European Journal of Operational Research 249 (2016) 177–187

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Building up time-consistency for risk measures and dynamic optimization ☆

Michel De Lara, Vincent Leclère*

Université Paris-Est, CERMICS (ENPC), 6–8 Avenue Blaise Pascal, Cité Descartes, F-77455 Marne-la-Vallée, France

ARTICLE INFO

Article history:
Received 2 April 2014
Accepted 30 March 2015
Available online 25 April 2015

Keywords:
Dynamic programming
Time-consistency
Dynamic risk measures

ABSTRACT

In stochastic optimal control, one deals with sequential decision-making under uncertainty; with dynamic risk measures, one assesses stochastic processes (costs) as time goes on and information accumulates. Under the same vocable of *time-consistency* (or *dynamic-consistency*), both theories coin two different notions: the latter is consistency between successive evaluations of a stochastic processes by a dynamic risk measure (a form of monotonicity); the former is consistency between solutions to intertemporal stochastic optimization problems. Interestingly, both notions meet in their use of dynamic programming, or nested, equations.

We provide a theoretical framework that offers (i) basic ingredients to jointly define dynamic risk measures and corresponding intertemporal stochastic optimization problems (ii) common sets of assumptions that lead to time-consistency for both. We highlight the role of time and risk preferences — materialized in one-step aggregators — in time-consistency. Depending on how one moves from one-step time and risk preferences to intertemporal time and risk preferences, and depending on their compatibility (commutation), one will or will not observe time-consistency. We also shed light on the relevance of information structure by giving an explicit role to a state control dynamical system, with a state that parameterizes risk measures and is the input to optimal policies.

1. Introduction

You come across *time-consistency* in two different mathematical fields. You are time-consistent if, as time goes on and information accumulates, you do not question the original assessment of stochastic processes (dynamic risk measures) or planning of policies (stochastic optimal control).

We propose a general mechanism to build up time-consistent dynamic risk measures, that serve as criteria for optimal control problems under uncertainty, which henceforth inherit time-consistency. We show how in a few words.

Consider two sets $\mathbb{T}_1$ and $\mathbb{T}_2$, representing sets of time periods ($\mathbb{T}_1 = \{1, 2, 3\}$, $\mathbb{T}_2 = \{4, 5\}$ for instance). Consider two sets $\mathbb{W}_1$ and $\mathbb{W}_2$, representing possible values of uncertainties. For any set $\mathbb{S}$, denote by $\mathcal{L}(\mathbb{S})$ the set of functions $\mathbb{S} \to \mathbb{R} \cup \{+\infty\}$, and by $\mathbb{G}_{\mathbb{S}} : \mathcal{L}(\mathbb{S}) \to \mathbb{R} \cup \{+\infty\}$ a mapping. You can assess any function $A : \mathbb{T}_1 \times \mathbb{T}_2 \times \mathbb{W}_1 \times \mathbb{W}_2 \to \mathbb{R} \cup \{+\infty\}$,

- either by block-aggregation: start by aggregating by time, yielding $\mathbb{G}_{\mathbb{T}_2}\mathbb{G}_{\mathbb{T}_1} : \mathbb{W}_1 \times \mathbb{W}_2 \to \mathbb{R} \cup \{+\infty\}$, then by uncertainty, yielding $\mathbb{G}_{\mathbb{W}_2}\mathbb{G}_{\mathbb{W}_1}\mathbb{G}_{\mathbb{T}_2}\mathbb{G}_{\mathbb{T}_1}A \in \mathbb{R} \cup \{+\infty\}$,
- or by nested-aggregation, yielding $\mathbb{G}_{\mathbb{W}_2}\mathbb{G}_{\mathbb{T}_2}\mathbb{G}_{\mathbb{W}_1}\mathbb{G}_{\mathbb{T}_1}A \in \mathbb{R} \cup \{+\infty\}$.

We will show that nested-aggregation produces both time-consistent dynamic risk measures and optimal control problems, and that so does block-aggregation when a commutation property holds true. For example, sum and integral are commuting operators and a block-aggregation is equivalent to a nested-aggregation as shown in the following equality

$$\iiint_{X\times Y\times Z} [c_1(x) + c_2(x, y) + c_3(x, y, z)]dxdydz$$
$$= \int_X \left[c_1(x) + \int_Y \left[c_2(x, y) + \int_Z c_3(x, y, z)dx\right]dy\right]dz.$$

Now, let us be more specific.

