Personalized Robo-Advising:
Enhancing Investment through Client Interaction

Agostino Capponi* Sveinn Ólafsson† Thaleia Zariphopoulou‡

Abstract

Automated investment managers, or robo-advisors, have emerged as an alternative to traditional financial advisors. The viability of robo-advisors crucially depends on their ability to offer personalized financial advice. We introduce a novel human-machine interaction framework, in which the robo-advisor solves an adaptive mean-variance portfolio optimization problem. The risk-return tradeoff dynamically adapts to the client's risk profile, which depends on idiosyncratic characteristics as well as on market performance and varying economic conditions. We characterize the optimal level of personalization in terms of a tradeoff faced by the robo-advisor between receiving client information in a timely manner and mitigating the effect of behavioral biases in the risk profile communicated by the client. We argue that the optimal portfolio's Sharpe ratio and return distribution improve if the robo-advisor counters the client's tendency to reduce portfolio risk during economic contractions, when the market risk-return tradeoff is more favorable.

1 Introduction

Automated investment managers, commonly referred to as robo-advisors, have gained widespread popularity in recent years. The value of assets under management by robo-advisors is the highest in the United States, exceeding $600 billion in 2019. Major robo-advising firms include Vanguard Personal Advisor Services, with about $140 billion of assets under management, Schwab Intelligent Portfolios ($40bn), Wealthfront ($20bn), and Betterment ($18bn). Robo-advisors are also on the rise in other parts of the world, managing over $100 billion in Europe, and exhibiting rapid growth in Asia, with assets under management exceeding $75 billion solely in China (Statista [2020]).

*Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, USA, ac3827@columbia.edu.

†Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, USA, so2570@columbia.edu

‡Departments of Mathematics and IROM, The University of Texas at Austin, and the Oxford-Man Institute, University of Oxford, zariphop@math.utexas.edu

The first robo-advisors were launched in 2008, in the wake of the financial crisis and the ensuing loss of trust in established financial services institutions. This included the two pioneering robo-advising firms, Betterment and Wealthfront, which began offering services formerly considered exclusive to the general public, including individuals who did meet the investment minima of traditional financial advisors. In the years that followed, industry incumbents - such as Vanguard and Charles Schwab - followed suit and began offering their own robo-advising services, taking advantage of their existing customer bases to quickly gain a large market share. With the focus of most robo-advisors being on passive long-term portfolio management, the rise of robo-advising since the financial crisis has also been compounded by the seismic shift towards passive investing and low cost exchange-traded funds.[1]

Robo-advising is a term that encompasses various forms of digital financial advice for investment management and trading. A recent survey paper by D'Acunto and Rossi [2020] distinguishes between digital tools and services that aid in active, short-term trading, with the client actively involved in the strategy implementation, and robo-advisors that focus on long-term passive investing, involving a higher level of delegation. The framework proposed in this paper belongs to the latter category, with the robo-advisor truly performing automated portfolio management, using quantitative methods and algorithms to construct and rebalance the client's portfolio.

The objective of the robo-advisor is to construct and manage a portfolio that is tailored to the dynamic risk profile of the client. The robo-advisor interacts repeatedly with the client to avoid making decisions based on stale information, and thus exposing itself to the risk of not acting in the client's best interest.[2] This stands in contrast to the approaches of existing robo-advising systems, which usually employ static portfolio optimization, with the portfolio reoptimized every time a change is made to the client's risk profile.

In our framework, the robo-advisor allocates the client's wealth between a risk-free asset, and a single risky asset which is representative of the overall stock market. The asset returns are assumed to follow a regime switching model (Hamilton [1989]), i.e., the dynamics of the market securities are modulated by an observable state process which captures the current state of the economy, and is assumed to be a finite state Markov chain. The market risk-return tradeoff, as measured by the market Sharpe ratio, depends on the economic state. Since the seminal work of Fama and French [1989], it has been extensively documented that the market risk-return tradeoff is countercyclical, i.e., higher at business cycle troughs than peaks.

The robo-advisor adopts a multi-period mean-variance investment criterion with a finite investment horizon. We remark that most robo-advising systems employ asset allocation methods based on mean-

[1] See https://www.cnbc.com/2018/09/14/the-trillion-dollar-etf-boom-triggered-by-the-financial-crisis.html

[2] Robo-advisors are fiduciaries under the Investment Advisers Act of 1940, and as such are subject to the duty of acting in the client's best interest.

variance analysis, primarily because of its tractability and the intuitive interpretation of the risk-return tradeoff (see, for instance, Beketov et al. [2018]). A central element in our optimization problem is a *stochastic* risk-return tradeoff coefficient, specifically designed by the robo-advisor to match the dynamic risk preferences of the client. This is in stark contrast to many other risk averse optimization problems, traditionally considered in the literature, where the risk-return tradeoff is assumed to be constant throughout the entire investment horizon.

The client's risk preferences are modelled via a risk aversion process, whose dynamics depend on several factors. First, we allow for a temporal component accounting for age effects, consistent with empirical research that has identified a positive and potentially nonlinear trend in risk aversion as a function of age (see, e.g., Brooks et al. [2018] and Hallahan et al. [2004]). Second, the client's risk aversion is impacted by idiosyncratic events such as a change in the client's disposable income, or an increase in her educational level or financial literacy (see, e.g, Hallahan et al. [2004] and Allgood and Walstad [2016]). These idiosyncratic shocks can also capture changes in unobservable factors driving the client's risk aversion, consistent with empirical studies that have shown a substantial variation in idiosyncratic risk preferences which is unexplained by consumer attributes and demographic characteristics (see, e.g., Guiso and Paiella [2008], Sahm [2012], and Van de Venter et al. [2012]). Third, and most importantly, the client's risk aversion may depend on realized market returns and changes in economic conditions. Indeed, it is well established that risk aversion varies with the business cycle. Specifically, risk aversion is countercyclical, just like the market Sharpe ratio.[3] This implies that at times when the reward per unit risk is high, the client's risk aversion is also high.

The robo-advisor constructs its own model of the client's risk aversion process and uses it as the risk-return tradeoff coefficient in the mean-variance optimization problem. The construction of this process depends on two sources of information. First, the robo-advisor observes both realized market returns and changes in the state of the economy, and continuously updates its model of the client's risk aversion based on this information. Second, at interaction times, the robo-advisor receives information about the idiosyncratic component of the client's risk aversion.[4] As a result, interaction with the client determines the information set of the robo-advisor and plays a key role in portfolio personalization.

We study the portfolio optimization problem solved by the robo-advisor and how it depends on the dynamics of the risk-return tradeoff coefficient. At any time, the optimal risky asset allocation has two

[3]On the theoretical side, the seminal work of Campbell and Cochrane [1999] provides an asset pricing framework that rationalizes these phenomena. See Lettau and Ludvigson [2010] for a survey of the literature and extensive empirical evidence. Recent empirical and experimental evidence on the countercyclical nature of risk aversion has been provided by Bucciol and Miniaci [2018], Cohn et al. [2015], and Guiso et al. [2018], among others.

[4]The majority of robo-advisors elicit the risk preferences of clients by means of online questionnaires (Beketov et al. [2018]). The client is presented with questions regarding, e.g., demographics, investment goals, education and financial literacy, and potential reactions to hypothetical gambles and market events. We refer to Charness et al. [2013] for an outline of the pros and cons of different methods used to assess risk preferences, and Cox and Harrion [2008] for a more comprehensive overview.

components. The first component depends on economic conditions and the risk-return tradeoff coefficient at the *current* time, which incorporates all available information about past market returns, economic states, as well as information communicated by the client. This component parallels the strategy of a myopic investor optimizing a single-period objective function. The second component adjusts this myopic strategy by effectively averaging over *future* paths of both economic conditions and the risk-return tradeoff coefficient implied by the robo-advisor's model.

We study the implications of interaction on portfolio personalization, which reveals the existence of two conflicting forces. First, if the client and the robo-advisor do not interact at all times, then the robo-advisor does not always have access to up-to-date information about the client's risk preferences. Second, because retail investors are subject to behavioral biases, the information communicated by the client at interaction times may not be representative of the client's true risk preferences. Specifically, if recent market returns have exceeded expectations, the client may feel overly exuberant and communicate a risk aversion value that is lower than her real risk aversion. Vice versa, following a market underperformance, the client may feel overly pessimistic and communicate a risk aversion value that is too high. These fluctuations in the risk aversion communicated by the client are representative of common behavioral biases, such as trend-chasing behavioral patterns (see Kahneman and Riepe [1998] for a discussion of such biases and strategies to overcome them). In our framework, the impact of portfolio losses on the risk aversion level communicated by the client is greater than the effect of portfolio gains, and a lower interaction frequency mitigates the effect of the client's biases because, over longer time periods, the average market return becomes closer to its expectation. These behavioral patterns are consistent with the notion of *myopic loss aversion*, introduced in Benartzi and Thaler [1995].

We introduce a measure of portfolio personalization to analyze the aforementioned tradeoff between higher and lower interaction frequencies. This measure is defined in terms of the difference between the client's actual risk aversion and the robo-advisor's model of it. We show that it is minimized by a uniquely determined interaction frequency, which strikes a balance between obtaining information about the client in a timely manner and ensuring that the communicated information is not too affected by recent market fluctuations. Our result is consistent with existing practices, as robo-advising firms encourage clients to refrain from making frequent changes to their risk aversion profiles, and even limit their ability to do so.[5] We also show that the optimal interaction frequency is decreasing in the level of the behavioral biases, i.e., in the client's sensitivity to abnormal market returns, and increasing in the rate at which the idiosyncratic

[5] For instance, Wealthfront provides the following guidelines on its website: "Our software limits our clients to one risk-score change per month. We encourage people who attempt more than three risk-score changes over the course of a year to try another investment manager."

component of the client's risk aversion changes.[6] Therefore, a higher level of personalization is achieved for clients with limited behavioral biases and clients with stable risk preferences.

We analyze the implications on the investment strategy of economic state transitions and the associated changes in the client's risk aversion. With the risky asset modeling a broad stock market index, we consider an economy consisting of two states, one corresponding to times of economic growth and the other corresponding to times of recessions.[7] Such a choice is supported by the National Bureau of Economic Research (NBER), whose methodology splits business cycles into periods of economic expansions and periods of economic contractions.[8] We compare analytically the Sharpe ratios of strategies that invest the same fixed proportion of wealth in the risky asset during times of economic growth, but differ in their risky asset allocation in a state of recession. Because both the client's risk aversion and the market Sharpe ratio are countercyclical, the client is inclined to shift wealth away from the risky asset in a state of recession, i.e., at times when the risk-return tradeoff is favorable and the benefit of investing is greater.

We establish that the Sharpe ratio of the portfolio in general increases if a greater proportion of wealth is allocated to the stock market when the risk-return tradeoff is high. However, we also find that the Sharpe ratio is concave in the change in allocations when the economy moves to a state of recession. In other words, the Sharpe ratio drops more if the allocation is reduced, compared to how much it increases in the opposite scenario. Hence, we observe that by simply rebalancing the portfolio to maintain prespecified weights throughout the business cycle, the portfolio's Sharpe ratio is satisfactory.

The robo-advising framework proposed herein is built on an interaction system between a human (the client) and a machine (the robo-advisor). The analysis of this *human-machine interaction system* presents a novel methodological contribution, as it gives rise to a new and quite nontrivial adaptive control problem in which the controlled system always maintains the same dynamics, but the optimality criterion changes in *real time* in response to data generated by the system. Specifically, it is the dynamics of the risk-return tradeoff coefficient that continuously adapts to incoming information. This contrasts with existing literature in adaptive control (see, for instance, Astrom and Wittenmark [1989]), where the optimality criterion remains the same while the system dynamics adapt to changes in the environment. To the best of our knowledge, this is the first time that such an adaptive control problem has been considered.

In our framework, the interaction times effectively divide the time interval from initiation and until the terminal date into subperiods, triggering a new optimization problem at the beginning of each subperiod.

[6] We do not account for noise in the information communicated by the client. As the level of noise is expected to increase with the interaction frequency, this would have the effect of further increasing the optimal time between consecutive interactions.

[7] For univariate stock market returns, two economic regimes are commonly identified. See, e.g., Guidolin and Timmermann [2006], who also show that modeling of the joint distribution of multiple asset classes may require a greater number of regimes.

[8] The work of Chauvet and Hamilton [2006] shows that a Markov switching model successfully identifies NBER's business cycle turning points.

This yields a time-inconsistent sequence of problems that are interlinked, because they all share the same terminal date, and because the initial condition of each problem is determined by the state of the system at the end of the previous subperiod. Notably, at each time, the optimal control depends on the future evolution of the risk-return tradeoff coefficient, until the terminal date, which in turn depends on the system dynamics and the frequency of interaction between the human and the machine.

The rest of this paper is organized as follows. In Section 2, we briefly review related literature. In Section 3, we introduce the main components of our modeling framework. In Section 4, we present the solution to the optimal investment problem. In Section 5, we introduce and study performance metrics for the interaction system formed by the client and the robo-advisor. Section 6 offers concluding remarks. Appendix A and Appendix B contain proofs related to the results in Sections 4 and 5, respectively. Appendix C discusses the computational complexity of the algorithm used to compute the optimal investment strategy.

2 Literature Review

From a methodological perspective, our work contributes to the literature on *time-inconsistent* stochastic control (Björk and Murgoci [2013]). Other related works include Li and Ng [2010] who solve a multi-period version of the classical Markowitz problem, and Basak and Chabakauri [2010] who solve a continuous-time version of the same problem. Björk et al. [2014] solve the dynamic mean-variance problem in continuous time, with the mean-variance utility function applied to the return on the client's wealth, consistent with the original single-period mean-variance analysis. A recent study of Dai et al. [2019] further develops a dynamic mean-variance framework based on log-returns, with the aim of generating investing policies conforming with conventional investment wisdom. In all of these works, the risk-return tradeoff is assumed to be constant throughout the investment horizon. By contrast, in our model, the risk-return tradeoff coefficient is stochastic, with explicitly modeled dynamics that can be used to generate investment policies tailored to the client's risk profile.

Our work is also related to the literature on portfolio optimization problems where the price dynamics are driven by a stochastic factor (see, for instance, Liu [2007]). Under this setup, the market is incomplete, because trading in the stock and the risk-free asset cannot perfectly hedge changes in the stochastic investment opportunity set. In our work, the stochastic factor is uncorrelated to the stock market dynamics, so the optimal stock allocation does not contain an intertemporal hedging term. However, in our study, the stochastic risk factor also impacts the dynamics of the investor's risk aversion, which in turn affects the optimal stock allocation. It is well established that risk aversion varies with the business cycle, along with the market risk-return tradeoff, as measured by the market Sharpe ratio (Lettau and Ludvigson [2010]). This modeling feature is of particular importance for robo-advisors, which face significant challenges when it comes to keeping clients invested in the stock market during economic downturns.

Our study also contributes to the growing literature on robo-advising. D'Acunto and Rossi [2020] describe the main components of robo-advising systems, and propose a classification in terms of four features: (i) portfolio personalization, (ii) client involvement, (iii) client discretion, and (iv) human interaction. We offer a quantitative robo-advising framework that aligns with this classification. In summary, our framework is consistent with that of the most prominent stand-alone robo-advising firms, which aim for high portfolio personalization, low client involvement, low and indirect investor discretion, and limited or no human interaction. We proceed to briefly describe each of the four features and how they fit into our framework.

(i) *Portfolio personalization* refers to the robo-advisor's ability to offer financial advice tailored to the client's needs. A common criticism of robo-advising is the lack of customization, with risk profiling based on information that is too limited.[9] In our framework, portfolio personalization is achieved through interacting with the client, but we do not model explicitly the transfer of information between the client and the robo-advisor. Hence, our framework and measure of portfolio personalization are not restricted to the relatively simple risk profiling of online questionnaires. Our results on the frequency of interaction can also be conjectured to extend to a setting where client-data is continuously collected from multiple sources, where there exists an optimal data-processing frequency that maximizes a signal-to-noise ratio.

(ii) *Client involvement* refers to the client's participation in the design of the investment strategy. At one end of the spectrum, we have robo-advisors that require the client to approve every single trading decision. In this case, the client is actively involved and the term robo-*advisor* is descriptive, because the client decides in what way to follow the *advice* provided. Our model lies at the other end of the spectrum, where the robo-advisor automatically manages the portfolio on the client's behalf and for which the term robo-*manager* is more descriptive.

(iii) *Client discretion* is related to client involvement, and measures the client's ability to override the robo-advisor's recommendation. In our framework, the client's behavioral biases can be viewed as the client overriding the robo-advisor's recommendations. This is considered a low level of discretion, and is consistent with the operations of robo-advisors focusing on long-term investing, where the client is allowed to adjust the level of portfolio risk, but has no control over which parts of the investment portfolio are modified.

[9] In principle, risk profiling is infinitely customizable. Barron's 2019 annual ranking of robo-advisors (available at `https://webreprints.djreprints.com/4642511400002.pdf`) reports the latest growth in robo-advising to be in cash management, with robo-advisors aiming to become everyday money managers of their clients, in addition to long-term investment planners. Robo-advising firms already offer services such as direct deposits of paychecks, automatic bill payments, and FDIC-insured checking and savings accounts. Such access to data on savings and spending behavior, and asset and liabilities, can be used by robo-advisors to improve investment recommendations.

(iv) *Human interaction* refers to the degree of interaction between the client and a human-advisor. In our framework there is no human contact, which is also the case for automated robo-advisors that strive to minimize operating costs.

The above classification is qualitative in nature, but the technical components of our framework are also constructed to resemble the operations of actual robo-advisors. Beketov et al. [2018] provide industry statistics and an industry overview based on the analysis of over 200 robo-advisors globally. Their study shows that a large majority of robo-advisors use an asset allocation framework based on mean-variance analysis, with the asset universe consisting of low cost exchange-traded funds. They also show that risk profiling of clients is primarily implemented via online questionnaires.

Rossi and Utkus [2019] conduct an extensive survey to study the "needs and wants" of individuals when they hire financial advisors. Their results lend support to the theoretical model of Gennaioli and Vishny [2015], indicating that traditionally advised individuals hire financial advisors largely to satisfy various needs other than portfolio return maximization. Namely, they argue that clients choose a traditional financial advisor primarily for the ability to interact with and receive financial advice from a human, emphasizing the role of *trust* in client-advisor relationships. They also show that for traditionally-advised clients, algorithm aversion and the inability to interact with a human are the main obstacles for switching to robo-advising, while robo-advised clients do not have the same need for trust and access to expert opinion. Our model lends theoretical support to the notion that robo-advisors are less suited to be the investment managers of algorithm averse client. The robo-advisor can improve the portfolio performance by going against the wishes of the client, i.e., investing in a way that may seem counterintuitive to the client. Doing so is more challenging for a robo-advisor than for a human-advisor, because market returns are random, and an algorithm averse client will be less forgiving to the robo-advisor in the event of adverse market returns.

3 Modeling Framework

The robo-advising framework consists of four main components: (i) a market model for the available investment securities, (ii) a mechanism of interaction between the client and the robo-advisor, (iii) a dynamic model for the client's risk preferences, and (iv) an optimal investment criterion. In the following sections, we describe each of these components in full generality, and then present a specific model fitting into this framework.

3.1 Market and Wealth Dynamics

The market consists of a risk-free money market account, $(B_n)_{n\geq 0}$, and a risky asset, $(S_n)_{n\geq 0}$, satisfying

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

The above dynamics are modulated by an observable state process, $(Y_n)_{n\geq 0}$, which captures macroeconomic conditions affecting interest rates and stock market returns. We assume $(Y_n)_{n\geq 0}$ to be a time-homogeneous Markov chain with transition matrix P, taking values in a finite set $\mathcal{Y} := \{1, 2, \dots, M\}$, for some $M \geq 1$.

Conditioned on $Y_n = y \in \mathcal{Y}$, the risk-free interest rate $r(y) \geq 0$ is constant, while the risky asset's return, $Z_{n+1}(y)$, admits a probability density function $f_{Z|y}$ which depends only on the current economic state y, and has mean $\mu(y) > r(y)$ and variance $0 < \sigma^2(y) < \infty$.

