thus ask questions about the clients' financial literacy and reduce the individuals' risk tolerance when financial literacy is low.

Robo-advisors typically propose their clients to pick a goal (for example, retirement, buying a house, a bequest to family members, a college/education fund or a safety net) among several possibilities during the risk profiling questionnaire. This goal can define

⁴In Europe, the MiFID regulation has set the objective of increased and harmonized individual investors' protection, according to their level of financial knowledge. MiFID I (2004/39 / 3C), implemented in November 2007, requires investment companies to send their clients a questionnaire to determine their level of financial knowledge, their assets and their investment objectives. MiFID I has been replaced in January 2018 by MiFID II (2014/65 / EU), which has demanded a strengthening of legislation in several areas, in particular in the requirements of advice independence and transparency (on costs, available offering, etc.).

⁵see Article 25(3) and 56.

the investment horizon or the risk capacity in the optimal portfolio allocation. Other robo-advisors allow their clients to name their goal before or outside of the risk profiling process, and do not necessarily incorporate it into the portfolio allocation. Finally, a few robo-advisors allow their clients to set multiple goals, thus offering their clients the ability to explicitly put their portfolio in a mental account (Das, Markowitz, Scheid and Statman, 2010). One limitation is that robo-advisors frequently lack a global view of an investor’s financial situation, as savings outside of the robo platform are rarely taken into account. Some of them have a broader view of the clients’ financial situation through partnerships with financial account aggregators or digital platforms of investment. For example, Wealthfront recently featured direct integrations with digital investment platforms (Venmo, Redfin, Coinbase), lending (Lending Club) and tax calculation (turbotax).

Asset Allocation

In a second step, the robo-advisor proposes to structure a portfolio by taking into account investment goals and the desired risk level. Beketov et al. (2018) analysed 219 robo advisors from 28 countries (30% in the US, 20% in Germany, 14% in the UK), that were founded between 1997 and 2017. As shown in Figure 1, a word count representing the occurrence of different methods used by robo advisors, a large variety of portfolio construction techniques are used. Beketov et al. (2018) show that most robo advisors use simple Markowitz optimisation or a variant of it, such as Black Litterman (40%), sample portfolios applying a pre-defined grid (27%) or constant portfolio weights (14%). A minority of robo advisors use alternative portfolio construction techniques such as liability driven investment, full-scale optimisation, risk parity and constant proportion portfolio insurance.

Figure 1: Word cloud representing the occurrence of different methods used by existing robo-advisors



Source: Beketov et al. (2018)

If most robo-advisors perform asset allocation by using a mean–variance analysis, or a variant of it, they rarely disclose information on how they choose their asset class investment universe or how they estimate variances and correlations between asset classes. They even more rarely disclose their expected return and risk parameters explicitly. Among

the dominant players in the US, Wealthfront is probably one of the few exceptions. It discloses on its website its portfolio optimisation method (Black-Litterman), but also its expected returns, volatilities and correlation matrices and the way these are estimated.⁶ Betterment is also relatively transparent. It provides justification and detail on the choice of its investment universe, its portfolio optimisation method (Black-Litterman) and the way it calculated expected returns and risk, without disclosing them explicitly.⁷ Schwab Intelligent Portfolios also discloses its portfolio optimisation method, a variant of the Markowitz approach (using Conditional Value at Risk instead of variance). However, it is less transparent on its Monte-Carlo simulation methodology and expected return hypotheses.⁸

Heterogeneity in the proposed asset allocations

In theory, these rigorous procedures and their systematic nature should make it possible to overcome the shortcomings of human advisors, by reducing unintentional biases and simplifying the interaction with the client. Rebalancing for example, is made easier through robo-advising platforms that implement this automatically or require a simple validation by the client. Also, if individual characteristics are measured with sufficient precision, robo-advising services should make it possible to offer investment recommendations that are tailored to each investor’s situation.

In practice, a large disparity in the proposed asset allocations has been documented, for the same investor’s profile. For example, Boreiko and Massarotti (2020) analyses 53 robo-advisors operating in the US and Germany in 2019. They show that a “moderate” profile invests an average 56% in equities, but the standard deviation of the proposed equity exposure is large (23%). Equity exposure can go from 14% to 100%, depending on the robo-advisor. Aggressive or conservative asset allocations have similar features, with an average equity exposure of 73% and 35% respectively, but range between 18% and 100% for aggressive allocations, and from 0 to 100% for conservative allocations.

This disparity in the proposed allocations can have several sources. It could come from different portfolio construction methodologies or different expected risk/return hypotheses. It may also reflect robo-advisors’ conflicts of interest. Boreiko and Massarotti (2020) show that the asset managers’ expertise in a given asset class (proxied as the percentage of funds in a given asset class across the total universe of funds proposed by the robo-advisor) is the main driver. Conflicts of interest were also demonstrated in the case of Schwab Intelligent Portfolio, which recommended that a significant portion of the clients’ portfolio be invested in money market funds. Lam (2016) argued that this unusually large asset allocation to cash allowed Schwab Intelligent Portfolios to delegate cash

⁶<https://research.wealthfront.com/whitepapers/investment-methodology/>

⁷On the investment universe, they excluded asset classes such as private equity, commodities and natural resources, since “estimates of their market capitalization is unreliable and there is a lack of data to support their historical performance”. Expected returns are derived from market weights, through a classical reverse optimisation exercise that uses the variance covariance matrix between all asset classes. An estimation of this covariance matrix is made using historical data, combined with a target matrix, and using the Ledoit and Wolf (2004) shrinkage method to reduce estimation error. Portfolios can also be tilted towards Fama and French (1992) value and size factors, the size of the tilt being freely parametrised by the confidence that Betterment has in these views. See <https://www.betterment.com/resources/betterment-portfolio-strategy/citations>.

⁸They simulate 10,000 hypothetical future realisations of returns, using fat-tailed assumptions for the distribution of asset returns, also allowing for changing correlations modeled with a Copula approach. See <https://intelligent.schwab.com/page/our-approach-to-portfolio-construction>.

management to Schwab Bank, allowing the firm to profit from the interest rate difference between lending rates and the paid rate of return (Fisch, Laboure and Turner, 2019).

3.2 Make Investors Better Off

As for many innovative financial services, a key promise of robo-advice is to make investors better off. Recent academic studies document that robo-advising services tend to improve investors' diversification and risk-adjusted returns. Such improvement can come from static changes in portfolio choices, for example by improving diversification and therefore reducing risk for a given level of expected returns. Or they may occur over time, by allowing investors to rebalance their portfolios in a way that stays closer to their target risk-return profile.

D'Acunto, Prabhala and Rossi (2019) study a Portfolio Optimizer targeting Indian equities, and find that robo-advice was beneficial to ex-ante under diversified investors, by increasing their portfolio diversification, reducing their risk and increasing their ex-post mean returns. However, the study also documents that not all investors are winners: the robo-advisor did not improve the performance of already-diversified investors. Rossi and Utkus (2019*b*) study the effects of a large U.S. robo-advisor on a population of previously self-directed investors. They find that, across all investors, robo-advice users reduced their money market investment and increased their bond holdings. The introduction of robo-advice also reduced idiosyncratic risk by lowering the holdings of individual stocks and active mutual funds and raising exposure to low-cost indexed mutual funds. It also reduced portfolios' home bias by significantly increasing international equity and fixed income diversification. The introduction of the robot increased individuals' overall risk-adjusted performance. In a different sample, Reher and Sun (2019) also pointed to a diversification improvement of robo-advice users generated by a large US robo-advisor. Bianchi and Brière (2021) study the introduction of a large French robo-advisor on employee savings's plans. They find that relative to self-managing, accessing the robo-advice services is associated with an increase in individuals' investment and risk-adjusted returns. Investors bear more risk, and rebalance their portfolios in a way that keeps their allocation closer to the target. This increased risk taking is also found by Hong, Lu and Pan (2020), studying a Chinese robo-advisor, and using unique account-level data on consumption and investments from Ant Group. Robo-adoption helped households move towards optimal risk-taking, reducing their consumption volatility.

What are the exact implications of the above evidence on investors' welfare remains an open question, one which is obviously difficult to test. Most of these studies check whether having access to the robot increases investors' returns, after having controlled for some measures of portfolio risk. Some key challenges need to be confronted if one wishes to venture into developing a broader welfare analysis. First, one needs to acquire a good understanding of investors' preferences, constraints and outside opportunities (for example, how they would otherwise have used the capital invested with the robo-advisor), as well as a broader picture of investors' assets.

