

Collegio Carlo Alberto



Tail optimality and preferences consistency for
stochastic optimal control problems

Elena Vigna

No. 502
July 2017

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

Tail optimality and preferences consistency for stochastic optimal control problems*

Elena Vigna[†]

April 12, 2018

Abstract

Given a stochastic optimal control problem over a time interval $[t_0, T]$ and a control map associated to it, we introduce the four notions of local and global tail optimality of the control map, and local and global preferences consistency of the optimizer. While the notion of tail optimality of a control map is not new in control problems, the notion of preferences consistency of an optimizer seems novel.

We prove that, in the case of a *linear* time-consistent problem where dynamic programming can be applied, the optimal control map is globally tail-optimal and the optimizer 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 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 approach there is neither local tail optimality nor local preferences consistency at any time with

*I thank Bjarne Højgaard and Luigi Montrucchio for useful comments.

[†]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.

respect to the original non-linear problem, but there is global tail optimality and global preferences consistency with respect to a different linear problem.

This 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, Nash perfect equilibrium, mean-variance portfolio selection.

JEL classification: C61, D81, G11.

1 Introduction

The aim of this paper is twofold: (i) shed light on the differences between *linear* stochastic optimal control problems, where dynamic programming can be applied, and *non-linear* stochastic optimal control problems, where dynamic programming cannot be applied; (ii) shed light on the differences among the three common approaches to non-linear problems, namely precommitment approach, dynamically optimal approach and consistent planning approach. 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 individual 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 (2010).

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

optimizer, sometimes for the optimization problem or criterion used, sometimes for the optimizer himself. 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’.*”. This confusion stems from the fact that when talking about time consistency the two notions of optimality of the map and consistency to one’s own preferences are often 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 stochastic optimal control problem into the two notions of *tail optimality of a control map* and *preferences consistency of the optimizer*. 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 feature of global tail optimality of a control map 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 optimizer wishing to solve an intertemporal stochastic optimal control problem is apparently novel. However, the idea to consistency to preferences is not new in economics literature and decision theory 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 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 optimizer 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 optimizer was already observed by Chew & Epstein (1990), see Sections 4.4 and 5.2.

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 map, 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 optimizer, 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 of the three current approaches to time inconsistency. In Section 5, we illustrate in detail the special case of mean-variance portfolio selection problem. Section 6 concludes.

2 Tail optimality of a control map

In this section we define the notion of tail optimality of a control map for an intertemporal optimization problem. We 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 map 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 map for the family of problems considered. It is important to recall that these notions are not new, and are strongly related to the Bellman’s optimality principle for optimization problems.

2.1 Setting

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

- the time horizon over which the optimization is done is fixed and is $[t_0, T]$;¹

¹For simplicity, we here assume $t_0 \leq T \in \mathbb{R}$, i.e., bounded time intervals, but the case of unbounded time interval could be considered too, by setting $T = +\infty$.

- the wealth² $X_s \in \mathbb{R}$ of the optimizer evolves according to the controlled stochastic differential equation:

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

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

- $\{u_s\}_{s \in [t_0, T]}$ is a control action that the controller can choose at any time s according to some criterion; 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 $u : [t_0, T] \times \mathbb{R}$;
- $\mathcal{U} \subseteq \mathbb{R}$ is some set of admissible controls.³

It is essential to highlight that the criterion selected by the controller in the optimization problem represents the *preferences* of the optimizer and is typically given by the combination of different utility functions. In particular, putting ourselves in the setting introduced by Björk & Murgoci (2010), we shall assume that the controller wants to solve the following optimization problem:

Problem \mathcal{P}_{t_0, x_0} :

$$\sup_{u \in \mathcal{U}} J(t_0, x_0, u) = \sup_{u \in \mathcal{U}} \left\{ \mathbb{E}_{t_0, x_0} \left[\int_{t_0}^T U^1(s, X_s, u_s) ds + U^2(X_T) \right] + U^3 \left[\mathbb{E}_{t_0, x_0}(X_T) \right] \right\} \tag{2}$$

where $U^1(\cdot)$, $U^2(\cdot)$ and $U^3(\cdot)$ are (possibly non-linear) utility functions that identify the optimizer'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 optimizer, to stress the crucial role played by the utility

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

³For 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 the identification of the individual's preferences. For notational convenience, in the following, we will simply refer to \mathcal{P}_{t_0, x_0} or $J(t_0, x_0, u)$.