For notational simplicity, we omit the dependence on the economic state. Hence, we use Z_{n+1} in place of $Z_{n+1}(Y_n)$, and denote its state-dependent mean and variance by $\mu_{n+1} := \mu(Y_n)$ and $\sigma^2_{n+1} := \sigma^2(Y_n)$, respectively. Similarly, we use r_{n+1} in place of $r(Y_n)$. We denote by $\widetilde{Z}_{n+1} := Z_{n+1} - r_{n+1}$ the excess return of the risky asset over the risk-free rate, which has state-dependent mean and variance given by $\tilde{\mu}_{n+1} := \mu_{n+1} - r_{n+1}$ and σ^2_{n+1}, respectively.

We denote by X_n the wealth of the client at time n, allocated between the risky asset and the money market account, and use π_n to denote the amount invested in the risky asset. For a given self-financing trading strategy $\pi := (\pi_n)_{n\geq 0}$, the wealth process $(X^\pi_n)_{n\geq 0}$ follows the dynamics

$$X^\pi_{n+1} = (1 + r_{n+1})X^\pi_n + (Z_{n+1} - r_{n+1})\pi_n =: R_{n+1}X^\pi_n + \widetilde{Z}_{n+1}\pi_n. \tag{3.1}$$

The initial state of the economy, Y_0, is assumed to be non-random. The random variables $(Y_n)_{n\geq 1}$ and $(Z_n)_{n\geq 0}$ are defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, which additionally supports a sequence $(\epsilon_n)_{n\geq 1}$ of independent real-valued random variables. This source of randomness captures idiosyncratic changes to the client's risk preferences and is independent of the variables $(Y_n)_{n\geq 0}$ and $(Z_n)_{n\geq 1}$. We use $(\mathcal{F}_n)_{n\geq 0}$ to denote the filtration generated by the three stochastic processes in our model:

$$\mathcal{F}_n := \sigma(Y_{(n)}, Z_{(n)}, \epsilon_{(n)}), \tag{3.2}$$

where $Y_{(n)} := (Y_0, \dots, Y_n)$, $Z_{(n)} := (Z_1, \dots, Z_n)$, and $\epsilon_{(n)} := (\epsilon_1, \dots, \epsilon_n)$. We will use analogous notation to denote the paths of other stochastic processes throughout the paper.

3.2 Interaction between Client and Robo-Advisor

A key component of the proposed framework is the dynamic interaction between the client and the robo-advisor. Following an initial interaction at the beginning of the investment process, the client and the robo-advisor interact *repeatedly* throughout the investment period.

At each interaction time, the client communicates her current risk preferences to the robo-advisor, which translates them into a numerical value, herein referred to as the *client's risk aversion parameter*. We abstract from the construction of such a mapping, effectively assuming that the client communicates directly a single risk aversion parameter to the robo-advisor. Note that between consecutive interaction times, the robo-advisor receives no input from the client.

The *interaction schedule*, denoted by $(T_k)_{k\geq 0}$, is an increasing sequence of stopping times with respect to the filtration $(\mathcal{F}_n)_{n\geq 0}$, defined in (3.2). That is, $T_0 = 0$, $T_k < T_{k+1}$, and $\{T_k \leq n\} \in \mathcal{F}_n$, for any $n \geq 0$. Hence, the interaction can be triggered by any combination of client-specific events, changes in the state of the economy, or market events, such as a cascade of negative market returns.

The rule used to determine the interaction schedule $(T_k)_{k\geq 0}$ is decided at the beginning of the investment process, and used throughout it. For future reference, we also define the process $(\tau_n)_{n\geq 0}$, where $\tau_n := \sup\{T_k : T_k \leq n\}$ is the last interaction time occurring up to and including time n.

3.3 Client's Risk Aversion Process

The client's risk aversion process, $(\gamma_n^C)_{n\geq 0}$, is an $\mathbb{R}_0^+$-valued stochastic process adapted to the filtration $(\mathcal{F}_n)_{n\geq 0}$ defined in (3.2), with a fixed initial value, γ_0^C. At time n, the client's risk aversion may have shifted away from its initial value due to changes in economic regimes $Y_{(n)}$ and realized market returns $Z_{(n)}$, as well as because of idiosyncratic shocks to the client's risk aversion, $\epsilon_{(n)}$.

The client's risk aversion process is only partially observed by the robo-advisor. We introduce the $(\mathcal{F}_n)_{n\geq 0}$-adapted process $(\xi_n)_{n\geq 0}$, where $\xi_n \in \mathbb{R}_0^+$, to track the risk aversion parameter communicated by the client at the most recent interaction time, τ_n. Observe that the process $(\xi_n)_{n\geq 0}$ depends on the schedule $(T_k)_{k\geq 0}$, and, by construction, it is constant between consecutive interaction times, with no new information coming from the client, i.e., $\xi_n = \xi_{\tau_n}$.

The robo-advisor constructs its *own model* of the client's risk aversion process, denoted by $(\gamma_n)_{n\geq 0}$, and uses it to solve the optimal investment problem. There are two main reasons for why the processes $(\gamma_n^C)_{n\geq 0}$ and $(\gamma_n)_{n\geq 0}$ do not necessarily always coincide. First, the client and the robo-advisor may not interact at all times, so the robo-advisor does not always have access to up-to-date information about the client's risk preferences. Second, information communicated by the client at interaction times may not be representative of the client's true risk preferences, due to the client's behavioral biases (see Section 3.5 for details).

The robo-advisor filtration, $(\mathcal{F}_n^R)_{n\geq 0}$, is generated by the random variables $(D_n)_{n\geq 0}$, defined as

$$D_n := (Y_{(n)}, Z_{(n)}, \tau_{(n)}, \xi_{(n)}) \in \mathcal{D}_n, \qquad \mathcal{D}_n := \mathcal{Y}^{n+1} \times \mathbb{R}^n \times \mathbb{N}^{n+1} \times (\mathbb{R}_0^+)^{n+1}. \tag{3.3}$$

The risk aversion process $(\gamma_n)_{n\geq 0}$ is a $\mathbb{R}_0^+$-valued stochastic process adapted to the robo-advisor filtration $(\mathcal{F}_n^R)_{n\geq 0}$. At time n, it may thus be written as

$$\gamma_n := \gamma_n(D_n) \in \mathcal{F}_n^R, \tag{3.4}$$

for a measurable function $\gamma_n : \mathcal{D}_n \mapsto \mathbb{R}_0^+$.

Remark 3.1. The random variable D_n in (3.3) can be decomposed as $D_n = (M_n, I_n)$, where $M_n := (Y_{(n)}, Z_{(n)})$, and $I_n := (\tau_{(n)}, \xi_{(n)})$. This shows that the robo-advisor filtration $(\mathcal{F}_n^R)_{n\geq 0}$ has two sources of information. First, the robo-advisor has the ability to process all available information about the market and the economy, so $M_n \in \mathcal{F}_n^R$. Second, the process $(I_n)_{n\geq 0}$ captures the interaction with the client, and contains both the history of interaction times, and the communicated risk aversion values. Under mild conditions, the robo-advisor filtration grows with the frequency of interaction, and, if interaction occurs at all times, $(\mathcal{F}_n^R)_{n\geq 0}$ becomes equal to the filtration $(\mathcal{F}_n)_{n\geq 0}$.[10] □

3.4 Investment Criterion

The robo-advisor's objective is to optimally invest the client's wealth, taking into account the stochastic nature of the client's risk preferences. For this purpose, we develop a dynamic version of the standard Markowitz [1952] mean-variance problem that adapts to the client's changing risk preferences. We proceed to introduce this criterion, which we refer to as an *adaptive mean-variance criterion.*

Let $T \geq 1$ be a fixed investment horizon. For each $n \in \{0, 1, \ldots, T-1\}$, $x \in \mathbb{R}$, $d \in \mathcal{D}_n$, and a control law $\pi := (\pi_n)_{n\geq 0}$, we consider the mean-variance functional

$$J_n(x, d; \pi) := \mathbb{E}_{n,x,d}[r_{n,T}^\pi] - \frac{\gamma_n(d)}{2} Var_{n,x,d}[r_{n,T}^\pi], \tag{3.5}$$

[10]For example, assume the interaction times to be given by $T_k = k\phi$, for some $\phi \geq 1$. If the sum of idiosyncratic shocks, $(\sum_{k=1}^n \epsilon_k)_{n\geq 0}$ is measurable with respect to the robo-advisor's filtration, then $(\mathcal{F}_n^R)_{n\geq 0}$ increases to $(\mathcal{F}_n)_{n\geq 0}$, as $\phi \downarrow 1$.

where $r^{\pi}_{n,T}$ is the simple return obtained by following the control law until the terminal date T,

$$r^{\pi}_{n,T} := \frac{X^{\pi}_T - X_n}{X_n}. \tag{3.6}$$

The initial condition in (3.5) is given by $X_n = x$ and $D_n = d$, and both the expectation and the variance are computed with respect to the probability measure $\mathbb{P}_{n,x,d}(\cdot) := \mathbb{P}(\cdot | X_n = x, D_n = d)$. The objective functionals in (3.5) and the risk aversion process (3.4) define a family of sequentially adaptive optimization problems, in the sense that, at each time, a new problem arises with properties that depend on realized market returns, economic state changes, and client-communicated information, but with the same initially specified terminal date. Moreover, observe that it is the robo-advisor that solves the optimization problem, and the initial condition fixes the value of the stochastic process that generates the filtration of the robo-advisor and the value of the client's wealth process.

The control law π in (3.5), also referred to as a strategy or allocation, is such that for each n, the control π_n is a measurable real-valued function of the state variables X_n and D_n. Additionally, admissible control laws are assumed to be self-financing, which results in the wealth dynamics in (3.1), and to satisfy the square-integrability condition $\mathbb{E}\Big[\sum_{n=0}^{T-1} \pi_n^2\Big] < \infty$.

Remark 3.2. We assume the investment horizon T to be fixed, but a stochastic terminal date can be captured through the client's risk aversion process $(\gamma_n)_{n\geq 0}$. For example, the possibility of the client dying can be modeled by introducing a time-dependent probability p_n of the risk aversion parameter γ_{n+1} becoming "infinite". The probability p_n would be increasing in n, and infinity would be an absorbing state reached only in the event of the client dying. In the death-state, the client's portfolio is liquidated in the sense that it shifts entirely to the risk-free asset. In the optimization problem of Section 3.4, a nonzero value of p_n would have the effect of tilting the risky asset allocation upward, to account for the possibility of death where no more investing is possible.

In the same vein, one can incorporate other rule-based portfolio constraints to account for the client's preferences and goals. For example, a constraint can be built into the risk aversion process so that the client's portfolio is liquidated (i.e., it becomes risk-free), if a given level of wealth is reached before the horizon T, or if a stop-loss threshold is exceeded. □

Through the wealth dynamics (3.1), the mean-variance functional J_n depends on the control law π restricted to the time points $\{n, n+1, \ldots, T-1\}$, and the robo-advisor chooses the control π_n *given* future control decisions $\pi_{n+1:T} := \{\pi_{n+1}, \pi_{n+2}, \ldots, \pi_{T-1}\}$. Any candidate optimal control law π^* is therefore such that for each $n \in \{0, 1, \ldots, T-1\}$, $x \in \mathbb{R}$, and $d \in \mathcal{D}_n$,

$$\sup_{\pi \in A^*_{n+1}} J_n(x, d; \pi) = J_n(x, d; \pi^*), \tag{3.7}$$

where $A^*_{n+1} := \{\pi : \pi_{n+1:T} = \pi^*_{n+1:T}\}$ is the set of control laws that coincide with π^* after time n. If a control law π^* satisfying (3.7) exists, we define the corresponding value function at time n as

$$V_n(x, d) := J_n(x, d; \pi^*). \tag{3.8}$$

The *risk-return tradeoff* at time n, in the mean-variance criterion (3.5), is given by the robo-advisor's model of the client's risk aversion, γ_n. If the process $(\gamma_n)_{n \geq 0}$ is constant, then (3.5) reduces to the classical mean-variance criterion. Herein, we allow for a richer structure to incorporate the dynamic nature of the client's risk preferences, and the effect of interaction between the client and the robo-advisor.

Observe that the risk-return tradeoff does not depend on the client's wealth. There are important reasons behind this assumption. First, this is the only choice that guarantees both the mean and variance terms to be unitless, and thus directly comparable. Second, this ensures that for the case of a single investment period, the original single-period mean-variance criterion is recovered, which was formulated in terms of returns with a constant risk aversion. Third, as we will see in Section 4, the optimal investment strategy turns out to be consistent with an investor with risk preferences exhibiting constant relative risk aversion (CRRA), i.e., an investor who chooses the same asset allocation at all levels of wealth.[11]

Remark 3.3. At time n, any optimal control π^*_n depends on the dynamics of the risk-return coefficient throughout the investment horizon, $\{\gamma_n, \gamma_{n+1}, \ldots, \gamma_{T-1}\}$, with only the current value γ_n being known. The expected value $\mathbb{E}_{n,x,d}[\cdot]$ averages over the distribution of this process. Specifically, at the initial time, the robo-advisor constructs a model for the future evolution of the client's risk aversion, throughout the investment horizon. This model is given by the process $(\gamma_n)_{0 \leq n < T}$, which depends on the client's initially communicated risk aversion, and the rule used to determine future times of interaction with the client. When seeking an optimal control π^*_0, the robo-advisor averages over future risk aversion paths implied by this model. At each subsequent time, the robo-advisor updates its model for the future evolution of the client's risk aversion. At non-interaction times, this update is based only on new information about market returns and economic states, while at times of interaction it also accounts for information received from the client. □

[11] Empirical evidence suggests that CRRA is a good description of microeconomic behavior (see, e.g., extensive panel data studies carried out in Brunnermeier and Nagel [2008], Chiappori and Paiella [2011], and Sahm [2012]). However, the evidence is not universal. For instance, Calvet and Sodini [2014] and Guiso and Paiella [2008] reject CRRA in favor of decreasing relative risk aversion.

3.5 A Robo-Advising Model

In this section, we present a model that fits into the robo-advising framework introduced in Sections 3.1-3.4.

Market Dynamics. The risky asset has conditionally Gaussian returns. That is, given the economic state $Y_n = y$ at time n, the return Z_{n+1} has a Gaussian distribution with mean $\mu(y)$ and variance $\sigma^2(y)$.

Interaction Schedule. The sequence of interaction times is deterministic and equally spaced. The interaction schedule $(T_k)_{k\geq 0}$ is then characterized by a fixed integer parameter $\phi \geq 1$, with $T_k = k\phi$. One extreme is $\phi = 1$, i.e., risk preferences communicated at all times. For a fixed investment horizon $T \geq 1$, the other extreme is $\phi \geq T$, i.e., risk preferences communicated at the outset and never after that.

Client's Risk Aversion. The client's risk aversion process, $(\gamma_n^C)_{n\geq 0}$, is of the form

$$\gamma_n^C = e^{\eta_n} \gamma_n^{id} \gamma_n^Y, \tag{3.9}$$

where $(e^{\eta_n})_{n\geq 0}$ is a temporal component of the form $e^{\eta_n} = e^{-\alpha(T-1-n)}$, for some $\alpha \geq 0$. The parameter α determines the rate of increase in the client's risk aversion due to aging.

Remark 3.4. A risk aversion process of the form $\gamma_n^C = e^{\eta_n}$ is consistent with the structure of so-called Target Date Funds (TDFs), which are age-based retirement portfolios that reduce equity risk as the client gets closer to retirement. This is similar to what was offered by the first generation of robo-advisors, which constructed "set-and-forget" portfolio strategies, with the allocation and risk tolerance changing over time based on the client's age. While TDFs have virtually stayed the same since their inception over a quarter century ago and offer no customization, the technology-driven personalized advice of robo-advisors has continued to evolve.

The second component, $(\gamma_n^{id})_{n\geq 0}$, depends on the client's personal circumstances, and has dynamics

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

where $0 \leq p_\epsilon \leq 1$, $\sigma_\epsilon > 0$, and $(W_n)_{n\geq 1}$ is an i.i.d. sequence of standard Gaussian random variables. This component captures idiosyncratic shocks to the client's risk aversion, that are unrelated to market dynamics. The multiplicative innovation terms $(e^{\epsilon_n})_{n\geq 0}$ are independent and with unit mean, making $(\gamma_n^{id})_{n\geq 0}$ a martingale, devoid of a predictable component. Observe that due to the convexity of the exponential function, an upward idiosyncratic shock (i.e., $\epsilon_n > 0$) has a greater impact on risk aversion than a downward shock (i.e., $\epsilon_n < 0$) of equal size.

The third component, $(\gamma_n^Y)_{n\geq 0}$, is given by

$$\gamma_n^Y := \sum_{y\in\mathcal{Y}} \bar{\gamma}(y)\mathbf{1}_{\{Y_n=y\}},$$

where $\bar{\gamma} : \mathcal{Y} \mapsto \mathbb{R}_0^+$ is a risk aversion coefficient that depends on the current economic conditions. Consistent with empirical evidence, which has shown both risk aversion and the market Sharpe ratio to be countercyclical, the function $\bar{\gamma}$ needs to satisfy the following relation

$$\frac{\mu(y) - r(y)}{\sigma(y)} \geq \frac{\mu(y') - r(y')}{\sigma(y)} \quad \text{implies} \quad \bar{\gamma}(y) \geq \bar{\gamma}(y'),$$

for any $y, y' \in \mathcal{Y}$. Under this condition, the client's risk aversion is higher precisely when the benefit of investing in the risky asset is higher, as measured by the expected reward per unit risk.

Robo-Advisor's Model of Client's Risk Aversion. The risk aversion process $(\gamma_n)_{n\geq 0}$ is of the form

$$\gamma_n = \xi_n e^{\eta_n - \eta_{\tau_n}} \frac{\gamma_n^Y}{\gamma_{\tau_n}^Y}. \tag{3.10}$$

The first component, ξ_n, is the risk aversion parameter communicated by the client to the robo-advisor at the most recent interaction time, τ_n. It is of the form

$$\xi_n = \gamma_{\tau_n}^C \gamma_{\tau_n}^Z := \gamma_{\tau_n}^C e^{-\beta\left(\frac{1}{\phi}\sum_{k=\tau_n-\phi}^{\tau_n-1}(Z_{k+1}-\mu_{k+1})\right)}, \tag{3.11}$$

where, for a fixed $\beta \geq 0$, the factor $\gamma_{\tau_n}^Z$ inflates or deflates the risk aversion value communicated to the robo-advisor at time τ_n, relative to the client's actual risk aversion, $\gamma_{\tau_n}^C$. This factor depends on recent market returns, and the sum in the exponent is the cumulative excess return of the risky asset over its expected value, in the time interval between the two most recent interaction times, $\tau_n - \phi$ and τ_n. Because previous market returns have no predictive value for future market returns, given the state of the economy, this is representative of common behavioral biases and the coefficient $\beta \geq 0$ determines the magnitude of this effect. We emphasize two important properties of the component $\gamma_{\tau_n}^Z$. First, the convexity of the exponential function allows us to capture loss aversion: the client's risk aversion increases more following a market underperformance, compared to its decrease when the market exceeds expectations by the same amount (Tversky and Kahneman [1979]). Second, in a model with a single economic state, $\mathbb{E}[\gamma_{\tau_n}^Z] = e^{\frac{\beta\sigma^2}{2\phi}}$, which shows that the average effect of behavioral biases decreases as the time between consecutive interactions increases (Benartzi and Thaler [1995]).[12]

The component ξ_n in (3.10), communicated by the client at time τ_n, stays constant until the following time of interaction, $\tau_n + \phi$. However, the risk aversion is adjusted for the passage of time through the factor $e^{\eta_n - \eta_{\tau_n}}$ and for changes in the state of the economy through the ratio $\gamma_n^Y / \gamma_{\tau_n}^Y$.

Remark 3.5. The component $\gamma_{\tau_n}^Z$ in (3.11) can also be viewed as the result of the client overriding the robo-advisor's decisions. That is, at an interaction time n, the robo-advisor proposes a portfolio tailored to the client's characteristics, γ_n^C, but the client makes changes to the proposed allocation based on recent market returns, resulting in a portfolio allocation consistent with a risk aversion coefficient $\gamma_n^C \gamma_n^Z$. The effect of this override, which is driven by the client's behavioral biases, prevails until the subsequent time of interaction, $n + \phi$, when the client again adjusts the portfolio allocation proposed by the robo-advisor. However, this time the adjustment is based on market returns realized in the time interval $[n, n + \phi]$. □

4 Optimal Investment with Client Interaction

In this section, we study the optimal investment problem solved by the robo-advisor. Section 4.1 analyzes the optimization problem defined by the objective functionals $(J_n)_{0 \leq n < T}$ in (3.5) and the optimality criterion in (3.7). Section 4.2 plots patterns of dependence of the optimal investment strategy on features of the client's risk aversion.

