European Journal of Operational Research 234 (2014) 743–750

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Time consistency and risk averse dynamic decision models: Definition, interpretation and practical consequences

Birgit Rudloff [a], Alexandre Street [b], Davi M. Valladão [c,*]

[a] Department of Operations Research and Financial Engineering, Princeton University, Sherrerd Hall, Charlton Street, Princeton, NJ 08544, USA
[b] Electrical Engineering Department, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Rua Marquês de São Vicente, 225, Gávea – Rio de Janeiro, RJ 22451-900, Brazil
[c] Industrial Engineering Department, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Rua Marquês de São Vicente, 225, Gávea – Rio de Janeiro, RJ 22451-900, Brazil

ARTICLE INFO

Article history:
Received 28 May 2012
Accepted 28 November 2013
Available online 4 December 2013

Keywords:
Time consistency
Dynamic stochastic programming
Risk aversion
Conditional Value-at-Risk (CVaR)
Portfolio selection

ABSTRACT

This paper aims at resolving a major obstacle to practical usage of time-consistent risk-averse decision models. The recursive objective function, generally used to ensure time consistency, is complex and has no clear/direct interpretation. Practitioners rather choose a simpler and more intuitive formulation, even though it may lead to a time inconsistent policy. Based on rigorous mathematical foundations, we impel practical usage of time consistent models as we provide practitioners with an intuitive economic interpretation for the referred recursive objective function. We also discourage time-inconsistent models by arguing that the associated policies are sub-optimal. We developed a new methodology to compute the sub-optimality gap associated with a time-inconsistent policy, providing practitioners with an objective method to quantify practical consequences of time inconsistency. Our results hold for a quite general class of problems and we choose, without loss of generality, a CVaR-based portfolio selection application to illustrate the developed concepts.

1. Introduction

Dynamic decisions under uncertainty are very common in financial planning and financial engineering problems. Based on Bellman's equations and on the well behaved properties of the risk neutral formulation, several models have been developed for different applications such as portfolio selection, asset and liability management, scheduling and energy planning. Indeed, some important works, for instance Pereira and Pinto (1991), Rockafellar and Wets (1991), Guigues and Römisch (2012a, 2012b), Philpott and de Matos (2012), Philpott, de Matos, and Finardi (2013), developed efficient algorithms to solve these problems. However, in order to have a risk aversion extension, such as in Ruszczyński (2010), Shapiro (2011), Shapiro, Tekaya, da Costa, and Soares (2013), one should choose carefully how to introduce the well known risk measures into these problems.

In this context, the Conditional Value at Risk (CVaR) became one of the most widely used risk measures for three reasons: first, it is a coherent risk measure (see Artzner, Delbaen, Eber, & Heath (1999)); second, it has a clear and suitable economic interpretation (see Rockafellar & Uryasev (2000) and Street (2009)); and last, but not least, it can be written as a linear stochastic programming model as shown in Rockafellar and Uryasev (2000). For these three reasons, the CVaR has been applied to static (see e.g. Krokhmal, Palmquist, & Uryasev (2002), Mansini, Ogryczak, & Speranza (2007), Sawik (2012) among many others in the context of portfolio optimization) and even to dynamic models. However, to choose a coherent risk measure as the objective function of a dynamic model is not a sufficient condition to obtain suitable optimal policies. In the recent literature, time consistency is shown to be one basic requirement to get suitable optimal decisions, in particular for multistage stochastic programming models. Papers on time consistency are actually divided in two different approaches: the first one focuses on risk measures and the second one on optimal policies.

The first approach states that, in a dynamic setting, if some random payoff A is always riskier than a payoff B conditioned to a given time $t+1$, than A should be riskier than B conditioned to t. It is well known that this property is achieved using a recursive setting leading to so called time consistent dynamic risk measures proposed by various authors, e.g., Bion-Nadal (2008), Detlefsen and Scandolo (2005), Riedel (2004), Cheridito, Delbaen, and Kupper (2006), Roorda and Schumacher (2007), Kovacevic and Pflug (2009). Other weaker definitions, like acceptance and rejection consistency, are also developed in these works (see Cheridito et al. (2006), Kovacevic & Pflug (2009) for details).