In stochastic optimal control, one deals with sequential decision-making under uncertainty; with dynamic risk measures, one assesses stochastic processes (costs) as time goes on and information accumulates. We discuss the definition of time-consistency in each setting one after the other (see also Rudloff, Street, & Valladão, 2014 for another analysis of links between both notion).

☆ This paper was processed by Guest-Editor prof. Ruediger Schultz for an uncompleted EJOR Feature Cluster on Stochastic Optimization.

* Corresponding author. Tel.: +33674627166.
E-mail address: vincent.leclere@cermics.enpc.fr (V. Leclère).

http://dx.doi.org/10.1016/j.ejor.2015.03.046
0377-2217/

In optimal control problems, we consider a dynamical process that can be influenced by exogenous noises as well as decisions made at every time step. The decision-maker (DM) wants to optimize a criterion (for instance, minimize a net present value) over a given time horizon. As time goes on and the system evolves, the DM makes observations. Naturally, it is generally more profitable for the DM to adapt his decisions to the observations. He is hence looking for policies (strategies, decision rules) rather than simple decisions: a policy is a function that maps every possible history of the observations to corresponding decisions.

The notion of "consistent course of action" (see Peleg & Yaari, 1973) is well-known in the field of economics, with the seminal work of (Strotz, 1955–1956): an individual having planned his consumption trajectory is consistent if, reevaluating his plans later on, he does not deviate from the originally chosen plan. This idea of consistency as "sticking to one's plan" may be extended to the uncertain case where plans are replaced by decision rules ("Do thus-and-thus if you find yourself in this portion of state space with this amount of time left", Richard Bellman cited in Dreyfus, 2002): Hammond (1976) addresses "consistency" and "coherent dynamic choice", Kreps and Porteus (1978) refer to "temporal consistency".

In this context, we loosely state the property of time-consistency in optimal control problems as follows (Carpentier, Chancelier, Cohen, De Lara, & Girardeau, 2012). The decision maker formulates an optimization problem at time t_0 that yields a sequence (planning) of optimal decision rules for t_0 and for the following increasing time steps $t_1, \ldots, t_N = T$. Then, at the next time step t_1, he formulates a new problem starting at t_1, that yields a new sequence of optimal decision rules from time steps t_1 to T. Suppose the process continues until time T is reached. The sequence of optimization problems is said to be time-consistent if the optimal strategies obtained when solving the original problem at time t_0 remain optimal for all subsequent problems. In other words, time consistency means that strategies obtained by solving the problem at the very first stage do not have to be questioned later on.

Now, we turn to dynamic risk measures. At time t_0, you assess, by means of a risk measure $\rho_{t_0,T}$, the "risk" of a stochastic process $\{\boldsymbol{A}_t\}_{t=t_0}^{t_N}$, that represents a stream of costs indexed by the increasing time steps $t_0, t_1, \ldots, t_N = T$. Then, at the next time step t_1, you assess the risk of the tail $\{\boldsymbol{A}_t\}_{t=t_1}^{t_N}$ of the stochastic process knowing the information obtained and materialized by a σ-field $\mathfrak{F}_{t_1}$. For this, you use a conditional risk measure $\rho_{t_1,T}$ with values in $\mathfrak{F}_{t_1}$-measurable random variables. Suppose the process continues until time T is reached. The sequence $\{\rho_{t,T}\}_{t=t_0}^{t_N}$ of conditional risk measures is called a dynamic risk measure.

Dynamic or time-consistency has been introduced in the context of risk measures (see Artzner, Delbaen, Eber, Heath, & Ku, 2007; Cheridito, Delbaen, & Kupper, 2006; Cheridito & Kupper, 2011; Detlefsen & Scandolo, 2005; Riedel, 2004 for definitions and properties of coherent and consistent dynamic risk measures). The dynamic risk measure $\{\rho_{t,T}\}_{t=t_0}^{t_N}$ is said to be time-consistent when the following property holds. Suppose that two streams of costs, $\{\underline{\boldsymbol{A}}_t\}_{t=t_0}^{t_N}$ and $\{\overline{\boldsymbol{A}}_t\}_{t=t_0}^{t_N}$, are such that they coincide from time t_i up to time $t_j > t_i$ and that, from that last time t_j, the risk of the tail stream $\{\underline{\boldsymbol{A}}_t\}_{t=t_j}^{t_N}$ is more than that of $\{\overline{\boldsymbol{A}}_t\}_{t=t_j}^{t_N}$ (i.e. $\rho_{t_j,T}(\{\underline{\boldsymbol{A}}_t\}_{t=t_j}^{t_N}) \geq \rho_{t_j,T}(\{\overline{\boldsymbol{A}}_t\}_{t=t_j}^{t_N})$). Then, the whole stream $\{\underline{\boldsymbol{A}}_t\}_{t=t_i}^{t_N}$ has higher risk than $\{\overline{\boldsymbol{A}}_t\}_{t=t_i}^{t_N}$ (i.e. $\rho_{t_i,T}(\{\underline{\boldsymbol{A}}_t\}_{t=t_i}^{t_N}) \geq \rho_{t_i,T}(\{\overline{\boldsymbol{A}}_t\}_{t=t_i}^{t_N})$).