4.1 Optimal Investment Strategy

In this section, we derive an extended Hamilton-Jacobi-Bellman (HJB) system of equations satisfied by the value function of the optimization problem, and then develop a recursive representation for the solution of this system. It is well known that, even if the risk-return tradeoff $(\gamma_n)_{n \geq 0}$ is constant through time, the family of optimization problems defined by (3.5) is time-inconsistent in the sense that the Bellman optimality principle does not hold. This means that if, at time n, the control law π^* maximizes the objective functional J_n, then, at time $n+1$, the restriction of π^* to the time points $\{n+1, n+2, \ldots, T-1\}$ may not maximize J_{n+1}. We refer to Björk and Murgoci [2013] and references therein for a general framework of time-inconsistent stochastic control in discrete time.

As standard in this literature, we view the optimization problem in (3.7) as a multi-player game, where the player at each time $n \in \{0, 1, \ldots, T-1\}$ is thought of as a future self of the client. Player n then wishes

[12] In a setup with multiple economic states, conditionally on $Y_{\tau_n - \phi} = y \in \mathcal{Y}$ we have $\mathbb{E}[\gamma_{\tau_n}^Z] = e^{\frac{\beta \sigma^2(y)}{2\phi}} + O(1 - P_{y,y}^{\phi})$.

to maximize the objective functional J_n, but decides only the strategy π_n at time n, while $\pi_{n+1}, \dots, \pi_{T-1}$ are determined by her future selves. The resulting optimal control strategy, π^*, is the subgame perfect equilibrium of this game, and can be computed using backward induction. At time $n = T-1$, the equilibrium control π^*_{T-1} is obtained by maximizing J_{T-1} over π_{T-1}, which is a standard single-period optimization problem. For $n < T-1$, the equilibrium control π^*_n is then obtained by letting player n choose π_n to maximize J_n, given that player n' will use $\pi^*_{n'}$, for $n' = n+1, n+2, \dots, T-1$.

Recall the return variables R_{n+1} and $\widetilde{Z}_{n+1}$ defined in Section 3.1. The following result formulates the optimal investment strategy for an interaction schedule $(T_k)_{k\geq 0}$ and a risk aversion process $(\gamma_n)_{n\geq 0}$ of the general form introduced in Sections 3.2 and 3.3.

Proposition 4.1. *The optimization problem (3.7) is solved by the control law*

$$\pi^*_n(x,d) = \tilde{\pi}^*_n(d)x_n, \quad 0 \leq n < T, \tag{4.1}$$

where, for $x \in \mathbb{R}$ and $d \in \mathcal{D}_n$, the optimal proportion of wealth allocated to the risky asset is given by

$$\tilde{\pi}^*_n(d) = \frac{1}{\gamma_n}\frac{\mu^{az}_n(d) - R_{n+1}\gamma_n(\mu^{bz}_n(d) - \mu^a_n(d)\mu^{az}_n(d))}{\mu^{bz^2}_n(d) - (\mu^{az}_n(d))^2}. \tag{4.2}$$

Setting $\mathbb{P}_{n,d}(\cdot) := \mathbb{P}(\cdot|D_n = d)$, the coefficients above are given by

$$\begin{aligned} &\mu^a_n(d) = \mathbb{E}_{n,d}[a_{n+1}(D_{n+1})], \quad \mu^{az}_n(d) = \mathbb{E}_{n,d}[a_{n+1}(D_{n+1})\widetilde{Z}_{n+1}], \\ &\mu^b_n(d) = \mathbb{E}_{n,d}[b_{n+1}(D_{n+1})], \quad \mu^{bz}_n(d) = \mathbb{E}_{n,d}[b_{n+1}(D_{n+1})\widetilde{Z}_{n+1}], \quad \mu^{bz^2}_n(d) = \mathbb{E}_{n,d}[b_{n+1}(D_{n+1})\widetilde{Z}^2_{n+1}]. \end{aligned} \tag{4.3}$$

The sequences $(a_n(d))_{0\leq n<T}$ and $(b_n(d))_{0\leq n<T}$ satisfy the recursions

$$\begin{aligned} a_n(d) &= \mathbb{E}_{n,d}[(R_{n+1} + \widetilde{Z}_{n+1}\tilde{\pi}^*_n(d))a_{n+1}(D_{n+1})], \\ b_n(d) &= \mathbb{E}_{n,d}[(R_{n+1} + \widetilde{Z}_{n+1}\tilde{\pi}^*_n(d))^2 b_{n+1}(D_{n+1})], \end{aligned} \tag{4.4}$$

with $a_T(d) = b_T(d) = 1$, for all $d \in \mathcal{D}_T$.

The quantities $a_n(d)$ and $b_n(d)$ in Proposition 4.1 are the first and second moments of the future value of one dollar invested optimally between time n and the terminal date T, given that $D_n = d$. That is,

$$a_n(d) = \mathbb{E}_{n,d}[1 + r^{\pi^*}_{n,T}], \qquad b_n(d) = \mathbb{E}_{n,d}[(1 + r^{\pi^*}_{n,T})^2],$$

with the simple return $r^{\pi^*}_{n,T}$ defined in (3.6). Because these values are independent of the wealth amount x, it then follows from (3.8) that the value function at time n is also independent of x, and given by

$$V_n(x,d) = V_n(d) := a_n(d) - 1 - \frac{\gamma_n(d)}{2}(b_n(d) - a_n^2(d)).$$

Furthermore, Proposition 4.1 shows that the optimal proportion of wealth allocated to the risky asset, $\tilde{\pi}_n^*$, is independent of the wealth variable.

The expected values in (4.3) admit integral representations. In the lemma below, we provide such a representation for $\mu_n^{az}(d)$, both in the general case, and in the specific case of the risk aversion model described in Section 3.5.

Lemma 4.2. *(a) Assume that given $D_n = d$, the random variable $\epsilon_{(n+1)}$ admits a generalized probability density function, $f_{\epsilon_{(n+1)}|d}$. Then,*

$$\mu_n^{az}(d) = \sum_{y' \in \mathcal{Y}} P_{y_n,y'} \int_{z' \in \mathbb{R}} \int_{\epsilon' \in \mathbb{R}^{n+1}} a_{n+1}(D_{n+1}) \tilde{z}' f_{Z|y_n}(z') f_{\epsilon_{(n+1)}|d}(\epsilon') dz' d\epsilon',$$

where $\tilde{z}' = z' - \mu(y_n)$, and $D_{n+1} = (Y_{(n+1)}, Z_{(n+1)}, \tau_{(n+1)}, \xi_{(n+1)}) \in \mathcal{D}_{n+1}$ is such that $Y_{(n+1)} = (y, y')$, $Z_{(n+1)} = (z, z')$, $\tau_{(n+1)} = (\tau, \tau_{n+1})$, and $\xi_{(n+1)} = (\xi, \xi_{n+1})$. The value of τ_{n+1} is determined by the triplet $(Y_{(n+1)}, Z_{(n+1)}, \epsilon_{(n+1)})$. If $\tau_{n+1} = \tau_n$, then $\xi_{n+1} = \xi_n$, while if $\tau_{n+1} = n+1$, ξ_{n+1} is determined by $(Y_{(n+1)}, Z_{(n+1)}, \epsilon_{(n+1)})$.

(b) For the risk aversion model in Section 3.5, we have

$$\mu_n^{az}(d) = \begin{cases} \sum_{y' \in \mathcal{Y}} P_{y_n,y'} \int_{z' \in \mathbb{R}} a_{n+1}(D_{n+1}) \tilde{z}' f_{Z|y_n}(z') dz', & \tau_{n+1} < n+1, \\ \sum_{y' \in \mathcal{Y}} P_{y_n,y'} \int_{z' \in \mathbb{R}} \int_{\epsilon' \in \mathbb{R}} a_{n+1}(D_{n+1}) \tilde{z}' f_{Z|y_n}(z') f_\epsilon^{(\phi)}(\epsilon') dz' d\epsilon', & \tau_{n+1} = n+1, \end{cases} \tag{4.5}$$

where $f_{Z|y_n}$ is the density of a Gaussian distribution with mean $\mu(y_n)$ and variance $\sigma^2(y_n)$, $f_\epsilon^{(\phi)}$ is the ϕ-fold convolution of the generalized density function of the i.i.d. sequence $(\epsilon_n)_{n \geq 1}$, and[13]

$$D_{n+1} = \begin{cases} \big((y, y'), (z, z'), (\xi, \xi_n)\big), & \tau_{n+1} < n+1, \\ \big((y, y'), (z, z'), (\xi, \xi_{n+1})\big), & \tau_{n+1} = n+1. \end{cases}$$

If $\tau_{n+1} = n+1$, then ξ_{n+1} is given by an explicit function of $((y, y'), (z, z'), \xi_n, \epsilon')$.

Equation (4.5) shows that there are two distinct cases, depending on whether or not the client's risk preferences are solicited. In all cases, the computation of $\mu_n^{az}(d)$ involves integrating over the distribution

[13] With a deterministic interaction schedule, we omit the interaction times τ_n from the state variable D_n.

of the market return Z_{n+1}, and the economic state Y_{n+1}. Additionally, if $\tau_{n+1} = n+1$, the client's risk preferences are solicited at time $n+1$, and $\mu_n^{az}(d)$ also involves integrating over the distribution of the cumulative idiosyncratic risk aversion shock $\sum_{k=n+1-\phi}^{n} \epsilon_{k+1}$. Hence, the probability distributions $f_{Z_{n+1}}$, $P_{y_n,\cdot}$, and $f_\epsilon^{(\phi)}$, link the optimal allocation at time n, in state d, to the optimal allocations at time $n+1$ in the random state D_{n+1}. Observe that if $\tau_{n+1} < n+1$, the dynamics are only based on incoming market and economic information, Z_{n+1} and Y_{n+1}, while in the case $\tau_{n+1} = n+1$ they also depend on information received from the client, $\sum_{k=n+1-\phi}^{n} \epsilon_{k+1}$. In Appendix C, we further discuss the computation of the optimal investment strategy, and the computational complexity associated with the dimensionality of the state variable D_n.

We next consider a special case of the model in Section 3.5, for which the optimal allocation strategy admits a more explicit expression. Specifically, we set $\beta = 0$, i.e., assume no behavioral biases, so that the risk aversion process is of the form

$$\gamma_n = e^{-\alpha(T-1-n)} e^{\sum_{i=1}^{\tau_n} \epsilon_i} \bar{\gamma}(Y_n). \tag{4.6}$$

Proposition 4.3. *Assume a deterministic interaction schedule $(T_k)_{k\geq 0}$ given by $T_k = k\phi$, for some $\phi \geq 1$, and that $(\gamma_n)_{n\geq 0}$ is of the form (4.6). The optimization problem (3.7) is then solved by the control law*

$$\pi_n^*(x,d) = \tilde{\pi}_n^*(d)x_n = \frac{1}{\gamma_n} \frac{\tilde{\mu}_{n+1}\mu_n^a(d) - R_{n+1}\gamma_n\tilde{\mu}_{n+1}(\mu_n^b(d) - (\mu_n^a(d))^2)}{(\tilde{\mu}_{n+1}^2 + \sigma_{n+1}^2)\mu_n^b(d) - (\tilde{\mu}_{n+1}\mu_n^a(d))^2} x_n, \quad 0 \leq n < T,$$

for each $x \in \mathbb{R}$ and $d \in \mathcal{D}_n$, and

$$\mu_n^a(d) = \mathbb{E}_{n,d}[a_{n+1}(D_{n+1})], \qquad \mu_n^b(d) = \mathbb{E}_{n,d}[b_{n+1}(D_{n+1})].$$

The sequences $(a_n(d))_{0\leq n<T}$ and $(b_n(d))_{0\leq n<T}$ satisfy the recursions

$$\begin{aligned} a_n(d) &= (R_{n+1} + \tilde{\mu}_{n+1}\tilde{\pi}_n^*(d))\mu_n^a(d), \\ b_n(d) &= ((R_{n+1} + \tilde{\mu}_{n+1}\tilde{\pi}_n^*(d))^2 + (\sigma_{n+1}\tilde{\pi}_n^*(d))^2)\mu_n^b(d), \end{aligned}$$

with $a_T(d) = b_T(d) = 1$, for all $d \in \mathcal{D}_T$.

The simplification of the formulas in the previous proposition are due to the expected values in (4.3) factoring into products. For example, $\mu_n^{az}(d)$ becomes equal to $\mu_n^a(d)\tilde{\mu}_{n+1}$, with

$$\mu_n^a(d) = \begin{cases} \sum_{y' \in \mathcal{Y}} P_{y_n,y'} a_{n+1}(D_{n+1}), & \tau_{n+1} < n+1, \\ \sum_{y' \in \mathcal{Y}} P_{y_n,y'} \int_{\epsilon' \in \mathbb{R}} a_{n+1}(D_{n+1}) f_\epsilon^{(\phi)}(\epsilon') d\epsilon', & \tau_{n+1} = n+1. \end{cases}$$

Observe that the optimal allocation in the proposition can be rewritten as[14]

$$\tilde{\pi}_n^*(d) = \frac{\tilde{\mu}_{n+1}}{\gamma_n \sigma_{n+1}^2} \frac{\mu_n^a(d) - R_{n+1}\gamma_n\big(\mu_n^b(d) - (\mu_n^a(d))^2\big)}{\mu_n^b(d) + \Big(\frac{\tilde{\mu}_{n+1}}{\sigma_{n+1}}\Big)^2 \big(\mu_n^b(d) - (\mu_n^a(d))^2\big)}. \tag{4.7}$$

This shows that the optimal investment strategy is proportional to $\tilde{\mu}_{n+1}/(\gamma_n \sigma_{n+1}^2)$, which is the optimal allocation of a myopic investor with a single period objective function (see Eq. (4.8)). This fraction depends on the *current* economic conditions, and the client's *current* risk aversion. The proportionality factor in (4.7) depends on the *future* return of the investment strategy, between times $n+1$ and T, which in turn depends on both future economic conditions and the client's future risk aversion dynamics. Consistently with intuition, this factor is increasing in $\mu_n^a(d)$, the first moment of the future value of one dollar invested under the optimal strategy, and decreasing in its second moment, $\mu_n^b(d)$. In the final time period, the proportionality factor is equal to one, and the optimal allocation is given explicitly by

$$\tilde{\pi}_{T-1}^*(d) = \frac{\tilde{\mu}_T}{\gamma_{T-1}\sigma_T^2}, \tag{4.8}$$

which coincides with the standard single-period Markowitz investment strategy.

4.2 Comparative Statics for the Optimal Investment Strategy

In this section, we consider the risk aversion model in Section 3.5 and visualize how the optimal allocation strategy depends on the probability of idiosyncratic shocks to the client's risk aversion (p_ϵ), the client's behavioral biases (β), and the probability of economic state transitions (P).

The left panel of Figure 1 shows the impact of idiosyncratic shocks on the optimal risky asset allocation. The parameter p_ϵ determines the intensity of such shocks, and it is clear from the figure that if the client's idiosyncratic risk preferences are less stable, i.e. if p_ϵ is larger, then the optimal allocation to the risky asset deviates more from the case where there are no idiosyncratic shocks. The right panel of the figure shows the effect of the client's behavioral biases. Observe that the impact on the optimal allocation of negative market returns ($\gamma_n^Z > 1$) is larger than that of positive market returns ($\gamma_n^Z < 1$). This is consistent with the notion

[14]An extension of formula (4.7) can be derived for the general model considered in Proposition 4.1 (see Appendix A).

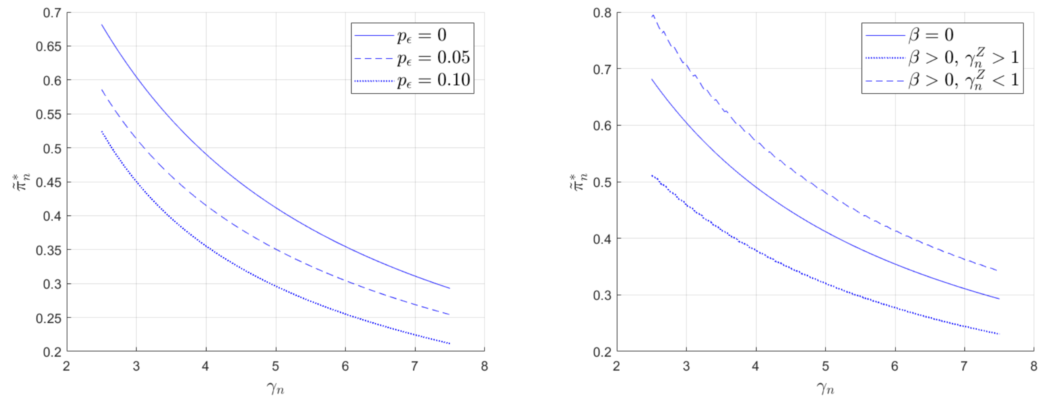

Figure 1: We consider the risk aversion model in Section 3.5 with a single economic state. The solid lines show the optimal allocations when $\alpha = p_\epsilon = \beta = 0$ (i.e., constant risk aversion process). In the left panel, we set $p_\epsilon \in \{0.05, 0.10\}$ and $\sigma_\epsilon = 0.64$. In the right panel, we set $\beta = 4$ and $\gamma_n^Z \in \{0.76, 1.30\}$, yielding a cumulative market return at an interaction time which is 1.3 standard deviations above/below its expected value. The annualized market parameters are $r = 0$, $\mu = 0.10$, and $\sigma = 0.20$. The investment horizon is $T = 36$ months, the time between interactions is $\phi = 3$ months. We show the allocations one year into the investment horizon, i.e., at $n = 13$.

of loss aversion.

Figure 2 shows how economic state transitions affect the optimal allocation in a two-state economy. If the economy is in the state with a lower market Sharpe ratio, then the amount invested in the risky asset is smaller if there is a positive probability that the economy will transit to the state with a higher market Sharpe ratio. Vice versa, if the economy is in the state with a higher market Sharpe ratio, the allocation to the risky asset is larger if there is a positive probability that the economy will transit to the state with a lower market Sharpe ratio. In other words, for a fixed risk aversion parameter, the optimal allocation adapts to the prevailing economic conditions, investing more (less) when the market Sharpe ratio is high (low).

5 Performance of the Robo-Advising System

We study the performance of the robo-advising model presented in Section 3.5. In Section 5.1, we study the effect of interaction on portfolio personalization, and show that a fundamental tradeoff arises between the rate of information arrival from the client and the client's level of behavioral biases. In Section 5.2, we study how the optimal investment strategy is affected by transitions between economic states and the resulting changes in the client's risk aversion. We also discuss how the robo-advisor can interfere to mitigate the effect of the client's behavioral biases.

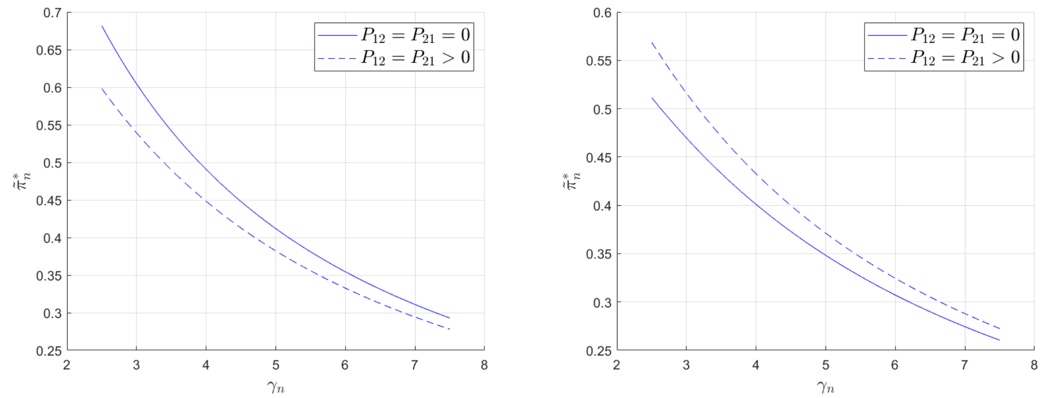

Figure 2: We consider the risk aversion model in Section 3.5, with two economic states and a state-independent risk aversion parameter ($\alpha = p_\epsilon = \beta = 0$, $\bar{\gamma} \equiv 1$). The solid lines show the optimal allocations when the probability of transitioning between states is zero. The dashed lines correspond to transition probabilities $P_{12} = P_{21} = 0.1$. The first state ($y = 1$, left panel) has the same market parameters as in Figure 1. For the second state ($y = 2$, right panel), we set $r(2) = 0$, $\mu(2) = 2\mu(1)$ and $\sigma(2) = \sqrt{2}\sigma(1)$. The investment horizon is $T = 36$ months. We graph the allocations one year into the investment horizon, i.e., at $n = 13$.

