Time consistency and moving horizons for risk measures

Samuel N. Cohen
University of Adelaide
samuel.cohen@adelaide.edu.au

Robert J. Elliott*
University of Adelaide and University of Calgary
robert.elliott@adelaide.edu.au

September 11, 2009

Abstract

Decision making in the presence of randomness is an important problem, particularly in finance. Often, decision makers base their choices on the values of 'risk measures' or 'nonlinear expectations'; it is important to understand how these decisions evolve through time. In this paper, we consider how these decisions are affected by the use of a moving horizon, and the possible inconsistencies that this creates. By giving a formal treatment of time consistency without Bellman's equations, we show that there is a new sense in which these decisions can be seen as consistent.

1 Introduction

An active area of study is that of risk. The management of uncertain outcomes, and decision making in this context, is of considerable importance. Much recent research has focussed around properties of 'coherent risk measures', as first discussed in [1], and 'convex risk measures', as defined by [10] and [11]. These are functionals $\rho : L^1(\mathcal{F}_T) \to \mathbb{R}$, where T is some future time, and $L^1(\mathcal{F}_T)$ is the space of integrable $\mathcal{F}_T$-measurable random variables. In the convex case, it is assumed that these functionals satisfy three assumptions, namely:

1. Monotonicity: $X \geq Y, \mathbb{P}$-a.s. $\Rightarrow \rho(X) \leq \rho(Y)$,
2. Translation invariance: $\rho(X + c) = \rho(X) - c$ for all $c \in \mathbb{R}$,
3. Convexity: $\rho(\lambda X + (1 - \lambda)Y) \leq \lambda\rho(X) + (1 - \lambda)\rho(Y)$ for all $\lambda \in [0, 1]$.

One significant flaw with these risk measures is that they are essentially static – they consider only one random outcome, and do not model the development

*Robert Elliott wishes to thank the Australian Research Council for support.

of information through time. Simply applying these risk measures to a multiple-period problem is insufficient, as there is no guarantee that they will lead to time-consistent decision making. In particular, there is no guarantee that Bellman's principle will be satisifed. Concrete examples of this can be found in [14] and [2].

More recently, the results in [2] showed that a particular expression of Bellman's principle is equivalent to a recursivity property of the risk measures, namely if $\rho_t(X)$ denotes the risk of X as considered at time t, then for any $s < t$, we have $\rho_s(X) = \rho_s(-\rho_t(X))$. In [15], an equivalent property, (given translation invariance), is considered. Specifically, in [15] a type of inter-temporal monotonicity is assumed, that is, for any times $s < t$, $\rho_t(X) \geq \rho_t(Y)$ $\mathbb{P}$-a.s. implies $\rho_s(X) \geq \rho_s(Y)$ $\mathbb{P}$-a.s. In this paper, we show that a version of this monotonicity condition is equivalent to a general form of Bellman's principle, see Theorem 3.2.

Much has been written on *dynamic risk measures*, that is, risk measures where a recursivity property is satisfied. See, for example, [18], [3], [9]. Similarly, a theory of 'time-consistent nonlinear expectations' has been developed. See particularly [16] and the references therein. These satisfy assumptions very similar to those of dynamic risk measures, the main difference being a sign change in each of the three assumptions above. To construct these functionals, a common tool is the theory of Backward Stochastic Differential Equations (BSDEs), and it is known that all nonlinear expectations, (satisfying some constraints), can be expressed as solutions of BSDEs, (see [8] and [12] in the continuous time case, [7] and [6] in the discrete time case).

To apply these methods, one must typically fix a distant point in the future, (possibly infinitely distant), at which all payoffs will be realised. In many applications, predicting even the distribution of extremely long-term behaviour is almost impossible, so one might hope to use a shorter-dated moving horizon, where payoffs at some fixed time into the future, (say, one-year from the present), are considered, but this horizon is allowed to move forward as time progresses. In this paper, we discuss the consistency properties of decision making under such a regime.

This work is also motivated by applications in economic regulation. In many risk management settings the 'risk' is calculated over some finite horizon, to ensure it does not exceed certain bounds. For example, in the Basel II Banking accords, regulators calculate a ten-day 99%-value-at-risk for market risk, and a one-year 99.9%-value-at-risk for credit and operational risks. See [13] for more details. Even if these risks are calculated using a dynamic risk measure, (which, as is well known, value-at-risk is not), the moving horizon will introduce inconsistencies into the analysis.

Some related ideas have already been discussed in [5], though in a different context, where the time-inconsistency comes about through direct modification of the classical Hamilton-Jacobi-Bellman equations. The analysis of [5] also approaches the issue of inconsistency very differently to this paper, by looking at the situation where decisions are made using a time-inconsistent value function, but future behaviour is taken into account. This paper presents a directly 'time-inconsistent' approach to the specific problem where inconsistency arises from a moving horizon, and where decisions are made without consideration of future behaviour except through the value function. In doing so, this paper also develops an understanding of time-inconsistency for general policy spaces, including those with commitment.

Working with a specific example, based on a classical financial market, we shall show that decision making with a moving horizon is not time-consistent. We shall then show that there exists a modified version of time-consistency which is satisfied, given certain assumptions on the possible policy space.

2 An investment policy model

We consider here a probability space based on a classical model of a financial market in discrete time.

We suppose that there are d risky assets $\{S^i\}$ defined on some filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. We also assume the existence of a 'risk-free' asset. However, for simplicity, we shall assume that the risk free interest rate is zero, or, equivalently, that all quantities have been appropriately discounted. We assume that there are no transaction costs.

We assume that all positions will be closed out at or before some distant, but finite, time T. As we are working in discrete time, we can, without loss of generality, assume that time is indexed by values in the set $\{0, 1, ..., T\}$. The values of the risky assets at these times will be the $\mathcal{F}_T$-measurable random variables S^i_T.

We now state the following general definition, due to Peng (eg [16], [17]).

Definition 2.1 (Nonlinear Expectations)**.** *A system of operators*

$$\mathcal{E}(\cdot|\mathcal{F}_t) : L^1(\mathcal{F}_T) \to L^1(\mathcal{F}_t)$$

is called an $\mathcal{F}_t$-consistent nonlinear expectation if it satisfies the following properties.

