Rutcor Research Report

Time Consistency Versus Law Invariance in Multistage Stochastic Optimization with Coherent Risk Measures: Multilevel Optimization Modeling and Computational Complexity

Jonathan Eckstein[a]

RUTCOR
Rutgers Center for
Operations Research
Rutgers University
640 Bartholomew Road
Piscataway, New Jersey
08854-8003
Telephone: 732-445-3804
Telefax: 732-445-5472
Email: rrr@rutcor.rutgers.edu
http://rutcor.rutgers.edu/∼rrr

RRR 7-2013, August 2013

[a]Department of Management Science and Information Systems (MSIS) and RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway, NJ 08854-8003

Rutcor Research Report
RRR 7-2013, August 2013

Time Consistency Versus Law Invariance in Multistage Stochastic Optimization with Coherent Risk Measures: Multilevel Optimization Modeling and Computational Complexity

Jonathan Eckstein

Abstract. Coherent risk measures have become a popular tool for incorporating risk aversion into stochastic optimization models. For dynamic models in which uncertainly is resolved at more than one stage, however, use of coherent risk measures within a standard single-level optimization framework presents the modeler with an uncomfortable choice between two desirable model properties, time consistency and law invariance. Prior published work has favored maintaining time consistency, but the absence of law invariance makes the resulting models unattractive to practical decision makers. This paper summarizes these issues and then presents an alternative multilevel optimization modeling approach that preserves law invariance, yet leads to *models* that are time-consistent even while using time-inconsistent risk measures. It argues that this approach should be the starting point for all multistage optimization modeling; however, when performing classical risk-neutral modeling, it simplifies to a more familiar single-objective form.

Unfortunately, this paper also shows that its proposed approach leads to $\mathcal{NP}$-hard models, even in the simplest imaginable setting in which it would be needed: three-stage linear problems on a finite probability space, using the standard mean-semideviation and average-value-at-risk measures. While not necessarily indicating that solution of such models is impractical, these results suggest that it will likely require approximation or implicit enumeration methods. We close with some preliminary computational results showing that high-quality local optimum solutions of models of the kind we propose are in fact practically computable, hence that the complexity results should not be taken as completely discouraging.

Acknowledgements: This work was supported by a grant from the Air Force Office of Scientific Research, award FA9550-11-1-0164. The author would like to thank Andrzej Ruszczyński for his very helpful assistance, and his student Jingnan Fan for performing the computational experiments described in Section 8. The results in Sections 5-7 also appeared in Eckstein (2012), an earlier, unpublished technical report.

1 Introduction

Coherent risk measures (Artzner et al. 1999, Rockafellar et al. 2006, Delbaen 2002) have become a popular tool for incorporating risk aversion into convex stochastic optimization models; if uncertainty is resolved at a single point in time, one may simply substitute a coherent risk measure for the classical expectation operator in the problem objective and convert a risk-neutral model to a risk-averse one, without losing any convexity that may have been present. If uncertainty is resolved in multiple stages, however, using coherent risk measures becomes more problematic. If one simply substitutes a general coherent risk measure for expectation in the objective function of a multistage stochastic model and its associated recourse problems, one will in general produce a system of models that violates the principle of time consistency. A definition of time consistency will be given below, but in practice the absence of time consistency in a multistage model has the following consequence: fully optimizing the objective function for the first stage may require assuming that in some particular scenario at some later stage, we will take an action that is suboptimal for the recourse problem that we would face in that particular later situation. In fact, as will be discussed in Section 9, Sections 6 and 7 will present examples of relatively simple models, using the most commonly employed risk measures in a natural manner, in which optimizing the first-stage objective requires that in one possible second-stage scenario, we must select the *worst* possible feasible course of action from the point of view of that scenario. Models with this kind of property are clearly indefensible and should not be used as the basis for decision making.

One way to avoid such difficulties is to employ time-consistent dynamic measures of risk (Scandolo 2003, Riedel 2004, Cheridito et al. 2006, Ruszczyński and Shapiro 2006a) whenever uncertainty is resolved at more than stage. These risk measures assign a scalar certainty equivalent not to a single random variable, but to a sequence of random variables evolving over time, and provide time-consistent objective functions that maintain any existing problem convexity. Regrettably, while time-consistent dynamic risk measures allow the modeler to maintain time consistency, they do so at the cost of another very desirable model property, law invariance. In brief, the law-invariance principle states that if the respective sums of two sequences of random variables have identical distributions, meaning that their final outcomes are essentially indistinguishable at the end of the time horizon, then the two sequences should be assessed as having the same level of risk. For sequences of random variables of length one, this principle simplifies to requiring that any two random variables having the same distribution be assessed as presenting the same level of risk. Aside from its intuitive appeal, this principle also accords with standard modeling practices: modelers tend to think of risk measures not as being specific to some particular abstract probability space, but instead as overall "risk attitudes" that may be applied to random variables defined on different spaces. Risk measures that are defined as a function of the distribution of the total "cost to go", and which are therefore law-invariant, are thus one of the most natural to use in modeling. All the standard examples of coherent single-period risk measures, such as the average-value-at-risk and mean-semidevation measures, are law-invariant.

Unfortunately, Shapiro (2012) has shown that time-consistent coherent dynamic risk measures cannot be law-invariant except in the the cases of complete risk neutrality or maximum risk aversion. Besides the counterintuitive assignment of different risk levels to variables with identical distributions, lack of law invariance also implies that the complicated nested formulas required to express dynamic risk measures cannot, except in two extreme cases, be simplified into forms that decision makers might find easier to understand. For this reason, it seems unlikely that multistage models involving time-consistent coherent dyamic risk measures will find much acceptance in practice.

This paper explores a different approach to modeling multistage stochastic optimization problems, an approach which allows law invariance to be maintained while permitting the modeler to select any desired objective function for every stage and scenario. It uses a multilevel optimization modeling approach in which one engages in a kind of Stackelberg game "against oneself". Section 3 will argue that this kind of approach should be the proper starting point for all multistage decision modeling; the traditional single-level optimization forms are simply convenient simplifications that occur when one uses a classical expected-value objective function at every decision point. The advantage of such a multilevel approach is that, even while using a time-inconsistent risk measure, one can construct models that have a reasonable form of internal time consistency and are thus defensible in practice.

Of course, the drawback of this approach is that multilevel optimization models tend to be far more challenging to solve than single-level models. It is well known that even bilevel linear programming models are $\mathcal{NP}$-hard (Jeroslow 1985, Ben-Ayed and Blair 1990, Bard 1991, Hansen et al. 1992). In Sections 6 and 7, it will be shown that the simplest imaginable application of the proposed modeling approach, while not necessarily leading to fully general bilevel linear programs, does indeed result in $\mathcal{NP}$-hard models. This finding does not mean that the proposed modeling approach is necessarily fruitless: in the practice of operations research, it is not uncommon to solve instances of $\mathcal{NP}$-hard problem classes, especially integer programs, essentially exactly. However, the $\mathcal{NP}$-hardness result does suggest that computing exact solutions to these models will require some form of implicit enumeration method, or that approximation methods may have to be used.

The rest of this paper is organized as follows: Section 2 lays out the formal underpinnings of the discussion, presenting the notion of coherent risk measures and a formal general multi-stage stochastic problem setting patterned after Ruszczyński and Shapiro (2006a). It also formally defines time consistency and law invariance, summarizing the basic incompatibility between them when using conventional single-level optimization modeling over multiple time periods.

Section 3 then describes how multilevel optimization provides a way to resolve the conflict between time consistency and law invariance while still yielding models that have a reasonable form of time consistency despite using time-inconsistent risk measures. Section 4 then specializes the very general model of Section 3 to the simplest setting in which single-level optimization modeling exhibits incompatibility between time consistency and law invariance; we call this model class *bilevel risk programming*. Section 5 next revisits the complexity theory of bilevel linear programming, and then Sections 6 and 7 apply this complexity theory to the bilevel risk programming models defined in Section 4, showing that they are $\mathcal{NP}$-hard when using mean-standard-deviation and average-value-at-risk risk measures. Contrasting with the somewhat negative conclusions of Sections 6 and 7, Section 8 describes some preliminary computational results indicating that it may be practical to compute high-quality local optima of the kind of multilevel models we propose. Section 9 presents some concluding remarks.

2 Risk Measures and Abstract Problem Setting

Consider a probability space $(\Omega, \mathcal{F}, \mathrm{P})$, where Ω is some set of outcomes, $\mathcal{F}$ is a σ-algebra over Ω, and P is a probability measure over $\mathcal{F}$. Also consider some linear space of $\mathcal{Z}$ of random variables over $(\Omega, \mathcal{F})$, that is, $\mathcal{F}$-measurable functions $Z : \Omega \to \mathbb{R}$. Specifically, we let $\mathcal{Z} = \mathcal{L}^p(\Omega, \mathcal{F})$ for some $p \geq 1$.

2.1 Risk measures on $\mathcal{Z}$

Here, we are concerned with a risk-averse decision maker using some function $\rho : \mathcal{Z} \to \mathbb{R}$ to express preferences between elements of $\mathcal{Z}$. We may think ρ as a function mapping a random variable to a scalar "certainty equivalent". Such a mapping ρ is called a *coherent risk measure* (Artzner et al. 1999, Delbaen 2002, Rockafellar et al. 2006) if it has the following properties:

Monotonicity. If $Z_1 \leq Z_2$ (that is, $Z_1(\omega) \leq Z_2(\omega)$ for all $\omega \in \Omega$), then we must have $\rho(Z_1) \leq \rho(Z_2)$.

Convexity. If $\alpha \in [0, 1]$, then $\rho\big(\alpha Z_1 + (1 - \alpha)Z_2\big) \leq \alpha\rho(Z_1) + (1 - \alpha)\rho(Z_2)$.

Positive homogeneity. If $\alpha \geq 0$, then $\rho(\alpha Z) = \alpha\rho(Z)$.

Translation invariance. For any $t \in \mathbb{R}$, we have $\rho(Z + t) = \rho(Z) + t$.

The classical expected-value mapping $\mathbb{E}[\,\cdot\,]$ satisfies these axioms, but so do many other functions. One common choice is the mean-semideviation

$$\mathrm{MSD}_\gamma : Z \mapsto \mathbb{E}[Z] + \gamma\mathbb{E}\big[[Z - \mathbb{E}[Z]]_+\big] \qquad \text{for } \gamma \in [0, 1], \tag{1}$$

where $[x]_+ = \max\{0, x\}$ denotes the positive part of a number. Another common choice of risk measure is *average value at risk* (Acerbi 2002, Föllmer and Schied 2004, Rockafellar and Uryasev 2000, Uryasev and Rockafellar 2001, Rockafellar and Uryasev 2002), also called *conditional value at risk*, computed in the general case by the formula

$$\mathrm{AVaR}_\alpha : Z \mapsto \frac{1}{\alpha}\int_{1-\alpha}^{1} F_Z^{-1}(\nu)\, d\nu \qquad \text{for } \alpha \in (0, 1), \tag{2}$$

where F_Z denotes the cumulative distribution function of Z and F_Z^{-1} denotes its "lower" inverse

$$F_Z^{-1}(\nu) = \inf\{x \mid F(x) \geq \nu\} = \inf\big\{x \mid \mathrm{P}\{Z \leq x\} \geq \nu\big\}. \tag{3}$$

When Z has a continuous distribution, an alternative and simpler expression for AVaR_α is

$$\mathrm{AVaR}_\alpha(Z) = \mathbb{E}\big[Z \;\big|\; Z \geq F_Z^{-1}(1-\alpha)\big],$$

that is, the expected value of Z given that its value is at least its $1-\alpha$ quantile; (2) properly generalizes this expression to the discrete and general cases.

Another, extreme example of a coherent risk measure is the "worst outcome" measure

$$\operatorname{ess\,sup} Z = \inf\big\{b \in \mathbb{R} \mid \mathrm{P}\{Z \leq b\} = 1\big\}.$$

Over a finite probability space, $\operatorname{ess\,sup} Z$ is simply the largest value of Z that can occur.

Note that the functions $\mathbb{E}[\,\cdot\,]$, MSD_γ, AVaR_α, and ess sup, being defined as functions of the distribution of their argument, are all law-invariant, meaning that if two random variables Z and W have identical distributions, then $\mathbb{E}[Z] = \mathbb{E}[W]$, $\mathrm{MSD}_\gamma(Z) = \mathrm{MSD}_\gamma(W)$, and so forth. In fact, we will use symbols such as MSD_γ and AVaR_α to refer not to a risk measure defined only for a particular space $(\Omega, \mathcal{F}, \mathrm{P})$, but to a distribution-based rule for generating a risk measure for any such space.

