1

Introduction to Part I
Robo-advising as a Technological Platform for Optimization and Recommendations

Lisa L. Huang[a]

It may be a self-evident truth that the financial services industry is driven by data and that data is increasing at an exponential pace. Robo-advisors are technological platforms that help individuals make better financial decisions, i.e., deliver 'advice', at scale using large disparate data sets. Advice may mean anything from investment portfolios, to consumption/savings rates, to financial goals, and to withdrawals in retirement, etc. The prefix 'robo' reflects the fact that advice is given most often algorithmically. This does not mean that there is not a human in the loop, and this is most often the case currently. Robo also implicitly means that advice can be delivered at scale. With this scale, the cost of advising can be lowered substantially, which leads naturally to the democratization of financial advice. With this platform, it's not hard to imagine a world where there is universal access to financial services which breaks down traditional economic, social, gender, and geographical barriers.

My own work helping to build one of the first robo-advisors in the world began in 2012 when I first learned of the mission that Betterment was founded upon. I joined Betterment the following year and built many of the foundational algorithms that deliver financial advice at scale to the many users on its platform during my years there.

The robo-advisor market is enormous, not measured in hundreds of billions, but in trillions of dollars. The robo-advisor market is also global, because the need to access financial services at scale is becoming more critical across the world. At the inception of robo-advising, advice was limited in scope to investment and portfolio management. Indexing can be seen as one of the first examples of robo-advice, which provided a ubiquitous and low-cost way for individuals to invest. The first wave of robo-advisors typically used mathematically sophisticated portfolio optimization tools, such as Modern Portfolio Theory (or extensions of it such as Black–Litterman), to create semi-customized solutions for retail investors. These tools for portfolio optimization were well known but not democratized at a cost that was accessible to the masses. These first robo-advisors helped change

[a] Head of AI Investment Management and Planning, Fidelity
Published in *Machine Learning And Data Sciences For Financial Markets*, Agostino Capponi and Charles-Albert Lehalle © 2023 Cambridge University Press.

that and indeed transformed an entire industry. In their chapter *Robo-Advisory: From Investing Principles and Algorithms to Future Development*, Adam Grealish and Petter N. Kolm give the readers a fantastic insider's view of the detailed blueprint of the traditional robo-advisor. What is striking is the simplicity and the elegance of advice when it is guided by a set of principles, as explained by the authors.

However, traditional robo-advising – that is rooted in investment management – is evolving. Since robo-advisors are, at their core, a technological platform for financial services, the scope of what can be achieved on that platform can broaden substantially. Robo-advisors as a platform can educate their users, correct for human bias, help to conceptualize the entire financial life cycle of an individual, and optimize every financial decision to maximize the 'happiness' of the user. The meaning of 'happiness' is a personal one but can in theory be captured algorithmically. With enough data, and allowing for feedback between users and algorithms, robo-advisors have opportunities to help users optimize every personal financial decision. In *New Frontiers of Robo-Advising: Consumption, Savings, Debt Management, and Taxes*, Francesco D'Acunto and Alberto G. Rossi outline the tantalising vision of the 'holistic robo-advisor'.

There are incredible challenges around realizing the full potential of robo-advisors. The most critical is data that gives a complete and holistic view of the financial life of a user. If partial data is available to the robo-advisor, then the algorithms will not be able to come up with the most optimal solutions for the user. Most users have a variety of financial relationships with different financial institutions. For example, they may have multiple bank accounts, brokerage accounts, retirement accounts, etc. Therefore, seeing a holistic picture is often non-trivial.

While the holy grail of robo-advising is personalization, the measurement of personal parameters that are needed for the robo-advisor is potentially fraught with uncertainty. High uncertainty in input parameters will lead to suboptimal outputs. One of these inputs is the 'risk tolerance' parameter. Loosely speaking, risk tolerance is a measure of the attitude toward investment risk. Different robo-advisors will try to access this number in different ways but most use a questionnaire to collect data from users. This is clearly insufficient because the definition of risk tolerance is unclear to begin with. It could be very customized for each user. In some implementations of robo investment advice, this risk number directly maps to a portfolio. Since the measurement of risk tolerance is imprecise, optimization of the portfolio only leads to a false sense of precision.