1. *(Monotonicity) If $Q \geq Q'$ $\mathbb{P}$-a.s. then $\mathcal{E}(Q|\mathcal{F}_t) \geq \mathcal{E}(Q'|\mathcal{F}_t)$, with $\mathcal{E}(Q|\mathcal{F}_t) = \mathcal{E}(Q'|\mathcal{F}_t)$ only if $Q = Q'$ $\mathbb{P}$-a.s.*
2. *(Constant invariance) For $Q \in L^1(\mathcal{F}_t)$, $\mathcal{E}(Q|\mathcal{F}_t) = Q$.*
3. *(Recursivity) For any $s \leq t$, $\mathcal{E}(\mathcal{E}(Q|\mathcal{F}_t)|\mathcal{F}_s) = \mathcal{E}(Q|\mathcal{F}_s)$ $\mathbb{P}$-a.s.*
4. *(Zero-one law) For any $A \in \mathcal{F}_t$, $I_A\mathcal{E}(Q|\mathcal{F}_t) = \mathcal{E}(I_AQ|\mathcal{F}_t)$.*

Remark 2.1. In [7] and [6], we have given a constructive representation result for these operators in discrete time using the theory of BSDEs. These results give a complete description of nonlinear expectations (and the more general class of nonlinear evalutations) in this context. In the present work, these results are of limited use, and we shall simply proceed by assuming that a nonlinear expectation is given.

Standard theory then tells us that, given that the risk-free interest rate is zero, the assets $\{S^i\}$ are arbitrage free if there exists a probability measure $\tilde{\mathbb{P}}$ equivalent to $\mathbb{P}$ under which the asset prices are all martingales. Equivalently, by the assumption that interest rates are zero, we can assume that prices are given by an expectation, which is simply a special case of the nonlinear expectations defined above.

If the market is incomplete, then there may be multiple equivalent martingale measures; for simplicity we shall take a single measure as given. (As all the assets we shall consider will be combinations of the S^i, this gives no contradiction.)

Let $\mathbb{U}$ denote the set of d-dimensional adapted processes. A firm wishing to invest in this market has a range of self-financing policies available, which is a subset $\mathcal{U} \subseteq \mathbb{U}$. For any policy $X \in \mathcal{U}$, $X_t(\omega)$ is a vector in $\mathbb{R}^d$, where each component indicates the amount invested in each risky asset, and hence for each policy X, an investor's wealth process V^X satisfies the stochastic difference equation

$$V_{t+1}^X = \langle X_t, \mathbf{S}_{t+1} - \mathbf{S}_t \rangle + V_t^X \tag{1}$$

where $\mathbf{S}$ denotes the vector of risky asset prices. (The risk-free asset could also be included, but as the risk-free interest rate is zero, it would not affect the dynamics of V^X.) For notational simplicity, we extend V beyond time T by setting $V_u^X = V_T^X$ for all $u > T$. Note that a policy X_t describes the choice to be made under every contingency, and is not required to be Markovian or of feedback form.

At each time t, an investor has a range of policies available. As an investor cannot modify their past behaviour, they clearly can only consider those policies which agree with what they have previously done. That is, if their past policy has been given by $X_s, s < t$, then the range of policies available at time t is

$$\mathcal{U}|_t^X = \{X' \in \mathcal{U} : X'_s = X_s \text{ for all } s < t\}.$$

For simplicity, we note that $\mathcal{U}|_0^{\hat{X}} = \mathcal{U}$.

Note that this definition requires that the past policy is matched both in the observed past and in all possible other pasts, (that is, for all ω). This requirement is needed to ensure that switching, at time t, from one policy X to another in $\mathcal{U}|_t^X$ results in a policy which is in $\mathcal{U}$. Furthermore, we shall assume that making different choices of policies in different states of the world results in a policy. More technically, for any time t, we have the 'pasting property'

$$I_A X + I_{A^c} X' \in \mathcal{U}|_t^{\hat{X}} \quad \text{for all} \quad X, X' \in \mathcal{U}|_t^{\hat{X}}, A \in \mathcal{F}_t. \tag{2}$$

Definition 2.2. *Let $\{X^t\}$ be a collection containing a policy choice $X^t \in \mathcal{U}$ for each time $t \leq T$. Let $\hat{X}$ denote that policy which is eventually chosen, that is $\hat{X}_u := X_u^u$. Then this policy choice is* viable *if, for every $s < t$,*

$$X^t \in \mathcal{U}|_s^{X^s}$$

or equivalently

$$X^t \in \mathcal{U}|_t^{\hat{X}} = \bigcap_{s<t} \mathcal{U}|_s^{X^s} = \mathcal{U}|_t^{X^{t-1}}.$$

Intuitively, we think of X^t as the policy which an investor intends to pursue, when making a selection at time t. A collection being viable ensures that $\hat{X}$, the policy that is finally chosen, does not involve an investor attempting to change their past actions at any time. The following result ensures that $\hat{X}$ is in fact a policy, that is, $\hat{X} \in \mathcal{U}$.

Lemma 2.1. *If $\{X^t\}$ is a viable policy choice, then $X^T = \hat{X}$. Hence $\hat{X} \in \mathcal{U}$.*

Proof. This is easily seen from the fact that, for all t, $X^t \in \mathcal{U}|_t^{\hat{X}}$, which implies that $X_u^t = \hat{X}_u$ for all $u < t$. Hence $X_u^T = \hat{X}_u$ for all $u < T$. By definition $X_T^T = \hat{X}_T$, the result follows as $X^T \in \mathcal{U}|_T^{\hat{X}} \subseteq \mathcal{U}$. □

At each time t, an investor wishes to choose the 'time-t-optimal' policy $X^t \in \mathcal{U}|_t^{\hat{X}}$. We shall model their decision as based on a value function given by a *dynamic* risk measure $\rho_t(\cdot)$, or equivalently, by an $\mathcal{F}_t$-consistent nonlinear expectation $\mathcal{E}(\cdot|\mathcal{F}_t)$. For each policy, at each time t, we shall consider the case where this nonlinear expectation evaluates the value of the portfolio at a future date $t+m$, where m is the 'horizon distance'. That is, the optimal policy is chosen to maximise the time-t value function

$$\mathcal{V}_t(X^t) = \mathcal{E}(V_{t+m}^{X^t}|\mathcal{F}_t) = -\rho_t(V_{t+m}^{X^t}).$$

We consider $\mathcal{V}_t$ as assigning a value to every policy X, where X details all past and future behaviour under each contingency. Note that for every policy X, the value function $\mathcal{V}_t(X)$ is a random variable in $L^1(\mathcal{F}_t)$. At each time t, we wish to maximise this random variable uniformly, that is, to find a policy $X \in \mathcal{U}|_t^{\hat{X}}$ such that, for any $X' \in \mathcal{U}|_t^{\hat{X}}$,

$$\mathcal{V}_t(X)(\omega) \geq \mathcal{V}_t(X')(\omega)$$

for almost all ω. The existence of such a policy will be given, under some conditions, by Theorem 3.1

We emphasise at this point that we have chosen our value function such that the time-inconsistency in this problem arises purely because of the short horizon. The nonlinear expectation itself is time-consistent, in the sense of [2]. However, the nonlinear expectation is not being evaluated on the terminal values, which, as we shall see, leads to inconsistencies.

