European Journal of Operational Research xxx (xxxx) xxx

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Interfaces with Other Disciplines

Long-term dynamic asset allocation under asymmetric risk preferences

Vasileios E. Kontosakos[a,b,*], Soosung Hwang[c], Vasileios Kallinterakis[d], Athanasios A. Pantelous[a]

[a] Department of Econometrics and Business Statistics, Monash Business School, Monash University, 20 Chancellors walk, Wellington Road, Clayton Campus, Victoria, Australia
[b] Group Risk, Allianz SE, Koeniginstrasse 28, Munich, 80802, Germany
[c] School of Economics, Sungkyunkwan University, 53 Myeongnyun-Dong 3-Ga Jongno-Gu, Seoul 110-745, Republic of Korea
[d] Durham University Business School, Durham University, UK

ARTICLE INFO

Article history:
Received 25 November 2021
Accepted 28 July 2023
Available online xxx

JEL classification:
G40
C61
G11

Keywords:
Decision analysis
Asset allocation
Asymmetric risk preferences
Parameter uncertainty
Simulation study

ABSTRACT

We examine the impact of return predictability and parameter uncertainty on long-term portfolio allocations when investors' utility function quantifies their asymmetric behaviour against expected gains and losses on risky assets. Allowing for different return generating systems and two investable assets, we examine the way portfolio allocation to the risky asset evolves over the course of the investment horizon in the presence of risk asymmetries. We find persisting horizon effects, with stocks appearing progressively more attractive at longer horizons as opposed to shorter ones. The role of parameter uncertainty also appears to be prominent in the portfolio choice problem. Accounting for this results in both significantly lowering the exposure to the risky asset and lessening the horizon effects driven by return predictability. An equally important aspect of this study relates to detecting a level of disappointment aversion below which it is optimal for investors to hold zero units of a risky asset. In this regard, our analysis has implications for the nonparticipation puzzle in stock markets.

1. Introduction

Asymmetric risk preferences in investors' decision making have been an integral part of the portfolio choice literature at least over the course of the past two decades (Bellemare, Kröger, & Sossou, 2020; Berkelaar, Kouwenberg, & Post, 2004; Dahlquist, Farago, & Tédongap, 2017; Gomes, 2005; Schmidt & Zank, 2005). Common among these studies is their departure from expected utility (Rabin, 2013, chapter 13) which implies equal treatment of gains and losses, and subsequently the consideration of frameworks that reflect asymmetries in the way individuals weigh expected profits and losses in their asset allocation decisions. The incorporation of asymmetric risk preferences in an asset allocation exercise is linked to two desirable effects. First, it leads to a theoretically more sound model that reflects well-documented behavioral biases of individuals in the face of uncertainty; second and a direct consequence of the aforementioned, it can generate portfolio allocations that are consistent with the composition of market portfolios which would also mirror investors' expectations on future movements of the participating asset classes.

The most prominent theoretical proposition that codifies individuals' attitude towards potential monetary outcomes is prospect theory (PT, hereafter), which has proved particularly successful in capturing frequently encountered traits of investors' behavior (Barberis & Huang, 2008; Bernard & Ghossoub, 2010; Dimmock & Kouwenberg, 2010; Kahneman & Tversky, 1979). In PT, investors are assumed to evaluate the performance of their investments by anchoring to some historical reference point, engaging in the computation of gains and losses relative to that point. Investors also respond asymmetrically to gains versus losses—courtesy of loss aversion—by exhibiting greater sensitivity to losses compared to gains.

A derivative of PT is the theory of disappointment aversion (DA, herafter), formally introduced by Gul (1991), which highlights the role of disappointment in decision making. Disappointment reflects the emotional response of individuals to outcomes falling below their expectations (Bell, 1982; 1985; Summers & Duxbury, 2012) and constitutes one of the most intense and frequently experienced negative emotions in real life (Schimmack & Diener, 1997; Weiner, Russell, & Lerman, 1979). From an asset allocation perspective, the impact of DA has been confirmed in a series of

* Corresponding author.
E-mail addresses: vasileios.kontosakos@allianz.com (V.E. Kontosakos), shwang@skku.edu (S. Hwang), Vasileios.Kallinterakis@durham.ac.uk (V. Kallinterakis), Athanasios.Pantelous@monash.edu (A.A. Pantelous).

https://doi.org/10.1016/j.ejor.2023.07.038

studies (see Ang, Bekaert, & Liu, 2005; Dahlquist et al., 2017; Fielding & Stracca, 2007; Gul, 1991) involving single-period, mostly static, settings. However, investors' focus need not be confined within a single-period, as it may well stretch across multiple periods involving a variety of horizons (be they shorter or longer ones), thus suggesting that the impact of DA over asset allocation decisions may exhibit horizon effects. This, in turn, would necessitate the study of DA in multi-period settings, something that has to date received little attention in the literature.

This paper aims at addressing this gap in the literature by formalizing a multi-period dynamic framework in partial equilibrium[1], which allows for sequential investing and reallocation of the available wealth under the impact of return predictability and uncertainty around the true values of the parameters used to model expected movements of the risky asset. The main tool for our analysis is Dynamic Programming (DP) which allows for recursively solving investor's portfolio allocation problem at each investment horizon, accommodating the stochastically-driven movements of the risky asset. Doing so would allow one to investigate how DA preferences affect the decision-making of people who seek to maximize their expected utility of wealth over long periods. In our setup, investors choose portfolio allocations for two investable assets, a risk-free, and a risky asset. Because it is impossible to, a priori, assert investors' beliefs as per the return-generation process in the market, we consider two possible data-generating processes (DGPs). The first one assumes returns are independent and identically distributed (i.i.d.), and the second one assumes a vector autoregression (VAR) that uses the dividend price ratio (i.e., the dividend yield) as its predictor variable. An issue arising from the above two DGPs pertains to the uncertainty inherent in their estimated parameters and how investors treat it. To assess the impact of parameter uncertainty, we compare asset allocations to the risky asset between cases in which parameter uncertainty is ignored and others where it is considered, for both DGPs.

Our main focus is on whether return predictability and parameter uncertainty generate horizon effects (i.e., different portfolio allocations at different investment horizons). We provide the theoretical foundation for the multi-period dynamic asset allocation problem along with a detailed simulation study where the effects of return predictability and parameter uncertainty are considered. This study presents evidence strongly supporting the role of DA in defining equity participation (and nonparticipation) regions, particularly in the long-run, thus contributing to the explanation of the *nonparticipation puzzle*. The latter, arises when individuals refrain from investing in risky assets, in favour of risk-free ones, despite the historical outperformance of the first over the second. We find that for every portfolio allocation and level of expected equity return, there is a critical value of DA below which it is optimal for a DA investor to hold zero units of the risky asset. More interestingly, we find that DA investors tend to allocate significantly less to equity compared to investors with isoelastic (power) utility. DA appears to be powerful in asset allocation at every horizon: we find that a small increase of the DA coefficient leads to a significant decrease in equity holdings in the case of an investor who accounts for predictability in stock returns.

Our study is comparable with Ang et al. (2005) in the sense that both deal with portfolio optimization in the presence of asymmetric risk preferences. Nevertheless, despite Ang et al. (2005) focusing on the single- or two-period portfolio problem, we are interested in multi-period dynamic portfolio optimization. In doing so we observe the variability in the composition of the optimal portfolio over time as a result of investors' risk preferences and the intertemporal evolution of the investment opportunity set. The main drivers for the latter are the consideration of asset return predictability and model parameter uncertainty as to the calibration of the DA portfolio model. Applications of parameter uncertainty in multi-period models are in general limited while, to our knowledge, parameter uncertainty has not been studied in a dynamic programming DA framework.

A knock-on effect of the use of DP in solving dynamic systems is the expansion of the state space (i.e., the universe of all possible solutions among which the optimal one lies too) that gives rise to the known *curse of dimensionality* of DP. This particularity is unlikely to emerge in a single- or two-period portfolio problem but it is relevant to our multi-period long-term dynamic one as the exponential increment of the state space size makes the problem numerically cumbersome. Remedy to this is provided by applying a dimensionality reduction technique discussed in Epstein & Zin (1989), Kreps & Porteus (1979) and Ang et al. (2005).

With this analysis, we produce a series of original contributions to the extant literature on investors' portfolio choices under asymmetric risk preferences, contextualized by predictability and parameter uncertainty. First, we extend the literature of investors' portfolio choices with DA utility by providing optimal equity participation conditions and nonparticipation regions for portfolio allocations. Second, we study the portfolio choice problem for a long-term buy-and-hold investor under return predictability and parameter uncertainty. Although our study primarily focuses on dynamic portfolio choice, revisiting the buy-and-hold asset allocation problem for very long investment horizons (up to 40 years) reveals a number of important implications for the different investment behaviors of a long versus a short-term DA buy-and-hold investor. To our knowledge, this specific focus has not attracted attention so far.[2] Third, we demonstrate how the incorporation of predictability into asset returns affects portfolio weights at different horizons for a dynamic investor and how this can give rise to horizon effects, in the sense that investors change their portfolio's composition taking into account the variability in investment opportunities. Finally, we complete our study constructing a Bayesian framework that incorporates both predictability and parameter uncertainty to investigate how each of the two properties affect portfolio compositions in a DA context. Here, the choice of the risky asset return generator is crucial; for example, the impact of parameter uncertainty on a dynamic strategy, where returns are i.i.d., is not as powerful as that in the case in which predictability is considered, leading to significantly different portfolio allocations.

2. Literature review

Our research is primarily motivated by extant evidence, according to which investors do not strictly adhere to the assumptions of expected utility theory in their decisions, instead being prone to viewing choices in a biased fashion, often under the influence of emotions and cognitive biases, such as mental accounting (i.e., the inclination to evaluate portfolio's assets in isolation; Barberis, Huang, & Santos, 2001) and framing which can prompt investors to choose an option that appears attractive on the background of less-attractive alternatives and not because it is the optimal option (for a more detailed discussion of the above, see Kahneman, Lovallo, & Sibony, 2011). Several studies depart from the expected utility framework to better explain investors' decision-making under risk, transforming probabilities into decision weights through nonlinear

[1] Our approach employs an exogenous price setting, making this study a partial equilibrium one. Modeling the cash flows can lead to an endogenous price setting, where equilibrium asset prices are attained and markets clear. For an example of an equilibrium study, see Lynch (2000).

[2] Dynamic portfolio allocation has, overall, been widely studied, in both discrete and continuous time (Aït-Sahalia, Cacho-Diaz, & Hurd, 2009; Brandt, Goyal, Santa-Clara, & Stroud, 2005; Campbell & Viceira, 2002), however none of these studies considers a utility function that accommodates asymmetric preferences.

probability functions (see Diecidue, Schmidt, & Zank, 2009; Gonzalez & Wu, 1999; Kilka & Weber, 2001 among others). This in turn generates skewed, instead of symmetric, probability distribution functions for the expected return of the risky assets which has implications for the portfolio allocation decisions of market participants. For example, in the case of PT, at odds with intuition, an indicative market-observed effect of loss aversion is investor's growing appetite for risk when in the domain of losses (they hold onto their loser stocks hoping for a price rebound) while they become more risk averse in the domain of gains (they sell their winner stocks to realize profits, while profits still exist). This in essence leads investors to sell their winning assets more quickly compared with their losing ones. Empirical evidence (Grinblatt & Keloharju, 2001; Haigh & List, 2005; Jin & Scherbina, 2010; Locke & Mann, 2001; Odean, 1998; Shapira & Venezia, 2001; Wermers, 2003) suggests that this pattern permeates both retail and institutional investors' behavior internationally, yet leads to sub-optimal performance. The latter has been ascribed to the effect of short-term momentum (Jegadeesh & Titman, 1993; 2001) in stock returns, according to which recent winners (losers) will continue outperforming (underperforming) in the near future; this finding, in turn, suggests that investors in the prospect theory setting should keep their winners (instead of quickly selling them) and sell their losers (instead of keeping them).

The derivative DA theory extends the expected utility theory by relaxing the independence axiom, while retaining the basic features of PT (asymmetric preferences, reference dependence, diminishing sensitivity, and probability weighting). Moreover, it provides us with better understanding in the way the certainty equivalent of wealth is chosen and updated. In a DA context, the certainty equivalent of wealth serves as a reference point for investor's wealth against which gains and losses are compared, and is updated in an endogenous way. On the other hand, in PT these points are set exogenously and are usually equal to the current wealth (the status quo; see Baillon, Bleichrodt, & Spinu, 2020; Werner & Zank, 2019, for potential ways reference points can be chosen in PT).

In terms of asset allocation, by deriving the optimal portfolio comprising a risky and a risk-free asset for an investor who uses the DA utility function, we extend the DA-related literature. Overall, empirical applications of the DA theory have been rather limited to date, a fact attributed by Abdellaoui & Bleichrodt (2007) to the theory lacking a method of formally extracting the DA coefficient. To that end, Abdellaoui & Bleichrodt (2007) proposed a trade-off method, which first derives the underlying utility function and then, based on the function, extracts the DA coefficient. In asset allocation setups, Dahlquist et al. (2017) employed DA preferences to derive analytical expressions for measures such as the effective risk aversion when studying higher moments of return distributions.

More importantly, we contribute to the growing portfolio choice literature, which discusses the incorporation of parameter uncertainty into the asset allocation topic. Relevant literature (Avramov & Zhou, 2010; Barberis, 2000; Bawa, Brown, & Klein, 1979; Kacperczyk & Damien, 2011; Kandel & Stambaugh, 1996) integrates several forms of uncertainty (model, parameter, or distribution) with asset allocation decision-making. Recently obtained evidence further corroborates the importance of predictability and parameter uncertainty for portfolio choices. Branger, Larsen, & Munk (2013) and DeMiguel, Martín-Utrera, & Nogales (2015) examine the construction of optimal portfolios under uncertainty about expected asset returns and find that parameter uncertainty is highly relevant to portfolio choice. Chen, Ju, & Miao (2014) study the dynamic portfolio choice problem when investors face uncertainty about the model's specification, incorporating learning to construct strategies that depart from the Bayesian approach. Hoevenaars, Molenaar, Schotman, & Steenkamp (2014) test the impact of different uninformative priors on both short- and long-term equity allocations, whereas Johannes, Korteweg, & Polson (2014) investigate the impact of predictability and parameter uncertainty in an expected utility framework mainly focusing on the impact of volatility on the portfolio choice problem.