The second approach, formally defined by Shapiro (2009), is on time consistency of optimal policies in multistage stochastic programming models. The interpretation of this property given by the author is the following: "at every state of the system, our

* Corresponding author.
E-mail addresses: brudloff@princeton.edu (B. Rudloff), street@ele.puc-rio.br (A. Street), davimichel@gmail.com, davimv@puc-rio.br (D.M. Valladão).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.11.037

optimal decisions should not depend on scenarios which we already know cannot happen in the future''. This interpretation is an indirect consequence of solving a sequence of problems whose objective functions can be written recursively as the former cited time consistent dynamic risk measures. It is shown in Shapiro (2009) for instance that if, for every state of the system, we want to minimize the CVaR of a given quantity at the end of the planning horizon, we would obtain a time *inconsistent* optimal policy. Indeed, this sequence of problems does not have recursive objective functions and the optimal decisions at particular future states might depend on scenarios that ''we already know cannot happen in the future''. However, if for $t = 0$ we want to minimize the CVaR of a given quantity at the end of the planning horizon and for $t > 0$ we actually follow the dynamic equations of the first stage problem, then we obtain a time *consistent* optimal policy even though it depends on those scenarios we already know cannot happen. On the other hand, one can argue that this policy is not reasonable because for $t > 0$ the objective function does not make any sense economically speaking.

In this paper, we use a conceptual definition for time consistency of optimal policies: *a policy is time consistent if and only if the future planned decisions are actually going to be implemented*. In the literature, time inconsistent optimal policies have been commonly proposed, in particular Bäuerle and Mundt (2009) in Sections 3 and 4.1 and Fábián and Veszprémi (2008) have developed portfolio selection models using CVaR in a time inconsistent way. In our work, we show with a numerical example that a time inconsistent CVaR based portfolio selection model can lead to a suboptimal sequence of implemented decisions and may not take risk aversion into account at some intermediate states of the system. Then, we propose a methodology to compute the sub-optimality gap as the difference of the objective function evaluated with two different policies: the one planned at our current stage and the one actually implemented in the future. We use the time consistent risk-averse dynamic stochastic programming model with a recursive objective function and compare its optimal policy to the time inconsistent one. Other alternatives have been proposed by Boda and Filar (2006) and Cuoco, He, and Issaenko (2008), however none of them used the recursive set up of time consistent dynamic risk measures.

Since the lack of a suitable economic interpretation for this recursive set up is one of the main reasons why it is not commonly used, we prove for a more general set of problems that this objective function is the certainty equivalent w.r.t. the time consistent dynamic utility defined as the composed form of one period preference functionals. Then, with a clear economic interpretation for the objective function of a general set of problems and consequently of the portfolio selection application, we discuss the developed results with a numerical example.

This paper aims at closing a conceptual gap between theory and practice regarding time-consistent risk-averse policies in a stochastic programming framework. The first contribution of this work is an intuitive economic interpretation for the complex time-consistent recursive objective function. Based on rigorous proofs provided in the paper, the proposed economic interpretation is entirely new in the dynamic risk-averse stochastic programming context. Additionally, we discourage time-inconsistent models by arguing that the associated policies are sub-optimal. Thus, the second contribution is the development of a new methodology to compute the sub-optimality gap associated with a time-inconsistent policy, providing practitioners with an objective method to quantify practical consequences of time inconsistency. Our results hold for a quite general class of problems and we choose, without loss of generality, a CVaR-based portfolio selection application to illustrate the developed concepts. Numerical results presented in the paper aim at reinforcing intuition and interpretation of the reader regarding theoretical results.