This problem belongs to the more general family of problems

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

where

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

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

2.2 Linear vs non-linear problems

In line with Björk & Murgoci (2010) we recall that the nature of Problem (2) strongly depends on the utility function $U^3(\cdot)$, and there are two possible cases:

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

2.3 Local and global tail optimality

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

Definition 2.1 (Local tail optimality). *Given the stochastic optimal control problem $\mathcal{P}_{t,x}$ as in (3), where the point $(t, x) \in [t_0, T] \times \mathbb{R}$ is fixed, we say that the control map*

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

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

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

if it exists.

Remark 2. 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 map $u_{t,x}^*(\cdot)$ is played from t until T , meaning that at each time $s \in [t, T]$ with wealth $X_s \in \mathbb{R}$ the optimizer 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 map u^* . Intuitively, the map 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 3. It is important to stress that the optimal control map $u_{t,x}^*(\cdot)$ is a function of time $s \geq t$ and wealth $y \in \mathbb{R}$, but it also depends on the initial point (t, x) . Indeed, two optimal control maps 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-wealths 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 map*

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

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 $u_{t_0, x_0}^(s, y)$ 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: does a control map 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.

2.4 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 (for conciseness reasons we omit all details and refer the interested reader to classical sources, e.g., Yong & Zhou (1999), Björk (1998)):

- consider the more general problem to be solved at time t with wealth x , Problem $\mathcal{P}_{t, x}^L$:

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

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

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

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 3 for the general case, the optimal control law $u_{t_0,x_0}^*(s, y)$ is optimal not only on $[t_0, T]$ but also on every subinterval $[\tau, T]$ with $\tau > t_0$ for the translated problem $\mathcal{P}_{\tau,x_\tau}^L$. This is the well-known Bellman's optimality principle, see Bertsekas (2012). This means that the optimal strategy for the new translated problem $\mathcal{P}_{\tau,x_\tau}^L$ at time τ with initial wealth x_τ coincides with the *restriction* on $[\tau, T]$ of the optimal strategy found at initial time t_0 :

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

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 lack of rigour,⁴ we shall simply denote it by $u^*(s, y)$:

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

In other words, and notably, (11) shows that for a linear problem the optimal control law *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 maps $u_{t,x}^*(s, y)$ of the problems $\mathcal{P}_{t,x}^L$ can be identified by the infinitely many *restrictions* of the control map $u^* : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ for each $(t, x) \in [t_0, T] \times \mathbb{R}$

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

⁴The maximum domain of the control map $u^*(s, y)$ is $[t_0, T] \times \mathbb{R}$. Clearly, the domain of the optimal control map of the translated problem $\mathcal{P}_{\tau,x_\tau}^L$ is restricted to $[\tau, T] \times \mathbb{R}$.

Given Definition 2.2 and the validity of the Bellman’s optimality principle for linear problems, we can now state the following obvious result.

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

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

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

Proposition 2.3 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 optimizer

While the notion of global tail optimality is not new and essentially coincides with the Bellman’s principle, the notion of preferences consistency for a stochastic optimal control 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: see Johnsen & Donaldson (1985), Epstein & Le Breton (1993), Chew & Epstein (1990). While the mentioned works consider choices among lotteries in a discrete framework, 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 2.1, imagine a decision maker 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). He 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 he does not change his preferences over time and that his 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 Problem \mathcal{P}_{t, x_t} as in (3), because of unchanged preferences. Intuitively, if this happens, we will say that the optimizer is preferences-consistent.

