Network for Studies on Pensions, Aging and Retirement

NETSPAR ACADEMIC SERIES

Time-consistency of optimal investment under smooth ambiguity

Anne Balter, Antje Mahayni, Nikolaus Schweizer

DP 01/2021-046

European Journal of Operational Research 293 (2021) 643–657

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Time-consistency of optimal investment under smooth ambiguity

Anne G. Balter[a,b], Antje Mahayni[c], Nikolaus Schweizer[a,*]

[a] Department of Econometrics & Operations Research, Tilburg University, Warandelaan 2, Tilburg 5037AB, the Netherlands
[b] Netspar, the Netherlands
[c] Mercator School of Management, Universität Duisburg-Essen, Lotharstr. 65, Duisburg D-47057, Germany

ARTICLE INFO

Article history:
Received 18 March 2020
Accepted 23 December 2020
Available online 5 January 2021

Keywords:
Finance
Optimal investment
Smooth ambiguity
Time-consistency
Robustness,

ABSTRACT

We study portfolio choice in a Black–Scholes world under drift uncertainty. Preferences towards risk and ambiguity are modeled using the smooth ambiguity approach under a double power utility assumption and a normal distribution assumption on the unknown drift. Optimal investment in this setting is time-inconsistent. Utility is maximized by a time-inconsistent pre-commitment strategy resembling the classical Merton solution. In contrast, the optimal dynamically consistent investment strategy accounts for variations in the perceived severity of drift uncertainty, increasing the riskiness of the strategy gradually over time. We provide a detailed comparative analysis of the mechanics and interplay of ambiguity, myopia and optimal decisions in this setting. We show that an investor who pre-commits will regret that decision from some time point onwards, wishing that she had followed the dynamically consistent strategy.

1. Introduction

We study the effects of ambiguity on the riskiness of investment decisions. Since Ellsberg's (1961) examples, there is a growing awareness that decision makers may have a non-neutral attitude towards ambiguity. Thus, preferences are decomposed into risk preferences (based on known probabilities) and preferences concerning the degree of uncertainty[1] about the (unknown) probabilities. In view of the uncertainty about the data generating processes that drive asset prices, ambiguity and ambiguity aversion are especially important in the context of portfolio planning. A recent literature suggests that ambiguity is at least as prominent as risk in making investment decisions, see Chen and Epstein (2002). In our analysis, we consider an investor who can invest in a risky asset and a risk–free asset. She ranks her investment opportunities in terms of her terminal wealth, i.e. her wealth at the end of her investment horizon. The investor believes in a diffusion model. She is confident about her knowledge of the diffusion coefficient but is uncertain about the asset price drift (i.e. excess return).

More precisely, our investor believes that the asset follows a geometric Brownian motion with constant drift as prescribed by the Black–Scholes model. Yet she acknowledges her uncertainty about the true drift by assuming a normally distributed drift where the expectation coincides with her best point estimate of the drift.

We model the investor's ambiguity aversion, i.e., the impact of model uncertainty on her preferences, by applying the smooth ambiguity model of Klibanoff, Marinacci, and Mukerji (2005).[2] This approach allows for a separation of risk and ambiguity. The impact of risk on decision making is described by a risk situation (a distribution) and a utility function characterizing the attitude towards risk. In our context, the risk situation is captured by the distribution of the terminal portfolio wealth under a fixed realization of the drift. A similar procedure now concerns the ambiguity situation and the ambiguity attitude. The ambiguity situation is described by the distribution the investor assumes for the drift parameter. In analogy to a utility function, an ambiguity function is introduced to capture the attitude towards ambiguity. Technically, this results in a double expectation. The inner expectation measures the expected utility with respect to each admissible risk model. The outer expectation accounts for the ambiguity with respect to the risk models. In particular, for vanishing drift ambiguity, i.e., in the Black–Scholes limit, the investor simply maximizes her expected utility of terminal wealth.

* Corresponding author.
E-mail addresses: a.g.balter@uvt.nl (A.G. Balter), antje.mahayni@uni-due.de (A. Mahayni), n.f.f.schweizer@uvt.nl (N. Schweizer).

[1] We use "ambiguity" and "uncertainty" interchangeably to describe situations with uncertain probabilities.

[2] For a review of decision making models under ambiguity we refer to Epstein and Schneider (2010).

https://doi.org/10.1016/j.ejor.2020.12.046

In order to derive tractable and easy to interpret solutions we make two simplifying assumptions which are discussed in more detail later on. First, we assume that both the risk aversion and ambiguity aversion of the investor are described by power functions. Second, for the main analysis we abstract from learning and restrict the investor to strategies which are deterministic functions of time. As a first result, we find that the optimal investment fraction is constant like the Merton fraction in the Black–Scholes setup that serves as our benchmark. Ambiguity leads to a reduction of the investment fraction and this reduction is more pronounced for longer time horizons. Consequently, when the investment horizon shrinks as time goes by, the investor wishes to revise her "optimal" investment decision. The solution is not time-consistent. Following Strotz (1955) who argues that an economic agent *should choose… the best plan among those that [s]he will actually follow*, we thus derive a closed-form expression for the optimal dynamically consistent investment strategy under ambiguity. Technically, the optimal dynamically consistent strategy is derived as the solution of the dynamic programming equation associated with the investment problem – even though this solution does not correspond to the optimal strategy derived before. The optimal solution can only be implemented if the investor finds a way to pre-commit on her investment strategy, e.g., by entering a non-negotiable long-term contract with an intermediary who invests on her behalf. We analyze the losses that occur when the investor cannot pre-commit and thus chooses the optimal dynamically consistent strategy.

On shorter investment horizons, the optimal pre-commitment strategy is more aggressive and thus closer to the Merton fraction. Intuitively, drift uncertainty is less relevant on shorter time spans. This also explains why under the dynamically consistent strategy, optimal investment is less aggressive in the beginning as the agent anticipates that she will want to invest more aggressively later on. We compare this dynamically consistent behavior to investments by a myopic investor who "restarts" the pre-commitment strategy at each instance. The myopic investor is thus ignorant to the future development of her perceived ambiguity.

The myopic strategy and the dynamically consistent strategy differ from pre-commitment in two crucial ways. Investment becomes more aggressive over time. Moreover, on average, investment is more aggressive than under the time-inconsistent pre-commitment strategy. Compared to pre-commitment, the strategies are thus suboptimal both in terms of overinvestment and in terms of a suboptimal investment schedule. We show that, essentially, the suboptimality of the myopic strategy is dominated by the overinvestment effect. In contrast, deviations between the utility from pre-commitment and the utility from dynamically consistent investment are largely due to suboptimal scheduling. Finally, we show that an investor who follows the pre-commitment strategy will always regret that decision from some time point onwards, wishing that she had followed the dynamically consistent strategy.

To understand the time-inconsistency problems in our stylized model setup better, we place them in a broader context. First, we show that these problems persist if we allow for stochastic rather than only deterministic strategies. Then, we demonstrate that time-inconsistency persists even under Bayesian learning if agents are ambiguity-averse (or ambiguity-loving). This stands in contrast to the fact that time-inconsistency does not arise for ambiguity-neutral agents in this setting (Rogers, 2001). We also demonstrate that the quantitative impact of incorporating learning into the model is small, further justifying our focus on a model without learning in the main analysis.[3] Finally, we discuss further possibilities for mitigating time-inconsistency from the literature, arguing that time-consistency is a knife-edge case that arises only when the relevant quantities cancel out suitably.

Literature review Model uncertainty has been investigated for the last 30 years in various fields under different names such as model risk, model misspecification, ambiguity and robustness. For pointers to the relevant economics literatures see Gilboa and Schmeidler (1989) and Hansen and Sargent (2008) while for the operations research literature overviews are given in Ben-Tal, El Ghaoui, and Nemirovski (2009) and Bertsimas, Brown, and Caramanis (2011). For an overview of the finance and economics literature on ambiguity in asset pricing and portfolio choice we refer to Guidolin and Rinaldi (2013) and Hansen and Marinacci (2016). The operations research literature on robust portfolio choice is surveyed by Kim, Kim, and Fabozzi (2014), Fabozzi, Huang, and Zhou (2010) and Pflug and Pohl (2018).

Unlike most of the literature we do not rely on a worst case approach to model ambiguity but rather on the utility-based smooth ambiguity model of Klibanoff et al. (2005) which assumes a probability distribution over possible models. Our dynamic utility-based approach thus stands in contrast to the static worst case, generalized mean-variance (i.e. risk-reward) frameworks studied in most of the literature on robust portfolio choice, e.g. Tütüncü and Koenig (2004). Some dynamic exceptions include Dantzig and Infanger (1993), Ben-Tal, Margalit, and Nemirovski (2000), Gülpınar and Rustem (2007) and Bertsimas and Pachamanova (2008). In the finance literature on portfolio choice under model uncertainty, dynamic frameworks have been more common. The case of a single risky asset (like in this paper) is considered e.g. by Maenhout (2004), and Biagini and Pınar (2017). For the multiple risky assets, we refer to Chen and Epstein (2002), Epstein and Miao (2003), Uppal and Wang (2003), Boyle, Garlappi, Uppal, and Wang (2012) and Raponi, Uppal, and Zaffaroni (2018). For instance, Boyle et al. (2012) relate ambiguity to the home bias towards familiar investment opportunities while Raponi et al. (2018) study how different types of model uncertainty behave in the limit as the number of possible assets goes to infinity.

A smaller number of contributions have studied portfolio planning under smooth ambiguity. Taboga (2005) focuses on a static mean variance setup but analyzes the joint uncertainty about means and variances. Gollier (2011) considers a static two-asset portfolio problem with one safe asset and one uncertain asset. Collard, Mukerji, Sheppard, and Tallon (2018) assess the quantitative impact of ambiguity on the historically observed equity premium. Ju and Miao (2012) propose a generalized recursive smooth ambiguity model and give applications to a consumption-based asset-pricing model. Suzuki (2018) generalizes their analysis from discrete to continuous time. In his setting, the instantaneous rates of returns are newly drawn every period. This makes it non-trivial to model ambiguity preferences in such a way that the impact of ambiguity is still seen in the continuous time limit, see also Skiadas (2013). This sharp distinction between continuous and discrete time does not apply in our more Black–Scholes-type setting where the unknown, constant rate of return is drawn only once in the beginning. The analogue of the classic Arrow-Pratt approximation of the certainty equivalent under model uncertainty is given in Maccheroni, Marinacci, and Ruffino (2013). Chen, Ju, and Miao (2014) apply, like us, a (doubly power function) smooth ambiguity model setup yet in discrete-time. They consider an optimal consumption and portfolio choice problem where the investor is confronted with two possibly misspecified models of stock returns, one model with i.i.d. returns and the other with return predictability.

To our knowledge, the discrepancy between pre-commitment and time-consistent optimal strategies has not been analyzed quantitatively in the literature on investment under ambiguity before. In this regard, our paper is closely related to Basak and

[3] This is in line with earlier results in the literature, see, e.g., Branger, Larsen, and Munk (2013) who find that the impact of learning is small if ambiguity aversion is substantial.

Chabakauri (2010) who study dynamic portfolio choice under mean-variance preferences. In the mean-variance setting, optimal investment is time-inconsistent even in the absence of ambiguity, leading to a similar distinction between pre-commitment, dynamically consistent and myopic strategies, see also Cong and Oosterlee (2016) and Pedersen and Peskir (2017). Following Basak and Chabakauri (2010), a mathematically rigorous game-theoretic theory of dynamically consistent behavior in time-inconsistent stochastic control problems was developed in Björk and Murgoci (2014) and Björk, Khapko, and Murgoci (2017). Within that stream of the literature, Pun (2018) is particularly related to our work as he adds model uncertainty to their setting and studies optimal portfolio choice with a mean-variance objective as an application. However, in his mean-variance setting the source of the time-inconsistency lies in the mean-variance preferences and not in the ambiguity which is modeled recursively like in Maenhout (2004). For a mean-variance objective, comparisons of investment strategies with and without pre-commitment are found, e.g., in Wang and Forsyth (2011) and in the recent papers Vigna (2020) and Van Staden, Dang, and Forsyth (2021). Wang and Forsyth (2011) develop a numerical approach for comparing the strategies under more realistic constraints than in Basak and Chabakauri (2010). Vigna (2020) characterizes, like us, a time point at which regret sets in. She shows that, on average from a time zero perspective, a mean-variance agent who cannot pre-commit and follows the dynamically consistent strategy will regret this decision from some point on, wishing she had followed the myopic strategy. Van Staden et al. (2021) focus on robustness to model uncertainty in the sense of misspecified parameter constellations. There is also a literature which studies how behavioral biases and bounded rationality lead to time-inconsistency in optimal investment, see e.g. Strub and Li (2020) and the references therein.

Outline The paper is organized as follows. In Section 2, we introduce the setting. Section 3 derives the optimal pre-commitment and dynamically consistent strategies and associated excess growth rates for the optimization problem under smooth ambiguity. In Section 4, we analyze and quantify the loss that arises without the possibility to pre-commit. In Section 5, we study the impact of learning on our analysis, showing that time-inconsistency remains a problem in general. Finally, in Section 6, we discuss our results in the context of alternative models of ambiguity. Section 7 concludes. All proofs are in the appendix.

2. The setting

In this section, we introduce our model of an investor's ambiguity towards the risks in a prospective investment. To capture both, ambiguity and risk, we rely on the smooth ambiguity framework of Klibanoff et al. (2005) where a class of stochastic models instead of just a single model is considered. Risk is modeled by the stochasticity within each of the models. Ambiguity is modeled by an additional layer of randomness concerning which of the models is the correct one.

