

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

Birgit Rudloff

Princeton University

Alexandre Street

Pontifical Catholic University - Rio de Janeiro

Davi M. Valladão

IBM Research - Brazil

Pontifical Catholic University - Rio de Janeiro

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.

Email addresses: brudloff@princeton.edu (Birgit Rudloff),
street@ele.puc-rio.br (Alexandre Street), davimv@br.ibm.com (Davi M. Valladão)

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.

Keywords: Time Consistency, Dynamic Stochastic Programming, Risk Aversion, Conditional Value-at-Risk (CVaR), Portfolio Selection

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 Römisich (2012a,b); Philpott and de Matos (2012); Philpott et al. (2012), developed efficient algorithms to solve these problems. However, in order to have a risk aversion extension, such as in Ruszczyński (2010); Shapiro (2011); Shapiro et al. (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 et al. (1999)); second, it has a clear and suitable economic interpretation (see Rockafellar and 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. Krokmal et al. (2002); Mansini et al. (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 et al. (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 and 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 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 Fábíán 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 et al. (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 $\mathcal{F}_0 \subseteq \dots \subseteq \mathcal{F}_T$, 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.

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

Table 1: Set and stochastic process notation

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 $\bar{\mathbf{r}}_{[1,t]} = (\bar{\mathbf{r}}_1, \dots, \bar{\mathbf{r}}_t)'$ of the asset returns. Then, we denote the conditional and unconditional expectations by $\mathbb{E}[W \mid \bar{\mathbf{r}}_{[1,t]}] = \mathbb{E}[W \mid \mathbf{r}_{[1,t]} = \bar{\mathbf{r}}_{[1,t]}]$ 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 and Pflug (2009) for

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 $\mathbf{x}_t(\omega) = (x_{1,t}(\omega), \dots, x_{A,t}(\omega))'$ and $\mathbf{x}_{[s,t]}(\omega) = (\mathbf{x}_s(\omega), \dots, \mathbf{x}_t(\omega))'$ for $s \leq t$

Table 2: State and decision variable notation

details) whose conditional and unconditional formulations are defined respectively as

$$\phi_t^\alpha(W, \bar{\mathbf{r}}_{[1,t]}) = -CVaR_\alpha(W \mid \bar{\mathbf{r}}_{[1,t]}) = \sup_{z \in \mathbb{R}} \left\{ z - \frac{\mathbb{E}[(W - z)^- \mid \bar{\mathbf{r}}_{[1,t]}]}{1 - \alpha} \right\} \quad (1)$$

and

$$\phi_0^\alpha(W) = -CVaR_\alpha(W) = \sup_{z \in \mathbb{R}} \left\{ z - \frac{\mathbb{E}[(W - z)^-]}{1 - \alpha} \right\}$$

at probability level $\alpha \in (0, 1)$ (respectively significance level $1 - \alpha$), where $x^- = -\min(x, 0)$. The supremum in (1) is attained by z^* being the conditional $(1 - \alpha)$ -quantile of W , given $\bar{\mathbf{r}}_{[1,t]}$.

Note that, $\mathbb{E}[\cdot \mid \bar{\mathbf{r}}_{[1,t]}]$, $\mathbb{E}[\cdot]$, $\phi_t^\alpha(\cdot, \bar{\mathbf{r}}_{[1,t]})$ and $\phi_t^\alpha(\cdot)$ are real valued functions, mapping from $L^\infty(\mathcal{F})$ to \mathbb{R} . It is also important to note that all constraints represented in this paper are defined for almost every $\omega \in \Omega$ (in the \mathbb{P} a.s. sense) that affects the objective function. For instance, if the objective function of a particular optimization problem is a conditional expectation $\mathbb{E}[\cdot \mid \bar{\mathbf{r}}_{[1,t]}]$, then the constraints of this problem are defined for almost every $\omega \in \{\bar{\omega} \in \Omega \mid \mathbf{r}_{[1,t]}(\bar{\omega}) = \bar{\mathbf{r}}_{[1,t]}\}$.

2. Time Consistency: Definition and Practical Consequences

The major reason for developing dynamic (multistage) models instead of static (two-stage) ones is the fact that we can incorporate the flexibility of dynamic decisions to improve our objective function. In other words, the possibility of changing a policy after the realization of some random variables increases the objective function (for a maximization problem) and allows the first stage decisions to be less conservative than their counterpart in the static case. However, it doesn't make any sense to incorporate this flexibility if the intermediate decisions are not actually going to be implemented.

In this context, we based our findings on the following definition of time consistency: *an optimal policy is time consistent if and only if the future planned decisions are actually going to be implemented.* Only under this property we guarantee that the flexibility and optimality of a dynamic policy will not be polluted by any spurious future planned decisions. Thus, the first stage decisions of a time inconsistent policy are in general suboptimal considering that this policy would not be followed in the future.

In a multistage stochastic programming context, a policy is a sequence of decisions for each stage and for each scenario (a realization of the uncertainty). As in Shapiro (2009), one has to define which (multistage) optimization problem should be solved when the current time is a particular stage $t \in \mathcal{H}$ of the planning horizon. When the current time is $t = 0$, we solve the corresponding optimization problem and obtain the first stage optimal decision and the future planned optimal policy. This policy is time consistent if and only if these future planned decisions for each scenario are also optimal for each problem when the current time is $t > 0$.

In order to motivate this discussion, we develop a CVaR-based portfolio selection model which incorporates the well known mean-risk trade-off pre-

sented by Markowitz (1952). As a coherent risk measure, the CVaR should be a suitable way to assess risk, however we want to point out the fact that if one chooses a dynamic model, time consistency should also be taken into account. Assessing risk in a time inconsistent way may lead to a time inconsistent policy and therefore to a suboptimal sequence of implemented decisions.

In subsections 2.1 and 2.2, we compare a time inconsistent formulation to the recursive time consistent alternative developed for a CVaR-based portfolio selection problem. Numerical results are presented to reinforce intuition and interpretation of the reader. For sake of simplicity, we consider a small instance with three stages and two assets where the solution is obtained with standard linear programming techniques. In subsection 2.3, we propose a new way of computing the sub-optimality gap induced by an inconsistent policy. We then perform in subsection 2.3 a sensitivity analysis of the sub-optimality gap by varying planning horizon T and risk aversion level λ for the same portfolio selection problem proposed in 2.1.