The last challenge that I will highlight here is a technical one. Many tasks that are universal in the financial lives of users do not yet have an accepted mathematical solution. One such problem that I helped solve during my time at Betterment was how to optimize the location of assets in a multi-account setting, in order to minimize taxes, given different tax treatments across multiple accounts. Surprisingly, the exact mathematical solution for this was not known when we began the work. We eventually solved this problem by mapping the asset location to the mathematical problem called the knapsack problem. However, the knapsack problem only solves the static allocation problem, but not the dynamic one, which is driven by any cash flow into accounts. The dynamic knapsack problem was one of many unsolved problems in financial planning. Another examples is finding the optimal way to save, given a multi-account setting with different risk tolerances for each account and different horizons with different priorities across those accounts? Most often, heuristics are relied upon to solve these mathematically complex problems. Milo Bianchi and Marie Briere, in their chapter, *Robo-Advising: Less AI and More XAI?*, delve into the nuanced nature of algorithmic advice and explore the challenges of how to generate trust in robo-advisors.

Since robo-advising is a technological platform, the users can be retail or institutional investors. In *Recommender Systems for Corporate Bond Trading*, Dominic Wright, Artur Henrykoswki, Jacky Lee and Luca Capriotti have created an application of robo-advising for corporate bond trading which leverages the recommender algorithms that are ubiquitous in retail businesses like Netflix and Facebook.

I will end here by referencing the title of a chapter, called called the *Investor's Worst Enemy*, from Ashwin B. Chhabra's book *The Aspirational Investor* (2015). This enemy, as many have pointed out, is the investor themselves. The promise of the robo-advisor is that the technology platform will help conquer the investor's worst enemy, to improve their financial decisions, and in turn, their lives.

2

New Frontiers of Robo-Advising: Consumption, Saving, Debt Management, and Taxes

Francesco D'Acunto[a] and Alberto G. Rossi[b]

Abstract

Traditional forms of robo-advice were targeted to help individuals make portfolio allocation decisions. Based on the balance-sheet view of households, the scope for robo-advising has been expanding to many other personal-finance choices, such as households' saving and consumption decisions, debt management, mortgage uptake, tax management, and lending. This sub-chapter reviews existing research on these new functions of robo-advising with a special emphasis on the questions that are still open for researchers across several disciplines. We also discuss the attempts to optimize jointly all personal-finance decisions, which we term "Holistic Robo-Advisors." We conclude by assessing fruitful avenues for research and practice in finance, computer science, marketing, decision science, information systems, law, and sociology.

2.1 Robo-advice and the balance-sheet view of the household

Robo-advice is any form of financial advice provided to human decision makers by algorithms. Even though many early applications of robo-advice were concentrated in the context of helping individual investors make portfolio allocation decisions, no inherent characteristic of algorithmic advice limits its application to that narrowly-specified context. And, indeed, the scope of robo-advice has broadened dramatically across all the areas of personal finance and more broadly to all contexts in which inexpert and often financially illiterate consumers need to make important choices that will affect their life-time wealth.

The breadth of applications of robo-advising are defined through the lens of the "balance-sheet view" of the household, which we depict schematically in Figure 2.1.

Under the balance-sheet view, households run dynamic budgets similar to those of firms: households have assets (left-hand side of Fig. 2.1), which include housing, durable goods, human capital, financial investments, and health. Households

[a] Carroll School of Management, Boston College
[b] McDonough School of Business, Georgetown University
Published in *Machine Learning And Data Sciences For Financial Markets*, Agostino Capponi and Charles-Albert Lehalle © 2023 Cambridge University Press.

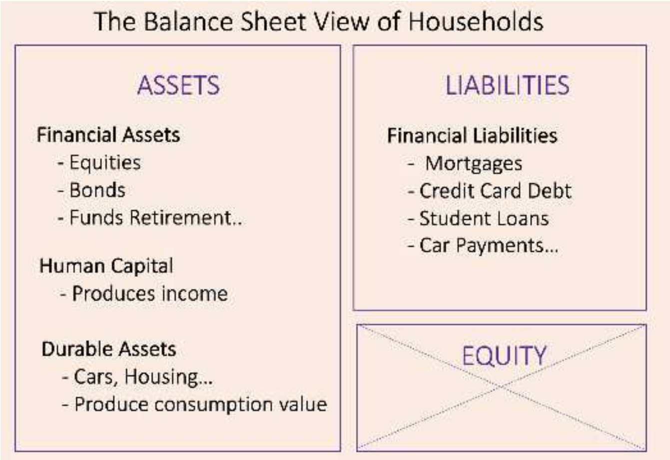

Figure 2.1 Balance Sheet View of the Household

also have to finance liabilities such as mortgages, credit-card debt, student loans, taxes, and insurance premiums. Households need to make decisions about all these budgetary items throughout their lifetime. Many such decisions will have enormous implications for their long-run wealth and financial sustainability.