5.1 Interaction and Optimal Personalization

In our framework, the client and robo-advisor do not interact at all times, and the robo-advisor is therefore not always immediately aware of changes that have occurred to the client's risk profile. While a higher interaction frequency is the only way to reduce this information asymmetry, it also increases the effect of the client's behavioral biases in the communicated information. The goal of this section is to analyze this tradeoff quantitatively.

We start by defining a measure of portfolio personalization. This measure depends on the relation between the client's risk aversion process, $(\gamma_n^C)_{n\geq 0}$, given by (3.9), and the robo-advisor's view of the client's risk aversion, $(\gamma_n)_{n\geq 0}$, given by (3.10). These two processes coincide if and only if $\beta = 0$ and $\phi = 1$, i.e., if there are no behavioral biases, and the client and the robo-advisor interact at all times.

Recall from Section 4 that at time $n = T-1$, the optimal risky asset allocation is proportional to the client's *risk tolerance*, $1/\gamma_n$, and approximately proportional to $1/\gamma_n$ for $n < T-1$ (see (4.2) and (4.8)). We therefore choose to define personalization in terms of the difference between the risk tolerance parameters $1/\gamma_n$ and $1/\gamma_n^C$, and for $\phi \geq 1$ and $\beta \geq 0$ we define the measure

$$\mathcal{R}(\phi, \beta) := \mathbb{E}\left[\frac{1}{T}\sum_{n=0}^{T-1}\left|\frac{\frac{1}{\gamma_n} - \frac{1}{\gamma_n^C}}{\frac{1}{\gamma_n^C}}\right|\right], \tag{5.1}$$

which is the expected relative distance of $1/\gamma_n$ from $1/\gamma_n^C$, averaged throughout the investment horizon. A lower value of $\mathcal{R}(\phi, \beta)$ implies a higher level of portfolio personalization. Although the notation does not explicitly highlight it, the value of $\mathcal{R}(\phi,\beta)$ also depends on the parameters $p_\epsilon \in [0,1]$ and $\sigma_\epsilon > 0$, which govern the distribution of idiosyncratic risk aversion shocks.

For a given value of β, our objective is to study how the measure $\mathcal{R}$ depends on the value of ϕ, and to that end we consider an approximation $\widetilde{\mathcal{R}}$ that is analytically more tractable than $\mathcal{R}$. In the following proposition, we characterize this approximation and show that it is minimized by a unique value of ϕ. Figure 3 shows graphically that $\mathcal{R}$ and $\widetilde{\mathcal{R}}$ are close, for various interaction frequencies and levels of behavioral bias.

Proposition 5.1. *Set* $\sigma_0 := \sigma(Y_0)$.

(i) The measure $\mathcal{R}(\phi,\beta)$ *satisfies the relation*

$$\mathcal{R}(\phi,\beta) = \widetilde{\mathcal{R}}(\phi,\beta) + O(p_\epsilon^2\sigma_\epsilon^2 + \beta^2) + O(1 - P^{T-1}_{y_0,y_0}), \tag{5.2}$$

where

$$\widetilde{\mathcal{R}}(\phi,\beta) := \sqrt{\frac{2}{\pi}}\Big(\frac{\beta\sigma_0}{\sqrt{\phi}}\Big(1 - \frac{\phi-1}{2}p_\epsilon\Big) + \sqrt{\frac{\beta^2\sigma_0^2}{\phi} + \sigma_\epsilon^2}\frac{\phi-1}{2}p_\epsilon\Big).$$

(ii) For β *and* p_ϵ *such that* $\beta + p_\epsilon > 0$, *there exists a unique value of* $\phi \geq 1$ *that minimizes* $\widetilde{\mathcal{R}}(\beta,\phi)$[15]*:*

$$\underset{\phi\geq 1}{\arg\min}\,\widetilde{\mathcal{R}}(\phi,\beta) = \begin{cases} 1, & \beta = 0,\\ \phi_0 \in [1,\infty), & \beta > 0,\ p_\epsilon > 0,\\ \infty, & p_\epsilon = 0,\end{cases}$$

where ϕ_0 *satisfies*

$$\frac{\partial\phi_0}{\partial\beta} > 0, \qquad \frac{\partial\phi_0}{\partial\sigma_0} > 0, \qquad \frac{\partial\phi_0}{\partial p_\epsilon} < 0, \qquad \frac{\partial\phi_0}{\partial\sigma_\epsilon} < 0. \tag{5.3}$$

The error term $O(1 - P^{T-1}_{y_0,y_0})$ in the approximation (5.2) comes from fixing the state of the economy throughout the investment horizon. Thus, the above result approximates the optimal interaction frequency for fixed economic conditions. Typically, transitions between economic states are infrequent (see Eq. (5.5) in Section 5.2), and the error term vanishes if the model consists of a single economic state. From the signs of the derivatives given in (5.3), it can be seen that if the economy transits to a state with a higher return volatility, i.e., σ_0 goes up, the optimal time elapsing between consecutive interactions increases. This is because the client's behavioral biases are based on market return fluctuations, which are magnified if

[15]For $\beta = p_\epsilon = 0$, the values $\mathcal{R}(\phi,\beta)$ and $\widetilde{\mathcal{R}}(\phi,\beta)$ are independent of ϕ.

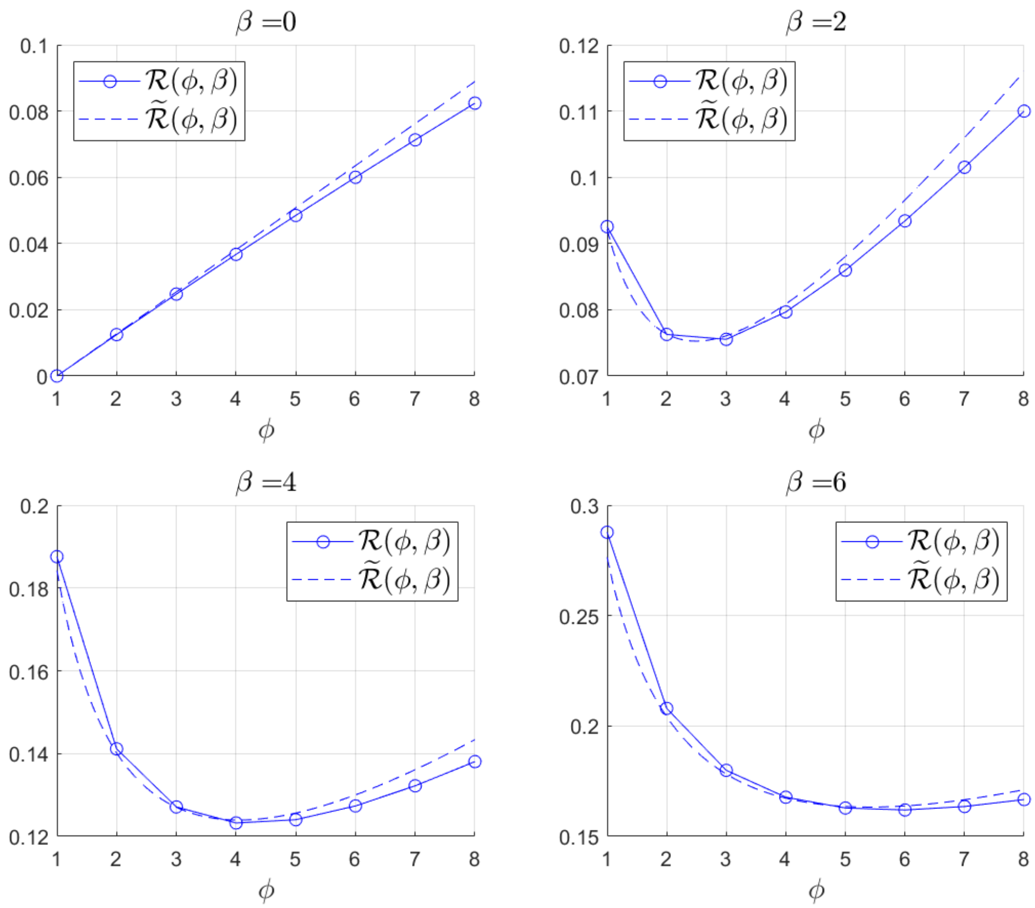

Figure 3: The graphs plot the personalization measure $\mathcal{R}(\phi, \beta)$ and its approximation $\widetilde{\mathcal{R}}(\phi, \beta)$ as a function of the time between interactions ϕ. Each panel corresponds to a different value of the parameter β. We consider a model consisting of a single economic state. The annualized market parameters are $r = 0$, $\mu = 0.10$, and $\sigma = 0.20$. Each time step corresponds to one month, and the risk aversion parameters are $\alpha = 0$, $p_\epsilon = 0.05$ and $\sigma_\epsilon = 0.64$.

volatility is higher.

The above proposition confirms two intuitive claims. First, in the absence of behavioral biases, i.e., $\beta = 0$, it is optimal for the robo-advisor to interact with the client at all times. Second, if there are no idiosyncratic risk aversion shocks, i.e., $p_\epsilon = 0$, then it is optimal to never interact. Most interestingly, if $\beta > 0$ and $p_\epsilon > 0$, then it may be suboptimal to interact at all times. This is because a higher frequency of interaction comes at the expense of increased behavioral bias in the communicated risk aversion parameter. Furthermore, the signs of the derivatives in (5.3) show that a larger value of β increases the optimal time between interactions, while larger values of p_ϵ and σ_ϵ, indicating a greater variance of idiosyncratic risk aversion shocks, push down the optimal time between interactions. In the proof of the proposition, we show that

$$\frac{p_\epsilon \sigma_\epsilon}{\beta \sigma_0} < 1 \quad \Longrightarrow \quad \frac{\partial \widetilde{\mathcal{R}}(\phi, \beta)}{\partial \phi}\Big|_{\phi=1} < 0,$$

which gives a sufficient condition for the optimal time between interactions to be greater than one.[16] In the above inequality, we compare the magnitude of the idiosyncratic component, measured by the product of the probability and volatility of the idiosyncratic risk aversion shocks, $p_\epsilon \sigma_\epsilon$, with the amount of behavioral bias, measured by the product of the client's sensitivity to abnormal returns and the volatility of returns, $\beta \sigma_0$. If the latter is greater, then the derivative of $\widetilde{\mathcal{R}}$ at $\phi = 1$ is negative, implying that it is suboptimal to interact at all times.

The robo-advisor's mean-variance criterion in (3.5) is parameterized by the risk aversion process $(\gamma_n)_{n \geq 0}$. To ensure high personalization, the robo-advisor needs to construct a process $(\gamma_n)_{n \geq 0}$ which is as close as possible to the client's actual risk aversion process $(\gamma_n^C)_{n \geq 0}$. The outcome of the mean-variance optimization is an investment strategy, and a related measure of personalization is the proximity of the investment strategies corresponding to $(\gamma_n)_{n \geq 0}$ and $(\gamma_n^C)_{n \geq 0}$. The latter is the benchmark strategy corresponding to full personalization, and is only attained if interaction happens at all times and there are no behavioral biases, in which case the two risk aversion processes also coincide.

Denote by $(\tilde{\pi}_n^*(\gamma_n))_{0 \leq n < T}$ and $(\tilde{\pi}_n^*(\gamma_n^C))_{0 \leq n < T}$ the optimal allocations corresponding to $(\gamma_n)_{0 \leq n < T}$ and $(\gamma_n^C)_{0 \leq n < T}$, respectively. In line with the definition of $\mathcal{R}(\phi, \beta)$ in (5.1), we then define

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

as a measure of the difference between $(\tilde{\pi}_n^*(\gamma_n))_{0 \leq n < T}$ and $(\tilde{\pi}_n^*(\gamma_n^C))_{0 \leq n < T}$. A smaller value of $\mathcal{S}(\phi, \beta)$ implies a higher level of personalization, while $\mathcal{S}(\phi, \beta)$ is equal to zero if and only if $\phi = 1$ and $\beta = 0$. In general, the difference between $\tilde{\pi}_n^*(\gamma_n)$ and $\tilde{\pi}_n^*(\gamma_n^C)$ is a nonlinear and rather complex function of the difference between γ_n and γ_n^C. Nevertheless, in Appendix B we show that

$$\mathcal{S}(\phi, \beta) = \mathcal{R}(\phi, \beta) + O((\phi - 1) p_\epsilon \sigma_\epsilon^2) + O\Big(\frac{\beta^2 \sigma^2}{\phi^2}\Big). \tag{5.4}$$

Hence, a small difference between $(\gamma_n)_{0 \leq n < T}$ and $(\gamma_n^C)_{0 \leq n < T}$ translates into a small difference between the corresponding investment strategies. In the above equation, the first error term is due to the idiosyncratic component of the client's risk aversion, and it vanishes if $\phi = 1$. The second error term is due to the client's

[16] It is also close to being a necessary condition. See Eq. (B.1) in Appendix B.

behavioral biases, and it is maximized at $\phi = 1$.

Figure 4 indicates that the value of ϕ that minimizes $\mathcal{R}(\phi, \beta)$ gives a lower bound for the value of ϕ that minimizes $\mathcal{S}(\phi, \beta)$. In other words, minimizing $\mathcal{R}(\phi, \beta)$ provides a conservative estimate for the time between interactions that minimizes the allocation difference $\mathcal{S}(\phi, \beta)$. This is because $\mathcal{R}(\phi, \beta)$ underestimates the impact of the client's behavior biases. Specifically, $\mathcal{R}(\phi, \beta)$ is equal to the average relative difference between $(\gamma_n^C)_{n\geq 0}$ and $(\gamma_n)_{n\geq 0}$ in a single time interval between two consecutive interactions, and its value is the same for all such intervals. Instead, the optimal allocation at any given time depends on all future allocations, and thus on the future path of the risk aversion process. As a result, the behavioral bias in investment decisions made at future interaction times feeds into the allocation decisions at earlier times, and the measure $\mathcal{S}(\phi, \beta)$ accounts for this effect.

Remark 5.2. A uniform interaction schedule yields a conservative estimate for the optimal interaction frequency. If the client were allowed to interact with the robo-advisor at times triggered by market conditions, then she is more likely to do so following a period of extreme returns. The client would either be overly exuberant following a period of positive returns, or overly pessimistic after a period of negative returns. This magnifies the impact of the client's behavioral bias on the investment decisions relative to a uniform schedule, and would result in a larger value for the optimal period ϕ between consecutive interactions. □

5.2 Impact of Economic Transitions on Risk Aversion and Sharpe Ratios

In this section, we study how the optimal investment strategy is affected by transitions between economic states and by the associated shifts in the client's risk aversion. We consider an economy consisting of two states, one corresponding to times of economic growth, and the other corresponding to times of recession. This choice is supported by the methodology of the National Bureau of Economic Research (NBER), which splits business cycles into periods of economic expansions and periods of economic contractions. Assuming each period to last one month, we set the transition matrix of $(Y_n)_{n\geq 0}$ to

$$P = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix} = \begin{bmatrix} 0.95 & 0.05 \\ 0.10 & 0.90 \end{bmatrix}, \tag{5.5}$$

and use $y = 1$ to denote the state of economic growth, and $y = 2$ to denote the state of recession.[17] We use values in accordance with Tang and Whitelaw [2011] for the state-dependent return parameters. They

[17] These transition probabilities are based on empirical values reported in Chauvet and Hamilton [2006], who also show that a regime switching model captures well what is being described by NBER's business cycle chronology. These values show that, on average, economic expansions last longer than contractions.

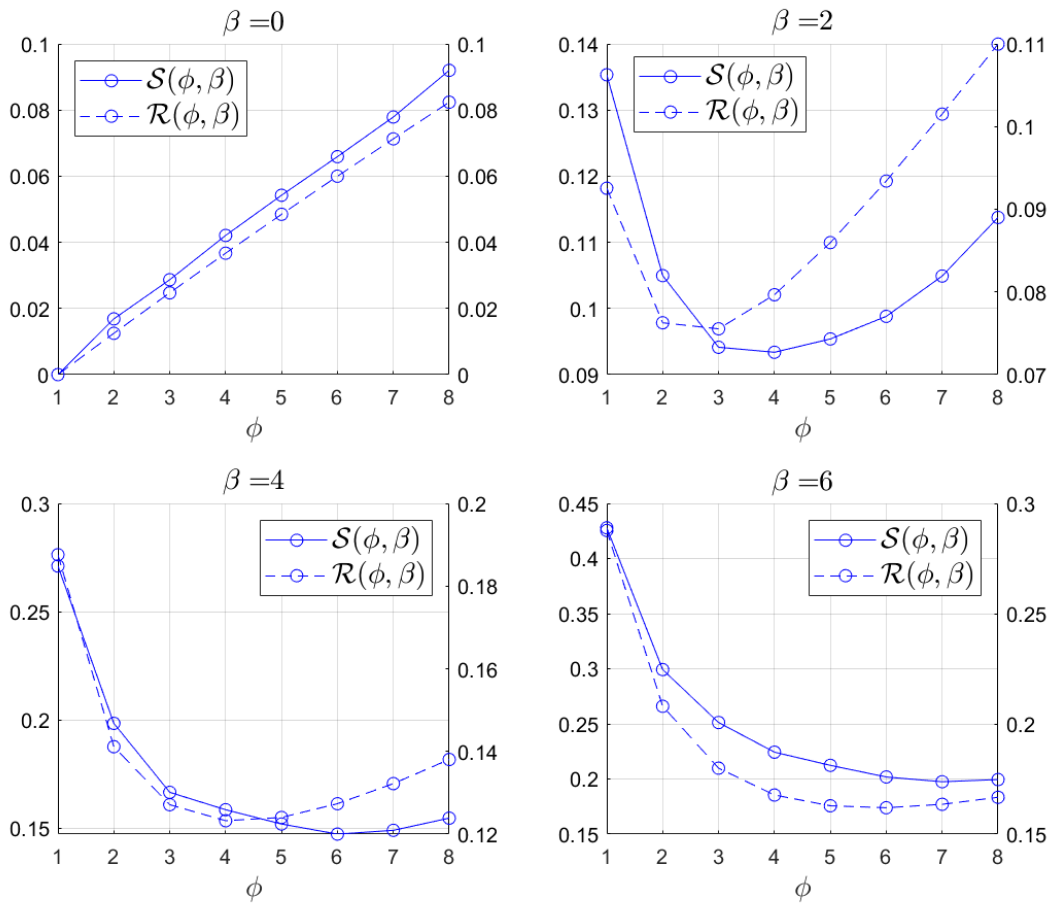

Figure 4: The graphs plot the measures $\mathcal{S}(\phi,\beta)$ (left y-axis) and $\mathcal{R}(\phi,\beta)$ (right y-axis) as a function of the time between interactions ϕ. Each panel corresponds to a different value of the parameter β. The market and risk aversion parameters are the same as in Figure 3.

report historical averages for changes in the mean and volatility of stock market returns, as well as changes in the market Sharpe ratio, between the peak of the business cycle and the subsequent trough. Specifically, we set the annual risk-free rate and the mean and volatility of market returns to

$$r = (0.015, 0), \qquad \mu = (0.081, 0.137), \qquad \sigma = (0.155, 0.173). \tag{5.6}$$

We also set $\tilde{\mu} := \mu - r$, and define the parameters $a := \tilde{\mu}(2)/\tilde{\mu}(1)$ and $b := \sigma(2)/\sigma(1)$, which quantify the relative change in the excess return moments between states.