2.1. Example of a time inconsistent policy

The portfolio selection problem is normally formulated to consider the mean-risk trade-off. Some models use the expected value as the objective function with a risk constraint while others minimize risk with a constraint on the expected value. In this paper, we follow the well known approach to combine these two methods by defining our objective function as a convex combination of the expected value and the acceptability measure previously stated. In other words, the investor wants to maximize its expected return and also minimize risk, given his current state. It is very important to note that the planning horizon is a fixed date in the future and, depending on the investor's current state, he / she solves a different optimization problem.

Then, we define the problem $Q_\tau (W_\tau, \bar{\mathbf{r}}_{[1,\tau]})$ solved by the investor, given

his / her current stage τ and the current realization $\bar{\mathbf{r}}_{[1,\tau]}$ of the random process, as

$$\underset{W_{[\tau+1,T]}, \mathbf{x}_{[\tau,T-1]}}{\text{maximize}} \quad (1 - \lambda) \mathbb{E} [W_T \mid \bar{\mathbf{r}}_{[1,\tau]}] + \lambda \phi_\tau^\alpha (W_T, \bar{\mathbf{r}}_{[1,\tau]}) \quad (2)$$

$$\text{subject to} \quad W_{t+1} = \sum_{i \in \mathcal{A}} (1 + r_{i,t+1}) x_{i,t}, \quad \forall t \in \mathcal{H}(\tau) \quad (3)$$

$$\sum_{i \in \mathcal{A}} x_{i,t} = W_t, \quad \forall t \in \mathcal{H}(\tau) \quad (4)$$

$$\mathbf{x}_t \geq 0, \quad \forall t \in \mathcal{H}(\tau), \quad (5)$$

where $\lambda \in [0, 1]$. The objective function in (2) is the convex combination of the expected value and the acceptability measure defined in (1). The parameter λ can be interpreted as a risk aversion level as it determines how much weight in the objective is given to minimizing risk as opposed to maximizing expected return. The extreme case $\lambda = 0$ corresponds to the risk neutral case. The higher λ , the higher the risk aversion of the investor as the impact of the risk measure in the objective function gets higher. The constraints consist of the budget constraint (4) stating that all the wealth W_t is invested into the A assets at time $t \in \mathcal{H}(\tau)$, where constraint (5) excludes short selling. The amount $x_{i,t}$ invested in asset i at time t generates the (random) wealth $(1 + r_{i,t+1})x_{i,t}$ in asset i at time $t + 1$, which leads to the (random) total wealth W_{t+1} by summing up over all assets in (3).

Using (1), the problem can be equivalently formulated as

$$\underset{W_{[\tau+1,T]}, \mathbf{x}_{[\tau,T-1]}, z}{\text{maximize}} \quad \mathbb{E} \left[(1 - \lambda) W_T + \lambda \left(z - \frac{(W_T - z)^-}{1 - \alpha} \right) \mid \bar{\mathbf{r}}_{[1,\tau]} \right]$$

$$\text{subject to} \quad W_{t+1} = \sum_{i \in \mathcal{A}} (1 + r_{i,t+1}) x_{i,t}, \quad \forall t \in \mathcal{H}(\tau)$$

$$\sum_{i \in \mathcal{A}} x_{i,t} = W_t, \quad \forall t \in \mathcal{H}(\tau)$$

$$\mathbf{x}_t \geq 0, \quad \forall t \in \mathcal{H}(\tau).$$

Note that the first stage problem $Q_0(W_0)$ is defined equivalently as:

$$\begin{aligned}
& \underset{W_{[1,T]}, \mathbf{x}_{[0,T-1]}, z}{\text{maximize}} && \mathbb{E} \left[(1 - \lambda) W_T + \lambda \left(z - \frac{(W_T - z)^-}{1 - \alpha} \right) \right] \\
& \text{subject to} && W_{t+1} = \sum_{i \in \mathcal{A}} (1 + r_{i,t+1}) x_{i,t}, \quad \forall t \in \mathcal{H} \\
& && \sum_{i \in \mathcal{A}} x_{i,t} = W_t, \quad \forall t \in \mathcal{H} \\
& && \mathbf{x}_t \geq 0, \quad \forall t \in \mathcal{H}(\tau).
\end{aligned} \tag{6}$$

In order to have a numerical example, let us assume our probability space to be represented by a discrete event tree. For instance, consider $T = 2$ and the tree represented in Figure 1, where the scenarios $\omega \in \Omega = \{1, 2, 3, 4\}$ are numbered by the terminal nodes. In our notation, a node is a subset of Ω , e.g., the root node is defined as $\Omega = \{1, 2, 3, 4\}$, the intermediate nodes as $\{1, 2\}$ and $\{3, 4\}$ and the terminal nodes as $\{1\}, \{2\}, \{3\}, \{4\}$. Now, let us denote \mathcal{N}_t the set of nodes at stage t and \mathcal{F}_t the σ -algebra generated by it. In our example, $\mathcal{N}_1 = \{\Omega\}$, $\mathcal{N}_2 = \{\{1, 2\}, \{3, 4\}\}$ and $\mathcal{N}_3 = \{\{1\}, \{2\}, \{3\}, \{4\}\}$.

For sake of simplicity, we consider a two-asset model, i.e., $\mathcal{A} = \{1, 2\}$, and a probability measure defined as $\mathbb{P}(\omega) = 0.25$ for all $\omega \in \Omega = \{1, 2, 3, 4\}$. The first asset indexed by $i = 1$ is risk free and it has null excess return for every state of the system, i.e, $r_{1,t}(\omega) = 0$ for all $t \in \{1, 2\}$, $\omega \in \Omega$. The second one is assumed to have iid returns given by

$$r_{2,t}(\omega) = \begin{cases} 1, & \forall t = 1, \omega \in \{1, 2\} \\ -0.5, & \forall t = 1, \omega \in \{3, 4\} \\ 1, & \forall t = 2, \omega \in \{1\} \\ -0.5, & \forall t = 2, \omega \in \{2\} \\ 1, & \forall t = 2, \omega \in \{3\} \\ -0.5, & \forall t = 2, \omega \in \{4\} \end{cases}$$

and is graphically represented in Figure 1. It is straightforward to see that the risky asset has greater expected return and higher risk than the




Figure 1: Return tree for $i = 2$

risk free one. This represents the mean-risk trade-off of a typical portfolio selection problem.