Our paper is structured as follows: Section 3 formally introduces the DA utility and the portfolio allocation framework along with the definition of the predictability and parameter uncertainty used in our empirical design. In Section 4, the case of equity non-participation is discussed followed by a simulation study for the empirical part of the study in Section 5. Concluding remarks are discussed in Section 6. A number of online available appendices can be found as supplementary material to this paper. The interested reader can have access to technical details for the problem formulation, the algorithmic procedure, the incorporation of parameter uncertainty and predictability in the DA asset allocation context and the performance of algorithmic implementation.

3. Extension of the DA framework

We introduce the theoretical components for the DA optimization framework under the present section. In specific, we start with the definition of the DA utility function (Section 3.1) and the necessary extensions to accommodate multi-period problems (Section 3.2). The case of parameter uncertainty is considered separately in Sections 3.5 and 3.6 after introducing the two DGPs in Section 3.4.

3.1. A classic DA setup

We define the DA utility employed (see, Ang et al., 2005) in this study as follows:

$$\mathcal{U}(\mu_W) = \frac{1}{K}\left(\int_{-\infty}^{\mu_W} U(W)dF(W) + A\int_{\mu_W}^{\infty} U(W)dF(W)\right), \qquad (1)$$

where A is the coefficient of DA, bounded between zero and one (i.e., $0 < A \leq 1$); $U(\cdot)$ is the constant relative risk aversion (CRRA) utility function defined by $U(W) = W^{1-\gamma}/(1-\gamma)$, where W denotes wealth; μ_W is the implicitly defined certainty equivalent of wealth; $F(\cdot)$ is the cumulative distribution function for wealth; and K is a scalar equal to $P(W \leq \mu_W) + AP(W > \mu_W)$.[3]

Eq. (1) was selected as the driver of the portfolio choice problem because it embodies the main features of the DA theory. First, it splits the outcomes based on whether they are smaller or larger than certainty equivalent μ_W; second, using the DA coefficient A it is clear that it penalizes portfolio wealth states better than μ_W, while it also skews the probability distribution of W with the scalar K.[4] It is expected that μ_W acts as the reference point in defining both the scalar factor K and the sum of the integrals in Eq. (1). The reference point is of no relevance to the portfolio allocation problem only when $A = 1$ as the DA preference reduces to CRRA. In practice, given that A cannot be larger than one, any outcome worse than μ_W is weighted more heavily than one better than it. Intuitively, asymmetries in risk preferences entail that downside portfolio movements occur *more frequently* (i.e., larger

[3] For the purpose of our study, the definition of P can deviate from that of the physical probability measure strictly defined in a probability triple. In practice, we use the quadrature probabilities provided by a Gauss-Hermite scheme as it is discussed in the following sections, which are transformed via K.

[4] In asset pricing, a slightly modified version of the original DA theory is used (Bonomo, Garcia, Meddahi, & Tédongap, 2011; Routledge & Zin, 2010) in which DA utility in Eq. (1) is extended so that an outcome signals *disappointment* only when it lies sufficiently below the certainty equivalent. The effect is captured by an additional coefficient that coexists with the DA parameter A.

probabilities are assigned) and as a result, they should be relevant to the DA investor.

Assume two assets, one risky asset and one risk-free asset, whose continuously compounded returns are denoted by e^y and e^r respectively. Then the investor's wealth is defined as $W = \alpha X + e^r$, where α is investment proportion in the risky asset; $X = e^y - e^r$ is the excess risky asset's return and the initial wealth is set to one, because the optimization problem is homogeneous in wealth under the CRRA utility function. When a DA investor allocates her wealth into assets in order to maximize the DA utility for a single period, the static optimization problem is

$$\max_{\alpha} U(\mu_W). \tag{2}$$

The above constitutes the asset allocation problem under the assumption of DA utility in a single-period setting.

3.2. Formulation of the dynamic allocation optimization problem with DA utility

We extend the static setup in Eq. (1) to an equation with time-dependent wealth, certainty equivalent and probability weighting for the potential wealth states as captured in K. The aim of this extension is to derive the expressions that will be converted to the objective functions used for the portfolio optimization problem. A dynamic optimization problem with DA utility in a multiple-period setting is considerably more complex than in a static one, because at every horizon the optimization routine should take into account the investment opportunity set for the whole remaining investment period (as opposed to a one-period forward-looking myopic strategy), while the certainty equivalent of wealth is itself a function of each horizon's optimal decision. The complexity of the optimization problem further increases by considering predictability, which leads to variable investment opportunity sets. We begin by building the optimization problem for a general utility function defined over wealth $U(W)$ and we then move to dynamic asset allocation with DA utility. In doing so, we have the chance to define variables that are also used in the DA optimization problem, in a simpler framework without the additional complexity that comes from the wealth0dependent utility function.

3.2.1. Dynamic asset allocation with general utility function

Assume the following dynamic asset allocation problem in discrete time in which an agent aims to maximize the expected utility of the end-of-period wealth W_T as follows:

$$\max_{\alpha_0,\alpha_1,\ldots,\alpha_{T-1}} \mathbb{E}_0[U(W_T)], \tag{3}$$

where $\alpha_0, \alpha_1, \ldots, \alpha_{T-1}$ are the investment proportions of the risky asset at times $t = 0, 1, \ldots, T-1$, respectively, and $U(W) = W^{1-\gamma}/1-\gamma$. In this problem, the investor allocates her wealth at time $t = 0$ for T periods, at $t = 1$ for $T-1$ periods and so on until she reaches time $t = T-1$, where she invests for a single period.[5] Wealth W_{t+1} is defined as $W_{t+1} = W_t R_{t+1}(\alpha_t)$, where $R_{t+1}(\alpha_t)$ and α_t represent the portfolio return over the period t to $t+1$ and the investment weight on the risky asset at time t, respectively. At time t when the investor seeks to allocate her available wealth optimally between the risky and the riskless asset in order to maximize her expected utility, the optimization problem becomes

$$\max_{\alpha_t} \mathbb{E}_t[U(W_{t+1}Q^*_{t+1,T})], \tag{4}$$

where $Q^*_{t+1,T} = R_T(\alpha^*_{T-1})R_{T-1}(\alpha^*_{T-2})\cdots R_{t+2}(\alpha^*_{t+1})$ represents the aggregate return-to-go over the investment horizon generated by the optimal risky asset allocation $\alpha_{t+1}, \alpha_{t+1}, \ldots, \alpha_{T-1}$ that maximizes investor's expected utility.

Using dynamic programming, we can solve the problem at time $t = T-1$ for the asset allocation decision for the period $T-1$ to T. Continuing recursively, we can solve the asset allocation subproblem at time $T-2$ using the solution to the problem at $T-1$, until we reach time t. This procedure derives a final solution for the portfolio allocation to the risky asset $\alpha_t, \alpha_{t+1}, \ldots, \alpha_{T-1}$ that will be optimal as guaranteed by the principle of optimality in dynamic programming.[6] For the power utility function, the objective function in (4) takes the form of

$$\max_{\alpha_t} \mathbb{E}_t\left[\frac{W_{t+1}^{1-\gamma}}{1-\gamma}(Q^*_{t+1,T})^{1-\gamma}\right]. \tag{5}$$

Backward induction suggests that $Q^*_{t+1,T}$ represents the optimal investment decision between times $t+1$ and T that maximizes the expected utility. We calculate the optimal investment proportions of the risky asset at every time step of the investment period as

$$\alpha^*_t = \arg\max_{\alpha_t} \mathbb{E}_t\left[W_{t+1}^{1-\gamma}(Q^*_{t+1,T})^{1-\gamma}\right]. \tag{6}$$

3.2.2. Dynamic asset allocation with DA utility

Having defined basic building blocks of the optimization setup, we now turn to the portfolio optimization problem under DA utility preference which is in-scope of the study. DA utility incorporates CRRA preferences as a special case in which $A = 1$, but the dynamic extension of the single-period problem for DA utility is far more complicated because of the so-called *curse of dimensionality*, i.e., the number of state variables exponentially increases with time.[7] We begin by first formulating the dynamic optimization problem between t and T.

Proposition 1. *For given $Q^*_{t+1,T} = R_T(\alpha^*_{T-1})R_{T-1}(\alpha^*_{T-2})\cdots R_{t+2}(\alpha^*_{t+1})$, the DA utility function for the dynamic asset allocation problem is given by*

$$U(\mu_t) = \frac{1}{K_t}\left[\mathbb{E}_t\left(U(W_{t+1}Q^*_{t+1,T})\mathbf{1}_{W_{t+1}Q^*_{t+1,T}\le\mu_t}\right) + A\mathbb{E}_t\left(U(W_{t+1}Q^*_{t+1,T})\mathbf{1}_{W_{t+1}Q^*_{t+1,T}>\mu_t}\right)\right], \tag{7}$$

where $W_{t+1}Q^*_{t+1,T} = W_T$, according to the recursive definition of wealth. The first-order condition (FOC) for optimizing the utility of the certainty equivalent return is given by

$$\mathbb{E}_t\left(\frac{dU(W_T)}{dW}Q^*_{t+1,T}R_{t+1}(\alpha_t)W_tX_{t+1}\mathbf{1}_{W_T\le\mu_t}\right) + A\mathbb{E}_t\left(\frac{dU(W_T)}{dW}Q^*_{t+1,T}R_{t+1}(\alpha_t)W_tX_{t+1}\mathbf{1}_{W_T>\mu_t}\right) = 0, \tag{8}$$

where $X_{t+1} = e^{y_{t+1}} - e^{r_t}$ is the excess return of the risky asset over the riskless asset.

Proof. See Appendix 9.1 □

[5] This problem mimics the optimization problem that pension fund managers face over multiple periods, (Xie, Hwang, & Pantelous, 2018, e.g.,).

[6] See Bertsekas (1995) for more details on that.

[7] As the state variables take a number of different values at each horizon, the state-space exponentially increases with time with every iteration of the algorithm. For example, a T-period problem with a state variable with s states produces s^T possible combinations. From an analytical perspective, this is not a big obstacle (as the problem still can be mathematically formulated), but computation-wise, the exponential increment of the state-space renders the use of algorithmic processes problematic.

For investment horizon T, the recursive formulation for the utility of the DA investor at time $T-1$ is defined as follows:

$$K_{T-1}U(\mu_{T-1}) = \mathbb{E}_{T-1}[U(W_T R_T(\alpha_{T-1}))\mathbf{1}_{W_T R_T(\alpha_{T-1})\le\mu_{T-1}}] + A\mathbb{E}_{T-1}[U(W_T R_T(\alpha_{T-1}))\mathbf{1}_{W_T R_T(\alpha_{T-1})>\mu_{T-1}}], \quad (9)$$

where $K_{T-1} = \mathbb{P}(W_T R_T(\alpha_{T-1}) \le \mu_{T-1}) + A\mathbb{P}(W_T R_T(\alpha_{T-1}) > \mu_{T-1})$ and $U(\mu) = \mu^{1-\gamma}/(1-\gamma)$. Accordingly, at $T-2$ the utility is defined as

$$K_{T-2}U(\mu_{T-2}) = \mathbb{E}_{T-2}[U(W_{T-1}R_T(\alpha^*_{T-1})R_{T-1}(\alpha_{T-2}))\mathbf{1}_{W_{T-1}R_T(\alpha^*_{T-1})R_{T-1}(\alpha_{T-2})\le\mu_{T-2}}] + A\mathbb{E}_{T-2}[U(W_{T-1}R_T(\alpha^*_{T-1})R_{T-1}(\alpha_{T-2}))\mathbf{1}_{W_{T-1}R_T(\alpha^*_{T-1})R_{T-1}(\alpha_{T-2})>\mu_{T-2}}], \quad (10)$$

where $K_{T-2} = \mathbb{P}(W_{T-1}R_T(\alpha^*_{T-1})R_{T-1}(\alpha_{T-2}) \le \mu_{T-2}) + A\mathbb{P}(W_{T-1}R_T(\alpha^*_{T-1})R_{T-1}(\alpha_{T-2}) > \mu_{T-2})$. By the same token, all intermediate optimization steps are defined recursively up to time $t=0$ where a single-period optimization problem is solved, similar to a that of a buy-and-hold investor. The main drawback with Proposition 1 is that recursive optimization exponentially enlarges the state space in $Q_{t+1,T}$ in order to take into account all the possible states for the return of the risky asset between times $t+1$ and T. Considering the time evolution of risky asset returns as a grid of multiple discrete states (equivalent to a binary tree in case of only two states, up and down), there is no reason to assume that returns will be recombining. As a result, the recursive formulation at $T-2$ requires tracking all states of risky asset return both at $T-2$ and $T-1$, hence they are both part of Eq. (10). To overcome the curse of dimensionality, we elaborate on the approach also met in Epstein & Zin (1989), by considering that future uncertainty about risky asset's returns is captured by the certainty equivalent. Under this approach, instead of carrying backward all the possible states for the equity return at each horizon, we pay attention to only one variable, next-period's certainty equivalent, keeping the dimension of the state space fixed over time. Let μ_t represent the certainty equivalent return for the utility at time $t+1$ with the optimal asset allocation:

$$\max_{\alpha_t} \mathbb{E}(U(W_{t+1})) = \max_{\alpha_t} U(W_t\mu_t(\alpha_t)). \quad (11)$$

Then we obtain the following result:

Proposition 2. *The utility of the certainty equivalent return at time $0 \le t < T-1$ is as follows:*

$$U(\mu_t) = \frac{1}{K_t}\left[\mathbb{E}_t\left(U(R_{t+1}(\alpha_t)\mu^*_{t+1}W_t\prod_{i=t+2}^{T-1}\mu^*_i)\mathbf{1}_{\{R_{t+1}(\alpha_t)\le\xi_t\}}\right) + A\mathbb{E}_t\left(U(R_{t+1}(\alpha_t)\mu^*_{t+1}W_t\prod_{i=t+2}^{T-1}\mu^*_i)\mathbf{1}_{\{R_{t+1}(\alpha_t)>\xi_t\}}\right)\right]. \quad (12)$$

The value of $U(\mu_t)$ for the boundary condition $t = T-1$ is given by

$$U(\mu_{T-1}) = \frac{1}{K_{T-1}}\left[\mathbb{E}_{T-1}\left(U(R_T(\alpha_{T-1})W_{T-1})\mathbf{1}_{\{R_T(\alpha_{T-1})\le\mu_{T-1}\}}\right) + A\mathbb{E}_{T-1}\left(U(R_T(\alpha_{T-1})W_{T-1})\mathbf{1}_{\{R_T(\alpha_{T-1})>\mu_{T-1}\}}\right)\right], \quad (13)$$

and the FOC for optimizing the utility of the certainty equivalent return is given by

$$\mathbb{E}_t\left(\frac{dU(R_{t+1}(\alpha_t))}{d\alpha_t}X_{t+1}\mu^*_{t+1}\mathbf{1}_{\{R_{t+1}(\alpha_t)\le\xi_t\}}\right) + A\mathbb{E}_t\left(\frac{dU(R_{t+1}(\alpha_t))}{d\alpha_t}X_{t+1}\mu^*_{t+1}\mathbf{1}_{\{R_{t+1}(\alpha_t)>\xi_t\}}\right) = 0, \quad (14)$$

where $\xi_t = \frac{\mu_t}{\mu^*_{T-1}\cdots\mu^*_{t+1}W_t}$, with μ^*'s as the optimal certainty equivalents between $t+1$ and $T-1$.