A second challenge is that it is notoriously difficult to define benchmarks against which evaluating performance if we recognize that investors' utility functions cannot always be reduced to, or well-approximated by, mean-variance preferences. Warren (2019) provide a practical guide on how to construct portfolios under different utility functions. Jondeau and Rockinger (2006) show the importance of higher moments when returns are far from a normal distribution. Barberis, Jin, Wang et al. (2020) show that when investors have

prospect theory preferences, quantitative analyses can be performed, but a number of extra parametric assumptions are needed.

A third aspect is that even if one can have reasonable approximations of how investors trade-off risk and returns, investors may care about other elements. For example, some investors may just use financial advice to acquire peace of mind. Gennaioli, Shleifer and Vishny (2015) propose a model in which a financial advisor acts as a "money doctor" and allows investors to effectively decrease their reluctance to take risk. Rossi and Utkus (2019a) document that acquiring peace of mind is one of the key drivers behind the demand for financial advice.

Despite those important challenges, adopting a more structured approach based on explicit assumptions on investors' utility and constraints is an area where in our view further research would be most useful.

3.3 Reach Under-served Investors

One of the most important promises of the fintech revolution is linked to financial inclusion. As mentioned previously, offering financial services often involves substantial fixed costs, which can make it unprofitable to serve poorer consumers. New technologies allow a dramatic decrease in transaction costs (Goldfarb and Tucker (2019) identify various ways through which this could happen). By reducing these costs, new technologies can reach those who have been traditionally under-served (Philippon (2019)).

Robo-advisors can be seen as part of this promise. First, they typically require lower initial capital to open an account. For example, Bank of America requires US\$25,000 to open an account with a private financial advisor, but only US\$5,000 to open an account with its robo-advisor. Some robo-advisors, such as Betterment, do not require a minimum investment at all. Second, they typically charge lower fees than human advisors. The automation of the advice process reduces the advice-related fixed costs. For example, a fully automated robo-advisor in the US typically charges a fee between 0.25% and 0.50% of assets managed (this is between 0.25% and 0.75% in Europe),⁹ whereas the fees for traditional human advisors rarely fall below 0.75% and can even reach 1.5% (Lopez, Babccic and De La Ossa, 2015; Better Better Finance, 2020).

Academic studies on robo advising and financial inclusion are scarce, but the initial results seem to support the above claims. D'Hondt, De Winne, Ghysels and Raymond (2020) perform a counterfactual analysis in which each investor is matched with an "AI alter-ego" that performs an investment strategy commonly employed by robo-advisors. The authors compute counterfactual returns on a large set of investors and show that investors with low income and low education have potentially the most to gain from robo-advice. Hong et al. (2020) show that the adoption of a popular fintech platform in China is associated with increased risk taking, and the effect is particularly large for households residing in areas with low financial service coverage. Reher and Sokolinski (2020) analyse the effects of the reduction of the account minimum from \$5,000 to \$500 by a major U.S. robo-advisor. They show that, thanks to this reduction, there was a 59% increase in the share of "middle class" participants (with wealth between \$1,000 and \$42,000), but no increase in participation by households with wealth below \$1,000. The majority of new middle-class robo-advice participants are also new to the stock market and, relative to upper class participants, they increase their risky share by 13 pps and their expected return by 1.2 pps. Bianchi and Brière (2021) also show that robo-advice

⁹We consider here management fees only, not underlying ETFs or funds' fees.

participants increase their risk exposure and their risk adjusted returns. Importantly, the increase in risk exposure is larger for investors with smaller portfolios and lower equity exposure at the baseline, and the increase in returns is larger for smaller investors and for investors with lower returns at the baseline. Finally, investors may also learn from the robo-advisory tool, as shown by Loos, Previtero, Scheurle and Hackethal (2020) who document an improvement in portfolio efficiency in the non-robo advised part of their portfolio. These results suggest that having access to a robo advisor may be particularly important for investors who are less likely to receive traditional advice and, as such, it can be seen as an important instrument towards financial inclusion.

4 Open Questions

4.1 Why Not More AI/Big Data?

There is growing academic interest in how AI can be used to enhance robo-advisors (Bartram et al., 2020; Xue, Liu, Li, Liu, Ye, Wang and Yin, 2018; D'Hondt et al., 2020). However, as previously mentioned, most robo-advisors in practice build on rather simple procedures both in terms of the information employed to profile the client and of how this information is used to construct an optimal portfolio. As emphasised in Beketov et al. (2018), modern portfolio theory remains dominant. At the same time, we often lack precise information on the actual use of AI algorithms by robo-advisors. This may seem surprising given the increased interest in AI and Big Data mentioned previously, and since robo-advisors are often presented as incorporating these trends. One may wonder why we fail to see more AI built into robo-advising.

A first reason may be that, while such inclusion would be desirable, it is not feasible due to technological or knowledge constraints. That is, finance theory has not advanced enough to be able to give recommendations on how to incorporate AI into finance models. Some scholars would not agree. Bartram et al. (2020) summarise the shortcomings of classical portfolio construction techniques and highlight how AI techniques improve the practice. In particular, they argue that AI can produce better risk-return estimates, solve portfolio optimisation problems with complex constraints, and yield better out-of-sample performance compared with traditional approaches.

A second reason may be that including more AI would violate regulatory constraints. According to the current discipline, as a registered investment advisor, a robo-advisor has a fiduciary duty to its clients. As discussed by Grealish and Kolm (2021), the fiduciary duty in the U.S. builds on the 1940 Advisers Act and it has been adapted by the SEC in 2017 so as to accommodate the specifics of robo-advising. In particular, robo-advisors are required to elicit enough information on the client, use properly tested and controlled algorithms, and fully disclose the algorithms' possible limitations.

Legal scholars debate how much a robo-advisor can and should be subject to fiduciary duty. Fein (2017) argues that robo-advisors cannot be fully considered as fiduciaries since they are programmed to serve a specific goal of the client, as opposed to considering their broader interest. As such, they cannot meet the standard of care of the prudent investor required for human advisers. Similarly, Strzelczyk (2017) stresses that robo-advisors cannot act as a fiduciary since they do not provide individualised portfolio analysis but rather base their recommendations on a partial knowledge of the client. On the other hand, Ji (2017) argues that robo-advisors can be capable of exercising the duty of loyalty to their clients so as to meet the Advisers Act's standards. In a similar vein, Clarke

(2020) argues that the fiduciary duty can be managed by basing recommendations on finance theory and by fully disclosing any possible conflict of interest.

A third reason may be that having more AI in robo-advice is simply not desirable. Incorporating AI would at least partly make these robots a black-box and would make it harder to provide investors with clear explanations of why certain recommendations are given. Patel and Lincoln (2019) identify three key sources of risk associated with AI applications: first, opacity and complexity; second, the distancing of humans from decision-making; and third, changing incentive structures (for example in data collection efforts). They consider the implications of these sources of risk in several domains, ranging from damaging trust in financial services, propagating biases, harming certain group of customers possibly in an unfair way. They also consider market level risks ranging from financial stability, cybersecurity and new regulatory challenges.

Algorithm complexity could be particularly problematic in bad times. Financial Stability Board (2017) argues that the growing use of AI in financial services can threaten financial stability. One reason is that AI can create new forms of interconnectedness between financial markets and institutions, since various institutions may employ previously unrelated data sources, for example. Moreover, the opacity of AI learning methods could become a source of macro-level risk due to their possibly unintended consequences.

Algorithm complexity is also particularly problematic for those with lower financial capabilities. Complex financial products have been shown to be particularly harmful for less sophisticated investors (see e.g. Bianchi and Jehiel (2020) for a theoretical investigation, Ryan, Trumbull and Tufano (2011) and Lerner and Tufano (2011) for historic evidence, and C el erier and Vall ee (2017) for more recent evidence). As for many (financial) innovations, the risk is that they do not reach those who would need it the most, or that they end up being misused.

In this way, some key promises of robo-advising, notably on improved financial inclusion and accountability, can be threatened by the widespread use of opaque models.

4.2 How Far Should We Go Into Personalisation?

The potential of robo-advice is to combine financial technology and artificial intelligence and offer to each investor personalised advice based on their objectives and preferences. One important difficulty lies in the precise measurement of investors' characteristics. A second issue relates to the sensitivity of the optimal asset allocation to these characteristics, which can be subject to a large degree of uncertainty. This can lead the estimated optimal portfolio to be substantially different from the truly optimal one, with dramatic consequences for the investor.