In contrast to firms, however, the typical household lacks the knowledge and experience needed to make such important choices. For instance, many households only make a decision about purchasing a house and hence borrowing money through mortgages once in a lifetime. Moreover, households usually only face the problem of which form of education to provide to their offspring and how to finance such education once per child. The disconnect between the importance of all these decisions for household budgets and the lack of knowledge and experience in making such decisions stresses the need and scope for advice. Indeed, there is a large literature showing that, when left to their own devices, households make significant and costly mistakes that limit their ability to accumulate wealth over time (see Odean, 1999, Agarwal et al., 2017, and Laibson et al., 1998).

Despite their limitations as economic decision makers, households still need to make decisions that shape their balance sheets both statically and dynamically. For instance, how much and what type of human capital to acquire. Or, what kind of durable goods to purchase – what car to use and what housing condition to live in. All these asset purchases have dramatic implications on the liability side, too. For example, car purchases or leases involve choosing only one out of the very many financing solutions and contracts available. The choice of acquiring human capital – obtaining college and/or graduate education – involves decisions on the ways in which such asset acquisition can be financed, for instance choosing appropriate student loan conditions or even planning on college funds many years before the offspring reaches college age. Also, think about what is possibly the most important choice households make, i.e. the purchase of a house, which requires choosing appropriate mortgage characteristics based on household members income paths and horizons, a decision-making problem under risk and uncertainty that is incredibly hard to solve even for experts.

Historically, whenever choosing how to manage their balance sheets, households had the option of hiring human advisors. This option is less than desirable, however. First, human advisors are relatively costly and have been shown to make suboptimal choices. Suboptimal choices could be due to conflicts of interest in principal-agent relationships with asymmetric information, such as advisors' incentive to propose high-fee financial products to their clients, who are often unaware of the differences across financial products. Behavioral and cognitive biases could drive suboptimal human-advisor decisions as well (Foerster et al., 2017; Linnainmaa et al., 2021). By relying on human advisors, households face at the same time a potentially high cost of advice paired with an often suboptimal quality of advice. Second, supply-side forces might also restrain the availability of human advice to households and especially to lower-income households, who tend to be the most vulnerable when making decisions about managing their balance sheets. Catering to individuals with low net worth might be unpalatable to human advisors due to the low prospective revenues such clients would generate over time (Reher and Sokolinski, 2020).

These severe limitations of human advice in a context in which potential advisees often lack the ability to understand, let alone solve, the decision-making problems they face has represented fertile ground for the swift diffusion of robo-advice, also known as algorithmic advice (see D'Acunto and Rossi, 2021, Rossi and Utkus, 2020a). Robo-advice eliminates the barriers to access advice represented by the cost of human advisers because, in contrast to human advisers, it can be scaled up without virtually any constraints. For this reason, providers of robo-advising services can reduce their fees to a fraction of those commanded by human advisers. Moreover, robo-advice has been shown to make better decisions than humans and experts in several contexts on both the assets and liabilities side of the household balance sheet, such as the allocation of financial investments (e.g., see D'Acunto et al., 2019f; Rossi and Utkus, 2020b) or the take-up of peer-to-peer (P2P) loans (e.g., see D'Acunto et al., 2020a).

In the rest of this chapter, we highlight important recent developments in the evolution of robo-advising services based on the balance-sheet view of the household. We discuss the institutional details of each form of advice as well as the findings of existing research on the characteristics and performance of robo-advice across various domains. In particular, we focus on robo-advice in the domains of households' consumption and savings decisions, borrowing decisions, tax management, and lending choices. Robo-advisors for lending choices are allowing consumers and households who need financing to obtain funds without the need to pay fees to intermediaries. Moreover, they allow households to use their own savings to finance other borrowers and hence reduce the scope for institutional financial intermediaries. For each area, we discuss open questions and opportunities for researchers. We then envision the possibility of forms of robo-advice that optimize jointly households' choices subject to their budget constraint across all the individual parts of households' balance sheets. We term these forms of robo-advice "Holistic Robo-Advisors." Throughout the subchapter, we discuss the challenges and opportunities these recent forms of robo-advice imply and how these challenges and opportunities can translate into fruitful avenues of future research for scholars in as disparate fields as finance, computer science, marketing, decision science, information systems, and sociology.