In particular, we will say that the optimizer 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

he plays at time t optimizes the translated 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 at time t_0 . In other words, in order to talk about preferences consistency at time $t > t_0$ one needs to know the original 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 optimizer 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 control that he plays at time t with wealth x coincides with the first control of the tail-optimal control map of the translated problem $\mathcal{P}_{t,x}$, i.e., if at time t with wealth x he plays $u_{t,x}^*(t, x)$, where $u_{t,x}^*(s, y)$ (for $(s, y) \in [t, T] \times \mathbb{R}$) is the tail-optimal control map for $\mathcal{P}_{t,x}$:*

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

Remark 4. Notice that the local preferences consistency at time t implies only that for every wealth x the controller plays the optimal control action for $\mathcal{P}_{t,x}$ at time t , but does not mean that he will continue to play the optimal map $u_{t,x}^*(s, y)$ also for $s > t$. In other words, the locally preferences-consistent optimizer plays the optimal control map for $\mathcal{P}_{t,x}$ only instantaneously at time t , i.e., he is only *instantaneously* optimal for the translated problem $\mathcal{P}_{t,x}$.

Quite naturally, if an optimizer is locally preferences-consistent at t for every t in a given interval, he 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 optimizer 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 he 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. According to their

definition, and roughly speaking, 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 map and its instantaneous optimality, we here stress the importance of the consistency of the optimizer to his *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 individual is globally preferences consistent.

As in the case of global tail optimality, the following existence issue can arise: does an optimizer 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 optimizer 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 optimizer are identified by the following linear problem:

$$\begin{aligned} & \text{Problem } \mathcal{P}_{t_0, x_0}^L : \\ & \sup_{u \in \mathcal{U}} J^L(t_0, x_0, u) = \sup_{u \in \mathcal{U}} \mathbb{E}_{t_0, x_0} \left[\int_{t_0}^T U^1(s, X_s, u_s) ds + U^2(X_T) \right], \end{aligned} \tag{12}$$

as in (4). Then, the optimizer can apply dynamic programming as explained in Section 2.4 to solve it and find the optimal control map $u_{t_0, x_0}^* : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}$. Imagine that the optimizer 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 he plays exactly the optimal control map of the translated Problem $\mathcal{P}_{\tau, x_\tau}^L$. This means that the optimizer is preferences-consistent at time $\tau > t_0$ with respect to his original preferences identified by Problem \mathcal{P}_{t_0, x_0}^L . Because this happens at every time $\tau \in [t_0, T]$ we conclude that the optimizer is globally preferences-consistent over $[t_0, T]$ with respect to the original problem (12).

This result is formalized in the following proposition.

Proposition 3.3. *Let the preferences of an optimizer be identified by the linear problem \mathcal{P}_{t_0, x_0}^L*

given by (12). If the optimizer plays the optimal control $u_{t_0, x_0}^*(s, X_s) = u^*(s, X_s)$ over $[t_0, T]$, where $u^*(\cdot)$ is found via dynamic programming, then he 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 map and global preferences consistency of the optimizer 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. 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 investor wants to solve the following non-linear problem:

$$\begin{aligned} & \text{Problem } \mathcal{P}_{t_0, x_0}^{NL} : \\ \sup_{u \in \mathcal{U}} J^{NL}(t_0, x_0, u) &= \sup_{u \in \mathcal{U}} \left\{ \mathbb{E}_{t_0, x_0} \left[\int_{t_0}^T U^1(s, X_s, u_s) ds + U^2(X_T) \right] + U^3 \left[\mathbb{E}_{t_0, x_0}(X_T) \right] \right\} \end{aligned} \quad (13)$$

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 this time inconsistency. We will see that none of the existing approaches available for a non-linear problem (13) keeps simultaneously both properties of global tail optimality and global preferences consistency. Nevertheless, it is possible to analyze the available approaches and see what are the properties characterizing each of them.

The three approaches currently available for the non-linear problem (13) are: (i) the pre-commitment approach, (ii) the dynamic optimality approach and (iii) the consistent plan-

ning, or game theoretical, or Nash-equilibrium approach.