Proof. See Appendix 9.2. □

Remark Notice that W_t will eventually not be part of the expressions for $U(\mu_t)$ in Eqs. (12) and (13) as moving backward in time we will have $W_t = W_0\prod_{i=1}^{t} R_i(\alpha_{i-1})$, where all uncertainty about $R_n(\alpha_{n-1})$, where $n \in \{t+1, t+2, \ldots, T\}$, will be captured by the certainty equivalent return μ^*_n, where $n \in \{t+1, t+2, \ldots, T\}$, and W_0 is set to one given wealth homogeneity. At each horizon, we need to track only the states for μ^*_n, keeping the dimension of the state space for μ^*_n fixed and allowing for horizon effects (hedging demands) in case the DA investor is not at horizon $T-1$.

Investor's gains or losses at time $t+1$ are now calculated with respect to ξ_t, that is, the certainty equivalent at time t for the optimal certainty equivalent from $t+1$ to T. Adopting the dynamic DA utility in Eq. (12), next period's optimal certainty equivalent μ^* is used to define this period's DA utility. To solve for the optimal certainty equivalent μ^* and portfolio α we use dynamic programming. As an example of the advantage of using the certainty equivalent, we can rewrite the FOC in Eq. (14) for power utility as follows:

$$\mathbb{E}_t\left(R^{-\gamma}_{t+1}(\alpha_t)X_{t+1}\mu^*_{t+1}\mathbf{1}_{R_{t+1}(\alpha_t)\le\xi_t}\right) + A\mathbb{E}_t\left(R^{-\gamma}_{t+1}(\alpha_t)X_{t+1}\mu^*_{t+1}\mathbf{1}_{R_{t+1}(\alpha_t)>\xi_t}\right) = 0. \quad (15)$$

By using certainty equivalent as each period's endogenously defined reference point, based on return expectations, the dimension of the state-space remains unchanged with time. The benefit from reducing the dimensionality of the state space is amplified as the investment horizon increases, since for longer horizons the optimization problem using numerical methods becomes intractable. Then at each time step we need to determine the optimal values for μ and α that simultaneously solve Eqs. (12) and (15). For this we adopt a Gaussian quadrature scheme (see Davis & Rabinowitz, 2007, for an in-depth review of numerical integration methods) as in Balduzzi & Lynch (1999) and Campbell & Viceira (1999) that allows us to discretize the continuous probability distributions of the risky asset's returns in the $\mathbb{E}(\cdot)$ in the two equations. In essence, the scheme allows us to track the states $\left\{R_{t+1}(\alpha_t)\mu^{s}_{t+1}{}^*\right\}_{s=1}^{N}\left(\prod_{i=t+2}^{T-1}\mu^*_i\right)$, where N is the number of quadrature states for the certainty equivalent return and R_{t+1} is the return given by the corresponding DGP.[8] We then solve the discretized expression of Eq. (12) (adjusted for power utility in place of $U(\cdot)$) in parallel with the FOC for the DA maximization problem in Eq. (15) recursively incorporating the calculations from periods $T-1$ to $t+1$. Due to the complexity of the notations and a necessary brief introduction of the discretization method, details on the solution of the system of equations are laid out in Appendix 8 of the supplementary online material.

3.3. Asset allocation with parameter uncertainty

The second theoretical result is the formulation of the DA portfolio optimization problem in the presence of parameter uncertainty. Investors who account for parameter uncertainty consider the true values of model parameters to be unknown. In fact, they acknowledge that by continuously updating risk and return expectations by incorporating new information over time, the values of model parameters are constantly changing, and as a result, their actual values are never known with complete certainty. In practice, the difference to the case where parameters are considered fixed

[8] Instead of quadrature-based methods, Monte-Carlo simulations or, even, regression-based methods, like in Brandt et al. (2005), can be used to calculate the expectations in Eq. (15). In practice however, the quadrature method offers sufficient accuracy and greater computational speed compared with the alternatives.

and known rests on whether investor treats model parameter as *inputs* or *outputs* of the portfolio model. Overall, parameter uncertainty has been discussed extensively in static portfolio setups. Part of the literature includes Bayesian frameworks that use asset pricing models (Pástor, 2000; Pástor & Stambaugh, 2000), diffuse priors (Bawa et al., 1979; Brown, 1978), robust optimization (Garlappi, Uppal, & Wang, 2007) and shrinkage methods (Kourtis, Dotsis, & Markellos, 2012; Wang, 2005). In terms of multi-period portfolio optimization, parameter uncertainty in mean-variance portfolio setups is studied among others in Barberis (2000) and DeMiguel et al. (2015). In an effort to accommodate parameter uncertainty in a multi-period DA framework and observe the horizon effects in the long-term portfolio allocation, we formulate the corresponding optimization problem.

First, as a general rule, one can investigate the effects of parameter uncertainty on asset allocation by allowing uncertainty in the parameter estimates (e.g., the mean and variance of asset returns and the correlation matrix as well in case of more advanced modelling) as opposed to the case where model parameters are treated as known. Faced with uncertainty, investors maximize the following utility function at time t :

$$\max_{\alpha} \int_{-\infty}^{\infty} \frac{W_{t+n}^{1-\gamma}}{1-\gamma} p(r_{t+n}|Y;\theta) dr_{t+n}, \tag{16}$$

where $W = \alpha X + e^r$ follows the recursive formulation in Proposition 2; n is the investment horizon, $U(\cdot)$ is the utility of wealth, and $p(r_{t+n}|Y;\theta)$ is the probability density function of the expected returns conditional on observed return data Y and the set of parameters θ (in our case the mean and variance of the risky asset's return). Uncertainty arises for θ, because these parameters become known only after the end of the investment horizon. A popular approach in the literature for manoeuvring the parameter uncertainty problem is to use a Bayesian framework that incorporates uncertainty in the parameters of θ. Integrating out θ in the prior distribution $p(r_{t+n}|Y;\theta)$, we obtain the posterior predictive distribution, which updates the distribution parameters by embodying the new data. A DA investor now maximizes

$$\max_{\alpha} \left[\int_{W_{t+n} \le \mu_W} U(W_{t+n}) p(r_{t+n}|Y) dr_{t+n} + A \int_{W_{t+n} > \mu_W} U(W_{t+n}) p(r_{t+n}|Y) dr_{t+n} \right], \tag{17}$$

in place of Eq. (16), in line with the DA utility definition in Eq. (1), where the distribution of the returns is now conditional on observed stock return data only, not on the set θ. To obtain optimal portfolios under parameter uncertainty, the optimization procedure as described in Section 3.2.2 is followed. Nevertheless, in order to account for parameter uncertainty in the generation of asset returns, a sampling algorithm sensitive to the choice of the underlying DGP needs to be applied. Hence, we take a detour to briefly introduce the two DGPs used in the study before we proceed with the remaining results.

3.4. Definition of the DGPs

We introduce the two DGPs used to derive optimal portfolio allocations under the theoretical models introduced in this section. We use return generators with the same properties throughout the portfolio optimisation exercise to allow comparison between the cases where parameter uncertainty is ignored and when it is considered in the calibration of the portfolio model. Following the extensive discussion in the finance literature, we consider two systems: one that assumes equity returns are i.i.d., and a vector autoregression that uses the equity price-dividend ratio as a predictor of equity return.

First, a straight-forward data generator comprising a single equation is presented. This DGP is able to model i.i.d. returns with no autocorrelation structure. The second data generator is a more complex vector autoregression (VAR) in which one period's asset return is related to (i.e., can be 'predicted' by) previous period's asset price dividend ratio. As opposed to the first one, the second model allows for some degree of correlation between the response (asset return) and the explanatory variable (dividend ratio). The purpose of introducing the i.i.d. return generator and producing a full-fledged portfolio analysis based on it (see Section 5), before deriving the same set of results for the VAR, is to have a basis that allows us to isolate the effects introduced by return predictability as modelled by the VAR.

3.4.1. i.i.d. returns

When investors ignore predictability in returns, they consider them to be i.i.d., and they use the following model to estimate next-period's excess equity return:

$$x_t = (\mu - r) + \epsilon_t. \tag{18}$$

In our case, x_t is the continuously compounded excess return in period t, μ is the mean asset return over some specified period, r is the yield of a riskless asset for the same period as defined by the modeller and ϵ_t are i.i.d. disturbance terms distributed as $\epsilon_t \sim \mathcal{N}(0,\sigma)$, where σ is the return volatility. By providing values for μ, r and σ one can trivially extract excess return values out of Eq. (18). This is the simplest way one can sample return values from a probability distribution as it requires only random sampling for the error terms in ϵ without paying any attention to autocorrelation structures.

3.4.2. Return predictability

In practice, asset returns are not i.i.d. Researchers have documented risk drivers that can be used to predict part of the variability in asset returns (Ang & Bekaert, 2007; Campbell & Yogo, 2006; Cochrane, 2008; Lettau & Ludvigson, 2001). Investors use available information to predict future returns for optimal asset allocation problems. In our study, we replicate the prediction process using a VAR model, where asset returns and the predictable variable are jointly considered.

This results in time-varying investment opportunity sets, as opposed to time-independent sets when returns are i.i.d., which are conditional on the predictor variable in the VAR model. Investors react to the variability by modifying the proportion of their current investment allocated to the risky asset. To determine a suitable driver for equity return, we examined a number of financial variables,[9] and chose the dividend yield, calculated as the dividend price ratio for of the risky asset, to drive next-period's equity return. The optimal number of lags was calculated as one, confirmed by both the Akaike and the Bayesian information criteria. We then model the dividend-adjusted log excess returns of the risky asset as a first-order VAR of the following form:

$$X_t = C + BX_{t-1} + E_t. \tag{19}$$

[9] To determine the variable that best fits our data, we test the following predictors: dividend yield (the sum of the dividends over a year divided by the level of the index at the end of the year; provided that the asset pays dividends); term spread (the difference between the 10-year Treasury bond and the 1-year Treasury bond); credit spread (the difference between Moody's BAA corporate bond yield and its AAA equivalent); the 3-month Treasury bill; and the 10-year Treasury bond. The criteria for selecting the best fit are (a) whether a variable enters the VAR as statistically significant and (b) how much of the risky asset's excess return variability it explains. More thorough analysis can be of course conducted for detecting the most suitable predictor or set of predictors but it is outside the scope of the current study. The objective of this exercise is to observe the effects introduced by return predictability compared to the case where returns are i.i.d., which in practice (see Section 5) is possible by using just one predictor variable.

Table 1
Summary statistics.

	S&P 500	3-month T-bill	Excess Return
	Annualized		
mean	0.1045	0.0344	0.0695
stdev	0.1625	0.0088	0.1644
	Quarterly		
mean	0.0251	0.0085	0.0166
stdev	0.0817	0.0044	0.0822

Annualized S&P 500 and T-bill summary statistics. Excess return is calculated by subtracting the 3-month T-bill rate from the value of the S &P 500 for the same period.

Table 2
Parameter estimates for the Data Generating Process (VAR).

Parameter	With predictability	Without predictability
c_1	0.1222	0.0128
	(0.0173)	(0.0178)
c_2	−0.0004	−0.0317
	(0.0119)	(0.0150)
b_{11}	0.0259	0.0
	0.1176	–
b_{12}	0.0220	0.0
	(0.1354)	–
b_{21}	−0.7068	0.0
	(0.0807)	–
b_{22}	0.9978	0.9932
	(0.0929)	(0.0912)
σ_{11}	0.0850	0.0856
	(0.0037)	(0.0042)
σ_{22}	0.0408	0.0752
	(0.0017)	(0.0029)
ρ	−0.5216	−0.2980
	(0.0021)	(0.0028)

This table shows VAR estimation and the corresponding standard errors of the parameters for the two systems (predictability/no predictability). We used maximum likelihood (MLE) to calculate the model in Eq. (19). For the nonpredictability system, the autoregressive coefficient matrix is set to zero, whereas, when we account for predictability in returns, all four coefficients are free to vary without restrictions. Parentheses include the standard errors of the estimated coefficients. The S&P 500 and dividend yield quarterly data for the period January 1934 to September 2016 are used in our calculations.

In the model of Eq. (19), $X_t = \begin{pmatrix} y_t - r_{t-1} \\ (d/p)_{t-1} \end{pmatrix}$ where $y_t - r_{t-1} = x_t$ is the excess equity return; r_t is the risk-free rate; $(d/p)_{t-1}$ is the dividend price ratio; B is the (2×2) matrix of the autoregression coefficients; C is a (2×1) vector of the constant terms; and E is a vector of *i.i.d.* normally distributed disturbance terms.