Now, we write an equivalent deterministic linear programming model for problem $Q_0(W_0)$ defined in (6) assuming, without loss of generality, that $W_0 = 1$. Then we have the following:

$$\begin{aligned}
& \underset{q, W_{[1,2]}, \mathbf{x}_{[0,1]}, z}{\text{maximize}} && \frac{1}{4} \sum_{\omega=1}^4 \left[(1 - \lambda) W_2(\omega) + \lambda \left(z - \frac{q(\omega)}{1 - \alpha} \right) \right] \\
& \text{subject to} && W_{t+1}(\omega) = \sum_{i \in \mathcal{A}} (1 + r_{i,t+1}(\omega)) x_{i,t}(\omega), \quad \forall t \in \mathcal{H}, \omega \in \Omega \\
& && \sum_{i \in \mathcal{A}} x_{i,t}(\omega) = W_t(\omega), \quad \forall t \in \mathcal{H}, \omega \in \Omega \\
& && \mathbf{x}_t(\omega) \geq 0, \quad \forall t \in \mathcal{H}, \omega \in \Omega \\
& && q(\omega) \geq z - W_2(\omega), \quad \forall \omega \in \Omega \\
& && q(\omega) \geq 0, \quad \forall \omega \in \Omega,
\end{aligned} \tag{7}$$

where \mathbf{x}_t is \mathcal{F}_t -adapted, i.e., $\mathbf{x}_0(1) = \mathbf{x}_0(2) = \mathbf{x}_0(3) = \mathbf{x}_0(4)$, $\mathbf{x}_1(1) = \mathbf{x}_1(2)$ and $\mathbf{x}_1(3) = \mathbf{x}_1(4)$, which are the well known non-antecipativity constraints. Note that q is a \mathcal{F}_T -adapted auxiliary variable to represent the CVaR as developed in Rockafellar and Uryasev (2000).

Solving this problem for $\alpha = 95\%$ and $\lambda = 0.5$, we have the following optimal solution:

$$x_{1,t}^*(\omega) = \begin{cases} 0.5, & \forall t = 0, \omega \in \Omega \\ 0, & \forall t = 1, \omega \in \{1, 2\} \\ 0.75, & \forall t = 1, \omega \in \{3, 4\} \end{cases} \quad x_{2,t}^*(\omega) = \begin{cases} 0.5, & \forall t = 0, \omega \in \Omega \\ 1.5, & \forall t = 1, \omega \in \{1, 2\} \\ 0, & \forall t = 1, \omega \in \{3, 4\}. \end{cases} \quad (8)$$

At the root node, it is optimal to split evenly the investment, while at node $\{1, 2\}$ everything is invested in the risky asset and at node $\{3, 4\}$ everything is invested in the risk free one.

Now, let us suppose one period has passed and the current state is at time $\tau = 1$ and at node $\{1, 2\}$. Let us write an equivalent deterministic problem for $Q_1(W_1, \bar{\mathbf{r}}_1)$, for $W_1 = 1.5$ and $\bar{\mathbf{r}}_1 = (0, 1)'$ as

$$\begin{aligned} & \underset{q, W_2, \mathbf{x}_1, z}{\text{maximize}} && \frac{1}{2} \sum_{\omega=1}^2 \left[(1 - \lambda) W_2(\omega) + \lambda \left(z - \frac{q(\omega)}{1 - \alpha} \right) \right] \\ & \text{subject to} && W_2(\omega) = \sum_{i \in \mathcal{A}} (1 + r_{i,2}(\omega)) x_{i,1}, \quad \forall \omega \in \{1, 2\} \\ & && \sum_{i \in \mathcal{A}} x_{i,1} = W_1 \\ & && \mathbf{x}_1 \geq 0 \\ & && q(\omega) \geq z - W_T(\omega), \quad \forall \omega \in \{1, 2\} \\ & && q(\omega) \geq 0, \quad \forall \omega \in \{1, 2\}. \end{aligned} \quad (9)$$

This problem reflects what the investor would do at $\tau = 1$ and at node $\{1, 2\}$ if the optimal decision \mathbf{x}_0^* in (8) had been implemented. In other words, given $x_{1,t}^*$ and $x_{2,t}^*$ for $t = 0$, the optimal solution of (9) is the decision implemented at $\tau = 1$ and at node $\{1, 2\}$ of an agent that maximizes the chosen acceptability measure of terminal wealth.

We want to show that the optimal solutions for this problem at node $\{1, 2\}$ are different from the ones in (8), meaning that at $t = 0$ the future planned decisions for $t = 1$ are different from the ones that are actually going to be implemented. It is also important to understand why it happens and

what kind of error a investor would do with this time inconsistent policy. The optimal solution of (9) is given by the following:

$$\tilde{x}_{1,t}^*(\omega) = 1.5 \quad \text{and} \quad \tilde{x}_{2,t}^*(\omega) = 0, \quad \forall t = 1, \omega \in \{1, 2\}. \quad (10)$$

The optimal planned strategy at node $\{1, 2\}$ obtained by solving (7) is to invest everything in the risky asset, while the solution of problem (9) (the one that is actually going to be implemented) is to invest everything in the risk free asset (see equation (10)). This happens because, in problem (7), $CVaR_{95\%}(W_2) = -W_2(4)$, which is the worst case loss (scenario $\omega = 4$). Then, at node $\{1, 2\}$, it is optimal in the first stage problem to choose the investment strategy with the highest expected return since this decision will not affect the terminal wealth at scenario $\omega = 4$.

This example points out that a time inconsistent policy may lead to a sequence of optimal decisions where a risk-averse decision maker shows a risk neutral preference at some intermediate state. In other words, risk aversion may not be taken into account at some intermediate states of the system. Therefore, we propose a time consistent alternative that has significant advantages over the time inconsistent one since it incorporates the flexibility of a dynamic decision model ensuring that the future planned decisions are actually going to be implemented.

2.2. Time consistent alternative

In this section, we propose an alternative to the previous time inconsistent policy. We base our formulation on Ruszczyński and Shapiro (2006) and develop dynamic equations. For $t = T - 1$, we define the problem

$V_{T-1}(W_{T-1}, \bar{\mathbf{r}}_{T-1})$ as follows:

$$\begin{aligned} & \underset{W_T, \mathbf{x}_{T-1}}{\text{maximize}} && (1 - \lambda) \mathbb{E} [W_T \mid \bar{\mathbf{r}}_{[1, T-1]}] + \lambda \phi_{T-1}^\alpha (W_T, \bar{\mathbf{r}}_{[1, T-1]}) \\ & \text{subject to} && W_T = \sum_{i \in \mathcal{A}} (1 + r_{i, T}) x_{i, T-1} \\ & && \sum_{i \in \mathcal{A}} x_{i, T-1} = W_{T-1} \\ & && \mathbf{x}_{T-1} \geq 0. \end{aligned}$$