3 Time-consistency and Bellman's Principle

We now digress slightly from this model, to present a general definition of time-consistency, which is essentially a formalisation of Bellman's Principle of Optimality. While taking Bellman's Principle as a useful basis for a definition of time consistency, we shall not assume that the value function is the solution of Bellman's equation and, hence, the problems considered may not be time-consistent.

In general, we assume that there is a set of allowable policies $\mathcal{U}$, and that policies selected must be viable, in the sense of Definition 2.2. They are selected to optimise some value function $\mathcal{V}$, which is in general a family of maps

$$\mathcal{V}_t : \mathcal{U} \to L^1(\mathcal{F}_t), \quad t \in \{0, 1, ...T\}.$$

For simplicity, we take higher values of $\mathcal{V}$ as better than lower.

Definition 3.1. *In general, we shall say that our problem is* standard *if*

- *The policy space $\mathcal{U}$ is a compact subset of $\mathbb{U}$, the class of adapted processes taking values in some (real) Hilbert space $\mathfrak{U}$. It follows that $\mathbb{U}$ is itself a Hilbert space with inner product*

$$(X, X') = \sum_{0 \leq t \leq T} E[\langle X_t, X'_t \rangle]$$

for $\langle \cdot, \cdot \rangle$ the inner product in $\mathfrak{U}$. We shall assume that $\mathbb{U}$ is separable.

- *For all t, the value function $\mathcal{V}_t : \mathcal{U} \to L^1(\mathcal{F}_t)$ is lower semicontinuous, except possibly for a countable set of policies $\mathcal{D} \subseteq \mathcal{U}$, that is, if X^n is a sequence in $\mathcal{U}$ converging to $X^\infty \notin \mathcal{D}$, and $\mathcal{V}_t(X^n) \leq \mathcal{V}_t(X^{n+1})$ $\mathbb{P}$-a.s. for all n, then $\lim_{n\to\infty} \mathcal{V}_t(X^n) = \mathcal{V}_t(X^\infty)$.*
- *For all t, $\mathcal{U}|_t^{\hat{X}}$ satisfies the pasting property (2) and $\mathcal{V}_t$ satisfies the zero-one law*

$$\mathcal{V}_t(I_A X + I_{A^c} X') = I_A \mathcal{V}_t(X) + I_{A^c} \mathcal{V}_t(X') \quad \text{for all} \quad X, X' \in \mathcal{U}|_t^{\hat{X}}, A \in \mathcal{F}_t. \tag{3}$$

Lemma 3.1. *If $\mathcal{U}$ is compact, then $\mathcal{U}|_t^{\hat{X}}$ is compact for all times t and all policies $\hat{X}$.*

Proof. As we are in a Hilbert space, it is sufficient to show that $\mathcal{U}|_t^{\hat{X}}$ is sequentially compact. For any sequence $\{X^n\}$ in $\mathcal{U}|_t^{\hat{X}}$, we know that $X_u^n = \hat{X}_u$ for all $u < t$. As $\mathcal{U}$ is compact, there exists a convergent subsequence of X^n. Given the inner product on $\mathbb{U}$, it is trivial to show that the limit of this subsequence X^∞ must satisfy $X_u^\infty = \hat{X}_u$ for all $u < t$, and hence $X^\infty \in \mathcal{U}|_t^{\hat{X}}$. Therefore $\mathcal{U}|_t^{\hat{X}}$ is sequentially compact, and hence compact. □

Theorem 3.1. *For a standard problem, for all t and any past policy $\hat{X}$, there exists a policy $X^t \in \mathcal{U}|_t^{\hat{X}}$ such that for any $X \in \mathcal{U}|_t^{\hat{X}}$*

$$\mathcal{V}_t(X^t) \geq \mathcal{V}_t(X), \qquad a.s.$$

In this case we say that X^t uniformly maximises $\mathcal{V}_t$ on $\mathcal{U}|_t^{\hat{X}}$.

Proof. As $\mathcal{U}|_t^{\hat{X}}$ is in $\mathbb{U}$, which is separable, it has a countable dense subset $\tilde{X}^n$. We assume without loss of generality that the (countable) set $\mathcal{D} \cap \mathcal{U}|_t^{\hat{X}}$ of points at which $\mathcal{V}_t$ is not lower semicontinuous satisfies $(\mathcal{D} \cap \mathcal{U}|_t^{\hat{X}}) \subset \{\tilde{X}^n\}$. By lower semicontinuity of $\mathcal{V}_t$, for almost any ω, it is clear that

$$\sup_{X \in \mathcal{U}|_t^{\hat{X}}} \mathcal{V}_t(X; \omega) = \sup_n \mathcal{V}_t(\tilde{X}^n; \omega). \tag{4}$$

We define a recursive sequence $\bar{X}^n$ with initial condition $\bar{X}^1 = \tilde{X}^1$, and

$$\bar{X}^{n+1} = I_{A_n} \bar{X}^n + I_{A_n^c} \tilde{X}^{n+1}$$

where

$$A_n = \{\omega : \mathcal{V}_t(\bar{X}^n; \omega) \geq \mathcal{V}_t(\tilde{X}^{n+1}; \omega)\}.$$

By measurability of $\mathcal{V}_t$, $A \in \mathcal{F}_t$, and so by the pasting property (2), $\{\bar{X}^n\} \subseteq \mathcal{U}|_t^{\hat{X}}$. As $\mathcal{V}_t$ satisfies the zero-one law (3), we have

$$\mathcal{V}_t(\bar{X}^{n+1}) = I_{A_n} \mathcal{V}_t(\bar{X}^n) + I_{A_n^c} \mathcal{V}_t(\tilde{X}^{n+1}) \geq \mathcal{V}_t(\bar{X}^n),$$

and by recursion, $\mathcal{V}_t(\bar{X}^n) \geq \mathcal{V}_t(\tilde{X}^n)$, both $\mathbb{P}$-a.s.