To focus exclusively on the implications of economic transitions on investment decisions, we make the client's risk aversion depend only on the current economic state. Specifically, we assume $\alpha = 0$, $p_\epsilon = 0$ and $\beta = 0$, i.e., no temporal component, idiosyncratic risk aversion shocks, or behavioral biases. The risk aversion function at time n then takes the form $\gamma_n = \bar{\gamma}(Y_n)$, where $\bar{\gamma} : \mathcal{Y} \mapsto \mathbb{R}_0^+$ is a state-dependent risk aversion coefficient. We allow for a slightly more general functional form, where $\bar{\gamma}$ is allowed to be time-dependent, and calibrated so that the resulting optimal allocation $(\tilde{\pi}_n^*)_{0 \le n < T}$ is time-homogeneous, given the current economic state. Specifically, in Appendix B (see Proposition B.1) we show that for a given $\bar{\pi} > 0$ and $\delta > -1$, there exists a unique risk aversion process such that the corresponding optimal allocation is given by

$$\tilde{\pi}_n^*(y) = \begin{cases} \bar{\pi}, & y = 1, \\ \bar{\pi}(1+\delta), & y = 2. \end{cases} \tag{5.7}$$

That is, in times of growth, the client's risky asset allocation is equal to $\bar{\pi}$,[18] and when the economy transits to the recessionary regime, the allocation becomes $\bar{\pi}(1+\delta)$, where the value of δ determines the allocation change. A negative value of δ captures a client who shifts wealth away from the risky asset in a state of recession, when the risk-return tradeoff is favorable. As discussed in Section 3.5, this behavior is typical of retail investors. A positive value of δ means that the client invests more in the risky asset when the return per unit risk is high. Finally, if $\delta = 0$, the client holds the same portfolio in both states of the economy.

The Sharpe ratio achieved by this investment strategy, defined as its expected excess return over the volatility of its excess returns, can be computed explicitly (see Appendix B) and is given by

$$s^{\pi^*}(\delta) = \frac{1}{\sqrt{\frac{\sigma^2(1)}{\tilde{\mu}^2(1)} \frac{1-\lambda+\lambda b^2(1+\delta)^2}{(1+\lambda(a(1+\delta)-1))^2} + \frac{1-\lambda+\lambda a^2(1+\delta)^2}{(1+\lambda(a(1+\delta)-1))^2} - 1}}, \tag{5.8}$$

where λ is the stationary probability of the recessionary state.[19]

[18] For instance, $\bar{\pi} = 0.60$ corresponds to the classical 60/40 portfolio composition. This strategy was popularized by Jack Bogle, the founder of Vanguard, and is commonly used as a benchmark in portfolio allocation.

[19] The Markov chain $(Y_n)_{n \ge 0}$ has a unique stationary distribution. For the transition matrix in (5.5), the stationary distri-

Figure 5 indicates that for a fixed value of δ, the Sharpe ratio (5.8) is increasing in λ, i.e., the long-run average time spent in the recessionary state. It is also increasing in a, which is the relative change in mean excess returns between the two states, and decreasing in b, i.e., in the relative change in the return volatility. These monotonicity properties turn out to hold under mild conditions, which are given in Lemma B.2 of Appendix B.

Note that a higher value of δ implies a larger proportion of wealth allocated to the risky asset when the market Sharpe ratio is high. Figure 5 shows that the portfolio's Sharpe ratio is, in general, increasing in δ, but as the third panel shows, a positive value of δ may result in a smaller Sharpe ratio than if $\delta = 0$. In Appendix B, we show that

$$\frac{\partial s^{\pi^*}(\delta)}{\partial \delta} > 0 \quad \Longleftrightarrow \quad 1 + \frac{\sigma^2(1)}{\tilde{\mu}^2(1)} > a\Big(1 + \frac{b^2\sigma^2(1)}{a^2\tilde{\mu}^2(1)}\Big)(1+\delta), \tag{5.9}$$

which, for $\delta = 0$, provides a condition for when the portfolio's Sharpe ratio increases by shifting the constant portfolio weights towards the risky asset, at times when the market Sharpe ratio is high. Whether or not this inequality is satisfied depends on the values of a and b, which determine the difference in market Sharpe ratios between the two economic states. In particular, for $\delta = 0$ and a fixed $b > 1$, the condition shows that there exists a threshold $a^*(b) > 1$ that the parameter a needs to exceed for the condition to be satisfied. We also observe that the right-hand side of the inequality is increasing in δ. Hence, for $\delta > 0$ a threshold larger than $a^*(b)$ is required for the condition to be satisfied, while a smaller threshold is sufficient for $\delta < 0$.

We observe that condition (5.9) is satisfied for parameter values observed in practice. For example, it is satisfied by the parameter values given in (5.6) if $\delta = 0$. That is, in general, the portfolio's Sharpe ratio increases if one allocates more to the risky asset when the market Sharpe ratio is high. However, Figure 5 indicates that the gain is modest, and that the Sharpe ratio decreases more when the allocation to the risky asset is decreased, compared to how much it increases when the allocation to the risky asset is increased by the same amount. We show in Appendix B that $s^{\pi^*}(\delta)$ is concave in a neighborhood of $\delta = 0$, i.e.,

$$\frac{\partial^2 s^{\pi^*}(\delta)}{\partial \delta^2}\bigg|_{\delta=0} \propto -\frac{\sigma^2(1)}{\tilde{\mu}^2(1)}\Big(a^2 + b^2\frac{\sigma^2(1)}{\tilde{\mu}^2(1)}\Big)\lambda + O(\lambda^2), \tag{5.10}$$

hence confirming that, for the Sharpe ratio of the portfolio to be satisfactory, it is more important not to reduce the risky asset allocation in a state of recession, than it is to increase it.

The portfolio's Sharpe ratio is determined by the first two moments of the portfolio returns, which may not characterize the return distribution. Comparing portfolios in terms of their Sharpe ratios therefore leaves

bution is $(P_{21}, P_{12})/(P_{12} + P_{21}) = (2/3, 1/3)$.

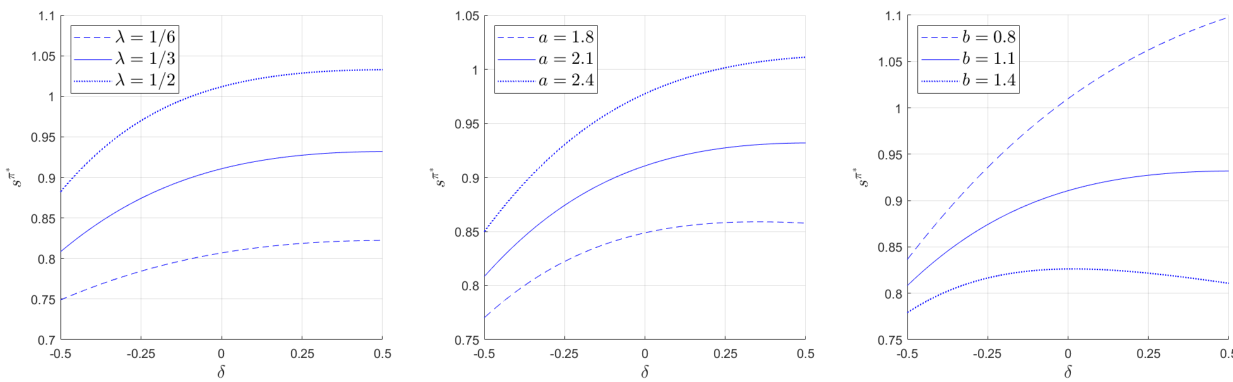

Figure 5: Annualized Sharpe ratios of the strategy given in Eq. (5.7), for $\bar{\pi} = 0.6$ and different values of δ. The market parameters are given in (5.5)-(5.6), and the portfolio is rebalanced monthly. These parameters imply $\lambda = 1/3$, $a = 2.1$, and $b = 1.1$.

out the effect of higher return moments. In a model with multiple economic states, the unconditional return distribution is leptokurtic and skewed, and we use simulation to compare the terminal distribution of the client's wealth, for different values of the parameter δ (see Remark 5.3).

Figure 6 shows the simulated distribution of the terminal return of the optimal investment strategy, and the corresponding annualized rate of return. It is clear from the figure that the skewness and kurtosis are higher when the risky asset allocation is maintained or increased in a state of recession, and that in those cases, the return upside is considerably higher, with a limited additional downside risk. The same can be observed from the summary statistics of the simulations, reported in Table 1. Interestingly, the sample mean divided by the sample standard deviation is the lowest for the strategy corresponding to $\delta > 0$, even though the sample mean return is the highest. This is because the strategy corresponding to $\delta > 0$ also has a higher standard deviation, which penalizes equally for upside and downside volatility.

Remark 5.3. The return distributions in Figure 6 are estimated using simulation, because the distribution of $r^{\pi^*}_{n,T}$ does not admit an explicit expression. Equation (4.4) shows how the first two moments of $1 + r^{\pi^*}_{n,T}$ can be computed recursively and, in the same way, the m-th finite moment, $\mu^{(m)}_n := \mathbb{E}_{n,d}[(1 + r^{\pi^*}_{n,T})^m]$, can be computed using the recursion

$$\mu^{(m)}_n = \mathbb{E}_{n,d}\big[\big(R_{n+1} + \widetilde{Z}_{n+1}\tilde{\pi}^*_n(d)\big)^m \mu^{(m)}_{n+1}\big].$$

Under mild conditions, a probability distribution with finite moments of all orders is uniquely determined by its moments (see Billingsley [1995], Sect. 30). In that case, an alternative to simulation is to recover the distribution of $r^{\pi^*}_{n,T}$ from a finite set of moments $\mu^{(1)}_n, \ldots, \mu^{(M)}_n$, for some $M \geq 1$ (see John et al. [2007] for an overview of different reconstruction approaches). $\square$

We note that modifying the portfolio allocation based on economic conditions, i.e., $\delta \neq 0$, is a form of active management, and robo-advising firms which focus on long-term investing in general do not engage in such market timing.[20] Rather, they urge clients to stay the course through changing market conditions, in order to reap the benefits of long-term investing. This corresponds to setting $\delta = 0$, and the return distributions in Table 1 and Figure 6 highlight that there are significant benefits of doing that, compared to reducing the risky asset allocation in recessions which corresponds to $\delta < 0$.

Even though the benefits are even greater if the risky asset allocation is increased during recessions, the question arises of how far the robo-advisor can reach "against the will" of the client. While higher returns will be obtained in the long run, the client may suffer from adverse market moves in the short run, and may not have an understanding of the long-term benefits. In a state of recession, with worsening economic outlook and risk aversion rising, the client may be particularly sensitive to what can be perceived as an investment mistake of the robo-advisor. The importance of this dilemma faced by robo-advisors is also emphasized by Rossi and Utkus [2019], who show empirically that algorithm aversion, i.e., the tendency of individuals to prefer a human forecaster over an algorithm, and to more quickly lose confidence in an algorithm than in a human after observing them make the same mistake (Dietvorst et al. [2015]), is one of the main obstacles for switching to robo-advising.

The above discussion indicates that a good middle ground for the robo-advisor is to encourage the client to simply maintain a fixed portfolio composition, which is consistent with the long-term investing principle of riding out business cycles. To that end, statistics such as those in Table 1 and Figure 6 can be presented to the client, showing that over time, higher returns can be earned with limited additional risk. While this may entail additional risk for the client, empirical research has found that retail investors are indeed willing to take above average risk in order to earn above average returns (Guiso et al. [2018]). Related to that, we remark that even though a "buy-and-hold" portfolio has constant relative allocations, it does imply a fluctuating level of risk, as measured by, e.g., the portfolio returns volatility. Specifically, the risk tends to be higher when the market Sharpe ratio is high, i.e., when the reward per unit risk is high, because of increased stock market volatility.[21]

[20] One of the claimed benefits of robo-advising is that by managing the portfolio on the client's behalf, the client is helped to resist the temptation of attempting to time the market. The robo-advising firms Betterment and Wealthfront also cite empirical work, such as Dalbar's annual *Quantitative Analysis of Investors* report, which shows that investors that try to time the market tend to do much worse than a buy-and-hold investor, and that the average investor significantly underperforms a broad stock market index. Betterment additionally argues that due to short-term capital gains taxes, a market timing investor needs to substantially beat the stock market index, just to break even.

[21] Higher market Sharpe ratio during recessions is associated with both a higher market risk premium and higher return volatility, with the former effect dominating (see Guo and Whitelaw [2006] and Tang and Whitelaw [2011]).

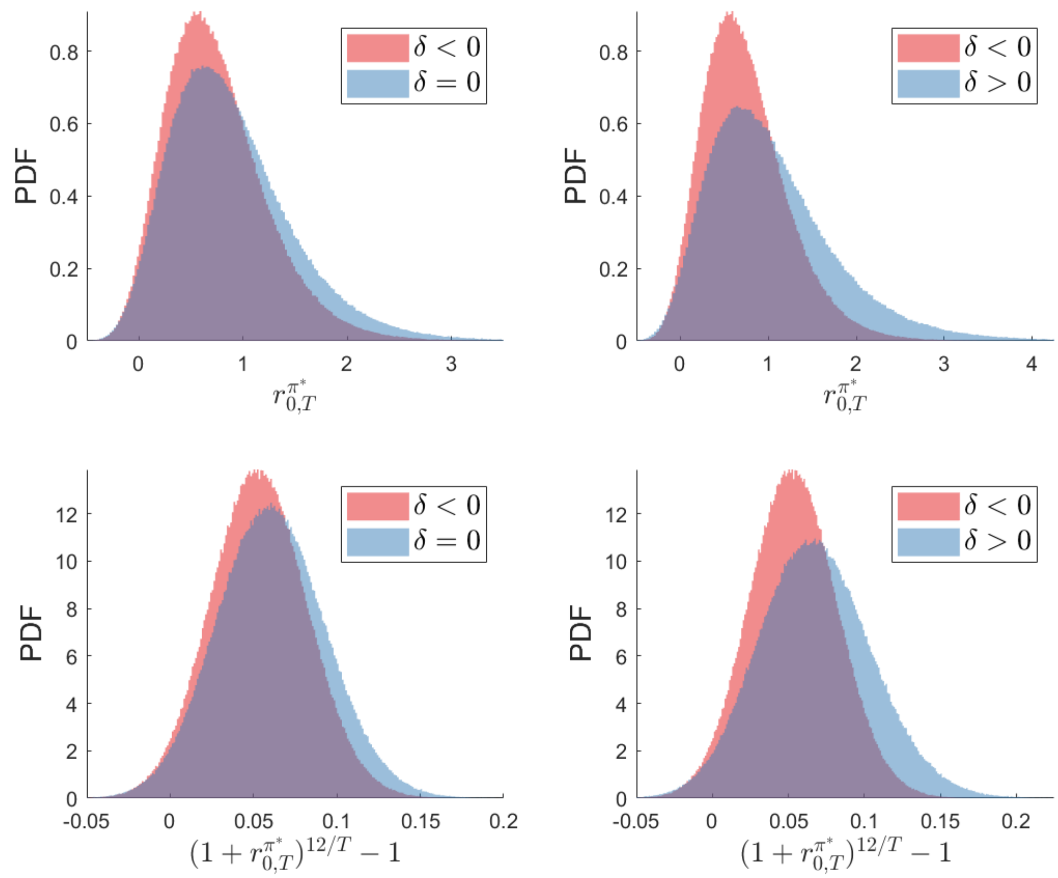

Figure 6: Simulated distributions of the return on one dollar invested in the optimal strategy between time 0 and time T (top panels), and the corresponding annualized rate of return (bottom panels). The market parameters are given in (5.5)-(5.6), the initial economic state is $Y_0 = 1$, and the investment horizon is $T = 120$ months. The risky asset allocation is given by (5.7) with $\bar{\pi} = 0.6$ and $\delta = -0.3$, $\delta = 0$, and $\delta = 0.3$.

	Mean	SD	Mean/SD	Skewness	Kurtosis	90% VaR	95% VaR	99% VaR
$\delta < 0$	0.740	0.485	1.526	0.838	4.257	-0.179	-0.067	0.117
$\delta = 0$	0.881	0.589	1.495	0.970	4.712	-0.213	-0.085	0.118
$\delta > 0$	1.036	0.735	1.409	1.231	5.934	-0.233	-0.091	0.134
$\delta < 0$	0.053	0.029	1.849	0.066	3.011	-0.017	-0.006	0.012
$\delta = 0$	0.061	0.032	1.871	0.092	3.016	-0.019	-0.008	0.012
$\delta > 0$	0.068	0.037	1.833	0.162	3.089	-0.021	-0.009	0.014

Table 1: Summary statistics for the simulated return distributions in Figure 6. The top half of the table reports the statistics for the return on one dollar invested in the strategy throughout the investment horizon. The bottom half of the table gives the statistics for the annualized rate of return.

6 Concluding Remarks

The past decade has witnessed the emergence of robo-advisors, which are investment platforms through which clients interact directly with an investment algorithm, with limited or no human intervention. Recent literature has provided empirical evidence on the characteristics of those investors that are more likely to switch to robo-advising, and how robo-advising tools affect the composition of investors' portfolios.

In this work, we build a novel framework that incorporates important features of modern robo-advising systems. This framework is in the form of a human-machine interaction system, where the robo-advisor (i.e., the machine) has the task of optimally investing the client's (i.e., the human's) wealth. The client has a risk profile that varies with time and to which the robo-advisor's investment performance criterion dynamically adapts. We derive and analyze optimal investment strategies, with the client's risk profile changing in response to market returns, economic conditions, and idiosyncratic events.

The frequency of human-machine interaction determines the level of portfolio personalization, which is defined in terms of how accurately the robo-advisor is able to model the client's risk profile. We show that there exists an optimal interaction frequency, which maximizes portfolio personalization by striking a balance between receiving information from the client in a timely manner, and mitigating the effect of the client's behavioral biases. We quantify how the Sharpe ratio of the optimal investment strategy depends on the allocation of wealth within economic regimes. We study the impact on portfolio performance resulting from the robo-advisor increasing the client's stock market allocation during periods of high market risk-return tradeoff. If left unassisted, the client is inclined to do the exact opposite due to the countercyclical nature of both risk aversion and the market risk-return tradeoff. We show that by simply rebalancing the portfolio to maintain constant weights throughout the business cycle, the portfolio's Sharpe ratio is near optimal and the portfolio return distribution significantly improved.

Our model can be extended along several directions, and aligned to the practices of operational robo-advising systems. First, the optimal allocation formula is based on a continuous range of risk aversion values. However, in practice, robo-advising firms manage a finite number of portfolios, grouping clients with similar risk profiles into risk categories (for instance, Wealthfront constructs a composite Risk Score ranging from 0.5, i.e., most risk averse, to 10.0, i.e., most risk tolerant, in increments of 0.5). To meet this requirement, the risk aversion process in our model can be structured to take values in a finite subset of the positive real numbers. Second, the market model can be extended to include multiple tradable assets and the portfolio rebalancing times can be allowed to be random. For instance, rebalancing could be triggered by the portfolio weights having drifted significantly from the target weights, due to moves in the prices of the underlying securities. Robo-advisors generally use such threshold updating rules to rebalance their portfolios (Beketov et al. [2018]), while also monitoring for rebalancing opportunities in order to minimize expense ratios and maximize tax efficiency.

The regime switching model can be extended to make the economic state factor hidden, in line with the fact that economic conditions are typically not directly observable. With multiple asset classes, empirical work has also shown that more than two economic regimes are needed to capture the joint distribution of asset returns (Guidolin and Timmermann [2006]). With a hidden factor, market timing based on economic conditions can only be done probabilistically, opening the door for misclassification of the prevailing economic conditions, and the robo-advisor acting against the client's best interests. This implies that in a hidden Markov model, there may be a stronger case for the buy-and-hold portfolio strategy.

Additionally, in the current setting, market returns are assumed to be conditionally independent of economic state transitions, but such transitions are commonly associated with large market moves, in particular when entering a recession.[22] This kind of market moves would effectively result in the client buying high and selling low, exacerbating the adverse effect of the client's risk aversion being positively correlated to the market Sharpe ratio, and making a market timing client even less likely to reach her investment goals. At the same time, asset dynamics of this nature would provide a further incentive for attempting to quickly detect changes in economic conditions. For example, if a market drop was likely to happen *after* a change in economic conditions, the portfolio's market exposure could be adjusted to account for the probability of a near term market drop.[23]

[22] This is also the case in the consumption-based asset pricing frameworks capable of providing countercyclical variation in risk aversion, such as Campbell and Cochrane [1999]. In their framework, worsening economic outlook leads to an increase in risk aversion, and investors requesting a higher reward for carrying risk, which leads to stock prices falling.