We use the lagged rate r_{t-1} to indicate that the value of the risk-free rate is known at the time of portfolio formation $t-1$, in contrast to the risky asset, whose return becomes known at time t only. When asset returns are not predictable (i.e., DGP of Eq. (18)), all elements of the matrix with the autoregressive coefficients B are not different from zero, and returns are assumed to be i.i.d. As a result, the VAR model reduces to the i.i.d. return generator of Eq. (18). We use maximum likelihood estimation (MLE) to calculate the VAR in Eq. (19), and Table 2 reports the results.

3.5. Parameter uncertainty with i.i.d. returns

Having defined the two asset return generators, we now examine how they affect investor's risk (i.e., variance) and return (i.e., mean) expectations compared to the case where these two variables are treated as known. The difference to the case where mean and variance are known is that investors are now faced with a *distribution* of different mean and variance values, namely the *posterior* distribution, instead of fixed values. The posterior distribution has the property that it absorbs information generated during the investment horizon, based on sampled asset price trajectories from a frequently calibrated DGP, which cannot be know at time $t = 0$. Accordingly, the starting mean and variance expectations, namely the *prior*, are shifted in order to reflect the newly absorbed information. Again, it was important to define the two DGPs prior to the current section because the choice of the DGP changes the way the sampling step mentioned in Section 3.3 is defined and implemented. Assuming investors are unaware of the true parameter values, we use an uninformative (diffuse) prior of the type

$$p(\mu, \sigma) d\mu d\sigma \propto \frac{1}{\sigma} d\mu d\sigma, \tag{20}$$

whereas the joint posterior of the mean return μ and volatility σ is

$$p(\mu, \sigma|Y) \propto p(\mu, \sigma) \times L(\mu, \sigma|Y), \tag{21}$$

where L is the likelihood function. The following lemmas report the results for the case of i.i.d. returns (Lemma 1) and predictive returns (Lemma 2), where the VAR is used.

Lemma 1. *The distribution of the posterior moments for the case of normally distributed i.i.d. returns is given by*

$$\sigma^2|Y \sim Inv - Gamma\left(\frac{N}{2}, \frac{1}{2}\sum_{i=1}^{N+1}(y_i - \overline{\mu})^2\right)$$

$$\mu|\sigma, Y \sim \mathcal{N}\left(\overline{\mu}, \frac{\sigma^2}{N}\right),$$

where Y is the observed asset return data; N is the sample size; and $\overline{\mu}$ is the sample mean return.

Proof. See Appendix 10.1. □

To construct the posterior predictive distribution for the i.i.d. returns of the risky asset, we first sample once from the marginal posterior distribution $p(\sigma^2|Y)$, and then from $p(\mu|\sigma, Y)$, which is now conditional on σ. We repeat this process to generate a sufficiently large number of pairs (μ, σ) to create return values and subsequently the posterior distribution for the returns of the risky asset, by sampling each pair (μ, σ) from the posterior distribution. Appendix 10.1 provides details about the sampling procedure from the derived distributions for the mean and variance.

3.6. Parameter uncertainty with return predictability

The VAR in Eq. (19) also can be written in the following compact form:

$$X = BZ + E, \tag{22}$$

where $X = (X_1 \ldots X_T)$ is a $(2 \times T)$ matrix with the number of observations T for the estimated variables; $Z = (z_0 \ldots z_T)$ a $(3 \times T)$ matrix; $B \equiv (c\,B)$ is a (2×3) matrix of the autoregressive coefficients and the constant terms; and the $E = (\epsilon_1 \ldots \epsilon_T)$ is a $(2 \times T)$ matrix with the serially uncorrelated disturbance terms. A suitable uninformative prior is the Jeffreys prior given by

$$p(B, \Sigma) \quad = p(B)p(\Sigma) \propto |\Sigma|^{-(m+1)/2}, \tag{23}$$

where $m = 2$ is the total number of regressors on the left-hand side of Eq. (22); $p(B)$ is constant and B is independent of Σ, the covariance matrix of the VAR error terms in E. We obtain the posterior density for the parameter matrix B and the covariance matrix of Eq. (22) by the following lemma.

Lemma 2. *The posterior distribution, $p(vec(B), \Sigma|X)$ for the coefficient matrix, B and the variance-covariance matrix, Σ conditional on normally distributed asset return data X is given by*

$$\Sigma|X \sim \mathcal{W}^{-1}((X - Z\hat{B})'((X - Z\hat{B}), T - n - 1)$$

$vec(B)|\Sigma, X \sim \mathcal{N}(vec(\hat{B}), \Sigma \otimes Z'Z^{-1}),$

where T is the number of observations and n is the number of predictor variables.

Proof. See Appendix (10.2). □

Again, to sample from $p(vec(B), \Sigma|X)$, we sample first from $p(\Sigma|X)$–the variance-covariance matrix–conditional on data set X and then from the posterior distribution $p(vec(B), \Sigma|X)$, which will give a draw for the matrix of the VAR coefficients. Appendix (10.2) presents the details of this process and the return-generating procedure.

Given the several portfolio optimization problems and the relative complexity of their solution, a flowchart with the overarching (high-level) solution process of the formulated optimization problems is presented below, see flowchart in Fig. 1.

4. Nonparticipation under DA utility

An interesting result, namely equity *non-participation*, is derived and discussed by showing that it is embedded in the context and theory of disappointment aversion. We derive a theorem which shows that there's always a level of disappointment aversion, let it be A^*, below which the optimal investment in risky asset(s) is zero (i.e., $\alpha^* = 0$). The theoretical finding is backed by numerical examples that track the relationship between α, A and A^* at different horizons.

The case for nonparticipation in risky assets has been a subject of considerable research to date. Mental accounting (Thaler & Sunstein, 2008, which assumes the nonfungibility of monetary resources allocated to each asset; see e.g.,) motivates narrow framing (Barberis & Huang, 2008), which prompts investors to perceive high-volatility assets as “risky” in isolation without assessing their contributions to the risk-return profile of their portfolio. Nonparticipation also can be promoted by the omission bias (Ritov & Baron, 1999), whereby omissions (e.g., not investing in stocks) are favored over equivalent commissions (investing in stocks), because commissions, unlike omissions, involve commitment to a course of action, thus entailing the possibility of a loss. Other alternative explanations proposed to account for nonparticipation include the familiarity bias (Huberman, 2001; Massa & Simonov, 2006, choosing more over less familiar assets, believing the latter to be riskier;), the recognition bias (Boyd, 2001, preferring more over less recognizable assets;) and limited cognition (Hirshleifer, 2008, when investors view risk diversification as a decision of enhanced complexity;).

Under CRRA preferences holding positive portfolio allocations to risky assets when the expected excess return is positive ($\mathbb{E}(X) > 0$) is always optimal. However, this is not always the case with DA utility preferences. Under DA preferences refraining from holding risky assets even if the expected excess return is positive can be optimal in certain cases. The following result shows that it is not optimal to hold risky assets whenever the DA coefficient lies below a critical value A^*.

Theorem 1. *Let $\mu = \mu_W(A, \alpha)$, with*

- $\mu(A, :) \in \mathbf{C}^1, \forall A \in [0, 1]$
- $\frac{d\mu(A,0)}{d\alpha} = \xi(A) \leq 0, \forall A \in [0, 1]$ [10]
- $\mathbb{E}(X) > 0$ *and* $\mathbb{E}(X\mathbf{1}_{W \geq \xi(A)}) > 0$, *where* $X = e^y - e^r$ *is the return of the risky asset in excess of the risk–free rate.*

Then, setting

$$A^* = \frac{\mathbb{E}(X\mathbf{1}_{W \geq \xi(A)})}{\mathbb{E}(X\mathbf{1}_{W < \xi(A)})}, \quad (24)$$

we have the following:

1. *For every $A \leq A^*$, $\alpha^* = 0$,*
2. *For every $A > A^*$, $\alpha^* > 0$,*

where α^* is the optimal investment proportion in the risky asset which maximizes $\mu(A, \alpha)$ for a given A. A^* is independent of the risk aversion parameter γ.

Proof. See section 9.3 □

Our Theorem 1 completes Proposition 2.1 of Ang et al. (2005), and can be intuitively presented in the following way: as DA increases, investors allocate less wealth to the risky asset for an arbitrary level of risk aversion γ. Given that the utility of wealth is a continuous function within the domain of A, there should be a level of A, let A^*, at which the optimal portfolio allocation to the risky asset, α^*, equals zero. Recalling the condition $d\mu(A, 0)/d\alpha \leq 0$, we see that a further decrease in the portfolio weight allocated to the risky asset α^* (e.g., because of short selling the risky asset) will result in a higher certainty equivalent return. When investment in the risky asset is nonzero, an increase in the investment in the risky asset decreases the certainty equivalent. This is intuitively correct, because, by increasing the portfolio allocation to equities to a nonzero level, investors become more willing to accept an amount of risk instead of holding only the risk-free security. This consequently implies that the monetary amount that can keep investors away from buying stocks should be lower. Subsequently, the following relationship will prevail:

$W = \alpha^* X + R > R,$

for $\alpha^* < 0$ and negative states ($X < 0$) of the excess equity return. Therefore, the optimal allocation for this critical level of the DA coefficient, A^*, is zero and $\alpha = \alpha^* = 0$.

Remark The critical level of DA, A^*, is estimated as the ratio between the states of the expected risk premium $\mathbb{E}(X)$ that could result in non-negative wealth ($W \geq \xi(A)$), and those that result in strictly negative wealth ($W < \xi(A)$). Given $\mathbb{E}(X) > 0$ and $A^* \in [0, 1]$ we have that

$$\mathbb{E}(X1_{W \geq \xi(A)}) = A^* \mathbb{E}(X1_{W < \xi(A)}). \quad (25)$$

Eq. (25) shows that a DA investor shapes her expectations of the elating outcomes – states of equity premium that lead to changes in wealth that exceed the negative impact of $\xi(A)$– by underweighting those outcomes that result in strictly negative states of wealth (RHS in Eq. (25)). The coefficient A^* expresses the extent to which the disappointing outcomes are underweighted. As A^* approaches zero, the expectations of the disappointing outcomes are very strongly underweighted, reflecting the low degree of DA and the high level of asymmetry in the DA utility function. Contrary to that, when $A^* \to 1^-$, the two expectations in Eq. (25) are approximately equal, and the investor does not account for DA in her decision making. Equivalently, the investor uses an almost fully symmetric utility function, similar to a power utility function.

A calculation of risky asset participation based on A^*, by using the historical mean equity return from 1934 to 2019 calculated at $\mu = 10.45\%$ as tracked by the S&P500 is presented in Fig. 2. To extract A^*, the level that results in allocating zero wealth to equity, a binary search algorithm is used. A typical binary search algorithm for our problem works by discretizing the state space of A (e.g., a linearly spaced vector within [0,1]) and performing sequential searches for the target value of A (i.e., the one that makes $\alpha = 0$). This is done by iteratively comparing the target value to the middle element of the state space and cutting the state space

[10] Positive risk premium when the end-of-period wealth exceeds the negative impact of the decrease in the certainty equivalent as the investment proportion of the risky asset increases. Suppose that the expected return of the risky asset is zero. The certainty equivalent decreases when the proportion of the risky asset increases. This occurs because for $\alpha < 0$ negative excess return states have higher wealth than R and hence are downweighted.

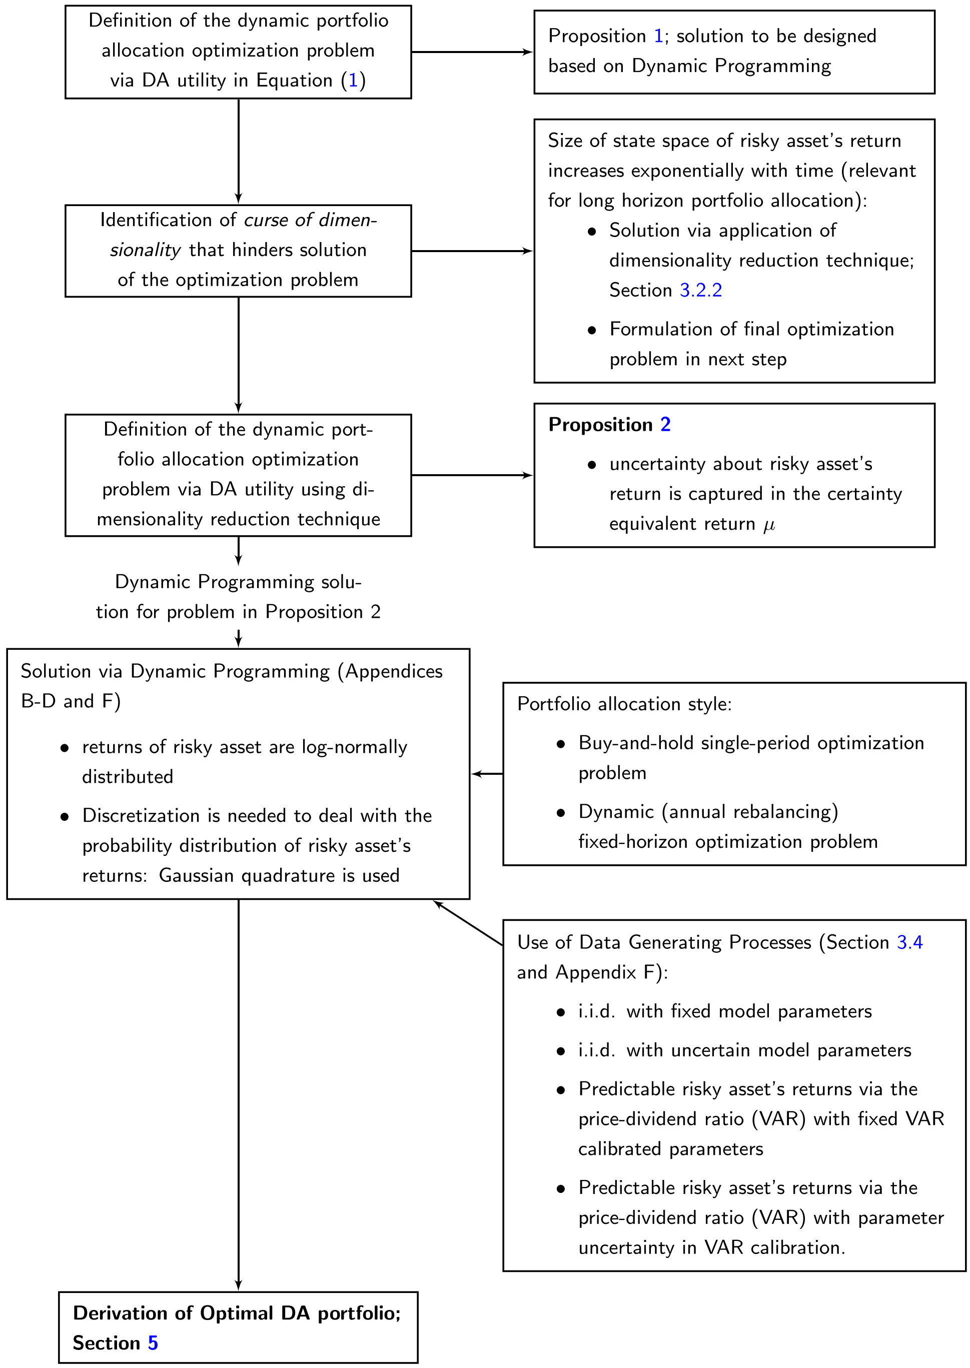

Fig. 1. Flowchart of the solution step for the optimal DA portfolio. Where available and to help readability, references to relevant parts of the paper are provided. For better understanding, the information in the flowchart can be combined with that in Appendix F in the supplementary online material.

in half with every iteration until the optimal value is detected (Sedgewick & Wayne, 2011, for the implementation details of the binary search algorithm, see sections 3.1 and 3.2 in). The reason the binary search algorithm fits in our application, is that in every iteration, that is for every halving of the input space A, the algorithm knows on which 'side' of the state space should look for the solution. The only condition that needs to be in place is that the relation between A and α be strictly monotonic.