There is a risky asset[4] with price process $(S_t)_t$ over a time horizon $[0, T]$ given by

$$S_t = S_0 e^{\sigma W_t - \frac{\sigma^2}{2} t + t\mu} \tag{1}$$

for positive constants σ and S_0 and a standard Brownian motion W_t. Thus, S_t is a geometric Brownian motion. In our model, risk is captured by the randomness of W_t. Throughout the paper, we assume that the volatility σ is a known constant. We introduce ambiguity, or model uncertainty, in the form of uncertainty about the drift μ. This distinction is in line with the well-known fact that volatilities are much easier to estimate than drifts, see Merton (1980).[5]

We take into account uncertainty about the drift by assuming that μ is random with known distribution. In particular, we assume that μ is normally distributed with known mean $\bar{\mu} > 0$ and variance $\bar{\sigma}^2$, $\mu \sim \mathcal{N}(\bar{\mu}, \bar{\sigma}^2)$ and independent of $(W_t)_t$. In the special case $\bar{\sigma}^2 = 0$, our model collapses to a standard Black–Scholes model with drift $\bar{\mu}$. We refer to this case as the *Black–Scholes benchmark*.

The investor in our model thus believes that the asset price follows a geometric Brownian motion with constant drift as prescribed by the Black–Scholes model. Yet she acknowledges that $\bar{\mu}$ is only the most likely value of μ but not necessarily the correct one. Besides its tractability, the assumption of a normal distribution for μ can be justified, e.g., by assuming that $\bar{\mu}$ is a drift estimate obtained from historical data and the normal distribution arises from an application of the central limit theorem. In particular, it is well-known that very long time series of historical data are necessary for reliably estimating the drift. To some extent, this justifies neglecting additional learning about the drift over the investment horizon $[0, T]$ unless T is large.

Our investor trades in the risky asset S and in the riskless asset B. We assume that the risk-free rate obtained on the bank account B is constant. We denote it by r and assume $\bar{\mu} > r \geq 0$. We assume that the investor evaluates risky and uncertain prospects according to the two-stage procedure introduced by Klibanoff et al. (2005). In the first step, the investor calculates her certainty equivalent of the risky prospect for each possible model, i.e., for each fixed realization of μ. In the second step, she treats model uncertainty as inducing a lottery over the calculated certainty equivalents.

To make these ideas precise, let V_T denote a payoff which depends only on the path $(S_t)_t$ and is realized at time T. The distribution of $(S_t)_t$ conditional on a fixed realization μ of the drift is indicated by P^μ. Correspondingly, E_{P^μ} is the expectation operator with respect to P^μ. In contrast, E_μ denotes the expectation with respect to the distribution of μ induced by the ambiguity. Then, the investor's time $t = 0$ certainty equivalent to receiving V_T is given by

$$v^{-1}\Big(E_\mu\Big[v\Big(u^{-1}\big(E_{P^\mu}[u(V_T)]\big)\Big)\Big]\Big) \tag{2}$$

for two increasing utility and ambiguity functions u and v. Throughout, we assume that both u and v are CRRA functions, i.e.,

$$u(x) = \begin{cases} \frac{x^{1-\gamma}}{1-\gamma} & \gamma > 1 \\ \ln x & \gamma = 1 \end{cases} \text{ and } v(x) = \begin{cases} \frac{x^{1-\eta}}{1-\eta} & \eta > 1 \\ \ln x & \eta = 1 \end{cases}$$

where γ and η capture the (constant) relative aversions towards risk and ambiguity.

If $\eta = \gamma$, the agent is ambiguity neutral in the sense that she has the same attitude towards fluctuations in wealth due to risk and towards fluctuations in wealth due to model uncertainty. In that case, the function $\phi = v \circ u^{-1}$ is linear so that the two expected values collapse into one. Similarly, an agent with $\eta > \gamma$ is ambiguity-averse in the sense that she is more averse towards fluctuations from uncertainty than towards fluctuations from risk. This corresponds to a concave ϕ. Finally, an agent with $\eta < \gamma$ is ambiguity-loving. Note however that we assume throughout that $\eta \geq 1$. An ambiguity-loving agent thus prefers ambiguity over risk – but she also prefers certainty over ambiguity.

[4] For simplicity, we call this asset "risky" throughout the paper even though it actually both risky and uncertain.

[5] In addition, Bollerslev, Chou, and Kroner (1992) give evidence for the predictability of variances. The effects of estimation errors of means and variances are analyzed in Chopra and Ziemba (1993).

2.1. Certainty equivalents

We now revisit the classical portfolio choice problem (Merton problem) under ambiguity. The investor maximizes the objective (2) where V_T is the terminal wealth from investing in a risky and ambiguous asset S with dynamics given by Eq. (1) and a riskless asset B. Denoting by m_t the fraction of wealth invested in the risky asset at time t, the dynamics of V is given by

$$dV_t = V_t\left(m_t\frac{dS_t}{S_t} + (1-m_t)\frac{dB_t}{B_t}\right)$$
$$= (m_t\mu + (1-m_t)r)V_t dt + \sigma m_t V_t dW_t$$

and thus

$$V_T = V_0 e^{rT+(\mu-r)\int_0^T m_t dt - \frac{1}{2}\sigma^2\int_0^T m_t^2 dt + \sigma\int_0^T m_t dW_t}.$$

Notice that, because of the double CRRA framework, we can normalize the initial investment to $V_0 = 1$. We assume throughout that the strategies $(m_t)_t$ are deterministic functions of time.

We distinguish between the overall certainty equivalent (investor's objective) $CE(V_T)$ defined by Eq. (2) and the certainty equivalent $CE^\mu(V_T)$ with respect to one admissible model of risk. $CE^\mu(V_T)$ thus denotes the certainty equivalent of the terminal wealth V_T under the measure P^μ corresponding to one admissible risk situation, i.e., a fixed realization of μ,

$$CE(V_T) = v^{-1}\left(E_\mu[v(CE^\mu(V_T))]\right) \quad \text{where} \quad CE^\mu(V_T) = u^{-1}(E_{P^\mu}[u(V_T)]).$$

$CE^\mu(V_T)$ thus captures a pure risk situation where the risk preference is given by a CRRA utility function u. The following lemma gives the certainty equivalents of a fixed deterministic strategy m_t under risk and under ambiguity.

Lemma 1 (Certainty equivalents)**.** *For a fixed deterministic strategy $(m_t)_t$ we have*

$$CE^\mu = e^{rT+(\mu-r)\int_0^T m_t dt - \frac{1}{2}\gamma\int_0^T m_t^2\sigma^2 dt} \quad \text{and}$$

$$CE = e^{rT+(\bar{\mu}-r)\int_0^T m_t dt - \frac{1}{2}\gamma\sigma^2\int_0^T m_t^2 dt - \frac{1}{2}(\eta-1)\bar{\sigma}^2\left(\int_0^T m_t dt\right)^2}. \tag{3}$$

3. Optimal investment strategies

In this section, we derive the optimal pre-commitment strategy, show why it is time-inconsistent and then derive the optimal dynamically consistent strategy.

3.1. The optimal pre-commitment strategy

Our first main result, Proposition 1 presents the strategy that maximizes the certainty equivalent under smooth ambiguity (3). Moreover, it provides an expression for the associated excess certainty equivalent growth rate. The excess certainty equivalent growth rate is the certain rate of return that leads to the same expected utility as the optimal investment strategy over the same time horizon in excess of the risk-free rate r. Here and in the following, we limit attention to strategies that prescribe deterministic investment fractions (m_t). Throughout our analysis, it turns out that deviations from the Black–Scholes benchmark are captured by a single, T-dependent parameter,

$$\tau = \omega T \quad \text{where} \quad \omega = \frac{(\eta-1)}{\gamma}\frac{\bar{\sigma}^2}{\sigma^2}.$$

τ can be interpreted as a risk- and ambiguity-weighted version of the investment horizon T where ω is the weighting factor. The weighting factor captures to which degree the situation is dominated either by ambiguity or by risk. In the Black–Scholes benchmark model, $\bar{\sigma} = 0$, the optimal investment strategy is given by the well-known Merton fraction

$$m^{BS} = \frac{\bar{\mu}-r}{\gamma\sigma^2}$$

while the associated excess certainty equivalent growth rate is given by

$$y^{BS} = \frac{1}{T}\ln CE^{BS} - r = \frac{(\bar{\mu}-r)^2}{2\gamma\sigma^2}.$$

Almost all of our subsequent results can be expressed in terms of m^{BS}, y^{BS} and τ.

Proposition 1 (Optimal pre-commitment investment strategy under ambiguity)**.** *The solution of*

$$m^{PC} = \text{argmax}_m\, v^{-1}\left(E_\mu[v(CE^\mu(V_T))]\right)$$

is given by the constant investment fraction

$$m^{PC} = \frac{\bar{\mu}-r}{\gamma\sigma^2 + (\eta-1)\bar{\sigma}^2 T} = m^{BS}\frac{1}{1+\tau}.$$

The associated optimal excess certainty equivalent growth rate $y^{PC} = \frac{1}{T}\ln CE^{PC} - r$ is

$$y^{PC} = \frac{(\bar{\mu}-r)^2}{2\left(\gamma\sigma^2 + (\eta-1)\bar{\sigma}^2 T\right)} = y^{BS}\frac{1}{1+\tau}.$$

The optimal strategy is thus constant over time. In the special $\eta = 1$ (a log-ambiguity function $v(x) = \ln x$), the solution coincides with the Merton solution with drift $\bar{\mu}$ and volatility σ. Otherwise ($\eta > 1$), the investor favors a lower investment fraction compared to the benchmark Merton solution, i.e. the investment fraction is the lower the higher the ambiguity aversion η is.

A simple calculation shows that if an investor follows the classical Black–Scholes strategy m^{BS} in the presence of ambiguity, the resulting excess certainty equivalent growth rate equals $y^{BS}(1-\tau)$. Not surprisingly this is always less than y^{BS}. In fact, for sufficiently long weighted time horizons, $\tau > 1$, the excess growth rate becomes negative. Abstaining from risky investment altogether is better than ignoring model uncertainty in this case: the constant strategy $m \equiv 0$ outperforms m^{BS}.

3.2. Dynamic inconsistency of optimal investment

The optimal investment fraction m^{PC} from Proposition 1 depends on the investment horizon T via $\tau = \omega T$. For all t in $[0, T]$, it corresponds to the investment strategy

$$m_t^{PC} = \frac{\bar{\mu}-r}{\gamma\sigma^2 + (\eta-1)\bar{\sigma}^2 T}.$$

In particular, since this is an investment fraction (rather than an invested amount), implementing this strategy requires constant rebalancing of the portfolio. If the investor can pre-commit herself to this strategy at time $t = 0$, she will do so – it is her optimal strategy. Yet at later dates, the strategy is no longer optimal. For this reason, we refer to m^{PC} as the pre-commitment strategy.

Concretely, at a subsequent date t_l, the investor prefers a pre-commitment to the strategy

$$m_{t,t_l}^{PC} = \frac{\bar{\mu}-r}{\gamma\sigma^2 + (\eta-1)\bar{\sigma}^2(T-t_l)}.$$

for all $t \in [t_l, T]$. She thus wants to invest more aggressively. Intuitively, her (cumulated) ambiguity over the remaining investment horizon decreases in time. Thus, as time evolves (and the investment horizon shrinks), the impact of ambiguity is lower. Such revisions of the investment strategy ultimately lead to the myopic, time-dependent strategy

$$m_t^{MY} = \frac{\bar{\mu}-r}{\gamma\sigma^2 + (\eta-1)\bar{\sigma}^2(T-t)} = m^{BS}\frac{1}{1+\omega(T-t)}$$

which yields an excess certainty equivalent growth rate of

$$y^{MY} = y^{BS}\left(2\frac{\ln(1+\tau)}{\tau} - \frac{1}{1+\tau} - \frac{(\ln(1+\tau))^2}{\tau}\right).$$

Let us clarify in which sense the strategy m^{MY} is myopic. The investment goal that motivates the strategy is utility from terminal wealth. Thus, in contrast to some classical notions of myopic behavior, our myopic investor's goals are not limited to today's consumption. The investor is myopic in the sense that she ignores the fact that the impact of ambiguity changes (in fact, decreases) over time.[6] This type of myopic behavior is also distinct from the one described in the seminal work of Campbell and Viceira (1999). They consider a setting with time-varying investment opportunities. In this setting, optimal investment strategies consist of two terms, a first term which corresponds to our Merton fraction m^{BS} and a second term, the so-called hedging demand, which hedges the agent against fluctuations in investment opportunities. In their terminology, an agent is myopic if her investment strategy neglects the hedging demand (which does not arise in our Black–Scholes setting).

3.3. Dynamically consistent strategies

Our agent perceives a decrease in ambiguity over time: the myopic investment fraction increases gradually, converging to the Merton fraction m^{BS} in the limit $t \to T$. Thus, our next step is to derive the optimal strategy that accounts for these predictable changes in ambiguity attitudes and the resulting dynamic inconsistency. The analysis of the type of dynamic inconsistency we find here goes back to Strotz (1955). Strotz emphasizes the distinction between optimal pre-commitment strategies and optimal dynamically consistent strategies.[7] A strategy is dynamically consistent if the agent has no incentive to deviate from the strategy as time evolves. Strotz argues that if pre-commitment is not possible, the optimal dynamically consistent strategy should be regarded as the "true" optimal strategy. More recently, this view has been reemphasized in Caplin and Leahy (2006). A seminal application in the context of mean-variance portfolio choice is given in Basak and Chabakauri (2010).