Difficulty of measuring an individual's characteristics

Lo (2016) calls for the development of smart indices, that could be tailored to individuals' circumstances and characteristics. Even if we are not there yet, robo-advisors could make a step in that direction, by helping to precisely define an investor's financial situation and goals (Gargano and Rossi, 2020). As it has been demonstrated by a large number of academic papers, optimal portfolio choices rely on various individual characteristics such as human capital (Viceira, 2001), housing market exposure (Kraft and Munk, 2011), time preference, risk aversion, ambiguity aversion (Dimmock, Kouwenberg, Mitchell and Peijnenburg, 2016; Bianchi and Tallon, 2019), etc. Individualisation

possibilities are much wider than what is currently implemented in robo-advice services.

For example, portfolio choice models with labor income risk advise that households account for the covariance between financial and non-financial income in their asset allocation. Labor income is an important source of heterogeneity across individuals. The usual hypothesis that human capital can be proxied by an inflation-linked bond has been challenged (Cocco, Gomes and Maenhout, 2005; Benzoni, Collin-Dufresne and Goldstein, 2007;). In some countries and for some categories of the population (typically, for very low and very high quantiles of the income distribution), income shocks display a positive correlation with equities (Guvenen, Karahan, Ozkan and Song, 2021). There is empirical evidence that households already account for labor income in their self-managed portfolios (Bagliano, Fugazza and Nicodano, 2019). Those whose occupations are more sensitive to the business cycle hold fewer high-beta stocks (Betermier, Calvet and Sodini, 2017). Robo-advisors would be well suited to offer this kind of personalization, but to the best of our knowledge, they currently do not do it.

One of the reasons for this lack of personalization of robo-advice services is that some individual characteristics are difficult to measure and subject to a large degree of uncertainty. Risk aversion is one of them. Different methods have been developed by economists and psychologists to measure individuals' risk aversion. Most of them are experimental measurements based on hypothetical choices. For example, the lotteries of Barseky, Juster, Kimball and Shapiro (1997) offer individuals the choice between employment with a risk-free salary, and a higher but risky salary. Other work (Holt and Laury, 2002; Kapteyn and Teppa, 2011; Weber, Weber and Nosić, 2013) measure preferences based on a series of risk/return trade-offs. The choice between a certain gain and a risky lottery is repeated, gradually increasing the prize until the subject picks a risky lottery.

One reason why it is difficult to measure risk aversion might be that people interpret outcomes as gains and losses relative to a reference point and are more sensitive to losses than to gains. Kahneman, Knetsch and Thaler (1990) or Barberis, Huang and Santos (2001) report experimental evidence of loss aversion. Loss aversion can also explain why many investors prefer portfolio insurance products offering capital guarantees (Calvet, Celerier, Sodini and Vallee, 2020).

In practice, robo-advisors frequently assess a clients' risk tolerance based on a self-declaration. People are asked to rate themselves in their ability to take risks on a scale of 1 to 10 (Dohmen, Falk, Huffman, Sunde, Schupp and Wagner, 2005). These ratings have the disadvantage of not being very comparable across individuals. Scoring techniques are also frequently used by robo-advisors. They ask the individual a large number of questions, covering different aspects of life (consumption, leisure, health, financial lotteries, work, retirement and family). Global scores are obtained by adding the scores across various dimensions, keeping only those questions which prove to be the most relevant ex-post to measure an individual's risk aversion, a statistical criterion which eliminates the least relevant questions (Arrondel and Masson, 2013).

In Europe, the implementation of MiFID regulation led to several academic studies assessing risk profiling questionnaires. European regulation does not impose a standardised solution, each investment company remains free to develop its questionnaire as it wishes, which explains the great heterogeneity of the questionnaires distributed in practice to clients. Marinelli and Mazzoli (2010) sent three different questionnaires used by banks to 100 potential investors to verify the consistency of the clients' risk profiles. Only 23% of individuals were profiled in a consistent way across the three questionnaires, a likely consequence of the differences in the contents and scoring methods of the questionnaires.

Other work carried out in several European countries (De Palma, Picard and Prigent, 2009; Marinelli and Mazzoli, 2010; Linciano and Soccorso, 2012) arrived to the same conclusion.

Algorithm sensitivity to parameter uncertainty

Optimal allocations are usually very sensitive to parameters (expected returns, covariance of assets' returns) which are hard to estimate. They also depend crucially on investor's characteristics (financial wealth, human capital, etc.) often known with poor accuracy. On one hand, there is a cost for suboptimal asset allocation (one size does not fit all) and substantial gains to individualise (see Dahlquist, Setty and Vestman, 2018; Warren, 2019). On the other hand, there is a risk of overreaction to extreme/time-varying individual characteristics, potentially leading to "extreme" asset allocations, as it has been shown by the literature on optimisation with parameter uncertainty (see for example Garlappi, Uppal and Wang, 2007). Blake, Cairns and Dowd (2009) claim that some standardisation is needed, like in the aircraft industry, to guarantee investors' security. How much customisation is needed depends largely on the trade-off between the gains to bring the portfolio closer to an individual's needs and the risk of estimating an individual's characteristics with a large degree of error.

How stable an individual characteristics are in practice also remains an open question. Capponi, Olafsson and Zariphopoulou (2019) show that if these risk profiles are changing through time (depending on idiosyncratic characteristics, market returns or economic conditions), the theoretical optimal dynamic portfolio of a robo-advisor should adapt to the client's dynamic risk profile, by adjusting the corresponding inter-temporal hedging demands. The robo-advisor faces a trade-off between receiving client information in a timely manner and mitigating behavioural biases in the risk profile communicated by the client. They show that with time-varying risk aversion, the optimal portfolio proposed by the robo-advisor should counter the client's tendency to reduce market exposure during economic contractions.

4.3 Can Humans Trust Robots?

In the interaction between humans and robo-advisors, a key ingredient is trust, determining the individual's willingness to use the service and to follow the robo recommendations. We review what creates trust in algorithms and discuss the impact of trust on financial decisions.

Trust is key for robo-advice adoption

Trust has been shown to be a key driver of financial decisions (see Sapienza, Toldra-Simats and Zingales (2013) for a review). For example, trustful investors are significantly more likely to invest in the stock market (Thakor and Merton, 2018). Trust is also a potential key driver of robo-advice adoption. As stated by Merton (2017), "What you need to make technology work is to create trust."

Trust has been studied across a variety of disciplines, including sociology, psychology and economics, to understand how humans interact with other humans, or more recently with machines. Trust is a "multidimensional psychological attitude involving beliefs and expectations about the trustee's trustworthiness, derived from experience and interactions

with the trustee in situations involving uncertainty and risk” (Abbass, Scholz and Reid, 2018). One can also see trust as a transaction between two parties: if A believes that B will act in A’s best interest, and accepts vulnerability to B’s actions, then A trusts B (Misztal, 2013). Importantly, trust exists to mitigate uncertainty and the risk of collaboration by enabling the trustor to anticipate that the trustee will act in the trustor’s best interests.

While trust has both cognitive and affective features, in the automation literature, cognitive (rather than affective) processes seem to play a dominant role. Trust in robots is multifaceted. It has been shown to depend on robot reliability, robustness, predictability, understandability, transparency, and fiduciary responsibility (Sheridan, 1989; Sheridan, 2019; Muir and Moray, 1996). One key feature of robo-advisors is their reliance on more or less complicated algorithms, in several steps of the advisory process. An algorithm is used to profile the investor, and then to define the optimal asset allocation. A client delegating the decision to the robot bears the risk that a wrong decision by the robot will lead to poor performance of their savings. Trust in these algorithms is thus key for robo-advisor adoption.

Algorithm aversion

Survey evidence (HSBC, 2019) shows that there is a general lack of trust in algorithms. While most people seem to trust their general environment and technology (68% of the survey respondents said they will trust a person until prove otherwise, 48% believe the majority of people are trustworthy and 76% that they feel comfortable using new technology), artificial intelligence is not yet trusted. Only 8% of respondents would trust a robot programmed by experts to offer mortgage advice, compared to 41% trusting a mortgage broker. As a comparison, 9% would be likely to use a horoscope to guide investment choices! 14% would trust a robot programmed by leading surgeons to conduct open heart surgery on them, while 9% would trust a family member to do an operation supported by a surgeon. Only 19% said they would trust a robo-advisor to help them make investment choices. There are large differences across countries however. The percentage of respondents who trust robo-advisors rises to 44% and 39% in China and India respectively, but it is only 9% and 6% in France and Germany.