2.2 Multiple stages of uncertainty resolution and dynamic risk-measure systems

In stochastic models, it is common for uncertainty to be resolved in a sequence of stages: for example, Ω may model the evolution of a stock index which is recomputed on a daily basis. To model such situations, we introduce a filtration (sequence of σ-algebras) $\{\emptyset, \Omega\} = \mathcal{F}_1 \subset \mathcal{F}_2 \subset \cdots \subset \mathcal{F}_T = \mathcal{F}$. We say that $\omega, \omega' \in \Omega$ are *distinguishable* at time τ if there exists an event $E \in \mathcal{F}_\tau$ that contains only ω and not ω' (and hence the reverse, by the rules of σ-algebras). Otherwise, we call ω and ω' *indistinguishable* at time τ, meaning that every event in $\mathcal{F}_\tau$ contains either both or neither of ω, ω'. For each $t = 1 \ldots, T$, let $\mathcal{Z}_t$ denote the subspace of $\mathcal{Z}$ consisting of functions that are $\mathcal{F}_t$-measurable, and consider sequences of random variables $Z_1, \ldots, Z_T$ adapted to the filtration, meaning that $Z_t \in \mathcal{Z}_t$ for each t. Note that since $\mathcal{F}_1 = \{\emptyset, \Omega\}$, an $\mathcal{F}_1$-measurable random variable is simply a deterministic quantity, and thus Z_1 is deterministic. To compare two such sequences of random variables, along with "tail" subsequences of the form $Z_\tau, Z_{\tau+1}, \ldots, Z_T$, where $1 \leq \tau < T$, we introduce the notion of a *dynamic risk measure* $\rho_{\tau,T} : \mathcal{Z}_\tau \times \mathcal{Z}_{\tau+1} \times \cdots \times \mathcal{Z}_T \to \mathcal{Z}_\tau$. Such a function converts a sequence of costs respectively measurable at times $\tau, \ldots, T$ to a "partial certainty equivalent" measurable at time τ. For the case $\tau = 1$, we have that $\rho_{1,T}(Z_1, \ldots, Z_T)$ must be $\mathcal{F}_1$-measurable, and so may be considered as being in $\mathbb{R}$. We make the following assumption, which any reasonably-chosen system of risk measures $\{\rho_{\tau,T}\}_{t=1}^T$ ought to obey:

Assumption 1 *For any $\omega \in \Omega$ and $1 \leq \tau \leq t \leq T$, the value of $\rho_{\tau,T}(Z_\tau, \ldots, Z_T)(\omega)$ is independent of the values of $Z_t(\nu)$ for which $\nu \in \Omega$ is distinguishable from ω at time τ.*

2.3 Abstract multistage single-level stochastic optimization models

In this paper, the main application of risk measures will be to situations in which uncertainty resolution and incremental decision-making are interleaved. That is, we suppose that at each time step t, we have control of an n_t-dimensional $\mathcal{F}_t$-measurable decision variable vector X_t that determines the costs Z_t. Specifically, we let $X_t \in \mathcal{Z}_t^{n_t}$, implying that X_t is an $\mathcal{F}_t$-measurable function $\Omega \to \mathbb{R}^{n_t}$, and suppose that we have functions $c_t : \mathbb{R}^{n_t} \to \mathbb{R}$ such that $Z_t = c_t(X_t)$, in the sense that $Z_t(\omega) = c_t\big(X_t(\omega)\big)$ for all $\omega \in \Omega$. While the X_t are formally "random variables" in the sense that they are functions of $\omega \in \Omega$, they are under the control of the decision maker, subject to $X_t \in \mathcal{Z}_t^{n_t}$ and hence the *nonanticipativity* constraint that each X_t is $\mathcal{F}_t$-measurable for each t. This constraint implies that X_t can only vary based on information available on time t: if two outcomes $\omega, \omega' \in \Omega$ are indistinguishable at time t, then $X_t(\omega') = X_t(\omega)$.

We also define a constraint structure linking successive decisions: for $t = 1, \ldots, T$, suppose that

$$G_t : \Omega \times \prod_{\tau=1}^{t-1} \mathbb{R}^{n_\tau} \rightrightarrows \mathbb{R}^{n_t}$$

is a closed-valued multifunction that is $\mathcal{F}_t$-measurable in the sense that for all possible fixed values $x_1 \in \mathbb{R}^{n_1}, \ldots, x_{t-1} \in \mathbb{R}^{n_{t-1}}$, we have $G_t(\omega, x_1, \ldots, x_{n-1}) = G_t(\omega', x_1, \ldots, x_{t-1})$ whenever $\omega, \omega' \in \Omega$ are indistinguishable at time t. We then propose the following very general constraint system:

$$X_t(\omega) \in G_t\big(\omega, X_1(\omega), \ldots, X_{t-1}(\omega)\big) \qquad \forall\, t = 1, \ldots, T,\ \omega \in \Omega.$$

Note that for $t = 1$, we simply have $G_1 : \omega \rightrightarrows \mathbb{R}^{n_1}$, and G_1 must be $\mathcal{F}_1$-measurable, that is, constant. Thus, the $t = 1$ constraints reduce to $X_1(\omega) \in \mathcal{X}_1$, where $\mathcal{X}_1 \subseteq \mathbb{R}^{n_1}$ is some fixed closed set. Furthermore, since X_1 must $\mathcal{F}_1$-measurable, it does not vary with ω, so the $t = 1$ constraint could simply be written $X_1 \in \mathcal{X}_1$. In general, the set of feasible choices for $X_t(\omega)$ will depend on ω in an $\mathcal{F}_t$-measurable way that can only depend on prior decisions in situations that could possibly lead to ω at time t.

We now formulate the following abstract, extremely general, and potentially infinite-dimensional stochastic programming problem, whose decision variables are $X_t(\omega)$, for $t = 1, \ldots, T$ and $\omega \in \Omega$:

$$\begin{array}{lll} \min & \rho_{1,T}\big(c_1(X_1), \ldots, c_T(X_T)\big) & \\ \text{ST} & X_t(\omega) \in G_t\big(\omega, X_1(\omega), \ldots, X_{t-1}(\omega)\big) & \forall\, t = 1, \ldots, T,\ \omega \in \Omega \\ & X_t \in \mathcal{Z}_t^{n_t} & \forall\, t = 1, \ldots, T. \end{array} \tag{4}$$

Letting $Z_t = c_t(X_t) \in \mathcal{Z}_t$, we could also denote the objective function as $\rho_{1,T}(Z_1, \ldots, Z_T)$. As is typical in stochastic programming, this model involves formulating a plan of action $X_t(\omega)$ for each point in time t and each eventuality ω that could possibly occur; the $\mathcal{F}_t$-measurability constraints $X_t \in \mathcal{Z}_t^{n_t}$ require these plans to be nonanticipative, meaning that actions at time t may be based only on information already available then, and not on knowledge of the future.

2.4 Recourse problems

Since the model (4) involves making a plan of action for every situation that can occur at each future decision point, it is natural to consider the planning problems that we would formulate if placed *de novo* in one of those situations (without advancing the final time horizon past T).

In the interest of clarity and simplicity, we henceforth assume that Ω is a finite set and without further loss of generality that each "atom" $\omega \in \Omega$ has $\mathrm{P}\big\{\{\omega\}\big\} > 0$. This simplifying assumption makes it possible to discuss an individual recourse problem in a particular system state at time stage τ, as opposed only to a functional system of recourse problems for each time stage. In the finite-Ω case, it is possible to identify a collection $\mathcal{E}_\tau \subset \mathcal{F}_\tau$ of *elementary events* in each $\mathcal{F}_\tau$, that is, events $E \in \mathcal{E}_\tau$ such that the only strict subset of E in $\mathcal{F}_\tau$ is $\emptyset$. Each $E \in \mathcal{E}_\tau$ then corresponds to a distinguishable scenario in stage τ. Note that any $\mathcal{F}_\tau$-measureable function, including $\rho_{\tau,T}(Z_\tau, \ldots, Z_T)$, is necessarily constant across each $E \in \mathcal{E}_\tau$; we will therefore write $\rho_{\tau,T}(Z_\tau, \ldots, Z_T)(E)$ for the common value of $\rho_{\tau,T}(Z_\tau, \ldots, Z_T)(\omega)$ for all $\omega \in E$. By Assumption 1, this value depends only on the values of $Z_t(\omega)$ for which $\omega \in E$. For each $t = \tau, \ldots, T$, we denote these values by $Z_t(E)$; formally, $Z_t(E)$ is the function $Z_t : \Omega \to \mathbb{R}$ with its domain restricted to E. We will use a similar restriction notation for any random variable.

Now consider the decision maker at time step τ, having taken the decisions $X_1, \ldots, X_{\tau-1}$ at the prior time steps, and knowing that elementary event $E \in \mathcal{E}_\tau$ has occurred. Now, the risk measure $\rho_{\tau,T}(\,\cdot\,, \ldots, \,\cdot\,)$ should in principle characterize our preferences at this point in time, and this function must be constant-valued across E. Therefore, we may try to optimize it with respect to the remaining decisions to be made, the only relevant portions of which, by Assumption 1, are constituted by $X_\tau(E), \ldots, X_T(E)$. This optimization problem, which we call the *recourse problem* in state E at time τ and denote by $R_\tau^E(X_1, \ldots, X_{\tau-1})$, is

$$\begin{array}{lll} \min & \rho_{\tau,T}\Big(c_\tau\big(X(E)\big), \ldots, c_T\big(X_T(E)\big)\Big)(E) & \\ \mathrm{ST} & X_t(\omega) \in G_t\big(\omega, X_1(\omega), \ldots, X_{t-1}(\omega)\big) & \forall\, t = \tau, \ldots, T,\ \omega \in E \\ & X_t(E) \in \mathcal{Z}_t^{n_t} \mid E & \forall\, t = \tau, \ldots, T, \end{array} \tag{5}$$

where $\mathcal{Z}_t^{n_t} \mid E$ denotes the space of functions in $Z_t^{n_t}$, restricted to the domain E. The decision variables of this model are $X_\tau(E), \ldots, X_T(E)$.

2.5 An intuitive choice of the dynamic risk measure system

We now consider how to choose the functions $\rho_{\tau,T}$ consistently, for $\tau = 1, \ldots, T-1$. To obtain a formulation equivalent to classical risk-neutral stochastic programming, we choose, for any $1 \le \tau < T$ and elementary event $E \in \mathcal{E}_\tau$,

$$\rho_{\tau,T}(Z_\tau, \ldots, Z_T)(E) = \mathbb{E}[Z_\tau + \cdots + Z_T \mid E]. \tag{6}$$

For $\tau = 1$, the only elementary event in $\mathcal{E}_1$ is Ω, so this formula reduces to $\rho_{1,T}(Z_1, \ldots, Z_T) = \mathbb{E}[Z_1 + \cdots + Z_T]$. It is natural to try to generalize this approach to other distribution-based rules for generating risk measures. Specifically, given any distribution-based rule ρ, we can for each elementary event $E \in \mathcal{E}_\tau$ at time τ evaluate the distribution of $Z_\tau + \cdots + Z_T$ conditional on the event E, and then apply ρ to this distribution to obtain a conditional certainty equivalent. For example, if we choose the MSD_γ rule, we then have

$$\rho_{\tau,T}(Z_\tau, \ldots, Z_T)(E) = \mathbb{E}\left[\sum_{t=\tau}^{T} Z_t \,\middle|\, E\right] + \gamma\, \mathbb{E}\left[\left[\sum_{t=\tau}^{T} Z_t - \mathbb{E}\left[\sum_{t=\tau}^{T} Z_t \,\middle|\, E\right]\right]_+ \,\middle|\, E\right]. \qquad (7)$$

For the AVaR_α rule we obtain

$$\rho_{\tau,T}(Z_\tau, \ldots, Z_T)(E) = \frac{1}{\alpha} \int_{1-\alpha}^{1} F^{-1}_{Z_\tau + \cdots + Z_T \,|\, E}(\nu)\, d\nu, \qquad (8)$$

where $F_{Z_\tau + \cdots + Z_T \,|\, E}$ denotes the distribution function of $Z_\tau + \cdots + Z_T$, conditioned on E.

This approach enforces what may be called *risk attitude consistency*: the decision maker's tolerance for risk at the time horizon T is consistently captured by applying the same distribution-based rule such as MSD_γ or AVaR_α to the anticipated distribution of total costs incurred by time T, with the distribution conditional upon the available information E known at the time of the first remaining decisions. Risk evaluations such as (7) or (8) are also necessarily law-invariant: if $Z_t, W_t \in \mathcal{Z}_t$, $t = \tau, \ldots, T$, are such that $Z_\tau + \cdots + Z_T$ and $W_\tau + \cdots + W_T$ have the same conditional distribution on E, then $Z_\tau, \ldots, Z_T$ and $W_\tau, \ldots, W_T$ should be assigned the same certainty equivalent within E, that is, $\rho_{\tau,T}(Z_\tau, \ldots, Z_T)(E) = \rho_{\tau,T}(W_\tau, \ldots, W_T)(E)$. In particular, if $Z_1 + \cdots + Z_T$ and $W_1 + \cdots + W_T$ have the same distribution — that is, the total costs incurred by time T have identical distributions — then $Z_1, \ldots, Z_T$ and $W_1, \ldots, W_T$ should have the same risk measure value. Note that while we have not explicitly included discount factors when combining costs possibly incurred at different times, such factors could easily be incorporated into the costs Z_t and W_t themselves.

2.6 Time consistency

Unfortunately, the apparently appealing approach to defining $\{\rho_{t,T}\}_{t=1}^T$ described above has a serious drawback if one uses standard single-level formulations of the main and recourse problems: except when using the respectively risk-neutral or maximally risk-averse rules $\mathbb{E}[\,\cdot\,]$ or ess sup, it runs afoul of considerations of *time consistency*. Ruszczyński (2010) defines time consistency of a system of risk measures $\{\rho_{t,T}\}_{t=1}^T$ as follows:

Definition 2 (Ruszczyński (2010)) *The system of dynamic risk measures* $\{\rho_{t,T}\}_{t=1}^T$ *is* time consistent *if for each* τ, θ *such that* $1 \le \tau < \theta \le T$, *all* $Z_\tau \in \mathcal{Z}_\tau, \ldots, Z_{\theta-1} \in \mathcal{Z}_{\theta-1}$, *and all* $Z_\theta, W_\theta \in \mathcal{Z}_\theta, \ldots, Z_T, W_T \in \mathcal{Z}_T$ *such that* $\rho_{\theta,T}(Z_\theta, \ldots, Z_T) \le \rho_{\theta,T}(W_\theta, \ldots, W_T)$, *we have*

$$\rho_{\tau,T}(Z_\tau, \ldots, Z_{\theta-1}, Z_\theta, \ldots, Z_T) \le \rho_{\tau,T}(Z_\tau, \ldots, Z_{\theta-1}, W_\theta, \ldots, W_T),$$

where all the inequalities are interpreted pointwise.

