

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Tail optimality and preferences consistency for intertemporal optimization problems

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1847264> since 2023-12-11T09:42:46Z

Published version:

DOI:10.1137/21M1435422

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Tail optimality and preferences consistency for intertemporal optimization problems *

Elena Vigna[†]

October 22, 2021

Abstract

Given an intertemporal optimization problem over a time interval $[t_0, T]$ and a control plan associated to it, we introduce the four notions of local and global tail optimality of the control plan, and local and global preferences consistency of the agent. While the notion of tail optimality of a control plan is not new, the main innovation of this paper is the definition of preferences consistency of an agent, that is a novel concept.

We prove that, in the case of a *linear* time-consistent problem where dynamic programming can be applied, the optimal control plan is globally tail-optimal and the agent is globally preferences-consistent. Opposite, in the case of a *non-linear* problem that gives rise to time inconsistency, we find that global tail optimality and global preferences consistency do not coexist. We analyze three common ways to attack a time-inconsistent problem: (i) precommitment approach, (ii) dynamically optimal approach, (iii) consistent planning approach. We find that none of the three approaches keeps simultaneously the desirable properties of global tail optimality and global preferences consistency: the existing approaches to time inconsistency are flawed in various ways. We also prove that if the performance criterion

*An earlier version of this paper has been circulated also under the title “Tail optimality and preferences consistency for stochastic optimal control problems”. I am indebted to Tiziano De Angelis, Bjarne Højgaard, Bertrand Lods, Massimo Marinacci and Luigi Montrucchio for useful comments that improved the paper.

[†]University of Torino and Collegio Carlo Alberto, Italy. Corso Unione Sovietica 218 bis, 10134, Torino, Italy. E-Mail: elena.vigna@unito.it. Tel. +39 011 670 5754.

includes a convex function of expected final wealth and a globally tail-optimal plan exists, then the three approaches coincide and the problem is linear.

The contribution of the paper is to disentangle the notion of time consistency into the two notions of tail optimality and preferences consistency. The analysis should shed light on the price to be paid in terms of tail optimality and preferences consistency with each of the three approaches currently available for time inconsistency.

Keywords. Time consistency, dynamic programming, Bellman’s optimality principle, time inconsistency, precommitment approach, game theoretical approach, dynamically optimal approach, mean-variance portfolio selection.

JEL classification: C61, D81, G11.

1 Introduction

The aim of this paper is twofold: (i) to shed light on the differences between *linear* intertemporal optimization problems, where dynamic programming can be applied, and *non-linear* intertemporal optimization problems, where dynamic programming cannot be applied; (ii) to shed light on the differences among the three common approaches to non-linear problems, namely precommitment, dynamically optimal and consistent planning. The first class of problems is said to be *time-consistent* (see Björk & Murgoci, 2010) or to produce a *time-consistent plan* (see Strotz, 1956), while problems belonging to the second class are said to be *time-inconsistent*.

The notion of time inconsistency for optimization problems dates back to Strotz (1956). Broadly speaking, time inconsistency arises in an intertemporal optimization problem when the optimal strategy selected at some time t is no longer optimal at time $s > t$. In other words, a strategy is time-inconsistent when the agent at future time $s > t$ is tempted to deviate from the strategy decided at time t . For an illuminating and clarifying formalization of the possible sources of time inconsistency in intertemporal optimization problems, see Björk & Murgoci (2014) for the discrete time framework and Björk, Khapko & Murgoci (2017) for the continuous-time framework.

In the context of intertemporal optimization problems, the term *time-consistent/inconsistent* is somehow ambiguous, for it is used sometimes for the plan or behaviour adopted by the agent,

sometimes for the optimization problem or criterion used, sometimes for the agent herself. The meaning of optimality becomes doubtful too: Björk & Murgoci (2010) notice that “*It is thus conceptually unclear what we mean by ‘optimality’ and even more unclear what we mean by ‘an optimal control law’.*” The idea of valid time-consistent candidate strategies is unclear too. Bensoussan, Wong & Yam (2019) stress that “*the ‘plan that he (would) actually follow’,¹ which is expected to be the very definition of time-consistent plans, has not yet been rigorously described.*”

This general confusion stems from the fact that when talking about time consistency the two notions of optimality of the strategy and consistency to one’s own preferences are merged together.

In an attempt to add some clarity to the picture, in this paper we disentangle the notion of time consistency for an intertemporal optimization problem over $[t_0, T]$ into the two notions of *tail optimality of the control plan* and *preferences consistency of the agent*, and we provide rigorous definitions. Because of the dynamic nature of intertemporal optimization problems, both definitions of tail optimality and preferences consistency are provided at *local* level (at a single time point) and at *global* level (over a whole time interval).

The notions of tail optimality and preferences consistency are defined in detail in Sections 2 and 3, so we refer the reader to these sections for the rigorous mathematical treatment. However, a rough intuition could be the following. A control map is locally tail-optimal at t with respect to a given optimization problem on $[t, T]$ if, whenever played from t to T , permits the achievement of the optimum, and is globally tail-optimal on $[t_0, T]$ if all its restrictions on every interval $[t, T]$ are locally tail-optimal. An agent is locally preferences-consistent at time t with respect to a given optimization problem if she is instantaneously optimal at t with respect to that problem, and is globally preferences-consistent on a time interval $[t_0, T]$ if she is locally preferences-consistent at every time over that time interval.

The feature of global tail optimality of a control plan for a linear stochastic optimal control problem where dynamic programming is applicable is not a new concept and is indeed due to the validity of the Bellman’s optimality principle. Opposite, the definition of preferences consistency of an agent wishing to solve an intertemporal stochastic optimal control problem is novel. However, the idea of consistency to preferences is not new in the economics and decision theory literature, in particular in the context of choice among lotteries in a discrete setup. Epstein & Le Breton (1993) formulate that preferences should be updated to new information in order to preserve their

¹See Strotz (1956).

dynamic consistency. Johnsen & Donaldson (1985) investigate what are the restrictions imposed by time-consistent planning on the agent’s preferences. A link between non-linear problems and preferences consistency of the agent is partly addressed by Chew & Epstein (1990), who investigate whether a behaviour driven by non-expected utility preferences can be intertemporally consistent.

A preview of our results is the following. Expectedly, global tail optimality and global preferences consistency occur simultaneously in the case of linear optimization problems where dynamic programming can be applied. However, they no longer hold together with a non-linear problem for which dynamic programming cannot be applied. In particular, we find that for the precommitment approach there is local tail optimality and local preferences consistency at initial time t_0 ; for the dynamically optimal approach there is global preferences consistency, but there is no local tail optimality at any time; for the consistent planning (or game theoretical) approach there is neither local tail optimality nor local preferences consistency at any time with respect to the original non-linear problem, but there is global tail optimality and global preferences consistency with respect to a different linear problem. It is worth noting that the inconsistency to preferences of the consistent planning agent was already observed by Chew & Epstein (1990), see Sections 4.3 and 6.2.

Although the three approaches to time inconsistency turn out to be flawed in different ways, in the literature on time-inconsistent problems these approaches have been extensively used without considering the issues discussed here. For instance, Bensoussan et al. (2019) notice that *“To tackle this time-inconsistency issue, the game-theoretic approach is widely used to recommend a time-consistent solution.”* Ekeland, Mbodji & Pirvu (2012) adopt the game theoretical approach and notice that in the presence of time inconsistency *“the optimal strategies are not implementable”*. Dai, Hin, Kou & Xu (2021) celebrate the equilibrium approach for the mean-variance problem introduced by Basak & Chabakauri (2010) as a breakthrough over earlier work based on the precommitment. Björk, Murgoci & Zhou (2014) place a mean-variance portfolio optimization in continuous time within a game theoretical framework simply because of its time inconsistency.

This paper introduces a new perspective and framework under which time-inconsistent problems should be considered, and sheds further light on some of the criticisms of the game theoretical approach to the mean-variance problem appeared in the literature. Some examples are: Wang & Forsyth (2011), who find that in some cases the efficient frontier of the constrained investor is higher than the frontier in the absence of constraints; Bensoussan et al. (2019) provide an analysis

of how the constrained strategy can actually outperform the unconstrained one arguing that time consistency can be seen as an additional constraint on admissible controls, that limits the flexibility of earlier players and imposes a sort of penalty on the value function; and Forsyth (2020), who notices that requiring time consistency changes the objective function and may produce strategies with undesirable characteristics.

The remainder of the paper is as follows. In Section 2, we formulate the notions of local tail optimality and global tail optimality of a control plan, and we prove global tail optimality in the case of a linear optimization problem. In Section 3, we formulate the notions of local preferences consistency and global preferences consistency of an agent, and we prove global preferences consistency in the case of a linear optimization problem. In Section 4, we extend the analysis to general non-linear time-inconsistent problems, analyzing tail optimality and preferences consistency for the three common approaches to time inconsistency. In Section 5 we prove that when the performance criterion includes a convex function of expected final wealth the existence of a globally tail-optimal control plan implies the linearity of the optimization problem and the coincidence of the three approaches. In Section 6, we illustrate in detail the special case of the mean-variance portfolio selection problem. Section 7 concludes.

2 Tail optimality of a control plan

In this section we introduce the problem's formulation and provide the notation which will be used throughout the paper. We define the notion of tail optimality of a control plan for an intertemporal optimization problem and provide two definitions. The first one is for one single optimization problem, and we will refer to it as *local tail optimality* of the control plan for the problem at hand. The second one applies to a family of problems, and we will refer to it as *global tail optimality* of a control plan for the family of problems considered. We recall that these concepts are not new, and are strongly related to the Bellman's optimality principle for optimization problems. It is, however, important to review them carefully for two reasons. First, we introduce the reader to the perspective of tail optimality (or lack thereof) that is a new way to frame and interpret existing ideas. Second, and more importantly, the very definition of preferences consistency (provided in Section 3) is based upon the notion of tail optimality that must be therefore defined accurately at an earlier stage.