[23] Portfolio allocation under regime switching is widely studied. Related to our work, Tang and Whitelaw [2011] show how simple market timing strategies can exploit the fact that Sharpe ratios are countercyclical, and Coudert and Gex [2008] show that risk aversion indices published by financial institutions are good leading indicators of stock market crises. Kritzman et al. [2012] show how a Markov switching model can be used to avoid large portfolio losses related to "event regimes" where asset price start deviating from their past dynamics, and during which market risk premium is low.

A Proofs of Results Related to Section 4

We start by proving a lemma that will be used in the proof of Proposition 4.1.

Lemma A.1. *Let $\pi = (\pi_n)_{n\geq 0}$ be an admissible control law of the form*

$$\pi_n(x, d) = \tilde{\pi}_n(d)x, \tag{A.1}$$

for any $x \in \mathbb{R}$, and $d \in \mathcal{D}_n$. Then,

$$\mathbb{E}_{n,x,d}\Big[\frac{X_T^\pi}{x}\Big] = a_n(d), \qquad \mathbb{E}_{n,x,d}\Big[\Big(\frac{X_T^\pi}{x}\Big)^2\Big] = b_n(d),$$

where the functions $a_n(d)$ and $b_n(d)$ do not depend on x, and satisfy the recursions (4.4).

Proof of Lemma A.1: For $n = T - 1$, it follows from the wealth dynamics (3.1) that

$$\mathbb{E}_{n,x,d}\Big[\frac{X_T^\pi}{x}\Big] = \mathbb{E}_{n,x,d}[R_T + \widetilde{Z}_T\tilde{\pi}_n(d)] = R_T + \tilde{\mu}_T\tilde{\pi}_n(d) =: a_n(d).$$

Next, let $n \in \{0, 1, \dots, T-2\}$ and assume that the result holds for $n+1, n+2, \dots, T-1$. Then,

$$\begin{aligned}
\mathbb{E}_{n,x,d}\Big[\frac{X_T^\pi}{x}\Big] &= \mathbb{E}_{n,x,d}\Big[(R_{n+1} + \widetilde{Z}_{n+1}\tilde{\pi}_n(d))\prod_{k=n+1}^{T-1}(R_{k+1} + \widetilde{Z}_{k+1}\tilde{\pi}_k(D_k))\Big] \\
&= \mathbb{E}_{n,x,d}\Big[(R_{n+1} + \widetilde{Z}_{n+1}\tilde{\pi}_n(d))\mathbb{E}_{n+1,X^\pi_{n+1},D_{n+1}}\Big[\prod_{k=n+1}^{T-1}(R_{k+1} + \widetilde{Z}_{k+1}\tilde{\pi}_k(D_k))\Big]\Big] \\
&= \mathbb{E}_{n,x,d}[(R_{n+1} + \widetilde{Z}_{n+1}\tilde{\pi}_n(d))a_{n+1}(D_{n+1})] \\
&= \mathbb{E}_{n,d}[(R_{n+1} + \widetilde{Z}_{n+1}\tilde{\pi}_n(d))a_{n+1}(D_{n+1})] \\
&=: a_n(d).
\end{aligned}$$

The assertion for the b_n-sequence is shown in the same way. □

Next, we derive in the following proposition the extended HJB system of equations necessarily satisfied by an optimal control law for the optimization problem defined by (3.5)-(3.7). To simplify notation, we introduce for each n the random variable

$$\widetilde{D}_n := (X_n, D_n) \in \widetilde{\mathcal{D}}_{n+1} := \mathbb{R} \times \mathcal{D}_n.$$

Furthermore, we write

$$\widetilde{D}^{\pi}_{n+1} = (X^{\pi}_{n+1}, D_{n+1}),$$

where X^{π}_{n+1} is obtained by applying the control law π to the wealth x at time n.

Proposition A.2. *Assume that an optimal control law π^* for the optimization problem (3.5)-(3.7) exists. Then, the value function (3.8) satisfies the recursive equation*

$$\begin{aligned} V_n(\tilde{d}) = \sup_{\pi} \Big\{ \mathbb{E}_{n,\tilde{d}}[V_{n+1}(\widetilde{D}^{\pi}_{n+1})] - \Big(\mathbb{E}_{n,\tilde{d}}[f_{n+1,n+1}(\widetilde{D}^{\pi}_{n+1}; \widetilde{D}^{\pi}_{n+1})] - \mathbb{E}_{n,\tilde{d}}[f_{n+1,n}(\widetilde{D}^{\pi}_{n+1}; \tilde{d})] \Big) \\ - \Big(\mathbb{E}_{n,\tilde{d}}\Big[\frac{\gamma_{n+1}(D_{n+1})}{2} \Big(\frac{g_{n+1}(\widetilde{D}^{\pi}_{n+1})}{X^{\pi}_{n+1}} \Big)^2 \Big] - \frac{\gamma_n(d)}{2} \Big(\mathbb{E}_{n,\tilde{d}}\Big[\frac{g_{n+1}(\widetilde{D}^{\pi}_{n+1})}{x} \Big] \Big)^2 \Big\}, \end{aligned} \tag{A.2}$$

for $0 \le n < T$ and $\tilde{d} = (x, d) \in \widetilde{\mathcal{D}}_n$, with the terminal condition

$$V_T(\tilde{d}) = 0, \quad \tilde{d} \in \widetilde{\mathcal{D}}_T.$$

Herein, for any fixed $k = 0, 1, \ldots, T$ and $\tilde{d}' = (x', d') \in \widetilde{\mathcal{D}}_k$, the function sequence $(f_{n,k}(\cdot; \tilde{d}'))_{0 \le n \le T}$, where $f_{n,k}(\cdot; \tilde{d}') : \widetilde{\mathcal{D}}_n \mapsto \mathbb{R}$, is determined by the recursion

$$\begin{aligned} f_{n,k}(\tilde{d}; \tilde{d}') &= \mathbb{E}_{n,\tilde{d}}[f_{n+1,k}(\widetilde{D}^{\pi^*}_{n+1}; \tilde{d}')], \qquad \tilde{d} \in \widetilde{\mathcal{D}}_n, \quad 0 \le n < T, \\ f_{T,k}(\tilde{d}; \tilde{d}') &= \frac{x_T}{x'} - 1 - \frac{\gamma_k(d')}{2} \Big(\frac{x_T}{x'} \Big)^2, \quad \tilde{d} \in \widetilde{\mathcal{D}}_T, \end{aligned}$$

and the function sequence $(g_n)_{0 \le n \le T}$, where $g_n : \widetilde{\mathcal{D}}_n \mapsto \mathbb{R}$, is determined by the recursion

$$\begin{aligned} g_n(\tilde{d}) &= \mathbb{E}_{n,\tilde{d}}[g_{n+1}(\widetilde{D}^{\pi^*}_{n+1})], \quad \tilde{d} \in \widetilde{\mathcal{D}}_n, \quad 0 \le n < T, \\ g_T(\tilde{d}) &= x, \quad \tilde{d} \in \widetilde{\mathcal{D}}_T. \end{aligned}$$

Furthermore, the function sequences admit the probabilistic representations

$$f_{n,k}(\tilde{d}; \tilde{d}') = \mathbb{E}_{n,\tilde{d}}\Big[\frac{X^{\pi^*}_T}{x'} - 1 - \frac{\gamma_k(d')}{2} \Big(\frac{X^{\pi^*}_T}{x'} \Big)^2 \Big], \qquad g_n(\tilde{d}) = \mathbb{E}_{n,d}[X^{\pi^*}_T].$$

Proof of Proposition A.2: We write the objective functions at time n as

$$J_n(\tilde{d}; \pi) = \mathbb{E}_{n,\tilde{d}}[F_n(\tilde{d}, X^{\pi}_T)] + G_n(\tilde{d}, \mathbb{E}_{n,\tilde{d}}[X^{\pi}_T]),$$

where, for $\tilde{d} \in \widetilde{\mathcal{D}}_n$, and $y \in \mathbb{R}$,

$$F_n(\tilde{d}, y) = \frac{y}{x} - 1 - \frac{\gamma_n(\tilde{d})}{2}\Big(\frac{y}{x}\Big)^2, \qquad G_n(\tilde{d}, y) = \frac{\gamma_n(\tilde{d})}{2}\Big(\frac{y}{x}\Big)^2. \tag{A.3}$$

The proof now consists of two parts. First, we derive the recursive equation satisfied by the sequence of objective functions for any admissible control law π. Then, we derive the system of equations that an optimal control law π^* must satisfy.

Step 1: Recursion for $J_n(\tilde{d}; \pi)$. For a given admissible control law π, and fixed $0 \leq k \leq T$ and $\tilde{d}' \in \widetilde{\mathcal{D}}_k$, we define the function sequences

$$\begin{aligned} f^{\pi}_{n,k}(\tilde{d}; \tilde{d}') &:= \mathbb{E}_{n,\tilde{d}}[F_k(\tilde{d}', X^{\pi}_T)], \quad && \tilde{d} \in \widetilde{\mathcal{D}}_n, \quad 0 \leq n \leq T, \\ g^{\pi}_n(\tilde{d}) &:= \mathbb{E}_{n,\tilde{d}}[X^{\pi}_T], && \tilde{d} \in \widetilde{\mathcal{D}}_n, \quad 0 \leq n \leq T, \end{aligned} \tag{A.4}$$

and note that, by the law of iterated expectations, we have

$$f^{\pi}_{n,k}(\tilde{d}; \tilde{d}') = \mathbb{E}_{n,\tilde{d}}[f^{\pi}_{n+1,k}(\widetilde{D}^{\pi}_{n+1}; \tilde{d}')], \qquad g^{\pi}_n(\tilde{d}) = \mathbb{E}_{n,\tilde{d}}[g^{\pi}_{n+1}(\widetilde{D}^{\pi}_{n+1})]. \tag{A.5}$$

For $\widetilde{D}_{n+1} \in \widetilde{\mathcal{D}}_{n+1}$, the objective function at time $n+1$ can then be written as

$$\begin{aligned} J_{n+1}(\widetilde{D}_{n+1}; \pi) &= \mathbb{E}_{n+1,\widetilde{D}_{n+1}}[F_{n+1}(\widetilde{D}_{n+1}, X^{\pi}_T)] + G_{n+1}(\widetilde{D}_{n+1}, \mathbb{E}_{n+1,\widetilde{D}_{n+1}}[X^{\pi}_T]) \\ &= f^{\pi}_{n+1,n+1}(\widetilde{D}_{n+1}; \widetilde{D}_{n+1}) + G_{n+1}(\widetilde{D}_{n+1}, g^{\pi}_{n+1}(\widetilde{D}_{n+1})). \end{aligned}$$

Taking expectations with respect to $\mathbb{P}_{n,\tilde{d}}$, and applying the control law π at time n yields

$$\mathbb{E}_{n,\tilde{d}}[J_{n+1}(\widetilde{D}^{\pi}_{n+1}; \pi)] = \mathbb{E}_{n,\tilde{d}}[f^{\pi}_{n+1,n+1}(\widetilde{D}^{\pi}_{n+1}; \widetilde{D}^{\pi}_{n+1})] + \mathbb{E}_{n,\tilde{d}}[G_{n+1}(\widetilde{D}^{\pi}_{n+1}, g^{\pi}_{n+1}(\widetilde{D}^{\pi}_{n+1})].$$

Adding and subtracting $J_n(\tilde{d}; \pi)$ gives

$$\begin{aligned} \mathbb{E}_{n,\tilde{d}}[J_{n+1}(\widetilde{D}^{\pi}_{n+1}; \pi)] = J_n(\tilde{d}; \pi) &+ \mathbb{E}_{n,\tilde{d}}[f^{\pi}_{n+1,n+1}(\widetilde{D}^{\pi}_{n+1}; \widetilde{D}^{\pi}_{n+1})] - \mathbb{E}_{n,\tilde{d}}[F_n(\tilde{d}, X^{\pi}_T)] \\ &+ \mathbb{E}_{n,\tilde{d}}[G_{n+1}(\widetilde{D}^{\pi}_{n+1}, g^{\pi}_{n+1}(\widetilde{D}^{\pi}_{n+1}))] - G_n(\tilde{d}, \mathbb{E}_{n,\tilde{d}}[X^{\pi}_T]). \end{aligned}$$

Using (A.4)-(A.5), we then obtain

$$
\begin{aligned}
J_n(\tilde{d};\pi) = \mathbb{E}_{n,\tilde{d}}[J_{n+1}(\widetilde{D}^{\pi}_{n+1};\pi)] - \Big(\mathbb{E}_{n,\tilde{d}}[f^{\pi}_{n+1,n+1}(\widetilde{D}^{\pi}_{n+1};\widetilde{D}^{\pi}_{n+1})] - \mathbb{E}_{n,\tilde{d}}[f^{\pi}_{n+1,n}(\widetilde{D}^{\pi}_{n+1};\tilde{d})]\Big) \\
- \Big(\mathbb{E}_{n,\tilde{d}}[G_{n+1}(\widetilde{D}^{\pi}_{n+1}, g^{\pi}_{n+1}(\widetilde{D}^{\pi}_{n+1})] - G_n(\tilde{d}, \mathbb{E}_{n,\tilde{d}}[g^{\pi}_{n+1}(\widetilde{D}^{\pi}_{n+1})])\Big).
\end{aligned}
\tag{A.6}
$$

Step 2: Recursion for $V_n(\tilde{d})$. Assume that there exists an optimal strategy π^* and consider a strategy π that coincides with π^* after time n, i.e., $\pi_k(\tilde{d}') = \pi^*_k(\tilde{d}')$, for any $k = n+1, \ldots, T-1$ and $\tilde{d}' \in \widetilde{\mathcal{D}}_k$. By definition, we then have

$$
J_n(\tilde{d};\pi) \le V_n(\tilde{d}) = J_n(\tilde{d};\pi^*), \quad \tilde{d} \in \widetilde{\mathcal{D}}_n.
$$

For the optimal strategy π^*, we recall the definition of the function sequences in (A.4), and let

$$
f_{n,k}(\tilde{d};\tilde{d}') := f^{\pi^*}_{n,k}(\tilde{d};\tilde{d}'), \qquad g_n(\tilde{d}) := g^{\pi^*}_n(\tilde{d}). \tag{A.7}
$$

Since π and π^* coincide after time n, for any $\widetilde{D}_{n+1} \in \widetilde{\mathcal{D}}_{n+1}$ we have that

$$
J_{n+1}(\widetilde{D}_{n+1};\pi) = V_{n+1}(\widetilde{D}_{n+1}), \quad f_{n+1,k}(\widetilde{D}_{n+1};\tilde{d}') = f^{\pi}_{n+1,k}(\widetilde{D}_{n+1};\tilde{d}'), \quad g_{n+1}(\widetilde{D}_{n+1}) = g^{\pi}_{n+1}(\widetilde{D}_{n+1}).
$$

In turn, using the recursion (A.6) for $J_n(d;\pi)$, we may write for $\tilde{d} \in \widetilde{\mathcal{D}}_n$,

$$
\begin{aligned}
V_n(\tilde{d}) = \sup_{\pi} \Big\{\mathbb{E}_{n,\tilde{d}}[V_{n+1}(\widetilde{D}^{\pi}_{n+1})] - \Big(\mathbb{E}_{n,\tilde{d}}[f_{n+1,n+1}(\widetilde{D}^{\pi}_{n+1};\widetilde{D}^{\pi}_{n+1})] - \mathbb{E}_{n,\tilde{d}}[f_{n+1,n}(\widetilde{D}^{\pi}_{n+1};\tilde{d})]\Big) \\
- \Big(\mathbb{E}_{n,\tilde{d}}[G_{n+1}(\widetilde{D}^{\pi}_{n+1}, g_{n+1}(\widetilde{D}^{\pi}_{n+1})] - G_n(\tilde{d}, \mathbb{E}_{n,\tilde{d}}[g_{n+1}(\widetilde{D}^{\pi}_{n+1})])\Big)\Big\},
\end{aligned}
$$

with the terminal condition $V_T(\tilde{d}) = 0$, for $\tilde{d} \in \mathcal{D}_T$. The recursions and probabilistic representations of $(f_{n,k}(\cdot;\tilde{d}'))_{0\le n\le T}$ and $(g_n)_{0\le n\le T}$ then follow from (A.3)-(A.5), and (A.7). □

Proof of Proposition 4.1: Assuming the existence of an optimal control law π^*, the value function at time $n+1$ satisfies

$$
V_{n+1}(\widetilde{D}_{n+1}) = f_{n+1,n+1}(\widetilde{D}_{n+1};\widetilde{D}_{n+1}) + \frac{\gamma_{n+1}(D_{n+1})}{2}\left(\frac{g_{n+1}(\widetilde{D}_{n+1})}{X_{n+1}}\right)^2.
$$

Plugging the above expression into the HJB equation (A.2) gives

$$V_n(\tilde{d}) = \sup_{\pi} \left\{ \mathbb{E}_{n,\tilde{d}}[f_{n+1,n}(\widetilde{D}^{\pi}_{n+1}; \tilde{d})] + \frac{\gamma_n(d)}{2} \left(\frac{\mathbb{E}_{n,\tilde{d}}[g_{n+1}(\widetilde{D}^{\pi}_{n+1})]}{x} \right)^2 \right\}. \tag{A.8}$$

Next, we look for a candidate optimal control law that is of the form (A.1), i.e., $\pi_n(\tilde{d}) = \tilde{\pi}_n(d)x$. It then follows from Lemma A.1, along with (A.3), (A.4), and (A.7) that, for this candidate optimal law,

$$\begin{aligned} f_{n+1,n}(\tilde{d}; \tilde{d}') &= a_{n+1}(d)\frac{x}{x'} - 1 - \frac{\gamma_n(\tilde{d}')}{2} b_{n+1}(d)\Big(\frac{x}{x'}\Big)^2, \\ g_{n+1}(\tilde{d}) &= a_{n+1}(d)x. \end{aligned}$$

Plugging the above expression into (A.8), using the wealth dynamics (3.1), and, for brevity, writing R for R_{n+1}, $\tilde{\pi}_n$ for $\tilde{\pi}_n(d)$, a_{n+1} for $a_{n+1}(D_{n+1})$, and b_{n+1} for $b_{n+1}(D_{n+1})$, gives

$$\begin{aligned} V_n(\tilde{d}) &= \sup_{\pi} \Big\{ \mathbb{E}_{n,\tilde{d}}\Big[a_{n+1}\frac{X^{\pi}_{n+1}}{x} - 1 - \frac{\gamma_n(d)}{2} b_{n+1}\Big(\frac{X^{\pi}_{n+1}}{x}\Big)^2 \Big] + \frac{\gamma_n(d)}{2}\Big(\mathbb{E}_{n,\tilde{d}}\Big[a_{n+1}\frac{X^{\pi}_{n+1}}{x}\Big]\Big)^2 \Big\} \\ &= \sup_{\pi} \Big\{ \mathbb{E}_{n,\tilde{d}}\Big[a_{n+1}(R + \widetilde{Z}_{n+1}\tilde{\pi}_n) - 1 - \frac{\gamma_n(d)}{2} b_{n+1}(R + \widetilde{Z}_{n+1}\tilde{\pi}_n)^2 \Big] + \frac{\gamma_n(d)}{2}\Big(\mathbb{E}_{n,\tilde{d}}[a_{n+1}(R + \widetilde{Z}_{n+1}\tilde{\pi}_n)]\Big)^2 \Big\} \\ &= \sup_{\pi} \Big\{ \mathbb{E}_{n,\tilde{d}}\Big[a_{n+1}(R + \widetilde{Z}_{n+1}\tilde{\pi}_n) - 1 - \frac{\gamma_n(d)}{2} b_{n+1}(R^2 + 2R\tilde{\pi}_n\widetilde{Z}_{n+1} + \tilde{\pi}_n^2\widetilde{Z}_{n+1}^2)\Big] \\ &\quad + \frac{\gamma_n(d)}{2}\Big(R^2(\mathbb{E}_{n,\tilde{d}}[a_{n+1}])^2 + 2R\tilde{\pi}_n\mathbb{E}_{n,\tilde{d}}[a_{n+1}]\mathbb{E}_{n,\tilde{d}}[a_{n+1}\widetilde{Z}_{n+1}] + \tilde{\pi}_n^2(\mathbb{E}_{n,\tilde{d}}[a_{n+1}\widetilde{Z}_{n+1}])^2\Big)\Big\}. \end{aligned}$$

