Individualization of Robo-Advice

Michael Faloon and Bernd Scherer

Michael Faloon
is chief operating officer at Standish Mellon Asset Management in Boston, MA.
mfaloon@standish.com

Bernd Scherer
is chief investment architect at Deutsche Asset Management in Frankfurt, Germany, and a research officer at EDHEC Risk in Nice, France.
drberndscherer@gmx.net

Automated asset management advisory firms, often called *robo-advisors*, assign risky portfolios to individual investors based on investment algorithms. These algorithms use investor characteristics such as age, net income, and assessments of individual risk aversion to recommend suitable asset allocations. Client interaction and delivery of portfolio advice are web-based and without human interaction. Robo-advice disintermediates the classical distribution model, which is now widely recognized as expensive, difficult to scale, and unacceptably heterogeneous (i.e., highly dependent on the individual advisor's skill level). Two catalysts are at work. First, lower interest rates make it increasingly difficult to justify high fee levels. Second, we see a shift from defined-benefit to defined-contribution plans around the world (a result of poor risk sharing between corporates and pension fund beneficiaries). Individual investors need to take responsibility for their own investments, and those with lower levels of wealth that do not meet minimum account limits are left stranded with little access to financial planning advice. This is where the robo-advisor comes in.

Sharpe [2008] outlined four guiding principles for investment advice: diversify (investments in a broad universe of assets spanning all available sources of excess returns), economize (awareness of transaction costs and fee layers), contextualize (allocation conditional on varying investment opportunity set), and finally, personalize (model household balance sheets and individual preferences). Following Sharpe's taxonomy, the main unifying value proposition of current robo-advisors is to provide cheap access to diversified beta. This is not surprising, given that most robo-advisors are deeply convinced of the advantages of passive investing. A focus on diversified exchange-traded fund (ETF) portfolios reduces client costs and has strong empirical and theoretical backing. It also keeps headcounts down on the robo-advisor side: No analysts are needed to screen actively managed funds or provide forecasts of time-varying investment opportunity sets. Whether this view is truly fiduciary or just driven by business economics is difficult to assess. The careful observer will note, however, there is a prevalence of passive investing among virtually all robo-advisors. The breadth of ETFs in terms of spanning available risk premiums differs among robo-advisors, and the number of ETFs that enter a given client portfolio is, on average, larger than what two-fund separation would suggest. Although the marginal costs of adding ETFs is low for the robo-advisor, a large number of ETFs increases the complexity costs and effort of replicating robo-advice. This stops clients of robo-advisors from replicating portfolio advice on a fraction of personal wealth shown to the robo-advisor, while replicating advice on the bulk of wealth free of percentage fees.

QUESTIONNAIRES

Robo-advisors raise the necessary inputs required for their investment algorithms via web-based questionnaires that users of their services need to answer. In this section, we present a set of questions in common use by most robo-advisors. This allows us to define a generic robo-advisor that we will use to describe the current state of robo-advice, rather than singling out a single supplier or getting lost in the idiosyncratic variations of an evolving industry. Without apology, we focus on the following set of questions:

1. Do you invest for retirement or to generate general savings?
2. What is your age?
3. What is your net income after taxes?
4. What is your savings rate?
5. What is the value of your current (liquid) investments?
6. When deciding how to invest your money, do you worry more about maximizing gains, minimizing losses, or both equally?
7. If your investment portfolio lost x% in a given month, would you liquidate your portfolio, just sell some investments, do nothing, or increase investments?

We can group these questions into three blocks: investment objectives (question 1) to decide on asset only or asset (retirement) liability optimizations; the ability to take risks (questions 2–5) to decide on risk capacity (economic suitability); and the willingness to take risks (questions 6 and 7) to measure how aggressive portfolios should become.

FROM QUESTIONNAIRES TO PORTFOLIOS

How is the preceding information used by robo-advisors? Two alternative methods exist. The first method is very ad hoc and driven more by lawyers trying to anticipate court rulings than by modeling client economics. In this approach, questions 1–7 (possibly supplemented by questions on investor experience and time horizon) are used in a scoring model that maps scores to points on a given efficient frontier. Scores and their weightings are invented rather than derived from an academically cross-validated decision-making model. Whoever determines the scoring model and its mapping has the biggest impact on client performance without being held accountable. It would be very difficult to argue for the validity of the scoring process relative to a model with strong decision theoretic foundations. The advantage of this model, however, is that it can include information provided by the legal department. If courts have ruled that investors with zero investment experience should hold a very low-risk portfolio, the algorithm could ensure clients without experience get low-risk portfolios by either assigning a large enough score to this answer to dominate all other scores or simply making this an overruling score.