2.1 Setting

To start fixing ideas, let us consider the following framework:

- the time frame over which the optimization is done is fixed and is $[t_0, T]$;
- the wealth² $X_s \in \mathbb{R}$ of the agent evolves according to the controlled stochastic differential equation (SDE):³

$$\begin{aligned} dX_s &= \mu(s, X_s, u_s)ds + \sigma(s, X_s, u_s)dW_s \\ X_{t_0} &= x_0 \end{aligned} \tag{1}$$

where W_s is a standard Brownian motion defined on a complete filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbf{P})$, with $\mathcal{F}_t = \sigma\{W(s) : s \leq t\}$;

- at any time $s \in [t_0, T]$ the agent can choose the *control action* u_s according to some criterion; the set of all control actions $\{u_s\}_{s \in [t_0, T]}$ is said to be a *control plan*; adopting the terminology common in economics, given the control plan $\{u_s\}_{s \in [t_0, T]}$ defined over the whole time interval $[t_0, T]$, its restriction on $[t, T]$ for $t \in (t_0, T)$, $\{u_s\}_{s \in [t, T]}$, is said to be the *continuation plan at* t ; we assume that $\{u_s\}_{s \in [t_0, T]}$ is a Markov control process, i.e., it is a deterministic function of time s and the wealth at that time: $u_s(\omega) = u(s, X_s(\omega))$ for some deterministic function $u : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$, also called feedback control map;
- \mathcal{U} is the set of admissible strategies, defined as the set of \mathbb{R} -valued stochastic processes $u = \{u_s\}_{s \in [t_0, T]}$ that are Markov control processes, \mathcal{F}_s -adapted and s.t. the SDE (1) has a unique strong solution.⁴

2.2 Optimal problem and preferences

It is essential to highlight that the criterion selected by the agent in the optimization problem represents the *preferences* of the agent and is typically given by the combination of different utility

²For simplicity, we here refer to wealth, but the controlled state equation X_s can be any quantity of interest to the agent.

³Sometimes in the paper we will use the notation X_s^u to denote the value of the state variable at time s under control u . Whenever possible, we shall try to keep the notation as simple as possible.

⁴For simplicity, we assume that the set of admissible controls does not change with time and wealth, i.e., $\mathcal{U}(t, x) = \mathcal{U}$ for every $(t, x) \in [t_0, T] \times \mathbb{R}$.

functions. In particular, putting ourselves in the setting introduced by Björk & Murgoci (2010), we shall assume that the agent wants to solve the following optimization problem:

$$\begin{aligned} & \text{Problem } \mathcal{P}_{t_0, x_0} : \\ \sup_{u \in \mathcal{U}} J(t_0, x_0, u) &= \sup_{u \in \mathcal{U}} \left\{ \mathbb{E}_{t_0, x_0} \left[\int_{t_0}^T U^1(s, X_s, u_s) ds + U^2(X_T) \right] + U^3 \left[\mathbb{E}_{t_0, x_0}(X_T) \right] \right\} \end{aligned} \quad (2)$$

where $U^1(\cdot)$, $U^2(\cdot)$ and $U^3(\cdot)$ are, possibly non-linear, utility functions that identify the agent's preferences.

Remark 1. To be more precise, we could denote by $\mathcal{P}_{t_0, x_0}^{\{U^1, U^2, U^3\}}$ or by $J^{\{U^1, U^2, U^3\}}(t_0, x_0, u)$ the performance criterion of the agent, to stress the crucial role played by the utility functions in the identification of the agent's preferences. For notational convenience, in the following, we will simply refer to \mathcal{P}_{t_0, x_0} or $J(t_0, x_0, u)$.

Problem \mathcal{P}_{t_0, x_0} belongs to the more general family of optimization problems

$$\{\mathcal{P}_{t,x}\}_{(t,x) \in [t_0, T] \times \mathbb{R}},$$

where

$$\begin{aligned} & \text{Problem } \mathcal{P}_{t,x} : \\ \sup_{u \in \mathcal{U}} J(t, x, u) &= \sup_{u \in \mathcal{U}} \left\{ \mathbb{E}_{t,x} \left[\int_t^T U^1(s, X_s, u_s) ds + U^2(X_T) \right] + U^3 \left[\mathbb{E}_{t,x}(X_T) \right] \right\} \end{aligned} \quad (3)$$

for $(t, x) \in [t_0, T] \times \mathbb{R}$. In the following, Problem $\mathcal{P}_{t,x}$ will be called the *continuation problem* starting from $t \geq t_0$.

2.3 Linear vs non-linear problems

In line with Björk & Murgoci (2010) we recall that the nature of Problem \mathcal{P}_{t_0, x_0} strongly depends on the utility function $U^3(\cdot)$, and there are two possible cases:

1. if $U^3(\cdot)$ is a linear function,⁵ then it can be incorporated into $U^2(\cdot)$ and we have a *linear*

⁵Clearly, the problem is linear also if $U^3(\cdot)$ is an affine function. With an abuse of terminology, in the following we shall often use the word *linear* to refer to an affine function and the word *non-linear* to refer to a non-affine one.

problem:

$$\begin{aligned} & \text{Problem } \mathcal{P}_{t_0, x_0}^L : \\ & \sup_{u \in \mathcal{U}} J^L(t_0, x_0, u) = \sup_{u \in \mathcal{U}} \left\{ \mathbb{E}_{t_0, x_0} \left[\int_{t_0}^T U^1(s, X_s, u_s) ds + U^2(X_T) \right] \right\}; \end{aligned} \quad (4)$$

2. if $U^3(\cdot)$ is a non-linear function, then it cannot be incorporated into $U^2(\cdot)$ and we have a *non-linear* problem:

$$\begin{aligned} & \text{Problem } \mathcal{P}_{t_0, x_0}^{NL} : \\ & \sup_{u \in \mathcal{U}} J^{NL}(t_0, x_0, u) = \sup_{u \in \mathcal{U}} \left\{ \mathbb{E}_{t_0, x_0} \left[\int_{t_0}^T U^1(s, X_s, u_s) ds + U^2(X_T) \right] + U^3[\mathbb{E}_{t_0, x_0}(X_T)] \right\}. \end{aligned} \quad (5)$$

According to Björk & Murgoci (2010) another possible source of non-linearity is the presence of t_0 or x_0 in the running utility $U^1(\cdot)$ or in the terminal utility $U^2(\cdot)$. We disregard this case due to space constraints, and in the remaining of the paper we shall assume that neither $U^1(\cdot)$ nor $U^2(\cdot)$ depend on the initial point (t_0, x_0) .

2.4 Local and global tail optimality

We are now ready to provide the definition of local tail optimality of a control plan.

Definition 2.1 (Local tail optimality). *Given $(t, x) \in [t_0, T] \times \mathbb{R}$ and the stochastic optimal control problem $\mathcal{P}_{t,x}$ as in (3), we say that the control plan (if it exists)*

$$u_{t,x}^* : [t, T] \times \mathbb{R} \rightarrow \mathbb{R} \quad (6)$$

is locally tail-optimal at t for $\mathcal{P}_{t,x}$ if

$$J(t, x, u_{t,x}^*) = \sup_{u \in \mathcal{U}} J(t, x, u). \quad (7)$$

The word “tail” of Definition 2.1 reflects the fact that in order to reach the supremum of the performance criterion $J(t, x, u)$ it is necessary that the control plan $u_{t,x}^*(\cdot)$ is played from t until T , meaning that at each time $s \in [t, T]$ with wealth $X_s \in \mathbb{R}$ the agent plays the control action

$u_{t,x}^*(s, X_s)$. For the definition of local tail optimality at time t , what has happened *before* t has no importance, but the control played *after* t must be determined by the optimal control plan u^* . Intuitively, the plan u^* is optimal in the right subinterval $[t, T]$ (after t), which can be seen as the right tail of the interval $[t_0, T]$.

Remark 2. It is important to stress that the optimal control plan $u_{t,x}^*(\cdot)$ is a function of time $s \geq t$ and wealth $y \in \mathbb{R}$, but it might depend also on the given initial point (t, x) (in Section 6 we will examine an example –the mean-variance portfolio selection problem– in which this happens). Indeed, two optimal control plans associated to two different initial time-wealth points are generally different on the same domain, i.e., if $(t, x) \neq (t_1, x_1)$ then, in general,

$$u_{t,x}^*(s, y) \neq u_{t_1, x_1}^*(s, y) \quad \text{for } (s, y) \in [t \wedge t_1, T] \times \mathbb{R}. \quad (8)$$

However, in some cases the above inequality holds as an equality for every couple of time-wealth points, and in this case the stronger feature of global tail optimality, defined below, holds.

Definition 2.2 (Global tail optimality). *Given the stochastic optimal control problem \mathcal{P}_{t_0, x_0} as in (2), we say that the control plan*

$$u_{t_0, x_0}^* : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$$

is globally tail-optimal over $[t_0, T]$ for Problem \mathcal{P}_{t_0, x_0} if for every $t \in [t_0, T]$ and every $x \in \mathbb{R}$ the restriction of u_{t_0, x_0}^ to $[t, T] \times \mathbb{R}$*

$$u_{t_0, x_0}^* : [t, T] \times \mathbb{R} \rightarrow \mathbb{R}$$

is locally tail-optimal at t for $\mathcal{P}_{t,x}$, where $\mathcal{P}_{t,x}$ is as in (3).

The following existence issue can arise:

Q1 Does a control plan exist that is globally tail-optimal for some stochastic optimal control problem?

The answer is positive, when considering the special case of linear stochastic optimal control problems.

2.5 Special case: tail optimality for linear problems

If $U^3(\cdot)$ is linear and the problem is linear as in (4), then dynamic programming is applicable. By dynamic programming, in order to approach Problem \mathcal{P}_{t_0, x_0}^L one should:⁶

- consider the more general problem to be solved at time t with wealth x , Problem $\mathcal{P}_{t, x}^L$, that is obtained by replacing t_0 with t and x_0 with x in Problem \mathcal{P}_{t_0, x_0}^L and Equation (4):

$$\begin{aligned} &\text{Problem } \mathcal{P}_{t, x}^L : \\ &\sup_{u \in \mathcal{U}} J^L(t, x, u) = \sup_{u \in \mathcal{U}} \mathbb{E}_{t, x} \left[\int_t^T U^1(s, X_s, u_s) ds + U^2(X_T) \right] \end{aligned} \quad (9)$$

for $(t, x) \in [t_0, T] \times \mathbb{R}$;

- write and solve (if possible) the associated Hamilton-Jacobi-Bellman (HJB) equation to find the value function

$$V(t, x) = \sup_{u \in \mathcal{U}} J^L(t, x, u),$$

and the optimal control law

$$u_{t, x}^* : [t, T] \times \mathbb{R} \rightarrow \mathbb{R}$$

as the maximizing control of the HJB equation.

Once Problem $\mathcal{P}_{t, x}^L$ is solved, the initial problem \mathcal{P}_{t_0, x_0}^L is also retrieved as a special case by replacing (t, x) with (t_0, x_0) . In this standard case, the Bellman's optimality principle holds: quite remarkably, and contrary to what observed in Remark 2 for the general case, the optimal control plan u_{t_0, x_0}^* is optimal not only on $[t_0, T]$ but also on every subinterval $[\tau, T]$ with $\tau > t_0$ for the continuation problem starting from τ , $\mathcal{P}_{\tau, x_\tau}^L$. This is the well-known Bellman's optimality principle (see Bertsekas, 2012). This means that the optimal strategy for the continuation problem $\mathcal{P}_{\tau, x_\tau}^L$ at time τ with current wealth x_τ coincides with the *restriction* on $[\tau, T]$ of the optimal strategy found at initial time t_0 (i.e., with the continuation strategy at τ):

$$\operatorname{argmax}_{u \in \mathcal{U}} J^L(\tau, x_\tau, u) = \{u_{\tau, x_\tau}^*(s, y)\}_{(s, y) \in [\tau, T] \times \mathbb{R}} = \{u_{t_0, x_0}^*(s, y)\}_{(s, y) \in [\tau, T] \times \mathbb{R}}. \quad (10)$$

⁶For brevity, we omit all details and refer the interested reader to classical sources, e.g., Yong & Zhou (1999), Björk (1998).