Taking the derivative with respect to $\tilde{\pi}_n$ yields the optimal allocation (4.1) and we easily conclude. □

Proof of Lemma 4.2: The formulas for $\mu^{az}_n(d)$ follow directly from (4.3)-(4.4) and the definition of the random variable D_{n+1}. In part (a), the value of τ_{n+1} is determined by $(Y_{(n+1)}, Z_{(n+1)}, \epsilon_{(n+1)})$, because the interaction schedule $(T_k)_{k\geq 0}$ is adapted to the filtration $(\mathcal{F}_n)_{n\geq 0}$ in (3.2). Conditionally on $\tau_{n+1} = n+1$, the value of ξ_{n+1} is determined by $(Y_{(n+1)}, Z_{(n+1)}, \epsilon_{(n+1)})$ for the same reason. In part (b), (3.9) and (3.11) yield that if $\tau_{n+1} = n+1$,

$$\xi_{n+1} = \xi_{\tau_n} e^{(n+1-\tau_n)\alpha} e^{\epsilon'} \frac{\gamma^Z_{n+1}}{\gamma^Z_{\tau_n}} \frac{\bar{\gamma}(Y_{n+1})}{\bar{\gamma}(Y_{\tau_n})} = \xi_{n+1-\phi} e^{\phi\alpha} e^{\sum_{k=\tau_n}^{n} \epsilon_{k+1}} \frac{\gamma^Z_{n+1}}{\gamma^Z_{n+1-\phi}} \frac{\bar{\gamma}(Y_{n+1})}{\bar{\gamma}(Y_{n+1-\phi})},$$

where γ^Z_{n+1} and $\gamma^Z_{n+1-\phi}$ depend on $\{Y_{n+1-2\phi}, \ldots, Y_n\}$ and $\{Z_{n+2-2\phi}, \ldots, Z_{n+1}\}$, and we used that the interaction time prior to $n+1$ was $n+1-\phi$. □

Proof of (4.7): The optimal allocation in Proposition 4.1 can be written as

$$\tilde{\pi}^*_n(d) = \frac{\tilde{\mu}_{n+1}}{\gamma_n(d)\sigma^2_{n+1}} \frac{\tilde{\mu}^{az}_n(d) - R_{n+1}\gamma_n(d)(\tilde{\mu}^{bz}_n(d) - \mu^a_n(d)\tilde{\mu}^{az}_n(d))}{\tilde{\mu}^{bz^2}_n(d) + \left(\frac{\tilde{\mu}_{n+1}}{\sigma_{n+1}}\right)^2 (\tilde{\mu}^{bz^2}_n(d) - (\tilde{\mu}^{az}_n(d))^2)},$$

where

$$\tilde{\mu}^{az}_n(d) := \frac{\mu^{az}_n(d)}{\tilde{\mu}_{n+1}}, \qquad \tilde{\mu}^{bz}_n(d) := \frac{\mu^{bz}_n(d)}{\tilde{\mu}_{n+1}}, \qquad \tilde{\mu}^{bz^2}_n(d) := \frac{\mu^{az}_n(d)}{\tilde{\mu}^2_{n+1} + \sigma^2_{n+1}}.$$

The above formula reduces to (4.7) if $a_{n+1}(D_{n+1})$ is independent of Z_{n+1}, because, in that case,

$$\tilde{\mu}^{az}_n(d) = \mathbb{E}_{n,d}\Big[a_{n+1}(D_{n+1})\frac{\widetilde{Z}_{n+1}}{\tilde{\mu}_{n+1}}\Big] = \mathbb{E}_{n,d}[a_{n+1}(D_{n+1})] = \mu^a_n(d),$$

with analogous results for $\tilde{\mu}^{bz}_n(d)$ and $\tilde{\mu}^{bz^2}_n(d)$. □

B Proofs of Results Related to Section 5

Proof of Proposition 5.1: From the definitions of γ_n and γ^C_n, it follows that

$$\mathcal{R}(\phi,\beta) = \frac{1}{T}\sum_{n=0}^{T-1}\mathbb{E}\Big[\Big|\frac{\gamma^{id}_n}{\gamma^{id}_{\tau_n}\gamma^Z_{\tau_n}} - 1\Big|\Big].$$

Let $Y_0 = y_0$, and extend the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to include an independent sequence of random variables $(Z_n)_{n\le 0}$, such that $Z_n \sim \mathcal{N}(\tilde{\mu}(y_0), \sigma^2(y_0))$. On the event $A := \{Y_{\tau_n-\phi} = \cdots = Y_{\tau_n-1} = y_0\}$, the economic state variable is fixed, and we have

$$\mathbb{E}\Big[\Big|\frac{\gamma^{id}_n}{\gamma^{id}_{\tau_n}\gamma^Z_{\tau_n}} - 1\Big|\Big] = \mathbb{E}\Big[\Big|e^{\beta\left(\frac{1}{\phi}\sum_{k=\tau_n-\phi}^{\tau_n-1} Z_{k+1}-\mu_{k+1}\right)+\sum_{k=\tau_n}^{n-1}\epsilon_{k+1}} - 1\Big|A\Big] + O\big(1 - P^{\tau_n-1}_{y_0,y_0}\big),$$

where $P^{\tau_n-1}_{y_0,y_0} := (P^{\tau_n-1})_{y_0,y_0}$ is the probability of staying in state y_0 from time zero to time $\tau_n - 1$. For the first term appearing in the exponent, we have

$$\beta\Big(\frac{1}{\phi}\sum_{k=\tau_n-\phi}^{\tau_n-1} Z_{k+1} - \mu_{k+1}\Big)\Big|A \sim \mathcal{N}\Big(0, \frac{\beta^2\sigma^2_0}{\phi}\Big).$$

For the second term, we define the random variable $J := |\{\tau_n \leq k \leq n-1 : \epsilon_{k+1} \neq 0\}| \sim \text{Bin}(n-\tau_n, p_\epsilon)$, which is equal to the number of idiosyncratic risk aversion jumps between times τ_n and n, so

$$\sum_{k=\tau_n}^{n-1} \epsilon_{k+1} \Big| \{J = j\} \sim \mathcal{N}\Big(-j\frac{\sigma_\epsilon^2}{2}, j\sigma_\epsilon^2\Big).$$

For $Z \sim \mathcal{N}(0,1)$, $a \in \mathbb{R}$, and $b \geq 0$, we have $\mathbb{E}[|e^{a+bZ}-1|] = b\mathbb{E}[|Z|] + O(|a|+b^2)$. Hence, by conditioning on the value of J and using the absolute Gaussian moment formula, we have

$$\begin{aligned}
&\mathbb{E}\Big[\Big|e^{\beta\left(\frac{1}{\phi}\sum_{k=\tau_n-\phi}^{\tau_n-1} Z_{k+1}-\mu_{k+1}\right)+\sum_{k=\tau_n}^{n-1}\epsilon_{k+1}} - 1\Big| A\Big] \\
&= (1-p_\epsilon)^{n-\tau_n}\sqrt{\frac{2}{\pi}}\frac{\beta\sigma_0}{\sqrt{\phi}} + np_\epsilon(1-p_\epsilon)^{n-\tau_n-1}\sqrt{\frac{2}{\pi}}\sqrt{\frac{\beta^2\sigma_0^2}{\phi}+\sigma_\epsilon^2} + O(p^2\sigma_\epsilon^2+\beta^2).
\end{aligned}$$

Then, assuming T is a multiple of ϕ,

$$\begin{aligned}
\mathcal{R}(\phi,\beta) &= \sqrt{\frac{2}{\pi}}\frac{1}{T}\Big(\frac{\beta\sigma_0}{\sqrt{\phi}}\sum_{n=0}^{T-1}(1-p_\epsilon)^{n-\tau_n} + \sqrt{\frac{\beta^2\sigma_0^2}{\phi}+\sigma_\epsilon^2}\sum_{n=0}^{T-1} np_\epsilon(1-p_\epsilon)^{n-\tau_n-1}\Big) + O(p_\epsilon^2\sigma_\epsilon^2+\beta^2) + O(1-P_{y_0,y_0}^{T-\phi-1}) \\
&= \sqrt{\frac{2}{\pi}}\frac{1}{\phi}\Big(\frac{\beta\sigma_0}{\sqrt{\phi}}\sum_{n=0}^{\phi-1}(1-p_\epsilon)^{n} + \sqrt{\frac{\beta^2\sigma_0^2}{\phi}+\sigma_\epsilon^2}\sum_{n=0}^{\phi-1} np_\epsilon(1-p_\epsilon)^{n-1}\Big) + O(p_\epsilon^2\sigma_\epsilon^2+\beta^2) + O(1-P_{y_0,y_0}^{T-\phi-1}) \\
&= \sqrt{\frac{2}{\pi}}\frac{1}{\phi}\Big(\frac{\beta\sigma_0}{\sqrt{\phi}}\frac{1-(1-p_\epsilon)^\phi}{p_\epsilon} + \sqrt{\frac{\beta^2\sigma_0^2}{\phi}+\sigma_\epsilon^2}\frac{1-(1-p_\epsilon)^\phi-\phi(1-p_\epsilon)^{\phi-1}p_\epsilon}{p_\epsilon}\Big) + O(p_\epsilon^2\sigma_\epsilon^2+\beta^2) + O(1-P_{y_0,y_0}^{T-\phi-1}) \\
&= \sqrt{\frac{2}{\pi}}\Big(\frac{\beta\sigma_0}{\sqrt{\phi}}\Big(1-\frac{\phi-1}{2}p_\epsilon\Big) + \sqrt{\frac{\beta^2\sigma_0^2}{\phi}+\sigma_\epsilon^2}\frac{\phi-1}{2}p_\epsilon\Big) + O(p_\epsilon^2\sigma_\epsilon^2+\beta^2) + O(1-P_{y_0,y_0}^{T-1}),
\end{aligned}$$

where the last equality uses a third-order Taylor approximation for $(1-p_\epsilon)^\phi$ and that $P_{y_0,y_0}^{T-\phi-1}$ is increasing in $\phi \geq 1$. Next, define the function

$$f(\phi) := \frac{\beta\sigma_0}{\sqrt{\phi}}\Big(1-\frac{\phi-1}{2}p_\epsilon\Big) + \sqrt{\frac{\beta^2\sigma_0^2}{\phi}+\sigma_\epsilon^2}\frac{\phi-1}{2}p_\epsilon.$$

Observe that $f \geq 0$ and that if $p_\epsilon = 0$ and $\beta = 0$, then $f \equiv 0$. If $p_\epsilon = 0$ and $\beta > 0$, then f is strictly decreasing for $\phi \geq 1$, and converges to zero as $\phi \to \infty$. If $\beta = 0$ and $p_\epsilon > 0$, then f is strictly increasing for $\phi \geq 1$, and equal to zero for $\phi = 1$. The derivative of f is

$$\begin{aligned}
f'(\phi) &= -\frac{1}{2}\frac{\beta\sigma_0}{\phi^{3/2}}\Big(1-\frac{\phi-1}{2}p_\epsilon\Big) - \frac{\beta\sigma_0}{\sqrt{\phi}}\frac{p_\epsilon}{2} - \frac{1}{2\sqrt{\frac{\beta^2\sigma_0^2}{\phi}+\sigma_\epsilon^2}}\frac{\beta^2\sigma_0^2}{\phi^2}\frac{\phi-1}{2}p_\epsilon + \sqrt{\frac{\beta^2\sigma_0^2}{\phi}+\sigma_\epsilon^2}\frac{p_\epsilon}{2} \\
&= \frac{1}{2}\frac{\beta\sigma_0 p_\epsilon}{\phi^2}\Big(-\frac{\sqrt{\phi}}{p_\epsilon} + \frac{1}{2}\Big(1-\frac{1}{\sqrt{1+\frac{\sigma_\epsilon^2\phi}{\beta^2\sigma_0^2}}}\Big)\sqrt{\phi}(\phi-1) + \Big(\sqrt{1+\frac{\sigma_\epsilon^2\phi}{\beta^2\sigma_0^2}}-1\Big)\phi^{3/2}\Big) =: \frac{1}{2}\frac{\beta\sigma_0 p}{\phi^2}g(\phi),
\end{aligned}$$

and it follows that the derivative at $\phi = 1$ satisfies

$$f'(1) = \frac{1}{2}\beta\sigma_0 p_\epsilon\Big(\sqrt{1+\frac{\sigma_\epsilon^2}{\beta^2\sigma_0^2}} - 1 - \frac{1}{p_\epsilon}\Big) \geq 0 \quad \iff \quad \frac{p_\epsilon\sigma_\epsilon^2}{\beta^2\sigma_0^2} \geq 2 + \frac{1}{p_\epsilon}.$$

Observe that, since the value p_ϵ is in general small, we have

$$\frac{p_\epsilon\sigma_\epsilon^2}{\beta^2\sigma_0^2} \geq 2 + \frac{1}{p_\epsilon} \quad \iff \quad \frac{p_\epsilon\sigma_\epsilon}{\beta\sigma_0} \geq \sqrt{1+2p_\epsilon} \approx 1, \tag{B.1}$$

which shows that $p_\epsilon\sigma_\epsilon/(\beta\sigma_0) > 1$ is approximately a sufficient and necessary condition for $f'(1) < 0$. Furthermore,

$$\begin{aligned} g'(\phi) &= -\frac{1}{2}\frac{\phi^{-1/2}}{p_\epsilon} + \Big[\frac{1}{2}\Big(1 - \frac{1}{\sqrt{1+\frac{\sigma_\epsilon^2\phi}{\beta^2\sigma_0^2}}}\Big)\sqrt{\phi}(\phi-1)\Big]' + \frac{3}{2}\Big(\sqrt{1+\frac{\sigma_\epsilon^2\phi}{\beta^2\sigma_0^2}} - 1\Big)\phi^{1/2} + \frac{1}{2\sqrt{1+\frac{\sigma_\epsilon^2\phi}{\beta^2\sigma_0^2}}}\frac{\sigma_\epsilon^2}{\beta^2\sigma_0^2}\phi^{3/2} \\ &\geq \frac{1}{2}\phi^{-1/2}\Big(-\frac{1}{p_\epsilon} + \Big(\sqrt{1+\frac{\sigma_\epsilon^2}{\beta^2\sigma_0^2}} - 1\Big)\phi\Big) \geq \frac{1}{2}\phi^{-1/2}g(1). \end{aligned}$$

It follows that if $g(1) \geq 0$, then $g'(\phi) \geq 0$ for all $\phi \geq 1$. Hence, $f'(\phi) \geq 0$ for all $\phi \geq 1$, and f is minimized for $\phi = 1$. If $g(1) < 0$, and thus $f'(1) < 0$, then it can be shown that g is strictly convex and $g(\phi) \to \infty$, as $\phi \to \infty$. Hence, there exists a unique point where $g(\phi) = 0$, and where f is minimized. Finally, from the form of the function f, it is easy to see the sign of the derivatives of the optimal value of ϕ with respect to $\beta, \sigma_0, p_\epsilon$, and σ_ϵ. □

Proof of (5.4): We denote the coefficient $\mu_n^a(D_n)$ in (4.3) by $\mu_n^a(\gamma_n)$ and $\mu_n^a(\gamma_n^C)$, for the two investment strategies, and show inductively that

$$\mathbb{E}[|\mu_n^a(\gamma_n) - \mu_n^a(\gamma_n^C)|] = O((\phi-1)p_\epsilon\sigma_\epsilon^2) + O\Big(\frac{\beta^2\sigma^2}{\phi^2}\Big).$$

Analogous identities can be obtained for μ_n^b, μ_n^{az}, $\mu_n^{az^2}$, and μ_n^{bz}. One may then inductively show that

$$\tilde{\pi}_n^*(\gamma_n) = \frac{\gamma_n^C}{\gamma_n}\tilde{\pi}_n^*(\gamma_n^C) + O((\phi-1)p_\epsilon\sigma_\epsilon^2) + O\Big(\frac{\beta^2\sigma^2}{\phi^2}\Big),$$

and the result follows. □

The following proposition shows that there exists a one-to-one relation between risk aversion processes and investment strategies that depend only on time and the economic state. That is, for a given $(\tilde{\pi}_n)_{0\leq n<T}$, where $\tilde{\pi}_n(y)$ is the risky asset allocation at time n in economic state $y \in \mathcal{Y}$, there exists a unique state- and

time-dependent risk aversion process such that the corresponding optimal strategy is given by $(\tilde{\pi}_n)_{0\le n<T}$. The proof shows how this risk aversion process can be constructed using backward induction.

Proposition B.1. *Fix a sequence $(\tilde{\pi}_n)_{0\le n<T}$, with $\tilde{\pi}_n : \mathcal{Y} \mapsto \mathbb{R}_0^+$. Then, there exists a unique risk aversion process $(\gamma_n)_{n\ge 0}$, with $\gamma_n : \mathcal{Y} \mapsto \mathbb{R}_0^+$, and such that $(\tilde{\pi}_n)_{0\le n<T}$ is the optimal allocation corresponding to $(\gamma_n)_{0\le n<T}$.*

Proof of Proposition B.1: We construct the process $(\gamma_n)_{0\le n<T}$ using backward induction. First, for $n = T-1$, we have $\gamma_n(y) = \mu(y)/(\tilde{\pi}_n(y)\sigma^2(y))$. For $n < T-1$, consider the function

$$h : \Big(0, \frac{1}{R(y)} \frac{\mu_n^a}{\mu_n^b - (\mu_n^a)^2}\Big) \mapsto (-\tilde{\pi}_n(y), \infty), \quad h(x) = \frac{\tilde{\mu}(y)}{x\sigma^2(y)} \frac{\mu_n^a - (1+r(y))x(\mu_n^b - (\mu_n^a)^2)}{\mu_n^b + \Big(\frac{\tilde{\mu}(y)}{\sigma(y)}\Big)^2 (\mu_n^b - (\mu_n^a)^2)} - \tilde{\pi}_n(y),$$

where μ_n^a and μ_n^b are determined by $\gamma_{n+1}, \dots, \gamma_{T-1}$. The value of $\gamma_n(y)$ is the unique root of h, which is strictly decreasing on its domain. Using Proposition 4.3 and (4.7), it is simple to check that $(\tilde{\pi}_n)_{0\le n<T}$ is the optimal allocation corresponding to the constructed process $(\gamma_n)_{0\le n<T}$. □

Proof of (5.8): The Sharpe ratio is defined as

$$s^{\pi^*} := \frac{\tilde{\mu}^{\pi^*}}{\sigma^{\pi^*}}, \tag{B.2}$$

where $\tilde{\mu}^{\pi^*}$ and σ^{π^*} are the mean and volatility of the excess return achieved by the strategy π^*. Recall that $r_{n,n+1}^{\pi^*}$ is the return of π^* and that r_{n+1} is the risk-free rate. Using that $\tilde{\pi}_n^*$ is independent of the investment horizon T and that Y_n converges in distribution to a random variable with distribution λ, we have

$$\begin{aligned}
\tilde{\mu}^{\pi^*} &= \lim_{T\to\infty} \mathbb{E}\Big[\frac{1}{T}\sum_{n=0}^{T-1}(r_{n,n+1}^{\pi^*} - r_{n+1})\Big] = \lim_{n\to\infty} \mathbb{E}[r_{n,n+1}^{\pi^*} - r_{n+1}] = \mathbb{E}[r_{0,1}^{\pi^*} - r_1 | Y_0 \sim \lambda],\\
(\sigma^{\pi^*})^2 &= \lim_{T\to\infty} \mathbb{E}\Big[\frac{1}{T}\sum_{n=0}^{T-1}(r_{n,n+1}^{\pi^*} - r_{n+1} - \tilde{\mu}^{\pi^*})^2\Big] = \lim_{n\to\infty} \mathbb{E}[(r_{n,n+1}^{\pi^*} - r_{n+1} - \tilde{\mu}^{\pi^*})^2] = \mathbb{E}[(r_{0,1}^{\pi^*})^2 | Y_0 \sim \lambda] - (\tilde{\mu}^{\pi^*})^2.
\end{aligned}$$