Using the definition of $\phi_{T-1}^\alpha (W, \bar{\mathbf{r}}_{[1, T-1]})$ given in (1), we rewrite the problem as follows:

$$\begin{aligned} & \underset{W_T, \mathbf{x}_{T-1}, z}{\text{maximize}} && \mathbb{E} \left[(1 - \lambda) W_T + \lambda \left(z - \frac{(W_T - z)^-}{1 - \alpha} \right) \mid \bar{\mathbf{r}}_{[1, T-1]} \right] \\ & \text{subject to} && W_T = \sum_{i \in \mathcal{A}} (1 + r_{i, T}) x_{i, T-1} \\ & && \sum_{i \in \mathcal{A}} x_{i, T-1} = W_{T-1} \\ & && \mathbf{x}_{T-1} \geq 0. \end{aligned}$$

For the last period, our proposed model is to maximize the convex combination of the expected terminal wealth and the acceptability measure $\phi_{T-1}^\alpha (W, \bar{\mathbf{r}}_{[1, T-1]})$. Now, for $t < T - 1$, we propose a nested value function, based on the conditional version of the same convex combination. Then, $V_t (W_t, \bar{\mathbf{r}}_{[1, t]})$, $\forall t = 0, \dots, T - 2$, is defined as follows:

$$\begin{aligned} & \underset{W_{t+1}, \mathbf{x}_t}{\text{maximize}} && (1 - \lambda) \mathbb{E} [V_{t+1} \mid \bar{\mathbf{r}}_{[1, t]}] + \lambda \phi_t^\alpha (V_{t+1}, \bar{\mathbf{r}}_{[1, t]}) \\ & \text{subject to} && W_{t+1} = \sum_{i \in \mathcal{A}} (1 + r_{i, t+1}) x_{i, t} \\ & && \sum_{i \in \mathcal{A}} x_{i, t} = W_t \\ & && \mathbf{x}_t \geq 0, \end{aligned} \tag{11}$$

where V_{t+1} stands for $V_{t+1} (W_{t+1}, \mathbf{r}_{[1, t+1]})$.

Equivalently to $t = T - 1$, we rewrite problem (11) as follows:

$$\begin{aligned}
& \underset{W_{t+1}, \mathbf{x}_t, z}{\text{maximize}} && \mathbb{E} \left[(1 - \lambda) V_{t+1} + \lambda \left(z - \frac{(V_{t+1} - z)^-}{1 - \alpha} \right) \middle| \bar{\mathbf{r}}_{[1,t]} \right] \\
& \text{subject to} && W_{t+1} = \sum_{i \in \mathcal{A}} (1 + r_{i,t+1}) x_{i,t} \\
& && \sum_{i \in \mathcal{A}} x_{i,t} = W_t \\
& && \mathbf{x}_t \geq 0.
\end{aligned} \tag{12}$$

For comparison purposes, we solve this model for the numerical example proposed in Section 2.1. To do so, we use the result shown in Blomvall and Shapiro (2006) that, for stage-wise independent returns such problem has a myopic optimal policy which is obtained as the solution of the following two-stage problem for $t \in \mathcal{H}$:

$$\begin{aligned}
& \underset{W_{t+1}, \mathbf{x}_t, z}{\text{maximize}} && \mathbb{E} \left[(1 - \lambda) W_{t+1} + \lambda \left(z - \frac{(W_{t+1} - z)^-}{1 - \alpha} \right) \right] \\
& \text{subject to} && W_{t+1} = \sum_{i \in \mathcal{A}} (1 + r_{i,t+1}) x_{i,t} \\
& && \sum_{i \in \mathcal{A}} x_{i,t} = W_t \\
& && \mathbf{x}_t \geq 0.
\end{aligned} \tag{13}$$

For $W_0 = 1$, the (time consistent) optimal policy obtained by solving problem (12) is the following:

$$x_{1,t}^*(\omega) = W_t = 1 \quad \text{and} \quad x_{2,t}^*(\omega) = 0, \quad \forall t \in \mathcal{H}, \omega \in \Omega.$$

The optimal policy is always to invest the total wealth in the risk free asset. Note that this strategy is more conservative compared to the time inconsistent one, because it takes risk into account at every state of the system. Since this difference may lead to a sub-optimal solution for the time inconsistent problem, we develop in the following section a systematic way of measuring this effect on the objective function.

2.3. The time inconsistent sub-optimality gap

In Section 2.1, we presented a time inconsistent example where planned and implemented policies are different. However, what we do want to know is how a time inconsistent policy impacts our objective function. The way of measuring it is computing the sub-optimality gap that concerns the disparity of the objective function we expect to obtain with our planned policy, denoted by OF_{plan} , and the one we actually get when evaluating the policy to be implemented in the future, denoted by OF_{imp} . Then, let us define

$$gap = \frac{OF_{plan} - OF_{imp}}{OF_{plan}}.$$

For instance, take the portfolio selection problem defined in (6), in particular the numerical example in subsection 2.1. The optimal value $Q_0(W_0)$ defines OF_{plan} , i.e., $OF_{plan} = Q_0(W_0)$. On the other hand, OF_{imp} is obtained by computing the wealth distribution \widehat{W}_T at $t = T$ using the implemented decisions and then evaluate the objective function

$$OF_{imp} = \mathbb{E} \left[(1 - \lambda) \widehat{W}_T + \lambda \left(z - \frac{(\widehat{W}_T - z)^-}{1 - \alpha} \right) \right]. \quad (14)$$

The terminal wealth \widehat{W}_T is obtained by the following procedure:

for $\tau \in \mathcal{H}$, $\omega \in \Omega$: **do**

$\widehat{\mathbf{x}}_\tau = (\widehat{x}_{1,\tau}, \dots, \widehat{x}_{A,\tau})' \leftarrow$ the first stage solution of $Q_\tau(\widehat{W}_\tau(\omega), \bar{\mathbf{r}}_{[1,\tau]}(\omega))$

Compute $\widehat{W}_{\tau+1}(\omega) = \sum_{i \in \mathcal{A}} (1 + r_{i,\tau+1}(\omega)) \widehat{x}_{i,\tau}$

end for

where, for $\tau = 0$, $\widehat{W}_\tau(\omega) = W_0$ and $Q_\tau(\widehat{W}_\tau(\omega), \bar{\mathbf{r}}_{[1,\tau]}(\omega)) = Q_0(W_0)$ for all $\omega \in \Omega$.