This definition says that if the $Z_\theta, \ldots, Z_T$ is preferable to $W_\theta, \ldots, W_T$ in all outcomes, this preference should remain after prepending some identical sequence of costs at earlier stages. This requirement is not only intuitive, but it also has the consequence of guaranteeing that a decision maker facing a recourse problem at stage τ has no incentive to make decisions that would invalidate the optimality of earlier decisions. To explain and clarify this claim, first suppose that $(X_1^*, \ldots, X_T^*)$ is an optimal solution to the problem (4), and consider some elementary event $E \in \mathcal{E}_\tau$ at stage $\tau \in \{2, \ldots, T\}$, along with the associated recourse problem $R_\tau^E(X_1^*, \ldots, X_{\tau-1}^*)$ as constructed in (5). Let $\big(\overline{X}_\tau(E), \ldots, \overline{X}_T(E)\big)$ be any optimal solution of this recourse problem. Suppose we construct a new solution $(\tilde{X}_1, \ldots, \tilde{X}_T)$ to the original problem as follows:

$$\tilde{X}_t(\omega) = \begin{cases} \overline{X}_t(\omega), & \text{if } t \geq \tau \text{ and } \omega \in E, \\ X_t^*(\omega), & \text{otherwise.} \end{cases} \tag{9}$$

In other words, we substitute the recourse-problem solution $\big(\overline{X}_\tau(E), \ldots, \overline{X}_T(E)\big)$ for the corresponding portions of the original solution. It is easily seen that $(\tilde{X}_1, \ldots, \tilde{X}_T)$ is feasible for (4). Letting $Z_t^* = c_t(X_t^*)$ and $\tilde{Z}_t = c_t(\tilde{X}_t)$ for $t = 1, \ldots, T$, we observe the following:

- $\rho_{\tau,T}(\tilde{Z}_\tau, \ldots, \tilde{Z}_T)(\omega) = \rho_{\tau,T}\Big(c_\tau\big(\overline{X}_\tau(E), \ldots, c_T(\overline{X}_T(E))\big)\Big)(\omega) \leq \rho_{\tau,T}(Z_\tau^*, \ldots, Z_T^*)(\omega)$ for $\omega \in E$ by the optimality of $\big(\overline{X}_\tau(E), \ldots, \overline{X}_T(E)\big)$ for the recourse problem.

- For $\omega \notin E$, Assumption 1 implies that $\rho_{\tau,T}(\tilde{Z}_\tau, \ldots, \tilde{Z}_T)(\omega) = \rho_{\tau,T}(Z_\tau^*, \ldots, Z_T^*)(\omega)$, since $\tilde{Z}_t(\omega)$ and $\mathcal{Z}_t^*(\omega)$ are identical for $\tau \leq t \leq T$.

- Combining the last two observations, we have $\rho_{\tau,T}(\tilde{Z}_\tau, \ldots, \tilde{Z}_T) \leq \rho_{\tau,T}(Z_\tau^*, \ldots, Z_T^*)$ pointwise.

- We have $\tilde{Z}_t(\omega) = Z_t^*(\omega)$ for all $1 \leq t < \tau$ and $\omega \in \Omega$ by construction.

- Because of the assumption that the system $\{\rho_{t,T}\}_{t=1}^T$ is time-consistent, we conclude that $\rho_{1,T}(\tilde{Z}_1, \ldots, \tilde{Z}_T) \leq \rho_{1,T}(Z_1^*, \ldots, \tilde{Z}_T^*)$.

- Since $(\tilde{X}_1, \cdots, \tilde{X}_T)$ is feasible for (4) and $X_1^*, \ldots, X_T^*$ is optimal for the same problem, $(\tilde{X}_1, \cdots, \tilde{X}_T)$ is also optimal for (4).

This chain of reasoning means that the formulation (4) and its recourse problems (5) have a property we call *strong time consistency*: any optimal solution we may select to a recourse problem can be combined with the preceding decisions as described in (9) without invalidating their optimality. Thus, if the system $\{\rho_{t,T}\}_{t=1}^T$ truly models the actual objective functions we wish to use at each point in time τ and information state E, there is no incentive to make decisions at later times that are incompatible with an optimal initial plan. This kind of model time consistency is essentially indispensible in multistage optimization settings. Without it, the validity of the first-stage optimal decisions X_1 rests on the assumption that we may take subsequent actions that will appear suboptimal when the time comes to implement them; such an assumption is generally not credible. Thus, if we use single-level optimization formulations such as (4) and (5), we should use a system of risk measures $\{\rho_{t,T}\}_{t=1}^{T}$ that is time-consistent.

Unfortunately, intuitive systems of risk measures such as those specified by (7) or (8) are not time-consistent, except in the case of the expected-value and worst-outcome (ess sup) risk-measure classes. This inconvenient fact can be seen as the consequence of two prior results: first, Ruszczyński (2010, p. 241) shows that coherent time-consistent risk measure systems necessarily have the nested structure

$$\rho_{\tau,T}(Z_\tau,\ldots,Z_T) = Z_\tau + \rho_\tau\Big(Z_{\tau+1} + \rho_{\tau+1}\big(Z_{\tau+2} + \cdots \rho_{T-1}(Z_T)\cdots\big)\Big), \tag{10}$$

where $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$, $t = 1,\ldots,T-1$, are *one-stage* risk measures satisfying axioms analogous to those for classic risk measures. Furthermore, Shapiro (2012) shows that dynamic risk measures of the form (10) cannot be law-invariant unless all the one-step risk measures are conditional expectations, or they are all worst-outcome operators. An immediate consequence of combining these two results is that the only possible law-invariant time-consistent risk measures are the worst-outcome measure and the classical expected value measures; Kupper and Schachermayer (2009) show a similar result. Thus, non-extreme law-invariant coherent risk measure systems such as those given by (7) and (8) cannot be time-consistent.

Thus, in systems of models like (4) and (5), time consistency and law invariance cannot coexist, unless we revert to classical expected-value stochastic programming or are maximally risk-averse — note that for the expected-value or worse-outcome operators, the expression (10) with $\tau = 1$ collapses to the the expected value or worst outcome of $Z_1+\cdots+Z_T$, respectively. Something similar also occurs in each possible recourse problem.

Furthermore, the result of Shapiro (2012) shows not only that using straighforward conditional analogues of standard risk measures like (7) or (8) as recourse problem objective functions would violate time consistency, but also asserts that time-consistent risk dynamic risk measures of the form (10), other than expected value or worst outcome, can never be simplified to more manageable, appealing expressions. For example, if one uses a variation of the MSD_γ or AVaR_α risk measure for each one-step risk measure ρ_t, the composed risk measure (10) cannot be simplified to any law-invariant expression applied to the distribution of $Z_1 + \cdots + Z_T$, and certainly not to something like $\mathrm{MSD}_\gamma(Z_1 + \cdots + Z_T)$ or $\mathrm{AVaR}_\alpha(Z_1 + \cdots + Z_T)$.

This proven impossibility of finding simpler expressions for dynamic risk measures of the form (10) presents a serious obstacle to using coherent risk measures to introduce risk aversion into multistage decision problems within the single-level optimization framework described by (4) and (5). Either we must formulate a time-inconsistent system of models, which is essentially indefensible, or we must use complicated, somewhat opaque objective functions whose form cannot be simplified, and which will thus have difficulty gaining acceptance with decision makers. Furthermore, we know that these objective functions must in general violate the appealing property of law invariance. Unless there is a way to resolve this incompatibility, multistage models using coherent risk measures are unlikely to be more than academic exercises.

3 Multilevel Optimization Modeling for Multistage Decision-Making

The way to resolve the incompatibility described in the previous section is to move beyond the traditional structure of single-level optimization in (4) and (5). This single-level optimization framework is familiar because it directly generalizes the standard structure of classical multistage expected-value stochastic programming, but it is also tailored to the expected-value approach's inherently time-consistent nature. Using it with more general dynamic risk measures requires the specific objective-function structure embodied in (10), or time inconsistency results.

This paper suggests a different approach: a different, multilevel constraint and recourse structure that makes risk measure time consistency unnecessary to obtain a time-consistent model.

We develop this system of models as follows: first, we select any system of dynamic risk measures $\{\rho_{t,T}\}_{t=1}^{T}$ that accurately models our risk preferences at each time step t. This system need not be time-consistent, and in particular we may choose to obtain the system by applying a distribution-based rule such as MSD_γ or AVaR_α to the conditional cost-to-go, as in (7) or (8).

To construct this system of models, we proceed recursively from time step $T-1$; the decision model at time step T is deterministic, so it is sufficient to start from step $T-1$. Consider some elementary event $E \in \mathcal{E}_{T-1}$. If the decision maker finds themselves in the state E at time step $T-1$, with $X_1, \ldots, X_{T-2}$ already determined, they would naturally wish to solve the recourse problem

$$\begin{array}{lll} \min & \rho_{T-1,T}\Big(c_{T-1}\big(X_{T-1}(E)\big), c_T\big(X_T(E)\big)\Big)(E) & \\ \mathrm{ST} & X_{T-1}(\omega) \in G_{T-1}\big(X_1(\omega), \ldots, X_{T-2}(\omega)\big) & \forall\, \omega \in E \\ & X_T(\omega) \in G_T\big(X_1(\omega), \ldots, X_{T-1}(\omega)\big) & \forall\, \omega \in E \\ & X_{T-1}(E) \in \mathcal{Z}_{T-1} \mid E. & \end{array} \tag{11}$$

Here, the last constraint just means that $X_{T-1}(\omega)$ should be identical for all $\omega \in E$. We will denote this problem by $\overline{R}_T^E(X_1, \ldots, X_{T-2})$ use the notation $\mathcal{X}^*_{T-1,E}(X_1, \ldots, X_{T-2})$ to denote its set of optimal solutions.

Next, consider time step $T-2$ and an elementary event $E \in \mathcal{E}_{T-2}$, assuming that $X_1, \ldots, X_{T-3}$ have already been selected. In this situation, the decision maker would like to minimize the objective function

$$\rho_{T-2,T}\Big(c_{T-2}\big(X_{T-2}(E)\big), c_{T-1}\big(X_{T-1}(E)\big), c_T\big(X_T(E)\big)\Big)(E),$$

subject to constraints on $X_{T-2}(E)$ given by

$$\begin{array}{ll} X_{T-2}(\omega) \in G_{T-2}\big(X_1(\omega), \ldots, X_{T-3}(\omega)\big) & \forall\, \omega \in E, \\ X_{T-2}(E) \in \mathcal{Z}_{T-2} \mid E. & \end{array}$$

For $X_{T-1}(E)$ and $X_T(E)$, now consider each elementary event $D \in \mathcal{E}_{T-1}$ such that $D \subseteq E$, that is, each event at stage $T-1$ that can follow from E. For each such event D, we impose the constraint that we can only select values of $\big(X_{T-1}(D), X_T(D)\big)$ that are optimal for the corresponding recourse problem $\overline{R}_T^D(X_1, \ldots, X_{T-2})$. That is, we can only select feasible solutions from which we would have no incentive to depart in the future, according to the model of preferences given by $\rho_{T-1,T}$. Thus, we obtain the model

$$
\begin{array}{lll}
\min & \rho_{T-2,T}\Big(c_{T-2}\big(X_{T-2}(E)\big), c_{T-1}\big(X_{T-1}(E)\big), c_T\big(X_T(E)\big)\Big)(E) & \\
\text{ST} & X_{T-2}(\omega) \in G_{T-2}\big(X_1(\omega), \ldots, X_{T-3}(\omega)\big) & \forall\, \omega \in E \\
& X_{T-2}(E) \in \mathcal{Z}_{T-2} \mid E & \\
& \big(X_{T-1}(D), X_T(D)\big) \in \mathcal{X}^*_{T-1,D}(X_1, \ldots, X_{T-2}) & \forall\, D \in \mathcal{E}_{T-1} : D \subseteq E.
\end{array}
\tag{12}
$$

We denote this model by the notation $\overline{R}_{T-2}^E(X_1, \ldots, X_{T-3})$, and its set of optimal solutions by $\mathcal{X}^*_{T-2,E}(X_1, \ldots, X_{T-3})$. At this point, we inductively define a recourse model $\overline{R}_\tau^E(X_1, \ldots, X_{\tau-1})$ for each time stage τ and elementary event $E \in \mathcal{E}_\tau$ by

$$
\begin{array}{lll}
\min & \rho_{\tau,T}\Big(c_\tau\big(X_\tau(E)\big), \ldots, c_T\big(X_T(E)\big)\Big)(E) & \\
\text{ST} & X_\tau(\omega) \in G_\tau\big(X_1(\omega), \ldots, X_{\tau-1}(\omega)\big) & \forall\, \omega \in E \\
& X_\tau(E) \in \mathcal{Z}_\tau \mid E & \\
& \big(X_{\tau+1}(D), \ldots, X_T(D)\big) \in \mathcal{X}^*_{\tau+1,D}(X_1, \ldots, X_\tau) & \forall\, D \in \mathcal{E}_{\tau+1} : D \subseteq E,
\end{array}
\tag{13}
$$

and let $\mathcal{X}^*_{\tau,E}(X_1, \ldots, X_{\tau-1})$ denote this problem's set of optimal solutions. Taking this process back to $\tau = 1$ and simplifying, we obtain the first-stage problem

$$
\begin{array}{lll}
\min & \rho_{1,T}\big(c_1(X_1), \ldots, c_T(X_T)\big) & \\
\text{ST} & X_1 \in \mathcal{X}_1 & \\
& \big(X_2(D), \ldots, X_T(D)\big) \in \mathcal{X}^*_{2,D}(X_1) & \forall\, D \in \mathcal{E}_2.
\end{array}
\tag{14}
$$

Using the notation already defined, this problem may be referred to as $\overline{R}_1^\Omega(\,)$.

In this modeling approach, the constraints impose a form of time consistency, regardless of the structure of the system of preferences $\{\rho_{t,T}\}_{t=1}^T$. By construction, the values chosen for $X_2, \ldots, X_T$ by optimally solving the problem (14) are necessarily optimal for any recourse problem that might be encountered. This form of time consistency imposed may be called *weak time-consistency*, since we only guarantee that the solution planned at prior stages is among the optimal ones for each recourse problem. It is not guaranteed in the case of nonunique optima that *any* optimal solution of the recourse problem could be substituted into the original plan without affecting its optimality, as shown for time-consistent risk-measure systems in Section 2.6. Nevertheless, if $\{\rho_{t,T}\}_{t=1}^T$ is an accurate model of the decision maker's evolving risk preferences, then the system of models $\{\overline{R}_t^E \mid t = 1, \ldots, T,\ E \in \mathcal{E}_t\}$ given by (14)-(13)-(11) is time-consistent in the sense that there should never be any positive incentive to select a recourse solution that would invalidate the optimality of earlier decisions. Of course, this property comes at the expense of a much more complex constraint structure involving multilevel optimization at all stages before $T-1$; by contrast, all the problems in the framework (4)-(5) involve only a single level of optimization.