Technically, the optimal dynamically consistent strategy is derived as the solution of the dynamic program associated with the investment problem (the "HJB equation") – even though this solution does *not* correspond to the overall optimal strategy m^{PC} under dynamic inconsistency. At time t the investor chooses the optimal dynamically consistent investment fraction m_t^{DC} given that the investment fractions at the later dates $(t, T]$ are already optimized by her future selves. The following proposition characterizes the optimal dynamically consistent strategy m_t^{DC} at time t and the associated excess certainty equivalent growth rate y^{DC} over $[0, T]$.

Proposition 2 (Optimal dynamically consistent strategy)**.** *The optimal dynamically consistent investment strategy is given by*

$$m_t^{DC} = \frac{\bar{\mu} - r}{\gamma\sigma^2} e^{\frac{-(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}(T-t)} = m^{BS} e^{-\omega(T-t)}.$$

The associated excess certainty equivalent growth rate $y^{DC} = \frac{1}{T}\ln CE^{DC} - r$ *is*

$$y^{DC} = \begin{cases} y^{BS}\frac{1}{2\tau}\left(1 - e^{-2\tau}\right) & \text{for } \eta > 1, \\ y^{BS} & \text{for } \eta = 1. \end{cases}$$

Notice that we have $\tau = \omega = 0$ in the boundary case $\eta = 1$ when the ambiguity function is logarithmic. In this case, the pre-commitment strategy is independent of T so that there is no time-inconsistency. Moreover, all three strategies m^{PC}, m^{MY} and m^{DC} then coincide with the constant Merton fraction m^{BS} throughout the entire investment horizon.

In general, m_t^{DC} is increasing in t just like the myopic strategy m_t^{MY}. The cumulated ambiguity decreases over time and thus the investment fraction increases. Compared to the Black–Scholes benchmark, the investment fraction implied by the optimal dynamically consistent strategy is the lower the longer the remaining investment horizon is. In the limit $t \to T$, there is no reason to deviate from the Black–Scholes solution, i.e., $\lim_{t\to T} m_t^{DC} = m^{BS}$. Fig. 1a compares the dynamically consistent strategy to its competitors. In the beginning, the investment fraction is the smallest, as the agent takes into account that at later dates a smaller perceived ambiguity will lead to a more aggressive investment behavior.

It is instructive to consider the limiting cases of very severe or almost negligible uncertainty, corresponding to large or small values of $\bar{\sigma}$. If uncertainty is very severe, the best thing the investor can do is to go out of the financial market, i.e. $\lim_{\bar{\sigma}\to\infty} m^{PC} = 0$. On the other hand, if there is no uncertainty we are back in the benchmark case as $\lim_{\bar{\sigma}\to 0} m^{PC} = m^{BS}$. For the dynamically consistent strategies the limiting cases are the same.

Finally, note that our focus on deterministic strategies is not restrictive when it comes to the dynamically consistent strategy. Inspecting the proof of Proposition 2, we see that current wealth is merely an additive term in the log-certainty equivalent that is optimized at each time point. Thus, even if we allowed for stochastic strategies in the optimization, the resulting optimal strategies are always deterministic as they do not depend on the evolution of wealth.[8]

From this observation about the dynamically consistent strategy, we can conclude that allowing for stochastic strategies cannot cure the time-inconsistency problems we observe. Vanishing time-inconsistency would imply that the optimal stochastic PC strategy coincides with the optimal stochastic DC strategy. We know that the optimal stochastic DC strategy is the deterministic strategy identified in Proposition 2. Yet the strategy from Proposition 2 cannot be the optimal stochastic PC strategy as it is outperformed by the optimal deterministic PC strategy of Proposition 1.[9]

4. Comparison of strategies

This section provides a more in depth comparison of the different investment strategies.

4.1. Quantitative comparisons

As a preliminary step in our comparisons, observe that all three strategies, the pre-commitment, the myopic, and the dynamically consistent strategies can be interpreted as time-discounted versions of the Black–Scholes strategy m^{BS}. The discount rate ω is always the same, but "interest rate convention" and the discounting horizon differ. The pre-commitment strategy $m^{PC} = m^{BS}/(1 + \omega T)$ can be interpreted as discounting the Black–Scholes strategy m^{BS} with simple compounding using the "interest rate" $\omega = \frac{(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}$ over the whole investment horizon. The rate ω which appears in this discount factor is unrelated to the actual interest rate. ω depends on σ, γ, $\bar{\sigma}$ and η, which are the parameters that quantify the amounts of risk and ambiguity and the attitudes towards them. The myopic strategy $m_t^{MY} = m^{BS}/(1 + \omega(T - t))$

[6] In the context of dynamic portfolio choice under mean-variance preferences in a Black–Scholes setting without ambiguity, a similar type of myopic investor was studied in Pedersen and Peskir (2017).

[7] In the literature, what we call "dynamically consistent" has also been referred to as "time consistent", "sophisticated" or "consistent planning".

[8] This is basically the same reasoning as the one that implies that the constant Merton fraction is optimal in a Black–Scholes setting with CRRA utility.

[9] A natural follow-up question is whether the time-inconsistency problems vanish when we allow for both stochastic strategies and Bayesian learning. This is discussed in Section 5.

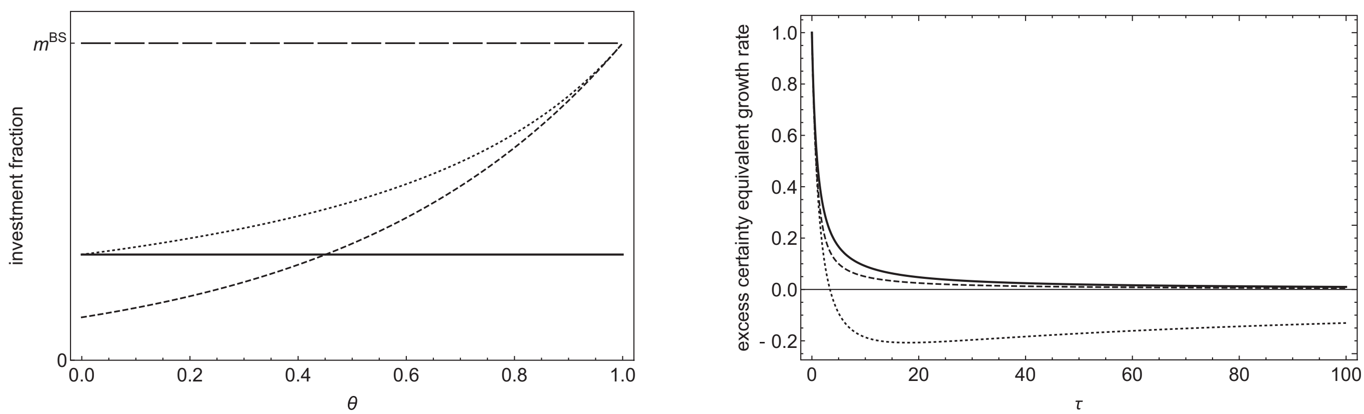

(a) Investment fractions. (b) Excess certainty equivalent growth rates.

Fig. 1. Strategies. For varying fractions $\theta = t/T$ of elapsed time, Panel (a) depicts the investment fractions implied by the pre-commitment (solid), myopic (dotted) and dynamically consistent (dashed) strategies over time for $\tau = 2$. The long-dashed line gives the corresponding Black-Scholes solution (without ambiguity) as a benchmark. Panel (b) compares the excess growth rate initiated at $t = 0$ for the pre-commitment investor (y^{PC}, solid), the myopic investor (y^{MY}, dotted) and the dynamically consistent investor (y^{DC}, dashed). Growth rates are stated as fractions of y^{BS}.

arises again by simple compounding over the remaining investment horizon. In contrast, the dynamically consistent strategy $m_t^{DC} = m^{BS}\exp(-\omega(T-t))$ results from *continuous* compounding with the same rate and horizons as in the myopic strategy. We now use these insights to derive some comparisons between the average investment fractions of the three strategies:

Proposition 3 (Comparison of average investment fractions)**.** *Define average investment fractions* $\bar{m}^{DC}$ *and* $\bar{m}^{MY}$ *by*

$$\bar{m}^{DC} = \frac{1}{T}\int_0^T m_t^{DC}\,dt = \frac{\bar{\mu}-r}{(\eta-1)\bar{\sigma}^2 T}\left(1 - e^{-\frac{(\eta-1)\bar{\sigma}^2 T}{\gamma\sigma^2}}\right)$$
$$= \frac{m^{BS}}{\tau}\left(1 - e^{-\tau}\right)$$

for the optimal dynamically consistent strategy and by

$$\bar{m}^{MY} = \frac{1}{T}\int_0^T m_t^{MY}\,dt = \frac{\bar{\mu}-r}{(\eta-1)\bar{\sigma}^2 T}\ln\left(1 + \frac{(\eta-1)\bar{\sigma}^2 T}{\gamma\sigma^2}\right)$$
$$= \frac{m^{BS}}{\tau}\ln(1+\tau)$$

for the myopic strategy. We have the ranking of average investment fractions

$$\bar{m}^{MY} \geq \bar{m}^{DC} \geq m^{PC}.$$

Moreover, $\bar{m}^{DC}/m^{PC} \to 1$ *in the limit for both* $T \to \infty$ *and* $T \to 0$*, and* $m_0^{DC} < m^{PC}$*. In contrast,* $\bar{m}^{MY}/m^{PC} \to \infty$ *as* $T \to \infty$ *whenever* $\omega > 0$*.*

On average, the investor relying on the dynamically consistent strategy invests too much in the risky asset compared to the pre-commitment solution. However, by taking into account how her ambiguity attitudes will influence her future decisions, she avoids investing as much as under the myopic strategy: the dynamically consistent investor invests less now because she anticipates that she will invest more later.

The difference between the optimal pre-commitment strategy and the optimal dynamically consistent strategy is illustrated in Fig. 1a. At the beginning of the investment horizon, the investment fraction implied by the dynamically consistent strategy (dashed line) is lower than the one of the pre-commitment strategy (solid line). However, the opposite is true when the end of the investment horizon is reached. There, the dynamically consistent investment fraction converges to the Merton solution (long-dashed line). Throughout, the myopic investment fraction (dotted line) is higher than both the pre-commitment and the dynamically consistent strategy. The horizon is fixed by $\tau = 2$, where τ corresponds to the weighted terminal time ωT. The horizontal axis shows the fraction of time that has elapsed, i.e. $\theta = t/T$.

In a way, dynamically consistent investment strikes a balance between the myopic and pre-commitment strategies. Visually, the dynamically consistent strategy resembles the myopic one with its convex and increasing investment fraction. Yet the average investment fraction of the dynamically consistent strategy remains, in a sense, close to the pre-commitment case while the myopic strategy is far too aggressive on long horizons. This can also be seen in the limit results provided in Proposition 3.

Fig. 1b shows the excess certainty equivalent growth rates as fractions of y^{BS} for varying τ. Notice that the longer the investment horizon, the larger the ambiguity and thus the lower the excess growth rate of both the pre-commitment and the dynamically consistent strategies. Due to its overall optimality, the pre-commitment strategy consistently outperforms the other two strategies. The excess growth rate implied by the myopic strategy becomes negative for sufficiently large τ. Myopic investing is thus dominated by not investing at all. Yet in the limit $\tau \to \infty$, y^{MY} converges to zero as the average invested amount vanishes. Finally, notice that this figure is not merely an illustration: all three curves in the picture depend on the model parameters only through τ.[10]

For a numerical illustration of the strategies, we assume an investment horizon of $T = 2$ year and model parameters $\bar{\mu} = 0.05$, $r = 0$, $\bar{\sigma} = 0.02$, and $\sigma = 0.15$. Table 1 shows the investment strategies for $\gamma = \{3, 4, 5\}$. If η is greater than one, then ambiguity leads to more conservative investment into the risky asset. However, there is not merely a reduction in risky investment. For the myopic and the dynamically consistent strategies, we also observe a different time pattern of optimal investment where the risk exposure increases gradually over time. For these two strategies, we thus report investment fractions at times $t = 0$, $t = 1$ and $t = 2$.

4.2. The value of commitment

The optimal dynamically consistent strategy is necessarily outperformed by the optimal pre-commitment strategy. Under dynamic consistency, the same objective is optimized over a smaller class of admissible strategies, namely, the dynamically consistent

[10] The same is true for Figs. 2 and 3.