For this numerical example, where $T = 2$, we first solve the three-stage problem $Q_0(W_0)$ to obtain: 1) the implemented decision at $\tau = 0$ under

the root node (Ω); 2) the planned decisions at $t = 1$ under nodes $\{1, 2\}$ and $\{3, 4\}$; and 3) the objective value considering the planned decisions $OF_{plan} = Q_0(W_0)$. After that, we compute for $\tau = 0$ the wealth $\widehat{W}_{\tau+1}(\omega)$ under nodes $\omega \in \{1, 2\}$ and $\omega \in \{3, 4\}$ given the implemented decisions $\widehat{\mathbf{x}}_\tau$ and the corresponding asset returns $r_{i,\tau+1}(\omega)$, $\forall i \in \mathcal{A}, \omega \in \Omega$. Now for $\tau = 1$, a two-stage problem for each node $\{1, 2\}$ and $\{3, 4\}$ is solved to obtain the implemented decisions at τ and compute the wealth $\widehat{W}_T(\omega)$ at $T = 2$ for all $\omega \in \Omega$. At last, we evaluate the objective function using the implemented decisions as in (14). In this numerical example where $\lambda = 0.5$ and $\alpha = 95\%$ we compute a gap of 9.09%.

For completeness, we run a sensitivity analysis on varying $\lambda \in \{0, 0.1, \dots, 0.9, 1\}$ and $T \in \{2, 3, \dots, 9, 10\}$. Note that, as before, we consider a two-asset model, i.e., $\mathcal{A} = \{1, 2\}$, and an event tree with 2^T scenarios. Note also that the first asset indexed by $i = 1$ is risk free and it has null excess return for every state of the system, i.e., $r_{1,t+1}(\omega) = 0$, $\forall t \in \mathcal{H}, \omega \in \Omega$. The second asset has iid returns where the conditional probability distribution is characterized by two equiprobable realizations, i.e., $\mathbb{P}(r_{2,t+1} = 100\% \mid r_{2,t}) = 1/2$ and $\mathbb{P}(r_{2,t+1} = -50\% \mid r_{2,t}) = 1/2$. We present the sub-optimality gap in Table 3 for differing values of λ and T to have a better assessment on how a time inconsistent policy would affect the decision process in practice.

We observe that the gap can be significantly different depending on planning horizon size T and risk aversion level λ . On the one hand, we observe a zero gap for $\lambda \in \{0, 1\}$ and $T \in \{2, \dots, 10\}$. For $\lambda = 0$ our results are validated by the fact that a risk neutral formulation ensures time consistency. For $\lambda = 1$, the preference function is too conservative in this particular example and, therefore, planned and implemented decisions are to invest everything in the risk free asset. On the other hand, for $\lambda \in \{0.4, 0.5, 0.6, 0.7, 0.9\}$ we




Figure 2: Sub-Optimality gap due to Time Inconsistency

observe significantly high gap values for different planning horizon sizes T . We illustrate this behavior in Figure 2 where we plot for each λ , the gap as a function of T . For a given λ , we observe that the sub-optimality gap increases with the planning horizon until it stabilizes around a certain value. Nonetheless the gap for a smaller λ increases faster with T , the case with a larger λ stabilizes at a higher gap value.

To sum up, for the most aggressive and the most conservative allocation strategies, time inconsistency would not be a problem. On the other hand, for $0 < \lambda < 1$, time consistency do matter and can have a significant impact on the objective function. For this reason, we reassure the importance of time consistency which can be achieved by the recursive formulation illustrated in Section 2.2. The major problem of utilizing this formulation in practice is the difficulty to economically interpret it. Therefore, we develop a suitable economic interpretation for this recursive objective function based on the certainty equivalent of the related preference function.

3. Economic interpretation

When choosing the proposed recursive set up, the problem is usually the lack of a suitable economic interpretation of the objective function. How can an investor choose a policy if he / she does not know what is actually going to be optimized? For this reason, we prove that the objective function is the certainty equivalent w.r.t. the time consistent dynamic utility generated by one period preference functionals.

Let us consider a generic one period preference functional $\psi_t : L^\infty(\mathcal{F}_{t+1}) \rightarrow L^\infty(\mathcal{F}_t)$ and, for a particular realization sequence of the uncertainty $\bar{\mathbf{r}}_{[1,t]}$, the related real valued function $\psi_t(\cdot | \bar{\mathbf{r}}_{[1,t]}) : L^\infty(\mathcal{F}_{t+1}) \rightarrow \mathbb{R}$. Moreover, let us define important properties and concepts of $\psi_t(\cdot | \bar{\mathbf{r}}_{[1,t]})$ that will be used to develop our main results:

Translation invariance: $\psi_t(W_{t+1} + m | \bar{\mathbf{r}}_{[1,t]}) = \psi_t(W_{t+1} | \bar{\mathbf{r}}_{[1,t]}) + m$, where $W_{t+1} \in L^\infty(\mathcal{F}_{t+1})$ and $m \in \mathbb{R}$.

Monotonicity: $\psi_t(X_{t+1} | \bar{\mathbf{r}}_{[1,t]}) \geq \psi_t(Y_{t+1} | \bar{\mathbf{r}}_{[1,t]})$ for all $X_{t+1}, Y_{t+1} \in L^\infty(\mathcal{F}_{t+1})$ with $X_{t+1}(\omega) \geq Y_{t+1}(\omega), \forall \omega \in \Omega$.

Normalization: $\psi_t(\cdot | \bar{\mathbf{r}}_{[1,t]})$ is normalized to zero, that is $\psi_t(0 | \bar{\mathbf{r}}_{[1,t]}) = 0$.

Definition 1. *The certainty equivalent $\tilde{C}_t(W_{t+1} | \bar{\mathbf{r}}_{[1,t]})$ of $W_{t+1} \in L^\infty(\mathcal{F}_{t+1})$ with respect to $\psi_t(W_{t+1} | \bar{\mathbf{r}}_{[1,t]})$ is a real value $m \in \mathbb{R}$ such that $\psi_t(m | \bar{\mathbf{r}}_{[1,t]}) = \psi_t(W_{t+1} | \bar{\mathbf{r}}_{[1,t]})$.*

Let us denote by $(U_t)_{t \in \mathcal{H}}$ the time consistent dynamic utility function generated by ψ_t (see Cheridito and Kupper (2009) for details). Formally speaking, $U_t : L^\infty(\mathcal{F}_T) \rightarrow L^\infty(\mathcal{F}_t)$ is defined as follows:

$$U_T(W_T) = W_T \quad \text{and} \quad U_t(W_T) = \psi_t(U_{t+1}(W_T)), \forall t \in \mathcal{H},$$

where $W_T \in L^\infty(\mathcal{F}_T)$. Note that we can also use a conditional version of U_t as follows:

$$U_t(W_T \mid \bar{\mathbf{r}}_{[1,t]}) = \psi_t(U_{t+1}(W_T) \mid \bar{\mathbf{r}}_{[1,t]}), \forall t \in \mathcal{H}.$$