Some academic studies have shown that decision makers are often averse to using algorithms, most of the time preferring less accurate human judgment. For example, professional forecasters have been shown not to use algorithms or give them insufficient weight (Fildes and Goodwin, 2007). Dietvorst, Simmons and Massey (2015) gave participants the choice of either exclusively using an algorithm’s forecasts or exclusively using their own forecasts during an incentivised forecasting task. They found that most participants chose to use the algorithm exclusively only when they had no information about the algorithm’s performance. However, when the experimenter told them it was imperfect, they were much more likely to choose the human forecast. This effect persisted even when they had explicitly seen the algorithm outperform the human’s forecasts. This tendency to irrationally discount advice that is generated and communicated by computer algorithms has been called “algorithm aversion”. In a later experimental study (Dietvorst, Simmons and Massey, 2018), participants were given the possibility to modify the algorithm. Participants were considerably more likely to choose the imperfect algorithm when they could modify its forecasts, even if they were severely restricted in the modifications they could make. This suggests that algorithm aversion can be reduced by giving people some control over an imperfect algorithm’s forecast.

Recent experimental evidence shows less algorithm aversion. Niszczoła and Kaszás (2020) tested if people exhibited algorithm aversion when asked to decide whether they would use human advice or an artificial neural network to predict stock price evolution. Without any prior information on the human versus robot’s performance, they found no general aversion towards algorithms. When it was made explicit that the performances of the human advisor was similar to that of the algorithm, 57% of the participants showed a preference for the human advice. In another experiment, subjects were asked to choose a human or a robo-advisor to exclude stocks that were controversial. Interestingly, people perceived algorithms as being less effective than humans when the tasks required a subjective judgment to be made, such as morality.

Germann and Merkle (2019) also found no evidence of algorithm aversion. In a laboratory experiment (mostly based on business or economics’ students), they asked participants to choose between a human fund manager and an investment algorithm. The selection process was repeated ten times, which allowed them to study the reaction to the advisor’s performance. With equal fees for both advisors, 56% of participants decided to follow the algorithm. When fees differed, most participants (80%) chose the advisor with the lower fees. Choices were strongly influenced by the cumulative past performance. But investors did not lose confidence in the algorithm more quickly after seeing forecasting errors. An additional survey provided interesting qualitative explanations to the results. Participants believed in the ability of the algorithm to be better able to learn than humans. They viewed humans as having a comparative advantage in using qualitative data and dealing with outliers. All in all, the algorithms are viewed as a complement rather than a competitor to a human advisor.

This reluctance of some clients to use purely automated platforms has to a few cases of hybrid advisors, in which robo-advisors also allow clients to speak with a human advisor. Scalable Capital launched in 2017 over-the-phone and face-to-face consultations for an additional fee charged to clients. Vanguard Personal Advisor Service also stands out as an example of such a hybrid advisor that relies on both automated and non-automated advice.

What creates trust in an algorithm?

Jacovi, Marasović, Miller and Goldberg (2020) distinguish two sources of trust in algorithm: intrinsic and extrinsic. Intrinsic trust can be gained when the observable decision process of the algorithm matches the user’s priors. Explanations of the decision process behind the algorithm can help create intrinsic trust.¹⁰ Additionally, an algorithm can become trustworthy through its actual behaviour: in this case, the source of trust is not the decision process of the model, but the evaluation of its output.

The European Commission (2019) recently listed a number of requirements for trustworthy algorithms. Related to intrinsic trust are the requirements of (1) the user’s agency and human oversight, (2) privacy and data governance, (3) transparency and the ability to explain the algorithm. Extrinsic trust can be increased by (4) the technical robustness and safety of the algorithm, (5) the ability to interpret its output, (6) its accountability and auditability. In addition, ethical and fairness considerations, such as (7) avoiding discrimination, promoting diversity and fairness or (8) encouraging societal

¹⁰For example, a robo-advisor may disclose its risk profiling methodology, its optimization method and risk/return hypotheses, or reveal the signals leading to portfolio rebalancing.

and environmental well-being are also considered to be a key component of trust.¹¹

Trust in algorithms also crucially depends on the perception of the expertise and reliability of the humans or institutions offering the service (Prahla and Van Swol, 2017). “Technology doesn’t create trust on its own” (Merton, 2017). People trust humans certifying a technology, not necessarily the technology itself. In the specific case of robo-advice, Lourenço, Dellaert and Donkers (2020) study the decision of consumers to adopt the service and show that this decision is clearly influenced by the for-profit versus not-for-profit orientation of the firm offering the service (for example private insurance and investment management firm versus pension fund or government-sponsored institution). Transparency, explainability and interpretability may not be sufficient by themselves for enhancing decisions and increasing trust. However, informing key hypotheses and potential shortcomings of algorithms when making certain decisions, may be a fundamental dimension to be worked on.

Trust in robots and financial decisions

Not everyone trusts robo-advisors. In a sample of 34,000 savers in French employee savings’ plans, Bianchi and Brière (2021) document that individuals who are young, male, and more attentive to their saving plans (measured by the time spent on the savings plan website), have a higher probability of adopting a robo-advice service. The probability of taking up the robo-advice option is also negatively related to the size of the investors’ portfolio, which suggests that the robo-advisor is able to reach less wealthy investors,¹² a result also confirmed by Brenner and Meyll (2020). Investors with smaller portfolios are also more likely to assign a larger fraction of their assets to the robot.

A unique feature of the robo-advice service studied by Bianchi and Brière (2021) allows them to analyse both “robo-takers” and the “robo-curious,” i.e., individuals who observe the robot’s recommendation without actually subscribing to it. Interestingly, the further away is the robot’s recommendation relative to the current allocation, the larger the probability that the investor will subscribe to the robot. This finding can be contrasted with the observation that human advisors tend to gain trust from their clients by being accommodating with clients (Mullainathan et al., 2012). Moreover, investors who are younger, female, those who have larger risk exposure and lower past returns, as well as less attentive investors are more likely to accept a larger increase in their exposure to risky assets, such as equities. These results confirm the common view that robo advising may develop as a popular investment choice for relatively young households. This may reflect a combination of their lower wealth as well as increased willingness to trust technology.

Trust can have a large impact on investor’s decisions. Bianchi and Brière (2021) and Hong et al. (2020) show evidence of increased risk taking, a result consistent with increased trust. For example, Bianchi and Brière (2021) document a 7% increase in equity exposure after robo-advice adoption (relative to an average 16% exposure). Hong et al. (2020) document a 14% increase (relative to an average risky exposure of 37% on

¹¹The Commission will make a regulatory proposal in 2021. This proposal will aim to safeguard user safety by obliging high-risk AI systems to meet mandatory requirements related to their trustworthiness. For example, ensuring there is human oversight, and clear information on the capabilities and limitations of AI.

¹²Conversely, wealthier investors are more likely to acquire information about the robot without subscribing to the service.

their sample of 50,000 Chinese consumer clients of Alibaba). Interestingly, Hong et al. (2020) additionally show that this result is likely not to be driven by an increase in the individual's risk tolerance driven by robot support. Rather, it seems to reflect a better alignment of the investment portfolio with the actual risk tolerance of the individual. In particular, they show that after robo-advice adoption, exposure to risky assets is more in line with the individual's risk tolerance estimated from their consumption growth volatility (Merton, 1971), measured from Alibaba's Taobao online shopping platform. The robo-advisor seems to help individuals to move closer to their optimal alignment of risk-taking and consumption. These results should, however, be used with caution, as both studies concentrate on a relatively short period of investment (absent any serious market crash) and lack a global view on the individuals' overall portfolios. More work would need to be done to document a long term impact.

4.4 Do Robots Substitute or Complement Human Decision-Making?

Autonomous systems are being developed across large areas of our everyday life. Understanding how humans will interact with them is a key issue. In particular, should we expect that robots will become substitutes to humans or rather be complementary? In the special case of financial advice, are they likely to replace human advisors?