The second method is in line with decision theoretic models on portfolio choice. These models typically include household assets and liabilities. From an investor's age, net income after taxes, and savings (questions 2–4), we calculate the present value of lifetime savings (human capital). All we need as additional information is a lifetime career path (e.g., from panel survey data), the retirement age, and an assumption on the nature of the investor's human capital. Typically, one would regard human capital as a mixture of equity and bond exposure, depending on the investor's profession (mostly fixed income for civil servants and mostly equities for the self-employed, as extremes). The nature of human capital influences both its covariance with financial assets as well as the discount rate used to derive its present value. The present value of future savings then sits as an exogenously given and nontradable shadow asset on the household balance sheet. Rather than directly asking for the savings rate, we could calibrate typical saving rates for different household types (double income investors without children will display a larger savings rate than single parents with four children). Some advisors opt for this form of householding instead.

Apart from shadow assets, the questionnaire also provides us with the investor's current investments, which we could use as financial assets. But there is more. With a projected lifetime career path, we have already estimated the investor's likely preretirement income. Assuming a typical replacement rate of 80% (the investor requires 80% of his or her preretirement income for consumption), we can calculate the present value of the investor's pension gap as a liability asset. From this information, we can build a simple ongoing concern household balance sheet for a household that is average in every respect apart from age and after-tax income. We show such a balance sheet in Exhibit 1. For simplicity, we normalize our balance sheet to unit length.

Assuming standard mean–variance preferences, in which risk information (from the answers in questions 6 and 7) is mapped into a numeric value for risk aversion, (see section on risk individualization), we arrive at the following solution for household portfolio choice:

$$w = \left(\frac{1}{\theta}\right)\lambda^{-1}\,\Omega_{aa}^{-1}\,\mu_a + \left(1-\frac{1}{\theta}\right)\Omega_{aa}^{-1}\,\Omega_{as} + \left(\frac{1}{\theta}\right)\left(\frac{1}{f}\right)\Omega_{aa}^{-1}\,\Omega_{al} \tag{1}$$

Exhibit 1
Generic Household Balance Sheet

Assets Total assets = Financial assets + Shadow assets	Liabilities
$\theta = \dfrac{\text{Financial assets}}{\text{Total assets}}$	$1 - \dfrac{1}{f}$
$1 - \theta$	$\dfrac{1}{f} = \dfrac{\text{Retirement liabilities}}{\text{Total assets}}$

Notes: The left side of the economic household balance sheet (assets) consists of financial assets (current investments from question 5) and shadow assets (human capital, derived from questions 2–4). We assume that shadow assets are exogenously given. The right side of the balance sheet (liabilities) consists of the retirement liabilities (discounted pension gap derived from questions 2–4 with the additional assumption of a replacement rate) and residual net wealth (household equity). All balances entries are rescaled to yield a balance sheet length of one; that is, all entries represent percentages rather than absolute values.

From Equation (1) we can infer that total demand for risky assets, w, can be decomposed into three portfolios plus cash (four-fund separation). Each portfolio represents a different form of demand given in Exhibit 2.

What implications arise from Exhibit 2? First, shadow assets create investors who are generally richer than their financial wealth alone suggests. Speculative demand therefore needs to be larger than in the case of a financial asset-only optimization. Although the direct effect of including shadow assets always makes investors more aggressive, the indirect effect on total demand for risky assets will depend on the covariance of shadow assets with financial assets. If financial assets co-vary strongly with shadow assets, they become unattractive as the volatility of net wealth increases for a nondiversifying financial asset.

Second, we see that the more leveraged a household becomes (smaller f), the less aggressively it should invest as the volatility of net wealth increases and liability hedging demand grows. If f were very large (assets vastly exceed liabilities), liabilities would play a minor role. More generally, particular securities are held if they show attractive (standalone) risk–reward tradeoffs, if they help to reduce fluctuations in shadow assets (negative covariance with shadow assets), or if they help hedge liability-related risks (positive covariance with liabilities). Although $\frac{1}{\theta}$ and $1-\frac{1}{\theta}$ always add to 1, this does not mean that assets add up to 1. We will always need cash (long or short) to complete the portfolio. For investors who are mainly interested in accumulating general savings (question 1), we can set $f = \infty$ to remove the liability hedging term from our solution. The same

