European Journal of Operational Research 249 (2016) 155–163

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Time-inconsistent multistage stochastic programs: Martingale bounds[☆,☆☆]

Georg Ch. Pflug[a], Alois Pichler[b,*]

[a] University of Vienna, Department of Statistics and Operations Research and International Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria
[b] Norwegian University of Science and Technololgy, NTNU, Alfred Getz vei, 7034 Trondheim, Norway

ARTICLE INFO

Article history:
Received 2 September 2014
Accepted 12 February 2015
Available online 20 February 2015

Keywords:
Stochastic optimization
Risk measure
Average Value-at-Risk
Dynamic programming
Time consistency

ABSTRACT

Multistage stochastic programs show time-inconsistency in general, if the objective is neither the expectation nor the maximum functional.

This paper considers distortion risk measures (in particular the Average Value-at-Risk) at the final stage of a multistage stochastic program. Such problems are not time consistent. However, it is shown that by considering risk parameters at random level and by extending the state space appropriately, the value function corresponding to the optimal decisions evolves as a martingale and a dynamic programming principle is applicable. In this setup the risk profile has to be accepted to vary over time and to be adapted dynamically. Further, a verification theorem is provided, which characterizes optimal decisions by sub- and supermartingales. These enveloping martingales constitute a lower and an upper bound of the optimal value function.

The basis of the analysis is a new decomposition theorem for the Average Value-at-Risk, which is given in a time consistent formulation.

1. Introduction

Coherent risk measures have been introduced by Artzner et al. in the pioneering papers (Artzner, Delbaen, & Heath, 1997) and (Artzner, Delbaen, Eber, & Heath, 1999). The set of axioms, which are proposed there, is widely accepted nowadays. Approximately 10 years later the same group of authors considered risk measures again in a multistage framework in Artzner, Delbaen, Eber, Heath, and Ku (2007). These authors notice that the Average Value-at-Risk, the most important risk measure, is not time consistent in the sense specified in their paper. Importantly, they relate time-consistency to the fundamental dynamic programming principle, known as Bellman's principle (cf. Fleming & Soner, 2006).

For giving a precise definition of time-consistency (time inconsistency, respectively) one has to distinguish between time-inconsistent risk measures and time-inconsistent decision problems.

Artzner et al. consider the following notion of time-inconsistency in Artzner et al. (2007): a risk measure ρ applied to a random variable Y is said to be time-consistent, if knowing the value $\rho(Y|\mathcal{F})$ for all conditional distributions for any conditioning σ-algebra $\mathcal{F}$ is sufficient to calculate its unconditional value $\rho(Y)$. Their counterexample, showing that the Average Value-at-Risk is time-inconsistent in this sense, is reconsidered and resolved in this paper below (Fig. 2).

The notion of time-consistent decision problems is related, but slightly different: a multistage stochastic decision problem is time-consistent, if resolving the problem at later stages (i.e., after observing some random outcomes), the original solutions remain optimal for the later stages.

Shapiro (2009, p. 144), referring to a tree-structured problem, remarks that for time-consistency of a problem the solution at each stage is not allowed to depend on random parameters, which cannot follow this stage (i.e., in the language of trees, lie in other subtrees): "It is natural to consider the conceptual requirement that an optimal decision at state ξ_t should not depend on states which do not follow ξ_t, i.e., cannot happen in the future. That is, optimality of our decision at state ξ_t should only involve future children nodes of state ξ_t. We call this principle time consistency."

Carpentier, Chancelier, Cohen, de Lara, and Girardeau (2012, p. 249) formulate the property as follows: "The sequence of optimization problems is said to be dynamically consistent if the optimal strategies obtained when solving the original problem at time t_0 remain optimal for all subsequent problems. In other words, dynamic

☆ A. Pichler acknowledges support from the Research Council of Norway (RCN) through project 207690.
☆☆ This paper was processed by Guest-Editor prof. Ruediger Schultz for an uncompleted EJOR Feature Cluster on Stochastic Optimization.
* Corresponding author.
E-mail address: alois.pichler@univie.ac.at, aloisp@ntnu.no (A. Pichler).

http://dx.doi.org/10.1016/j.ejor.2015.02.033
0377-2217/

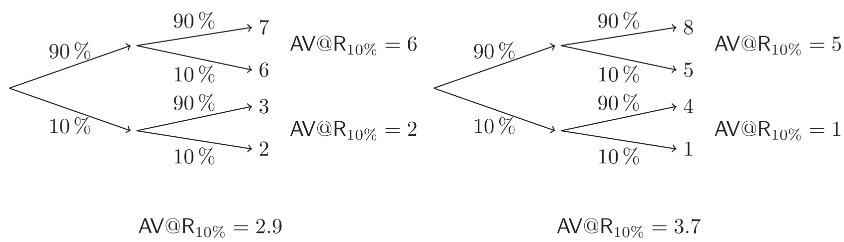

Fig. 1. Counterexample: the Average Value-at-Risk, conditioned on the subtrees, is higher on the left tree. But the Average Value-at-Risk at the initial stage is higher on the right tree.

consistency means that strategies obtained by solving the problem at the very first stage do not have to be questioned later on."

In order to enforce time-consistency for decision problems significant efforts and investigations have been initiated to identify classes of multiperiod risk measures which lead to time-consistency and allow stagewise decompositions. To this end, nested conditional risk measures have been proposed by Shapiro and Ruszczyński (Ruszczyński, 2010; Ruszczyński & Shapiro, 2006; Shapiro, 2009) (cf. as well Shapiro, 2012) and the recent paper by Carpentier et al. (2012). Here, we follow a different path.