1.1. Assumptions and notation

In this paper, we assume a multistage setting with a finite planning horizon T. We consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with a related filtration $\mathcal{F}_0 \subseteq \ldots \subseteq \mathcal{F}_T$, where $\mathcal{F}_0 = \{\emptyset, \Omega\}$ and $\mathcal{F} = \mathcal{F}_T$.

Since our application is on portfolio selection, we use a unique notation for all models developed here. The sets, stochastic processes, decision and state variables involved are defined in Tables 1 and 2.

Without loss of generality, we assume that there is a risk free asset, indexed by $i = 1$, with null excess return for each state of the system, i.e., $r_{1,t}(\omega) = 0$, for all $t \in \mathcal{H} \cup \{T\}$ and all $\omega \in \Omega$. Moreover, we assume that $W_t, r_{i,t}, x_{i,t} \in L^\infty(\mathcal{F}_t)$, for all $t \in \mathcal{H} \cup \{T\}$, where $L^\infty(\mathcal{F}_t)$ denotes the linear space of the equivalence classes of almost surely bounded $\mathcal{F}_t$-measurable random variables on $(\Omega, \mathcal{F}, \mathbb{P})$.

Let W be an $\mathcal{F}$ measurable function and consider a realization sequence $\bar{\mathbf{r}}_{[1,t]} = (\bar{\mathbf{r}}_1, \ldots, \bar{\mathbf{r}}_t)'$ of the asset returns. Then, we denote the conditional and unconditional expectations by $\mathbb{E}\left[W|\bar{\mathbf{r}}_{[1,t]}\right] = \mathbb{E}\left[W|\mathbf{r}_{[1,t]} = \bar{\mathbf{r}}_{[1,t]}\right]$ and $\mathbb{E}[W]$, respectively.

We also use the negative of the CVaR developed by Rockafellar and Uryasev (2000) as an ''acceptability'' measure (see Kovacevic & Pflug (2009) for details) whose conditional and unconditional formulations are defined respectively as

Table 1
Set and stochastic process notation.

Sets	
$\mathcal{A} = \{1, \ldots, A\}$:	Index set of the $A \geqslant 1$ assets
$\mathcal{H} = \{0, \ldots, T-1\}$:	Set of stages
$\mathcal{H}(\tau) = \{\tau, \ldots, T-1\}$:	Set of stages starting from τ
Stochastic Process	
$r_{i,t}(\omega)$:	Excess return of asset $i \in \mathcal{A}$, between stages $t \in \{1, \ldots, T\}$ and $t-1$, under scenario $\omega \in \Omega$, where $\mathbf{r_t}(\omega) = (r_{1,t}(\omega), \ldots, r_{A,t}(\omega))'$ and $\mathbf{r}_{[\mathbf{s},\mathbf{t}]}(\omega) = (\mathbf{r_s}(\omega), \ldots, \mathbf{r_t}(\omega))'$ for $s \leqslant t$
$\bar{\mathbf{r}}_{[s,t]} = (\bar{\mathbf{r}}_\mathbf{s}, \ldots, \bar{\mathbf{r}}_\mathbf{t})'$:	Realization sequence of the asset returns for $s \leqslant t$

Table 2
State and decision variable notation.

State variables	
$W_t(\omega)$:	Wealth at stage $t \in \mathcal{H} \cup \{T\}$ under scenario $\omega \in \Omega$
Decision variables	
$x_{i,t}(\omega)$:	Amount invested in asset $i \in \mathcal{A}$, at stage $t \in \mathcal{H}$ under scenario $\omega \in \Omega$, where $\mathbf{x_t}(\omega) = (x_{1,t}(\omega), \ldots, x_{A,t}(\omega))'$ and $\mathbf{x}_{[\mathbf{s},\mathbf{t}](\omega)} = (\mathbf{x_s}(\omega), \ldots, \mathbf{x_t}(\omega))'$ for $s \leqslant t$