We observe that both notions of time-consistency look quite different: the latter is consistency between successive risk assessments of a stochastic process by a dynamic risk measure (a form of monotonicity); the former is consistency between solutions to intertemporal stochastic optimization problems. We now stress the role of information accumulation in both notions of time-consistency, because it highlights how the two notions can be connected. For dynamic risk measures, the flow of information is materialized by a filtration $\{\mathfrak{F}_t\}_{t=t_1}^{t_N}$. In stochastic optimal control, an amount of information more modest than the past of exogenous noises is often sufficient to make an optimal decision. In the seminal work of (Bellman, 1957), the minimal information necessary to make optimal decisions is captured in a *state variable* (see Whittle, 1982 for a more formal definition). Moreover, the famous Bellman or *Dynamic Programming Equation (DPE)* provides a theoretical way to find optimal strategies (see Bertsekas, 2000 for a broad overview on *Dynamic Programming (DP)*).

Interestingly, time-consistency in optimal control problems and time-consistency for dynamic risk measures meet in their use of DPEs. On the one hand, in optimal control problems, it is well known that the existence of a DPE with state x for a sequence of optimization problems implies time-consistency when solutions are looked after as feedback policies that are functions of the state x. On the other hand, proving time-consistency for a dynamic risk measure appears rather easy when the corresponding conditional risk measures can be expressed by a *nested* formulation. In both contexts, such nested formulations are possible only for proper information structures. In optimal control problems, a sequence of optimization problems may be consistent for some information structure while inconsistent for a different one (see Carpentier et al., 2012). For dynamic risk measures, time-consistency appears to be strongly dependent on the underlying information structure (filtration or scenario tree). Moreover, in both contexts, nested formulations and the existence of a DPE are established under various forms of decomposability of operators that display monotonicity and commutation properties.

Our objective is to provide a theoretical framework that offers (i) basic ingredients to jointly define dynamic risk measures and corresponding intertemporal optimization problems under uncertainty (ii) common sets of assumptions that lead to time-consistency for both. We wish to highlight the role of time and risk preferences, materialized in *one-step aggregators*, in time-consistency. Depending on how you move from one-step time and risk preferences to intertemporal time and risk preferences, and depending on their compatibility (commutation), you will or will not observe time-consistency. We also shed light on the relevance of information structure by giving an explicit role to a dynamical system with state x.

The paper is organized as follows. In Section 2, we define dynamic uncertainty criteria ("cousins" of dynamic risk measures) and their time-consistency. Then, we introduce the notions of time and uncertainty-aggregators, define their composition, and show two ways to craft a dynamic uncertainty criterion from one-step aggregators: in the nested-aggregation case, we prove time-consistency; in the block-aggregation case, we have to add a commutation property for this. In Section 3, we introduce the basic material to formulate intertemporal optimization problems under uncertainty from dynamic uncertainty criteria, and define their time-consistency. In the nested-aggregation case, we prove time-consistency by displaying a DPE; in the block-aggregation case, we have to add a commutation property for this. We end with applications in Section 4, before concluding.

Notations

We fix notations used throughout the paper:

- $[\![a, b]\!]$ is the set of integers between a and b (included);
- $\mathcal{F}(E, F)$ is the set of functions mapping E into F;
- $\{u_t\}_0^T$ is the sequence $\{u_0, \ldots, u_T\}$;
- $\bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$;
- $\mathbb{W}_{[0:s]}$ is the Cartesian product $\mathbb{W}_0 \times \cdots \times \mathbb{W}_s$;
- $\mathbb{G}$ is used to refer to an aggregator with respect to uncertainty;
- ϕ is used to refer to an aggregator with respect to time.

Furthermore, the superscript notation indicates that the domain of the mapping $\mathbb{G}^{[t:s]}$ is $\mathcal{F}(\mathbb{W}_{[t:s]}; \bar{\mathbb{R}})$ (not to be confused with $\mathbb{G}_{[t:s]} = \{\mathbb{G}_r\}_{r=t}^{s}$).