As $\{\bar{X}^n\} \subseteq \mathcal{U}|_t^{\hat{X}}$, which is compact by Lemma 3.1, there exists a convergent subsequence of $\bar{X}^n$, with limit denoted X^t. By lower semicontinuity of $\mathcal{V}_t$, it follows that $\mathcal{V}_t(X^t) \geq \mathcal{V}_t(\bar{X}^n)$ $\mathbb{P}$-a.s. for all n. It follows from (4) that $\mathcal{V}_t(X^t) \geq \mathcal{V}_t(X)$ for all $X \in \mathcal{U}|_t^{\hat{X}}$, for almost all ω. □

Definition 3.2. *Let $\{X^t\}$ be a viable collection containing a policy choice for each time t. This generates a realised policy $\hat{X}$, defined by $\hat{X}_t = X_t^t$ for all t. This collection is called* optimal *if*

(i) for any t, the policy X^t uniformly maximises $\mathcal{V}_t(X)$ for $X \in \mathcal{U}|_t^{\hat{X}}$,

and time consistent *if*

(ii) for any time t, we have

$$\mathcal{V}_t(X^t) = \mathcal{V}_t(\hat{X}), \qquad \mathbb{P}\text{-}a.s.$$

Remark 3.1. This definition corresponds closely to Bellman's statement:

> "Principle of Optimality: An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision." [4, p. 83]

Unlike most interpretations of Bellman's principle, this definition is 'forward looking', and does not, in general, lead to Bellman's equation. Definition 3.2 directly allows initial behaviour to affect future behaviour in complex ways. (This idea is embedded in the assumption of viability and the freedom to specify which policies are in $\mathcal{U}$.) The definition, therefore, looks for consistent behaviour contingent on what has already been done.

Essentially, if we choose an optimal policy today, we simply need to check that, in the future, we shall continue to follow a policy which we consider equivalent to the optimal choice today. In some sense, a policy is time consistent if it leads to a '*commitment to previous decisions*'.

This definition has the distinct disadvantage of not requiring us to ensure that our decisions today will make us happy in the future. The policies selected as optimal in the future only need to lie in $\mathcal{U}|_t^{\hat{X}}$, that is, in the space of policies we have left ourselves to choose from. A simple example of this is when the space $\mathcal{U}$ consists only of 'buy-and-hold' policies. Here we make a decision at time zero, and are unable to modify it at any point in the future – the policy X^0 chosen at time zero is the only policy in $\mathcal{U}|_t^{X^0} = \mathcal{U}|_t^{\hat{X}}$ for all $t > 0$. Hence this decision is time consistent, as we will not be able to change our minds at any point in the future!

Remark 3.2. It is important to note that, if $\{X^t\}$ is optimal, for any t, as X^t maximises $\mathcal{V}_t$ and

$$\hat{X} \in \mathcal{U}|_t^{\hat{X}},$$

Property (iii) of Definition 3.2 can only ever fail through a future decision appearing sub-optimal today, that is, it is always true that

$$\mathcal{V}_t(X^t) \geq \mathcal{V}_t(\hat{X}), \qquad \mathbb{P}\text{-a.s.}$$

It will prove useful to extend these notions to stopping times. In this context, we shall assume that all stopping times τ take values in the set $\{0, 1, ..., T\}$ $\mathbb{P}$-a.s.

Definition 3.3. *Let τ be a stopping time. We can then extend $\mathcal{V}$ pathwise, yielding for any fixed policy X, the value evaluated at time τ, $\mathcal{V}_\tau(X) = \mathcal{V}_{\tau(\omega)}(X;\omega)$,* *and for any policy choice $\{X^t\}$, the value evaluated at time τ with the corresponding policy choice, $\mathcal{V}_\tau(X^\tau) = \mathcal{V}_{\tau(\omega)}(X^{\tau(\omega)};\omega)$.*

Furthermore, we can say that $X^\tau \in \mathcal{U}|_\tau^{\hat{X}}$ if $X^{\tau(\omega)} \in \mathcal{U}|_{\tau(\omega)}^{\hat{X}}$ for almost all ω. With an abuse of terminology, we shall call X^τ a policy if $X^{\tau(\omega)} \in \mathcal{U}$ $\mathbb{P}$-a.s.

We now give a relation between time consistency in the sense of Definition 3.2 and a type of intertemporal monotonicity for the value function. For simplicity, we write $[s,t[$ for the discrete collection of times $\{s, s+1, ..., t-1\}$, and similarly for stopping times.

Theorem 3.2. *The following statements are equivalent*

(i) For every initial policy set $\mathcal{U}$, $\mathcal{V}$ is such that every optimal policy choice, (i.e. a choice satisfying Properties (i) and (ii) of Definition 3.2), is also time-consistent, (Property (iii)).

(ii) For any policies $X, X' \in \mathbb{U}$ and any stopping times $\sigma < \tau$, if $X_u = X'_u$ for all $u \in [0,\tau[$ and $\mathcal{V}_\tau(X) \geq \mathcal{V}_\tau(X')$ a.s. then $\mathcal{V}_\sigma(X) \geq \mathcal{V}_\sigma(X')$ a.s.

In this case, we shall say that $\mathcal{V}_\sigma$ is nondecreasing with respect to the future values.

(iii) For all stopping times $\sigma < \tau$, every past policy $\hat{X}$, when restricted to $X \in \mathcal{U}|_\sigma^{\hat{X}}$, $\mathcal{V}_\sigma(X)$ can be written as a functional of $\{X_u\}_{u\in[\sigma,\tau[}$ and $\mathcal{V}_\tau(X)$. This functional is nondecreasing in $\mathcal{V}_\tau(X)$, in the sense of (ii).

Proof. We interpret all (in-)equalities as $\mathbb{P}$-a.s.