First, we will maximize the acceptability rather than minimize risk. This is for ease of presentation of the results. Since acceptability measures can be seen as the negatives of risk measures this is no restriction. Secondly, we stick to the problem of maximizing the acceptability of the final outcome, where the acceptability is measured by a plain distortion measure, not by a nested one. We believe that nested risk or acceptability measures are difficult to interpret and are not what decision makers would understand under multistage risk or acceptability.

As a measure for acceptability we use initially the Average Value-at-Risk (AV@R) defined below in (1), and explain later how the results extend to general distortion measures. As illustration of time-inconsistency of decision problems involving AV@R consider the example displayed in Fig. 1. Suppose that the decision has to made, which out of the two tree processes is to be selected for the criterion to maximize the final $\mathsf{AV@R}_{10\,\text{percent}}$. It can be easily seen that looking at the conditional AV@R's at intermediate time 1, the left tree has to be preferred to the right one. Looking at the same criterion, but at time 0, the preference is opposite. Therefore, this introductory example confirms the fact that in general the maximization of the AV@R of the final outcome leads to time-inconsistency of decisions.

This paper is based on a new decomposition of the AV@R and related measures. The decomposition measures risk on conditional level only, and it recovers the initial risk measure by collecting the conditional risk measures via an expectation. In this setup the risk profile has to be adapted conditionally, such that the conditional risk profile is not static any longer.

Additional information changes the perception of risk. The adaptive choice of appropriate measures of risk complies with the course of action of a risk manager who adjusts the preferences whenever additional information is available. The decision maker is less reluctant, if an observation reveals that the future will be bright, but conversely she or he will be more strict if losses at the end become more likely. This gives rise to defining an extended notion of a conditional risk measure, which is not just the same risk measure applied to conditional distributions, but which may be a different functional for different conditional distributions, depending on its respective history.

By involving adapted conditional risk measures it is possible to recover dynamic programming principles for multistage stochastic programs. Moreover verification theorems, which are central in dynamic control, are established here for multistage stochastic programs. The dynamic programming equations presented are based on the dual representation of the risk measure, and different to those provided by Shapiro (2009). The presented approach allows a characterization of optimal solutions of a multistage stochastic program in terms of enveloping sub- and supermartingales. It is shown that a solution of a multistage problem evolves as a martingale over time, where different risk measures are encountered at each stage.

Dynamic programming equations notably cannot remove the time inconsistency, which is inherent to these problems. But these equations come along with verification theorems, and it is their purpose to enable checking, if a given policy is optimal. By assessing the enveloping sub- and supermartingales it is moreover possible to provide upper and lower bounds, such that the quality of a given multistage policy can be assessed with these sub- and supermartingales as well.

1.1. Outline of the paper

Section 2 provides the setting for the Average Value-at-Risk, as this risk measure is basic for the presentation. Next, the conditional version is considered. The decomposition theorem, the central statement of this paper, is contained in Section 4 and Section 5 characterizes its properties. Section 6 introduces the multistage optimization problem. Section 7 exposes the dynamic programming formulation, while the subsequent Section 8 introduces the martingale representations, which are in line with dynamic programming.

2. Representations of the genuine risk measure

We reduce the conceptual description of the problem to the Average Value-at-Risk, AV@R. This reduction is justified, as more general coherent risk measures—distortion risk measures—are composed in a linear way of Average Value-at-Risks at different levels. Further, Kusuoka's theorem provides all version independent (also known as law invariant) risk measures via distortion risk measures (cf., for example, Pflug & Römisch, 2007), such that this reduction is without loss of generality.

The Average Value-at-Risk is considered in its *concave* variant involving the lower quantiles of the distribution function F_Y of the random variable Y,

$$\mathsf{AV@R}_\alpha(Y) := \frac{1}{\alpha}\int_0^\alpha F_Y^{-1}(u)\,\mathrm{d}u \quad (0 < \alpha \le 1), \tag{1}$$

where α is called *level*. In this setting AV@R accounts for profits, which are subject to maximization. Throughout this paper we shall assume that the profit variable Y is a $\mathbb{R}$-valued random variable defined on a general, filtered probability space $(\Omega,\ (\mathcal{F}_t)_{t\in\{0,1,\dots T\}},\ P)$. For convenience of presentation we assume that $Y \in L^\infty(\mathcal{F}_T, P)$ ($L^1(\mathcal{F}_T, P)$ could be chosen in many, but not in all situations).

The *dual representation* of the Average Value-at-Risk at level α is

$$\mathsf{AV@R}_\alpha\,(Y) = \inf\left\{\mathbb{E}\,(YZ)\ :\ 0 \le Z,\ \alpha Z \le 1 \text{ and } \mathbb{E}\,(Z) = 1\right\}, \tag{2}$$

where the expectation is with respect to the measure P, the infimum in (2) is among all positive random variables $Z \ge 0$ with expectation $\mathbb{E}(Z) = 1$ (i.e., Z are densities), satisfying the additional truncation constraint $\alpha Z \le 1$, as indicated. The infimum is attained if $\alpha > 0$, and in this case the optimal random variable Z in (2) is coupled in an anti-monotone way with Y (cf. Nelsen, 1998).