5. Asset allocation with DA preference: Simulation study

The theoretical results in Sections 3 and 4 are used in a numerical simulation study split in two parts. The first deals with the remaining analyses for equity non-participation in Section 4 and the second with the main topic of the paper, that is asset allocation in the long-run, derived as the solution to the optimization problem presented in Section 3. Before we discuss the results of the simulation study, we briefly describe the data and the calibration results for the two DGPs, that model either i.i.d. or predictable returns (VAR).

5.1. Data and model calibration

To study the portfolio choice problem, we use quarterly data from the U.S. market from January 1934 to September 2019 for the S&P 500 index (index returns and dividend-price ratios), the risky asset, and the 3-month Treasury bill yield as risk-free asset. As mentioned in Section 3.4.2 the dividend-price ratio (dividend

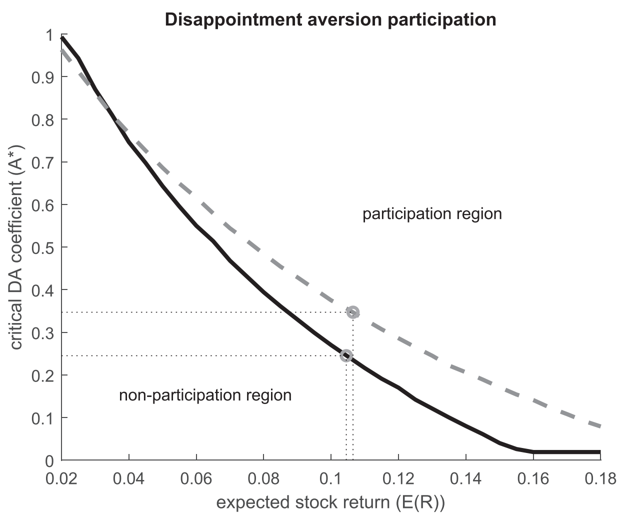

Fig. 2. Stock market participation/nonparticipation regions with DA preferences. The graph shows how the expected level of stock returns (stated annually) affects the critical level of the DA coefficient (A^*). Two lines are presented: the solid one corresponds to the critical DA coefficients for the data set used in our study (1934–2019), and the dashed line plots the critical DA values for the data set used in Ang et al. (2005). The gray circles represent the critical DA level (A^*, which induces non-participation), which corresponds to the historical mean of the equity return for the two data samples.

yield) serves as the predictor variable for the VAR. The data sets related to the S&P 500 returns and the 3-month T-bill rates can be easily acquired by a number of sources as they are readily available online.[11]

5.1.1. Calibration of i.i.d. return generator

Based on the data described above, the following summary statistics were derived:

To link the parameter estimates with the i.i.d. model in Eq. (18), $\mu = 0.02515$, $r = 0.00854$, and $\sigma = 0.08175$, all of which are given in Table 1.

5.1.2. Calibration of the VAR

The MLE of the VAR in Eq. (19) produced the following values:

Simulating asset return trajectories under the assumption that the dividend yield at time t can forecast asset returns at time $t+1$, we match the first two moments of the historical returns' distribution up to two to three significant digits. All coefficients of the matrix with the autoregressive parameters B are statistically significant at the 5% level, and both series (dividend yield and excess asset log returns) are stationary. It cannot but catch one's attention the relatively weak coefficient and statistical insignificance of the explanatory variable. Although a reasonable reaction to this would be to admit that predictability is simply not there, in our view, this effect is intertwined with the presence of parameter uncertainty. Facing uncertainty about the actual calibrated parameters of the VAR cannot but be reflected in the predictive capacity of the model, as – in Bayesian portfolio theory – different asset price realizations would lead to different models of different predictive capacity. As a result, instead of discarding the case of predictability, we opt for studying it in a parameter uncertainty DA setup already discussed in Section 3.

5.2. Equity non-participation in the long-run

We simulate asset return trajectories under the i.i.d. assumption and using the DGP with predictability in a Monte Carlo setup, in order to estimate the excess return and the corresponding return volatility. We examine two relationships: first, how A^* changes with time depending on whether returns are i.i.d. or predictable (Fig. 3), and second, how allocation to the risky asset is affected by the choice of A, in Fig. 4. In the second, investment horizon is fixed at one or ten years. To extract A^* we use the binary search algorithm presented in Section 4 while to determine α at different levels of A the optimization problem introduced in Section 3 is solved.

The left graph of Fig. 3 plots the critical level of the coefficient of DA (A^*) across investment horizons for a buy-and-hold DA investor. The DA coefficient (A^*) is critical, because a decreasing A^* within these setting results in larger market participation, as a lower A^* implies that the investor has to be more disappointment averse to refrain from holding the risky asset. For a longer than a 5-year investment horizon, a DA investor who follows a buy-and-hold strategy will hold risky assets regardless of the DGP assumed for equity returns.

The right graph of Fig. 3 reports critical levels of A^* for dynamic asset allocation strategies for various investment horizons ($T-t$, where t is the current horizon). In the case of i.i.d. returns (dashed line), the critical DA coefficient remains constant regardless of the investment horizon as a result of the invariable opportunity set. The solid line corresponds to predictable returns using the VAR to forecast the next-period's equity return as a function of the dividend price ratio. Contrary to the case of i.i.d. returns, where A^* remains constant, investors' participation increases at longer horizons as a result of the decreasing A^*.

To obtain more insight into non-participation, we additionally plot portfolio allocation to the risky asset with respect to different levels of disappointment aversion. This analysis serves also as an introduction to the simulation study for the optimal portfolio construction discussed in the next section as it tracks all the four different investment profiles for which optimal portfolios will be derived. Moreover, effects observed in Fig. 4 are relevant for the analyses in the following section.

Fig. 4 shows primarily two behaviours. First, that the 10-year line always stands above the 1-year line; this implies that investors become less disappointment averse the longer the horizon they invest for. Second and more interesting, the vertical distance between the two allocation lines (i.e., 1-year vs 10-year) is linked to the size of the horizon effect (hedging demand) generated as a result of the different allocation between short- and long-term investment horizons. Investing for shorter periods typically entails smaller portfolio allocation to the risky asset, a strategy that serves as a hedge against adverse market movements. As discussed in the next section, a buy-and-hold investor who uses the VAR to model asset returns (top row, second plot of Fig. 4) invests very differently at short and long horizons. This results in observing the largest vertical distance between the 1-year and 10-year allocation lines among the four plots presented.

Said result is of importance. It suggests that the DA framework can generate equity non-participation conditions by using a sufficiently low A (or equivalently a large enough A^*). In addition, values of A around 0.5 can result in relatively low portfolio allocations to the risky asset, especially when investment horizon is short. Reasonably low values of DA can thus lead to portfolios with limited risky asset participation, a pattern frequently observed in practice

5.3. Portfolio allocation for buy-and-hold strategies

We first investigate the asset allocation problem at different investment horizons for buy-and-hold strategies. Here, agents choose a portfolio allocation for a specific investment horizon determined at time $t=0$. This strategy results in the same allocation regard-

[11] Our sources are the online platform of Bloomberg Professional Services (for the data on S&P 500 returns) and the Federal Reserve (for the risk free asset).

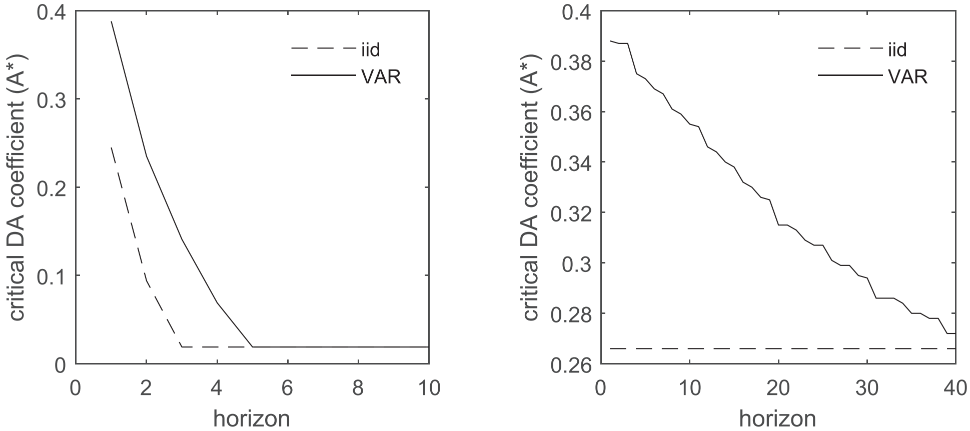

Fig. 3. Critical DA level (A^*) **that induces nonparticipation in the stock market.** The plot on the left refers to a buy-and-hold investor while that on the right to a dynamic investor. The dashed line corresponds to the case of i.i.d. returns (normality and nonpredictability), and the solid line corresponds to the case of predictable returns. Investors invest in the stock market when their DA coefficient lies in the area above the lines. To display the graphs more clearly, the one on the left (buy-and-hold) plots the A^* for a period up to 10 years, as beyond that point A^* remains constant and very close to zero.

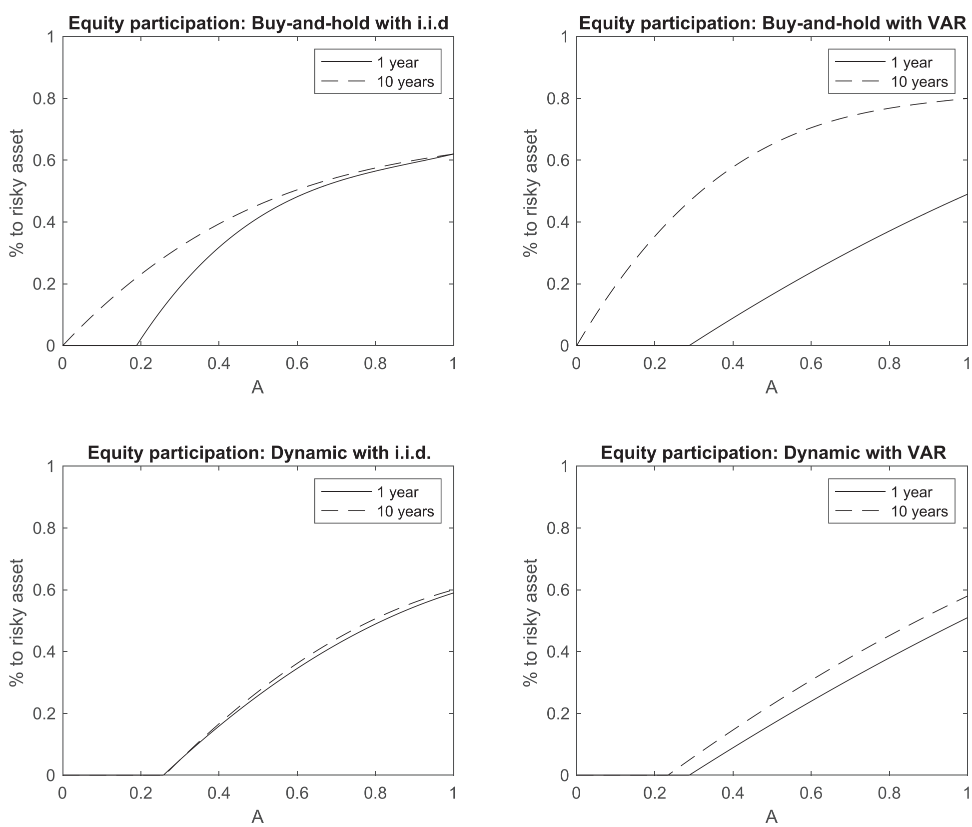

Fig. 4. Participation in the risky asset for different levels of disappointment aversion. The graph plots portfolio allocation to the risky asset for varying levels of disappointment aversion, A for the different cases considered in the study. As per Theorem 1 there's always the marginal level A^* beyond which (i.e., $A < A^*$) no equity participation is predicted by the model.

less of the investment horizon for an investor with power utility when returns follow the i.i.d. process in Eq. (18). Our goal is to explore the effects of the DA utility function in conjunction with parameter uncertainty on the optimal asset allocation. We mainly focus on whether parameter uncertainty in a DA framework induces horizon effects (that is whether long-term allocation to the risky asset is different than short-term allocation. Fig. 5 shows the optimal buy-and-hold portfolio allocations to the risky asset for a DA investor ($A = 0.44$ or $A = 0.30$) and an investor with power utility ($A = 1$; solid line) when returns are i.i.d. and parameter uncertainty about μ and σ is either considered (solid line) or ignored (dashed line). A DA investor who acknowledges parameter uncertainty will decrease her portfolio allocation to the risky asset with the investment horizon compared to the one with the power utility who will hold the same portfolio regardless of the horizon. This comes as the result of the evolution of cumulative return variance