4.1 Precommitment approach

In order to solve the non-linear problem (13) 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 (13), the strategy \hat{u} that maximizes $J^{NL}(t_0, x_0, u)$, i.e., the control map*

$$\hat{u}_{t_0, x_0} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R} \quad (14)$$

such that

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

if it exists, is called the precommitment strategy.

Because in this kind 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 control map $\hat{u}_{t_0, x_0}(s, y)$ (for $(s, y) \in [\tau, T] \times \mathbb{R}$), is, in general, not optimal for the translated criterion $J^{NL}(\tau, x_\tau, u)$. In other words,

$$\operatorname{argmax}_{u \in \mathcal{U}} J^{NL}(\tau, x_\tau, u) = \hat{u}_{\tau, x_\tau}(s, y) \neq \hat{u}_{t_0, x_0}(s, y) \quad \text{for } (s, y) \in [\tau, T] \times \mathbb{R}, \quad (15)$$

while there would be equality with validity of the Bellman's principle (see equation (10)). In other words, the precommitment strategy (14) *depends essentially on the initial point* (t_0, x_0) .

This is the reason why the strategy is named precommitment strategy: the precommitted decision-maker standing at time t_0 should "precommit" himself to follow the strategy $\hat{u}_{t_0, x_0}(s, y)$ from t_0 to T , even if he knows that at later point in time τ he is still solving the original problem $\mathcal{P}_{t_0, x_0}^{NL}$, but *not* the translated problem $\mathcal{P}_{\tau, x_\tau}^{NL}$. Indeed, due to (15), the control that the precommitted investor plays at every time $\tau > t_0$ is, in general, not equal to the first optimal control for the translated problem $\mathcal{P}_{\tau, x_\tau}^{NL}$. Therefore, the precommitted

investor is locally preferences-consistent at time t_0 with respect to $\mathcal{P}_{t_0, x_0}^{NL}$ (because at time t_0 he plays $\hat{u}_{t_0, x_0}(t_0, x_0)$ that is the first action of the optimal map for $\mathcal{P}_{t_0, x_0}^{NL}$), but 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 optimizer be identified by the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (13), and let assume that there exists the precommitment strategy $\hat{u}_{t_0, x_0}(s, y)$ for Problem $\mathcal{P}_{t_0, x_0}^{NL}$. If the optimizer plays the precommitment strategy $\hat{u}_{t_0, x_0}(s, X_s)$ over $[t_0, T]$, then he 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 map. 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 stochastic optimal control problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (13), the precommitment strategy $\hat{u}_{t_0, x_0}(s, y)$ given by (14) for $(s, y) \in [t_0, T] \times \mathbb{R}$, 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 map and local preferences consistency at initial time t_0 of the optimizer who adopts the precommitment strategy is all what the precommitment approach can offer. In general, for a non-linear problem the precommitment optimal control map $\hat{u}_{t_0, x_0}(s, y)$ is not globally tail-optimal and the precommitted optimizer 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 5.

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 (2016). The problem of precommitment is about preferences inconsistency after t_0 : the precommitted decision-maker only cares about initial time t_0 and final time T , disregarding that he 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 Markovitz framework than to the continuous-time intertemporal 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 investor is consistent to his 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 (2016)), 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 (13) in 4 steps.⁵

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 fixed initial point (t_0, x_0) the precommitment strategy maximizing the criterion $J^{NL}(t_0, x_0, u)$ exists and is given by (see (14)):

$$\hat{u}_{t_0, x_0} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}. \quad (16)$$

Step 3. Define the new control map

$$\tilde{u}(s, y) = \hat{u}_{s, y}(s, y), \quad \text{for } (s, y) \in [t_0, T] \times \mathbb{R}, \quad (17)$$

where the right hand side of (17) is obtained by replacing t_0 with s and x_0 with y in the function (16).

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

Remark 5. Unlike the precommitment strategy (14), the dynamically optimal investment strategy (17) 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 map of a linear optimization problem $u^*(s, y)$ 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 individual.