Using a representative sample of US investor, Brenner and Meyll (2020) investigate whether robo-advisors, reduce investors' demand for human financial advice offered by financial service providers. They document a large substitution effect and show that this effect is driven by investors who fear to be victimised by investment fraud or worry about potential conflicts of interest. In practice however, a number of platforms that were entirely digital decided to reintroduce human advisors. For example, Scalable Capital, the European online robo-advice company backed by BlackRock, or Nutmeg, reintroduced over-the-phone and face-to-face consultations after finding that a number of clients preferred talking to human advisors rather than solely answering online questionnaires.

Another related question is how people will interact with robots. Will they delegate the entire decision to the robot or will they keep an eye on it, to monitor the process and intervene if necessary? In certain experiments, users put too much faith in robots. Robynette, Li, Allen, Howard and Wagner (2016) designed an experiment where participants were asked choose to follow, or not to follow, a robot's instructions in an emergency. All participants followed the robot during the emergency, even though half of the participants observed the same robot perform poorly in a non-emergency navigation guidance task just a few minutes before. Even when the robot pointed to a dark room with no discernible exit the majority of people did not choose to safely exit the way they had entered. Andersen, Köslich, Pedersen, Weigelin and Jensen (2017) expand on this work and show that such over-trust can also affect human/robot interactions that are not set in an emergency situation.

In the context of financial decisions, Bianchi and Brière (2021) document that robo-advisor adoption leads to significantly increase attention on savings plans, during the months following the adoption. Individuals are in general more attentive to their saving plan, particularly when they receive variable remuneration and need to make an investment decision (in their context, saving plans are not automatically rebalanced). This seems to indicate that people do not use the robot as a substitute for their own attention.

5 The Next Generation of Robo-Advisors

It is not clear which generation of robo-advisors we are currently facing. Beketov et al. (2018) focus on robots of third and fourth generation, which differ from earlier generations as they use more automation and more sophisticated methods to construct and rebalance portfolios. One possibility is that the next generation of robots would continue the trend of using more data and more complex models. One may, however, imagine an alternative path. As discussed previously, incorporating more complex AI into robo-advice (and more generally into financial services) faces three key challenges. Firstly, while highly personalised asset allocations have the great potential of accommodating an individual's needs, they are also more exposed to measurement errors of relevant individual characteristics and to parameter uncertainty. Secondly, to the extent that increased AI is associated with increased opacity, the risk is to miss some key promises of increased accountability and financial inclusion. Third, trust is key for technology adoption, even more so in the domain of financial advice. These challenges, in our view, call for algorithms that can be easily interpreted and evaluated. Toreini, Aitken, Coopamootoo, Elliott, Zelaya and van Moorsel (2020) discuss how developing trust in (machine learning) technologies requires them to be fair, explainable, accountable, and safe (FEAS).

From this perspective, recent advances in so-called XAI (explainable artificial intelligence), can be particularly useful when thinking about the future of robo-advisors (Molnar, 2020). "Explainability" refers to the possibility of explaining a given prediction or recommendation, even if based on a very complicated model, for example by evaluating the sensitivity of the prediction when changing one of the inputs. It also refers to how much a given model can itself be explained. Explanations can help humans perform a given task and, at the same time, evaluate a given model (see e.g. Biran and Cotton (2017) for a recent survey). As discussed in Doshi-Velez and Kim (2017), explainability can be considered a desiderata both in itself, in relation to the issues of trust and accountability expressed above, and also as a tool to assess whether other desiderata, such as fairness, privacy, reliability, robustness, causality, usability, are met.

There is a large amount of academic literature examining whether explainable artificial intelligence can improve human decision-making. How much explainability is needed for the actual functioning of an automated system remains an open question and it is often debated in the context of self-driving cars, for example. On the one hand, psychological research regarding decision-making suggests that when decisions involve complex reasoning, ignoring part of the available information and using heuristics can help to deal more robustly with uncertainty than relying on resource-intensive processing strategies (Gigerenzer and Brighton, 2009). On the other hand, experimental studies show that providing the driver with information on why and how an autonomous vehicle acts, is important to maintain a safe driving experience (Koo, Kwac, Ju, Steinert, Leifer and Nass, 2015). This information is particularly key in emergency situations. Drivers receiving such information tend to trust the car less and are faster to take control of the car when a dangerous situation occurs (Helldin, Falkman, Riveiro and Davidsson, 2013). One should also be particularly attentive to the risk of information overload. An algorithm is easier to interpret and to use when it focuses on a few features, it is also easier to correct in case of mistakes (Poursabzi-Sangdeh, Goldstein, Hofman, Vaughan and Wallach, 2018).

In the context of robo-advisors, explainability is not an easy task. Evaluating the

performance of a robo-recommendation is not straightforward, especially if one uses AI to move towards fully personalised allocations to be evaluated against fully personalised benchmarks (as described in Lo (2016)). Even more difficult for the client is to build counterfactuals of performance. And probably even more difficult is to appreciate the underlying finance model which governs the algorithm, especially if one wishes to serve less experienced investors who may lack financial literacy.

In that respect, the quest is not for full transparency of the potentially complicated algorithm underlying the robo-advice process, disclosing for example all the details of the portfolio optimisation methodology or the covariance matrix estimates. It would probably be more effective to disclose, for example, which economic scenarios may cause the algorithm to perform less accurately, possibly proving ex-post sub-optimal, and informing clients about the potential limitations of the algorithm.

Acquiring reliable information about clients' characteristics and preferences is a key step towards an adequate customization of financial advice, but this remains challenging. Recently, some robo-advisors have incorporated automated financial planning platforms, which allow recovering more precise information on clients' dynamic financial decisions, such as borrowing and consumption. These robo-advisors would seem well-suited to benefit from AI. For example, one promising development could be to improve the quality of the individual's information set available to the robo-advisor through, say, imputation of unobserved characteristics, such as unknown risk-aversion, or even more basic information about the household's balance sheet (e.g., debt, non-financial assets, non-financial income), from actual financial behavior (Hong et al., 2020). These are promising avenues for future academic research and applications.

Another potentially interesting development would be to strengthen the interactions with clients. For example, some robo-advisors send alerts when a clients' portfolio deviates significantly from the target asset allocation (see e.g. Bianchi and Brière (2021)). These alerts could also be seen as an opportunity to interact with the client. For example, alerts could be used to explain why a deviation occurred (market movements, change in personal characteristics, etc.) or why rebalancing is recommended. Another example is that one could elicit customers' perceptions regarding the quality of the response provided by the algorithm and integrate this feedback as part of the evaluation of the robo-advice service (Dupont, 2020).

These issues are not new in AI. Biran and Cotton (2017) discuss earlier approaches of explainability of decisions in expert systems in the 1970s and more recently in recommender systems. One may argue, however, that today models are probably more complex, more autonomous and they span a larger set of decisions across a larger set of agents (including possibly less sophisticated ones), which make these issues particularly relevant in current debates. Indeed, improving transparency is also central to the policy domain, such as the recent EU regulation on data protection (GDPR). As discussed in Goodman and Flaxman (2017), the law defines a right to explanation, whereby users can inquire about the logic involved in an algorithmic decision affecting them (say, through profiling), and this calls for algorithms which are as explainable as they are efficient.

Some prominent scholars argue that the AI revolution has not happened yet. Instead of mimicking human interactions or more sophisticated human thinking, the AI revolution will happen when new forms of intelligence are considered (Jordan (2019a)). In this context, importing insights from social sciences seems crucial. AI needs psychology to capture how humans actually think and behave, or, to reference Lo (2019), to include forms of "artificial stupidity". Insights from philosophy, psychology and cognitive sciences

are also key to informing how explanations are and should be communicated. Miller (2019) reviews the large amount of literature in these fields and emphasises the importance of providing selective explanations, based on causal relations and counterfactuals rather than likely statistical relations, and of allowing a social dimension in which explainers and “explainees” may interact. AI also needs economics not only to help address causality and discuss counterfactuals, but also to help design new forms of collective intelligence. These new forms may go beyond a purely anthropocentric approach, and build on some understanding of how markets function and how they may fail (Jordan (2019b)). We share the enthusiasm of these scholars when imagining advances in these directions, we look forward to seeing more social sciences in the next generation of robo-advisors!