Next consider what would occur within this framework if we select the classical expected-value system of risk-neutral risk measures (6). For any $E \in \mathcal{E}_{T-2}$, consider the recourse problem $\overline{R}^{E}_{T-2}(X_1, \ldots, X_{T-3})$ given by (12). Suppose next that we were to substitute for each constraint $\big(X_{T-1}(D), X_T(D)\big) \in \mathcal{X}^*_{T-1,D}(X_1, \ldots, X_{T-2})$ the simple feasibility conditions of the following recourse model $\overline{R}^{D}_{T-1}(X_1, \ldots, X_{T-2})$ given by (11), namely

$$\begin{aligned} X_{T-1}(\omega) &\in G_{T-1}\big(X_1(\omega), \ldots, X_{T-2}(\omega)\big) && \forall\, \omega \in D \\ X_T(\omega) &\in G_T\big(X_1(\omega), \ldots, X_{T-1}(\omega)\big) && \forall\, \omega \in D \\ X_{T-1}(D) &\in \mathcal{Z}_{T-1} \mid D. \end{aligned}$$

We thus obtain the model

$$\begin{array}{lll} \min & \mathbb{E}[Z_{T-2} + Z_{T-1} + Z_T \mid E] & \\ \text{ST} & X_t(\omega) \in G_t\big(\omega, X_1(\omega), \ldots, X_{t-1}(\omega)\big) & \forall\, t = T-2, T-1, T,\ \omega \in E \\ & X_t(E) \in \mathcal{Z}_t^{n_t} \mid E & t = T-2, T-1. \end{array} \tag{15}$$

Due to our standing assumption on $(\Omega, \mathcal{F}, \mathrm{P})$, which guarantees that all nonempty events have positive probability, and the "nesting" property of conditional expected values, which in this case reduces to

$$\mathbb{E}[Z_{T-2} + Z_{T-1} + Z_T \mid E] = \mathbb{E}[Z_{T-2} \mid E] \ + \sum_{\substack{D \in \mathcal{E}_{T-1}: \\ D \subseteq E}} \mathrm{P}\{D \mid E\}\, \mathbb{E}[Z_{T-1} + Z_T \mid D],$$

implies that any optimal solution obtained for (15) would still be optimal for the original formulation (11): the substitute model (15) cannot select values of X_{T-1}, X_T that are not composed entirely of optimal solutions to the recourse problems $\overline{R}^{D}_{T-1}(X_1, \ldots, X_{T-2})$, $D \in \mathcal{F}_{T-1}$, $D \subseteq E$, since doing so would clearly make the solution of (15) suboptimal.

We can now repeat this logic recursively back to time step 1, establishing that in the case of (6), the proposed multilevel optimization framework (14)-(13)-(11) is equivalent to the classical framework of risk-neutral stochastic programming. Now, in the case of (6), the simpler framework (4)-(5) also reduces to the same standard risk-neutral framework. When other choices of $\{\rho_{t,T}\}_{t=1}^{T}$ are made, however, the two frameworks behave differently. The simple framework (4)-(5), which tries to stay as close as possible to the familiar structure of recourse problems from classical stochastic programming, leads to indefensible, time-inconsistent models unless we use a time-consistent system of risk measures. This requirement, in turn, requires that the risk measures have the nested form (10), which resembles one way of representing the expected-value function, and which we have seen has serious drawbacks in cases other than expected value. The framework (14)-(13)-(11) is much more robust from the modeling point of view: it can accept essentially any system of risk measures without creating time inconsistency in the model. Risk-measure choice thus becomes a full part of the modeling process, where it should ideally belong. In particular, one is free to select rules such as (7) or (8) that ascribe a consistent law-invariant risk profile to the total costs incurred at the time horizon T. ¿From a modeling perspective, (14)-(13)-(11) is thus the preferred way to generalize classical stochastic programming to potentially risk-averse situations.

Despite its apparent complexity, the framework (14)-(13)-(11) — or generalization thereof — is really the natural starting point for multistage optimization models, because, at each stage, it includes a model of the decision maker's future preferences and future optimal courses of action. It explicitly models a set of decisions made over time, with each decision taking into account future actions. By choosing a system of risk measures based on conditional expected value, it is in fact compatible with standard risk-neutral multistage models, but in that case it happens to reduce to a simpler form that only requires single-level optimization, and is thus also special cases of the less robust framework (4)-(5).

Finally, it may seem natural to think that if one adopts any time-consistent, nested-structure system of risk measures of the form (10), then the multilevel optimization framework (14)-(13)-(11) would reduce to the simpler framework (4)-(5), but this is not exactly the case. The reason is that it is sometimes possible for optimal solutions to (4)-(5) to contain sub-solutions that are suboptimal for certain recourse problems, something which cannot happen in the framework (14)-(13)-(11). Time consistency guarantees that any optimal recourse solution could be substituted into the top-level solution without affecting its optimality, as shown in Section 2.6, but, in general, some optimal solutions may contain portions which are "ignored" in the objective function and may thus contain anything feasible. To illustrate this phenomenon, suppose we select each one-step risk measure ρ_t in (10) to be of the form AVaR_α, and that we have $T = 3$ and $\Omega = \{\omega_1, \omega_2, \omega_3\}$. Suppose further that $\mathcal{E}_2 = \big\{\{\omega_1, \omega_2\}, \{\omega_3\}\big\}$, $\mathrm{P}\{\omega_3\} = \alpha$, and that the minimal cost in outcome ω_3 is some large value K that is higher than the objective value of all feasible solutions in outcomes ω_1 or ω_2. Then the optimal value of the entire problem will be K, no matter what values are chosen for $X_2\big(\{\omega_1, \omega_2\}\big)$ and $X_3\big(\{\omega_1, \omega_2\}\big)$; any feasible values can be selected for those variables in an optimal solution to the top-level problem (4). Figure 1 illustrates this construction.

4 A Simple Class of Problems

While it may be the preferred approach from a modeling perspective, the framework proposed in Section 3 has the obvious drawback of leading to multilevel optimization problems, which are in general far more difficult than single-level problems. We now consider this increased difficulty from the perspective of computational complexity theory. We will show that even in the most straightforward imaginable special cases, the modeling framework described in Section 3 leads to $\mathcal{NP}$-hard problem classes. The subclass of models we consider has the following special form:

- Ω is a finite set (as already assumed).
- $T = 3$; there are three stages.

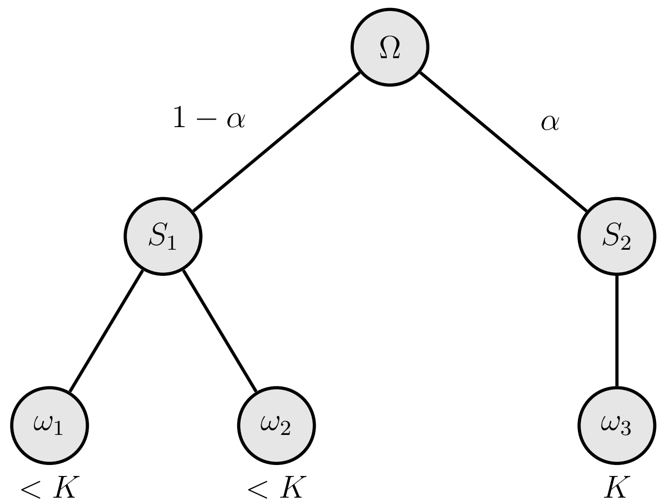

Figure 1: Scenario tree illustrating non-equivalence of the problem frameworks (4)-(5) and (14)-(13)-(11) for a general time-consistent system of risk measures. S_1 and S_2 denote the elementary events in $\mathcal{E}_2$.

- We denote the members of $\mathcal{E}_2$ by S, standing for "scenario". Each scenario denotes a distinguishable state of the system at time 2.
- The cost functions $c_t(X_t)$ are all linear; hence, we write them as $c_t^\top X_t$, for $t = 1, 2, 3$.
- All the constraints are linear:
  - The first-stage constraints may be expressed as $A_{11}X_1 \le b_1$; here, A_{11} and b_1 are an $\mathcal{F}_1$-measurable (that is, constant) matrix and vector, respectively.
  - The second-stage constraints are $A_{21}X_1 + A_{22}X_2 \le b_2$, interpreted pointwise, where the random matrix A_{21}, random matrix A_{22}, and random vector b_2 are all $\mathcal{F}_2$-measurable.
  - The third-stage constraints are $A_{31}X_1 + A_{32}X_2 + A_{33}X_3 \le b_3$, again interpreted pointwise, where A_{31}, A_{32}, and A_{33} are random matrices over $(\Omega, \mathcal{F}, \mathrm{P})$, and b_3 is a random vector over $(\Omega, \mathcal{F}, \mathrm{P})$.
- We select some law-invariant risk-measure family ρ and set

$$\rho_{1,3}(Z_1, Z_2, Z_3) = \rho(Z_1 + Z_2 + Z_3) \quad \rho_{2,3}(Z_2, Z_3)(S) = \rho(Z_2 + Z_3 \mid S) \text{ for each } S \in \mathcal{E}_2,$$

  that is, $\rho_{2,3}$ is obtained by applying the law-invariant risk measure family ρ to the distribution of $Z_2 + Z_3$ conditional on the scenario S which occurs at stage 2.

Since $T = 3$, the framework (14)-(13)-(11) reduces to a bilevel optimization problem. The leader problem consists of (14), and the follower problem consists of $|\mathcal{E}_2|$ problems of the form (11). Specifically, we obtain the leader problem

$$\begin{array}{lll} \min & \rho(c_1^\top X_1 + c_2^\top X_2 + c_2^\top X_2) & \\ \mathrm{ST} & A_{11}X_1(\omega) \le b_1 & \forall\, \omega \in \Omega \\ & X_1(\omega) = X_1(\omega') & \forall\, \omega, \omega' \in \Omega \\ & \big(X_2(S), X_3(S)\big) \in \mathcal{X}^*_{2,S}(X_1) & \forall\, S \in \mathcal{E}_2, \end{array} \tag{16}$$

where $\mathcal{X}^*_{2,S}(X_1)$ denotes the set of optimal solutions to

$$\begin{array}{lll} \min & \rho(c_2^\top X_2 + c_2^\top X_2 \mid S) & \\ \mathrm{ST} & A_{21}(S)X_1(\omega) + A_{22}(S)X_2(\omega) \le b_2(S) & \forall\, \omega \in S \\ & A_{31}(\omega)X_1(\omega) + A_{32}(\omega)X_2(\omega) + A_{33}(\omega)X_3(\omega) \le b_3(\omega) & \forall\, \omega \in S \\ & X_2(\omega) = X_2(\omega') & \forall\, \omega, \omega' \in S; \end{array} \tag{17}$$

here, $A_{21}(S)$ denotes the common value of $A_{21}(\omega)$ for all $\omega \in S$, with $A_{22}(S)$ and $b_2(S)$ defined similarly. Formally, we define a parameterized class of problems as follows:

Problem Class BLRP(ρ) **: Bilevel Risk Programming**

Parameter:	A law-invariant coherent risk-measure rule ρ.
Input:	All expressed over the rational numbers $\mathbb{Q}$: • A finite probability space Ω, along with a partition $\mathcal{E}_2$ and positive probabilities $\mathrm{P}\{\omega\}$ for each $\omega \in \Omega$ • Vectors c_1 and b_1, and a matrix A_{11} • For each $S \in \mathcal{E}_2$, vectors $c_2(S)$ and $b_2(S)$, and matrices $A_{21}(S)$, $A_{22}(S)$ • For each $\omega \in \Omega$, vectors $c_3(\omega)$ and $b_3(\omega)$, and matrices $A_{31}(\omega)$, $A_{32}(\omega)$, $A_{33}(\omega)$.
Output:	Any optimal solution (X_1, X_2, X_3) to the problem (16)-(17).

When the coherent risk measure parameter ρ itself has a parameter, as in the case of MSD_γ and AVaR_α, we use a set-valued parameter to denote the version of BLRP(ρ) in which this risk-measure parameter, restricted to the rationals, is encoded as part of the problem input. Thus, BLRP($\mathrm{MSD}_{(0,1]}$) denotes the class of all BLRP(MSD_γ) problems, with $\gamma \in (0,1] \cap \mathbb{Q}$ appended to the problem input, and BLRP($\mathrm{AVaR}_{(0,1)}$) denotes the class of all BLRP(AVaR_α) problems, with $\alpha \in (0,1) \cap \mathbb{Q}$ appended to the problem input.

We will analyze the computational complexity of the problem class BLRP(ρ) in the cases of $\rho = \mathrm{MSD}_\gamma$, $\gamma \in (0,1)$, and $\rho = \mathrm{AVaR}_\alpha$, $\alpha \in (0,1)$. To do so, we first revisit the complexity theory of bilevel linear programming.

5 Bilevel Linear Programming Complexity Revisited

Even the simplest form of bilevel programming, bilevel linear programming, has long been known to be $\mathcal{NP}$-hard (Jeroslow 1985, Ben-Ayed and Blair 1990, Bard 1991, Hansen et al. 1992). This problem class takes the form

$$\begin{array}{lll} \min & f_1^\top y_1 + f_2^\top y_2 & \\ \text{ST} & y_2 \in \ \text{Arg min} & \tilde{f}_2^\top y_2 \\ & \qquad\ \ \text{ST} & B_1 y_1 + B_2 y_2 \le r. \end{array} \tag{18}$$

Even for risk measures ρ that can be expressed in a linear-programming form, such as MSD_γ and AVaR_α, it is not immediately clear whether the $\text{BLRP}(\rho)$ problem class can model a completely general bilevel linear program; in particular, the objective functions of the leader and follower in $\text{BLRP}(\rho)$ appear very strongly correlated. The known $\mathcal{NP}$-hardness proofs for (18) employ reductions from various combinatorial problems to the special case $\tilde{f}_2 = -f_2$, a problem subclass that may be called *oppositional programming*. We begin by focusing on this special case, with the additional restriction that the y_2 must lie in a bounded set:

Problem Class BOLP **: Bounded Oppositional Linear Programming**

Input: Vectors $f_1 \in \mathbb{Q}^{n_1}$, $f_2 \in \mathbb{Q}^{n_2}$, and $r \in \mathbb{Q}^m$, matrices $B_1 \in \mathbb{Q}^{m \times n_1}$ and $B_2 \in \mathbb{Q}^{m \times n_2}$, and $\zeta \in \mathbb{Q}_+$.