By construction, at generic time $t \in [t_0, T]$ with wealth x_t the dynamically optimal individual faces the problem $\mathcal{P}_{t, x_t}^{NL}$ and solves it with the precommitment approach for $\mathcal{P}_{t, x_t}^{NL}$,

⁵Pedersen & 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 (13).

as if (t, x_t) was the initial time-wealth point. In fact, he plays the *first* control of the precommitment strategy for \mathcal{P}_{t,x_t}^{NL} , because he plays $\hat{u}_{t,x_t}(t, x_t)$. Because the initial preferences of the individual are given by Problem $\mathcal{P}_{t_0,x_0}^{NL}$ and because at time t he plays the first control of the optimal strategy for the translated problem \mathcal{P}_{t,x_t}^{NL} , by definition he is locally preferences-consistent at time t with respect to $\mathcal{P}_{t_0,x_0}^{NL}$. Not only: because this happens for every $t \in [t_0, T]$, he is globally preferences-consistent over $[t_0, T]$ with respect to $\mathcal{P}_{t_0,x_0}^{NL}$.

This result is formalized in the next proposition.

Proposition 4.4. *Let the preferences of an optimizer be identified by the non-linear problem $\mathcal{P}_{t_0,x_0}^{NL}$ as in (13), and let assume that there exists the precommitment strategy $\hat{u}_{t_0,x_0}(s, y)$ for $\mathcal{P}_{t_0,x_0}^{NL}$. If the optimizer plays the dynamically optimal strategy $\tilde{u}(s, X_s)$ over $[t_0, T]$, where $\tilde{u}(s, y)$ is given by (17), then he is globally preferences-consistent over $[t_0, T]$ with respect to $\mathcal{P}_{t_0,x_0}^{NL}$.*

Remark 6. Regarding the relationship between precommitment approach and dynamically optimal approach, as already mentioned we see that at each $t \in [t_0, T]$ with wealth x the dynamically optimal strategy $\tilde{u}(t, x)$ coincides with the *first* control of the precommitment strategy solution to $\mathcal{P}_{t,x}^{NL}$, i.e., it coincides with the first control of the control map $\hat{u}_{t,x}(s, y)$ ($(s, y) \in [t, T] \times \mathbb{R}$) selected by the investor 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' the dynamically optimal strategy coincides with the first control of the precommitted strategy for Problem $\mathcal{P}_{t',x_{t'}}^{NL}$. Therefore, the dynamically optimal investor can be seen as the *continuous reincarnation* of the precommitted investor. Moreover, even if this strategy has been formalized and deeply studied by Pedersen & Peskir (2017),⁶ the dynamically optimal individual is similar to the continuous version of the naive individual described by Pollak (1968). We notice that, rather than being naive, the dynamically optimal decision maker 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

⁶Related work can be found also in Karnam et al. (2016).

this paper the time interval over which the optimization is done shrinks when time passes, while it remains fixed in those problems).

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

From Definition 2.1 and Remark 2, we see that a control map 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 map $\tilde{u}(s, y)$ 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 Problem $\mathcal{P}_{t_0, x_0}^{NL}$, there exists no such $t \in [t_0, T]$ that makes the dynamically optimal strategy being locally tail-optimal at t for $\mathcal{P}_{t, x}^{NL}$.

4.3 Consistent planning, Nash equilibrium approach

According to the consistent planning approach, in order to solve the non-linear problem (13), 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 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), Proposition 4.1 and Theorem 4.1. Like the optimal control law of a linear problem, also 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:

$$\bar{u} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R}. \quad (18)$$

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 $\bar{u}(s, y)$ associated to the non-linear problem (13) coincides with the optimal control law found via dynamic programming solution to the linear problem

$$\begin{aligned} & \text{Problem } \mathcal{P}_{t_0, x_0}^{L-ass-NL} : \\ & \sup_{u \in \mathcal{U}} \mathbb{E}_{t_0, x_0} \left[\int_{t_0}^T U^4(s, X_s, u_s) ds + U^5(X_T) \right]. \end{aligned} \tag{19}$$

For instance, in the case of the mean-variance preferences, where $U^1(\cdot) \equiv 0$, $U^2(x) = x - \alpha x^2$ and $U^3(x) = \alpha x^2$, it is easy to show that $U^4(\cdot) \equiv 0$ while $U^5(\cdot)$ is the exponential utility function, see Section 5. This equivalence between mean-variance and exponential preferences is noted also by Basak & Chabakauri (2010).