Because this happens for every $\tau \in [t_0, T]$ and every $x_\tau \in \mathbb{R}$, the optimal control law is the same regardless of the initial time-wealth point, and with some abuse of notation,⁷ we shall simply denote it by $u^*(s, y)$:

$$u_{\tau, x_\tau}^*(s, y) = u_{t_0, x_0}^*(s, y) = u^*(s, y). \quad (11)$$

This equality shows that for a linear problem the optimal control plan *does not* depend on the initial time-wealth point (t, x) : it is simply a function of time $s \in [t_0, T]$ and wealth $y \in \mathbb{R}$. Therefore, the infinitely many optimal control plans $u_{t,x}^*(s, y)$ of the continuation problems $P_{t,x}^L$ can be identified by the infinitely many *restrictions* of the control plan $u^* : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ for each $(t, x) \in [t_0, T] \times \mathbb{R}$

$$u_{t,x}^*(s, y) = u^*(s, y) \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}.$$

Given Definition 2.2 and the validity of the Bellman's optimality principle for linear problems, we can now state the following known result, that answers question **Q1**.

Proposition 2.3. *Given the linear problem \mathcal{P}_{t_0, x_0}^L as in (4), the optimal control plan*

$$u^* : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$$

found via dynamic programming is globally tail-optimal over $[t_0, T]$ for Problem \mathcal{P}_{t_0, x_0}^L .

This proposition is nothing but a rephrasing of the Bellman's optimality principle (see Bertsekas, 2012) for stochastic optimal control problems in continuous time.

3 Preferences consistency of an agent

While the notion of global tail optimality is not new and essentially coincides with the Bellman's principle, the notion of preferences consistency for an intertemporal optimization problem deserves special care. As mentioned in Section 1, the link between the agent's preferences and a time-consistent behaviour have been addressed in some papers in decision theory economic literature:

⁷The maximum domain of the control plan $u^*(s, y)$ is $[t_0, T] \times \mathbb{R}$, whereas the domain of the optimal control plan of the continuation problem $\mathcal{P}_{\tau, x_\tau}^L$ is restricted to $[\tau, T] \times \mathbb{R}$.

see Johnsen & Donaldson (1985), Epstein & Le Breton (1993), Chew & Epstein (1990). The mentioned works consider choices among lotteries in a discrete framework, whereas in this paper we consider an intertemporal optimization problem in continuous time. In this setup, we provide a novel definition of preferences consistency.

As in Section 2.2, imagine an agent who sits at initial time t_0 with wealth x_0 , and optimizes over the time horizon $[t_0, T]$ with preferences identified by the utility functions U^1, U^2 and U^3 (see also Remark 1). She then wants to solve the initial optimization problem \mathcal{P}_{t_0, x_0} as in (2). In this paper we disregard the case of time-varying preferences: we here make the assumption that she does not change her preferences over time and that her preferences are represented by U^1, U^2 and U^3 also over $[t, T]$ for any $t \in (t_0, T]$. Therefore, no matter what happens between t_0 and $t > t_0$, we assume that the agent at time $t > t_0$ with wealth x_t will be solving the continuation problem \mathcal{P}_{t, x_t} as in (3), because of unchanged preferences. Intuitively, if this happens, we will say that the agent is preferences-consistent.

In particular, we will say that the agent who was solving Problem \mathcal{P}_{t_0, x_0} driven by $\{U^1, U^2, U^3\}$ at initial time t_0 is preferences-consistent at time $t > t_0$ if the action that she plays at time t optimizes the continuation problem \mathcal{P}_{t, x_t} , still driven by $\{U^1, U^2, U^3\}$. It is evident that the notion of preferences consistency at time t needs a reference point, that consists in the initial preferences $\{U^1, U^2, U^3\}$ at time t_0 .

The notion of local preferences consistency is formalized by the following definition.

Definition 3.1 (Local preferences consistency). *An agent whose initial preferences at time t_0 are described by the optimization problem \mathcal{P}_{t_0, x_0} as in (2) is **locally preferences-consistent at t with respect to \mathcal{P}_{t_0, x_0}** , where $t > t_0$ is fixed, if for every $x \in \mathbb{R}$ the current control action that she chooses at time t with wealth x coincides with the first control action of the tail-optimal control plan of the continuation problem $\mathcal{P}_{t, x}$ as in (3), i.e., if at time t with wealth x she chooses $u_{t, x}^*(t, x)$, where $u_{t, x}^*(s, y)$ (for $(s, y) \in [t, T] \times \mathbb{R}$) is the tail-optimal control plan for $\mathcal{P}_{t, x}$:*

$$J(t, x, u_{t, x}^*) = \sup_{u \in \mathcal{U}} J(t, x, u).$$

In other words, being a locally preferences-consistent agent at t means being *instantaneously* optimal for the continuation problem $\mathcal{P}_{t, x}$. Notice that the local preferences consistency at time t implies only that for every wealth x the agent plays the optimal control action for $\mathcal{P}_{t, x}$ at time t ,

but does not mean that she will continue to play the optimal plan $u_{t,x}^*(s, \cdot)$ also for $s > t$.

Quite naturally, if an agent is locally preferences-consistent at t for every t in a given interval, she is globally preferences-consistent over the interval. The notion of global preferences consistency is formalized by the following definition.

Definition 3.2 (Global preferences consistency). *An agent whose initial preferences at time t_0 are described by the optimization problem \mathcal{P}_{t_0, x_0} as in (2) is **globally preferences-consistent over $[t_0, T]$ with respect to \mathcal{P}_{t_0, x_0}** if she is locally preferences-consistent at t with respect to \mathcal{P}_{t_0, x_0} for every $t \in [t_0, T]$.*

Definition 3.2 has a strong connection with Definition 2 of dynamical optimality given by Pedersen & Peskir (2017) in the case of mean-variance preferences. Roughly speaking, according to their definition, a control is dynamically optimal if, for every fixed t and x , it coincides with the first control of the optimal strategy at (t, x) . While Pedersen & Peskir (2017) focus on the control strategy and its instantaneous optimality, we here stress the importance of the consistency of the agent to her *initial preferences*, that are described by the original optimization problem \mathcal{P}_{t_0, x_0} . The link between the two definitions will become clear in Section 4, where we show that the dynamically optimal agent is globally preferences-consistent.

As in the case of global tail optimality, the following existence issue can arise:

Q2 Does an agent exist who is globally preferences-consistent over a time interval $[t_0, T]$ with respect to some initial preferences?

The answer is again positive, by considering again the special case of the agent of a linear stochastic optimal control problem.

3.1 Special case: preferences consistency for linear problems

Let us assume that the original preferences of the agent are identified by the linear problem \mathcal{P}_{t_0, x_0}^L as in (4). Then, the agent can apply dynamic programming as explained in Section 2.5 to solve it and find the optimal control plan $u_{t_0, x_0}^* : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$. Imagine that the agent plays the optimal control $u_{t_0, x_0}^*(s, X_s)$ over $[t_0, T]$. Then, because of the Bellman's optimality principle and equation (10), it turns out that at time τ with wealth x_τ she plays exactly the first control action

of the tail-optimal control plan of the continuation problem $\mathcal{P}_{\tau, x_\tau}^L$. This means that the agent is preferences-consistent at time $\tau > t_0$ with respect to her original preferences identified by \mathcal{P}_{t_0, x_0}^L . Because this happens at every time $\tau \in [t_0, T]$ we conclude that the agent is globally preferences-consistent over $[t_0, T]$ with respect to the original problem \mathcal{P}_{t_0, x_0}^L .

This result is formalized by the following proposition, that answers question **Q2**.

Proposition 3.3. *Let the preferences of an agent be identified by the linear problem \mathcal{P}_{t_0, x_0}^L . If the agent plays the optimal control plan $u_{t_0, x_0}^*(s, X_s) = u^*(s, X_s)$ over $[t_0, T]$, where $u^*(\cdot)$ is found via dynamic programming, then she is globally preferences-consistent over $[t_0, T]$ with respect to \mathcal{P}_{t_0, x_0}^L .*

4 Non-linear problems

Propositions 2.3 and 3.3 show that in the ideal world of linear problems where dynamic programming can be applied, the two desirable features of global tail optimality of the control plan and global preferences consistency of the agent take place simultaneously. The coexistence of global tail optimality and global preferences consistency is a consequence of the validity of the Bellman's principle and the applicability of dynamic programming.

The situation becomes more complicated in the case of non-linear problems, when the bequest function includes also a non-linear function of expected final wealth.⁸ In this case, the non-applicability of dynamic programming and the non-validity of the Bellman's principle prevent the simultaneous occurrence of global tail optimality and global preferences consistency.

Let us suppose that an agent wants to solve the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (5), where $U^3(\cdot)$ is a non-linear utility function.

It is well known (see Björk & Murgoci, 2010) that the presence of the non-linear term $U^3 [\mathbb{E}_{t_0, x_0}(X_T)]$ prevents the straight use of dynamic programming. According to the current literature, this problem gives rise to time inconsistency, and there are different approaches to deal with it. We will see that none of the existing approaches keeps simultaneously both properties of global tail optimality and global preferences consistency. Nevertheless, it is possible to analyze them and see what are the properties characterizing each of them.

⁸In a standard expected utility linear problem (4) the bequest function is given by U^2 only. In the general problem (2) the bequest function includes also U^3 .

The three approaches currently available for the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ are:

1. precommitment approach;
2. dynamic optimality approach;
3. consistent planning (also known as game theoretical, or Nash-equilibrium) approach.

4.1 Precommitment approach

To solve the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ with the precommitment approach, one fixes the initial point (t_0, x_0) and finds, if it exists, the control law \hat{u} that maximizes only $J^{NL}(t_0, x_0, u)$, i.e., the precommitment strategy. This is formalized by the following definition.

Definition 4.1. *Given the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (5), the strategy \hat{u} that maximizes $J^{NL}(t_0, x_0, u)$, i.e., the control plan*

$$\hat{u}_{t_0, x_0} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R} \tag{12}$$

such that

$$J^{NL}(t_0, x_0, \hat{u}_{t_0, x_0}) = \sup_{u \in \mathcal{U}} J^{NL}(t_0, x_0, u)$$

if it exists, is called the precommitment strategy for $\mathcal{P}_{t_0, x_0}^{NL}$.