Exhibit 2
Four-Fund Separation

Demand	Leverage	Portfolio	Inputs
Speculative	$\frac{1}{\theta}$	$\lambda^{-1}\Omega_{aa}^{-1}\mu_a$	μ_a (excess returns), λ (risk aversion) Ω_{aa} (covariance of excess returns)
Diversification	$1-\frac{1}{\theta}$	$\Omega_{aa}^{-1}\,\Omega_{as}$	Ω_{as} (covariance with shadow assets) $\Omega_{aa}^{-1}\,\Omega_{ao}$ (shadow asset beta)
Hedging	$\frac{1}{\theta}\,\frac{1}{f}$	$\Omega_{aa}^{-1}\,\Omega_{al}$	Ω_{as} (covariance with retirement liabilities) $\Omega_{aa}^{-1}\,\Omega_{al}$ (retirement liability beta)

Note: Speculative demand, diversification demand, and hedging demand for investors with balance sheets according to Exhibit 1.

robo-advisor could also (given the information raised via its questionnaire) project the evolution in θ and f over time to arrive at a generic glide path representing portfolio allocations as a function of time. Inputs (asset means and covariances) into Equation (1) are usually derived from historical data, in which expected returns are calculated from reverse optimization, that is, from asking the question: What are the expected returns that would make a given observed portfolio (typically a cap-weighted market portfolio) optimal? This is sometimes connected wrongly to Bayesian models of portfolio choice, but rather represents a mathematical identity. Optimized portfolios are typically rebalanced using trigger rebalancing. As market performance affects a given portfolio, rebalancing takes place only when the actual portfolio differs too much in risk or security equivalent from the advised portfolio.

In summary, the portfolio advice provided by our generic robo-advisor is suitable for investors who are average in every aspect of their economic lives apart from age, net income, and current investments.

ECONOMIC SUITABILITY?

The previously described generic robo-advisor provides canned advice. Personalized information about the individual investor is not raised. A vast number of investors differ materially in their ability to take risks, although they share the same age, net income, and savings rate. Portfolio proposals based on such limited information in our opinion neither meet economic suitability criteria nor represent advice that is in the best interest of the client. To support this claim, we make a few observations.

First, the nature of human capital assumed previously represents the average investor only. Easily calibrated differences for the self-employed, blue-collar workers, white-collar workers, or civil servants—with different earning profiles and typically different asset and liability structures each—are not taken into account. Neither will the industry or sector (private or public) in which a potential investor is employed matter to our generic advisor. A young banker is already considerably exposed to equities (potentially value stocks), whereas a young civil servant is not. In other words, human capital is a highly individual asset on each investor's balance sheet. This necessitates individualization.

Second, investors are typically characterized by many more assets and liabilities on their economic balance sheets. The most obvious example is leveraged real estate, which represents a typical background risk. It will reduce the investor's appetite for risky assets, let alone further investments into real estate. This is an important driver of asset allocation over time because younger households typically hold larger fractions of leveraged real estate (relative to their total wealth), whereas older households already have paid off large parts of their mortgage.

Third, robo-advisors make little attempt to better understand the sources of retirement income (payoffs from defined-benefit or defined-contribution plans, life insurance, etc.). Incorporating all these balance sheet items requires us to use a much richer covariance matrix. This leads to a vast set of solutions that would no longer lie on a single efficient frontier. The current practice of offering precalculated portfolios in web portals (to avoid delays in presenting solutions that are created in real time) would cease to be practical, and the number of portfolio solutions would become very large.

Fourth, rebalancing is realistically driven by more than relative asset performance. Balance sheet changes through time (aging) or lifetime events have a major impact on optimal portfolios.

Fifth, separate mortality tables for men and women would allow further individualization. This would create even more realistic allocations because differences in longevity need to be reflected in investor portfolios.

In summary, we believe that generic robo-advice is unsatisfactory because of an obvious lack of adequately modeling the investor's risk-taking capacity. We believe that individualization will become a main differentiator among advisors.

RISK INDIVIDUALIZATION?

So far we have talked about the lack of individualization in deriving the client's risk capacity, but our generic robo-advisor also takes a very casual look at how it derives the client's willingness to take on risks (broadly described as *risk aversion*). This is surprising, given that risk aversion is the very input variable with the largest impact on allocations suggested by robo-advisory firms. One difficulty is that risk aversion is a latent (unobservable) variable. We can either look for observable client characteristics that are highly correlated with risk aversion (answers to ad hoc questions or psychometric questionnaires) or derive risk preferences from observable experimental or investment decisions (i.e., backing out risk aversion parameters or utility functions from a decision theoretic framework). Whatever framework is taken, we must ensure that portfolios are optimized using the same decision theoretic framework under which preferences have been derived.