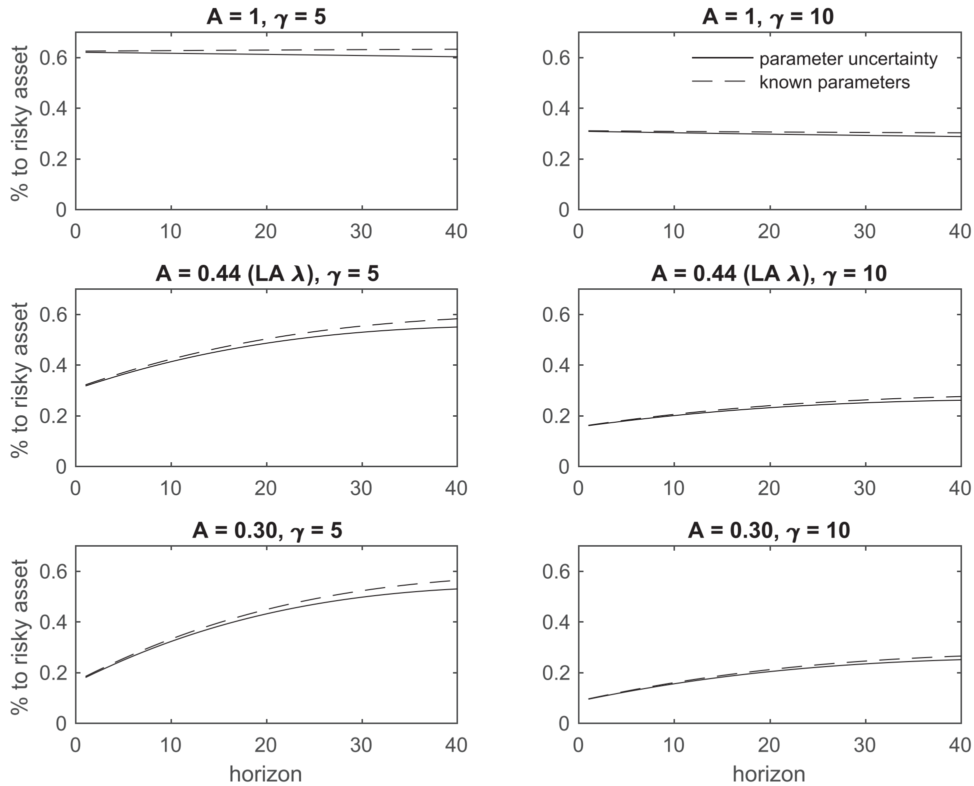

Fig. 5. Optimal portfolio allocation for a buy-and-hold investment strategy. The investor uses the i.i.d. return generator, and either incorporates (optimization problem in 17) or ignores (optimization problem in 2) uncertainty in model parameters. The investor in the top row uses a CRRA (i.e., power) utility function with two levels of risk aversion, whereas the other two cases (middle and bottom rows) make use of the DA utility function with two different values for the DA coefficient. $A = 0.44$ is equivalent to the value of the loss aversion (LA) parameter calculated in Tversky & Kahneman (1992), that is, $DA = 1/\lambda = 0.44$. We observe that a DA investor holds a significantly different portfolio from one who uses a power utility function.

at different horizons. In the following, the key components of our model, namely disappointment aversion, predictability and parameter uncertainty are discussed.

5.3.1. The effect of disappointment aversion

Incorporating DA drastically changes the portfolio composition over different investment horizons compared to the CRRA case ($A = 1$). A DA investor ($A = 0.44$ or $A = 0.30$) will increase her investment proportion to the risky asset when allocating wealth for longer periods. The effect of DA appears to be more powerful at short horizons ($T < 10$), as a DA investor holds significantly less equity compared to one with power utility. For example, Fig. 5 shows that a DA investor invests 20% to 50% of her wealth in the risky asset when her investment horizon is shorter than ten years (between 60% and 20% less equity compared to one with power utility), whereas an even more DA investor ($A = 0.30$) will hold no more than 10% to 40% equity for the same horizon. However, investors with DA utility will allocate similar to those with power utility as the investment horizon increases. A DA investor appears to be very conservative in the short run, whereas, longer investment horizons, even a very DA investor ($A = 0.30$, i.e., for whom losses in her utility function are weighed more than 3 times than gains) is willing to accept the additional risk in anticipation of higher terminal wealth, because of the lower volatility as a result of the longer investment horizon.

5.3.2. Effects of predictability

Predictability is critical in the case of a buy-and-hold investor. Investors who take predictability into account will hold significantly larger weights in equity for longer investment horizons. The reason is that volatility does not grow in proportion to asset returns. This results in lower long-term volatility, compared to the short-term, thus making equities appear more attractive to an investor with a long-term outlook. Fig. 7 displays optimal allocations to the risky asset for three levels of risk aversion (the ones most commonly used in relevant studies) and four levels of DA, among which is the value of $\frac{1}{\lambda}$, where λ is the loss aversion coefficient equal to $\lambda = 2.25$, as calculated in Tversky & Kahneman (1992). As expected, both risk aversion and DA affect the asset allocation to the risky asset as the more risk averse or disappointment averse an investor becomes, the lower the allocation in the risky asset will be.

The horizon effects we report for the buy-and-hold investor who uses the VAR to forecast equity returns can be traced to the evolution of return volatility. Long-term volatility is lower than in the case of i.i.d. returns due to the correlation between the predictor variable and the predicted equity return.[12]

As a result, the long-term volatility for a buy-and-hold investor who uses the VAR is much smaller than that for the investor who uses the i.i.d. return generator, growing slower than linearly. In

[12] More specifically, when we model returns as i.i.d., the two-period variance is equal to

$$var_{r_1,r_2} = var_{r_1} + var_{r_2} \Leftrightarrow \sigma_{r_1,r_2} = \sqrt{var_{r_1} + var_{r_2}}.$$

When returns are predictable, the covariance between equity returns and the predictor variable should be taken into consideration as well. The two-period variance is now equal to

$$var_{r_1,r_2} = var_{r_1} + var_{r_2} + 2cov(r_1, r_2).$$

Given that the covariance term in our VAR estimation is negative (see ρ, σ_{11}, and σ_{22} in the first column of Table 2), the following holds:

$$var_{r_1} + var_{r_2} + 2cov(r_1, r_2) < var_{r_1} + var_{r_2}.$$

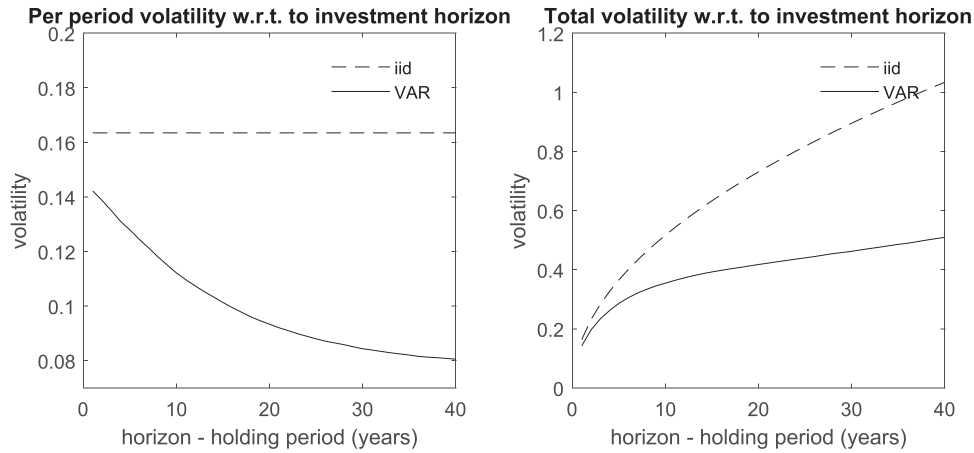

Fig. 6. Evolution of per-period and long-term volatility for the risky asset. The dotted line corresponds to the case of an investor who models returns as i.i.d., whereas the solid line shows the volatility for an investor who uses the VAR to forecast equity returns.

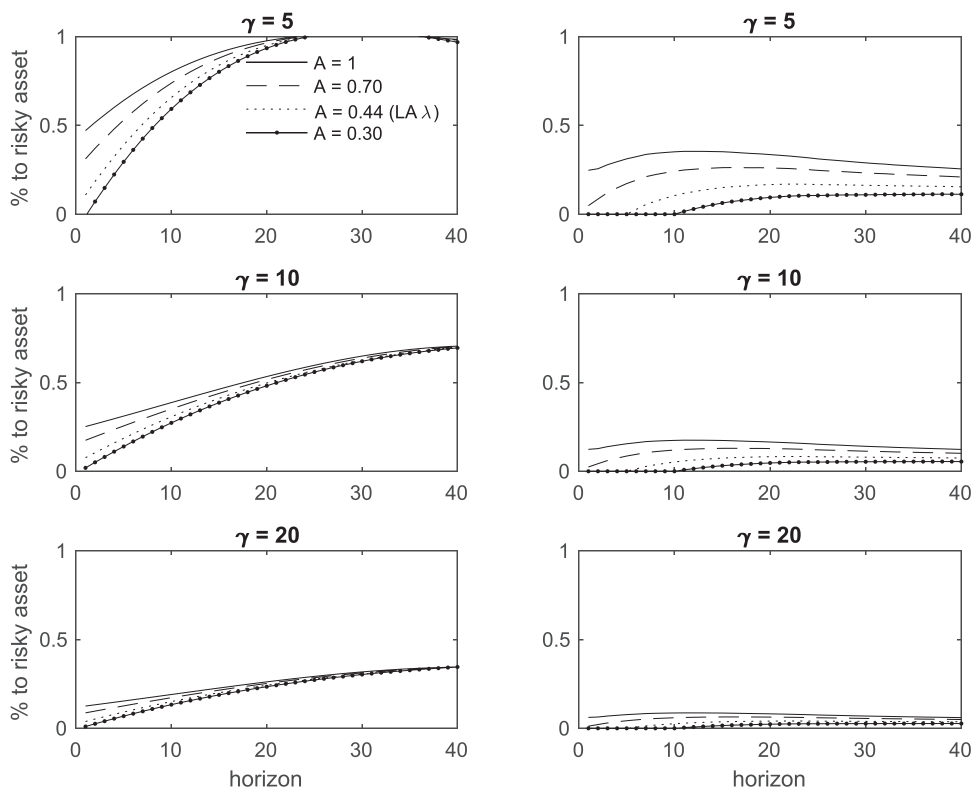

Fig. 7. Optimal portfolio allocation when the VAR is used. The investor follows a buy-and-hold strategy by choosing the portfolio allocation to the risky asset in the beginning of the investment period. $A = 0.44$ is equivalent to the value of the loss aversion (LA) parameter calculated in Tversky & Kahneman (1992), that is, $DA = 1/\lambda = 0.44$. The graphs on the left column ignore parameter uncertainty, whereas those on the right account for this. Three levels of risk aversion and four levels of disappointment aversion are represented. Accordingly, the results are solutions to the optimization problems in 2 and 17 (parameter uncertainty) when the VAR of Eq. (19) is used.

particular, under i.i.d. returns, the 40-year total volatility equals $0.1625\sqrt{40} = 1.02$, compared to 0.5091 when the VAR is used; that is, it is half as much (see Fig. 6). This result shows how the investment allocation in stocks can be affected (i.e., increase) by using a variable believed to predict stock returns.

The intuition behind this effect is twofold. On the one hand, when the dividend yield decreases, the asset price will increase, in effect, disproportionately compared to the dividend yield. This signals that the current price is too high or equivalently that the expected return is too low. The too high current price mean-reverts, resulting in the negative association ($\rho < 0$; see Table 2) between the dividend yield and the future realized return, which, in turn, reduces the rate of increase of the variance, thus rendering equity more attractive at longer horizons. On the other hand, investors relying on a given strategy (in our case, the dividend yield) could develop an illusion of control if they grow overly confident in the strategy's ability to generate precise predictions of future returns. Overreliance is bound to boost investors' overconfidence levels, and lead them to assume higher risk in their investments by increasing their equity exposure (Gervais & Odean, 2001; Odean, 1998). Overconfidence is expected to be further encouraged by the fact that investors whose outlook involves long horizons and/or buy-and-hold strategies monitor their investments less frequently; neglect leads them to experience feelings of regret and/or disappointment equally less frequently, prompting them to view equity as less risky (because longer horizons experience fewer price fluctuations than do shorter ones) and thus tacitly encourage them to increase their exposure to risk (Benartzi & Thaler, 1995).

5.3.3. The effect of parameter uncertainty

For $A = 1$, when parameter uncertainty is ignored, return and variance grow linearly with time and the choice of the investment horizon becomes irrelevant to the solution of the asset allocation problem. Under parameter uncertainty, variance grows faster than linearly and equity is not as attractive as when predictability is ignored. Factoring this understanding in the decision making typically results in lower portfolio allocation to the risky asset. In Lemma 1 we see that the magnitude of the horizon effects depends on the available data incorporated into the model, in the following way: given σ, the variance of μ is inversely proportional to N (the sample size of risky asset return); subsequently, the larger the N (i.e., longer investment horizon), the lower the variance of μ and, equivalently, the smaller the uncertainty around its true value. A shorter investment horizon would result in a significantly lower allocation to the risky asset for an investor who considers uncertainty compared to one who ignores it, especially for longer horizons.

When returns are generated according to the VAR and parameter uncertainty is incorporated (right column of Fig. 7), a DA investor who accounts for predictability will allocate a smaller proportion of her investment to the risky asset compared to an investor who ignores parameter uncertainty. In that case, equities do not appear as attractive as when parameter uncertainty is ignored, because of the higher volatility of equity returns; the latter is due to uncertainty dampening the correlation between the predictor variable and the dependent variable (i.e., equity return), which, in turn, increases the volatility. Expressing uncertainty about the model parameters is, in essence, equivalent to expressing uncertainty about the forecasting capacity of the predictor variable (i.e., the dividend price ratio). This uncertainty, in turn, can prompt investors to start viewing the VAR process as potentially misspecified, thus rendering them more ambiguity averse and leading them to reduce their exposure to equity investments Under parameter uncertainty a DA investor will, in general, still hold larger weights for longer horizons compared to shorter ones, but they will be significantly lower than those allocated when parameter uncertainty is ignored.