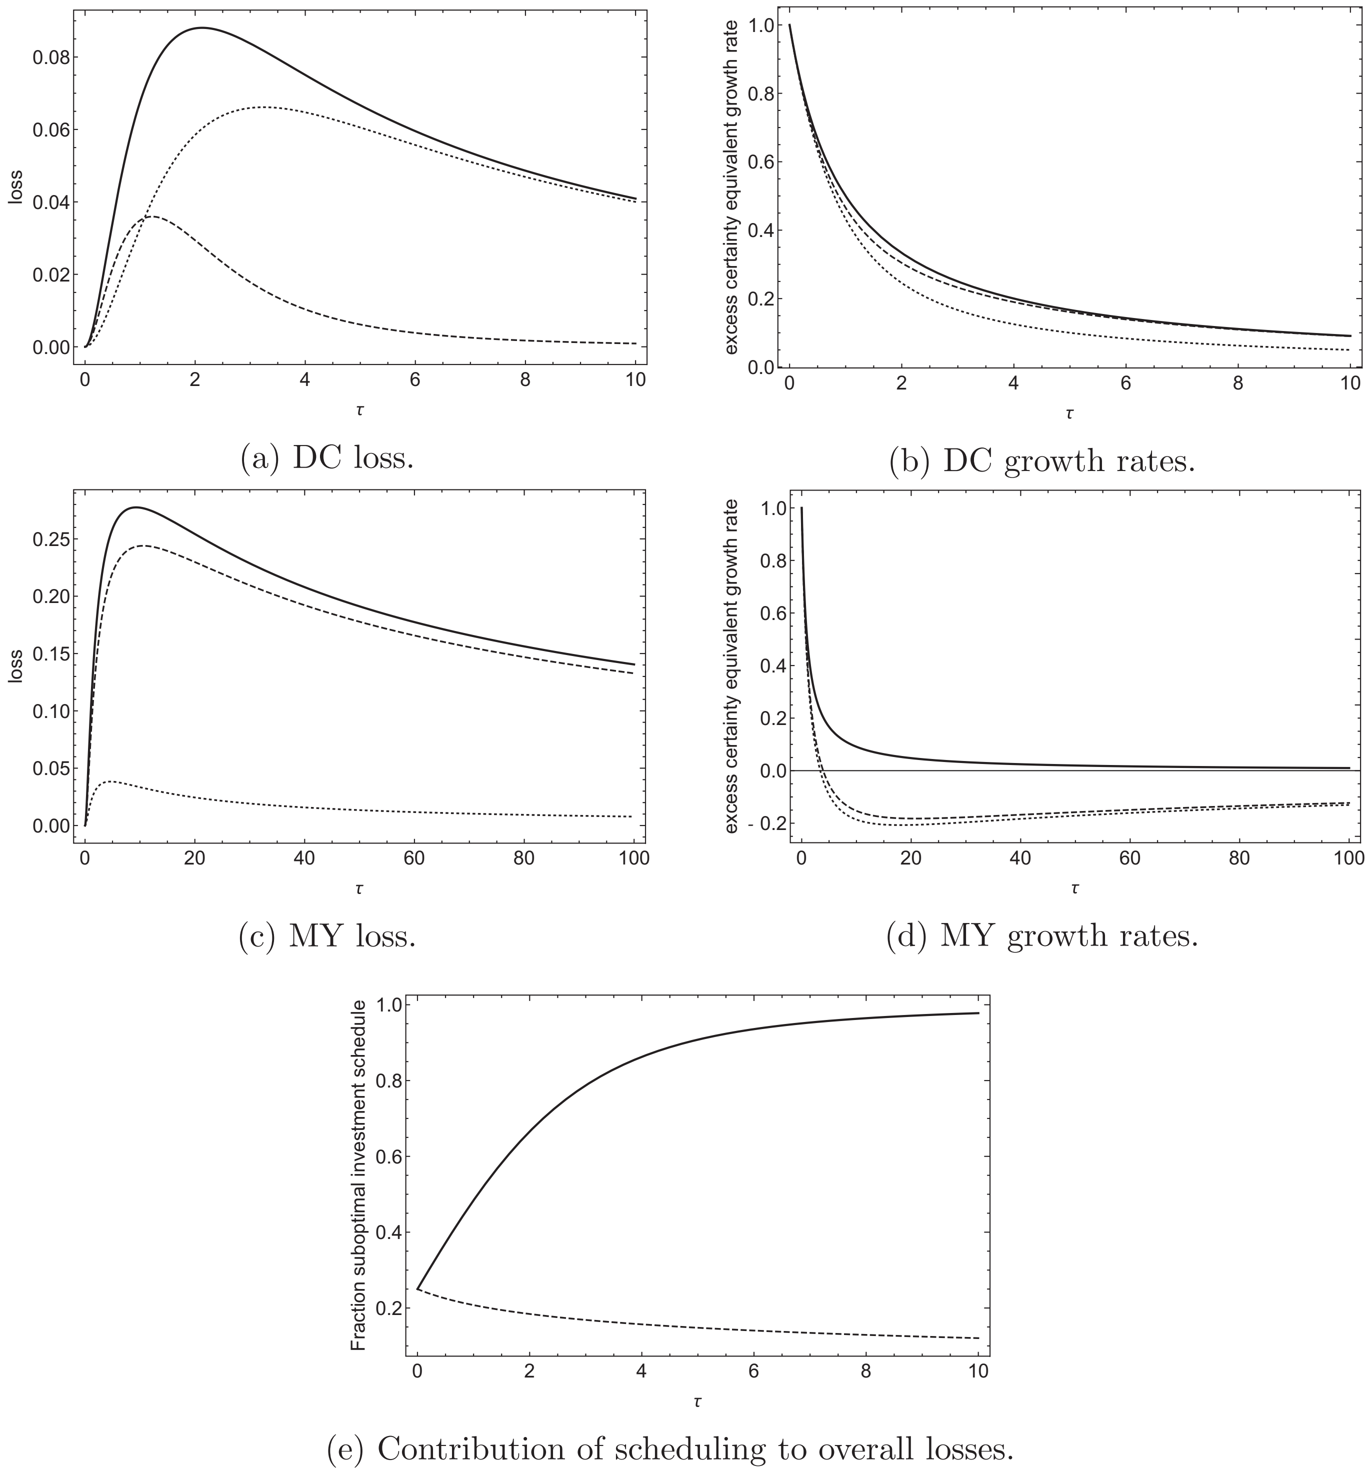

Fig. 2. Growth rates and loss decomposition. In panels (a) and (c) the total losses l^{DC} and l^{MY} are depicted by the solid lines, the losses due to overinvestment l_o^{DC} and l_0^{MY} are depicted by the dashed line and the losses due to suboptimal investment scheduling l_s^{DC} and l_s^{MY} are depicted by the dotted lines. In panels (b) and (d) $\tilde{y}^{DC}$ and $\tilde{y}^{MY}$ are, respectively, depicted by the dashed lines, y^{PC} is depicted by the solid line and y^{DC} and y^{MY} are depicted by the dotted lines. All three excess growth rates are stated as fractions of y^{BS}. Panel (e) shows the contribution of the suboptimal investment scheduling effect to the total loss l_s/l, for both the dynamically consistent (solid line) and the myopic strategies (dashed line).

ones. A crucial question is how much the investor loses if she is limited to dynamically consistent strategies or, conversely, how much she gains if pre-commitment is possible. To quantify this value of commitment we define the loss l^{DC} as

$$l^{DC} = \frac{y^{PC} - y^{DC}}{y^{BS}}.$$

Here we divide by the Black–Scholes benchmark excess growth rate, so that the loss is expressed as a fraction of this quantity. This implies in particular that the loss is purely a function of τ which is depicted as the solid curve in Fig. 2a.

Notice that the loss rate is first increasing in the investment horizon but, after some critical level of horizon, it is decreasing. In

Table 1
Optimal investment strategies.

	$\eta = 1$			$\eta = 4$			$\eta = 7$		
	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$
m^{BS}	0.74	0.56	0.44	0.74	0.56	0.44	0.74	0.56	0.44
m^{PC}	0.74	0.56	0.44	0.16	0.15	0.14	0.09	0.09	0.08
m_0^{MY}	0.74	0.56	0.44	0.16	0.15	0.14	0.09	0.09	0.08
m_1^{MY}	0.74	0.56	0.44	0.27	0.24	0.22	0.16	0.15	0.14
m_2^{MY}	0.74	0.56	0.44	0.74	0.56	0.44	0.74	0.56	0.44
m_0^{DC}	0.74	0.56	0.44	0.02	0.04	0.05	0.00	0.00	0.01
m_1^{DC}	0.74	0.56	0.44	0.13	0.15	0.15	0.02	0.04	0.05
m_2^{DC}	0.74	0.56	0.44	0.74	0.56	0.44	0.74	0.56	0.44

particular, the ability to pre-commit on a strategy is most valuable for some intermediate time horizons. Intuitively, this is explained by the limit behavior of investment fractions from Proposition 3. For very long horizons, ambiguity is so large that there is relatively little risky investment regardless of the strategy. For short horizons, the differences between the strategies are too small. Finally, note that the overall losses are quantitatively rather small. In fact, from Fig. 2a we can conclude an upper bound on the maximal loss rate, showing that it always amounts to less than 10% of the optimal excess growth rate in the Black–Scholes benchmark. This upper bound on losses[11] $l^{DC} \leq 0.09$ for all $\tau > 0$. The maximal loss bound is approximately attained at $\tau \approx 2.13$ and is independent of the model parameters otherwise.

4.3. Decomposing the losses

Basically, both the myopic strategy and the dynamically consistent strategy differ from the pre-commitment strategy in the same two ways. First, as seen in Proposition 3, both strategies lead to overinvestment in the risky asset on average. Second, both strategies replace the constant investment fraction of the pre-commitment strategy by a time-dependent investment fraction. In the following, we study for both of these strategies to what extent the losses compared to the optimum are due to (i) overinvestment and due to (ii) suboptimal investment scheduling. Our decompositions are based on artificial strategies whose constant investment fractions correspond to the average investment fractions of, respectively, the dynamically consistent and the myopic strategy. These benchmark strategies $\bar{m}^{DC}$ and $\bar{m}^{MY}$ thus combine the overinvestment of m^{DC} and m^{MY} with an optimal investment schedule. The proof of Proposition 1 shows that for any given average investment fraction the highest excess growth rate is achieved with a constant investment fraction.

4.3.1. Losses of the dynamically consistent strategy

We begin by decomposing the loss rate for the dynamically consistent strategy. By Proposition 3, the constant strategy $\bar{m}^{DC}$ with the same average investment fraction over $[0, T]$ as m_t^{DC} is given by

$$\bar{m}^{DC} = \frac{m^{BS}}{\tau}\left(1 - e^{-\tau}\right).$$

The strategy $\bar{m}^{DC}$ can be interpreted as the optimal pre-commitment strategy of an agent who is constrained to match the average investment fraction of the dynamically consistent strategy. The excess certainty equivalent growth rate over $[0, T]$ associated with the strategy $\bar{m}^{DC}$ is given by

$$\bar{y}^{DC} = y^{BS}\frac{(1 - e^{-\tau})}{\tau}\left(2 - (1 - e^{-\tau})\left(\frac{1}{\tau} + 1\right)\right).$$

[11] For large values of τ, not depicted in Fig. 2a follows directly from $l^{DC} \leq y^{PC}/y^{BS} = \frac{1}{1+\tau}$.

To disentangle the impacts of overinvestment and suboptimal investment scheduling, we thus compare $\bar{y}^{DC}$ with y^{DC} and with the optimum y^{PC}.

Proposition 4. (Loss decomposition for m^{DC}).
The total loss l^{DC} is the sum of the loss due to overinvestment

$$l_o^{DC} = \frac{y^{PC} - \bar{y}^{DC}}{y^{BS}} = \frac{e^{-2\tau}(1 - e^{\tau} + \tau)^2}{\tau^2(1 + \tau)} \tag{4}$$

and the loss due to suboptimal scheduling

$$l_s^{DC} = \frac{\bar{y}^{DC} - y^{DC}}{y^{BS}} = \frac{e^{-2\tau}(e^{\tau} - 1)(2 + e^{\tau}(\tau - 2) + \tau)}{2\tau^2}. \tag{5}$$

The decomposition of the losses is depicted in Figs. 2a and 2b. We see that the suboptimal scheduling effect is considerably larger than the overinvestment effect. This holds especially for larger values of τ where the difference between $\bar{y}^{DC}$ and the best-possible value y^{PC} vanishes quickly. The picture is reversed only for very short investment horizons where the effect of overinvestment dominates the effect of suboptimal scheduling.

One can further show that $\bar{y}^{DC} < 2 \cdot y^{DC}$. Thus, in the worst case, suboptimal scheduling reduces the excess growth rate by 50%. The worst-case impact of the overinvestment that comes with the dynamically consistent strategy is much smaller. One can show that $y^{PC} < 1.1 \cdot \bar{y}^{DC}$ so that adjusting the constant investment fraction from $\bar{m}^{DC}$ to m^{PC} increases the excess growth rate by at most 10%. This bound is approximately attained for $\tau \approx 1.8$.

4.3.2. Losses of the myopic strategy

We now study an analogous decomposition of losses for the myopic strategy. Here, we define the overall loss l^{MY} compared to the pre-commitment optimum as

$$l^{MY} = \frac{y^{PC} - y^{MY}}{y^{BS}}.$$

By Proposition 3, the constant investment fraction that equals the average investment fraction of the myopic strategy is now given by

$$\bar{m}^{MY} = \frac{m^{BS}}{\tau}\ln(1 + \tau).$$

For the associated excess growth rate, we find that

$$\bar{y}^{MY} = y^{BS}\frac{\ln(1 + \tau)}{\tau}\left(2 - \ln(1 + \tau)\left(1 + \frac{1}{\tau}\right)\right).$$

The decomposition into a loss l_o^{MY} from overinvestment and a loss l_s^{MY} from suboptimal scheduling now looks as follows:

Proposition 5. (Loss decomposition for m^{MY}).
The total loss l^{MY} is the sum of the loss due to overinvestment

$$l_o^{MY} = \frac{y^{PC} - \bar{y}^{MY}}{y^{BS}} = \frac{(\tau - (1 + \tau)\ln(1 + \tau))^2}{\tau^2(1 + \tau)} \tag{6}$$

and the loss due to suboptimal scheduling

$$l_s^{MY} = \frac{\bar{y}^{MY} - y^{MY}}{y^{BS}} = \frac{\tau^2 - (1+\tau)\ln(1+\tau)^2}{\tau^2(1+\tau)}. \quad (7)$$

Fig. 2e compares the relative contribution of suboptimal scheduling to overall losses between the myopic and the dynamically consistent strategies. It thus depicts l_s^{MY}/l^{MY} and l_s^{DC}/l^{DC} as functions of τ. For small τ, suboptimal scheduling contributes 25% to the overall loss of *both* strategies while the remaining 75% are due to overinvestment. For larger τ, most of the losses of the dynamically consistent strategy are due to scheduling. In contrast, more and more of the losses of the myopic strategy are due to overinvestment rather than suboptimal scheduling as τ increases.