References

- Abbass, H. A., Scholz, J. and Reid, D. J. (2018), *Foundations of trusted autonomy*, Springer Nature.
- Abraham, F., Schmukler, S. L. and Tessada, J. (2019), ‘Robo-advisors: Investing through machines’, *World Bank Research and Policy Briefs* (134881).
- Agnew, J., Balduzzi, P. and Sunden, A. (2003), ‘Portfolio choice and trading in a large 401 (k) plan’, *American Economic Review* **93**(1), 193–215.
- Andersen, K. E., Kösllich, S., Pedersen, B. K. M. K., Weigelin, B. C. and Jensen, L. C. (2017), Do we blindly trust self-driving cars, in ‘Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-robot Interaction’, pp. 67–68.
- Arrondel, L. and Masson, A. (2013), Measuring savers’ preferences how and why?, Technical report, HAL Working Paper.
- Bach, L., Calvet, L. E. and Sodini, P. (2020), ‘Rich pickings? risk, return, and skill in household wealth’, *American Economic Review* **110**(9), 2703–47.
- Bagliano, F. C., Fugazza, C. and Nicodano, G. (2019), ‘Life-cycle portfolios, unemployment and human capital loss’, *Journal of Macroeconomics* **60**, 325–340.
- Barber, B. M. and Odean, T. (2013), The behavior of individual investors, in ‘Handbook of the Economics of Finance’, Vol. 2, Elsevier, pp. 1533–1570.
- Barberis, N. C., Jin, L. J., Wang, B. et al. (2020), ‘Prospect theory and stock market anomalies’, *NBER Working Paper* (27155).
- Barberis, N., Huang, M. and Santos, T. (2001), ‘Prospect theory and asset prices’, *The Quarterly Journal of Economics* **116**(1), 1–53.
- Barsky, R. B., Juster, F. T., Kimball, M. S. and Shapiro, M. D. (1997), ‘Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study’, *The Quarterly Journal of Economics* **112**(2), 537–579.
- Bartram, S. M., Branke, J. and Motahari, M. (2020), *Artificial Intelligence in Asset Management*, number 14525, CFA Institute Research Foundation.

- Beketov, M., Lehmann, K. and Wittke, M. (2018), ‘Robo advisors: quantitative methods inside the robots’, *Journal of Asset Management* **19**(6), 363–370.
- Benartzi, S. and Thaler, R. (2007), ‘Heuristics and biases in retirement savings behavior’, *Journal of Economic Perspectives* **21**(3), 81–104.
- Benzoni, L., Collin-Dufresne, P. and Goldstein, R. S. (2007), ‘Portfolio choice over the life-cycle when the stock and labor markets are cointegrated’, *Journal of Finance* **62**(5), 2123–2167.
- Beshears, J., Choi, J. J., Laibson, D. and Madrian, B. C. (2018), Behavioral household finance, in ‘Handbook of Behavioral Economics: Applications and Foundations 1’, Vol. 1, Elsevier, pp. 177–276.
- Betermier, S., Calvet, L. E. and Sodini, P. (2017), ‘Who are the value and growth investors?’, *Journal of Finance* **72**(1), 5–46.
- Better Finance (2020), Robo-advice 5.0: Can consumers trust robots?, Technical report.
- Bianchi, M. (2018), ‘Financial literacy and portfolio dynamics’, *Journal of Finance* **73**(2), 831–859.
- Bianchi, M. and Brière, M. (2021), ‘Robo-advising for small investors’, *Working Paper available at SSRN* (3751620).
- Bianchi, M. and Jehiel, P. (2020), ‘Bundlers’ dilemmas in financial markets with sampling investors’, *Theoretical Economics* **15**(2), 545–582.
- Bianchi, M. and Tallon, J.-M. (2019), ‘Ambiguity preferences and portfolio choices: Evidence from the field’, *Management Science* **65**(4), 1486–1501.
- Bilias, Y., Georgarakos, D. and Haliassos, M. (2010), ‘Portfolio inertia and stock market fluctuations’, *Journal of Money, Credit and Banking* **42**(4), 715–742.
- Biran, O. and Cotton, C. (2017), Explanation and justification in machine learning: A survey, in ‘IJCAI-17 workshop on explainable AI (XAI)’, Vol. 8(1), pp. 8–13.
- Blake, D., Cairns, A. and Dowd, K. (2009), ‘Designing a defined-contribution plan: What to learn from aircraft designers’, *Financial Analysts Journal* **65**(1), 37–42.
- Bodie, Z. (2003), ‘An analysis of investment advice to retirement plan participants’, *The pension challenge: Risk transfers and retirement income security* pp. 19–32.
- Boreiko, D. and Massarotti, F. (2020), ‘How risk profiles of investors affect robo-advised portfolios how risk profiles of investors affect robo-advised portfolios’, *Frontiers in Artificial Intelligence* **3**, 60.
- Brenner, L. and Meyll, T. (2020), ‘Robo-advisors: A substitute for human financial advice?’, *Journal of Behavioral and Experimental Finance* **25**, 100275.
- Buchanan, B. (2019), ‘Artificial intelligence in finance’, *Alan Turing Institute Working Paper available at <http://doi.org/10.5281/zenodo.2612537>*.

- Buisson, P. (2019), ‘Pure robo-advisors have become viable competitors in the US’, *Amundi Digibook* .
- Calvet, L. E., Celerier, C., Sodini, P. and Vallee, B. (2020), ‘Can security design foster household risk-taking?’, *Working Paper available at SSRN* (3474645).
- Campbell, J. Y. (2006), ‘Household finance’, *Journal of Finance* **61**(4), 1553–1604.
- Campbell, J. Y., Jackson, H. E., Madrian, B. C. and Tufano, P. (2011), ‘Consumer financial protection’, *Journal of Economic Perspectives* **25**(1), 91–114.
- Capponi, A., Olafsson, S. and Zariphopoulou, T. (2019), ‘Personalized robo-advising: Enhancing investment through client interaction’, *arXiv preprint arXiv:1911.01391* .
- C  l  rier, C. and Vall  e, B. (2017), ‘Catering to investors through security design: Headline rate and complexity’, *The Quarterly Journal of Economics* **132**(3), 1469–1508.
- Choi, J. J., Laibson, D. and Madrian, B. C. (2009), ‘Mental accounting in portfolio choice: Evidence from a flypaper effect’, *American Economic Review* **99**(5), 2085–95.
- Clarke, D. (2020), ‘Robo-advisors-market impact and fiduciary duty of care to retail investors’, *Working Paper available at SSRN* (3539122).
- Cocco, J. F., Gomes, F. J. and Maenhout, P. J. (2005), ‘Consumption and portfolio choice over the life cycle’, *Review of Financial Studies* **18**(2), 491–533.
- D’Acunto, F., Prabhala, N. and Rossi, A. G. (2019), ‘The promises and pitfalls of robo-advising’, *Review of Financial Studies* **32**(5), 1983–2020.
- D’Acunto, F. and Rossi, A. G. (2020), ‘Robo-advising’, *CESifo Working Paper 8225* .
- Dahlquist, M., Setty, O. and Vestman, R. (2018), ‘On the asset allocation of a default pension fund’, *Journal of Finance* **73**(4), 1893–1936.
- Das, S., Markowitz, H., Scheid, J. and Statman, M. (2010), ‘Portfolio optimization with mental accounts’, *Journal of Financial and Quantitative Analysis* **45**(2), 311–334.
- De Palma, A., Picard, N. and Prigent, J.-L. (2009), ‘Prise en compte de l’attitude face au risque dans le cadre de la directive mifid’, *HAL Working Paper* .
- D’Hondt, C., De Winne, R., Ghysels, E. and Raymond, S. (2020), ‘Artificial intelligence alter egos: Who might benefit from robo-investing?’, *Journal of Empirical Finance* **59**, 278–299.
- Dietvorst, B. J., Simmons, J. P. and Massey, C. (2015), ‘Algorithm aversion: People erroneously avoid algorithms after seeing them err.’, *Journal of Experimental Psychology: General* **144**(1), 114.
- Dietvorst, B. J., Simmons, J. P. and Massey, C. (2018), ‘Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them’, *Management Science* **64**(3), 1155–1170.