(i implies ii.) Assume our policy space is given by $\mathcal{U} = \{I_A X + I_{A^c} X' : A \in \mathcal{F}_t\}$. Note that as $X_u = X'_u$ for all $u \in [0,\tau[$, we have $\mathcal{U}_\tau^{\hat{X}} = \mathcal{U}_\tau^X = \mathcal{U}_\tau^{X'}$. Then at time τ, if $\mathcal{V}_\tau(X) \geq \mathcal{V}_\tau(X')$ we will find $X^\tau = X$ is an optimal policy. This implies that $\hat{X} = X$, as $X_u = X'_u$ for $u \in [0,\tau[$. Hence by time consistency,

$$\mathcal{V}_\sigma(X') \leq \mathcal{V}_\sigma(X^\sigma) = \mathcal{V}_\sigma(\hat{X}) = \mathcal{V}_\sigma(X).$$

(ii implies i.) Let σ be the stopping time given by the first time that

$$\mathcal{V}_t(X^t) = \mathcal{V}_t(\hat{X})$$

for all $t > \sigma$. By Lemma 2.1, $\sigma < T$.

As $\{X^t\}$ is optimal, and $X^{\sigma+1} \in \mathcal{U}|_\sigma^{X^\sigma}$, we know that

$$\mathcal{V}_{\sigma+1}(X^\sigma) \leq \mathcal{V}_{\sigma+1}(X^{\sigma+1}) = \mathcal{V}_{\sigma+1}(\hat{X})$$

By (ii), this implies that

$$\mathcal{V}_\sigma(X^\sigma) \leq \mathcal{V}_\sigma(\hat{X}).$$

As X^σ is optimal, this implies $\mathcal{V}_\sigma(X^\sigma) = \mathcal{V}_\sigma(\hat{X})$.

Therefore, if $\mathcal{V}_t(X^t) = \mathcal{V}_t(\hat{X})$ for all $t > \sigma$, it follows that $\mathcal{V}_\sigma(X^\sigma) = \mathcal{V}_\sigma(\hat{X})$. By induction, this must hold for all times, that is, the optimal choice is consistent.

(ii implies iii.) By a double application of (ii), we know that for a fixed $s < t$, if $X_u = X'_u$ for $u \in [0,t[$, and $\mathcal{V}_t(X) = \mathcal{V}_t(X')$, then $\mathcal{V}_s(X) = \mathcal{V}_s(X')$. Consequently, given $\{X_u\}_{u\in[0,t[}$ and $\mathcal{V}_t(X)$, it is possible to determine $\mathcal{V}_s(X)$

completely. We have restricted our attention to $X \in \mathcal{U}|_s^{\hat{X}}$, which determines X_u for $u < s$, therefore we can determine $\mathcal{V}_s(X)$ from $\{X_u\}_{u\in[s,t[}$ and $\mathcal{V}_t(X)$.

(iii implies ii.) For any two policies with $X_u = X'_u$ for all $u \in [0, \tau[$, we know that $X, X' \in \mathcal{U}|_\sigma^X$. Restricted to this subset of policies, $\mathcal{V}_\sigma$ can be written as a functional of $\{X_u\}_{u\in[\sigma,\tau[}$ and $\mathcal{V}_\tau$, and this functional is nondecreasing in $\mathcal{V}_\tau$. As $X_u = X'_u$ for $u \in [\sigma, \tau[$, we can conclude that if $\mathcal{V}_\tau(X) \geq \mathcal{V}_\tau(X')$ a.s. then $\mathcal{V}_\sigma(X) \geq \mathcal{V}_\sigma(X')$ a.s., as desired. □

Statement (ii) is related to the definition of ('strong') time-consistency considered in [15]. Statement (iii) shows, (as [15] does, but in a different setting), that this is equivalent to a type of recursivity property of the value function, similar to that in Bellman equations. In fact, this gives the following corollary.

Corollary 3.2.1. *The value function given by Bellman's equation is time consistent for* any *initial policy set* $\mathcal{U}$.

Proof. Rephrased into our context, Bellman's equation gives the value function

$$\mathcal{V}_s(X) = f(\omega, s, X_s) + E[\mathcal{V}_{s+1}(X)|\mathcal{F}_s],$$

where f is the negative of the (expected) cost process. By recursion, given past policy $\hat{X}$, this clearly implies that $\mathcal{V}_s(X)$ is a functional only of $\{X_u\}_{u\in[s,t[}$ and $\mathcal{V}_t(X)$, and hence, statement (iii) of Theorem 3.2 is satisfied. □

Corollary 3.2.2. *The value functions given by dynamic risk measures and nonlinear expectations are time consistent for* any *initial policy set* $\mathcal{U}$.

Proof. In this context, the policy X chosen can be assumed to only affect the value $\mathcal{V}_t(X)$ through the terminal condition V_T^X. Hence, by the recursivity and monotonicity properties of nonlinear expectations, we can write $\mathcal{V}_s(X)$ as a nondecreasing functional of the future values $\mathcal{V}_t(X)$. Hence statement (ii) or (iii) of Theorem 3.2 is satisfied. □

4 An inconsistent example

We shall now show, through a simple concrete example, that the moving horizon problem considered in Section 2 is not time-consistent for some policy spaces. Let $d = k = 2$, that is, S^1, S^2 are two independent risky assets, both satisfying the stochastic difference equations

$$\frac{S_u^i - S_{u-1}^i}{S_u^i} = \mu^i(u) + B_u^i,$$

where B^1, B^2 are two independent martingales with $P[B_u^i > 1|\mathcal{F}_{u-1}] > 0$ a.s. for all u. Now assume $\mu^1(u) = 0.05$ for all u and

$$\mu^2(u) = \begin{cases} 0.10 & u < 10 \\ -0.99 & u \geq 10. \end{cases}$$

We shall write μ_u for the vector $(\mu^1(u), \mu^2(u))^*$.