The observation that the losses of a myopic investor are mainly due to overinvestment is confirmed in Fig. 2c and 2d which depict the excess growth rates and the losses as functions of τ. Notice in particular that even under optimal scheduling the overinvestment effect is so strong that $\bar{y}^{MY}$ eventually becomes negative.[12]

4.4. The point of regret

So far, we have focused on comparisons of the investment performance over the entire investment horizon $[0, T]$. In these comparisons, the pre-commitment strategy is the best-possible strategy by construction. Yet we have also seen that later the investor would prefer to update her commitment. In this section, we study whether there is a time point from which on the investor regrets her pre-commitment. More precisely, we ask whether there is a time point from which on the excess certainty equivalent growth rate of the dynamically consistent strategy is higher than the excess growth rate of an "old" pre-commitment strategy? We call this time point the "point of regret" t^R.

We need to introduce a richer notation for this discussion. We write $y^{PC}_{t_0,t_1,t_2}$ for the excess certainty equivalent growth rate over the period $[t_1, t_2]$ based on the optimal pre-commitment strategy that was initiated at time t_0. For example, the previous excess growth rate y^{PC} becomes $y^{PC}_{0,0,T}$. The excess growth rate of the pre-commitment strategy initiated at $t = 0$ over $[t, T]$ is thus obtained via

$$y^{PC}_{0,t,T} = \frac{1}{T-t} \ln\left(CE^{PC}_{0,t,T} / V_t\right) - r$$

where

$$CE^{PC}_{0,t,T} = V_t\, e^{rT + (\bar{\mu} - r)\int_t^T m_0^{PC} ds - \frac{1}{2}\gamma\sigma^2 \int_t^T m_0^{PC\,2} ds - \frac{1}{2}(\eta - 1)\bar{\sigma}^2 \left(\int_t^T m_0^{PC} ds\right)^2}$$

with $m_0^{PC} = m^{BS}\frac{1}{1+\omega T}$. It follows that

$$y^{PC}_{0,t,T} = y^{BS}\left(\frac{1}{1+\omega T} + \frac{\omega t}{(1+\omega T)^2}\right).$$

The excess growth rate of the dynamically consistent strategy over $[t, T]$ initiated at 0 is given by

$$y^{DC}_{0,t,T} = y^{BS}\left(\frac{1 - e^{-2\omega(T-t)}}{2\omega(T-t)}\right).$$

For the dynamically consistent strategy, the time of initiation is irrelevant. Regardless of when investment started, this strategy will look exactly the same over the interval $[t, T]$. We can thus drop the first subscript and write $y^{DC}_{t,T} = y^{DC}_{0,t,T}$. We say that t^R is the point of regret if $y^{DC}_{t,T} > y^{PC}_{0,t,T}$ for $t > t^R$ and $y^{DC}_{t,T} < y^{PC}_{0,t,T}$ for $t < t^R$.

[12] Note that the plotted range of τ is different for the myopic and the dynamically consistent strategies. For the dynamically consistent case, a small range is necessary to show how the loss from overinvestment dominates on short time horizons. For the myopic case, we need a broader range to display the full picture.

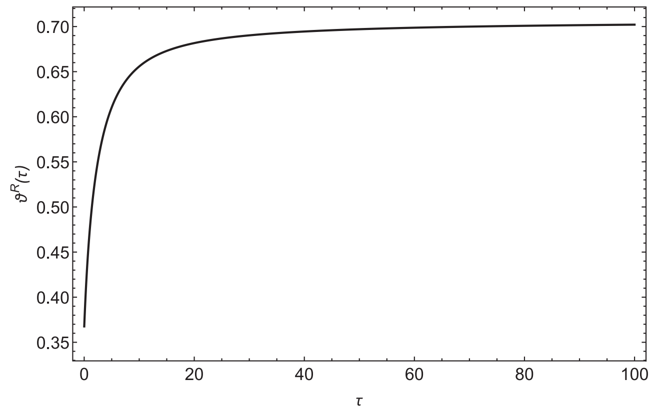

Fig. 3. Point of regret.

Proposition 6 shows that there always exists a unique point of regret. Moreover, there is a function $\theta^R : \mathbb{R} \to [0, 1]$, such that $t^R/T = \theta^R(\tau)$. The relative location of t^R in the overall investment horizon thus depends only on τ.

Proposition 6 (Location of the point of regret)**.** *There exists a unique point of regret $t^R \in (0, T)$ which is given by $t^R = \theta^R(\tau)T$. Here, the function $\theta^R(\tau)$ is the unique solution of*

$$\frac{1}{1+\tau} + \frac{\tau\theta^R(\tau)}{(1+\tau)^2} = \frac{1 - e^{-2\tau(1-\theta^R(\tau))}}{2\tau(1-\theta^R(\tau))} \quad (8)$$

for all $\tau > 0$. Moreover, we have the limiting values

$$\lim_{\tau\to 0} \theta^R(\tau) = \frac{\sqrt{3}-1}{2} \approx 36.6\% \quad \text{and} \quad \lim_{\tau\to\infty} \theta^R(\tau) = \frac{1}{\sqrt{2}} \approx 70.7\%.$$

The condition (8) characterizing $\theta^R(\tau)$ is equivalent to $y^{DC}_{t^R,T} = y^{PC}_{0,t^R,T}$. Fig. 3 shows that θ^R increases in τ. From a relative perspective, regret thus comes later on longer time horizons. Based on the figure and the two limit results, we see that θ^R always lies between 36.6% and 70.7%. Consequently, the investor never regrets her pre-commitment during the first 36% of the investment horizon. Yet she always regrets it during the final 29%.[13]

5. Learning

One main restriction of the previous analysis is that we assume that our agent does not dynamically adjust her beliefs by learning from the stock price trajectory she sees. A main justification for this is that learning about the drift coefficient tends to be very slow. As argued, e.g., in Rogers (2001), hundreds of years of data are necessary to obtain a precise estimate of the drift. This motivates us to assume that the impact of learning over a few years can safely be ignored.

Nevertheless, one may ask to what extent the time-consistency problems we find are driven by the absence of learning. For the ambiguity neutral case $\eta = \gamma$, a version of our problem with Bayesian learning was studied by Brennan (1998) and Rogers (2001). In this case, there is no time-inconsistency once learning is taken into account: Brennan's solution derived by dynamic programming coincides with Rogers' direct solution approach via the

[13] The limits of $\tau = \omega T$ taken in the proposition can be interpreted as varying the time horizon T for fixed ω, or they can be interpreted as varying the impact of ambiguity ω for a fixed time horizon T. In the latter case, as ambiguity vanishes, $\omega \to 0$, the PC and DC strategies converge to each other, while the point of regret t^R converges to 36.6%T. Compared to the limiting case $\omega = 0$ itself, there is a discontinuity. In the limiting case, both strategies equal the Merton fraction and there is no (point of) regret.

martingale method. What makes this case special (and also relatively tractable) is that the double expectations from the smooth ambiguity model turns into a single expected value. Due to the fact that the agent treats ambiguity and risk in the same way, the smooth ambiguity model collapses to an expected utility model.

In this section, we consider a simplified version of our baseline model with Bayesian learning. We assume that the agent is restricted to strategies with constant investment fractions that can only be adjusted once at a given intermediate time point. There are three time points $t_0 = 0,\ t_1 \in (0,T)$ and $t_2 = T$. At times t_0 and t_1 the agent chooses investment fractions m_{t_0} and m_{t_1} which are then kept constant until the next time point t_1 or t_2. We show that for general η and γ the problem remains time-inconsistent in the presence of learning. This is not unexpected as the agent's objective for $\eta \neq \gamma$ depends in a non-linear way on conditional expectations formed at intermediate time-points. Moreover, we find that the utility gains due to learning are very moderate.

As in our baseline model, we assume that at time 0 the agent has a Gaussian prior with parameters $\bar{\mu}$ and $\bar{\sigma}$ for the realization of the drift μ while the volatility σ is known. Numerical illustrations are based on a $T = 2$ year investment horizon and model parameters $\bar{\mu} = 0.05,\ r = 0,\ \bar{\sigma} = 0.02,$ and $\sigma = 0.15$.

We assume that the agent applies Bayesian learning. As derived in Rogers (2001), p. 146, the agent's time-t updated belief about the realization of μ is given by a Gaussian distribution with mean $\bar{\mu}_t$ and standard deviation $\bar{\sigma}_t$,

$$\bar{\mu}_t = \frac{1}{1+\frac{\bar{\sigma}^2}{\sigma^2}t}\left(\mu t\frac{\bar{\sigma}^2}{\sigma^2} + W_t\frac{\bar{\sigma}^2}{\sigma} + \bar{\mu}\right) \quad \text{and} \quad \bar{\sigma}_t^2 = \frac{\bar{\sigma}^2}{1+\frac{\bar{\sigma}^2}{\sigma^2}t}.$$

Even though this formula for $\bar{\mu}_t$ depends on the unobserved, true drift parameter μ, it can be evaluated based on the quantity $\mu t + \sigma W_t$ which can be computed from the stock price given the agent's knowledge of σ.

We denote by $y_{0,T}^{PC,L}$ and $y_{0,T}^{DC,L}$ the excess certainty equivalent growth rates from a time-0 perspective under, respectively, the optimal pre-commitment strategy and the optimal dynamically consistent strategy. The DC rate $y_{0,T}^{DC,L}$ is straightforward to compute. Clearly, at the last decision time t_1, the agent chooses the optimal *pre-commitment* strategy given her current knowledge for the remaining time horizon, i.e.

$$m_{t_1}^{DC,L} = \frac{\bar{\mu}_{t_1} - r}{\gamma\sigma^2 + (\eta-1)\bar{\sigma}_{t_1}^2(T-t_1)}$$

relying on the fact that she can pre-commit from t_1 to $t_2 = T$. In order to find $y_{0,T}^{DC,L}$ numerically, we just have to optimize over the constant investment fraction m_{t_0} at time $t_0 = 0$. This is illustrated by Fig. 4. For varying investment fractions m_{t_0}, the figure depicts the excess certainty equivalent (CE) growth rate given that the investment fraction $m_{t_1}^{DC,L}$ is fixed from time t_1 onwards. The highest excess growth rate is obtained for $m_{t_0}^{DC,L} = 0.7153$ when $\gamma = 3$ and $\eta = 4$.

Unfortunately, $y_{0,T}^{PC,L}$ is harder to compute. To show that the problem is time-inconsistent in general, we thus provide an example of a (possibly suboptimal) pre-commitment strategy that yields a higher excess growth rate $y_{0,T}^{exPC,L}$ than the dynamically consistent strategy, $y_{0,T}^{PC,L} \geq y_{0,T}^{exPC,L} > y_{0,T}^{DC,L}$. In particular, we consider the following example of a pre-commitment strategy $m^{exPC,L}$ which accounts for learning: $m_{t_0}^{exPC,L} = m_{t_0}^{DC,L}$ and $m_{t_1}^{exPC,L} = a \cdot m_{t_1}^{DC,L}$. Thus, the example of a pre-commitment strategy coincides with the dynamically consistent strategy (under learning) if $a = 1$. Otherwise, the agent pre-commits to adjusting her optimal dynamically consistent strategy at t_1 with the scaling parameter a. In order to show that time inconsistency is still a problem under learning, it is sufficient to find some $a \neq 1$ which implies a higher excess growth rate (seen from t_0) than implied by $a = 1$. A numerical illustration is given in Fig. 5. Notice that in the ambiguity-averse case $\gamma = 3 < \eta = 4$, a higher excess growth rate is obtained by choosing $a < 1$ than for $a = 1$. In contrast, the opposite is true for an ambiguity-loving agent with $\gamma = 5 > \eta = 4$. For an ambiguity neutral agent with $\gamma = \eta = 4$ the maximum excess growth rate is obtained at $a = 1$. This is in line with the continuous-time result of Rogers (2001), that time inconsistency vanishes in the special case $\gamma = \eta$.

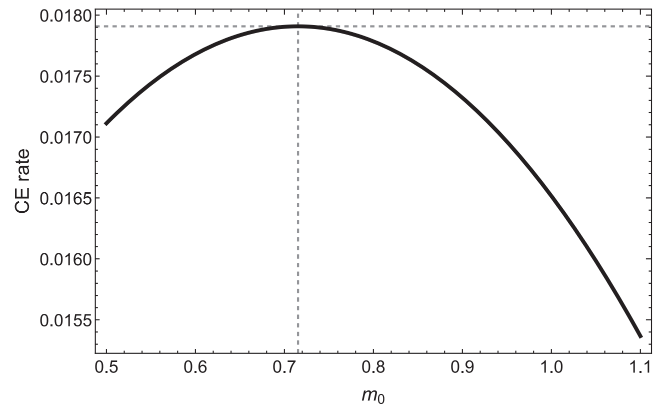

Fig. 4. Dynamically consistent case with learning. The figure depicts the excess certainty equivalent growth rate for fixed $m_{t_1}^{DC,L}$ and varying m_{t_0} with $t_0 = 0$ for $\gamma = 3$ and $\eta = 4$.