- Dimmock, S. G., Kouwenberg, R., Mitchell, O. S. and Peijnenburg, K. (2016), ‘Ambiguity aversion and household portfolio choice puzzles: Empirical evidence’, *Journal of Financial Economics* **119**(3), 559–577.
- Dohmen, T. J., Falk, A., Huffman, D., Sunde, U., Schupp, J. and Wagner, G. G. (2005), ‘Individual risk attitudes: New evidence from a large, representative, experimentally-validated survey’, *Working Paper available at SSRN* (807408).
- Doshi-Velez, F. and Kim, B. (2017), ‘Towards a rigorous science of interpretable machine learning’, *arXiv preprint arXiv:1702.08608*.
- Dupont, L. (2020), ‘Gouvernance des algorithmes d’intelligence artificielle dans le secteur financier: Analyse des réponses à la consultation’, *Pôle finance innovation de l’ACPR* (December).
- European Commission (2019), ‘Ethics guidelines for trustworthy ai’, *Available at <https://ec.europa.eu/futurium/en/ai-alliance-consultation>*.
- Fagereng, A., Guiso, L., Malacrino, D. and Pistaferri, L. (2020), ‘Heterogeneity and persistence in returns to wealth’, *Econometrica* **88**(1), 115–170.
- Fein, M. L. (2017), ‘Are robo-advisors fiduciaries?’, *Working Paper available at SSRN* (3028268).
- Fildes, R. and Goodwin, P. (2007), ‘Against your better judgment? how organizations can improve their use of management judgment in forecasting’, *Interfaces* **37**(6), 570–576.
- Financial Stability Board (2017), ‘Artificial intelligence and machine learning in financial services: Market developments and financial stability implications’.
- Fisch, J. E., Laboure, M. and Turner, J. A. (2019), ‘The emergence of the robo-advisor’, *The Disruptive Impact of FinTech on Retirement Systems* **13**.
- Foà, G., Gambacorta, L., Guiso, L. and Mistrulli, P. E. (2019), ‘The supply side of household finance’, *Review of Financial Studies* **32**(10), 3762–3798.
- Foerster, S., Linnainmaa, J. T., Melzer, B. T. and Previtero, A. (2017), ‘Retail financial advice: does one size fit all?’, *Journal of Finance* **72**(4), 1441–1482.
- Gargano, A. and Rossi, A. G. (2020), ‘There’s an app for that: Goal-setting and saving in the fintech era’, *Working Paper available at SSRN* (3579275).
- Garlappi, L., Uppal, R. and Wang, T. (2007), ‘Portfolio selection with parameter and model uncertainty: A multi-prior approach’, *Review of Financial Studies* **20**(1), 41–81.
- Gennaioli, N., Shleifer, A. and Vishny, R. (2015), ‘Money doctors’, *Journal of Finance* **70**(1), 91–114.
- Germann, M. and Merkle, C. (2019), ‘Algorithm aversion in financial investing’, *Working Paper available at SSRN* (3364850).

- Gigerenzer, G. and Brighton, H. (2009), ‘Homo heuristicus: Why biased minds make better inferences’, *Topics in cognitive science* **1**(1), 107–143.
- Goetzmann, W. N. and Kumar, A. (2008), ‘Equity portfolio diversification’, *Review of Finance* **12**(3), 433–463.
- Goldfarb, A. and Tucker, C. (2019), ‘Digital economics’, *Journal of Economic Literature* **57**(1), 3–43.
- Goodman, B. and Flaxman, S. (2017), ‘European union regulations on algorithmic decision-making and a ”right to explanation”’, *AI magazine* **38**(3), 50–57.
- Grealish, A. and Kolm, P. N. (2021), Robo-advisory: From investing principles and algorithms to future developments, *in* ‘Machine Learning in Financial Markets: A Guide to Contemporary Practice’, Cambridge University Press.
- Greenwood, R. and Nagel, S. (2009), ‘Inexperienced investors and bubbles’, *Journal of Financial Economics* **93**(2), 239 – 258.
- Grinblatt, M., Keloharju, M. and Linnainmaa, J. (2011), ‘Iq and stock market participation’, *Journal of Finance* **66**(6), 2121–2164.
- Guiso, L. and Sodini, P. (2013), Household finance: An emerging field, *in* H. R. Stoll, G. M. Constantinides, M. Harris and R. Stulz, eds, ‘Handbook of the Economics of Finance’, Vol. 2, Elsevier, pp. 1397–1532.
- Guvenen, F., Karahan, F., Ozkan, S. and Song, J. (2021), ‘What do data on millions of us workers reveal about life-cycle earnings dynamics?’, *Econometrica* (forthcoming).
- Helldin, T., Falkman, G., Riveiro, M. and Davidsson, S. (2013), Presenting system uncertainty in automotive uis for supporting trust calibration in autonomous driving, *in* ‘Proceedings of the 5th international conference on automotive user interfaces and interactive vehicular applications’, pp. 210–217.
- Holt, C. A. and Laury, S. K. (2002), ‘Risk aversion and incentive effects’, *American Economic Review* **92**(5), 1644–1655.
- Hong, C. Y., Lu, X. and Pan, J. (2020), ‘Fintech adoption and household risk-taking’, *NBER Working Paper* (28063).
- HSBC (2019), Trust in technology, Technical report.
- Jacovi, A., Marasović, A., Miller, T. and Goldberg, Y. (2020), ‘Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in ai’, *arXiv preprint arXiv:2010.07487*.
- Ji, M. (2017), ‘Are robots good fiduciaries: Regulating robo-advisors under the investment advisers act of 1940’, *Colum. L. Rev.* **117**, 1543.
- Jondeau, E. and Rockinger, M. (2006), ‘Optimal portfolio allocation under higher moments’, *European Financial Management* **12**(1), 29–55.
- Jordan, M. I. (2019a), ‘Artificial intelligence - the revolution hasn’t happened yet’, *Harvard Data Science Review* **1**(1).

- Jordan, M. I. (2019b), ‘Dr. ai or: How i learned to stop worrying and love economics’, *Harvard Data Science Review* **1**(1).
- Kahneman, D., Knetsch, J. L. and Thaler, R. H. (1990), ‘Experimental tests of the endowment effect and the coase theorem’, *Journal of Political Economy* **98**(6), 1325–1348.
- Kapteyn, A. and Teppa, F. (2011), ‘Subjective measures of risk aversion, fixed costs, and portfolio choice’, *Journal of Economic Psychology* **32**(4), 564–580.
- Koo, J., Kwac, J., Ju, W., Steinert, M., Leifer, L. and Nass, C. (2015), ‘Why did my car just do that? explaining semi-autonomous driving actions to improve driver understanding, trust, and performance’, *International Journal on Interactive Design and Manufacturing (IJIDeM)* **9**(4), 269–275.
- Kraft, H. and Munk, C. (2011), ‘Optimal housing, consumption, and investment decisions over the life cycle’, *Management Science* **57**(6), 1025–1041.
- Lerner, J. and Tufano, P. (2011), ‘The consequences of financial innovation: a counterfactual research agenda’, *Annual Review of Financial Economics* **3**(1), 41–85.
- Linciano, N. and Soccorso, P. (2012), ‘Assessing investors’ risk tolerance through a questionnaire’, *Working Paper available at SSRN* (2207958).
- Linnainmaa, J. T., Melzer, B. T. and Previtero, A. (2020), ‘The misguided beliefs of financial advisors’, *Journal of Finance* .
- Lo, A. W. (2016), ‘What is an index?’, *Journal of Portfolio Management* **42**(2), 21–36.
- Lo, A. W. (2019), ‘Why artificial intelligence may not be as useful or as challenging as artificial stupidity’, *Harvard Data Science Review* **1**(1).
- Loos, B., Previtero, A., Scheurle, S. and Hackethal, A. (2020), ‘Robo-advisers and investor behavior’, *University of Technology of Sydney Working Paper* .
- Lopez, J. C., Babcic, S. and De La Ossa, A. (2015), ‘Advice goes virtual: how new digital investment services are changing the wealth management landscape’, *Journal of Financial Perspectives* **3**(3).
- Lourenço, C. J., Dellaert, B. G. and Donkers, B. (2020), ‘Whose algorithm says so: The relationships between type of firm, perceptions of trust and expertise, and the acceptance of financial robo-advice’, *Journal of Interactive Marketing* **49**, 107–124.
- Lusardi, A., Michaud, P.-C. and Mitchell, O. S. (2017), ‘Optimal financial knowledge and wealth inequality’, *Journal of Political Economy* **125**(2), 431–477.
- Lusardi, A. and Mitchell, O. S. (2014), ‘The economic importance of financial literacy: Theory and evidence’, *Journal of Economic Literature* **52**(1), 5–44.
- Mankiw, N. G. and Zeldes, S. P. (1991), ‘The consumption of stockholders and nonstockholders’, *Journal of Financial Economics* **29**(1), 97–112.