Our value function is simply given by the conditional expectation of the wealth process at the horizon,

$$\mathcal{V}_t(X) = E[V^X_{t+m}|\mathcal{F}_t].$$

Now suppose that the policy space $\mathcal{U}$ consists of all adapted processes X in $\mathbb{R}^2$ which are constant from time 5 onwards, and take nonnegative values uniformly bounded by $\|X_u\| \leq 1$. It is easy to show that, given a policy X, the expected value of V^X_{t+m} is given by

$$E[V^X_{t+m}|\mathcal{F}_t] = V_t \cdot E\left[\prod_{u=t+1}^{t+m}(1 + \langle X_u, \mu_u\rangle)\middle|\mathcal{F}_t\right].$$

Now assume $m = 10$. At time $t = 0$, the value function is uniquely maximised by taking $X^0_u = e_2$ for all u, where e_i denotes the ith standard basis vector in $\mathbb{R}^2$. However,the policy chosen at time $t = 5$ must be committed to for all times $t > 5$, and the significant downward trend of S^2 which will occur from time 10 onwards is sufficient to ensure that $X^5_u = e_1$ for $u > 5$. It is straightforward to see that this results in the realised policy

$$\hat{X}_u = \begin{cases} e_2 & u < 5 \\ e_1 & u \geq 5. \end{cases}$$

Therefore, we have

$$\mathcal{V}_0(X^0) > \mathcal{V}_0(\hat{X}),$$

and hence the optimal policy selection is not time-consistent.

Remark 4.1. Note that, in some sense, the inconsistency in this example is due to the commitment to a policy from time 5 onwards. This counterintuitive behaviour, where commitment leads to inconsistency, is typical of decision making under a moving horizon. See Remark 6.6.

5 An equivalence for policies

We now return specifically to the moving-horizon situation described in Section 2. As the above example demonstrates, one cannot, in general, assume that this moving-horizon problem will have time-consistent optimal policies. To obtain a positive result, we have the following useful lemmata.

Lemma 5.1. *We define*

$$I_{[0,t+m]}\mathcal{U}|^{\hat{X}}_t = \{I_{[0,t+m]}X | X \in \mathcal{U}|^{\hat{X}}_t\},$$

that is, the set of policies in $\mathcal{U}|^{\hat{X}}_t$, with the value zero assigned from time $t+m$ onwards. Then, at any time t, for any policy $X \in I_{[0,t+m]}\mathcal{U}|^{\hat{X}}_t$ we have the identity

$$\mathcal{V}_t(X) = \mathcal{E}(V^X_{t+m}|\mathcal{F}_t) = \mathcal{E}(V^X_T|\mathcal{F}_t) \qquad \mathbb{P}\text{-a.s.}$$

Proof. As $X \in I_{[0,t+m]}\mathcal{U}|^{\hat{X}}_t$, we know that $X_u = \mathbf{0}$ for $u > t+m$. Hence $V^X_u = V^X_{u-1}$ for $u > t+m$, and therefore $V^X_{t+m} = V^X_T$ $\mathbb{P}$-a.s. The result follows. □

Lemma 5.2. *For any policy $X \in \mathcal{U}$, any time t,*

$$\mathcal{V}_t(X) = \mathcal{V}_t(I_{[0,t+m]}X) \quad \mathbb{P}\text{-}a.s.$$

Proof. As $X_u = I_{[0,t+m]}X_u$ for all $u < t+m$, it follows

$$V^X_{t+m} = V^{I_{[0,t+m]}X}_{t+m}, \qquad \mathbb{P}\text{-a.s.}$$

the result follows. □

Remark 5.1. We can now consider the 'moving horizon problem' in two distinct ways. We can either consider taking the value function $\mathcal{V}_t(X) = \mathcal{E}(V^X_{t+m}|\mathcal{F}_t)$, in which case we have a time-inconsistent problem. Alternatively, we can equivalently consider taking the value function $\tilde{\mathcal{V}}_t(X) = \mathcal{E}(V^X_T|\mathcal{F}_t)$, and then requiring that our selection X^t must lie in the set $I_{[0,t+m]}\mathcal{U}|^{\hat{X}}_t$ for each t. By Lemma 5.2, we can assume, without loss of generality, that the policy X which maximises $\mathcal{V}_t$ will lie in this set, and by Lemma 5.1, for all such policies we have $\mathcal{V}_t(X) = \tilde{\mathcal{V}}_t(X)$.

That is, we can consider the moving horizon in terms of a restriction on the policy space, rather than in terms of evaluating the wealth process V at the moving horizon. The values associated with each policy under these alternative approaches will be identical. We shall show that, in some sense, this latter problem is time-consistent.

6 Dependable decisions

We now propose a new type of 'time consistency', which we call 'dependability'. One can characterise classical time-consistency through the statement 'a policy X is time consistent if the time-t-optimal policies chosen, for future t, pasted together with X, give *the same value today* as X does.'

Our new definition would then read, 'a policy X is *dependable* if the time-t-optimal policies chosen, for future t, pasted together with X, give *higher values today* than policy X does'. In some sense, dependable policies are those that form a lower bound on the value function, irrespective of future decisions. This is expressed formally by the following definition.

Definition 6.1. *Suppose we have a standard problem, but the policies considered are, at each time $t \geq 0$, restricted to some compact subset $\tilde{\mathcal{U}}|^{\hat{X}}_t \subseteq \mathcal{U}|^{\hat{X}}_t$. Assume $\tilde{\mathcal{U}}|^{\hat{X}}_t$ satisfies the pasting property (2).*

Let $\{X^t\}$ be a viable collection containing a policy choice $X^t \in \tilde{\mathcal{U}}|^{\hat{X}}_t$ for each time t. (Note the definition of viability is not *restricted to $\tilde{\mathcal{U}}$.) This collection is called $\tilde{\mathcal{U}}$*-optimal *if*

(i) for any t, the policy X^t uniformly maximises $\tilde{\mathcal{V}}_t(X)$ for $X \in \tilde{\mathcal{U}}|^{\hat{X}}_t$,

and dependable *if*

(ii) for any time t, we have

$$\tilde{\mathcal{V}}_t(X^t) \leq \tilde{\mathcal{V}}_t(\hat{X}), \qquad \mathbb{P}\text{-}a.s.$$

Remark 6.1. As highlighted by Remark 3.2, when $\tilde{\mathcal{U}}|_t^{\hat{X}} = \mathcal{U}|_t^{\hat{X}}$, this will degenerate into the usual time-consistency properties. Here, on the other hand, our restricted set of policies $\tilde{\mathcal{U}}|_t^{\hat{X}}$, over which we optimise at each time point, can make our problem time-inconsistent.