In order to obtain some quantitative understanding of the value of commitment under learning, $y_{0,T}^{PC,L} - y_{0,T}^{DC,L}$, we derive an upper bound on $y_{0,T}^{PC,L}$. To this end, we consider an agent who knows that at the second decision stage t_1 she will "meet a visionary" who informs her about the true value of μ. This agent has to base her first investment decision on the prior while in the second investment decision Bayesian learning is replaced by full information. We denote by $y_{0,T}^{PC,MV}$ and $y_{0,T}^{DC,MV}$ the optimal excess growth rates associated with this information structure. Clearly, we have $y_{0,T}^{PC,MV} > y_{0,T}^{PC,L}$. The next lemma shows that for the agent who meets the visionary there is no time-consistency problem and the optimal strategy is easy to derive.

Lemma 2. *The optimal dynamically consistent strategy for the agent who meets the visionary is to choose the Merton strategy*

$$m_{t_0}^{DC,MV} = \frac{\mu - r}{\gamma\sigma^2}$$

from t_0 to t_1 and the "pre-commitment strategy"

$$m_{t_1}^{DC,MV} = \frac{\bar{\mu} - r}{\gamma\sigma^2 + (\eta-1)\bar{\sigma}^2 T}$$

from t_1 to t_2. This dynamically consistent strategy is also optimal under pre-commitment and the resulting excess growth rates satisfy $y_{0,T}^{PC,MV} = y_{0,T}^{DC,MV}$.

Based on the lemma, we can easily obtain numerical values of $y_{0,T}^{PC,MV}$. These values can be used to establish upper bounds on the value of commitment under learning

$$y_{0,T}^{PC,L} - y_{0,T}^{DC,L} < y_{0,T}^{PC,MV} - y_{0,T}^{DC,L}$$

and on the value of learning under pre-commitment

$$y_{0,T}^{PC,L} - y_{0,T}^{PC,NL} < y_{0,T}^{PC,MV} - y_{0,T}^{PC,NL}.$$

Here, $y_{0,T}^{PC,NL}$ denotes the excess certainty equivalent growth under pre-commitment without learning as derived in Proposition 1. A quantification of the impact of learning is provided by Table 2. In addition to the upper bounds stated above, we consider the difference of the certainty equivalent rates $y_{0,T}^{DC,L} - y_{0,T}^{PC,NL}$. Recall that,

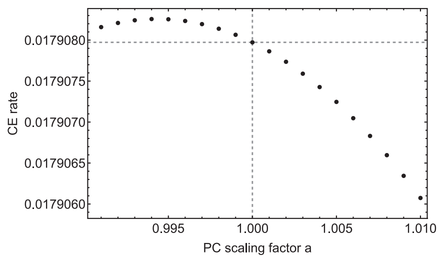

(a) $\gamma = 3$.

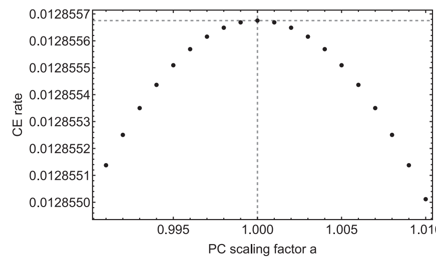

(b) $\gamma = 4$.

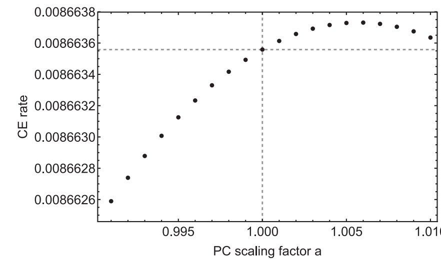

(c) $\gamma = 5$.

Fig. 5. Time inconsistency under learning and ambiguity. The figures depict excess certainty equivalent growth rates of $m^{exPC,L}$ as functions of the scaling factor a. The left panel shows the ambiguity-averse case, $\gamma = 3 < \eta = 4$. The middle panel is the ambiguity neutral case, $\gamma = \eta = 4$, and the right panel illustrates the ambiguity-loving case, $\gamma = 5 > \eta = 4$.

Table 2
Impact of learning for $\eta = 4$.

	$y_{0,T}^{PC,MV}$	$y_{0,T}^{DC,L}$	$y_{0,T}^{PC,NL}$	$y_{0,T}^{PC,MV} - y_{0,T}^{DC,L}$	$y_{0,T}^{PC,MV} - y_{0,T}^{PC,NL}$	$y_{0,T}^{DC,L} - y_{0,T}^{PC,NL}$
γ=3	0.01935	0.01791	0.01788	0.00144	0.00147	0.00003
γ=4	0.01463	0.01355	0.01353	0.00108	0.00110	**0.00002**
γ=5	0.01176	0.01089	0.01088	0.00087	0.00088	0.00002

in the special case of $\gamma = \eta$, it holds $y_{0,T}^{DC,L} - y_{0,T}^{PC,NL} = y_{0,T}^{PC,L} - y_{0,T}^{PC,NL}$. Here, the gain which is obtained by learning is only 0.00002. Generally, the quantitative impact of learning is small in this example.

6. Discussion

How robust is the time-inconsistency we find in our model? In the previous section we showed that time inconsistency persists even if we include learning. Denote by $m_{s,t,T}^{PC}$ the investment of the pre-commitment strategy at time t if pre-commitment was initiated at time s,

$$m_{s,t,T}^{PC} = m^{BS} \frac{1}{1 + (T-s)\frac{\eta-1}{\gamma}\frac{\bar{\sigma}^2}{\sigma^2}}.$$

If we want the time-inconsistency to vanish in a variation of the model then we have to modify it in such a way that $m_{s,t,T}^{PC}$ no longer depends on s. In a sense, this already shows that our problem of time-inconsistency is fairly robust. Possible variations of our model may lead to different pre-commitment strategies $m_{s,t,T}^{PC}$. Yet the time-inconsistency persists unless the factor $(T-s)$ cancels out for one reason or another. In addition, we can expect the model to behave qualitatively similar to ours as long $m_{s,t,T}^{PC}$ increases in s. This monotonicity implies that the myopic strategy $m_{t,T}^{MY} = m_{t,t,T}^{PC}$ leads to overinvestment on average, and we conjecture that dynamically consistent investment behavior will balance this overinvestment in a similar way as in our model.

In the following we discuss the scope of three possible approaches for making the model time-consistent. In particular, we briefly discuss (i) changes in the model of ambiguity attitudes, (ii) changes in the model of ambiguity itself, and (iii) the introduction of recursive preferences.

It is sometimes argued that ambiguity attitudes vary over time. One can thus make η time-dependent in such a way the pre-commitment strategy no longer depends on the length of the investment horizon. While somewhat artificial, such a variation of the model does not suffer from time-inconsistency. Yet the resolution of the consistency problem is based on an assumption that two a priori distinct effects cancel out each other *exactly*.[14]

Another approach for achieving this type of exact cancellation between potential sources of time-inconsistency is to vary the ambiguity model, i.e., the distributional assumptions on the drift μ. Intuitively, what is needed is that the perceived variance of μ at time s is inversely proportional to the length of the time horizon $T-s$. In the following, we present a model variation in the spirit of Chen and Epstein (2002)'s "i.i.d.-ambiguity" which has this property. Consequently, optimal investment is indeed time-consistent in this model.

We assume now that the drift μ is no longer constant over $[0,T]$. Instead, we assume that $(\mu_t)_t$ is a stochastic process which satisfies the following assumption: the integrals $\int_{t_1}^{t_2}\mu_s ds$ and $\int_{t_3}^{t_4}\mu_s ds$ are stochastically independent from each other and from $(W_t)_t$ for any $0 \le t_1 < t_2 \le t_3 < t_4 \le T$. Moreover,

$$\frac{1}{t_2 - t_1}\int_{t_1}^{t_2}\mu_s ds \sim \mathcal{N}\left(\bar{\mu}, \frac{\bar{\sigma}^2}{t_2 - t_1}\right). \tag{9}$$

This model captures the beliefs of an investor who thinks that while the Black–Scholes benchmark is not literally correct, it is a reasonably accurate approximation of the asset's long run behavior. As can be seen from (9), the average drift over a time interval $[s,t]$ has mean $\bar{\mu}$ and a variance which vanishes as $t-s$ grows large. The agent thus knows that the excess return on very long horizons is close to $\bar{\mu}$. Uncertainty is mostly a short-run phenomenon.

In a sense, this distributional assumption is the polar opposite of our previous assumption of an unknown but constant drift. This can be seen from the following microfoundation for (9) in the spirit of Chen and Epstein (2002)'s i.i.d.-ambiguity. Suppose that the investor believes that the instantaneous drifts μ_t are independent random variables which are normally distributed with mean $\bar{\mu}$ and variance $\bar{\sigma}^2$. Then, a law-of-large-numbers-type argument implies that average drifts behave as indicated by (9). To phrase this argument slightly more formally, the assumption of Gaussian i.i.d. instantaneous drifts is the same as assuming that the instantaneous drifts are Gaussian white noise, i.e., there exists a Brownian motion $W_t^{\perp}$ such that the aggregate drift over a time interval

[14] The general idea of making an investment problem tractable by introducing time variation in the strength of ambiguity attitudes is found, e.g., in Maenhout (2004) in a loosely related model. Maenhout shows that introducing ambiguity by means of a penalty increases the effective risk aversion. Balter and Horvath (2018) show that Maenhout's model of behavior under uncertainty is equivalent to the behavior in a pure risk situation with a more impatient agent, who evaluates his expectations under a distorted probability measure. Thus, dynamic consistency of the optimization problem is inherited from the corresponding pure risk setting.

$[t_1, t_2]$ can be written as

$$\int_{t_1}^{t_2} \mu_s ds = \bar{\mu}(t_2 - t_1) + \bar{\sigma}(W_{t_2}^{\perp} - W_{t_1}^{\perp}). \quad (10)$$

Formula (10) immediately yields (9). It turns out that the time-behavior of model uncertainty in (9) is *exactly* what is needed in order to render our investment problem time-consistent. For a fixed strategy $(m_t)_t$ which depends on time but not on the asset dynamics we have[15]

$$CE^{\mu} = e^{rT + \int_0^T m_t(\mu_t - r)dt - \frac{1}{2}\gamma\sigma^2 \int_0^T m_t^2 dt} \quad \text{and thus}$$

$$CE = e^{rT + (\bar{\mu} - r)\int_0^T m_t dt - \frac{1}{2}(\gamma\sigma^2 + (\eta - 1)\bar{\sigma}^2)\int_0^T m_t^2 dt}. \quad (11)$$

The final certainty equivalent is equal to that in a standard Merton problem with $\gamma\sigma^2 + (\eta - 1)\bar{\sigma}^2$ in place of $\gamma\sigma^2$. Consequently, the optimal pre-commitment strategy is now given by

$$m^{PC} = \frac{\bar{\mu} - r}{\gamma\sigma^2 + (\eta - 1)\bar{\sigma}^2}.$$

This strategy does not depend on the investment horizon and is thus time-consistent.

A final common route towards eliminating time-inconsistency is to directly define preferences as recursive.[16] We could replace our preference functional by the dynamic program that leads to the dynamically consistent strategy. This approach automatically promotes the dynamically consistent strategy from “best feasible without commitment” to “unique optimum”. Yet in our view, such a recursive view is not very natural for the problem like investing towards retirement. It feels artificial to claim that the agent optimizes how she will feel about retirement tomorrow – instead of thinking about retirement directly.

To summarize, one can come up with model variations that do not suffer from time-inconsistency. Yet it seems that time-consistency is a knife-edge case that arises only when the relevant quantities cancel out in a suitable way. Investors and regulators should be aware of the fact that time-inconsistency is closely tied to ambiguity or model uncertainty. Thus, it is important to understand and quantify the benefits of pre-commitment and the impact of time inconsistency on investment behavior.

7. Conclusion

In a simple model, we study the effects of ambiguity on the riskiness of investment decisions. Intuitively, an aversion against ambiguity makes optimal investment decisions less aggressive compared to a suitable benchmark without ambiguity. In addition, time inconsistency arises naturally, as the impact of ambiguity depends on the length of the investment horizon. Due to the time-inconsistency, there is a distinction between the optimal dynamically consistent strategy, and the optimal pre-commitment strategy. When comparing the excess certainty equivalent growth rates over the entire investment horizon, the dynamically consistent strategy suffers from two sources of suboptimality. Part of the comparative loss is due to overinvestment into the risky asset. The other part is due to a suboptimal investment schedule that arises as the investor has to balance her current behavior against later changes in perceived ambiguity, which she already anticipates. Pre-commitment thus beats the dynamically consistent strategy from a time 0 perspective. Yet due to the continuously decreasing perceived ambiguity, there will always be a point of regret when the investor wishes she had *not* entered a (now suboptimal) pre-commitment in the beginning.

Appendix A. Proofs

A1. Proof of Lemma 1

Since the term $\int_0^T m_t \mu dt$ is assumed to be deterministic under P^{μ}, the same calculations as in the standard risk situation yield

$$CE^{\mu} = u^{-1}(E_{P^{\mu}}[u(V_T)]) = e^{rT + (\mu - r)\int_0^T m_t dt - \frac{1}{2}\gamma\sigma^2 \int_0^T m_t^2 dt}.$$

Using that $\mu \sim \mathcal{N}(\bar{\mu}, \bar{\sigma}^2)$ now yields the expression for $CE = v^{-1}\big(E_{\mu}[v(CE^{\mu}(V_T))]\big)$.