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

By noting that the Nash equilibrium control map $\bar{u}(s, y)$ coincides with the optimal control map 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 map is globally tail-optimal for Problem $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$. Similarly, using Proposition 3.3 it is immediate to conclude that the individual who plays the Nash equilibrium strategy over $[t_0, T]$ is globally preferences-consistent over $[t_0, T]$ with respect to Problem $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$.

These results are formalized in the following propositions.

Proposition 4.5. *Given the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (13), and given the linear problem*

$$\mathcal{P}_{t_0, x_0}^{L-ass-NL} \tag{20}$$

as in (19) 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 map

$$\bar{u} : [t_0, T] \times \mathbb{R} \rightarrow \mathbb{R} \tag{21}$$

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 $\mathcal{P}_{t_0, x_0}^{L-ass-NL}$.

Proposition 4.6. *Let the preferences of an optimizer be identified by the non-linear problem $\mathcal{P}_{t_0, x_0}^{NL}$ as in (13), and let assume that there exists the Nash equilibrium strategy $\bar{u}(s, y)$ for $\mathcal{P}_{t_0, x_0}^{NL}$. If the optimizer plays $\bar{u}(s, X_s)$ over $[t_0, T]$, then he 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 Bjrk & Murgoci (2010).*

To sum up, the consistent planning optimizer, or Nash equilibrium optimizer, is globally preferences-consistent with respect to the associated linear problem, and he plays a strategy that is globally tail-optimal for the associated linear problem. In general, he is not preferences consistent to his original non-linear preferences and the map that he plays is not tail-optimal for the original non-linear problem.

Regarding their surprising result in Proposition 5.1, Bjrk & 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 this result from a different angle. For a non-linear problem (13) the Nash equilibrium approach is equivalent to apply the solution to the associated linear problem (19). This means that in order to be time-consistent in the consistent-planning sense, the investor has to choose a different objective functional, in other words, different preferences. For the mean-variance problem, the investor 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.

4.4 Considerations

In the previous sections we have showed that each of the three approaches currently available for a non-linear intertemporal optimization problem (13) carries only some of the two desirable properties of tail optimality and preferences consistency.

By Propositions 4.2 and 4.3, the precommitted optimizer is locally preferences-consistent at time t_0 with respect to his initial preferences given by Problem (13), and the precommitment strategy is locally tail-optimal at t_0 for Problem (13). Therefore, the precommitment approach keeps local tail optimality and local preferences consistency at initial time t_0 .

By Proposition 4.4, the dynamically optimal investor is globally preferences-consistent with respect to his initial preferences given by Problem (13).

By Propositions 4.5 and 4.6, the Nash equilibrium strategy is globally tail-optimal for the linear problem that is associated to Problem (13), and the Nash equilibrium optimizer is globally preferences-consistent with respect to the linear problem that is associated to Problem (13).

In general, the precommitment strategy is never locally tail-optimal at time $t > t_0$ for the original problem (13) and the precommitted optimizer is never locally preferences-consistent at time $t > t_0$ with respect to the original problem (13). In general, the dynamically optimal strategy is never tail-optimal, not even locally, for the original problem (13). In general, the Nash equilibrium strategy is never tail-optimal, not even locally, for the original problem (13) and the Nash equilibrium optimizer is never preferences-consistent, not even locally, with respect to the original problem (13). As mentioned in Section 1, the preferences inconsistency of the Nash equilibrium optimizer 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”*. Proving these results in a general framework is far from trivial. However, this can be done with a case example, namely, the mean-variance portfolio selection problem.