Some firms develop "sophisticated" questionnaires designed by psychologists to uncover individual risk aversion. As a result, investors are assigned a risk index score that resembles a normalized measure (not unlike IQ) across the population of participants. We see two main problems with this approach: validity (does the test score measure what it claims to measure?) and usability (how can we translate the test score into a utility function or even a risk aversion parameter?). Most robo-advisors do not even follow this approach. Instead, they ask ad hoc questions that are assumed (hoped) to be related to risk-taking behavior. Answers to questions 6 and 7 are transformed into an aggregate score and then mapped into a numerical value for risk aversion. This is essentially psychometrics without psychologists and with even larger doubts regarding validity.

One recently developed way forward is to use experience sampling. The main idea is that investors observe via (historical or forward-looking) simulations what could happen to a chosen portfolio over time and how end-of-period wealth can vary in good and adverse environments. Essentially, this is fast forwarding capital markets and might work as a substitute for accumulated market and investment experience. A series of papers (Kaufmann, Weber, and Haisley [2013]; Bachmann, Hens, and Stössel [2014]; Bradbury, Hens, and Zeisberger [2014]) offers vast cross-validated evidence showing that investors who were exposed to experience sampling chose, on average, riskier portfolios with less regret (portfolios are less likely to be revised after an adverse return shock) than those who were exposed to alternative risk profiling methods. This is ideal for digital asset management advice in which no handholding (by client relationship officers) takes place.

COMMON VARIATIONS

The set of generic questions introduced at the beginning of our article describes the modeling efforts in common among most robo-advisors. Here, we will describe the two most common approaches to achieve more individualization: taxation and goalification. *Taxation* is one area of true individualization in robo-advice. This is highly valuable to investors because the timing of individual contributions and redemptions has a large impact on the optimal rebalancing of investor portfolios. *Goalification* creates portfolios dedicated to personalized investment objectives (retirement, college education, new car, new house, etc.). Each of these mental accounts is optimized in isolation with potentially very different risk aversions. Investors are claimed to have difficulties in describing their overall risk aversion but are much better at determining specific goals and their specific risk aversions. The narrative is that investors find it much less important to miss their savings target for their second car than to miss their retirement liabilities. We believe the dangers of such mental accounting are obvious because a very inefficient portfolio can arise upon aggregation.

Suppose the robo-advisor expects a negative return on fixed income (not implausible at the time of writing). Further suppose an investor wants to build two equally sized portfolios: one for retirement purposes, in which he or she feels a burning desire to reach the investment target as surely as possible, and a second savings portfolio in which he or she is willing to take a lot more risk. According to preference, the investor will mainly invest in bonds (large positive liability hedging demand relative to small negative speculative demand due to high risk aversion, as can be seen from Exhibit 2) in his or her retirement portfolio. He or she would also like to short bonds (large negative speculative demand due to negative return expectations paired with low risk aversion) in the speculative portfolio but is not permitted to do because of a no shorting constraint. Goalification would force our investor to hold a large fraction of his or her wealth in bonds (aggregating two long-only portfolios), whereas the joint optimization of these equally sized portfolios would likely result in holding no bonds at all (positive hedging demand in the retirement portfolio and negative speculative demand in the savings portfolio would cancel out). Although goalification is certainly attractive from a marketing perspective, conceptual problems and efficiency losses should not be ignored. The latter become even larger if separate investment universes are used for each investment goal.

HOUSEHOLD BALANCE SHEETS

Individualization of portfolio advice requires us to identify those household characteristics that are likely to

EXHIBIT 3
Extended Household Balance Sheet

Assets	Liabilities
Total assets = Financial assets + Shadow assets + $\sum_{i=1}^{n}$ Outside assets	
Financial assets	Net wealth
Shadow assets	Retirement liabilities
Outside assets$_1$ (e.g., real estate)	Liabilities$_1$ (e.g., mortgage)
Outside assets$_2$ (e.g., art)	Liabilities$_2$ (e.g., education)
Outside assets$_3$ (e.g., human capital)	Liabilities$_3$ (e.g., pension gap)
⋮	⋮

Notes: Assets of the economic household balance sheet (assets) consist of financial assets, shadow assets, and a variety of outside assets. Outside wealth can come in many forms. Examples are real assets (art, sports car collection, real estate) and business stakes (family-owned business). Liabilities for our balance sheet (liabilities) can also consist of retirement liabilities plus different forms of financial liabilities (mortgage, overdraft, consumer credit) or other ongoing concern liabilities (prospective tuition fees, saving goals). Remaining net wealth (also surplus, household equity, or discretionary wealth) is a residual value to equate both sides of the balance sheet.