Because in these kinds of problems dynamic programming cannot be applied and the Bellman's principle does not hold, by adopting \hat{u}_{t_0, x_0} one disregards the fact that at a later point in time $\tau \in (t_0, T]$ with wealth x_τ the continuation plan $\hat{u}_{t_0, x_0}(s, y)$ at τ (for $(s, y) \in [\tau, T] \times \mathbb{R}$), is, in general, *not* optimal for the continuation criterion $J^{NL}(\tau, x_\tau, u)$. In other words,

$$\operatorname{argmax}_{u \in \mathcal{U}} J^{NL}(\tau, x_\tau, u) = \{\hat{u}_{\tau, x_\tau}(s, y)\}_{(s, y) \in [\tau, T] \times \mathbb{R}} \neq \{\hat{u}_{t_0, x_0}(s, y)\}_{(s, y) \in [\tau, T] \times \mathbb{R}}, \tag{13}$$

while there would be equality with validity of the Bellman's principle, see Equation (10). In other words, the precommitment strategy for $\mathcal{P}_{t_0, x_0}^{NL}$ (12) *depends essentially on the initial point (t_0, x_0) .*

This is the reason why the strategy is named precommitment strategy: the precommitted agent standing at time t_0 should “precommit” herself to follow the strategy $\hat{u}_{t_0, x_0}(s, y)$ from t_0 to T , even

if she knows that at later point in time τ she will still be solving the original problem $\mathcal{P}_{t_0, x_0}^{NL}$, but *not* the continuation problem $\mathcal{P}_{\tau, x_\tau}^{NL}$. Indeed, due to (13), the control action that the precommitted agent plays at every time $\tau > t_0$ is, in general, not equal to the first optimal control action for the continuation problem $\mathcal{P}_{\tau, x_\tau}^{NL}$. Therefore, the precommitted agent is locally preferences-consistent at time t_0 with respect to $\mathcal{P}_{t_0, x_0}^{NL}$ (because at time t_0 she plays $\hat{u}_{t_0, x_0}(t_0, x_0)$ that is the first action of the optimal plan for $\mathcal{P}_{t_0, x_0}^{NL}$), but, in general, is *not* preferences-consistent at any time $\tau > t_0$ with respect to $\mathcal{P}_{t_0, x_0}^{NL}$. Supported by Definition 3.1, we can formalize this result in the next proposition.

Proposition 4.2. *Let the preferences of an agent be identified by the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (5), and let us assume that there exists the precommitment strategy \hat{u}_{t_0, x_0} for Problem $\mathcal{P}_{t_0, x_0}^{NL}$. If the agent plays the precommitment strategy $\hat{u}_{t_0, x_0}(s, X_s)$ over $[t_0, T]$, then she is locally preferences-consistent at t_0 with respect to $\mathcal{P}_{t_0, x_0}^{NL}$.*

Let us now turn to the feature of local and global tail optimality of the control plan. By Definitions 2.1 and 4.1, it is clear that, if it exists, the precommitment strategy is locally tail-optimal at initial time t_0 for Problem $\mathcal{P}_{t_0, x_0}^{NL}$.

Proposition 4.3. *Given the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (5), the precommitment strategy \hat{u}_{t_0, x_0} given by (12), if it exists, is locally tail-optimal at t_0 for $\mathcal{P}_{t_0, x_0}^{NL}$.*

Local tail optimality at initial time t_0 of the control plan and local preferences consistency at initial time t_0 of the precommitted agent is all that the precommitment approach can offer. In general, for a non-linear problem the precommitment strategy \hat{u}_{t_0, x_0} is not globally tail-optimal and the precommitted agent is not globally preferences-consistent. While a proof of this result in general is far from trivial, this can be easily proven in the important case of the mean-variance portfolio selection problem, see Section 6.

Clearly, the precommitment strategy is the best strategy standing at time t_0 with the aim of optimizing $J^{NL}(t_0, x_0, u)$, see also Vigna (2020). The problem of precommitment is about preferences inconsistency after t_0 : the precommitted agent only cares about initial time t_0 and final time T , disregarding that she will be preferences-inconsistent at any time $t \in (t_0, T)$. In other words, the precommitment approach is closer in spirit to the single-period Markowitz framework than to the dynamic continuous-time setup: only t_0 and T matter, what happens at any time $t \in (t_0, T)$ does not matter. The interval (t_0, T) goes into a black box and the agent is consistent

to her own preferences only at initial time t_0 . In this respect, the name “static” given by some authors to identify the precommitment strategy (Pedersen & Peskir, 2017) or the optimization problem as defined in (t_0, x_0) only (Karnam, Ma & Zhang, 2017), could not be more appropriate.

4.2 Dynamic optimality approach

We illustrate the construction of the dynamically optimal strategy introduced by Pedersen & Peskir (2017) for a non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ in 4 steps.⁹

Step 1. A family of non-linear problems $\{\mathcal{P}_{t,x}^{NL}\}_{(t,x) \in [t_0, T] \times \mathbb{R}}$, with $\mathcal{P}_{t,x}^{NL}$ as in (3), is given.

Step 2. Assume that for the initial time-wealth point (t_0, x_0) the precommitment strategy \hat{u}_{t_0, x_0} maximizing the criterion $J^{NL}(t_0, x_0, u)$ exists and is given by (12).

Step 3. Define the new control plan

$$\tilde{u}(s, y) = \hat{u}_{s,y}(s, y), \quad \text{for } (s, y) \in [t_0, T] \times \mathbb{R}, \quad (14)$$

where the right hand side of (14) is obtained by replacing t_0 with s and x_0 with y in the function (12).

Step 4. The strategy $\tilde{u} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$, is called the *dynamically optimal strategy*.¹⁰

Remark 3. Unlike the precommitment strategy (12), the dynamically optimal strategy (14) does not depend on the initial time-wealth point: it is a simple function of time $s \in [t_0, T]$ and wealth $y \in \mathbb{R}$. In this respect, it looks similar to the optimal control plan of a linear optimization problem u^* as in (11). For this reason, the dynamically optimal strategy is known to be time-consistent (see Pedersen & Peskir, 2017).

Let us analyze the preferences consistency of the dynamically optimal agent.

By construction, at generic time $t \in [t_0, T]$ with wealth x the dynamically optimal agent faces the problem $\mathcal{P}_{t,x}^{NL}$ and solves it with the precommitment approach for $\mathcal{P}_{t,x}^{NL}$, as if (t, x) was the initial

⁹Pedersen & Peskir (2017) introduce the dynamically optimal strategy in order to solve the mean-variance portfolio selection problem. Clearly, their approach can be extended to any non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$.

¹⁰As a practical example of the construction of the dynamically optimal policy from the precommitment one as illustrated by Equation (14), we refer to the mean-variance portfolio selection case in Section 6: the dynamically optimal policy (25) is derived by replacing (t_0, x_0) with (s, y) in the precommitment policy (24).

time-wealth point. In fact, she plays the *first* control action of the precommitment strategy for $\mathcal{P}_{t,x}^{NL}$, for she plays $\hat{u}_{t,x}(t, x)$. Because the initial preferences of the agent are given by Problem $\mathcal{P}_{t_0,x_0}^{NL}$ and because at time t she plays the first control action of the optimal strategy for the continuation problem $\mathcal{P}_{t,x}^{NL}$, by definition she is locally preferences-consistent at time t with respect to $\mathcal{P}_{t_0,x_0}^{NL}$. This happens for every $t \in [t_0, T]$, and therefore she is globally preferences-consistent over $[t_0, T]$ with respect to $\mathcal{P}_{t_0,x_0}^{NL}$.

This result is formalized by the next proposition, that, together with Proposition 3.3, also answers question **Q2**.

Proposition 4.4. *Let the preferences of an agent be identified by the non-linear problem $\mathcal{P}_{t_0,x_0}^{NL}$ as in (5), and let us assume that there exists the precommitment strategy \hat{u}_{t_0,x_0} for $\mathcal{P}_{t_0,x_0}^{NL}$. If the agent plays the dynamically optimal strategy $\tilde{u}(s, X_s)$ over $[t_0, T]$, then she is globally preferences-consistent over $[t_0, T]$ with respect to $\mathcal{P}_{t_0,x_0}^{NL}$.*

Remark 4. Regarding the relationship between precommitment approach and dynamically optimal approach, we see that by construction at each $t \in [t_0, T]$ with wealth x the control action of the dynamically optimal strategy $\tilde{u}(t, x)$ coincides with the *first* control action of the precommitment strategy for $\mathcal{P}_{t,x}^{NL}$, i.e., it coincides with the first control action of the control plan $\hat{u}_{t,x}(s, y)$ $((s, y) \in [t, T] \times \mathbb{R})$ selected by the agent who wants to solve $\mathcal{P}_{t,x}^{NL}$ with the precommitment approach. But it deviates from it immediately after, at time $t' = t + dt$, because at time t' with wealth x' the dynamically optimal strategy coincides with the first control action of the precommitted strategy for $\mathcal{P}_{t',x'}^{NL}$. Therefore, the dynamically optimal agent can be seen as the *continuous reincarnation* of the precommitted agent. Moreover, even if this strategy has been formalized and deeply studied by Pedersen & Peskir (2017), the dynamically optimal agent is similar to the continuous version of the naive agent described by Pollak (1968). We notice that the dynamically optimal naive agent is the only one to be globally preferences-consistent in the presence of a non-linear optimization problem. Moreover, the dynamically optimal approach has strong similarities with the receding horizon procedure or the model predictive control (see Powell, 2011), that are well established methods of repeated optimization over a rolling horizon for engineering optimization problems with an infinite time horizon (although in the problem considered in this paper the time interval over which the optimization is done shrinks when time passes, while it remains fixed in those problems).

Let us now turn to the question of tail optimality of the dynamically optimal strategy.

From Definition 2.1, we see that a control plan is locally tail-optimal at time t for an optimization problem if, whenever played from t to the time horizon T , it reaches the supremum of the performance criterion. The dynamically optimal strategy is a collection of infinitely many *first* optimal control actions for infinitely many problems. As such, there is no problem for which it is locally tail-optimal at time t . Indeed, as Pedersen & Peskir (2017) notice, the control plan \tilde{u} is instantaneously optimal at each $t \in [t_0, T]$, so it is instantaneously optimal for infinitely many non-linear problems. Therefore, unlike the precommitment strategy that is locally tail-optimal at the initial time point t_0 —and only at t_0 — for $\mathcal{P}_{t_0, x_0}^{NL}$, there exists no such $t \in [t_0, T]$ that makes the dynamically optimal strategy locally tail-optimal at t for $\mathcal{P}_{t, x}^{NL}$.