5 A notable example: the mean-variance problem

Perhaps, the most famous example of non-linear time-inconsistent problem in finance is the mean-variance portfolio selection problem. 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.

5.1 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:

$$\begin{aligned} dX^u(t) &= X^u(t) [u(t)(\lambda - r) + r] dt + X^u(t)u(t)\sigma dW(t), \\ X^u(t_0) &= x_0 \geq 0. \end{aligned} \tag{22}$$

The investor is a mean-variance optimizer and his aim is to solve the problem

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

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

where $\alpha > 0$ and \mathcal{U} is some set of admissible strategies. It is easy to see that Problem (23) is a non-linear problem as in (13) with $U^1(x) = 0$, $U^2(x) = x - \alpha x^2$ and $U^3(x) = \alpha x^2$.

By 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 $\hat{u}_{t_0, x_0}(s, y)$ 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 $\tilde{u}(s, y)$ 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 $\bar{u}(s, y)$ is (see Basak & Chabakauri (2010) and Björk & Murgoci (2010)):

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

5.2 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

$$\{ P_{t,x}^{MV} \}_{(t,x) \in [t_0, T] \times \mathbb{R}} \quad (27)$$

where

$$\begin{aligned} & \text{Problem } P_{t,x}^{MV} : \\ & \sup_{u \in \mathcal{U}} J^{MV}(t, x, u) = \sup_{u \in \mathcal{U}} \{ \mathbb{E}_{t,x}(X^u(T)) - \alpha \mathbb{V}_{t,x}(X^u(T)) \}. \end{aligned} \quad (28)$$

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

Proposition 5.1. (i) For every $(t, x) \in (t_0, T] \times \mathbb{R}$, the precommitment strategy

$$\hat{u}_{t_0, x_0}(s, y) = \frac{\delta}{\sigma y} \left[x_0 e^{r(s-t_0)} - y + \frac{1}{2\alpha} e^{\delta^2(T-t_0) - r(T-s)} \right], \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}, \quad (29)$$

given by the restriction of (24) to $[t, T] \times \mathbb{R}$, is not locally tail-optimal at t for $P_{t,x}^{MV}$ given by (28).

(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} given by (23) for any $t \in (t_0, T]$.

Proof. The proof is in the Appendix. □

Proposition 5.2. For every $(t, x) \in [t_0, T] \times \mathbb{R}$, the dynamically optimal strategy

$$\tilde{u}(s, y) = \frac{\delta}{\sigma y} \frac{1}{2\alpha} e^{(\delta^2 - r)(T-s)} \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}, \quad (30)$$

given by the restriction of (25) to $[t, T] \times \mathbb{R}$, is not locally tail-optimal at t for $P_{t,x}^{MV}$ given by (28).

Proof. The proof is in the Appendix. □

Proposition 5.3. (i) For every $(t, x) \in [t_0, T] \times \mathbb{R}$, the Nash equilibrium strategy

$$\bar{u}(s, y) = \frac{\delta}{\sigma y} \frac{1}{2\alpha} e^{-r(T-s)} \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}, \quad (31)$$

given by the restriction of (26) to $[t, T] \times \mathbb{R}$, is not locally tail-optimal at t for $P_{t,x}^{MV}$ given by (28).

(ii) The Nash equilibrium investor who adopts the Nash equilibrium strategy $\bar{u}(s, X_s)$ (where \bar{u} is given by (26)) over $[t_0, T]$ is not locally preferences-consistent at t with respect to P_{t_0, x_0}^{MV} given by (23) for any $t \in [t_0, T]$.

Proof. The proof is in the Appendix. □

Proposition 5.3 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 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

$$\text{Problem } \mathcal{P}_{t_0, x_0}^{L-ass-MV} : \sup_{u \in \mathcal{U}} \mathbb{E}_{t_0, x_0} \left[-\frac{1}{2\alpha} e^{-2\alpha X_T} \right] \quad (32)$$

coincides with the Nash-equilibrium strategy (26) (see also Basak & Chabakauri (2010), Remark 1). In other words, $U^4(x) = 0$ and $U^5(x) = -1/(2\alpha)e^{-2\alpha x}$ (the CARA utility function). Therefore, Propositions 4.5 and 4.6 hold considering the linear problem (32) and its obvious version $\mathcal{P}_{t, x}^{L-ass-MV}$ at time t with wealth x .