A2. Proof of Proposition 1

To see that optimal strategies are given by constant investment fractions, consider all strategies for which the average fraction is constant, $\int_0^T m_t dt = \kappa$. It suffices to show that for each fixed κ the optimal strategy within this class of strategies is constant and then to optimize over κ. Observe that by Lemma 1, the certainty equivalents can then be expressed in terms of κ, κ^2 and $\int_0^T m_t^2 dt$ and that certainty equivalents are larger if $\int_0^T m_t^2 dt$ is small.[17] It follows just like in the standard situation under risk that

$$\min_{(m_t)_t} \int_0^T m_t^2 dt \ \text{ s.t. } \int_0^T m_t dt = \kappa$$

is solved by $m_t = \kappa/T$. Differentiating the (log)-certainty equivalents from Lemma 1 with respect to a constant strategy $m_t \equiv m$ yields the first order condition $\bar{\mu} - r - m\gamma\sigma^2 - m(\eta - 1)\bar{\sigma}^2 = 0$. Solving for m yields the optimal strategy and, after a few manipulations, the excess certainty equivalent growth rate.

A3. Proof of Proposition 2

We now consider investment fractions of the form $m_t^{DC} = \alpha(t)$, i.e. the strategies under consideration may depend deterministically on time. Following Lemma 1, the log-certainty equivalent as seen from time t is given by

$$\ln CE = \ln(V_t) + rT + (\bar{\mu} - r)\int_t^T \alpha(u)\,du - \frac{1}{2}\gamma\sigma^2 \int_t^T \alpha^2(u)\,du + \frac{1}{2}(1 - \eta)\bar{\sigma}^2\left(\int_t^T \alpha(u)\,du\right)^2$$

Dynamic consistency implies that the investor optimizes the investment fraction at time t conditional on her (then) optimal decisions made at later times. Thus, we consider the optimization problem

$$\ln(V_t) + rT + (\bar{\mu} - r)\left(\alpha(t)\epsilon + \int_{t+\epsilon}^T \alpha(u)\,du\right) - \frac{1}{2}\gamma\sigma^2\left(\alpha^2(t)\epsilon + \int_{t+\epsilon}^T \alpha^2(u)\,du\right) + \frac{1}{2}(1 - \eta)\bar{\sigma}^2\left(\alpha(t)\epsilon + \int_{t+\epsilon}^T \alpha(u)\,du\right)^2 \to \max_{\alpha(t)},$$

[15] To see this quickly, one can write (10) as $\mu_t dt = \bar{\mu} dt + \bar{\sigma} dW_t^{\perp}$, substitute this into (11) and then compute the certainty equivalent similarly to the baseline model.

[16] In the context of smooth ambiguity, this is done, e.g. in Ju and Miao (2012) and Klibanoff, Marinacci, and Mukerji (2009). Similar ideas of formulating preferences directly as a dynamic program or an HJB equation have also been used in combination with other models of ambiguity such as the multiplier preferences discussed in Anderson, Hansen, and Sargent (2003) and Maenhout (2004), see e.g. Liu (2010), Yi, Li, Viens, and Zeng (2013), Escobar, Ferrando, and Rubtsov (2015) and Gu, Viens, and Yi (2017).

[17] Recall that we restrict attention to $\gamma, \eta \geq 1$. Otherwise, degenerate situations are possible where the dependence on $\int_0^T m_t^2 dt$ is reversed.

i.e., optimizing the current decision $\alpha(t)$ at time t while holding later decisions after $t+\epsilon$ fixed (and then letting $\epsilon > 0$ vanish). The first order condition is

$$(\bar{\mu}-r)-\gamma\sigma^2\alpha(t)+(1-\eta)\bar{\sigma}^2\left(\int_{t+\epsilon}^{T}\alpha(u)\,du+\alpha(t)\epsilon\right)=0$$

such that, in the optimum, it holds that

$$\alpha(t)=\lim_{\epsilon\to 0}\frac{(\bar{\mu}-r)-(\eta-1)\bar{\sigma}^2\int_{t+\epsilon}^{T}\alpha(u)\,du}{\gamma\sigma^2+(\eta-1)\bar{\sigma}^2\epsilon}=\frac{(\bar{\mu}-r)-(\eta-1)\bar{\sigma}^2\int_{t}^{T}\alpha(u)\,du}{\gamma\sigma^2}$$

i.e.

$$\alpha(t)=\frac{(\bar{\mu}-r)-(\eta-1)\bar{\sigma}^2A(t)}{\gamma\sigma^2}\quad\text{where } A(t)=\int_t^T\alpha(u)\,du.$$

Recall that $A'(t)=-\alpha(t)$, i.e.

$$A'(t)=-\left(\frac{(\bar{\mu}-r)-(\eta-1)\bar{\sigma}^2A(t)}{\gamma\sigma^2}\right)=-\frac{\bar{\mu}-r}{\gamma\sigma^2}+\frac{(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}A(t).$$

Observe that $f'(t)=a+bf(t)$ implies $f(t)=\kappa e^{bt}-\frac{a}{b}$ and thus

$$A(t)=\frac{\bar{\mu}-r}{(\eta-1)\bar{\sigma}^2}+\kappa e^{\frac{(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}t}=\frac{\bar{\mu}-r}{(\eta-1)\bar{\sigma}^2}+\tilde{\kappa}e^{\frac{-(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}(T-t)}.$$

Obviously, we have $A(T)=0$ such that $\tilde{\kappa}=-\frac{\bar{\mu}-r}{(\eta-1)\bar{\sigma}^2}$, i.e.

$$A(t)=\frac{\bar{\mu}-r}{(\eta-1)\bar{\sigma}^2}-\frac{\bar{\mu}-r}{(\eta-1)\bar{\sigma}^2}e^{\frac{-(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}(T-t)}.$$

It immediately follows that

$$m_t^{DC}=\alpha(t)=-A'(t)=\frac{\bar{\mu}-r}{\gamma\sigma^2}e^{\frac{-(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}(T-t)}.$$

To derive the excess certainty equivalent growth rate of the optimal time-consistent strategy, notice that

$$\int_0^T m_u^{DC}\,du=\frac{\bar{\mu}-r}{(\eta-1)\bar{\sigma}^2}\left(1-e^{-\frac{(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}T}\right)$$

and $\int_0^T {m_u^{DC}}^2\,du=\frac{1}{2}\frac{\bar{\mu}-r}{\gamma\sigma^2(\eta-1)\bar{\sigma}^2}\left(1-e^{-2\frac{(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}T}\right)$

which immediately implies

$$\ln\left(v^{-1}\left(E_\mu\left[v\left(u^{-1}(E_{P^\mu}[u(V_T)])\right)\right]\right)\right)=rT+\frac{1}{4}\frac{(\bar{\mu}-r)^2}{(\eta-1)\bar{\sigma}^2}\left(1-e^{-2\frac{(\eta-1)\bar{\sigma}^2}{\gamma\sigma^2}T}\right).$$

A4. Proof of Proposition 3

The formulas for the average investment fractions follow by direct calculation. To compare the investment fractions, we have to show that

$$\frac{1}{1+\tau}\le\frac{1}{\tau}(1-e^{-\tau})\le\frac{1}{\tau}\ln(1+\tau).$$

The first inequality follows from

$$\frac{1}{1+\tau}\le\frac{1}{\tau}(1-e^{-\tau})\Leftrightarrow e^{\tau}\ge\tau+1$$

where the final inequality is well-known. We prove the second inequality by showing that $h^{DC}(\tau)\le h^{MY}(\tau)$ where $h^{DC}(\tau)=1-e^{-\tau}$ and $h^{MY}(\tau)=\ln(1+\tau)$. Since $h^{DC}(0)=h^{MY}(0)=0$ it suffices to compare the derivatives,

$$h^{DC\prime}(\tau)\le h^{MY\prime}(\tau)\Leftrightarrow e^{-\tau}\le\frac{1}{1+\tau}\Leftrightarrow e^{\tau}\ge 1+\tau.$$

The inequality $e^{-\tau}\le\frac{1}{1+\tau}$ also implies $m_0^{DC}<m^{PC}$. The limits as $T\to 0$ and $T\to\infty$ follow directly.

A5. Proof of Proposition 4

We saw in the proof of Proposition 1 that the optimal (precommitment) strategy with an arbitrary fixed average investment fraction is constant over time. The excess certainty equivalent growth rate implied by the constant investment fraction that equals the average investment of the dynamically consistent strategy is $\bar{y}^{DC}$. This is obtained by plugging in $\bar{m}^{DC}$ for m_t in (3), which yields for any constant investment fraction $\bar{m}$

$$\bar{y}=\frac{1}{T}\ln CE-r=(\bar{\mu}-r)\bar{m}-\frac{1}{2}\gamma\sigma^2\bar{m}^2-\frac{1}{2}(\eta-1)\bar{\sigma}^2\bar{m}^2T. \qquad (A.1)$$

Since

$$y^{PC}=y^{BS}\frac{1}{1+\tau}\quad\text{and}$$

$$\bar{y}^{DC}=y^{BS}\frac{(1-e^{-\tau})}{\tau}\left(2-(1-e^{-\tau})\left(\frac{1}{\tau}+1\right)\right)$$

the loss due to overinvestment is

$$\frac{y^{PC}-\bar{y}^{DC}}{y^{BS}}=\frac{1}{1+\tau}-\frac{(1-e^{-\tau})}{\tau}\left(2-(1-e^{-\tau})\left(\frac{1}{\tau}+1\right)\right)$$

which simplifies to (4). The loss due to suboptimal scheduling is obtained by

$$\frac{\bar{y}^{DC}-y^{DC}}{y^{BS}}=\frac{(1-e^{-\tau})}{\tau}\left(2-(1-e^{-\tau})\left(\frac{1}{\tau}+1\right)\right)-\frac{\left(1-e^{-2\tau}\right)}{2\tau}$$

which simplifies to (5).

A6. Proof of Proposition 5

Plugging $\bar{m}^{MY}$ into (A.1) gives the expression for the excess certainty equivalent growth rate. Subsequently, the loss is decomposed into the part due to overinvestment by

$$\frac{y^{PC}-\bar{y}^{MY}}{y^{BS}}=\frac{1}{1+\tau}-\frac{\ln(1+\tau)}{\tau}\left(2-\ln(1+\tau)\left(1+\frac{1}{\tau}\right)\right)$$

and due to suboptimal scheduling by

$$\frac{\bar{y}^{MY}-y^{MY}}{y^{BS}}=\frac{\ln(1+\tau)}{\tau}\left(2-\ln(1+\tau)\left(1+\frac{1}{\tau}\right)\right)-\left(2\frac{\ln(1+\tau)}{\tau}-\frac{1}{1+\tau}-\frac{(\ln(1+\tau))^2}{\tau}\right)$$

which lead to Eqs. (6) and (7) respectively.

A7. Proof of Proposition 6

To describe the properties of $\theta^R(\tau)$, we introduce the functions

$$g(\theta,\tau)=\frac{1-e^{-2\tau(1-\theta)}}{2\tau(1-\theta)}\quad\text{and}\quad h(\theta,\tau)=\frac{1}{1+\tau}+\theta\frac{\tau}{(1+\tau)^2}.$$

The equation characterizing $\theta^R(\tau)$ is then $g(\theta,\tau)-h(\theta,\tau)=0$, i.e., for each τ the value $\theta^R(\tau)$ is given by the intersection of $g(\theta,\cdot)$ and $h(\theta,\cdot)$. In a first step, we show that $\tau g(\theta,\tau)$ is increasing and convex in θ and that $\tau h(\theta,\tau)$ is increasing and linear in θ to show uniqueness of θ^R. In a second step, we prove that $\tau g(0,\tau)\le\tau h(0,\tau)$ and $\tau g(1,\tau)\ge\tau h(1,\tau)$ to show existence. The first derivative of $\tau g(\theta,\tau)$ is

$$\frac{\partial\tau g(\theta,\tau)}{\partial\theta}=\frac{e^{-2\tau(1-\theta)}\left(e^{2\tau(1-\theta)}-1-2\tau(1-\theta)\right)}{2(1-\theta)^2}\ge 0$$

where the inequality holds because $e^x\ge 1+x$. This proves that the function is increasing. The second derivative is

$$\frac{\partial^2\tau g(\theta,\tau)}{\partial\theta^2}=\frac{e^{-2\tau(1-\theta)}}{(1-\theta)^3}\left(e^{2\tau(1-\theta)}-1-2\tau(1-\theta)-2\tau^2(1-\theta)^2\right)\ge 0$$

where the inequality holds because $e^x \geq 1 + x + \frac{x^2}{2}$. This proves that the function is convex. Moreover, $\tau h(\theta, \tau)$ is clearly linear and increasing due to $\frac{\tau^2}{(1+\tau)^2} \geq 0$. To prove that the two functions cross, it suffices to show that $g(\theta, \tau)$ starts below the linear function and ends above. Hence, we evaluate both function at $\theta = 0$ and $\theta = 1$,

$$\tau g(0, \tau) = \frac{1 - e^{-2\tau}}{2}, \quad \lim_{\theta \uparrow 1} \tau g(\theta, \tau) = \lim_{\theta \uparrow 1} \frac{-2\tau e^{-2\tau(1-\theta)}}{-2} = \tau,$$

and

$$\tau h(0, \tau) = \frac{\tau}{1 + \tau}, \quad \tau h(1, \tau) = \frac{\tau(1 + 2\tau)}{(1 + \tau)^2}.$$

It thus remains to show the two inequalities $\tau g(0, \tau) \leq \tau h(0, \tau)$ and $\tau g(1, \tau) \geq \tau h(1, \tau)$. The second inequality is straightforward, namely

$$\tau g(1, \tau) \geq \tau h(1, \tau) \Leftrightarrow \tau \geq \frac{\tau(1 + 2\tau)}{(1 + \tau)^2} \Leftrightarrow \tau^2 \geq 0.$$