Output: Any optimal solution (y_1, y_2) of the problem

$$\begin{array}{lll} \min & f_1^\top y_1 + f_2^\top y_2 & \\ \text{ST} & y_2 \in \ \text{Arg min} & -f_2^\top y_2 \\ & \qquad\ \ \text{ST} & B_1 y_1 + B_2 y_2 \le r \\ & & \|y_2\|_\infty \le \zeta. \end{array}$$

In principle, the $\mathcal{NP}$-hardness of BOLP may be inferred by remarking that the existing proofs of the $\mathcal{NP}$-hardness of general bilevel linear programming all reduce various $\mathcal{NP}$-hard or $\mathcal{NP}$-complete combinatorial problems to the special case $\tilde{f}_2 = -f_2$, with all decision variables bounded. For completeness, we give a new proof that is similar in basic spirit to Hansen et al. (1992), but is simpler and involves reduction from a less complicated decision problem, although it does not demonstrate *strong* $\mathcal{NP}$-hardness as in Hansen et al. (1992).

Proposition 3 BOLP *is* $\mathcal{NP}$*-hard.*

Proof. We proceed by reduction from the number partition problem (NPP), one of the classical $\mathcal{NP}$-complete decision problems (Garey and Johnson 1979):

Problem Class NPP **: Number Partition**

Input:	$a_1, \ldots, a_n \in \mathbb{Z}$.
Output:	"Yes" if there exists $J \subseteq \{1, \ldots, n\}$ such that $\sum_{i \in J} a_i = \frac{1}{2} \sum_{i=1}^{n} a_i$, and otherwise "no".

Given an instance $a_1, \ldots, a_n \in \mathbb{Z}$ of NPP, consider the following bilevel program with leader variables $u_0, u_1, \ldots, u_n$ and follower variables $w = (w_1, \ldots, w_n)$:

$$\begin{array}{ll} \min & u_0 + \sum_{i=1}^{n} w_i \\ \text{ST} & u_0 \geq \sum_{i=1}^{n} a_i u_i - \frac{1}{2} \sum_{i=1}^{n} a_i \\ & u_0 \geq \frac{1}{2} \sum_{i=1}^{n} a_i - \sum_{i=1}^{n} a_i u_i \\ & 0 \leq u_1, \ldots, u_n \leq 1 \\ & w \in \begin{array}[t]{lll} \operatorname{Arg\,min} & -\sum_{i=1}^{n} w_i & \\ \text{ST} & w_i \leq u_i & i = 1, \ldots, n \\ & w_i \leq 1 - u_i & i = 1, \ldots, n. \end{array} \end{array} \tag{19}$$

We claim that (19) has an optimal value of zero if and only if the answer to the partition instance is "yes". To prove this claim, we note that the optimality of w for the follower program is equivalent to $w_i = \min\{u_i, 1 - u_i\}$, $i = 1, \ldots, n$, and the optimal value of u_0 is $\left|\sum_{i=1}^{n} a_i u_i - \frac{1}{2} \sum_{i=1}^{n} a_i\right|$, so (19) is equivalent to the (nonconvex) single-level optimization problem

$$\begin{array}{ll} \min & \left|\sum_{i=1}^{n} a_i u_i - \frac{1}{2} \sum_{i=1}^{n} a_i\right| + \sum_{i=1}^{n} \min\{u_i, 1 - u_i\} \\ \text{ST} & 0 \leq u_1, \ldots, u_n \leq 1. \end{array} \tag{20}$$

Both terms in the objective function of this problem are nonnegative, the first being zero whenever $\sum_{i=1}^{n} a_i u_i = \frac{1}{2} \sum_{i=1}^{n} a_i$ and the second being zero whenever $u = (u_1, \ldots, u_n)$ is a binary vector. Thus, if answer to the partition problem is "yes", then there must exist $u \in \{0, 1\}^n$ making the objective of (20) zero, which must be optimal. On the other hand, if the answer to the partition problem is "no", then all choices of $u \in [0, 1]^n$ make the objective of (20) positive, either because $\sum_{i=1}^{n} a_i u_i \neq \frac{1}{2} \sum_{i=1}^{n} a_i$ or because at least one of $u_1, \ldots, u_n$ is fractional. Since (20) involves the optimization of a continuous function over a compact set, it must acheive its minimum, and hence the minimum value must be positive. This establishes the claim; it then remains only to observe that problem (19) can be put into the BOLP form by appropriate choices of f_1, f_2, B_1, B_2, and r, with $\zeta = 1$, and that number of bits needed to encode such an instance of BOLP is polynomially bounded in the number of bits needed to encode $a_1, \ldots, a_n$. $\square$

Direct reduction of BOLP or classic bilevel linear programming to problem classes of the form BRLP(ρ) appears to be an intricate task. We instead perform a two-stage reduction, first considering another class of restricted bilevel linear problems that we show is also $\mathcal{NP}$-hard by reduction from BOLP. In this class of problems, we break the follower variables into two blocks, with the only difference between the leader and follower objectives being that the coefficients for one of the two blocks are scaled by a rational parameter $\beta \neq 1$.

Problem Class BBSBLP(β) **:**
Bounded Block-Scaled Bilevel Linear Programming

Parameter: $\beta \in \mathbb{Q}\backslash\{1\}$.

Input: Vectors $g_1 \in \mathbb{Q}^{n_1}$, $g_2 \in \mathbb{Q}^{n_2}$, $g_3 \in \mathbb{Q}^{n_3}$, and $t \in \mathbb{Q}^m$, and matrices $C_1 \in \mathbb{Q}^{m\times n_1}$, $C_2 \in \mathbb{Q}^{m\times n_2}$, and $C_3 \in \mathbb{Q}^{m\times n_3}$ and $\eta_2, \eta_3 \in \mathbb{Q}_+$.

Output: Any optimal solution (y_1, y_2, y_3) of the problem

$$\begin{array}{lll} \min & g_1^\top x_1 + g_2^\top x_2 + g_3^\top x_3 & \\ \text{ST} & (x_2, x_3) \in \ \operatorname{Arg\,min} & g_2^\top x_2 + \beta g_3^\top x_3 \\ & \qquad\qquad\quad \text{ST} & C_1x_1 + C_2x_2 + C_3x_3 \leq t \\ & & \|x_2\|_\infty \leq \eta_2 \\ & & \|x_3\|_\infty \leq \eta_3. \end{array}$$

We now show that although the leader and follower objectives of BBSBLP(β) may appear very similar — if $\beta > 0$, in particular, then all corresponding leader and follower objective coefficients have the same sign — the possibly small difference between the two objectives is enough to make the problem class $\mathcal{NP}$-hard. The technique used is inspired by an analysis of Marcotte and Savard (1991) showing that solutions of bilevel programs need not be Pareto-optimal for the leader and follower objectives as long as they are not colinear.

Proposition 4 *For any rational $\beta \neq 1$, the problem class* BBSBLP(β) *is $\mathcal{NP}$-hard.*

Proof. We proceed by reduction from BOLP. Consider any instance $(f_1, f_2, r, B_1, B_2, \zeta)$ of BOLP. We create a corresponding instance of BBSBLP(β) as follows: first, we set

$$g_1 = f_1 \qquad\qquad g_2 = \left(1 - \tfrac{2}{1-\beta}\right) f_2 \qquad\qquad g_3 = f_2. \tag{21}$$

We then set C_1, C_2, C_3, and t to be equivalent to the constraints

$$B_1x_1 + B_2x_2 \leq r \qquad\qquad x_3 = \left(\tfrac{2}{1-\beta}\right)x_2. \tag{22}$$

Specifically, we may induce such an equivalence by setting

$$C_1 = \begin{bmatrix} B_1 \\ 0 \\ 0 \end{bmatrix} \qquad C_2 = \begin{bmatrix} B_2 \\ -\left(\tfrac{2}{1-\beta}\right)I \\ \left(\tfrac{2}{1-\beta}\right)I \end{bmatrix} \qquad C_3 = \begin{bmatrix} 0 \\ I \\ -I \end{bmatrix} \qquad t = \begin{bmatrix} r \\ 0 \\ 0 \end{bmatrix}. \tag{23}$$

Finally, we set $\eta_2 = \zeta$ and $\eta_3 = \left|\frac{2}{1-\beta}\right|\zeta$. Using (21) and that $x_3 = \left(\frac{2}{1-\beta}\right)x_2$ in any feasible solution, the follower problem objective $g_2^\top x_2 + g_3^\top x_3$ may be rewritten as

$$\left(1 - \tfrac{2}{1-\beta}\right) f_2^\top x_2 + \beta f_2^\top \left(\tfrac{2}{1-\beta}\right) x_2 = \left(\tfrac{-\beta-1}{1-\beta} + \tfrac{2\beta}{1-\beta}\right) f_2^\top x_2 = \left(\tfrac{\beta-1}{1-\beta}\right) f_2^\top x_2 = -f_2^\top x_2.$$

The constraint $\|x_3\|_\infty \le \eta_3$ is equivalent to $\left\|\left(\frac{2}{1-\beta}\right)x_2\right\|_\infty \le \left|\frac{2}{1-\beta}\right|\zeta$, which is exactly the same as the constraint $\|x_2\|_\infty \le \eta_2 = \zeta$, so the follower problem may be written

$$\begin{array}{ll} \min & -f_2^\top x_2 \\ \text{ST} & B_1 x_1 + B_2 x_2 \le r \\ & \|x_2\|_\infty \le \zeta. \end{array} \tag{24}$$

Next consider the leader objective, which we may rewrite using (21) and $x_3 = \left(\frac{2}{1-\beta}\right)x_2$ as

$$f_1^\top x_1 + \left(1 - \tfrac{2}{1-\beta}\right) f_2^\top x_2 + f_2^\top \left(\tfrac{2}{1-\beta}\right) x_2 = f_1^\top x_1 + \left(\tfrac{-\beta-1}{1-\beta} + \tfrac{2}{1-\beta}\right) f_2^\top x_2 = f_1^\top x_1 + f_2^\top x_2.$$

Combining this observation with the form of the follower problem in (24), it follows that the constructed BBSBLP(β) instance is completely equivalent to the BOLP instance. It remains only to observe that the number of bits needed to encode the BBSBLP(β) instance is clearly bounded by a polynomial function of the number of bits needed to encode the BOLP instance. □

Note: the respective proportionality factors $1 - \frac{2}{1-\beta}$ and $\frac{2}{1-\beta}$ in (21) and (22) were chosen because they form the solution (ν_1, ν_2) to the simultaneous equations $\nu_1 + \nu_2 = 1$ and $\nu_1 + \beta\nu_2 = -1$.

6 Complexity of Bilevel MSD Risk Models

We now consider the complexity of the problem class BLRP(MSD$_\gamma$), for $\gamma \in (0,1) \cap \mathbb{Q}$. To avoid the replicated, redundant X_1 and X_2 variables used in the formulations introduced in Sections 2 and 3, we change (without altering the substance of the problem formulation) the decision variable notation to

First stage: $x_1 \in \mathbb{R}^{n_1}$

Second stage: $x_2(S) \in \mathbb{R}^{n_2}$ for each $S \in \mathcal{E}_2$

Third stage: $x_3(\omega) \in \mathbb{R}^{n_3}$ for each $\omega \in \Omega$. For any $S \in \mathcal{E}_2$, the notation $x_3(S)$ denotes the concatenation of all $x_3(\omega)$, $\omega \in S$.

We also introduce some additional "helper" variables $u_1 \in \mathbb{R}$ and $u_2(S) \in \mathbb{R}$, $S \in \mathcal{E}_2$; then, for $\rho = \text{MSD}_\gamma$, the problem system (16)-(17) may be expressed as

$$\begin{array}{lll} \min & c_1^\top x_1 + u_1 + \gamma \sum_{S\in\mathcal{E}_2} \sum_{\omega\in S} \mathrm{P}\{\omega\} \left[c_2(S)^\top x_2(S) + c_3(\omega)^\top x_3(\omega) - u_1\right]_+ & \\ \text{ST} & A_{11} x_1 \le b_1 & \\ & u_1 = \sum_{S\in\mathcal{E}_2} \mathrm{P}\{S\}\, c_2(S)^\top x_2(S) + \sum_{\omega\in\Omega} \mathrm{P}\{\omega\}\, c_3(\omega)^\top x_3(\omega) & \\ & \big(x_2(S), x_3(S)\big) \in \widetilde{\mathcal{X}}_S^*(x_1) & \forall\, S \in \mathcal{E}_2, \end{array} \tag{25}$$

where, for each $S \in \mathcal{E}_2$, $\widetilde{\mathcal{X}}_S^*(x_1)$ denotes the set of $\big(x_2(S), x_3(S)\big)$ portions of all optimal solutions $\big(x_2(S), x_3(S), u_2(S)\big)$ to the scenario-S follower problem

$$\begin{array}{lll} \min & c_2(S)^\top x_2(S) + u_2(S) + \gamma \frac{\mathrm{P}\{\omega\}}{\mathrm{P}\{S\}} \left[c_3(\omega)^\top x_3(\omega) - u_2(S) \right]_+ & \\ \mathrm{ST} & A_{21}(S)x_1 + A_{22}(S)x_2 \le b_2(S) & \\ & A_{31}(\omega)x_1 + A_{32}(\omega)x_2(S) + A_{33}(\omega)x_3(\omega) \le b_3(\omega) & \forall\, \omega \in S \\ & u_2(S) = \sum_{\omega \in S} \big(\frac{\mathrm{P}\{\omega\}}{\mathrm{P}\{S\}}\big) c_3(\omega)^\top x_3(\omega). & \end{array} \tag{26}$$

We now show how to construct a subclass of BLRP(MSD$_\gamma$) problems that is very similar to BBSBLP(β) for an appropriate choice of β. Consider a similar probability space to the one considered at the end of Section 3, with $\Omega = \{\omega_1, \omega_2, \omega_3\}$, partitioned into two second-stage scenarios $S_1 = \{\omega_1, \omega_2\}$ and $S_2 = \{\omega_3\}$. For two parameters $p_1, p_2 \in (0,1) \cap \mathbb{Q}$, we set up the remainder of the BLRP(MSD$_\gamma$) problem instance as follows, using the data of the given BBSBLP(β) instance, and as illustrated in Figure 2:

- The stage-one variables are $x_1 \in \mathbb{R}^{n_1}$, with corresponding cost coefficients $h_1 \in \mathbb{Q}^{n_1}$.
- Scenario S_1 has probability p_1, and hence scenario S_2 has probability $1 - p_1$.
- In scenario S_1, the recourse decision variables are $x_2 \in \mathbb{R}^{n_2}$; here, we omit the "(S_1)" from the original notation $x_2(S_1)$ for brevity, because $x_2(S_2)$ will be essentially fixed. The cost coefficient vector for x_2 is $h_2 \in \mathbb{Q}^{n_2}$, and the constraints are $\|x_2\|_\infty \le \eta_2$. Given that S_1 occurs, we choose the conditional probability of ω_1 to be p_2, so the conditional probablity of ω_2 given S_1 is $1 - p_2$.
  - For outcome ω_1, the final recourse decision variables are $x_3 \in \mathbb{R}^{n_3}$; we omit the "(ω_1)" following x_3 for brevity, because $x_3(\omega_2)$ and $x_3(\omega_3)$ will be essentially fixed. The cost coefficient vector for x_3 is $h_3 \in \mathbb{Q}^{n_3}$, and the constraints are

    $$C_1 x_1 + C_2 x_2 + C_3 x_3 \le t \qquad\qquad \|x_3\|_\infty \le \eta_3.$$

  - For outcome ω_2, the final stage incurs a fixed cost of $K_2 = \eta_3 \|h_3\|_1$.
- Scenario S_2, from which the only possible final-stage consequence is ω_3, incurs a fixed cost of

$$K_1 = \eta_2 \|h_2\|_1 + \left(1 + \tfrac{2p_1p_2}{1-p_1}\right) K_2. \tag{27}$$

Clearly, it is possible to configure the input data of a BLRP(MSD$_\gamma$) instance so that the above arrangement is acheived, and the space required is polynomial in the space required to encode $(h_1, h_2, h_3, C_1, C_2, C_3, \eta_2, \eta_3)$. For example, to make outcome ω_2 incur a fixed cost of K_2, we may set $c_3(\omega_2) = [K_2 \;\; 0^\top]$, and

$$A_{31}(\omega_2) = 0 \qquad A_{32}(\omega_2) = 0 \qquad A_{33}(\omega_2) = \begin{bmatrix} 1 & 0^\top \\ -1 & 0^\top \end{bmatrix} \qquad b_3(\omega_2) = \begin{bmatrix} 1 \\ -1 \end{bmatrix}.$$

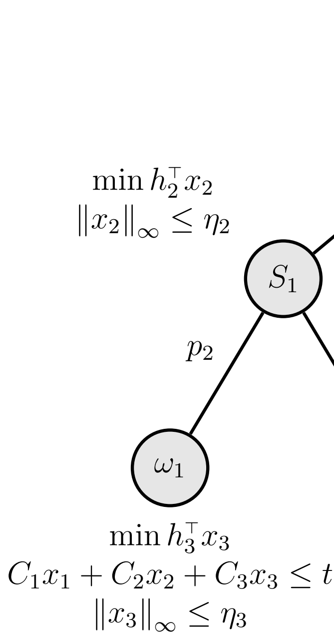

Figure 2: Scenario tree for reduction of BBSBLP(β) instances to BLRP(MSD$_\gamma$) instances.

The intent of this construction is that K_2 is sufficiently large that the conditional expected value $u_2(S_1)$ of the stage-three costs given that scenario S_1 occurs must always be worse than for outcome ω_1 for all feasible values of the decision variables. Similarly, although the analysis is more complicated, K_1 is taken sufficiently large that the expected value u_1 at the root of the scenario tree must always be worse than the objective for either of outcomes ω_1 and ω_2 for all feasible settings of the decision variables. We will now show that these properties mean that the resulting BLRP(MSD$_\gamma$) problem is equivalent to a problem very similar to BBSBLP(β) for an appropriate choice of β.

First, for any rational $\gamma \in (0, 1]$, formulating an instance (25)-(26) of BLRP(MSD$_\gamma$) as just described yields the following problem, where we abbreviate $u_2(S_1)$ to simply u_2 since $u_2(S_2)$ is a constant:

$$
\begin{array}{lll}
\min & h_1^\top x_1 + u_1 + \gamma\big(p_1 p_2 \left[h_2^\top x_2 + h_3^\top x_3 - u_1\right]_+ \\
& \qquad\qquad + p_1(1-p_2)\left[h_2^\top x_2 + K_2 - u_1\right]_+ + (1-p_1)\left[K_1 - u_1\right]_+\big) \\
\text{ST} & u_1 = p_1 p_2 (h_2^\top x_2 + h_3^\top x_3) + p_1(1-p_2)(h_2^\top x_2 + K_2) + (1-p_1)K_1 \\
& (x_2, x_3) \in \begin{array}[t]{ll} \text{Arg min} & h_2^\top x_2 + u_2 + \gamma\left(p_2\left[h_3^\top x_3 - u_2\right]_+ + (1-p_2)\left[K_2 - u_2\right]_+\right) \\ \text{ST} & u_2 = p_2 h_3^\top x_3 + (1-p_2)K_2 \\ & C_1x_1 + C_2x_2 + C_3x_3 \leq t \\ & \|x_2\|_\infty \leq \eta_2 \\ & \|x_3\|_\infty \leq \eta_3. \end{array}
\end{array}
\tag{28}
$$

Now consider the follower problem of (28). By the choice of K_2, we have that for any feasible value of x_3,

$$h_3^\top x_3 \le \|h_3\|_1 \|x_3\|_\infty \le \eta_3 \|h_3\|_1 = K_2,$$

and, since u_2 is a convex combination of $h_3^\top x_3$ and K_2, we therefore always have $h_3^\top x_3 \le u_2 \le K_2$. Hence, the first $[\,\cdot\,]_+$ term in the follower objective is always zero, and the second $[\,\cdot\,]_+$ term may be written as

$$K_2 - u_2 = K_2 - \big(p_2 h_3^\top x_3 + (1-p_2)K_2\big) = p_2\big(K_2 - h_3^\top x_3\big).$$

Substituting for u_2 and the $[\,\cdot\,]_+$ terms, we obtain the equivalent follower objective function

$$\begin{aligned}
& h_2^\top x_2 + p_2 h_3^\top x_3 + (1-p_2)K_2 + \gamma p_2 \cdot 0 + \gamma(1-p_2)p_2\big(K_2 - h_3^\top x_3\big) \\
= \;& h_2^\top x_2 + \big(p_2 - \gamma(1-p_2)p_2\big)h_3^\top x_3 + \big((1-p_2) + \gamma(1-p_2)p_2\big)K_2 \\
= \;& h_2^\top x_2 + p_2(1-\gamma+\gamma p_2)h_3^\top x_3 + (1-p_2)(1+\gamma p_2)K_2.
\end{aligned}$$

Discarding the constant term $(1-p_2)(1+\gamma p_2)K_2$ from the objective, the follower problem is thus equivalent to

$$\begin{array}{ll}
\min & h_2^\top x_2 + p_2(1-\gamma+\gamma p_2)h_3^\top x_3 \\
\text{ST} & C_1 x_1 + C_2 x_2 + C_3 x_3 \le t \\
& \|x_2\|_\infty \le \eta_2 \\
& \|x_3\|_\infty \le \eta_3.
\end{array} \tag{29}$$

We now consider the leader problem in (28). We claim that the choice of K_1 in (27) is sufficiently large for all feasible values of (x_2, x_3) that we have

$$u_1 \ge h_2^\top x_2 + K_2 \ge h_2^\top x_2 + h_3^\top x_3. \tag{30}$$

The second inequality in (30) follows immediately from $h_3^\top x_3 \le \|h_3\|_1\|x_3\|_\infty \le \eta_3\|h_3\|_1 = K_2$, so it remains to prove the first inequality. We note that

$$\begin{aligned}
u_1 &= p_1\big(h_2^\top x_2 + p_2 h_3^\top x_3 + (1-p_2)K_2\big) + (1-p_1)K_1 \\
&= p_1\big(h_2^\top x_2 + p_2 h_3^\top x_3 + (1-p_2)K_2\big) + (1-p_1)\left(\eta_2\|h_2\|_1 + \left(1 + \tfrac{2p_1p_2}{1-p_1}\right)K_2\right) \\
&= \big[p_1 h_2^\top x_2 + (1-p_1)\eta_2\|h_2\|_1\big] + p_1p_2h_3^\top x_3 + p_1(1-p_2)K_2 + (1-p_1+2p_1p_2)K_2.
\end{aligned} \tag{31}$$

Since $h_2^\top x_2 \le \|h_2\|_1 \|x_2\|_\infty \le \eta_2 \|h_2\|_1$, we have

$$p_1 h_2^\top x_2 + (1-p_1)\eta_2\|h_2\|_1 \ge h_2^\top x_2. \tag{32}$$

Next, we observe that

$$h_3^\top x_3 \ge -\|h_3\|_1 \|x_3\|_\infty \ge -\|h_3\|_1 \eta_3 = -K_2. \tag{33}$$

Substituting (32) and (33) into (31), we have

$$\begin{aligned}
u_1 &\ge h_2^\top x_2 - p_1p_2K_2 + p_1(1-p_2)K_2 + (1-p_1+2p_1p_2)K_2 \\
&= h_2^\top x_2 + \big(-p_1p_2 + p_1 - p_1p_2 + 1 - p_1 + 2p_1p_2\big)K_2 \\
&= h_2^\top x_2 + K_2,
\end{aligned}$$

establishing (30). It follows immediately that the first two $[\,\cdot\,]_+$ terms in the leader objective of (28) are always zero. Noting that

$$K_1 > \eta_2 \left\|h_2\right\|_1 + K_2 \geq h_2^\top x_2 + K_2 \geq h_2^\top x_2 + h_3^\top x_3$$

for any feasible (x_2, x_3), it follows from u_1 being a convex combination of $h_2^\top x_2 + h_3^\top x_3$, $h_2^\top x_2 + K_2$, and K_1 that $K_1 \geq u_1$. Therefore, $[K_1 - u_1]_+ = K_1 - u_1$, and the leader objective of (28) may be written, where "$\simeq$" denotes equivalence up to a constant among functions of (x_1, x_2, x_3), as

$$\begin{aligned}
&h_1^\top x_1 + u_1 + \gamma(1-p_1)(K_1 - u_1) \\
\simeq\ & h_1^\top x_1 + (1-\gamma+\gamma p_1)u_1 \\
=\ & h_1^\top x_1 + (1-\gamma+\gamma p_1)\big(p_1 h_2^\top x_2 + p_1 p_2 h_3^\top x_3 + p_1(1-p_2)K_2 + (1-p_1)K_1\big) \\
\simeq\ & h_1^\top x_1 + (1-\gamma+\gamma p_1)\big(p_1 h_2^\top x_2 + p_1 p_2 h_3^\top x_3\big) \\
=\ & h_1^\top x_1 + p_1(1-\gamma+\gamma p_1)\big(h_2^\top x_2 + p_2 h_3^\top x_3\big).
\end{aligned}$$

Combining this form of the leader objective with (29), we may express the entire problem as

$$\begin{array}{lllll}
\min & h_1^\top x_1 + p_1(1-\gamma+\gamma p_1)h_2^\top x_2 + p_1 p_2(1-\gamma+\gamma p_1)h_3^\top x_3 & & & \\
\text{ST} & (x_2, x_3) \in & \text{Arg min} & h_2^\top x_2 + p_2(1-\gamma+\gamma p_2)h_3^\top x_3 & \\
& & \text{ST} & C_1 x_1 + C_2 x_2 + C_3 x_3 \leq t & \qquad (34) \\
& & & \left\|x_2\right\|_\infty \leq \eta_2 & \\
& & & \left\|x_3\right\|_\infty \leq \eta_3. &
\end{array}$$

For an appropriate choice of β, this problem form has exactly the same constraint structure as BBSBLP(β). We can now exploit the differing relative scaling of the $h_2^\top x_2$ and $h_3^\top x_3$ terms in the two objective functions of (34) to reduce BBSBLP(β) to BLRP(MSD_γ) for an appropriate choice of β, thus proving that BLRP(MSD_γ) is $\mathcal{NP}$-hard.