The concept of certainty equivalent can also be developed for the time consistent dynamic utility U_t :

Definition 2. *The certainty equivalent $C_t(W_T \mid \bar{\mathbf{r}}_{[1,t]})$ of $W_T \in L^\infty(\mathcal{F}_T)$ with respect to $U_t(W_T \mid \bar{\mathbf{r}}_{[1,t]})$ is a real value $m \in \mathbb{R}$ such that $U_t(m \mid \bar{\mathbf{r}}_{[1,t]}) = U_t(W_T \mid \bar{\mathbf{r}}_{[1,t]})$.*

Now, let us define the following dynamic stochastic programming model where the value function at time t depends on the decisions at $t-1$ and the realization sequence of the uncertainty until t . Thus, for $t = T$ we define it as follows:

$$\mathcal{V}_T(\mathbf{x}_{T-1}, \bar{\mathbf{r}}_{[1,T]}) = W_T(\mathbf{x}_{T-1}, \bar{\mathbf{r}}_{[1,T]}),$$

where $W_T = W_T(\mathbf{x}_{T-1}, \bar{\mathbf{r}}_{[1,T]})$ is a real valued function.

For $t \in \mathcal{H}$, we define

$$\mathcal{V}_t(\mathbf{x}_{t-1}, \bar{\mathbf{r}}_{[1,t]}) = \sup_{\mathbf{x}_t \in \mathcal{X}_t} \psi_t(\mathcal{V}_{t+1}(\mathbf{x}_t, \mathbf{r}_{[1,t+1]}) \mid \bar{\mathbf{r}}_{[1,t]}),$$

where $\mathcal{X}_t = \mathcal{X}_t(\mathbf{x}_{t-1}, \bar{\mathbf{r}}_{[1,t]})$ is the feasible set for each time t . Note that for $t = 0$, we have

$$\mathcal{V}_0 = \sup_{\mathbf{x}_0 \in \mathcal{X}_0} \psi_0(\mathcal{V}_1(\mathbf{x}_0, \mathbf{r}_1)),$$

where \mathcal{X}_0 is a deterministic set. Then, we develop the following results.

Proposition 3. *If ψ_t is a translation invariant, monotone functional normalized to zero, then for $t \in \mathcal{H}$ the value function can be written as*

$$\mathcal{V}_t(\mathbf{x}_{t-1}, \bar{\mathbf{r}}_{[1,t]}) = \sup_{\mathbf{x}_\tau \in \mathcal{X}_\tau, \forall \tau=t, \dots, T-1} C_t(W_T \mid \bar{\mathbf{r}}_{[1,t]}),$$

where $C_t(W_T \mid \bar{\mathbf{r}}_{[1,t]})$ is the certainty equivalent of W_T w.r.t. U_t conditioned on the realization sequence $\bar{\mathbf{r}}_{[1,t]}$.

Proof. By definition we have

$$\begin{aligned} \mathcal{V}_t(\mathbf{x}_{t-1}, \bar{\mathbf{r}}_{[1,t]}) &= \sup_{\mathbf{x}_t \in \mathcal{X}_t} \psi_t(\mathcal{V}_{t+1}(\mathbf{x}_t, \mathbf{r}_{[1,t+1]}) \mid \bar{\mathbf{r}}_{[1,t]}) \\ &= \sup_{\mathbf{x}_t \in \mathcal{X}_t} \psi_t \left(\dots \sup_{\mathbf{x}_{T-1} \in \mathcal{X}_{T-1}} \psi_{T-1}(W_T) \mid \bar{\mathbf{r}}_{[1,t]} \right). \end{aligned}$$

Using the monotonicity of ψ_t and the definition of U_t we have:

$$\begin{aligned} \mathcal{V}_t(\mathbf{x}_{t-1}, \bar{\mathbf{r}}_{[1,t]}) &= \sup_{\mathbf{x}_\tau \in \mathcal{X}_\tau, \forall \tau=t, \dots, T-1} \psi_t(\dots \psi_{T-1}(W_T) \mid \bar{\mathbf{r}}_{[1,t]}) \quad (15) \\ &= \sup_{\mathbf{x}_\tau \in \mathcal{X}_\tau, \forall \tau=t, \dots, T-1} U_t(W_T \mid \bar{\mathbf{r}}_{[1,t]}). \end{aligned}$$

By the certainty equivalent definition, Definition 2, we have that $C_t(W_T \mid \bar{\mathbf{r}}_{[1,t]})$ satisfies $U_t(C_t(W_T \mid \bar{\mathbf{r}}_{[1,t]}) \mid \bar{\mathbf{r}}_{[1,t]}) = U_t(W_T \mid \bar{\mathbf{r}}_{[1,t]})$. It is easy to show that $U_t(\cdot \mid \bar{\mathbf{r}}_{[1,t]})$ is translation invariant and normalized to zero, since its generators ψ_t have the same properties. Then, $U_t(C_t(W_T \mid \bar{\mathbf{r}}_{[1,t]}) \mid \bar{\mathbf{r}}_{[1,t]}) = C_t(W_T \mid \bar{\mathbf{r}}_{[1,t]})$ and consequently, $U_t(W_T \mid \bar{\mathbf{r}}_{[1,t]}) = C_t(W_T \mid \bar{\mathbf{r}}_{[1,t]})$.

Finally we have that

$$\mathcal{V}_t(\mathbf{x}_{t-1}, \bar{\mathbf{r}}_{[1,t]}) = \sup_{\mathbf{x}_\tau \in \mathcal{X}_\tau, \forall \tau=t, \dots, T-1} C_t(W_T \mid \bar{\mathbf{r}}_{[1,t]}).$$

□

Corollary 4. *If ψ_t is a translation invariant, monotone functional normalized to zero, then for $t \in \mathcal{H}$ the value function can be written as*

$$\mathcal{V}_t(\mathbf{x}_{t-1}, \bar{\mathbf{r}}_{[1,t]}) = \sup_{\mathbf{x}_\tau \in \mathcal{X}_\tau, \forall \tau=t, \dots, T-1} \tilde{C}_t \left(\dots \tilde{C}_{T-1}(W_T) \mid \bar{\mathbf{r}}_{[1,t]} \right),$$

where \tilde{C}_t and $\tilde{C}_t(\cdot \mid \bar{\mathbf{r}}_{[1,t]})$ are the certainty equivalent w.r.t. ψ_t and $\psi_t(\cdot \mid \bar{\mathbf{r}}_{[1,t]})$, respectively.