Let λ be the stationary probability of the state $y = 2$. We then have

$$\begin{aligned}
\tilde{\mu}^{\pi^*} &= (1-\lambda)\tilde{\mu}(1)\bar{\pi} + \lambda a \tilde{\mu}(1)\bar{\pi}(1+\delta) = \tilde{\mu}(1)\bar{\pi}\big(1 + \lambda(a(1+\delta)-1)\big),\\
(\sigma^{\pi^*})^2 &= (1-\lambda)(\tilde{\mu}^2(1) + \sigma^2(1))\bar{\pi}^2 + \lambda(a^2\tilde{\mu}^2(1) + b^2\sigma^2(1))\bar{\pi}^2(1+\delta)^2 - \tilde{\mu}^2(1)\bar{\pi}^2\big(1+\lambda(a(1+\delta)-1)\big)^2\\
&= \bar{\pi}^2\Big(\tilde{\mu}^2(1)\big(1-\lambda+\lambda a^2(1+\delta)^2 - \big(1+\lambda(a(1+\delta)-1)\big)^2\big) + \sigma^2(1)\big(1-\lambda+\lambda b^2(1+\delta)^2\big)\Big),
\end{aligned}$$

and (5.8) follows by plugging these values into (B.2) and viewing the Sharpe ratio as a function of δ. □

Lemma B.2. *Let $a, b > 1$, $0 < \lambda < 1$, $\delta > -1$, and $u := (\sigma(1)/\tilde{\mu}(1))^2$. We have the following characterizations for the sensitivities of the Sharpe ratio $s^{\pi^*}(\delta)$ defined in (5.8):*

$$\begin{aligned}
\frac{\partial s^{\pi^*}(\delta)}{\partial a} > 0 \quad &\iff \quad a < \frac{1+u}{1+\delta} + \frac{\lambda}{1-\lambda} u b^2 (1+\delta), \\
\frac{\partial s^{\pi^*}(\delta)}{\partial b} < 0 \quad &\iff \quad b > 1, \\
\frac{\partial s^{\pi^*}(\delta)}{\partial \lambda} > 0 \quad &\iff \quad \lambda > \frac{(a^2+ub^2)(1+\delta)^2 - (1+u)(2a(1+\delta)-1)}{(a(1+\delta)-1)((a^2+ub^2)(1+\delta)^2-(1+u))}.
\end{aligned}$$

Furthermore,

$$\lim_{a\to\infty} s^{\pi^*}(\delta) = \sqrt{\frac{\lambda}{1-\lambda}}, \qquad \lim_{b\to\infty} s^{\pi^*}(\delta) = 0, \qquad s^{\pi^*}(\delta)\big|_{\lambda=0} = \frac{\tilde{\mu}(1)}{\sigma(1)}, \qquad s^{\pi^*}(\delta)\big|_{\lambda=1} = \frac{a\tilde{\mu}(1)}{b\sigma(1)}.$$

Discussion of parameter conditions: The conditions for the parameters a and λ in the previous lemma are not restrictive. For instance, let us consider the values of the derivatives in the case $\delta = 0$, which corresponds to equal market allocation in both states. We then have

$$\frac{\partial s^{\pi^*}(\delta)}{\partial a}\bigg|_{\delta=0} > 0 \quad \iff \quad a < 1 + u + \frac{\lambda}{1-\lambda} u b^2 < 1 + u,$$

which is trivially satisfied for realistic parameter values. For example, for the parameter values in (5.6), $a \approx 2.1$ and $u \approx 5.5$.[24] For the parameter λ, we consider two cases,

$$\frac{\partial s^{\pi^*}(\delta)}{\partial \lambda}\bigg|_{p=0} > 0 \quad \iff \quad \lambda > \frac{a^2+ub^2-(1+u)(2a-1)}{(a-1)(a^2+ub^2-(1+u))} = \begin{cases} \frac{1}{a+1}, & b = a, \\ \frac{a-(1+u)}{a(a+u)-(1+u)}, & b = \sqrt{a}. \end{cases}$$

First, if $b = a > 1$, the market Sharpe ratio is the same in both states, and for the portfolio's Sharpe ratio to increase, the stationary probability λ needs to be large enough. Intuitively, if little time is spent in the state, the higher volatility reduces the portfolio's Sharpe ratio. On the other hand, if sufficient time is spent in the state, then the effect of the higher returns dominates. The case $b = \sqrt{a}$ corresponds to a higher market Sharpe ratio in the recessionary state, which is consistent with empirical evidence. In this case, the portfolio's Sharpe ratio is increasing in λ for $0 < \lambda < 1$, because, as mentioned above, the condition $a < 1 + u$ is in general satisfied. Note that for the parameter values in (5.6), $b \approx a^{0.15} < a^{1/2}$, and the portfolio's Sharpe ratio is increasing for $0 < \lambda < 1$.

[24] These values are based on annual returns. The value of u corresponding to monthly returns is approximately obtained by multiplying by 12, while the value of a corresponding to monthly returns is approximately the same.

Proof of Lemma B.2: The limit as $a \to \infty$ follows directly from the formula for $s^{\pi^*}(\delta)$ in (5.8). Let

$$h(a) := \frac{(1-\lambda)(1+u) + \lambda(a^2+ub^2)(1+\delta)^2}{(1+\lambda(a(1+\delta)-1))^2} = \frac{(1-\lambda)(1+u) + \lambda ub^2(1+\delta)^2 + \lambda a^2(1+\delta)^2}{(1-\lambda+\lambda a(1+\delta))^2} = \frac{c_0 + c_1 a^2}{(d_0+d_1 a)^2},$$

and note that $s^{\pi^*}(\delta)$ is increasing in a if and only if $h(a)$ is decreasing. We have

$$h'(a) = \frac{2c_1 a(d_0+d_1a)^2 - 2(d_0+d_1a)d_1(c_0+c_1a^2)}{(d_0+d_1a)^4} = \frac{2c_1a(d_0+d_1a) - 2d_1(c_0+c_1a^2)}{(d_0+d_1a)^3},$$

and the a-inequality is obtained by solving $h'(a) < 0$ for a. The result for λ is obtained in the same way by noting that $s^{\pi^*}(\delta)$ is increasing in λ if and only if

$$h(\lambda) := \frac{(1-\lambda)(1+u) + \lambda(a^2+ub^2)(1+\delta)^2}{(1+\lambda(a(1+\delta)-1))^2} = \frac{(1-\lambda)(1+u) + \lambda(a^2+ub^2)(1+\delta)^2}{(1+\lambda(a(1+\delta)-1))^2}$$

is decreasing. The results for b follow directly from the formula for $s^{\pi^*}(\delta)$ in (5.8). □

Proof of (5.9): Let $u := (\sigma(1)/\tilde{\mu}(1))^2$ and

$$h(\delta) := \frac{(1-\lambda)(1+u) + \lambda(a^2+b^2u)(1+\delta)^2}{(1+\lambda(a(1+\delta)-1))^2} - 1.$$

Then,

$$\frac{\partial h(\delta)}{\partial \delta} = 2\lambda(1-\lambda)\frac{(a^2+b^2u)(1+\delta) - a(1+u)}{(1+\lambda(a(1+\delta)-1))^3}.$$

The left-hand side of (5.9) is equivalent to the derivative of h being negative, and since $1+\lambda(a(1+\delta)-1) > 0$ for $\delta > -1$, the right-hand side of (5.9) follows directly. □

Proof of (5.10): Let $u := (\sigma(1)/\tilde{\mu}^2(1))^2$. The Sharpe ratio formula (5.8) can be rewritten as

$$s^{\pi^*}(\delta) = \frac{1+\lambda(a-1)+\lambda a\delta}{\sqrt{(\lambda+u)(1-\lambda) + \lambda\big(a^2(1-\lambda)+ub^2\big)(1+\delta)^2 + 2a\lambda(\lambda-1)(1+\delta)}} =: \frac{d_0+d_1\delta}{\sqrt{c_0+c_1(1+\delta)+c_2(1+\delta)^2}}.$$

Let $\kappa(\delta) := \sqrt{c_0+c_1(1+\delta)+c_2(1+\delta)^2}$. Then,

$$\frac{\partial s^{\pi^*}(\delta)}{\partial \delta} = \frac{d_1\kappa(\delta) - (d_0+d_1\delta)\frac{c_1+2c_2(1+\delta)}{2\kappa(\delta)}}{\kappa^2(\delta)} = \frac{d_1\kappa^2(\delta) - \frac{1}{2}(d_0+d_1\delta)(c_1+2c_2(1+\delta))}{\kappa^3(\delta)},$$

and

$$\begin{aligned}\frac{\partial^2 s^{\pi^*}(\delta)}{\partial \delta^2} &= \frac{\big(d_1(c_1+2c_2(1+\delta)) - \frac{d_1}{2}(c_1+2c_2(1+\delta)) - c_2(d_0+d_1\delta)\big)\kappa^3(\delta)}{\kappa^6(\delta)} \\ &\quad - \frac{\frac{3}{2}\big(d_1\kappa^2(\delta) - \frac{1}{2}(d_0+d_1\delta)(c_1+2c_2(1+\delta))\big)(c_1+2c_2(1+\delta))\kappa(\delta)}{\kappa^6(\delta)} \\ &= -\frac{(d_0+d_1\delta)\big(c_2\kappa^2(\delta) - \frac{3}{4}(c_1+2c_2(1+\delta))^2\big) + d_1(c_1+2c_2(1+\delta))\kappa^2(\delta)}{\kappa^5(\delta)}.\end{aligned}$$

For $\delta = 0$, we have

$$\begin{aligned}\frac{\partial^2 s^{\pi^*}(\delta)}{\partial \delta^2}\bigg|_{\delta=0} &= -d_0\big(c_0c_2 - \frac{3}{4}c_1^2 - 2c_1c_2 - 2c_2^2\big) - d_1(c_1+2c_2)(c_0+c_1+c_2) \\ &= \Big(\frac{s^\pi(0)}{d_0}\Big)^5\Big(-(a^2u+b^2u^2)\lambda + O(\lambda^2)\Big),\end{aligned}$$

and the result follows. □

C Computational Complexity of the Optimal Investment Strategy

The optimal investment strategy in Proposition 4.1 is computed using backward induction, which requires discretization of the state variable D_n. This variable is high-dimensional, containing the history of economic states and market returns, the history of interaction times, and the communicated risk aversion parameters. However, the dimensionality of the state variable needed to compute the optimal investment strategy depends on the complexity of the risk aversion process, $(\gamma_n)_{n\geq 0}$, and the interaction schedule, $(T_k)_{k\geq 0}$.

For instance, assuming a single economic state, the state variable needed for the risk aversion model in Section 3.5 has three components:[25]

$$\widetilde{D}_n := \Big(\xi_n, \sum_{k=\tau_n-\phi}^{\tau_n-1} Z_{k+1}, \sum_{k=\tau_n}^{n-1} Z_{k+1}\Big) \in \mathbb{R}_0^+ \times \mathbb{R}^2.$$

That is, $\widetilde{D}_n$ consists of the most recently communicated risk aversion parameter, the cumulative market return in the previous updating interval, and the running cumulative market return in the current updating interval. With step size $\Delta\gamma$ used in the discretization of the risk aversion parameter, step size Δz used for each of the return variables, and allowing for $|\mathcal{Y}| \geq 1$ economic states, the number of points in the

[25]The functional J_n in (3.5) depends on $d \in \mathcal{D}_n$, with the risk-return tradeoff given by $\gamma_n(d)$. In this example, the risk aversion γ_n is measurable w.r.t. $\sigma(\widetilde{D}_n)$, which is a sub-σ-algebra of $\sigma(D_n)$. That is, for a fixed $\tilde{d} \in \mathbb{R}_0^+ \times \mathbb{R}^2$, the value of J_n is the same for any $d \in \mathcal{D}_n$ such that $\widetilde{D}_n = d$.

discretization grid is of order

$$\frac{T|\mathcal{Y}|}{\Delta\gamma\Delta z}\Big(\frac{1}{\Delta z}\mathbf{1}_{\{\phi>1\}}+\mathbf{1}_{\{\phi=1\}}\Big).$$

Each point requires a numerical integration to compute each of the expected values in (4.3) and a double integration at interaction times.

The model in Section 3.5 has a deterministic interaction schedule, but a stochastic interaction schedule can add to the dimensionality of the discretization grid. For instance, if interaction is triggered by a market drawdown exceeding a certain threshold, then the state variable needs to include information about the cumulative return since the previous time of interaction.

One may use properties of the risk aversion process to reduce the computational complexity. For example, in the computation of μ_n^{az} in Lemma 4.2, the numerical integration can be avoided by evaluating Z_{n+1} at its expected value, μ_{n+1}. This leads to the approximation

$$\begin{aligned}\mu_n^{az}(d) &= \mathbb{E}_{n,d}[a_{n+1}(D_{n+1})\tilde{Z}_{n+1}]\\ &= \mathbb{E}_{n,d}[a_{n+1}(D_{n+1})\tilde{Z}_{n+1}|Z_{n+1}=\mu_{n+1}] + O\Big(\Big(\sum_{k=n+1}^{T-1}\mathbb{E}\big[(\tilde{\pi}_k^*(\gamma_k)-\tilde{\pi}_k^*(\gamma_k'))^2\big]\Big)^{1/2}\Big).\end{aligned}$$

In the above error term, $(\tilde{\pi}^*(\gamma_k'))_{k=n+1}^{T-1}$ are the future optimal allocations corresponding to $(\gamma_k')_{k=n+1}^{T-1}$, which is the risk aversion process resulting from setting $Z_{n+1}=\mu_{n+1}$. In particular, if future risk aversion values are independent of Z_{n+1}, then $(\gamma_k')_{k=n+1}^{T-1}=(\gamma_k)_{k=n+1}^{T-1}$, and the error term vanishes. In general, the effect of a single market return on the future risk aversion path is limited - either temporary, or quickly diluted by the effect of future returns. For instance, in the risk aversion model of Section 3.5, the return Z_{n+1} impacts the risk aversion process only in a single interaction interval.

The computational complexity can also be reduced by using properties of the economic state transition matrix. In particular, for a single time step, the transition probability out of the current state is generally small (see (5.5)). In the computation of $\mu_n^{az}(d)$, this leads to the approximation

$$\mu_n^{az}(d)=\mathbb{E}_{n,d}[a_{n+1}(D_{n+1})\tilde{Z}_{n+1}]=\sum_{y'\in\widetilde{\mathcal{Y}}}P_{y_n,y'}\mathbb{E}_{n,d}\big[a_{n+1}(D_{n+1})\tilde{Z}_{n+1}\big|Y_{n+1}=y'\big]+O\Big(1-\sum_{y'\in\widetilde{\mathcal{Y}}}P_{y_n,y'}\Big),$$

where $\widetilde{\mathcal{Y}}\subseteq\mathcal{Y}$. For instance, $\widetilde{\mathcal{Y}}=\{y_n\}$ conditions on staying in the current economic state, while $\widetilde{\mathcal{Y}}=\{y_n-1,y_n,y_n+1\}$ conditions on staying in the current state or transitioning to neighboring states.

References

S. Allgood and W. Walstad. The Effects of Perceived and Actual Financial Literacy on Financial Behaviors. *Economic Inquiry*, 54(1), 675–697, 2016.

K. Astrom and B. Wittenmark. *Adaptive Control.* MA, Reading:Addison Wesley, 1989.

S. Basak and G. Chabakauri. Dynamic Mean Variance Asset Allocation. *Review of Financial Studies*, 23(8), 2970–3016, 2010.

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

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

P. Billingsley. *Probability and Measure*, 3rd edition. John Wiley & Sons, 1995.

T. Björk and A. Murgoci. A Theory of Markovian Time Inconsistent Stochastic Control in Discrete Time. *Finance & Stochastics*, 18(3), 545–592, 2013.

T. Björk, A. Murgoci, and X. Zhou. Mean-Variance Portfolio Optimization with State-Dependent Risk Aversion. *Mathematical Finance*, 24(1), 1–24, 2014.

C. Brooks, I. Sangiorgi, C. Hillenbrand, and K. Money. Why are older investors less willing to take financial risks? *International Review of Financial Analysis*, 56, 52—72, 2018.

M. Brunnermeier and S. Nagel. Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation. *American Economic Review*, 98(3), 713–736, 2008.

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

L. Calvet and P. Sodini. Twin Picks: Disentangling the Determinants of Risk-Taking in Household Portfolios. *Journal of Finance*, 69(2), 867–906, 2014.

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

G. Charness, U. Gneezy, and A. Imas. Experimental Methods: Eliciting Risk Preferences. *Journal of Economic Behavior & Organization*, 87, 43–51, 2013.

M. Chauvet and J. Hamilton. Dating Business Cycle Turning Points. In Nonlinear Time Series Analysis of Business Cycles, ed. Costas Milas, Philip Rothman, and Dick van Dijk, 1–54. Amsterdam and San Diego: Elsevier. 2006

P-A. Chiappori and M. Paiella. Relative Risk Aversion Is Constant: Evidence From Panel Data. *Journal of the European Economic Association*, 9(6), 1021–1052, 2011.

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

J. Cox and G. Harrison. *Risk Aversion in Experiments*, Vol. 12 of Research in Experimental Economics. Emerald Group Publishing Limited, United Kingdom, 2008.

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

M. Dai, H. Jin, S. Kou, and Y. Xu. A Dynamic Mean-Variance Analysis for Log Returns. Forthcoming in *Management Science*, 2020.

B. Dietvorst, J. Simmons, and C. Massey. Algorithm Aversion: People Erroneously Avoid Algorithms after Seeing Them Err. *Journal of Experimental Psychology: General*, 144(1), 114–126, 2015.

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

N. Gennaioli, A. Shleifer, and R. Vishny. Money Doctors. *Journal of Finance*, 70(1), 91–114, 2015.

M. Guidolin and A. Timmermann. An Econometric Model of Nonlinear Dynamics in the Joint Distribution of Stock and Bond Returns. *Journal of Applied Econometrics*, 21(1), 1–22, 2006.

V. Coudert and M. Gex. Does Risk Aversion Drive Financial Crises? Testing the Predictive Power of Empirical Indicators. *Journal of Empirical Finance*, 15(2), 167–184, 2008.

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

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

H. Guo and R. Whitelaw. Uncovering The Risk-Return Relation In The Stock Market. *Journal of Finance*, 61(3), 1433–1463, 2006.

T. Hallahan, R. Faff, and M. McKenzie. An Empirical Investigation of Personal Financial Risk Tolerance. *Financial Services Review*, 13, 57—78, 2004.

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

V. John, I. Angelov, A. Öncül, and D. Thévenin. Techniques for the Reconstruction of a Distribution From a Finite Number of its Moments. *Chemical Engineering Science*, 62, 2890–2904, 2007.

D. Kahneman and M. Riepe. Aspects of Investor Psychology: Beliefs, Preferences, and Biases Investment Advisors Should Know about. *The Journal of Portfolio Management*, 24(4) 52–65, 1998.

M. Kritzman, S. Page, and D. Turkington. Regime Shifts: Implications for Dynamic Strategies. *Financial Analysts Journal*, 68(2), 2012.

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

D. Li and W-L. Ng. Optimal Dynamic Portfolio Selection: Multiperiod Mean-Variance Formulation. *Mathematical Finance*, 10(3), 387–406, 2000.

J. Liu. Portfolio Selection in Stochastic Environments. *Review of Financial Studies*, 20, 1–39, 2007.

H. Markowitz. Portfolio Selection. *Journal of Finance*, 7(1), 77–91, 1952.

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

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

Statista. FinTech Report 2020 - Personal Finance. Available at `https://www.statista.com/study/41710/fintech-report-personal-finance/`.

Y. Tang and R. Whitelaw. Time-Varying Sharpe Ratios and Market Timing. *Quarterly Journal of Finance*, 1(3), 465–493, 2011.

A. Tversky and D. Kahneman. Prospect Theory: An Analysis of Decision under Risk, *Econometrica*, 47(4), 263–291, 1979.

G. Van de Venter, D. Michayluk, and G. Davey. A Longitudinal Study of Financial Risk Tolerance. *Journal of Economic Psychology*, 33, pp. 794–800, 2012.