Proposition 5 *For any $\gamma \in (0,1] \cap \mathbb{Q}$, the problem class* BLRP(MSD_γ) *is $\mathcal{NP}$-hard.*

Proof. The proof is by reduction from BBSBLP($1-\gamma/2$); note that since $\gamma > 0$, it follows that $1-\gamma/2 \neq 1$, and thus that BBSBLP($1-\gamma/2$) is $\mathcal{NP}$-hard by Proposition 4. Also, since $\gamma \leq 1$, we have $1-\gamma/2 > 0$. Now consider any instance $(g_1, g_2, g_3, C_1, C_2, C_3, t, \eta_2, \eta_3)$ of BBSBLP($1-\gamma/2$), fix $p_1 = p_2 = 1/2$, and set

$$\begin{aligned}
h_1 &= g_1 \\
h_2 &= \left(\frac{1}{p_1(1-\gamma+\gamma p_1)}\right) g_2 = \left(\frac{2}{1-\gamma/2}\right) g_2 \\
h_3 &= \left(\frac{1}{p_1 p_2(1-\gamma+\gamma p_1)}\right) g_3 = \left(\frac{4}{1-\gamma/2}\right) g_3.
\end{aligned}$$

We now use $(h_1, h_2, h_3, C_1, C_2, C_3, t, \eta_2, \eta_3)$ to construct a BLRP(MSD$_\gamma$) problem instance of the form (28); the space required to encode (h_1, h_2, h_3) is polynomial in the space required to encode (g_1, g_2, g_3), so the size of the resulting BLRP(MSD$_\gamma$) instance is polynomial in the encoding sized of the BBSBLP$(1-\gamma/2)$ instance $(g_1, g_2, g_3, C_1, C_2, C_3, t, \eta_2, \eta_3)$. From the analysis above, the resulting BLRP(MSD$_\gamma$) instance is equivalent to (34). Substituting the above choices of h_1, h_2, and h_3 into the leader objective of (34), along with $p_1 = p_2 = 1/2$, we obtain $g_1^\top x_1 + g_2^\top x_2 + g_3^\top x_3$ as the leader objective, exactly as in BBSBLP$(1-\gamma/2)$. Making the same substitutions into the follower objective of (34), we obtain

$$\begin{aligned} h_2^\top x_2 + p_2(1-\gamma+\gamma p_2)h_3^\top x_3 &= \left(\frac{2}{1-\gamma/2}\right) g_2^\top x_2 + \left(\frac{(1/2)(1-\gamma/2)\cdot 4}{1-\gamma/2}\right) g_3^\top x_3 \\ &= \left(\frac{2}{1-\gamma/2}\right) g_2^\top x_2 + 2 g_3^\top x_3. \end{aligned}$$

Applying the positive scaling factor $(1-\gamma/2)/2$ to both terms in its objective does not make any difference to the solution set of follower problem, so we may equivalently use the follower objective $g_2^\top x_2 + (1-\gamma/2) g_3^\top x_3$. In summary, the BLRP(MSD$_\gamma$) instance we have constructed is equivalent to the problem

$$\begin{array}{lllll} \min & g_1^\top x_1 + g_2^\top x_2 + g_3^\top x_3 & & \\ \text{ST} & (x_2, x_3) \in & \text{Arg min} & g_2^\top x_2 + (1-\gamma/2) g_3^\top x_3 \\ & & \text{ST} & C_1 x_1 + C_2 x_2 + C_3 x_3 \le t \\ & & & \|x_2\|_\infty \le \eta_2 \\ & & & \|x_3\|_\infty \le \eta_3, \end{array}$$

precisely the BBSBLP$(1-\gamma/2)$ instance encoded by $(g_1, g_2, g_3, C_1, C_2, C_3, t, \eta_2, \eta_3)$. Since the BLRP(MSD$_\gamma$) instance encoding size is polynomial in the size of the BBSBLP$(1-\gamma/2)$ instance, existence of a polynomial-time solution algorithm for BLRP(MSD$_\gamma$) would imply polynomial-time algorithm for the $\mathcal{NP}$-hard problem class BBSBLP$(1-\gamma/2)$. □

Corollary 6 *The problem class* BLRP(MSD$_{(0,1]}$)*, with* γ *encoded as part of the problem input, is also* $\mathcal{NP}$*-hard.*

Proof. Consider any instance of the problem class BLRP(MSD$_{1/2}$), which is $\mathcal{NP}$-hard by Proposition 5. Appending $\gamma = 1/2$ to the encoding of this instance only increases the problem size by a constant, so a polyomial-time algorithm for BLRP(MSD$_{(0,1]}$) would imply a polynomial-time algorithm for BLRP(MSD$_{1/2}$). □

7 Complexity of Bilevel AVaR Risk Models

We now consider the complexity of bilevel models using the AVaR$_\alpha$ risk measure instead of the MSD$_\gamma$ risk measure; the overall analysis technique is similar to the MSD$_\gamma$ case, but involves a reduction from BBSBLP(0), regardless of the value of α.

To set up the analysis, we construct a simple scenario tree similar to that of Section 6 and the end of Section 3, but with different probabilities, all based on the parameter α, and with outcome ω_3 representing a highly desirable result rather than an highly undesirable one:

- The stage-one variables are $x_1 \in \mathbb{R}^{n_1}$, with corresponding cost coefficients $h_1 \in \mathbb{Q}^{n_1}$.
- Scenario S_1 has probability α, and hence scenario S_2 has probability $1-\alpha$.
- Scenario S_1 has the recourse decision variables $x_2 \in \mathbb{R}^{n_2}$, with corresponding cost coefficient vector $h_2 \in \mathbb{Q}^{n_2}$ and subject to the constraint $\|x_2\|_\infty \le \eta_2$. Given that S_1 occurs, the conditional probability of outcome ω_1 is $1-\alpha$; hence, the conditional probability of outcome ω_2 given S_1 is α.
  - For outcome ω_1, the final recourse decision variables are $x_3 \in \mathbb{R}^{n_3}$, with corresponding cost coefficients $h_3 \in \mathbb{Q}^{n_3}$ and subject to the constraints
  $$C_1x_1 + C_2x_2 + C_3x_3 \le t \qquad\qquad \|x_3\|_\infty \le \eta_3.$$
  - For outcome ω_2, the final stage incurs a fixed cost of $K_2 = \eta_3\|h_3\|_1 + 1$.
- Scenario S_2, from which the only possible final-stage consequence is ω_3, incurs a fixed cost of $-K_1$ (that is, a benefit), where $K_1 = \eta_2\|h_2\|_1 + K_2 = \eta_2\|h_2\|_1 + \eta_3\|h_3\|_1 + 1$.

This slightly modified scenario structure is shown in Figure 3. As long as $\alpha \in \mathbb{Q}$, the space required to express the resulting BLRP(AVaR_α) problem instance is polynomially bounded in the space required to express $(h_1, h_2, h_3, C_1, C_2, C_2, t, \eta_2, \eta_3)$. To express this BLRP($\text{AVaR}_\alpha$) instance in a compact manner, we define two random variables parameterized by x_2 and x_3:

$$W_2(x_2,x_3) = \begin{cases} -K_1, & \text{with probability } 1-\alpha \\ h_2^\top x_2 + h_3^\top x_3, & \text{with probability } \alpha-\alpha^2 \\ h_2^\top x_2 + K_2, & \text{with probability } \alpha^2 \end{cases}$$

$$W_3(x_3) = \begin{cases} h_3^\top x_3, & \text{with probability } 1-\alpha \\ K_2, & \text{with probability } \alpha. \end{cases}$$

This BLRP(AVaR_α) instance then has the form

$$\begin{array}{lll} \min & h_1^\top x_1 + \text{AVaR}_\alpha\big(W_2(x_2,x_3)\big) & \\ \text{ST} & (x_2,x_3) \in \ \text{Arg min} & h_2^\top x_2 + \text{AVaR}_\alpha\big(W_3(x_3)\big) \\ & \qquad\qquad\ \ \text{ST} & C_1x_1 + C_2x_2 + C_3x_3 \le t \\ & & \|x_2\|_\infty \le \eta_2 \\ & & \|x_3\|_\infty \le \eta_3, \end{array} \tag{35}$$

As before, we write x_2 instead of $x_2(S_1)$, since the value of $x_2(S)$ is only of importance for $S = S_1$, and similarly we write x_3 instead of $x_3(\omega_1)$.

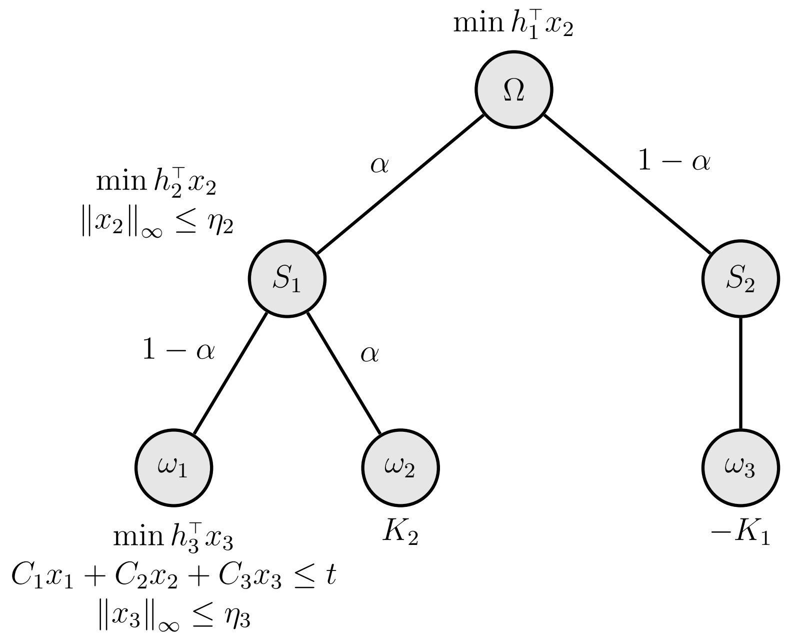

Figure 3: Scenario tree for reduction of BBSBLP(0) instances to BLRP(AVaR_α) instances.

To provide some intuition, the problem instance has been constructed so that all quantiles of $W_3(x_3)$ above the $1-\alpha$ quantile are simply K_2 for any feasible value of x_3, so the $\text{AVaR}_\alpha(\,\cdot\,)$ term in the follower objective of (35) is equivalent to a constant. We now verify this claim using the formal definition (2) of the AVaR_α risk measure. Noting that for any feasible value of x_3, we have

$$h_3^\top x_3 \le \|h_3\|_1 \|x_3\|_\infty \le \eta_3 \|h_3\|_1 < \eta_3 \|h_3\|_1 + 1 = K_2, \tag{36}$$

the lower inverse cumulative function $F^{-1}_{W_3(x_3)}(\,\cdot\,)$ as defined through (3) takes the form

$$F^{-1}_{W_3(x_3)}(\nu) = \begin{cases} h_3^\top x_3, & \text{if } \nu \in (0, 1-\alpha] \\ K_2, & \text{if } \nu \in (1-\alpha, 1], \end{cases}$$

hence we have

$$\text{AVaR}_\alpha\big(W_3(x_3)\big) = \frac{1}{\alpha}\int_{1-\alpha}^{1} F^{-1}_{W_3(x_3)}(\nu)\, d\nu = \frac{1}{\alpha}\int_{1-\alpha}^{1} K_2\, d\nu = \tfrac{1}{\alpha}(\alpha K_2) = K_2,$$

since the $F^{-1}_{W_3(x_3)}(\nu) = K_2$ throughout $[1-\alpha, 1]$ except on the singleton $\{1-\alpha\}$, which has measure zero. Thus, the follower objective of (35) may be replaced, without any change to the solution set of the follower problem, by $h_2^\top x_2 + K_2$, or, dropping the constant, equivalently simply $h_2^\top x_2$.

Consider now the $\text{AVaR}_\alpha\big(W_3(x_2, x_3)\big)$ term in the leader objective of (35). We note that

$$-K_1 = -\eta_2 \|h_2\|_1 - \eta_3 \|h_3\|_1 - 1 < h_2^\top x_2 + h_3^\top x_3$$

for any feasible values of x_2 and x_3. Further, as we have already established in (36) that $h_3^\top x_3 < K_2$ for all feasible values of x_3, we have that

$$-K_1 < h_2^\top x_2 + h_3^\top x_3 < h_2^\top x_2 + K_2$$

for all feasible values of x_2 and x_3. Thus, the lower inverse cumulative function $F^{-1}_{W_2(x_2,x_3)}(\,\cdot\,)$ takes the form

$$F^{-1}_{W_2(x_2,x_3)}(\nu) = \begin{cases} -K_1, & \text{if } \nu \in (0, 1-\alpha] \\ h_2^\top x_2 + h_3^\top x_3, & \text{if } \nu \in (1-\alpha, 1-\alpha^2] \\ h_2^\top x_2 + K_2, & \text{if } \nu \in (1-\alpha^2, 1]. \end{cases}$$

Applying the definition of the AVaR_α risk measure, we have

$$\begin{aligned} \mathrm{AVaR}_\alpha\big(W_2(x_2,x_3)\big) &= \frac{1}{\alpha}\int_{1-\alpha}^{1} F^{-1}_{W_2(x_2,x_3)}(\nu)\,d\nu \\ &= \frac{1}{\alpha}\left(\int_{1-\alpha}^{1-\alpha^2}(h_2^\top x_2 + h_3^\top x_3)\,d\nu + \int_{1-\alpha^2}^{1}(h_2^\top x_2 + K_2)\,d\nu\right) \\ &= \frac{1}{\alpha}\Big(\big(1-\alpha^2-(1-\alpha)\big)(h_2^\top x_2 + h_3^\top x_3) + \alpha^2(h_2^\top x_2 + K_2)\Big) \\ &= \tfrac{1}{\alpha}\big(\alpha h_2^\top x_2 + (\alpha-\alpha^2)h_3^\top x_3 + \alpha^2 K_2\big) \\ &= h_2^\top x_2 + (1-\alpha)h_3^\top x_3 + \alpha K_2. \end{aligned}$$

Substituting this equation into the leader objective of (35), discarding the resulting constant αK_2, and also substituting the already-established equivalent follower objective $h_2^\top x_2$, we arrive at the equivalent problem

$$\begin{array}{lllll} \min & h_1^\top x_1 + h_2^\top x_2 + (1-\alpha)h_3^\top x_3 & & & \\ \mathrm{ST} & (x_2,x_3) \in & \mathrm{Arg\ min} & h_2^\top x_2 & \\ & & \mathrm{ST} & C_1x_1 + C_2x_2 + C_3x_3 \le t & \\ & & & \|x_2\|_\infty \le \eta_2 & \\ & & & \|x_3\|_\infty \le \eta_3. & \end{array} \qquad (37)$$

This problem is essentially identical in form to BBSBLP(0), a property that we now exploit.