Remark 6.2. As $\tilde{\mathcal{U}}|_t^{\hat{X}}$ is compact and satisfies the pasting property (2), the argument of Theorem 3.1 applies at each time t to show the existence of a policy X^t uniformly maximising $\mathcal{V}_t$ on $\tilde{\mathcal{U}}|_t^{\hat{X}}$.

Remark 6.3. Under this definition, it is perfectly reasonable that a naïve policy may be selected early on. However, when it is reconsidered later, this decision might be changed. The difference is that this decision is 'dependable' if, had we been allowed to initially consider the decision with the later change, we would have preferred it to the policy initially chosen.

Remark 6.4. Note that as we have now restricted the set of policies which we can consider at any time point, the results of Theorem 3.2 no longer apply.

Definition 6.2. *For a given horizon m, we say that $\mathcal{U}$ is closed under truncation if, for all times t, all past policies $\hat{X}$,*

$$I_{[0,t+m]}\mathcal{U}|_t^{\hat{X}} \subseteq \mathcal{U}|_t^{\hat{X}}.$$

Remark 6.5. This definition essentially states that one can choose to switch at the horizon from a nonzero policy to a zero policy, independently of what policy has been taken up to that point.

We can now give the following positive result for the moving horizon problem.

Theorem 6.1. *Suppose that $\mathcal{U}$ is closed under truncation. Then any policy choice $\{X^t\}$ chosen to uniformly maximise*

$$\tilde{\mathcal{V}}_t(X^t) = \mathcal{E}\left(V_T^{X^t}\middle|\mathcal{F}_t\right), \qquad \text{for } X^t \in I_{[0,t+m]}\mathcal{U}|_t^{\hat{X}} =: \tilde{\mathcal{U}}_t^{\hat{X}}$$

is $\tilde{\mathcal{U}}$-optimal and dependable.

By Lemmas 5.1 and 5.2, as $\mathcal{U}$ is closed under truncation, this will give the same values and policy choices at all times as when using a moving horizon.

Proof. As we are chosing $\{X^t\}$ to maximise $\tilde{\mathcal{V}}_t$ for $X \in \tilde{\mathcal{U}}|_t^{\hat{X}}$, the fact our choice is $\tilde{\mathcal{U}}$-optimal is true by definition. We now show this choice is dependable.

For each t, we choose X^t to maximise $\tilde{\mathcal{V}}_t(X^t) = \mathcal{E}(V_T^{X^t}|\mathcal{F}_t)$, for $X^t \in \tilde{\mathcal{U}}|_t^{\hat{X}}$. We also know that, for any $X \in \tilde{\mathcal{U}}|_{t+1}^{\hat{X}} = \tilde{\mathcal{U}}|_{t+1}^{X^t}$, we have $\tilde{\mathcal{V}}_{t+1}(X) \leq \tilde{\mathcal{V}}_{t+1}(X^{t+1})$. Hence, by the monotonicity of nonlinear expectations, we know that

$$\tilde{\mathcal{V}}_t(X) \leq \tilde{\mathcal{V}}_t(X^{t+1}) \text{ for all } X \in \tilde{\mathcal{U}}|_{t+1}^{\hat{X}}.$$

Specifically, this implies

$$\tilde{\mathcal{V}}_t(X^t) \leq \tilde{\mathcal{V}}_t(X^{t+1}).$$

Similarly, it follows that

$$\tilde{\mathcal{V}}_{t-1}(X^{t-1}) \leq \tilde{\mathcal{V}}_{t-1}(X^t) \leq \tilde{\mathcal{V}}_{t-1}(X^{t+1}).$$

where the last inequality is again by monotonicity of nonlinear expectations.

By induction, this argument shows that for any times $s < t$,

$$\tilde{\mathcal{V}}_s(X^s) \leq \tilde{\mathcal{V}}_s(X^t).$$

Hence, for all $s \leq T$, as by Lemma 2.1 $\hat{X} = X^T$, we have the result

$$\tilde{\mathcal{V}}_s(X^s) \leq \tilde{\mathcal{V}}_s(\hat{X}).$$

□

Remark 6.6. Note that the requirement on $\mathcal{U}$ is that, in some sense, it does *not* enforce commitment, specifically that one can always choose to 'quit at the horizon', that is, to take the truncated policy $I_{[0,t+m]}X$.

Remark 6.7. This construction does not allow situations where warning must be given before quitting, for example, where an investor must signal their intent to sell stocks prior to their sale. This is because such a signal would need to be incorporated into the policy X, and hence the policy $I_{[0,t+m]}X \notin \mathcal{U}$ for any X without the required signal.

7 A dependable but inconsistent example

To demonstrate the usefulness of these results, we give a simple, if contrived, example of a situation where the moving horizon approach is inconsistent, but the equivalent approach using a modified policy space is dependable.

Suppose our market contains only one asset S. The policy space $\mathcal{U}$ consists of those processes X of the form

$$X_u = \begin{cases} 1 & u < \sigma \\ 0 & u \geq \sigma, \end{cases}$$

where σ is a stopping time.

Let $T = 3$, and suppose that values are given by the nonlinear expectation

$$\mathcal{E}(Q|\mathcal{F}_t) = -10 \log E[e^{-Q/10}|\mathcal{F}_t].$$

This is evaluated on a horizon two periods from the present, that is, $m = 2$, and

$$\mathcal{V}_t(X) = \mathcal{E}\left(V^X_{t+m}\middle|\, \mathcal{F}_t\right).$$

Let S follow a non-recombining binomial tree, with independent increments given by

$$S_0 = 20$$

$$S_1 - S_0 = \begin{cases} 1 & \text{w.p. } 0.5 \\ -0.1 & \text{w.p. } 0.5 \end{cases}$$

$$S_2 - S_1 = \begin{cases} 0.1 & \text{w.p. } 0.5 \\ -10 & \text{w.p. } 0.5 \end{cases}$$

$$S_3 - S_2 = \begin{cases} 100 & \text{w.p. } 0.5 \\ -0.1 & \text{w.p. } 0.5 \end{cases}$$

Here w.p. denotes 'with probability'.

It is then easy to see that, in every state of the world ω, the policy chosen at each time will be:

$$X_t^0 = \begin{cases} 1 & t = 0 \\ 0 & t > 0 \end{cases}$$
$$X_t^1 = X_t^2 = X_t^3 = 1 \quad \text{for all } t$$

and therefore $\hat{X}_t = 1$ for all t. Comparing these at time 0, we have

$$\mathcal{V}_0(X^0) = 0.188867122 > -2.492618614 = \mathcal{V}_0(\hat{X}),$$

and so our optimal solution is not time-consistent.