Proof. By Definition 1, we have that $\tilde{C}_t(\cdot \mid \bar{\mathbf{r}}_{[1,t]})$ satisfies $\psi_t(\tilde{C}_t(\cdot \mid \bar{\mathbf{r}}_{[1,t]}) \mid \bar{\mathbf{r}}_{[1,t]}) = \psi_t(\cdot \mid \bar{\mathbf{r}}_{[1,t]})$ and using the assumption that ψ_t is translation invariant and normalized to zero, we have $\psi_t(\tilde{C}_t(\cdot)) = \tilde{C}_t(\cdot)$. Note that this property also holds true for the conditional version. Then, from equation (15) we have:

$$\begin{aligned} \mathcal{V}_t(\mathbf{x}_{t-1}, \bar{\mathbf{r}}_{[1,t]}) &= \sup_{\mathbf{x}_\tau \in \mathcal{X}_\tau, \forall \tau=t, \dots, T-1} \psi_t(\dots \psi_{T-1}(W_T) \mid \bar{\mathbf{r}}_{[1,t]}) \\ &= \sup_{\mathbf{x}_\tau \in \mathcal{X}_\tau, \forall \tau=t, \dots, T-1} \tilde{C}_t(\dots \tilde{C}_{T-1}(W_T) \mid \bar{\mathbf{r}}_{[1,t]}). \end{aligned}$$

□

Note that we could also include intermediate “costs” as in Ruszczyński and Shapiro (2006) and our results would still hold true for a more general set of problems. It is worth mentioning that we define the feasible sets \mathcal{X}_t for $t \in \mathcal{H}$, and the terminal wealth function $W_T(x_{T-1}, r_{[1,T]})$ generically depending on the application. For the portfolio selection problem, we define them to fit the original constraints. Then, we have that

$$\begin{aligned} \mathcal{X}_t(\mathbf{x}_{t-1}, \bar{\mathbf{r}}_{[1,t]}) &= \{\mathbf{x}_t \in \mathbb{R}^A : \sum_{i \in \mathcal{A}} x_{i,t} = \sum_{i \in \mathcal{A}} (1 + \bar{r}_{i,t}) x_{i,t-1}\}, \\ \mathcal{X}_0 &= \{x_0 \in \mathbb{R}^A : \sum_{i \in \mathcal{A}} x_{i,0} = W_0\}, \\ W_T(\mathbf{x}_{T-1}, \bar{\mathbf{r}}_{[1,T]}) &= \sum_{i \in \mathcal{A}} (1 + \bar{r}_{i,T}) x_{i,T-1}. \end{aligned}$$

For the proposed portfolio selection model, we define our one period translation invariant, monotone and normalized utility functional ψ_t as the convex combination of the expected value and the CVaR based acceptability measure, formally defined as $\psi_t(\mathcal{V}_{t+1}) = (1 - \lambda) \mathbb{E}[\mathcal{V}_{t+1} \mid \mathbf{r}_{[1,t]}] + \lambda \phi_t^\alpha(\mathcal{V}_{t+1}, \mathbf{r}_{[1,t]})$, which is again a coherent acceptability measure. As before, $\mathcal{V}_{t+1} \in L^\infty(\mathcal{F}_{t+1})$ and we can write the conditional version as the real valued function $\psi_t(\mathcal{V}_{t+1} \mid \bar{\mathbf{r}}_{[1,t]}) = (1 - \lambda) \mathbb{E}[\mathcal{V}_{t+1} \mid \bar{\mathbf{r}}_{[1,t]}] + \lambda \phi_t^\alpha(\mathcal{V}_{t+1}, \bar{\mathbf{r}}_{[1,t]})$.

The objective function of the proposed model at t is the certainty equivalent w.r.t. the time consistent dynamic utility function generated by the

one period preference functional of the investor. This recursive formulation ensures time consistent optimal policies and it is also motivated by Corollary 4. The objective at $t = T - 1$ is to maximize the certainty equivalent (CE) of terminal wealth w.r.t. the one period preference functional ψ_{T-1} . Indeed, we can interpret the optimal CE as the portfolio value since it is the deterministic amount of money the investor would accept instead of the (random) terminal wealth obtained by his / her optimal trading strategy. At $t = T - 2, \dots, 0$, the preference functional ψ_t is applied to the (random) portfolio value whose realizations are given by all possible optimal CE's at $t + 1$. Thus, one obtains the following meaningful economic interpretation of the recursive objective function: *the problem at time t is to maximize the certainty equivalent of the portfolio value w.r.t. the one period preference functional ψ_t of the investor.*

For instance, in our numerical example the (random) portfolio value at $t = 1$ is given by the realizations v_1 and v_2 in Figure 3 obtained by solving problem (13) for nodes $\{1, 2\}$ and $\{3, 4\}$, respectively. The portfolio value v_0 (see also Figure 3) obtained by solving (13) for $t = 0$ is the optimal certainty equivalent of the (random) portfolio value at $t = 1$.

4. Conclusions

This paper contributes on closing the gap between theory and practice regarding time consistency and stochastic programming. It motivates and resolves one of the main conceptual issues regarding practical usage of time consistent risk-averse dynamic stochastic programming models: The complex objective function given by the recursive formulation that ensures time consistency has no clear/direct interpretation. Practitioners resist on using time consistent models because they have no intuition about what is being opti-

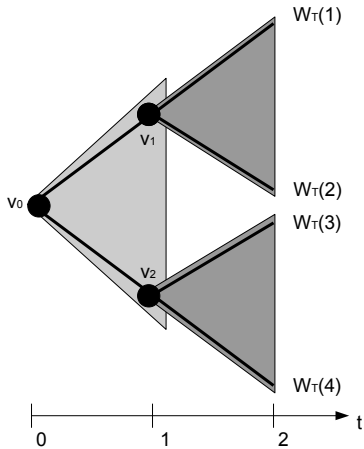


Figure 3: Conditional certainty equivalents

mized. They rather choose a simpler and more intuitive formulation, even though it may lead to a time inconsistent and, consequently, sub-optimal policy. Therefore, we developed an intuitive economic interpretation for the complex time-consistent recursive formulation as well as a new methodology to compute the sub-optimality gap of a time inconsistent policy. Our results hold for a quite general class of problems and we choose, with loss of generality, a portfolio selection application to illustrate the developed concepts. Numerical results presented in this paper aims at reinforcing intuition and interpretation of the reader regarding our theoretical results.

We prove that the recursive objective function is the certainty equivalent with respect to the time consistent dynamic utility function. We also prove that this objective is the composed form of certainty equivalents with respect to one-period preference functionals. This result gives us the intuition that at stage t the agent is maximizing the certainty equivalent with respect to his / her one-period preference functional.