Proposition 7 *For any $\alpha \in (0,1) \cap \mathbb{Q}$, the problem class* $\mathrm{BLRP}(\mathrm{AVaR}_\alpha)$ *is $\mathcal{NP}$-hard.*

Proof. Consider any instance $(g_1, g_2, g_3, C_1, C_2, C_3, t, \eta_2, \eta_3)$ of BBSBLP(0), and set

$$h_1 = g_1 \qquad\qquad h_2 = g_2 \qquad\qquad h_3 = \left(\tfrac{1}{1-\alpha}\right) g_3.$$

Now construct a BLRP(AVaR_α) instance of the form (35), which by the immediately preceding analysis is equivalent to (37). Substituting the definitions of h_1, h_2, h_3 above into (37), we obtain the equivalent problem

$$\begin{array}{lllll} \min & g_1^\top x_1 + g_2^\top x_2 + g_3^\top x_3 & & & \\ \text{ST} & (x_2, x_3) \in & \text{Arg min} & g_2^\top x_2 & \\ & & \text{ST} & C_1 x_1 + C_2 x_2 + C_3 x_3 \le t & \\ & & & \|x_2\|_\infty \le \eta_2 & \\ & & & \|x_3\|_\infty \le \eta_3, & \end{array}$$

which is precisely the original BBSBLP(0) problem instance. The space needed to encode the BLRP(AVaR_α) instance is polynomially bounded in the space to encode the BBSBLP(0) instance, so existence of a polynomial-time algorithm for BLRP(AVaR_α) would imply the existence of a polynomial-time algorithm for BBSBLP(0), which is $\mathcal{NP}$-hard. □

Corollary 8 *The problem class* BLRP($\text{AVaR}_{(0,1)}$)*, with the quantile parameter α encoded as part of the problem input, is also $\mathcal{NP}$-hard.*

Proof. Similar to the proof of Corollary 6. □

8 Some Example Computational Experiments: Local Optima of Bilevel Risk Models

While the results of the last two sections are negative, they do not necessarily imply that multilevel models of the kind proposed in Section 3 should be completely avoided. Instances of $\mathcal{NP}$-hard problems are frequently solved in practice, particularly in the field of mixed integer programming, so $\mathcal{NP}$-hardness does not necessarily mean that an entire class of models must be avoided. It does, however, suggest that it may be useful to explore approximation methods, or if exact solutions are desired, some form of implicit enumeration (branch-and-bound) algorithm. The models we propose have a rich structure that could possibly be exploited by customized implicit enumeration or approximation algorithms.

To get a preliminary feeling for how hard the models we propose might be in a practical setting, we experimented with some realistic three-stage stochastic programming models created using the problem generator described by Collado et al. (2012, Section 8). These models have the following structure:

- There are P_1 parts.
- There are P_2 products constructed from overlapping subsets of these parts.
- Initially, we decide the quantity of each part we wish to order.
- After the initial decision, the short-term demand for each product becomes known.

- We then decide the quantity of each product to build. Each product may be sold up to its short-term demand, producing revenue. Products that we build beyond their short-term demand will still be sold eventually, but after incurring an as-yet-unknown "storage cost"
- Finally, the storage costs become known and the final profit may be calculated.

Throughout, we modeled a moderate level of risk aversion by using the $\mathrm{MSD}_{0.3}$ risk measure, that is, a mean-semideviation measure with $\gamma = 0.3$.

For our test datasets, we tried $(P_1, P_2) = (5, 5), (10, 5), (10, 10)$. We used either 5 or 10 first-stage scenarios, each leading to a further 5 or 10 second-stage outcomes, for a total of between 25 and 100 total outcomes. For each combination of parameters, we generated 5 different test data sets. Finally, we formulated two different optimization problems for each dataset:

1. The first version, which we denote "TC", used a time-consistent nested risk measure system of the form (10), in which the one-step risk measures ρ_1 and ρ_2 both derived from $\mathrm{MSD}_{0.3}$. We formulated this model as a single linear program expressed in AMPL (Fourer et al. 1991), and solved it to optimality with the LOQO Newton barrier solver (Vanderbei 1999, Benson et al. 2002).
2. The second version, which we denote "NTC", was an instance of $\mathrm{BLRP}(\mathrm{MSD}_{0.3})$, that is, a bilevel problem using the non-time-consistent risk-measure system given by (7) with $\gamma = 0.3$. Using the `complements` keyword in AMPL, we reformulated this model as a mathematical program with equlibrium constraints (MPEC); see for example Luo et al. (1996). We then used LOQO to identify local optima satisfying the MPEC necessary optimality conditions, using the capabilities described in Benson et al. (2006).

LOQO was able to identify local optimal solutions to all the models tested (in the case of the TC models, these solutions are also global optima). Solution times were much longer for the NTC models than for the TC models, but never exceeded 30 seconds on a single core of a 3GHz Xeon X5472 processor.

The results are summarized in Table 1. The first column is in the format P_1-P_1-s_1-$s_2(k)$, where P_1 is the number of parts, P_2 is the number of products, s_1 is the number of second-stage scenarios (that is, $s_1 = |\mathcal{E}_2|$), and s_2 is the number of third-stage outcomes per scenario (that is, each scenario $S \in \mathcal{E}_2$ has $|S| = s_2$, and $|\Omega| = s_1 s_2$). For each value of k, we generated a problem instance using a different random-number seed. The next four columns respectively show the following objective values:

1. The optimal objective value of the bilevel NTC model
2. The non-time-consistent first-stage objective function evaluated at the optimal solution of the TC model
3. The time-consistent objective function evaluated at the optimal solution of the bilevel NTC model

Table 1: Computational results for example problems.

Dataset	Obj-NTC Sol=NTC	Obj-NTC Sol=TC	Obj-TC Sol=NTC	Obj-TC Sol=TC	Δ_1	Δ_2
5-5-5-5(1)	–821.164	–821.164	–789.897	–789.897	0.00%	0.00%
5-5-5-5(2)	–1454.26	–1454.22	–1429.49	–1430.16	0.37%	1.12%
5-5-5-5(3)	–1454.26	–1454.22	–1429.49	–1430.16	0.37%	1.12%
5-5-5-5(4)	–980.502	–979.678	–950.320	–952.151	1.02%	3.71%
5-5-5-5(5)	–2733.06	–2733.06	–2706.21	–2706.21	0.00%	0.00%
10-5-5-10(1)	–13097.4	–13084.1	–12989.8	–12999.5	6.93%	11.99%
10-5-5-10(2)	–1353.85	–1352.58	–1311.13	–1311.27	0.47%	7.13%
10-5-5-10(3)	–1353.85	–1352.58	–1311.13	–1311.27	0.47%	7.13%
10-5-5-10(4)	–3839.42	–3839.42	–3766.54	–3766.54	3.27%	0.00%
10-5-5-10(5)	–4978.18	–4978.10	–4892.19	–4892.19	0.03%	0.09%
10-10-10-10(1)	–11953.8	–11951.9	–11886.3	–11887.0	2.88%	5.47%
10-10-10-10(2)	–8777.65	–8766.31	–8692.34	–8703.97	2.36%	8.75%
10-10-10-10(3)	–5984.11	–5978.23	–5938.50	–5941.64	1.80%	11.42%
10-10-10-10(4)	–9601.78	–9600.26	–9533.19	–9534.04	0.62%	4.07%
10-10-10-10(5)	–11862.4	–11862.0	–11765.9	–11766.0	0.08%	0.30%

4. The optimal objective value of the TC model.

Note that all the optimal objective values are negative because the formulation converts a profit maximization problem into the minimization form which is standard in the coherent risk literature.

The last two columns of the table summarize the differences in the two models' first- and second-stage solution vectors, respectively. The Δ_1 columns displays the percentage difference in the first-stage decisions, as calculated by

$$\Delta_1 = 100 \cdot \frac{\left\| x_1^{\mathrm{NTC}} - x_1^{\mathrm{TC}} \right\|_2}{\| x_1^{\mathrm{TC}} \|_2},$$

where the "TC" and "NTC" superscripts indicate the solutions obtained for the TC and NTC models, respectively. For the second stage, Δ_2 reports the maximum of a similar measure over all second-stage scenarios:

$$\Delta_2 = 100 \cdot \max_{S \in \mathcal{E}_2} \left\{ \frac{\left\| x_2^{\mathrm{NTC}}(S) - x_2^{\mathrm{TC}}(S) \right\|_2}{\| x_2^{\mathrm{TC}}(S) \|_2} \right\}.$$

While the solutions obtained for the NTC models are not guaranteed to be globally optimal by LOQO, it appears that they are of high quality. While the solutions values can vary — Δ_1 approaches 7% for one model and Δ_2 is almost 12% for another — the differences in objective values are small.

While they include only one class of examples, these computational results suggest that it is at least possible to closely approximate the solution of a bilevel model by a more tractable single-level, time-consistent model, for the eventual purpose either of simple approximation or incorporation into a branch-and-bound framework.

9 Concluding Remarks

Finally, we make some brief concluding observations. As already noted, the theoretical complexity results above should not be interpreted as indicating that any attempt to obtain optimal or provably near-optimal solutions to problem instances of the form BLRP(ρ) should be abandoned. However, the results do suggest that research into solution methods should not focus on algorithms with polynomial worst-case run-time guarantees (except perhaps for heuristics), and that some form of implicit enumeration is likely to be needed to compute exact solutions. As the results of Section 8 suggest, it may at least sometimes be possible to tightly approximate a multilevel model by a single-level one, a phenomenon that may prove useful in deriving efficient methods for approximately or exactly solving the proposed multilevel models.

Although our $\mathcal{NP}$-hardness results cover only two specific common families of risk measures, it seems reasonable to conjecture that they will extend to any family of risk measures which has a polyhedral dual form (Artzner et al. 1999, Delbaen 2002, Ruszczyński and Shapiro 2006b); this subject is a matter for further research.

Finally, we note that reducibility of the "oppositional" problem form BOLP respectively through BBSBLP($1-\gamma/2$) or BBSBLP(0) to BLRP(MSD_γ) or BLRP(AVaR_α) indicates that it is possible to contrive three-stage stochastic programming problems instances in which the law-invariant use of the MSD_γ or AVaR_α risk measures of the form (7) or (8) breaks time consistency in a particularly dramatic way. Specifically, BLRP(MSD_γ) and BLRP(AVaR_α) instances constructed using such two-step reductions from BOLP possess a scenario which, if it is revealed to have occured, effectively reverses the preference order among the feasible solutions available at the second stage, as compared to the perspective of the first stage. This phenonenon underscores that when using conditioned measures of risk like (7) or (8), the revelation of partial information can at least in theory dramatically change a decision maker's preferences among the remaining courses of action.

References

Acerbi, C. 2002. Spectral measures of risk: a coherent representation of subjective risk aversion. *J. Banking Finance* **26**.

Artzner, P., F. Delbaen, J.-M. Eber, D. Heath. 1999. Coherent measures of risk. *Math. Finance* **9** 203–228.

Bard, J. F. 1991. Some properties of the bilevel programming problem. *J. Optim. Theory Appl.* **68**(2) 371–378.

Ben-Ayed, O., C. E. Blair. 1990. Computational difficulties of bilevel linear programming. *Oper. Res.* **38**(3) 556–560.

Benson, H. Y., A. Sen, D. F. Shanno, R. J. Vanderbei. 2006. Interior-point algorithms, penalty methods and equilibrium problems. *Comput. Optim. Appl.* **34**(2) 155–182.

Benson, H. Y., R. J. Vanderbei, D. F. Shanno. 2002. Interior-point methods for nonconvex nonlinear programming: filter methods and merit functions. *Comput. Optim. Appl.* **23**(2) 257–272.

Cheridito, P., F. Delbaen, M. Kupper. 2006. Dynamic monetary risk measures for bounded discrete-time processes. *Electron. J. Prob.* **11** 57–106.

Collado, R. A., D. Papp, A. Ruszczyński. 2012. Scenario decomposition of risk-averse multistage stochastic programming problems. *Ann. Oper. Res.* **200** 147–170.

Delbaen, F. 2002. Coherent risk measures on general probability spaces. *Advances in finance and stochastics*. Springer, Berlin, 1–37.

Eckstein, Jonathan. 2012. Complexity of bilevel coherent risk programming. RUTCOR Research Report RRR 17-2012, Rutgers University.

Föllmer, H., A. Schied. 2004. *Stochastic finance: an introduction in discrete time, second edition*. DeGruyter, Berlin.

Fourer, R., D. M. Gay, B. W. Kernighan. 1991. *AMPL: A Modeling Language for Mathematical Programming*. Cengage Learning.

Garey, M. R., D. S. Johnson. 1979. *Computers and intractability: a guide to the theory of $\mathcal{NP}$-completeness*. W. H. Freeman, San Francisco.

Hansen, P., B. Jaumard, G. Savard. 1992. New branch-and-bound rules for linear bilevel programming. *SIAM J. Sci. Statist. Comput.* **13**(5) 1194–1217.

Jeroslow, R. G. 1985. The polynomial hierarchy and a simple model for competitive analysis. *Math. Programming* **32**(2) 146–164.

Kupper, M., W. Schachermayer. 2009. Representation results for law invariant time consistent functions. *Math. Financ. Econ.* **2**(3) 189–210.

Luo, Z. Q., J. S. Pang, D. Ralph. 1996. *Mathematical Programs with Equilibrium Constraints*. Cambridge University Press.

Marcotte, P., G. Savard. 1991. A note on the Pareto optimality of solutions to the linear bilevel programming problem. *Comput. Oper. Res.* **18**(4) 355–359.

Riedel, F. 2004. Dynamic coherent risk measures. *Stoch. Processes and Their Appl.* **112** 185–200.

Rockafellar, R. T., S. Uryasev. 2000. Optimization of conditional value-at-risk. *J. Risk* **2** 21–42.

Rockafellar, R. T., S. Uryasev. 2002. Conditional value-at-risk for general loss distributions. *J. Banking Finance* **26** 1442–1471.

Rockafellar, R. T., S. Uryasev, M. Zabarankin. 2006. Generalized deviations in risk analysis. *Finance Stoch.* **10**(1) 51–74.

Ruszczyński, A. 2010. Risk-averse dynamic programming for Markov decision processes. *Math. Programming* **125**(2) 235–261.

Ruszczyński, A., A. Shapiro. 2006a. Conditional risk mappings. *Math. Oper. Res.* **31**(3) 544–561.

Ruszczyński, A., A. Shapiro. 2006b. Optimization of convex risk functions. *Math. Oper. Res.* **31**(3) 433–452.

Scandolo, G. 2003. Risk measures in a dynamic setting. Ph.D. thesis, Università degli Studi di Milano, Milan, Italy.

Shapiro, A. 2012. Time consistency of dynamic risk measures. Optimization Online preprint, Georgia Institute of Technology.

Uryasev, S., R. T. Rockafellar. 2001. Conditional value-at-risk: optimization approach. *Stochastic optimization: algorithms and applications (Gainesville, FL, 2000)*, *Appl. Optim.*, vol. 54. Kluwer, Dordrecht, 411–435.

Vanderbei, R. J. 1999. LOQO: an interior point code for quadratic programming. *Optim. Methods Softw.* **11/12**(1-4) 451–484.