- Marinelli, N. and Mazzoli, C. (2010), ‘Profiling investors with the mi-fid: current practice and future prospects’, *Research Paper available at <https://www.ascosim.it/public/19Ric.pdf>*.
- Merton, R. (1971), ‘Optimal portfolio and consumption rules in a continuous-time model’, *Journal of Economic Theory* **3**(4), 373–413.
- Merton, R. C. (2017), ‘The future of robo-advisors’, *Video available at <https://www.cnn.com/2017/11/05/mit-expert-robert-merton-on-the-future-of-robo-advisors.html>*.
- Miller, T. (2019), ‘Explanation in artificial intelligence: Insights from the social sciences’, *Artificial Intelligence* **267**, 1–38.
- Misztal, B. (2013), *Trust in modern societies: The search for the bases of social order*, John Wiley & Sons.
- Molnar, C. (2020), *Interpretable machine learning*, Leanpub.
- Muir, B. M. and Moray, N. (1996), ‘Trust in automation. part ii. experimental studies of trust and human intervention in a process control simulation’, *Ergonomics* **39**(3), 429–460.
- Mullainathan, S., Noeth, M. and Schoar, A. (2012), ‘The market for financial advice: An audit study’, *NBER Working Paper* (17929).
- Niszczota, P. and Kaszás, D. (2020), ‘Robo-investment aversion’, *Plos one* **15**(9), e0239277.
- Odean, T. (1999), ‘Do investors trade too much?’, *American Economic Review* **89**(5), 1279–1298.
- Patel, K. and Lincoln, M. (2019), ‘It’s not magic: Weighing the risks of AI in financial services’, *Centre for the Study of Financial Innovation Working Paper*.
- Philippon, T. (2019), ‘On fintech and financial inclusion’, *NBER Working Paper* (26330).
- Poursabzi-Sangdeh, F., Goldstein, D. G., Hofman, J. M., Vaughan, J. W. and Wal-lach, H. (2018), ‘Manipulating and measuring model interpretability’, *arXiv preprint [arXiv:1802.07810](https://arxiv.org/abs/1802.07810)*.
- Prahl, A. and Van Swol, L. (2017), ‘Understanding algorithm aversion: When is advice from automation discounted?’, *Journal of Forecasting* **36**(6), 691–702.
- Reher, M. and Sokolinski, S. (2020), ‘Does fintech democratize investing?’, *Working Paper available at SSRN* (3515707).
- Reher, M. and Sun, C. (2019), ‘Automated financial management: Diversification and account size flexibility’, *Journal of Investment Management* **17**(2), 1–13.
- Robinette, P., Li, W., Allen, R., Howard, A. M. and Wagner, A. R. (2016), Overtrust of robots in emergency evacuation scenarios, in ‘2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)’, IEEE, pp. 101–108.

- Rossi, A. and Utkus, S. (2019a), ‘The needs and wants in financial advise: Humans versus robo-advising’, *George Washington University Working Paper* .
- Rossi, A. and Utkus, S. (2019b), ‘Who benefits from robo-advising’, *Working Paper available at SSRN* (3552671).
- Ryan, A., Trumbull, G. and Tufano, P. (2011), ‘A brief postwar history of us consumer finance’, *Business History Review* **85**(03), 461–498.
- Sapienza, P., Toldra-Simats, A. and Zingales, L. (2013), ‘Understanding trust’, *The Economic Journal* **123**(573), 1313–1332.
- SEC (2006), ‘Questions advisers should ask while establishing or reviewing their compliance programs’, *Available at https://www.sec.gov/info/cco/advise_compliance_questions.htm* .
- SEC (2019), ‘Commission interpretation regarding standard of conduct for investment advisors’, *Available at https://www.sec.gov/info/cco/adviser_compliance_questions.htm* .
- Sheridan, T. (1989), Trustworthiness of command and control systems, *in* ‘Analysis, Design and Evaluation of Man–Machine Systems 1988’, Elsevier, pp. 427–431.
- Sheridan, T. B. (2019), ‘Individual differences in attributes of trust in automation: measurement and application to system design’, *Frontiers in Psychology* **10**, 1117.
- Strzelczyk, B. E. (2017), ‘Rise of the machines: the legal implications for investor protection with the rise of robo-advisors’, *DePaul Business and Commercial Law Journal* **16**, 54.
- Thakor, R. T. and Merton, R. C. (2018), ‘Trust in lending’, *NBER Working Paper* (24778).
- Toreini, E., Aitken, M., Coopamootoo, K., Elliott, K., Zelaya, C. G. and van Moorsel, A. (2020), The relationship between trust in ai and trustworthy machine learning technologies, *in* ‘Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency’, pp. 272–283.
- Viceira, L. M. (2001), ‘Optimal portfolio choice for long-horizon investors with nontradable labor income’, *Journal of Finance* **56**(2), 433–470.
- Vissing-Jorgensen, A. (2004), Perspectives on behavioral finance: Does ”irrationality” disappear with wealth? evidence from expectations and actions, *in* M. Gertler and K. Rogoff, eds, ‘NBER Macroeconomics Annual 2003, Volume 18’, MIT Press, pp. 139–208.
- Warren, G. J. (2019), ‘Choosing and using utility functions in forming portfolios’, *Financial Analysts Journal* **75**(3), 39–69.
- Weber, M., Weber, E. U. and Nosić, A. (2013), ‘Who takes risks when and why: Determinants of changes in investor risk taking’, *Review of Finance* **17**(3), 847–883.

Xue, J., Liu, Q., Li, M., Liu, X., Ye, Y., Wang, S. and Yin, J. (2018), ‘Incremental multiple kernel extreme learning machine and its application in robo-advisors’, *Soft Computing* **22**(11), 3507–3517.

Chief Editors

Pascal BLANQUÉ

Chief Investment Officer

Philippe ITHURBIDE

Senior Economic Advisor

DISCLAIMER

In the European Union, this document is only for the attention of “Professional” investors as defined in Directive 2004/39/EC dated 21 April 2004 on markets in financial instruments (“MIFID”), to investment services providers and any other professional of the financial industry, and as the case may be in each local regulations and, as far as the offering in Switzerland is concerned, a “Qualified Investor” within the meaning of the provisions of the Swiss Collective Investment Schemes Act of 23 June 2006 (CISA), the Swiss Collective Investment Schemes Ordinance of 22 November 2006 (CISO) and the FINMA’s Circular 08/8 on Public Advertising under the Collective Investment Schemes legislation of 20 November 2008. In no event may this material be distributed in the European Union to non “Professional” investors as defined in the MIFID or in each local regulation, or in Switzerland to investors who do not comply with the definition of “qualified investors” as defined in the applicable legislation and regulation. This document is not intended for citizens or residents of the United States of America or to any «U.S. Person» , as this term is defined in SEC Regulation S under the U.S. Securities Act of 1933.

This document neither constitutes an offer to buy nor a solicitation to sell a product, and shall not be considered as an unlawful solicitation or an investment advice.

Amundi accepts no liability whatsoever, whether direct or indirect, that may arise from the use of information contained in this material. Amundi can in no way be held responsible for any decision or investment made on the basis of information contained in this material. The information contained in this document is disclosed to you on a confidential basis and shall not be copied, reproduced, modified, translated or distributed without the prior written approval of Amundi, to any third person or entity in any country or jurisdiction which would subject Amundi or any of “the Funds”, to any registration requirements within these jurisdictions or where it might be considered as unlawful. Accordingly, this material is for distribution solely in jurisdictions where permitted and to persons who may receive it without breaching applicable legal or regulatory requirements.

The information contained in this document is deemed accurate as at the date of publication set out on the first page of this document. Data, opinions and estimates may be changed without notice. These data have been analyzed anonymously for scientific, statistical or historical research purposes.

You have the right to receive information about the personal information we hold on you. You can obtain a copy of the information we hold on you by sending an email to info@amundi.com. If you are concerned that any of the information we hold on you is incorrect, please contact us at info@amundi.com

Document issued by Amundi Asset Management, “société par actions simplifiée”- SAS with a capital of €1,086,262,605 - Portfolio manager regulated by the AMF under number GP04000036 - Head office: 90 boulevard Pasteur - 75015 Paris - France - 437 574 452 RCS Paris - www.amundi.com

Photo credit: iStock by Getty Images - monsitj/Sam Edwards