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} as in (23). This is formalized by the following corollary.

Corollary 5.4. *The control map*

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

given by (29) is not globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} .

The control map

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

given by (30) is not globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} .

The control map

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

given by (31) is not globally tail-optimal over $[t_0, T]$ for P_{t_0, x_0}^{MV} .

Proof. The proof is obvious. □

6 Concluding remarks

When an intertemporal stochastic optimal control problem over a time period $[t_0, T]$ 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 decision-maker who adopts the optimal strategy is globally preferences-consistent over $[t_0, T]$ with respect to his initial preferences.

When an intertemporal stochastic optimal control problem 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 decision makers. The price is different depending on the approach selected.

With the precommitment approach, the investor solves a kind of static Markovitz 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 his initial preferences.

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

In general, 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 and for all non-linear optimization problems. Rather the opposite: each approach has its own pro and contra, and the appropriate strategy depends not only on the non-linear optimizing

criterion but also on other subjective factors, such as the attitude towards tail optimality and consistency to one's own preferences. A normative approach that pretends to be universal fails to provide convincing arguments that hold for all individuals, while we believe that a philosophical approach to discuss appropriateness of each approach is more suitable.

Appendix

Proof of Proposition 5.1

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

$$\hat{u}_{t,x}(s, y) = \frac{\delta}{\sigma y} \left[x e^{r(s-t)} - y + \frac{1}{2\alpha} e^{\delta^2(T-t) - r(T-s)} \right], \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}. \quad (33)$$

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 (29) 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} he should play the first control action of the control map $\hat{u}_{t,x}(s, y)$ given by (33). 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

Proof of Proposition 5.2

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

$$\hat{u}_{t,x}(s, y) = \frac{\delta}{\sigma y} \left[x e^{r(s-t)} - y + \frac{1}{2\alpha} e^{\delta^2(T-t) - r(T-s)} \right], \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}. \quad (34)$$

At time t with wealth x the dynamically optimal strategy coincides with the optimal map (34): $\tilde{u}(t, x) = \hat{u}_{t,x}(t, x)$. However, after time t there is no longer coincidence between dynamically optimal strategy and optimal map (34): 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

Proof of Proposition 5.3

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

$$\hat{u}_{t,x}(s, y) = \frac{\delta}{\sigma y} \left[x e^{r(s-t)} - y + \frac{1}{2\alpha} e^{\delta^2(T-t) - r(T-s)} \right], \quad \text{for } (s, y) \in [t, T] \times \mathbb{R}. \quad (35)$$

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 (31) 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} he should play the first control action of the control map $\hat{u}_{t,x}(s, y)$. 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

References

- Basak, S. & Chabakauri, G. (2010). Dynamic mean-variance asset allocation, *Review of Financial Studies* **23**: 2970–3016.
- Bertsekas, D. (2012). *Dynamic Programming and Optimal Control*, Athena Scientific.
- Björk, T. (1998). *Arbitrage Theory in Continuous Time*, Oxford University Press.
- Björk, T. & Murgoci, A. (2010). A general theory of markovian time inconsistent stochastic control problems, *Preprint, SSRN: <http://ssrn.com/abstract=1694759>*.
- 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.

- Epstein, L. & Le Breton, M. (1993). Dynamically consistent beliefs must be Bayesian, *Journal of Economic theory* **61**: 1–22.
- 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. (2016). Dynamic approaches for some time inconsistent problems, *arXiv:1604.03913 [math.OC]*.
- 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.
- Vigna, E. (2016). On time consistency for mean-variance portfolio selection, *Collegio Carlo Alberto Notebook 476*.
- 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.