4.3 Consistent planning, game theoretical, Nash equilibrium approach

According to the consistent planning approach, in order to solve the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$, one should choose “the best plan among those that he will actually follow”. The construction of this strategy is based on the game theoretic interpretation that to each point in time t is associated a player who can choose the control at time t . At time $s > t$ there is another player who chooses the control at time s . The key of this approach is to search for a Nash subgame perfect equilibrium among the continuum of players $[t_0, T]$. A strategy \bar{u} is an equilibrium strategy if, given that all players in $(t, T]$ will play \bar{u} then also player t finds it optimal to play \bar{u} . The equilibrium strategy is found by solving an extended Hamilton-Jacobi-Bellman equation for the value function, see Björk & Murgoci (2010). Like the optimal control law of a linear problem, the Nash equilibrium strategy \bar{u} does not depend on the initial time-wealth point and is a function of time s and wealth y only:

$$\bar{u} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}. \quad (15)$$

This is the reason why it is known to be time-consistent.

Notably, Björk & Murgoci (2010) also prove that to each time-inconsistent non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ it is possible to associate a standard time-consistent linear problem $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$ such that (i) the optimal value function of the linear problem is equal to the equilibrium value function of the time-inconsistent non-linear problem; (ii) the optimal control law of the linear problem is equal to the equilibrium strategy of the time-inconsistent non-linear problem, see Björk & Murgoci (2010),

Proposition 5.1.

This remarkable result implies that there exist utility functions $U^4(\cdot)$ and $U^5(\cdot)$ (not necessarily easy to find) such that the Nash equilibrium strategy \bar{u} associated to the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ coincides with the optimal control law found via dynamic programming solution to the linear problem

$$\begin{aligned} & \text{Problem } \mathcal{P}_{t_0, x_0}^{L-ass-NL} : \\ & \sup_{u \in \mathcal{U}} \mathbb{E}_{t_0, x_0} \left[\int_{t_0}^T U^4(s, X_s, u_s) ds + U^5(X_T) \right]. \end{aligned} \tag{16}$$

For instance, in the case of the mean-variance preferences, where $U^1(x) = 0$, $U^2(x) = x - \alpha x^2$ and $U^3(x) = \alpha x^2$, it is easy to show that $U^4(x) = 0$ while $U^5(\cdot)$ is the exponential utility function, see Section 6.

We use this important result to analyze the consistent planning approach under the two criteria of tail optimality of the control plan and preferences consistency of the agent.

By noting that the Nash equilibrium control plan \bar{u} coincides with the optimal control plan of the associated linear problem $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$, using Proposition 2.3 it is straightforward to conclude that the Nash equilibrium control plan is globally tail-optimal for Problem $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$, which is the result of the next proposition.

Proposition 4.5. *Given the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (5), and given the linear problem $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$ as in (16) associated to the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ in the sense of Proposition 5.1 of Björk & Murgoci (2010), the control plan*

$$\bar{u} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R} \tag{17}$$

that is the equilibrium strategy of $\mathcal{P}_{t_0, x_0}^{NL}$ found via the consistent planning approach, is globally tail-optimal over $[t_0, T]$ for the linear problem $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$.

Similarly, using Proposition 3.3 it is immediate to conclude that the agent who plays the Nash equilibrium strategy over $[t_0, T]$ is globally preferences-consistent over $[t_0, T]$ with respect to Problem $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$, which is the result of the next proposition.

Proposition 4.6. *Let the preferences of an agent be identified by the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (5), and let us assume that there exists the Nash equilibrium strategy \bar{u} for $\mathcal{P}_{t_0, x_0}^{NL}$. If the agent*

plays $\bar{u}(s, X_s)$ over $[t_0, T]$, then she is globally preferences-consistent over $[t_0, T]$ with respect to the linear problem $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$ associated to $\mathcal{P}_{t_0, x_0}^{NL}$ in the sense of Proposition 5.1 of Björk & Murgoci (2010).

To sum up, the consistent planning agent is globally preferences-consistent with respect to the associated linear problem, and she plays a strategy that is globally tail-optimal for the associated linear problem. In general, she is not preferences-consistent to her original non-linear preferences and the plan that she plays is not tail-optimal for the original non-linear problem.

Regarding their surprising result in Proposition 5.1, Björk & Murgoci (2010) comment that there is no gain by enlarging the class of consumer behaviour to time-inconsistent preferences, because every time-inconsistent strategy can be replicated by some time-consistent utility function. We comment on this result from a different angle. For a non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ the Nash equilibrium approach is equivalent to applying the solution to the associated linear problem (16). This means that in order to be time-consistent in the consistent-planning sense, the agent has to choose a different objective functional, in other words, different preferences, see also Forsyth (2020). For the mean-variance problem, the agent who chooses the Nash-equilibrium approach applies a strategy that is optimal according to a *different* criterion than the mean-variance one, namely the exponential preferences. The price to be paid in order to be time-consistent in the consistent-planning sense consists in changing preferences. As mentioned in Section 1, the preferences inconsistency of the Nash equilibrium agent was already observed by Chew & Epstein (1990) who write: “*The equilibrium represents a time-consistent form of behaviour, even though preferences are not intertemporally consistent*”.

4.4 Discussion

In the previous sections we have showed that each of the three approaches currently available for a non-linear dynamic optimization problem $\mathcal{P}_{t_0, x_0}^{NL}$ presents only some of the two desirable properties of tail optimality and preferences consistency.

By Propositions 4.2 and 4.3, the precommitted agent is locally preferences-consistent at time t_0 with respect to her initial preferences given by Problem $\mathcal{P}_{t_0, x_0}^{NL}$, and the precommitment strategy is locally tail-optimal at t_0 for Problem $\mathcal{P}_{t_0, x_0}^{NL}$. Therefore, the precommitment approach keeps local tail optimality and local preferences consistency at initial time t_0 .

By Proposition 4.4, the dynamically optimal agent is globally preferences-consistent with respect to her initial preferences given by Problem $\mathcal{P}_{t_0, x_0}^{NL}$.

By Propositions 4.5 and 4.6, the Nash equilibrium strategy is globally tail-optimal for the linear problem that is associated to Problem $\mathcal{P}_{t_0, x_0}^{NL}$, and the Nash equilibrium agent is globally preferences-consistent with respect to the linear problem that is associated to Problem $\mathcal{P}_{t_0, x_0}^{NL}$.

In general, the precommitment strategy is never locally tail-optimal at time $t > t_0$ for the original problem $\mathcal{P}_{t_0, x_0}^{NL}$ and the precommitted agent is never locally preferences-consistent at time $t > t_0$ with respect to the original problem $\mathcal{P}_{t_0, x_0}^{NL}$. In general, the dynamically optimal strategy is never tail-optimal, not even locally, for the original problem $\mathcal{P}_{t_0, x_0}^{NL}$. In general, the Nash equilibrium strategy is never tail-optimal, not even locally, for the original problem $\mathcal{P}_{t_0, x_0}^{NL}$ and the Nash equilibrium agent is never preferences-consistent, not even locally, with respect to the original problem $\mathcal{P}_{t_0, x_0}^{NL}$.

5 On global tail optimality and linearity

The lack of local tail optimality for $t > t_0$ of the three strategies analyzed above seems to suggest that a globally tail-optimal strategy does not exist if the problem is non-linear. In other words, one may think that if a globally tail-optimal strategy exists, then the problem must be linear. This is indeed the case if the objective function includes a convex function $U^3(\cdot)$ of expected final wealth: this result is proven in Theorem 5.2 below, which – for convex functions – is the converse of Proposition 2.3.

Throughout this section, we shall work in the setting introduced in Section 2.1 and we shall also add the following technical assumption on the nature of the control plan u given a convex function h :

Assumption 5.1. *Let X_t^u be a controlled diffusion as in Section 2.1, where $u \in \mathcal{U}$. Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be a convex function, and let $I^h \subset \mathbb{R}$ be the set where h is strictly convex. We will say that the control plan u satisfies this assumption with the function h if either (i) $I^h = \emptyset$, or (ii) by setting $Y_t^u := \mathbb{E}[X_T^u | \mathcal{F}_t]$, we have $\text{Supp}(Y_{t'}^u | \mathcal{F}_t) \cap I^h \neq \emptyset$ for all $t_0 \leq t < t' \leq T$.*

Remark 5. This assumption, which will be used to prove Lemma 5.3, holds for instance if either

$I^h = \mathbb{R}$ or $\text{Supp}(Y_{t'}^u | \mathcal{F}_t) = \mathbb{R}$. The former will be the case of the mean-variance portfolio selection problem, see Section 6.

Theorem 5.2. *In the setting of Section 2.1, consider the optimization problem*

$$\begin{aligned} & \text{Problem } \mathcal{P}_{t_0, x_0} : \\ \sup_{u \in \mathcal{U}} J(t_0, x_0, u) &= \sup_{u \in \mathcal{U}} \left\{ \mathbb{E}_{t_0, x_0} \left[\int_{t_0}^T U^1(s, X_s, u_s) ds + U^2(X_T) \right] + U^3 \left[\mathbb{E}_{t_0, x_0}(X_T) \right] \right\}, \end{aligned}$$

where $U^3(\cdot)$ is a convex function. Suppose that the control plan

$$u_{t_0, x_0}^* : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$$

is globally tail-optimal over $[t_0, T]$ for Problem \mathcal{P}_{t_0, x_0} and satisfies Assumption 5.1 with $h = U^3$. Then, $U^3(x) = ax + b$ for some $a, b \in \mathbb{R}$, and the problem \mathcal{P}_{t_0, x_0} is linear. Moreover, the precommitment strategy, the dynamically optimal strategy and the Nash equilibrium strategy coincide.

In order to prove Theorem 5.2 we need the following lemma.

Lemma 5.3. *Let X_t^u be a scalar diffusion process as in Section 2.1 and let $\{u_s\}_{s \in [t_0, T]}$ satisfy Assumption 5.1 with the convex function $h : \mathbb{R} \rightarrow \mathbb{R}$. If $h(\cdot)$ is non-affine, then there are two times $t_0 \leq t < t' \leq T$ such that, conditional on \mathcal{F}_t :*¹¹

$$h(\mathbb{E}[X_T^u | \mathcal{F}_t]) < \mathbb{E}[h(\mathbb{E}[X_T^u | \mathcal{F}_{t'}]) | \mathcal{F}_t] \quad (18)$$

Proof. Let

$$Y_t^u =: \mathbb{E}[X_T^u | \mathcal{F}_t] \quad (19)$$

If h is strictly convex, for any two times $t_0 \leq t < t' \leq T$, by the Jensen's inequality we have that, conditional on \mathcal{F}_t

$$h(\mathbb{E}[Y_{t'}^u | \mathcal{F}_t]) < \mathbb{E}[h(Y_{t'}^u) | \mathcal{F}_t] \quad (20)$$

However Y_t^u is a martingale, therefore

$$h(\mathbb{E}[Y_{t'}^u | \mathcal{F}_t]) = h(\mathbb{E}[X_T^u | \mathcal{F}_t])$$

¹¹By “conditional on \mathcal{F}_t ” we mean outside the sets with zero conditional probability.