2.2 Robo-advising for consumption-saving choices

A fundamental factor that determines a household's ability to accumulate wealth throughout the life cycle is the choice of how much to consume and save out of household income in each period in which income is earned. Computing the optimal saving rate requires solving a complicated optimization problem (D'Acunto et al., 2019a) that can prove challenging even for experienced economists. Non-economists are at a further disadvantage, because they often lack a clear understanding of the status of their finances, they cannot assess their own budget constraints, and they do not understand the implications of macroeconomic shocks for their individual consumption-saving decisions (see Agarwal et al., 2009; Agarwal and Mazumder, 2013; Christelis et al., 2010; D'Acunto et al., 2019d). Most households may find it hard to merely conceptualize this problem, even intuitively (see D'Acunto et al., 2019e), let alone to assess the optimal behavior throughout the life-cycle path and subject to budget constraints.

And, indeed, unsurprisingly many households fail to choose a saving rate during their working years that allows them to maintain a lifestyle comparable to the one they enjoyed before retirement (e.g., see Banks et al., 1998; Bernheim et al., 2001; Lusardi and Mitchell, 2007, among many others). This phenomenon represents not only a problem for individual households, but also produces negative externalities for society as a whole as the average tax payer needs to contribute higher taxes to maintain minimal living standards for the undersavers.

Even if potentially less problematic under the societal point of view, the opposite mistake in households' consumption-saving choices has also been detected: several US and European households tend to save large amounts based on perceived rather than actual precuationary savings motives (D'Acunto et al., 2020b). This phenomenon has been detected even during retirement – the phase of their life in which they should be engaging in the process of "decumulation" (See Mitchell and Utkus, 2004) – even when bequest motives are absent. Households' use of rules of thumb based on cultural norms, which substitute for their inability to understand and solve the dynamic optimization problem, have been proposed to explain this type of decisions (e.g., see D'Acunto, 2015). Households' consumption-saving choices are also at the heart of the balance-sheet view of the household discussed above, because the allocation of income across these two alternative types of assets has substantial dynamic implications in terms of long-run net worth.

Pairing the importance of the consumption-saving choice for individual households with the widespread inability of households to conceptualize and optimize such choice represents fertile ground for robo-advising applications. In this context, robo-advising applications might solve two different types of needs. First, they should provide households with information about their own balance sheet, size of assets and liabilities, and budget constraints, in a unified and simple format so that households can understand the parameters of the decision-making problem they face. This information role of robo-advising is especially important for households who have irregular income inflows or those who are self-employed and business owners, and hence whose income streams are irregular and not always easy to forecast.

Second, robo-advising applications to consumption-saving decisions should provide suggestions and advice to households on how to improve their choices as well as easy implementation of such advice. Suggestions can cover several aspects of decision-making such as the choice of which credit card(s) to use, which share of income to save each month based on projections of future values of saved amounts, as well as potential nudges to increase households' incentives to save rather than spend, which would be especially helpful for households who tend to spend more than what the permanent-income hypothesis implies at each point in time.

Real-world applications of robo-advising to the consumption-saving choice based on the criteria discussed above abound. In particular, one class or robo-advisors known as "income aggregators" fulfils this role (e.g., see Olafsson and Pagel, 2017, 2018). As the name suggests, income aggregators are a class of robo-advisors that covers the first scope of robo-advsing in the consumption-saving choice, i.e. providing households with clear and easy-to-grasp information about their own balance sheet and constraints.

Income aggregators require users to provide access to their asset and liability accounts. Asset accounts might include checking, saving, and other forms of financial investment accounts, such as brokerage accounts and retirement accounts. Liability accounts include mortgages, student debt, credit cards, and other forms of debt. In this way, robo-advisors collect information from the households' accounts, typically at the level of the individual transaction. By collecting this large amount of big data across accounts that would otherwise be unlinked, income aggregators are able to construct the balance sheet of each household following the balance-sheet view of the household discussed above. The accuracy of the information income aggregators produce depends on whether users link all their accounts to the robo-advising platform. For this reason, users have a strong incentive to link all their accounts.

The information income aggregators produce has a set of unique characteristics. First of all, income aggregators provide a just-in-time holistic representation of an household's balance sheet, which the household can check at any point in time. This feature is especially compelling for households who have substantial wealth invested in financial markets, the volatility of whose returns might be high. Moreover, income aggregators display information about households' balance sheet and budget constraints vividly in simple graphical forms that are intuitive for households and allow them to grasp basic concepts of household finance even without being trained, such as the balancing of budgets or the sustainability of assets and liabilities accounts. Having access to such intuitive display of information about one's own finances is crucial to create awareness in investors' mind and was shown to have a major impact in helping individuals make better financial decisions (Olafsson and Pagel, 2017, 2018).

However, advising individuals on how much to consume, what items to purchase, and how to split spending between durable and non-durable consumption is more complicated than helping individuals form well-diversified investment portfolios, because an algorithm would need to input specific information regarding individuals' preferences over all possible consumption bundles as well as their beliefs about a large range of future outcomes.