On the other hand, at any time t, the permitted polices allow the choice $X_u = 0$ for $u > t$. That is, $\mathcal{U}$ is closed under truncation, in the sense of Theorem 6.1. Hence we know that this decision is dependable, under an equivalent value function. To show this empirically, we define

$$\tilde{\mathcal{V}}_t(X) = \mathcal{E}\left(V_T^X \middle| \mathcal{F}_t\right)$$

and instead consider, at each time, policies in the restricted set $I_{[0,t+m]}\mathcal{U}|_t^{\hat{X}} = \tilde{\mathcal{U}}_t^{\hat{X}}$. On this set, we know $\mathcal{V}_t(X) = \tilde{\mathcal{V}}_t(X)$, and that the policy which maximises $\mathcal{V}_t$ will lie in this set. We obtain exactly the same optimal policies, but have the values

$$\tilde{\mathcal{V}}_0(X^0) = 0.188867122 < 0.474056692 = \tilde{\mathcal{V}}_0(\hat{X}),$$

and so see that, (given $X^1 = X^2 = X^3 = \hat{X}$), our choice is dependable. Note that $\mathcal{V}_0(X^0) = \tilde{\mathcal{V}}_0(X^0)$, as expected.

8 Conclusions

We have discussed the theory of time-consistency, and have given a definition for a new type of property, that of 'dependability'. We have shown that, for a simple model of a financial market, under some assumptions on the allowable policies, the optimal decision reached using a moving horizon approach is dependable.

This result gives a partial justification for using a moving horizon approach in risk management. Assume that one can always decide to stop investing, (that is, to take the policy $I_{[0,t]}X$). Then one can be sure that the optimal policy today, considering only a finite horizon, will only be improved by future decisions.

This analysis still assumes that the underlying value function used is recursive up to the horizon, in particular, that it is an $\mathcal{F}_t$-consistent nonlinear expectation. This could be weakened to assuming that it is simply a nonlinear evaluation, and with appropriate adaptation of the arguments involved, we can also remove the assumption that interest rates are zero or deterministic. However, if the value function used is not recursive, for example, as with Coherent Value at Risk, these results would not apply. Essentially this is because these value functions introduce different types of time inconsistency, apart from the issues of moving horizons.

Often the question is whether a policy is acceptable, that is, has risk above or below a certain level. This is equivalent to comparing a nonzero policy with a zero policy, looking over a finite horizon. Provided one can always choose to switch from the nonzero policy to the zero policy at the horizon, independently of the choice of policy up to the horizon, our results show that the choice of policy with optimal switching time will be dependable. Therefore, if a given policy is acceptable under the finite horizon, one can be confident that, in some sense, this policy will also be acceptable over the long term.

Given the extreme uncertainties that may be faced when attempting to model asset dynamics in the very long term, it may be appropriate to use a moving horizon approach. At the same time, if decisions involve commitment beyond the horizon, (and hence the policy space is not closed under truncation, in the sense of Theorem 6.1), consideration of the longer term is necessary.

References

[1] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.

[2] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, David Heath, and Hyejin Ku. Coherent multiperiod risk adjusted values and Bellman's principle. *Annals of Operations Research*, 152:5–22, July 2007.

[3] Pauline Barrieu and Nicole El Karoui. Optimal derivatives design under dynamic risk measures. In George Yin and Qin Zhang, editors, *Mathematics of Finance*, volume 351 of *Contemporary Mathematics*, pages 13–26, Providence, RI., 2004. A.M.S. Proceedings.

[4] Richard E. Bellman. *Dynamic Programming.* Princeton University Press, Princeton, 1957.

[5] Tomas Björk and Agatha Murgoci. A general theory of markovian time inconsistent stochastic control problems. forthcoming.

[6] Samuel N. Cohen and Robert J. Elliott. Backward stochastic difference equations and nearly-time-consistent nonlinear expectations. forthcoming, 2009.

[7] Samuel N. Cohen and Robert J. Elliott. A general theory of finite state backward stochastic difference equations. forthcoming, 2009.

[8] François Coquet, Ying Hu, Jean Memin, and Shige Peng. Filtration consistent nonlinear expectations and related g-expectations. *Probability Theory and Related Fields*, 123(1):1–27, May 2002.

[9] Freddy Delbaen, Shige Peng, and Emanuela Rosazza Gianin. Representation of the penalty term of dynamic concave utilities. *Working Paper*, 2008. Available at *http://arxiv.org/abs/0802.1121.*

[10] Hans Föllmer and Alexander Schied. *Stochastic Finance: An introduction in discrete time.* Studies in Mathematics 27. de Gruyter, Berlin-New York, 2002.

[11] Marco Frittelli and Emanuela Rosazza Gianin. Putting order in risk measures. *Journal of Banking & Finance*, 26(7):1473–1486, 2002.

[12] Ying Hu, Jin Ma, Shige Peng, and Song Yao. Representation theorems for quadratic $\mathcal{F}$-consistent nonlinear expectations. *Stochastic Processes and their Applications*, 118:1518–1551, September 2008.

[13] John C. Hull. *Risk Management and Financial Institutions*. Prentice Hall, 2nd edition, 2009.

[14] Boda Kang and Jerzy A. Filar. Time consistent dynamic risk measures. *Mathematical Methods of Operations Research*, 63:169–186, 2006.

[15] Susanne Klöppel and Martin Schweizer. Dynamic indifference valuation via convex risk measures. *Mathematical Finance*, 17(4):599–627, October 2007.

[16] Shige Peng. *Backward Stochastic Differential Equations*, chapter 9: Backward SDE and related g-expectation, pages 141–159. Pitman Research Notes in Mathematics. Longman, 1997.

[17] Shige Peng. *Stochastic Methods in Finance*, chapter 4: Nonlinear Expectations, Nonlinear Evaluations and Risk Measures, pages 165–254. Springer, Berlin-Heidelberg-New York, 2004.

[18] Emanuela Rosazza Gianin. Risk measures via $g-$expectations. *Insurance Mathematics and Economics*, 39:19–34, 2006.