Thus, for every two times $t_0 \leq t < t' \leq T$, conditional on \mathcal{F}_t

$$h(\mathbb{E}[X_T^u | \mathcal{F}_t]) < \mathbb{E}[h(\mathbb{E}[X_T^u | \mathcal{F}_{t'}]) | \mathcal{F}_t] \quad (21)$$

that is (18). Let now h be convex but not strictly convex and not affine. Then there exists a set $I^h \subset \mathbb{R}$ s.t. h is strictly convex on I^h . By Assumption 5.1 for any $t_0 \leq t < t' \leq T$ conditional on \mathcal{F}_t the random variable $Y_{t'}^u = \mathbb{E}[X_T^u | \mathcal{F}_{t'}]$ has support with non-empty intersection with I^h , i.e. $\text{Supp}(Y_{t'}^u | \mathcal{F}_t) \cap I^h \neq \emptyset$. Then, in this case, conditional on \mathcal{F}_t ,

$$h(\mathbb{E}[Y_{t'}^u | \mathcal{F}_t]) < \mathbb{E}[h(Y_{t'}^u) | \mathcal{F}_t]$$

meaning that (18) holds. □

We can now prove Theorem 5.2.

Proof of Theorem 5.2. See the Appendix. □

Remark 6. Theorem 5.2 is proven for U^3 being a convex function. It is also possible to prove a weaker version of the theorem for U^3 concave. However, for concave functions Assumption 5.1 should be satisfied not only by the globally optimal control plan but also by all the admissible control plans belonging to \mathcal{U} . In this paper, we have preferred to present the stronger version of the theorem, that is also consistent with the classical example of non-linear problem in finance, namely, the mean-variance portfolio selection problem, that is the subject of the next section.

6 A notable example: the mean-variance problem

To better illustrate the theoretical results of the previous sections, we analyze a notable case example, the mean-variance portfolio selection problem, that is probably the most famous example of a non-linear time-inconsistent problem in finance. Its time inconsistency is due to the presence of the variance of final wealth in the performance criterion.

In the simplest framework, the mean-variance problem can be formalized as follows.

6.1 Formulation of the mean-variance portfolio selection problem

An investor has a wealth $x_0 > 0$ at time t_0 , and wants to solve a portfolio selection problem on the time horizon $[t_0, T]$. The financial market is the Black-Scholes model (see e.g. Björk (1998)): it consists of two assets, a riskless one, whose price $B(t)$ follows the dynamics:

$$dB(t) = rB(t)dt,$$

where $r > 0$, and a risky asset, whose price dynamics $S(t)$ follows a geometric Brownian motion with drift $\lambda \geq r$ and volatility $\sigma > 0$:

$$dS(t) = \lambda S(t)dt + \sigma S(t)dW(t),$$

where $W(t)$ is a standard Brownian motion defined on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbf{P})$, with $\mathcal{F}_t = \sigma\{W(s) : s \leq t\}$ the natural filtration. The proportion of portfolio invested in the risky asset at time t is denoted by $u(t)$. The fund at time t under control u , $X^u(t)$, grows according to the following SDE:

$$\begin{aligned} dX^u(t) &= X^u(t) [u(t)(\lambda - r) + r] dt + X^u(t)u(t)\sigma dW(t), \\ X^u(t_0) &= x_0. \end{aligned} \tag{22}$$

The investor is a mean-variance agent and her aim is to solve the problem

Problem $\mathcal{P}_{t_0, x_0}^{MV}$:

$$\sup_{u \in \mathcal{U}} J^{MV}(t_0, x_0, u) = \sup_{u \in \mathcal{U}} \{\mathbb{E}_{t_0, x_0}(X^u(T)) - \alpha \mathbb{V}_{t_0, x_0}(X^u(T))\}, \tag{23}$$

where $\alpha > 0$ and where $\mathbb{V}(\cdot)$ is the variance operator. It is easy to see that Problem (23) is a non-linear problem as in (5) with $U^1(x) = 0$, $U^2(x) = x - \alpha x^2$ and $U^3(x) = \alpha x^2$.

By the results in Section 4, there are three approaches for the mean-variance problem: (i) precommitment, (ii) dynamic optimality, and (iii) consistent planning.

The precommitment strategy \hat{u}_{t_0, x_0} is (see Zhou & Li, 2000):

$$\hat{u}_{t_0, x_0}(s, y) = \frac{\delta}{\sigma y} \left[x_0 e^{r(s-t_0)} - y + \frac{1}{2\alpha} e^{\delta^2(T-t_0) - r(T-s)} \right], \quad \text{for } (s, y) \in [t_0, T] \times \mathbb{R}, \quad (24)$$

where $\delta = (\lambda - r)/\sigma$.

The dynamically optimal policy \tilde{u} is (see Pedersen & Peskir, 2017):

$$\tilde{u}(s, y) = \frac{\delta}{\sigma y} \frac{1}{2\alpha} e^{(\delta^2 - r)(T-s)}, \quad \text{for } (s, y) \in [t_0, T] \times \mathbb{R}. \quad (25)$$

The consistent planning, Nash equilibrium policy \bar{u} is (see Basak & Chabakauri, 2010, and Björk & Murgoci, 2010):

$$\bar{u}(s, y) = \frac{\delta}{\sigma y} \frac{1}{2\alpha} e^{-r(T-s)}, \quad \text{for } (s, y) \in [t_0, T] \times \mathbb{R}. \quad (26)$$

6.2 Tail optimality and preferences consistency for mean-variance

In order to discuss tail optimality and preferences consistency for the three approaches to the mean-variance problem, we need to define the family of mean-variance problems

$$\{P_{t,x}^{MV}\}_{(t,x) \in [t_0, T] \times \mathbb{R}}$$

where

Problem $P_{t,x}^{MV}$:

$$\sup_{u \in \mathcal{U}} J^{MV}(t, x, u) = \sup_{u \in \mathcal{U}} \{ \mathbb{E}_{t,x}(X^u(T)) - \alpha \mathbb{V}_{t,x}(X^u(T)) \}.$$

We can now prove the results mentioned in Section 4.4 for the mean-variance problem.

Proposition 6.1. (i) For every $(t, x) \in (t_0, T] \times \mathbb{R}$, the precommitment strategy

$$\hat{u}_{t_0, x_0}(s, y) = \frac{\delta}{\sigma y} \left[x_0 e^{r(s-t_0)} - y + \frac{1}{2\alpha} e^{\delta^2(T-t_0) - r(T-s)} \right], \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}, \quad (27)$$

given by the restriction of (24) to $[t, T] \times \mathbb{R}$, is not locally tail-optimal at t for $P_{t,x}^{MV}$.

(ii) The precommitted investor who adopts the precommitment strategy $\hat{u}_{t_0, x_0}(s, X_s)$ (where \hat{u}

is given by (24)) over $[t_0, T]$ is not locally preferences-consistent at t with respect to P_{t_0, x_0}^{MV} for any $t \in (t_0, T]$.

Proof. (i) Let $(t, x) \in (t_0, T) \times \mathbb{R}$. By Definition 4.1 and Equation (24), the control plan that maximizes $J^{MV}(t, x, u)$ is given by

$$\hat{u}_{t,x}(s, y) = \frac{\delta}{\sigma y} \left[x e^{r(s-t)} - y + \frac{1}{2\alpha} e^{\delta^2(T-t) - r(T-s)} \right], \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}. \quad (28)$$

Because $\hat{u}_{t_0, x_0}(s, y) \neq \hat{u}_{t,x}(s, y)$ for $(s, y) \in [t, T] \times \mathbb{R}$, the precommitment strategy (27) is not locally tail-optimal at t for $P_{t,x}^{MV}$.

(ii) The precommitted investor who adopts the precommitment strategy $\hat{u}_{t_0, x_0}(s, y)$ over $[t_0, T]$, at time t with wealth x plays $\hat{u}_{t_0, x_0}(t, x)$. In order to be locally preferences-consistent with respect to P_{t_0, x_0}^{MV} she should play the first control action of the control plan $\hat{u}_{t,x}(s, y)$ given by (28). Because $\hat{u}_{t_0, x_0}(t, x) \neq \hat{u}_{t,x}(t, x)$, the precommitted investor is not locally preferences-consistent at t with respect to P_{t_0, x_0}^{MV} . \square

Proposition 6.2. *For every $(t, x) \in [t_0, T] \times \mathbb{R}$, the dynamically optimal strategy*

$$\tilde{u}(s, y) = \frac{\delta}{\sigma y} \frac{1}{2\alpha} e^{(\delta^2 - r)(T-s)} \quad \text{for } (s, y) \in [t, T] \times \mathbb{R},$$

given by the restriction of (25) to $[t, T] \times \mathbb{R}$, is not locally tail-optimal at t for $P_{t,x}^{MV}$.

Proof. Let $(t, x) \in [t_0, T] \times \mathbb{R}$. By Definition 4.1 and Equation (24), the control plan that maximizes $J^{MV}(t, x, u)$ is given by (28).

At time t with wealth x the dynamically optimal control action coincides with the optimal control action of (28): $\tilde{u}(t, x) = \hat{u}_{t,x}(t, x)$. However, after time t there is no longer coincidence between dynamically optimal strategy and optimal control plan (28): for $(s, y) \in (t, T) \times \mathbb{R}$, $\tilde{u}(s, y) = \hat{u}_{s,y}(s, y) \neq \hat{u}_{t,x}(s, y)$. Hence, the dynamically optimal strategy is not locally tail-optimal at t for $P_{t,x}^{MV}$. \square

Proposition 6.3. (i) *For every $(t, x) \in [t_0, T] \times \mathbb{R}$, the Nash equilibrium strategy*

$$\bar{u}(s, y) = \frac{\delta}{\sigma y} \frac{1}{2\alpha} e^{-r(T-s)} \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}, \quad (29)$$

given by the restriction of (26) to $[t, T] \times \mathbb{R}$, is not locally tail-optimal at t for $P_{t,x}^{MV}$.

(ii) The Nash equilibrium investor who adopts the Nash equilibrium strategy $\bar{u}(s, X_s)$ over $[t_0, T]$ is not locally preferences-consistent at t with respect to P_{t_0, x_0}^{MV} for any $t \in [t_0, T]$.

Proof. (i) Let $(t, x) \in [t_0, T] \times \mathbb{R}$. By Definition 4.1 and Equation (24), the control plan that maximizes $J^{MV}(t, x, u)$ is given by (28).

Because $\bar{u}(s, y) \neq \hat{u}_{t,x}(s, y)$ for $(s, y) \in [t, T] \times \mathbb{R}$, the Nash equilibrium strategy (29) is not locally tail-optimal at t for $P_{t,x}^{MV}$.