To overcome these limitations, innovative FinTech Apps have proposed alternative ways to help individuals by providing them with simple rules of thumb. A recent example is the US application *Status Money*. Status Money is an income aggregator, and hence as discussed above can compute users' net worth and observe all their transactions, including spending transactions. The unique feature of this App, which provides advice in the form of a rule of thumb, is providing users with information about peers' spending, where peers are defined as individuals observed in a US-representative sample outside the App and who are similar to users based on a set of demographic characteristics. Upon subscribing to the App, users fill in a form about demographic characteristics that include their annual income, age, home-ownership status, location of residence, and location type.

Based on this information Status Money assigns a peer group to each users and provides users with information about the average spending, assets, debts, and net worth of such peers. In this way, users can calibrate their spending to the spending of individuals who look similar to them. This rule of thumb is based on the notion of *the wisdom of the crowd*, whereby agents might obtain valuable signals about their (unknown to the user and to the robo-advisor) optimal spending and saving rate based on the average values of these ratios in a large population of decision makers that look similar to them (Chen et al., 2014; Da and Huang, 2020). Delivering information about crowds through media outlets has been shown to be effective in persuading consumers to change their behavior through the management of their subjective beliefs (Barone et al., 2015). Another channel behind this form of advice is *peer pressure*, whereby it is especially those users who spend substantially more than their peers – and hence are likely to spend more than their own optimal rate – who feel more compelled to react to the peer information and converge to peers' spending than those who spend less than their peers (Rosenberg, 2011). This potential asymmetric reaction to peers' spending information based on users' position relative to their peers would be valuable because overspending, and hence accumulating fewer savings and lower wealth for retirement, is a mistake that creates more issues for individual households and society than underspending.

D'Acunto et al. (2019b) study the effectiveness and the mechanisms behind this form of robo-advice. They find that providing salient peer information through the App has a large effect on users' consumption behavior. Users who were overspending with respect to their peer group at the time of sign-up ended up reducing their spending after signing up for the App. Those individuals who underspent instead, increased their spending but the reaction was much less pronounced for underspenders. D'Acunto et al. (2019b) also show that the informativeness of the peer group plays an important role in explaining users' changes in consumption. The authors conclude that FinTech Apps can provide valuable advice to individuals by collecting and summarizing in an unbiased fashion the decisions made by others and exploiting mechanisms such as the *the wisdom of the crowd* and *peer pressure*.

Another form in which income aggregators provide robo-advice for spending and saving decisions is through nudges, which are based on App notifications and reminders (Acquisti et al., 2017). Notifications and reminders from Apps are becoming ubiquitous and have proven useful in motivating individuals to stay active and eat healthy, among other outcomes. In the context of income aggregators, recent studies have documented the importance and effectiveness of these notifications. For example, Lee (2019) studies individuals' responses to overspending alerts, which are based on the robo-advising algorithm of an income aggregator that compares a users' own spending over time and identifies unusual patterns of spending within their spending history. Lee (2019) finds that users who receive overspending alerts reduce their spending 5.4% more than users who do not receive them. These changes in spending affect long-run cumulative spending. Lee (2019) also finds that the effect of nudges vary across the user population, with older, more financially-savvy, and more educated users adjusting their spending more after receiving overspending notifications, which suggests that more sophisticated users, rather than the least sophisticated, find notifications about their own unusual spending patterns useful. This result encourages further research on how robo-advising could be used to reach to the least sophisticated parts of the population, whose consumption, saving, and education choices tend to be stickier over time than those of the highly educated (D'Acunto, 2014).

Whereas the robo-advising income aggregators discussed so far provide advice on users' spending decisions, another class of robo-advisors target users' saving choices. Consumption and saving choices are obviously strongly interlinked, but the principles extant robo-advisors use to provide advice on these two dimensions are quite different. For example, Apps such as Acorn in the US and Gimme5 in Italy provide robo-advice to their users by helping them to set saving goals and reach such goals using nudges (Gargano and Rossi, 2020).

Goal setting exploits a behavioral mechanism that is not contemplated in standard life-cycle consumption-saving models. According to such models, agents should care about their overall savings but not about the specific objectives for which a certain amount is saved. This is because, for the most part, savings are fungible – they can be used for any purpose at any time (Browning and Crossley, 2001). However, setting budgets and goals is a common feature of agents' daily life, because as a large literature in experimental social psychology shows, agents are intrinsically motivated by goals and work hard to achieve them (Locke and Latham, 1991, 2002, 2006).

Using data from the robo-advisor for saving Gimme5, Gargano and Rossi