5.4. Portfolio allocation for dynamic strategies

We now present the results for the case of a DA investor who follows a dynamic strategy and reallocates her available wealth at the beginning of each period between the risk-free and the risky asset.[13] An investor who dynamically allocates wealth considers the investment opportunity set for the whole investment horizon $T - t$ and assigns the optimal weight to the risky asset knowing that she will have the chance to revise her strategy by the end of the next period in case her expectations of the risky asset's return and volatility change.

5.4.1. The effect of disappointment aversion

The dynamic case entails that investors reallocate their available wealth at the end of each year, considering the optimal solutions from the solved subproblems at each horizon. For the same level of risk aversion, the more disappointment averse an investor grows, the less she allocates to equities. The horizon effect of DA is visible by measuring the equity allocation at a short- and a long-term investment horizon $T - t$. The dynamic allocation to the risky asset drops as the investment horizon becomes shorter as a result of the lower per-period volatility for longer investment horizons shown in Fig. 6. A moderately DA investor will still be heavily invested in equity even at very short horizons (dashed line in Fig. 9), whereas a more DA investor will almost refrain from holding any units of the risky asset even when having a relatively low level of risk aversion.

5.4.2. Effects of predictability

With i.i.d. returns, an investor who dynamically allocates the available wealth at each horizon uses the i.i.d. asset return generator with parameters equal to the historical annual mean and volatility of the S&P 500, seen in Table 1. As expected she has the same investment opportunity set at every horizon, and the allocation to the risky asset does not change at different horizons (dashed line in Fig. 8).

Contrary to i.i.d. returns, the left column of Fig. 9 reports optimal portfolio allocations for four different levels of the DA coefficient A and three levels of the risk aversion coefficient γ at different horizons $T - t$ when the VAR is used. The four levels of DA are the same as those used in the buy-and-hold case. When investors believe returns to be forecastable, they use the VAR to predict next period's equity return and allocation drops for the investment horizon for all four different values of A. As the investment horizon $T - t$ shortens, a DA investor who follows a dynamic strategy allocates a smaller proportion of her wealth to the risky asset, whereas a DA and risk-averse investor will hold no units of the risky asset as $T - t$ approaches zero. Again, dynamically investing in the risky asset in the short run is not as attractive as in the long run given the higher volatility per period of the former. As a consequence, the more disappointment averse an investor is, the more likely she is to be affected by short-run volatility. This gives rise to horizon effects when investors try to hedge their portfolios at shorter horizons.

5.4.3. The effect of parameter uncertainty

Let us assume an investor who uses the i.i.d. return generator and considers uncertainty in parameters. In this case she will exhibit slightly different portfolio allocations compared to when parameters are treated as known. Fig. 8 shows that both a DA investor and one who uses the power utility function will slightly increase their portfolio allocation to the risky asset with the investment horizon (solid line) to eventually hold a portfolio position very similar to an investor who ignores parameter uncertainty (dashed line). Investing for a longer horizon appears to be less risky than holding the risky asset in the short run as a result of the lower per-period volatility of the latter. As a result, an investor who dynamically invests with a shorter-term outlook will hold slightly less equity in their portfolio compared to an investor who invests for a longer horizon.

Turning to the case of predictability, the right column of Fig. 9 reports results that reflect optimal allocations to the risky asset for investors who rebalance their portfolios annually by predicting asset returns based on the dividend yield when parameter uncertainty is accounted for. These plots mainly reveal two facts: first, equity allocation is, in general, lower compared with the case of an investor who ignores parameter uncertainty, and, second, the impact of parameter uncertainty on hedging demands is stronger compared to the case where this is ignored due to the changes in the opportunity set. When we express uncertainty about the parameters of the VAR, we use the posterior predictive distribution in Lemma 2 in place of the VAR model with fixed parameters as stated in Eq. (19). In this case, instead of simulating future return

[13] See appendix for execution times of the algorithm for each case studied in this section.

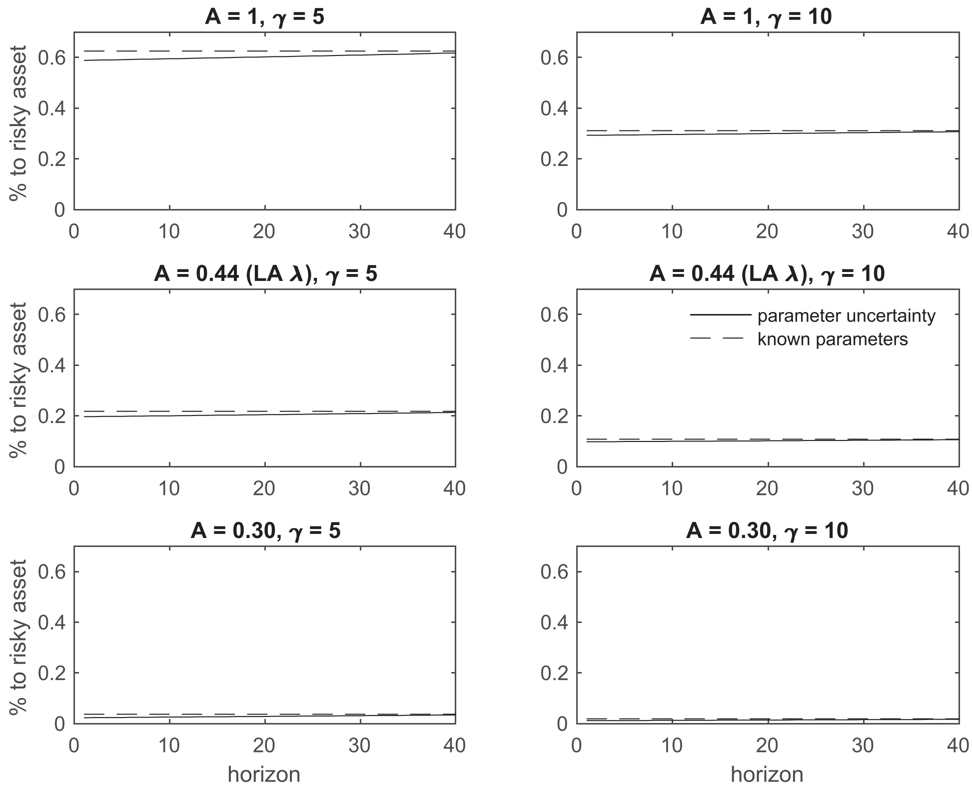

Fig. 8. Dynamic portfolio allocation when returns are i.i.d. Investor chooses between the risky and the riskless asset and uses the i.i.d. return generator for the risky asset. The objective of this exercise is to show how the portfolio allocation to the risky asset changes for an investor who acknowledges parameter uncertainty (optimization problem in 17) compared with one who ignores (optimization problem in 11) it and holds the same portfolio throughout the investment horizon. $A = 0.44$ is equivalent to the value of the loss aversion (LA) parameter calculated in Tversky & Kahneman (1992), that is, $DA = 1/\lambda = 0.44$.

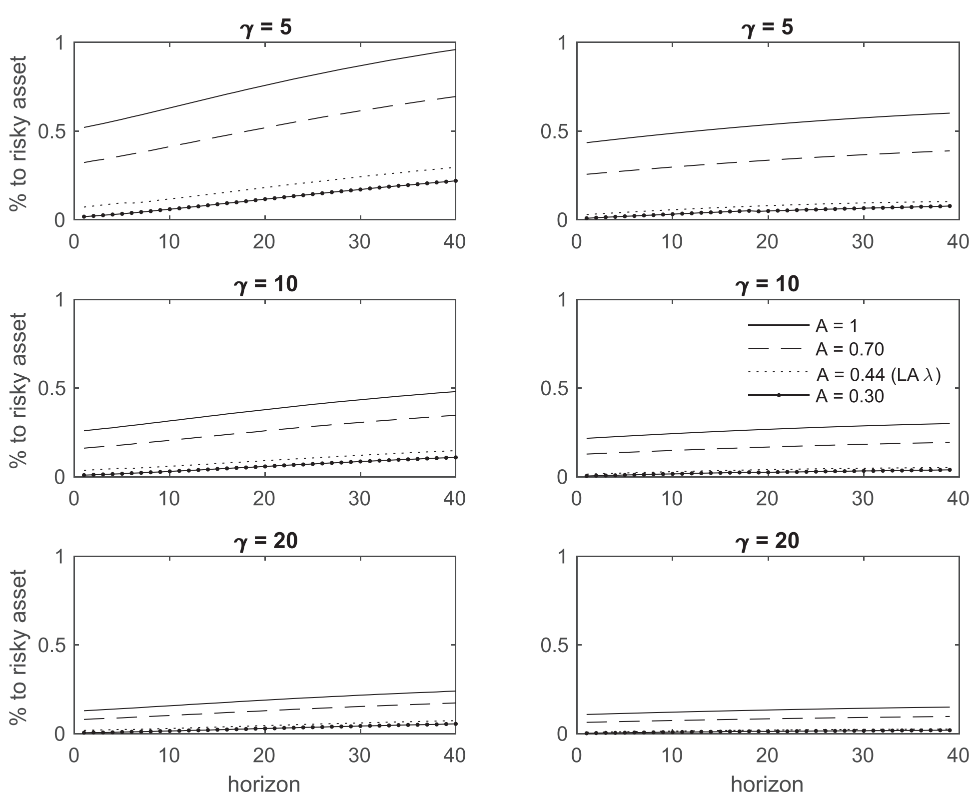

Fig. 9. Optimal portfolio allocation for a dynamic strategy. Investor uses the VAR to forecast returns. The left columns report results when parameter uncertainty is ignored (optimization problem in 11 with predictable returns as in 19), whereas the one on the right accounts for parameter uncertainty (problem in 17). Each line corresponds to a different level of the DA coefficient (A) as follows: solid line, A = 1; dashed line, A = 0.70; dotted line, A = 0.44; solid/dotted line, A = 0.30. $A = 0.44$ is equivalent to the value of the loss aversion (LA) parameter calculated in Tversky & Kahneman (1992), that is, $DA = 1/\lambda = 0.44$.

paths conditioning on fixed values for the model parameters (constant terms, matrix of AR coefficients, and variance-covariance matrix), we sample from their posterior distributions, each time obtaining a new set of parameters that is conditional on observed data only.

The results exhibit a pattern similar to the one in the left column of Fig. 9. The more disappointment averse and risk averse an investor grows, the lower the equity allocation will be at different investment horizons. Still, similar to the case of a DA investor who follows a buy-and-hold strategy, parameter uncertainty mitigates the magnitude of the observed hedging demands. The underlying cause for this behavior can be explained by the way the mean return and variance change over time. Investors' uncertainty about the predictive capacity of the dividend yield results in higher long-term per-period volatility, which explains investors' lower allocation to the risky asset compared to allocation in the left column of Fig. 9, where parameter uncertainty is ignored. Similar to a buy-and-hold investor, considering parameter uncertainty in dynamic investing makes investors aware of changes in the investment opportunity, as a result of the weakened predictive capacity of the dividend price ratio. Subsequently, investors doubt that higher or lower equity allocations will result in more optimal portfolios. In these cases, portfolio allocations will not change significantly among different investment horizons, moderating therefore the observed horizon effects.

6. Concluding remarks

Risk asymmetries implied by disappointment aversion (DA) can decisively change intertemporal portfolio choices, especially when risk asymmetry is studied together with predictability and parameter uncertainty. Our contribution suggests that a DA investor would allocate lower weights to equity compared to an investor who uses a standard CRRA power utility function. What is more, DA introduces horizon effects for a buy-and-hold investor regardless of whether she employs either of the return generators and accounts for or ignores parameter uncertainty. When the latter is considered, equity allocation at longer horizons is significantly lower to when parameter uncertainty is ignored, and there could be cases where parameter uncertainty changes the return distribution by this much that long-term equity allocation is lower to shorter term one.

A dynamic investor who accounts for predictability will hold a completely different portfolio to one who uses an i.i.d. return generator as the distribution of the future returns generated by the VAR is significantly different from that of i.i.d. returns, because of the correlation between the dividend price ratio and the return of the risky asset. As opposed to the latter, a dynamic allocation will exhibit horizon effects as a result of the time-varying investment opportunity set. Furthermore, the incorporation of parameter uncertainty in the DA framework with predictability can drastically change equity allocations over time. Although a portfolio that ignores parameter uncertainty and is based on the i.i.d. return generator will be no materially different to one that accounts for parameter uncertainty, this is no longer when investing under the VAR. There, equity allocation would be much lower for the portfolio that accounts for parameter uncertainty as a result of the additional estimation risk in the VAR model parameters.

Overall, it is beneficial to examine parameter uncertainty as a special case in a portfolio model, as frameworks that do not account for this may generate portfolios with too large equity allocations. When model parameters are taken as uncertain, a DA investor will still allocate larger weights to stocks at longer horizons. Nevertheless, the difference between a long-term and a short-term equity weight is smaller compared to the case in which parameter uncertainty is ignored, and as a result of the doubts investors cast on the predictive power of the dividend yield.

Our results should be of particular interest to policy makers, as they indicate that DA, conditional on its magnitude, tacitly fosters limited-to-no participation in equity investing. To the extent that DA is likely to affect individual investors more (Barber, Odean, & Zhu, 2009, given their lower sophistication levels,), financial literacy programs could raise awareness of DA, while training people to assess their investments from a longer-term perspective, regardless of price movements in the short run (where the effects of DA are more likely to be felt). This, in turn, will help enhance the participation of retail investors in equity turnover (thus benefiting market liquidity), while ensuring that those that invest in equities are less likely to exit the market because of disappointment-related reasons. Our results are also relevant to finance practitioners, in particular brokers and financial advisors, who, by virtue of their profession, tend to engage with retail investors on a regular basis. For these practitioners, accounting for DA in their clients' risk profiling and overall day-to-day interactions would considerably help inform their professional practice, by permitting practitioners additional insight into their clients' trading decisions. Such insight could allow them to educate their clients about the role of DA in trading, thus helping them potentially improve their trading decisions. From an academic perspective, and to the extent that disappointment stems from prior investment experience, our results also offer an alternative explanation of previously documented evidence (Seru, Shumway, & Stoffman, 2010; Strahilevitz, Odean, & Barber, 2011) of the reluctance of investors to reenter the market if they have exited it previously at a loss.