(ii) The Nash equilibrium investor who adopts the Nash equilibrium strategy $\bar{u}(s, y)$ over $[t_0, T]$, at time t with wealth x plays $\bar{u}(t, x)$. In order to be locally preferences-consistent with respect to P_{t_0, x_0}^{MV} she should play the first control action of the control plan $\hat{u}_{t,x}(s, y)$. Because $\bar{u}(t, x) \neq \hat{u}_{t,x}(t, x)$, the Nash equilibrium investor is not locally preferences-consistent at t with respect to P_{t_0, x_0}^{MV} . \square

Remark 7. Proposition 6.3 is a rigorous example of the remark written by Chew & Epstein (1990) on the fact that the Nash equilibrium approach leads to a time-consistent form of behaviour, but is driven by preferences that are not intertemporally consistent. As mentioned in the Introduction, this proposition sheds also further light on some of the criticisms to the game theoretical approach to the mean-variance problem appeared in the literature, e.g. Wang & Forsyth (2011), Bensoussan et al. (2019) and Forsyth (2020).

As mentioned in Section 4, the linear optimization problem associated to the mean-variance problem in the sense of Proposition 5.1 of Björk & Murgoci (2010) is well known. Indeed, the optimal solution to the linear stochastic optimal control problem

$$\text{Problem } \mathcal{P}_{t_0, x_0}^{L-ass-MV} : \sup_{u \in \mathcal{U}} \mathbb{E}_{t_0, x_0} \left[-\frac{1}{2\alpha} e^{-2\alpha X_T} \right] \quad (30)$$

coincides with the Nash-equilibrium strategy (26) (see also Basak & Chabakauri, 2010, Remark 1, and Vigna, 2014, Remark 3). In other words, $U^4(x) = 0$ and $U^5(x) = -1/(2\alpha)e^{-2\alpha x}$. Therefore, Propositions 4.5 and 4.6 hold considering the linear problem (30) and its obvious version $\mathcal{P}_{t,x}^{L-ass-MV}$ at time t with wealth x .

Finally, the lack of local tail optimality for $t > t_0$ of the three possible strategies for the mean-variance problem implies that the precommitment, the dynamically optimal and the Nash-

equilibrium strategies are not globally tail-optimal over $[t_0, T]$ for the mean-variance problem P_{t_0, x_0}^{MV} . This result, that is formalized in the following corollary, is an obvious consequence of Propositions 6.1, 6.2 and 6.3, but it can be proved also as a corollary of Theorem 5.2.

Corollary 6.4. *The control plans*

$$\hat{u}_{t_0, x_0} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R},$$

$$\tilde{u} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$$

and

$$\bar{u} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$$

given by (24), (25) and (26), respectively, are not globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} .

Proof. We provide two proofs.

First proof, as a corollary to Theorem 5.2. In Problem P_{t_0, x_0}^{MV} , $U^3(x) = \alpha x^2$, thus U^3 is convex and Assumption 5.1 is trivially satisfied by the control plans \hat{u}_{t_0, x_0} , \tilde{u} and \bar{u} , because $I^{U^3} = \mathbb{R}$. Moreover, $U^3(\cdot)$ is *not* in the form $ax + b$. Therefore, due to Theorem 5.2, none of the three control plans is globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} .

Second proof, as a corollary to Propositions 6.1, 6.2 and 6.3. Due to Propositions 6.1, 6.2 and 6.3, none of the three control plans is locally tail-optimal at t for $P_{t, x}^{MV}$ for every $t \in (t_0, T]$. Therefore, by definition, none of the three control plans is globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} . \square

Not only the three control plans \hat{u}_{t_0, x_0} , \tilde{u} and \bar{u} are not globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} . As a corollary of Theorem 5.2, there exists no control plan that is globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} .

Corollary 6.5. *There exists no control plan that is globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} .*

Proof. In Problem P_{t_0, x_0}^{MV} , $U^3(x) = \alpha x^2$, thus U^3 is convex and Assumption 5.1 is trivially satisfied by every control plan, because $I^{U^3} = \mathbb{R}$. Moreover, $U^3(\cdot)$ is *not* in the form $ax + b$. Therefore, due to Theorem 5.2, there exists no control plan that is globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} . \square

6.3 Further discussion on the mean-variance problem

In this paper the three approaches to the mean-variance portfolio selection problem have been analyzed under the new perspective of tail optimality and preferences consistency. In previous literature on mean-variance this comparison has been done by several authors exploiting different ideas. The three alternatives turn out to be meaningfully different in economics terms, and the choice among them can be driven by subjective factors, for example the distribution of final wealth or the behaviour of the investment strategy. Some results from the current literature are as follows.

An illuminating paper is van Staden, Dang & Forsyth (2021a), who assume that the agent is agnostic about the philosophical differences underlying the three approaches and compare the distribution of terminal wealth with each approach by equating the expectation of final wealths; they find that the precommitment final wealth has the lowest variance and the highest median value, but this advantage comes at the cost of increased left tail risk for the investor, due to negative skewness and large kurtosis; the game theoretical strategy produces lower variance of final wealth than the dynamically optimal one, and dominates the dynamically optimal one from a first-order stochastic order dominance when considering final wealth outcomes below the expected value.

The analysis performed by van Staden et al. (2021a) is done at initial time t_0 . A different angle is taken by Vigna (2020), who compares the three approaches not only at initial time t_0 but also at every time $t \in [t_0, T]$ with an intertemporal time- t -reward function that measures the happiness of the agent at every intermediate time; she finds that, while the precommitment agent dominates all the other agents if the relevant point of view is t_0 only, there is a unique break even point $t^* \in (t_0, T)$ such the Nash equilibrium agent beats the dynamically optimal one from t_0 to t^* and is dominated by him from t^* to T .

Another point of interest to the agent could be the investment strategy. A straight comparison between (25) and (26) clearly shows that the dynamically optimal strategy is riskier than the Nash equilibrium one. Menoncin & Vigna (2020) investigate the precommitment and the dynamically optimal strategies in a defined contribution pension scheme; they find that on average the precommitment portfolio contains less risky asset than the dynamically optimal one but the optimal share invested in the risky asset is highly more volatile in the precommitment case than in the dynamically optimal case; moreover, they also find that under extreme scenarios for market returns

the dynamically optimal strategy allows a more effective reaction to the market conditions, that is consistent with the results of van Staden et al. (2021a); the reason for such a better reaction lies in the fact that the precommitment strategy makes the final wealth as close as possible to a constant target, while the dynamically optimal strategy makes the final wealth as close as possible to a time-varying target, that adjusts to market returns.

Another useful criterion for the comparison is the impact of misspecification errors (considered both as model misspecification and parameter misspecification) on the dynamic mean-variance optimization: van Staden, Dang & Forsyth (2021b) investigate on it and find that without constraints the precommitment strategy is less robust than the Nash equilibrium strategy, but the opposite holds when constraints are applied.

7 Concluding remarks

When an intertemporal optimization problem over a time frame $[t_0, T]$ is linear and can be solved using dynamic programming, then, thanks to the Bellman's optimality principle, two important desirable features occur simultaneously. First, the optimal strategy is globally tail-optimal over $[t_0, T]$ for the considered problem. Second, the agent who adopts the optimal strategy is globally preferences-consistent over $[t_0, T]$ with respect to her initial preferences.

When an intertemporal optimization problem is not linear and does not permit application of dynamic programming, then the two features described above do not hold simultaneously. According to the existing literature, we say that the problem gives rise to time inconsistency.

The non-applicability of dynamic programming and the violation of the Bellman's optimality principle imposes an unavoidable price to be paid by agents. The price is different depending on the approach selected.

With the precommitment approach, the investor solves a kind of static Markowitz problem over $[t_0, T]$ and therefore keeps both properties of tail optimality and preferences consistency, but only at initial time t_0 : the precommitment strategy is locally tail-optimal at time t_0 (only) for the considered problem and the precommitted investor is locally preferences-consistent at time t_0 (only) with respect to her initial preferences.

With the dynamically optimal approach, the investor keeps the second property but not the first one, i.e., she is globally preferences-consistent with respect to her initial preferences, but, in

general, the dynamically optimal strategy is not locally tail-optimal at any time $t \in [t_0, T]$ for the considered problem.

With the Nash-equilibrium approach, the investor keeps none of the properties, i.e., the Nash equilibrium strategy is not locally tail-optimal at any time $t \in [t_0, T]$ for the considered problem and the investor who adopts it is not locally preferences-consistent at any time $t \in [t_0, T]$ with respect to her initial preferences.

Finally, if the objective function of an optimization problem includes a convex function of expected final wealth, a globally tail-optimal control plan exists if and only if the optimization problem is linear and dynamic programming is applicable.

In general, when dealing with non-linear problems, it seems quite hard to argue that one of the three approaches to time inconsistency currently available should be unambiguously preferable to the others for all agents. If the agent is agnostic about the philosophical differences underlying the three different approaches, a meaningful help can be provided by the comparison of the distribution of the final state variable or the behaviour of the optimal policy. If instead the agent is keen on playing a tail-optimal plan and on being preferences-consistent, then each approach has its own pros and cons and the appropriate strategy will depend on the subjective weight given to tail optimality and preferences consistency. The awareness of the price to be paid in terms of (lack of) these two criteria for each strategy might be of help to the agent.

References

- Basak, S. & Chabakauri, G. (2010). Dynamic mean-variance asset allocation, *Review of Financial Studies* **23**: 2970–3016.
- Bensoussan, A., Wong, K. & Yam, S. (2019). A paradox in time-consistency in the mean-variance problem?, *Finance and Stochastics* **23**: 173–207.
- Bertsekas, D. (2012). *Dynamic Programming and Optimal Control*, Athena Scientific.
- Björk, T. (1998). *Arbitrage Theory in Continuous Time*, Oxford University Press.
- Björk, T., Khapko, M. & Murgoci, A. (2017). On time-inconsistent stochastic control in continuous time, *Finance and Stochastics* **21**: 331–360.

- Björk, T. & Murgoci, A. (2010). A general theory of markovian time inconsistent stochastic control problems, *Preprint, SSRN: <http://ssrn.com/abstract=1694759>*.
- Björk, T. & Murgoci, A. (2014). A theory of markovian time-inconsistent stochastic control in discrete time, *Finance and Stochastics* **18**: 545–592.
- Björk, T., Murgoci, A. & Zhou, X. (2014). Mean variance portfolio optimization with state dependent risk aversion, *Mathematical Finance* **24**: 1–24.
- Chew, S. & Epstein, L. (1990). Nonexpected utility preferences in a temporal framework with an application to consumption-savings behaviour, *Journal of Economic theory* **50**: 54–81.
- Dai, M., Hin, H., Kou, S. & Xu, Y. (2021). A dynamic mean-variance analysis for log returns, *Management Science* **67**: 1093–1108.
- Ekeland, I., Mbodji, O. & Pirvu, T. (2012). Time-consistent portfolio management, *SIAM Journal on Financial Mathematics* **3**: 1–32.
- Epstein, L. & Le Breton, M. (1993). Dynamically consistent beliefs must be Bayesian, *Journal of Economic theory* **61**: 1–22.
- Forsyth, P. (2020). Multiperiod mean conditional value at risk asset allocation: Is it advantageous to be time consistent?, *SIAM Journal on Financial Mathematics* **11**: 358–384.
- Johnsen, T. & Donaldson, J. (1985). The structure of intertemporal preferences under uncertainty and time consistent plans, *Econometrica* **53**: 1451–1458.
- Karnam, C., Ma, J. & Zhang, J. (2017). Dynamic approaches for some time inconsistent problems, *The Annals of Applied Probability* **27**: 3435–3477.
- Menoncin, F. & Vigna, E. (2020). Mean-variance dynamic optimality for DC pension schemes, *European Actuarial Journal* **10**: 125–148.
- Pedersen, J. & Peskir, G. (2017). Optimal mean-variance portfolio selection, *Mathematics and Financial Economics* **11**: 137–160.
- Pollak, R. A. (1968). Consistent planning, *The Review of Economic Studies* **35**: 201–208.