EXHIBIT 4
Consolidated Household Balance Sheet

Assets	Liabilities
Total assets = Financial assets + Shadow assets + $\sum_{m=1}^{M}$ Outside assets$_m$	
Financial assets	Household equity
$\theta = \frac{\text{Financial assets}}{\text{Total assets}}$	$1 - \frac{1}{f}$
Outside assets and shadow assets	Liabilities
$1 - \theta$	$\frac{1}{f} = \frac{\text{Retirement liabilities} + \sum_{k=1}^{K} \text{Liabilities}_k}{\text{Total assets}}$

result in different investment decisions. What are these characteristics? What makes households different? In our view, household balance sheets create household-specific hedging and diversification demands, determining the household ability to take speculative risks. We can now match the variety of household characteristics with an equal variety in offered portfolio advice. This is a clear break with capital asset pricing model (CAPM)-based two-fund separation, precisely because we introduce investor heterogeneity that is ignored in a CAPM setting. Balance sheets have both a cross-sectional (balance sheets differ between investors of the same age) as well as a time dimension (balance sheets for the same investor differ across age). This creates individualized portfolio advice including individualized glide paths. Consequently, the poor individualization of the described generic robo-advisor can be easily healed by broadening the economic balance sheet to include a much richer set of personalized assets (alternative sources of income as well as retirement income, real estate, etc.) and liabilities, shown in Exhibit 3. We reduce the complexity of the household balance sheet in Exhibit 4 by standardizing the length of our balance sheet to one and consolidating outside assets and financial liabilities into composite assets.

The left side of the extended economic household balance sheet (assets) consists of financial assets, shadow assets, and a variety of outside assets. Outside wealth can come in many forms. Examples are real assets (art, sports car collection, real estate) or business stakes (family-owned business). We assume that outside assets are not subject to decision-making because they are reasonably illiquid, nontradable, or simply not something the household is willing to sell. The right side of the balance sheet (liabilities) can also consist of retirement liabilities plus different forms of financial liabilities (mortgage, overdraft, consumer credit) or other ongoing concern liabilities (prospective tuition fees, investment goals, etc.). In our framework, there is no restriction on the number of subdivisions for outside wealth or liabilities as long as we can find market value estimates for a given asset as well as a (proxy) return series to calculate risk return characteristics. The remaining net wealth (also surplus, household equity, or discretionary wealth) is a residual value to equate both sides of the balance sheet. If investors were to differ only with respect to θ and f, little individualization would be achieved. What opens the robo-advisor to a whole range of new solutions (making advice highly personalized) is client-specific differences in $\Omega_{as,i}$ and $\Omega_{al,i}$. Optimal allocations are now highly investor specific and driven by different balance sheet compositions (differences in the nature of human capital or outside assets) and differences in liabilities.

CONCLUSION

Robo-advisors represent a new, exciting development in the digitalization of asset management. They have every potential to establish a new standard in fiduciary client advice and become the model of choice for passive investors. In their current form, however, the quality of their advice is difficult for clients to assess. Does my robo-advisor offer a diversified set of risk premiums or just traditional beta? Are risk premiums missing? Does the algorithm know my personalized balance sheet? Does it correctly identify my risk aversion? In other words, can the robo-advisor demonstrate the validity of its approach to determine my individual willingness and capacity to take risks? Do changes in inputs also lead to expected changes in allocations—both in size and direction—or are the inputs merely used to know more about the user? There is an obvious difference between asking individualized questions and using the information to model the client's decision-making problem.

In summary, we believe most robo-advisors provide generic and poorly individualized advice. The legal system offers little help because the U.S. Advisors Act does not even impose the minimum information required to provide what is called appropriate advice. Any defense offered by poorly individualized robo-advice offerings that they meet all legal criteria therefore looks self-serving rather than client-serving. Robo-advisors in 2016 look a lot more like the Tin Man in the Wizard of Oz than R2D2 in Star Wars.

REFERENCES

Bachmann, K., T. Hens, and R. Stössel. "Which Measures Predict Risk Taking in a Multi-Stage Controlled Decision Process?" 2014, http://ssrn.com/abstract=2535859.

Bradbury, M., T. Hens, and S. Zeisberger. "Improving Investment Decisions with Simulated Experience." *Review of Finance*, Vol. 19, No. 3 (2014), pp. 1019-1052.

Kaufmann, C., M. Weber, and E.C. Haisley. "The Role of Simulated Experience and Graphical Displays on One's Investment Risk Appetite and Comprehension." *Management Science*, Vol. 59, No. 2 (2013), pp. 323-340.

Sharpe, W.F. *Investors and Markets: Portfolio Choices, Asset Prices and Investment Advice.* Princeton, NJ: Princeton University Press, 2008.