There are a number of interesting avenues for future research. Exploration of behavioral utility functions, particularly the DA theory, in the context of an asset allocation optimization problem is far from complete. In practice, investors' portfolios contain riskless and numerous asset classes of risky assets. Hence, a natural extension of the current framework seems to be building one able to deal with multi-asset portfolios, with additional asset classes, and incorporating consumption and trading costs. Moreover, given that prices of financial assets are determined by the forces of supply and demand, which stem from the trading and investment decisions of market participants, it would be of great interest to depart from partial equilibrium and study this interplay in a general equilibrium context. Doing so would allow us to investigate the interaction between investors with CRRA preferences and those using a DA utility function and, eventually, to reveal the role that DA plays in both defining asset prices and determining trading behaviors. Finally, another important direction for future exploitation is to focus on portfolio and consumption problems, rather than just a maximization of terminal wealth as it is performed in Ang et al. (2005) and our study. Thus, like for the case of CRRA utility, portfolio choice under DA might also result in dependence of the portfolio decision on future consumption-wealth ratios.

Acknowledgments

We are grateful to Emanuele Borgonovo and three anonymous referees whose comments improved the paper significantly. The authors would like to thank Nicholas Barberis, Stephen J. Brown, Lauren H. Cohen, Victor DeMiguel, David Hirshleifer, Steven Kou, Roy Kouwenberg, Jun Liu, David H. Solomon, Athanasios Yannacopoulos, Stavros Zenios and participants at Quantitative Finance and Risk Analysis (QFRA2017) symposium, Athens University of Economics and Business, University of Patras, Bayes (former Cass) Business School, University of Aalborg, University of Liverpool, University of Loughborough, Monash Business School, Shanghai University, Southwestern University of Finance and Economics and

Sungkyunkwan University for helpful comments. Any remaining errors are ours.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.ejor.2023.07.038.

References

Abdellaoui, M., & Bleichrodt, H. (2007). Eliciting Gul's theory of disappointment aversion by the tradeoff method. *Journal of Economic Psychology, 28*(6), 631–645.
Aït-Sahalia, Y., Cacho-Diaz, J., & Hurd, T. R. (2009). Portfolio choice with jumps: A closed-form solution. *Annals of Applied Probability, 19*(2), 556–584.
Ang, A., & Bekaert, G. (2007). Stock return predictability: Is it there? *Review of Financial Studies, 20*(3), 651–707.
Ang, A., Bekaert, G., & Liu, J. (2005). Why stocks may disappoint. *Journal of Financial Economics, 76*(3), 471–508.
Avramov, D., & Zhou, G. (2010). Bayesian portfolio analysis. *Annual Reviews of Financial Economics, 2*(1), 25–47.
Baillon, A., Bleichrodt, H., & Spinu, V. (2020). Searching for the reference point. *Management Science, 66*(1), 93–112.
Balduzzi, P., & Lynch, A. W. (1999). Transaction costs and predictability: Some utility cost calculations. *Journal of Financial Economics, 52*(1), 47–78.
Barber, B. M., Odean, T., & Zhu, N. (2009). Systematic noise. *Journal of Financial Markets, 12*(4), 547–569.
Barberis, N. (2000). Investing for the long run when returns are predictable. *Journal of Finance, 55*(1), 225–264.
Barberis, N., & Huang, M. (2008). Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review, 98*(5), 2066–2100.
Barberis, N., Huang, M., & Santos, T. (2001). Prospect theory and asset prices. *The Quarterly Journal of Economics, 116*(1), 1–53.
Bawa, V. S., Brown, S. J., & Klein, R. W. (1979). *Estimation risk and optimal portfolio choice* p. 190. North-Holland Publishing Co., N.Y..
Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research, 30*(5), 961–981.
Bell, D. E. (1985). Disappointment in decision making under uncertainty. *Operations Research, 33*(1), 1–27.
Bellemare, C., Kröger, S., & Sossou, K. M. (2020). Optimal frequency of portfolio evaluation in a choice experiment with ambiguity and loss aversion. *Journal of Econometrics*.
Benartzi, S., & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *The Quarterly Journal of Economics, 110*(1), 73–92.
Berkelaar, A. B., Kouwenberg, R., & Post, T. (2004). Optimal portfolio choice under loss aversion. *Review of Economics and Statistics, 86*(4), 973–987.
Bernard, C., & Ghossoub, M. (2010). Static portfolio choice under cumulative prospect theory. *Mathematics and Financial Economics, 2*(4), 277–306.
Bertsekas, D. P. (1995). *Dynamic programming and optimal control*: Vol. 1. Athena Scientific Belmont, MA USA.
Bonomo, M., Garcia, R., Meddahi, N., & Tédongap, R. (2011). Generalized disappointment aversion, long-run volatility risk, and asset prices. *Review of Financial Studies, 24*(1), 82–122.
Boyd, M. (2001). On ignorance, intuition, and investing: A bear market test of the recognition heuristic. *The Journal of Psychology and Financial Markets, 2*(3), 150–156.
Brandt, M. W., Goyal, A., Santa-Clara, P., & Stroud, J. R. (2005). A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Review of Financial Studies, 18*(3), 831–873.
Branger, N., Larsen, L. S., & Munk, C. (2013). Robust portfolio choice with ambiguity and learning about return predictability. *Journal of Banking & Finance, 37*(5), 1397–1411.
Brown, S. J. (1978). The portfolio choice problem: Comparison of certainty equivalence and optimal Bayes portfolios. *Communications in Statistics-Simulation and Computation, 7*(4), 321–334.
Campbell, J. Y., & Viceira, L. M. (1999). Consumption and portfolio decisions when expected returns are time varying. *The Quarterly Journal of Economics, 114*(2), 433–495.
Campbell, J. Y., & Viceira, L. M. (2002). *Strategic asset allocation: Portfolio choice for long-term investors*. USA: Oxford University Press.
Campbell, J. Y., & Yogo, M. (2006). Efficient tests of stock return predictability. *Journal of Financial Economics, 81*(1), 27–60.
Chen, H., Ju, N., & Miao, J. (2014). Dynamic asset allocation with ambiguous return predictability. *Review of Economic Dynamics, 17*(4), 799–823.
Cochrane, J. H. (2008). The dog that did not bark: A defense of return predictability. *Review of Financial Studies, 21*(4), 1533–1575.
Dahlquist, M., Farago, A., & Tédongap, R. (2017). Asymmetries and portfolio choice. *Review of Financial Studies, 30*(2), 667–702.
Davis, P. J., & Rabinowitz, P. (2007). *Methods of numerical integration*. Courier Corporation.
DeMiguel, V., Martín-Utrera, A., & Nogales, F. J. (2015). Parameter uncertainty in multiperiod portfolio optimization with transaction costs. *Journal of Financial and Quantitative Analysis, 50*(06), 1443–1471.
Diecidue, E., Schmidt, U., & Zank, H. (2009). Parametric weighting functions. *Journal of Economic Theory, 144*(3), 1102–1118.
Dimmock, S. G., & Kouwenberg, R. (2010). Loss-aversion and household portfolio choice. *Journal of Empirical Finance, 17*(3), 441–459.
Epstein, L. G., & Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica: Journal of the Econometric Society, 57*(4), 937–969.
Fielding, D., & Stracca, L. (2007). Myopic loss aversion, disappointment aversion, and the equity premium puzzle. *Journal of Economic Behavior & Organization, 64*(2), 250–268.
Garlappi, L., Uppal, R., & Wang, T. (2007). Portfolio selection with parameter and model uncertainty: A multi-prior approach. *The Review of Financial Studies, 20*(1), 41–81.
Gervais, S., & Odean, T. (2001). Learning to be overconfident. *Review of Financial Studies, 14*(1), 1–27.
Gomes, F. J. (2005). Portfolio choice and trading volume with loss-averse investors. *Journal of Business, 78*(2), 675–706.
Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology, 38*(1), 129–166.
Grinblatt, M., & Keloharju, M. (2001). How distance, language, and culture influence stockholdings and trades. *Journal of Finance, 56*(3), 1053–1073.
Gul, F. (1991). A theory of disappointment aversion. *Econometrica: Journal of the Econometric Society, 59*(3), 667–686.
Haigh, M. S., & List, J. A. (2005). Do professional traders exhibit myopic loss aversion? An experimental analysis. *Journal of Finance, 60*(1), 523–534.
Hirshleifer, D. (2008). Psychological bias as a driver of financial regulation. *European Financial Management, 14*(5), 856–874.
Hoevenaars, R. P. P. M., Molenaar, R. D. J., Schotman, P. C., & Steenkamp, T. (2014). Strategic asset allocation for long-term investors: Parameter uncertainty and prior information. *Journal of Applied Econometrics, 29*(3), 353–376.
Huberman, G. (2001). Familiarity breeds investment. *Review of Financial Studies, 14*(3), 659–680.
Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance, 48*(1), 65–91.
Jegadeesh, N., & Titman, S. (2001). Profitability of momentum strategies: An evaluation of alternative explanations. *Journal of Finance, 56*(2), 699–720.
Jin, L., & Scherbina, A. (2010). Inheriting losers. *Review of Financial Studies, 24*(3), 786–820.
Johannes, M., Korteweg, A., & Polson, N. (2014). Sequential learning, predictability, and optimal portfolio returns. *Journal of Finance, 69*(2), 611–644.
Kacperczyk, M., & Damien, P. (2011). Asset allocation under distribution uncertainty. *McCombs Research Paper Series No. IROM-01-11*.
Kahneman, D., Lovallo, D., & Sibony, O. (2011). Before you make that big decision. *Harvard Business Review, 89*(6), 50–60.
Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society, 47*(2), 263–291.
Kandel, S., & Stambaugh, R. F. (1996). On the predictability of stock returns: An asset-allocation perspective. *Journal of Finance, 51*(2), 385–424.
Kilka, M., & Weber, M. (2001). What determines the shape of the probability weighting function under uncertainty? *Management Science, 47*(12), 1712–1726.
Kourtis, A., Dotsis, G., & Markellos, R. N. (2012). Parameter uncertainty in portfolio selection: Shrinking the inverse covariance matrix. *Journal of Banking & Finance, 36*(9), 2522–2531.
Kreps, D. M., & Porteus, E. L. (1979). Dynamic choice theory and dynamic programming. *Econometrica: Journal of the Econometric Society, 47*(1), 91–100.
Lettau, M., & Ludvigson, S. (2001). Consumption, aggregate wealth, and expected stock returns. *Journal of Finance, 56*(3), 815–849.
Locke, P. R., & Mann, S. C. (2001). House money and overconfidence on the trading floor. In *AFA 2002 atlanta meetings*.
Lynch, A. W. (2000). Portfolio choice with many risky assets, market clearing and cash flow predictability,.
Massa, M., & Simonov, A. (2006). Hedging, familiarity and portfolio choice. *Review of Financial Studies, 19*(2), 633–685.
Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance, 53*(5), 1775–1798.
Pástor, L. (2000). Portfolio selection and asset pricing models. *The Journal of Finance, 55*(1), 179–223.
Pástor, L., & Stambaugh, R. F. (2000). Comparing asset pricing models: An investment perspective. *Journal of Financial Economics, 56*(3), 335–381.
Rabin, M. (2013). Risk aversion and expected-utility theory: A calibration theorem. In *Handbook of The fundamentals of financial decision making: Part I* (pp. 241–252). World Scientific.
Ritov, I., & Baron, J. (1999). Protected values and omission bias. *Organizational Behavior and Human Decision Processes, 79*(2), 79–94.
Routledge, B. R., & Zin, S. E. (2010). Generalized disappointment aversion and asset prices. *Journal of Finance, 65*(4), 1303–1332.
Schimmack, U., & Diener, E. (1997). Affect intensity: Separating intensity and frequency in repeatedly measured affect. *Journal of Personality and Social Psychology, 73*(6), 1313–1329.
Schmidt, U., & Zank, H. (2005). What is loss aversion? *Journal of Risk and Uncertainty, 30*(2), 157–167.
Sedgewick, R., & Wayne, K. (2011). *Algorithms*. Addison-Wesley Professional.
Seru, A., Shumway, T., & Stoffman, N. (2010). Learning by trading. *Review of Financial Studies, 23*(2), 705–739.
Shapira, Z., & Venezia, I. (2001). Patterns of behavior of professionally managed and independent investors. *Journal of Banking & Finance, 25*(8), 1573–1587.

Strahilevitz, M. A., Odean, T., & Barber, B. M. (2011). Once burned, twice shy: How naïve learning, counterfactuals, and regret affect the repurchase of stocks previously sold. *Journal of Marketing Research, 48*(SPL), S102–S120.

Summers, B., & Duxbury, D. (2012). Decision-dependent emotions and behavioral anomalies. *Organizational Behavior and Human Decision Processes, 118*(2), 226–238.

Thaler, R. H., & Sunstein, C. H. (2008). Nudge: Improving decisions about health, wealth, and happiness.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5*(4), 297–323.

Wang, Z. (2005). A shrinkage approach to model uncertainty and asset allocation. *Review of Financial Studies*, 673–705.

Weiner, B., Russell, D., & Lerman, D. (1979). The cognition–emotion process in achievement-related contexts. *Journal of Personality and Social Psychology, 37*(7), 1211–1220.

Wermers, R. (2003). Is money really ‘smart’? New evidence on the relation between mutual fund flows, manager behavior, and performance persistence,.

Werner, K. M., & Zank, H. (2019). A revealed reference point for prospect theory. *Economic Theory, 67*(4), 731–773.

Xie, Y., Hwang, S., & Pantelous, A. A. (2018). Loss aversion around the world: Empirical evidence from pension funds. *Journal of Banking & Finance, 88*, 52–62.