Here, the division by τ in the second step leaves the inequality unchanged since $\tau \geq 0$ by assumption. The first inequality can be written as

$$\tau g(0, \tau) \leq \tau h(0, \tau) \Leftrightarrow \frac{1 - e^{-2\tau}}{2} \leq \frac{\tau}{1 + \tau}$$
$$\Leftrightarrow 2\tau \geq (1 + \tau)(1 - e^{-2\tau}).$$

To see the final inequality, we show that $(1 + \tau)(1 - e^{-2\tau})$ is concave in τ and that the derivative evaluated at $\tau = 0$ is equal to the derivative of 2τ evaluated at $\tau = 0$. Hence, the linear function 2τ is the tangent of the concave function at this point and thus the inequality holds. To see this, it suffices to consider the relevant derivatives of the right hand side,

$$\frac{\partial (1 + \tau)(1 - e^{-2\tau})}{\partial \tau}\bigg|_{\tau=0} = e^{-2\tau}\left(1 + e^{2\tau} + 2\tau\right)\bigg|_{\tau=0} = 2$$

and $$\frac{\partial^2 (1 + \tau)(1 - e^{-2\tau})}{\partial \tau^2} = -4e^{-2\tau}\tau < 0.$$

To find the limiting value of θ^R for $\tau \to 0$, we perform a second order Taylor expansion of the equation

$$\frac{1}{1 + \tau} + \frac{\tau\theta}{(1 + \tau)^2} - \frac{1 - e^{-2\tau(1-\theta)}}{2\tau(1 - \theta)} = 0$$

around $\tau = 0$. This gives

$$\frac{1}{3}(1 - 2\theta - 2\theta^2)\tau^2 + \mathcal{O}(\tau^3) = 0 \Rightarrow 1 - 2\theta^R(0) - 2\theta^R(0)^2 = 0.$$

The unique positive solution of this equation is $\frac{1}{2}(\sqrt{3} - 1)$. For the limit $\tau \to \infty$, we multiply the characterizing equation by $(1 + \tau)$ to obtain

$$1 + \theta \frac{\tau}{1 + \tau} = \frac{1 + \tau}{\tau} \frac{1 - e^{-2\tau(1-\theta)}}{2(1 - \theta)}.$$

In the limit $\tau \to \infty$, this converges to

$$1 + \theta = \frac{1}{2(1 - \theta)} \quad \Leftrightarrow \theta^2 = \frac{1}{2}.$$

A8. Proof of Lemma 2

The dynamically consistent strategy is obtained backwards in time. Since at t_1 the true drift μ is observed, the agent is not uncertain and thus the remaining optimization problem reduces to the classical Merton problem without uncertainty. Therefore, under full information, the optimal investment decision from t_1 onwards is $m_{t_1} = \frac{\mu - r}{\gamma\sigma^2}$. Next, the optimal dynamically consistent strategy at t_0 is derived by maximizing the expected utility based on the ambiguity attitude described by $v(\cdot)$, the Gaussian prior with $\bar{\mu}$ and $\bar{\sigma}$ at time t_0, and the already made decision for m_{t_1}. The time-t_0 certainty equivalents in the stylized example under learning reduce to

$$CE^{\mu}(V_{t_2}) = e^{rt_2 + t_1(\mu - r)m_{t_0} - t_1\frac{1}{2}\gamma m_{t_0}^2\sigma^2 + (t_2 - t_1)(\mu - r)m_{t_1} - (t_2 - t_1)\frac{1}{2}\gamma m_{t_1}^2\sigma^2}$$

and

$$CE(V_{t_2}) = v^{-1}\left(E_{\mu}[v(CE^{\mu}(V_{t_2}))]\right)$$
$$= e^{rt_2 + \frac{(\bar{\mu}-r)t_1 + 2t_1 m_{t_0}\gamma(\bar{\mu}-r)\sigma^2 - m_{t_0}^2\gamma\sigma^2 t_1(\gamma\sigma^2 + 2t_2(\eta-1)\bar{\sigma}^2)}{2(\gamma\sigma^2 + (t_2 - t_1)(\eta-1)\bar{\sigma}^2)}}$$
$$\left(\frac{\eta - 1}{\gamma}\frac{\bar{\sigma}^2}{\sigma^2}(t_2 - t_1) + 1\right)^{\frac{1}{2(\eta-1)}}.$$

Thus the first-order condition with respect to m_{t_0} implies $m_{t_0} = \frac{\bar{\mu} - r}{\gamma\sigma^2 + t_2(\eta-1)\bar{\sigma}^2}$. It remains to show that the dynamically consistent strategy we just derived is also pre-commitment optimal. To this end, consider an agent at time t_0 who is choosing the optimal investment strategy between t_1 and t_2 while the strategy m_{t_0} between t_0 and t_1 is fixed at some value. This second-period strategy can, in principle, depend on μ, on S_{t_1}, and on the investment strategy between t_0 and t_1. However, due to the multiplicative structure of the CRRA utility, the utility contribution of the second period return conditional on μ, S_{t_1} and m_{t_0} does not depend on S_{t_1} and m_{t_0}.[18] In particular, the optimal strategy does not depend on S_{t_1} and m_{t_0}. This implies that the optimization problem for the second period strategy is equivalent between the pre-commitment problem and the backward solution approach so the second period strategies coincide. Consequently, also the first period strategies coincide between the dynamically consistent and the pre-commitment approach as the dynamically consistent first-period strategy is an optimal response to the second-period strategy.

References

Anderson, E. W., Hansen, L. P., & Sargent, T. J. (2003). A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association, 1*(1), 68–123.

Balter, A. G., & Horvath, F. (2018). Homothetic robust preferences. *Technical report*.

Basak, S., & Chabakauri, G. (2010). Dynamic mean-variance asset allocation. *Review of Financial Studies, 23*(8), 2970–3016.

Ben-Tal, A., El Ghaoui, L., & Nemirovski, A. (2009). *Robust optimization*: 28. Princeton University Press.

Ben-Tal, A., Margalit, T., & Nemirovski, A. (2000). Robust modeling of multi-stage portfolio problems. In H. Frenk, K. Roos, T. Terlaky, & S. Zhang (Eds.), *High performance optimization* (pp. 304–328). Boston, MA.: Springer US.

Bertsimas, D., Brown, D. B., & Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM Review, 53*(3), 464–501.

Bertsimas, D., & Pachamanova, D. (2008). Robust multiperiod portfolio management in the presence of transaction costs. *Computers & Operations Research, 35*(1), 3–17.

Biagini, S., & Pınar, M. C (2017). The robust Merton problem of an ambiguity averse investor. *Mathematics and Financial Economics, 11*(1), 1–24.

Björk, T., Khapko, M., & Murgoci, A. (2017). On time-inconsistent stochastic control in continuous time. *Finance and Stochastics, 21*(2), 331–360.

Björk, T., & Murgoci, A. (2014). A theory of Markovian time-inconsistent stochastic control in discrete time. *Finance and Stochastics, 18*(3), 545–592.

Bollerslev, T., Chou, R. Y., & Kroner, K. F. (1992). ARCH modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics, 52*(1), 5–59.

Boyle, P., Garlappi, L., Uppal, R., & Wang, T. (2012). Keynes meets Markowitz: The trade-off between familiarity and diversification. *Management Science, 58*(2), 253–272.

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.

Brennan, M. J. (1998). The role of learning in dynamic portfolio decisions. *European Finance Review, 1*(3), 295–306.

Campbell, J. Y., & Viceira, L. M. (1999). Consumption and portfolio decisions when expected returns are time varying. *Quarterly Journal of Economics, 114*(2), 433–495.

Caplin, A., & Leahy, J. (2006). The recursive approach to time inconsistency. *Journal of Economic Theory, 131*(1), 134–156.

[18] This is due to our assumption of perfect learning after meeting the visionary. Under Bayesian learning, beliefs at t_1 are correlated with returns between t_0 and t_1.

Chen, H., Ju, N., & Miao, J. (2014). Dynamic asset allocation with ambiguous return predictability. *Review of Economic Dynamics, 17*(4), 799–823.
Chen, Z., & Epstein, L. (2002). Ambiguity, risk, and asset returns in continuous time. *Econometrica, 70*(4), 1403–1443.
Chopra, V. K., & Ziemba, W. T. (1993). The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management, 19*(2), 6–11.
Collard, F., Mukerji, S., Sheppard, K., & Tallon, J.-M. (2018). Ambiguity and the historical equity premium. *Quantitative Economics, 9*(2), 945–993.
Cong, F., & Oosterlee, C. W. (2016). On pre-commitment aspects of a time-consistent strategy for a mean-variance investor. *Journal of Economic Dynamics and Control, 70*, 178–193.
Dantzig, G. B., & Infanger, G. (1993). Multi-stage stochastic linear programs for portfolio optimization. *Annals of Operations Research, 45*(1), 59–76.
Epstein, L. G., & Miao, J. (2003). A two-person dynamic equilibrium under ambiguity. *Journal of Economic Dynamics and Control, 27*(7), 1253–1288.
Epstein, L. G., & Schneider, M. (2010). Ambiguity and asset markets. *Annual Review of Financial Economics, 2*, 315–334.
Escobar, M., Ferrando, S., & Rubtsov, A. (2015). Robust portfolio choice with derivative trading under stochastic volatility. *Journal of Banking & Finance, 61*, 142–157.
Fabozzi, F. J., Huang, D., & Zhou, G. (2010). Robust portfolios: Contributions from operations research and finance. *Annals of Operations Research, 176*(1), 191–220.
Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics, 18*(2), 141–153.
Gollier, C. (2011). Portfolio choices and asset prices: The comparative statics of ambiguity aversion. *Review of Economic Studies, 78*(4), 1329–1344.
Gu, A., Viens, F. G., & Yi, B. (2017). Optimal reinsurance and investment strategies for insurers with mispricing and model ambiguity. *Insurance: Mathematics and Economics, 72*, 235–249.
Guidolin, M., & Rinaldi, F. (2013). Ambiguity in asset pricing and portfolio choice: A review of the literature. *Theory and Decision, 74*(2), 183–217.
Gülpınar, N., & Rustem, B. (2007). Worst-case robust decisions for multi-period mean–variance portfolio optimization. *European Journal of Operational Research, 183*(3), 981–1000.
Hansen, L. P., & Marinacci, M. (2016). Ambiguity aversion and model misspecification: An economic perspective. *Statistical Science, 31*(4), 511–515.
Hansen, L. P., & Sargent, T. J. (2008). *Robustness*. Princeton University Press.
Ju, N., & Miao, J. (2012). Ambiguity, learning, and asset returns. *Econometrica, 80*(2), 559–591.
Kim, J. H., Kim, W. C., & Fabozzi, F. J. (2014). Recent developments in robust portfolios with a worst-case approach. *Journal of Optimization Theory and Applications, 161*(1), 103–121.
Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica, 73*(6), 1849–1892.
Klibanoff, P., Marinacci, M., & Mukerji, S. (2009). Recursive smooth ambiguity preferences. *Journal of Economic Theory, 144*(3), 930–976.
Liu, H. (2010). Robust consumption and portfolio choice for time varying investment opportunities. *Annals of Finance, 6*(4), 435–454.
Maccheroni, F., Marinacci, M., & Ruffino, D. (2013). Alpha as ambiguity: Robust mean-variance portfolio analysis. *Econometrica, 81*(3), 1075–1113.
Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *Review of Financial Studies, 17*(4), 951–983.
Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics, 8*(4), 323–361.
Pedersen, J. L., & Peskir, G. (2017). Optimal mean-variance portfolio selection. *Mathematics and Financial Economics, 11*(2), 137–160.
Pflug, G. C., & Pohl, M. (2018). A review on ambiguity in stochastic portfolio optimization. *Set-Valued and Variational Analysis, 26*(4), 733–757.
Pun, C. S. (2018). Robust time-inconsistent stochastic control problems. *Automatica, 94*, 249–257.
Raponi, V., Uppal, R., & Zaffaroni, P. (2018). *Portfolio choice with model misspecification: A foundation for alpha and beta portfolios*.
Rogers, L. C. G. (2001). The relaxed investor and parameter uncertainty. *Finance and Stochastics, 5*(2), 131–154.
Skiadas, C. (2013). Smooth ambiguity aversion toward small risks and continuous-time recursive utility. *Journal of Political Economy, 121*(4), 775–792.
Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies, 23*(3), 165–180.
Strub, M. S., & Li, D. (2020). Failing to foresee the updating of the reference point leads to time-inconsistent investment. *Operations Research, 68*(1), 199–213.
Suzuki, M. (2018). Continuous-time smooth ambiguity preferences. *Journal of Economic Dynamics and Control, 90*, 30–44.
Taboga, M. (2005). Portfolio selection with two-stage preferences. *Finance Research Letters, 2*(3), 152–164.
Tütüncü, R. H., & Koenig, M. (2004). Robust asset allocation. *Annals of Operations Research, 132*(1), 157–187.
Uppal, R., & Wang, T. (2003). Model misspecification and underdiversification. *Journal of Finance, 58*(6), 2465–2486.
Van Staden, P. M., Dang, D.-M., & Forsyth, P. A. (2021). The surprising robustness of dynamic mean-variance portfolio optimization to model misspecification errors. *European Journal of Operational Research, 289*, 774–792.
Vigna, E. (2020). On time consistency for mean-variance portfolio selection. *International Journal of Theoretical and Applied Finance*, 2050042.
Wang, J., & Forsyth, P. A. (2011). Continuous time mean variance asset allocation: A time-consistent strategy. *European Journal of Operational Research, 209*(2), 184–201.
Yi, B., Li, Z., Viens, F. G., & Zeng, Y. (2013). Robust optimal control for an insurer with reinsurance and investment under Heston's stochastic volatility model. *Insurance: Mathematics and Economics, 53*(3), 601–614.