- Powell, W. (2011). *Approximate Dynamic Programming*, John Wiley and Sons.
- Strotz, R. H. (1956). Myopia and inconsistency in dynamic utility maximization, *The Review of Economic Studies* **23**: 165–180.
- van Staden, P., Dang, D.-M. & Forsyth, P. (2021a). On the distribution of terminal wealth under dynamic mean-variance optimal strategies, *SIAM Journal on Financial Mathematics* **12**: 566–601.
- van Staden, P., Dang, D.-M. & Forsyth, P. (2021b). The surprising robustness of dynamic mean-variance portfolio optimization to model misspecification errors, *European Journal of Operational Research* **289**: 774–792.
- Vigna, E. (2014). On efficiency of mean variance based portfolio selection in defined contribution pension schemes, *Quantitative Finance* **14**: 237–258.
- Vigna, E. (2020). On time consistency for mean-variance portfolio selection, *International Journal of Theoretical and Applied Finance* **23**: 1–22.
- Wang, J. & Forsyth, P. A. (2011). Continuous time mean variance asset allocation: A time-consistent strategy, *European Journal of Operational Research* **209**: 184–201.
- Yong, J. & Zhou, X. Y. (1999). *Stochastic Controls: Hamiltonian Systems and HJB Equations*, Springer Verlag, New York.
- Zhou, X. & Li, D. (2000). Continuous-time mean-variance portfolio selection: A stochastic LQ framework, *Applied Mathematics and Optimization* **42**: 19–33.

Appendix

Proof of Theorem 5.2

If the control plan

$$u_{t_0, x_0}^* : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R} \quad (31)$$

is globally tail-optimal over $[t_0, T]$ for Problem \mathcal{P}_{t_0, x_0} , then its restriction to $[t, T] \times \mathbb{R}$

$$u_{t_0, x_0}^* : [t, T] \times \mathbb{R} \rightarrow \mathbb{R} \quad (32)$$

is locally tail-optimal for Problem $\mathcal{P}_{t, x}$ for every $(t, x) \in [t_0, T] \times \mathbb{R}$, where $\mathcal{P}_{t, x}$ is given by (3).

In particular, the plan (31) is locally tail-optimal at t_0 for \mathcal{P}_{t_0, x_0} , and therefore the precommitment strategy for Problem \mathcal{P}_{t_0, x_0} , $\hat{u}_{t_0, x_0} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$, exists and

$$\hat{u}_{t_0, x_0}(s, y) = u_{t_0, x_0}^*(s, y) \quad \text{for all } (s, y) \in [t_0, T] \times \mathbb{R}. \quad (33)$$

Then, for every $(t, x) \in [t_0, T] \times \mathbb{R}$ the precommitment strategy for the continuation problem $\mathcal{P}_{t, x}$, i.e., the locally tail-optimal strategy for $\mathcal{P}_{t, x}$, is given by

$$\hat{u}_{t, x} : [t, T] \times \mathbb{R} \rightarrow \mathbb{R}.$$

But because the continuation plan (32) is also locally tail-optimal for $\mathcal{P}_{t, x}$, then

$$\hat{u}_{t, x}(s, y) = u_{t_0, x_0}^*(s, y) \quad \text{for all } (s, y) \in [t, T] \times \mathbb{R}. \quad (34)$$

Then, (33) and (34) yield

$$\hat{u}_{t_0, x_0}(s, y) = \hat{u}_{t, x}(s, y) \quad \text{for all } (s, y) \in [t_0, T] \times \mathbb{R}. \quad (35)$$

Because (35) holds for every $(t, x) \in [t_0, T] \times \mathbb{R}$, we deduce that the precommitment strategy \hat{u}_{t_0, x_0} does not depend on the initial time-wealth point (t_0, x_0) :

$$\hat{u}_{t_0, x_0}(s, y) = \hat{u}(s, y) \quad \text{for all } (s, y) \in [t_0, T] \times \mathbb{R}$$

and because of (33) the globally tail-optimal strategy too does not depend on (t_0, x_0) :

$$u_{t_0, x_0}^*(s, y) = u_{t, x}^*(s, y) = u^*(s, y).$$

As a consequence, the dynamically optimal strategy also coincides with the globally tail-optimal

strategy. Indeed, by definition

$$\tilde{u}(t, x) = \hat{u}_{t,x}(t, x),$$

therefore

$$\tilde{u}(t, x) = \hat{u}_{t,x}(t, x) = \hat{u}(t, x) = u^*(t, x) \quad \text{for all } (t, x) \in [t_0, T] \times \mathbb{R}.$$

Furthermore, because $\{u^*(s, y)\}_{(s,y) \in [t_0, T] \times \mathbb{R}}$ is globally tail-optimal over $[t_0, T]$ for Problem \mathcal{P}_{t_0, x_0} , by definition the Bellman's optimality principle holds, and therefore the recursive equation for the value function

$$V(t, x) = \sup_{u \in \mathcal{U}} \mathbb{E}_{t,x} \left[\int_t^{t'} U^1(s, X_s, u_s) ds + V(t', X_{t'}) \right] \quad \text{for } t' \in [t, T]$$

holds. Now notice that, using the tower property of conditional expectation, for every $t_0 \leq t < t' \leq T$ we have (in the following we shall write X_t^* in the place of $X_t^{u^*}$)

$$\begin{aligned}
V(t, x) &= J(t, x, u^*) = \mathbb{E}_{t,x} \left[\int_t^{t'} U^1(s, X_s^*, u_s^*) ds + \int_{t'}^T U^1(s, X_s^*, u_s^*) ds + U^2(X_T^*) \right] + U^3(\mathbb{E}_{t,x}(X_T^*)) = \\
&= \mathbb{E} \left[\int_t^{t'} U^1(s, X_s^*, u_s^*) ds + \int_{t'}^T U^1(s, X_s^*, u_s^*) ds + U^2(X_T^*) \middle| \mathcal{F}_t \right] + U^3(\mathbb{E}(X_T^* | \mathcal{F}_t)) = \\
&= \mathbb{E} \left[\int_t^{t'} U^1(s, X_s^*, u_s^*) ds + \int_{t'}^T U^1(s, X_s^*, u_s^*) ds + U^2(X_T^*) + U^3(X_T^*) \middle| \mathcal{F}_t \right] + \\
&+ U^3(\mathbb{E}(X_T^* | \mathcal{F}_t)) - \mathbb{E}[U^3(X_T^*) | \mathcal{F}_t] = \\
&= \mathbb{E} \left\{ \int_t^{t'} U^1(s, X_s^*, u_s^*) ds + \mathbb{E} \left[\int_{t'}^T U^1(s, X_s^*, u_s^*) ds + U^2(X_T^*) + U^3(X_T^*) \middle| \mathcal{F}_{t'} \right] \middle| \mathcal{F}_t \right\} + \\
&+ U^3(\mathbb{E}(X_T^* | \mathcal{F}_t)) - \mathbb{E}[U^3(X_T^*) | \mathcal{F}_t] = \\
&= \mathbb{E} \left\{ \int_t^{t'} U^1(s, X_s^*, u_s^*) ds + \mathbb{E} \left[\int_{t'}^T U^1(s, X_s^*, u_s^*) ds + U^2(X_T^*) \middle| \mathcal{F}_{t'} \right] + U^3(\mathbb{E}[X_T^* | \mathcal{F}_{t'}]) \right. \\
&\left. + \mathbb{E}[U^3(X_T^*) | \mathcal{F}_{t'}] - U^3(\mathbb{E}[X_T^* | \mathcal{F}_{t'}]) \middle| \mathcal{F}_t \right\} + U^3(\mathbb{E}(X_T^* | \mathcal{F}_t)) - \mathbb{E}[U^3(X_T^*) | \mathcal{F}_t] = \\
&= \mathbb{E} \left[\int_t^{t'} U^1(s, X_s^*, u_s^*) ds + V(t', X_{t'}^*) \middle| \mathcal{F}_t \right] + U^3(\mathbb{E}(X_T^* | \mathcal{F}_t)) - \mathbb{E}[U^3(\mathbb{E}[X_T^* | \mathcal{F}_{t'}]) | \mathcal{F}_t]
\end{aligned}$$

where we have used the fact that $\mathbb{E}[\mathbb{E}[U^3(X_T^*) | \mathcal{F}_{t'}] | \mathcal{F}_t] = \mathbb{E}[U^3(X_T^*) | \mathcal{F}_t]$ and the fact that $J(t', X_{t'}^*, u^*) = V(t, X_{t'}^*)$. If U^3 is non-affine, then due to Lemma 5.3, there are $t_0 \leq t < t' \leq T$ s.t.

$$U^3(\mathbb{E}(X_T^* | \mathcal{F}_t)) - \mathbb{E}[U^3(\mathbb{E}[X_T^* | \mathcal{F}_{t'}]) | \mathcal{F}_t] < 0$$

yielding

$$\begin{aligned}
V(t, x) &< \mathbb{E} \left[\int_t^{t'} U^1(s, X_s^*, u_s^*) ds + V(t', X_{t'}^*) \middle| \mathcal{F}_t \right] \\
&\leq \sup_{u \in \mathcal{U}} \mathbb{E} \left[\int_t^{t'} U^1(s, X_s^*, u_s^*) ds + V(t', X_{t'}^*) \middle| \mathcal{F}_t \right] = V(t, x)
\end{aligned}$$

Therefore, $U^3(x) = ax + b$ for some $a, b \in \mathbb{R}$ and Problem \mathcal{P}_{t_0, x_0} is linear. Finally, if $U^3(x) = ax + b$, the extended Hamilton-Jacobi-Bellman equation needed to find the equilibrium strategy in the consistent planning approach (see Björk & Murgoci, 2010) is the classical Hamilton-Jacobi-Bellman equation of dynamic programming, and the equilibrium strategy coincides with the optimal strategy $u^*(\cdot)$:

$$\bar{u}(s, y) = u^*(s, y) \quad \text{for all } (s, y) \in [t_0, T] \times \mathbb{R}.$$

This is what we needed to prove. □