Moreover, we consider the conceptual definition: *a policy is time consistent if and only if the future planned decisions are actually going to be*

implemented. Based on that, we discouraged time-inconsistent models by arguing that the associated first-stage decision is, in general, sub-optimal. With the purpose of providing practitioners with an objective way to quantify time inconsistency, we proposed a methodology to compute the related sub-optimality gap as the difference of the objective function evaluated using planned and implemented policies.

For the proposed CVaR based portfolio selection problem, we compared a time-consistent formulation to an inconsistent one for the sake of showing that an inconsistent policy may not take risk aversion into account at some intermediate states of the system. We show a sensitivity analysis on the sub-optimality gap with respect to planning horizon size and risk-aversion level. We concluded that, for this example, the most aggressive and the most conservative allocation strategies, time inconsistency would not be a problem. On the other hand, for intermediate cases, time consistency does matter depending on the planning horizon.

Finally, it is worth mentioning that future research regarding the assessment of the sub-optimality gap deserves attention. Applying this methodology in different modeling frameworks, applications, risk measures, and in the context of large-scale approximative algorithms, constitute relevant extensions to this subject.

Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.

Bäuerle, N. and Mundt, A. (2009). Dynamic mean-risk optimization in a binomial model. *Mathematical Methods of Operations Research*, 70:219–239.

- Bion-Nadal, J. (2008). Dynamic risk measures: Time consistency and risk measures from BMO martingales. *Finance and Stochastics*, 12(2):219–244.
- Blomvall, J. and Shapiro, A. (2006). Solving multistage asset investment problems by the sample average approximation method. *Mathematical Programming*, 108(2):571–595.
- Boda, K. and Filar, J. (2006). Time consistent dynamic risk measures. *Mathematics of Operations Research*, 63:169–186.
- Cheridito, P., Delbaen, F., and Kupper, M. (2006). Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability*, 11(3):57–106.
- Cheridito, P. and Kupper, M. (2009). Recursiveness of indifference prices and translation-invariant preferences. *Mathematics and Financial Economics*, 2(3):173–188.
- Cuoco, D., He, H., and Issaenko, S. (2008). Optimal. *Optimal Dynamic Trading Strategies with Risk Limits*, 56(2):358–368.
- Detlefsen, K. and Scandolo, G. (2005). Conditional and dynamic convex risk measures. *Finance and Stochastics*, 9(4):539–561.
- Fábián, C. and Veszprémi, A. (2008). Algorithms for handling cvar constraints in dynamic stochastic programming models with applications to finance. *Journal of Risk*, 10(3):111–131.
- Guigues, V. and Römisch, W. (2012a). Sampling-based decomposition methods for multistage stochastic programs based on extended polyhedral risk measures. *SIAM Journal on Optimization*, 22(2):286–312.

- Guigues, V. and Römisch, W. (2012b). SDDP for multistage stochastic linear programs based on spectral risk measures. *Operations Research Letters*, 40(5):313–318.
- Kovacevic, R. and Pflug, G. (2009). Time consistency and information monotonicity of multiperiod acceptability functionals. *Radon Series on Computational and Applied Mathematics*, 8:1–24.
- Krokhmal, P., Palmquist, J., and Uryasev, S. (2002). Portfolio optimization with conditional value-at-risk objective and constraints. *The Journal of Risk*, 4(2):11–27.
- Mansini, R., Ogryczak, W., and Speranza, M. G. (2007). Conditional value at risk and related linear programming models for portfolio optimization. *Annals of Operations Research*, 152:227–256.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1):77–91.
- Pereira, M. V. F. and Pinto, L. M. V. G. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52:359–375. 10.1007/BF01582895.
- Philpott, A. B. and de Matos, V. L. (2012). Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research*, 218(2):470–483.
- Philpott, A. B., de Matos, V. L., and Finardi, E. (2012). On solving multistage stochastic programs with coherent risk measures. *Optimization online*.
- Riedel, F. (2004). Dynamic coherent risk measures. *Stochastic Processes and their Applications*, 112(2):185–200.

- Rockafellar, R. T. and Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42.
- Rockafellar, R. T. and Wets, R. J.-B. (1991). Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research*, 16(1):pp. 119–147.
- Roorda, B. and Schumacher, J. (2007). Time consistency conditions for acceptability measures, with an application to tail value at risk. *Insurance: Mathematics and Economics*, 40(2):209–230.
- Ruszczynski, A. (2010). Risk-averse dynamic programming for markov decision processes. *Mathematical Programming*, 125(2):235–261.
- Ruszczynski, A. and Shapiro, A. (2006). Conditional risk mappings. *Mathematics of Operations Research*, 31(3):544–561.
- Sawik, B. (2012). Bi-criteria portfolio optimization models with percentile and symmetric risk measures by mathematical programming. *Przegląd Elektrotechniczny (Electrical Review)*, 88(10):176–180.
- Shapiro, A. (2009). On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters*, 37(3):143–147.
- Shapiro, A. (2011). Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209(1):63–72.
- Shapiro, A., Tekaya, W., da Costa, J., and Soares, M. (2013). Risk neutral and risk averse stochastic dual dynamic programming method. *European Journal of Operational Research*, 224(2):375–391.
- Street, A. (2009). On the conditional value-at-risk probability-dependent utility function. *Theory and Decision*, 68(1-2):49–68.

Table 3: Results of the sensitivity analysis

λ	T									
	2	3	4	5	6	7	8	9	10	
0	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
0.1	0.00%	0.00%	0.12%	0.00%	0.00%	0.02%	0.00%	0.00%	0.00%	0.01%
0.2	0.00%	0.78%	0.63%	0.19%	0.09%	0.11%	0.04%	0.02%	0.02%	0.03%
0.3	2.60%	2.18%	1.80%	0.63%	0.40%	0.32%	0.13%	0.09%	0.09%	0.08%
0.4	13.64%	19.42%	20.85%	21.20%	20.36%	20.24%	20.26%	20.09%	20.09%	20.06%
0.5	9.09%	22.97%	27.95%	29.24%	28.99%	28.40%	28.35%	28.24%	28.24%	28.12%
0.6	0.00%	14.29%	32.98%	40.49%	42.83%	42.94%	42.54%	42.61%	42.61%	42.57%
0.7	0.00%	5.62%	20.90%	40.78%	49.37%	51.66%	51.95%	51.73%	51.73%	51.80%
0.8	0.00%	0.00%	2.54%	10.81%	29.67%	53.15%	63.03%	66.62%	66.62%	67.58%
0.9	0.00%	0.00%	0.00%	0.00%	0.41%	3.85%	14.17%	36.41%	36.41%	63.09%
1	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%