RUTCOR
RESEARCH
REPORT

Multilevel Optimization Modeling for Risk-Averse Stochastic Programming

Jonathan Eckstein[a] Deniz Eskandani[b]
Jingnan Fan[c]

RRR 3-2014, Aug 2014,

RUTCOR
Rutgers Center for
Operations Research
Rutgers University
100 Rockafeller Road
Piscataway, New Jersey
08854
Email: rrr@rutcor.rutgers.edu
http://rutcor.rutgers.edu/∼rrr

[a]Department of Management Science and Information Systems (MSIS) and RUTCOR, Rutgers University, 100 Rockafeller Road, Piscataway, NJ 08854

[b]Doctoral concentration in Management Science, Rutgers University, 100 Rockafeller Road, Piscataway, NJ 08854

[c]Doctoral program in Operations Research, Rutgers University, 100 Rockafeller Road, Piscataway, NJ 08854

Rutcor Research Report
RRR 3-2014, Aug 2014,

Multilevel Optimization Modeling for Risk-Averse Stochastic Programming

Jonathan Eckstein Deniz Eskandani Jingnan Fan

Abstract. Coherent risk measures have become a popular tool for incorporating risk aversion into stochastic optimization models. For dynamic models in which uncertainty is resolved at more than one stage, however, using coherent risk measures within a standard single-level optimization framework becomes problematic. To avoid severe time-consistency difficulties, the current state of the art is to employ risk measures of a specific nested form, which unfortunately have some undesirable and somewhat counterintuitive modeling properties. For one thing, this technique requires increasing risk aversion as risks and reward recede in time. Further, it produces objective functions that cannot be law invariant with respect to the total incurred costs and rewards, meaning that two solutions with identical probability distributions of final wealth may be assigned different levels of risk, and the nested form of the objective function cannot be simplified. These properties deter practical acceptance of such models, and are particularly undesirable for situations with close final time horizons.

This paper summarizes these issues and then presents an alternative multilevel optimization modeling approach that enforces a form of time consistency through constraints rather than by restricting the modeler's choice of objective function. This technique leads to *models* that are time-consistent even while using time-inconsistent risk measures, and can easily be formulated to be law invariant with respect to the final wealth if so desired. We argue that this approach should be the starting point for all multistage optimization modeling. When used with time-consistent objective functions, we show its multilevel optimization constraints become redundant and the associated models thus simplify 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.

Finally, we show that for a simple but reasonably realistic test application, that the kind of models we propose, while being drawn from an $\mathcal{NP}$-hard family and certainly more time consuming to solve than those obtained from the nested-objective approach, are readily solvable to global optimality using a standard commercial MILP solver. Therefore, there seems some promise of the modeling approach being useful despite its computational complexity properties.

Acknowledgements: This work was supported by a grant from the Air Force Office of Scientific Research, award FA9550-11-1-0164. The authors would like to thank Andrzej Ruszczyński for his assistance and support.

Note: Sections 4-7 of this report essentially duplicate material presented in the earlier reports RRR 7-2013 and (using different notation) RRR 17-2012. The remaining sections contain extensive new material and results.

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 is 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 problem may require assuming that one will take a course of action that is suboptimal for the recourse problem arising in one of the later-stage scenarios. In fact, as will be discussed in Section 9, Sections 6 and 7 will present examples of three-stage linear models, using the most commonly employed risk measures in a natural manner, in which optimizing the first-stage objective requires choosing what would later appear to be the *worst* feasible course of action in a particular second-stage 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, providing 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 also introduce some undesirable and in some cases counterintuitive model properties. In particular, they generally make it impossible to formulate models that conform to the principle of law invariance with respect to the final wealth (that is, the sum of all costs and rewards). For a single-period risk measure, the principle of law invariance states that two random variables having identical probability distributions should be assigned the same level of risk; all the standard examples of coherent single-period risk measures, such as the average-value-at-risk and mean-semidevation measures, are law invariant. The simplest way to extend the law-invariance principle to sequences of random variables is to require that the assessed risk levels of two sequences whose respective sums have identical distributions should also be identical. When modeling situations in which costs and rewards accumulate over long time spans, one may argue that this notion of law invariance is overly restrictive, but for close time horizons — up to a few months or perhaps years — it is a natural and intuitive requirement.

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 total outcome variables with identical distributions, lack of law invariance also implies that the somewhat complicated nested formulas required to express nested dynamic risk measures cannot be simplified into forms that decision makers might find easier to understand, unless one wishes to revert to classical expected-value or worse-outcome models.

For cases in which law invariance with respect to final wealth may not be required or desirable, such as with very long time horizons, it is in many cases natural to decrease one's level of risk aversion as costs and rewards become increasingly distant in time. Another drawback of the nested form of time-consistent dynamic risk measures is that it essentially forces the opposite preference: increasing risk aversion with time. Thus, there are several reasons to seek alternatives to formulating multistage stochastic programming problems using time-consistent dynamic risk measures.

One possible alternative would be to use risk measures that are not coherent. Kupper and Schachermayer (2009) have shown that a certain class of "entropic" risk measures are the only possible ones that are both time consistent and law invariant. While use of such risk measures definitely warrants further study, they are are not coherent, failing to meet the scale-invariance property of coherent risk measures (see below). While lack of scale invariance may not always be a serious concern and may in fact be desirable in some applications, the entropic risk measures form a somewhat restrictive family, and popular measures such as average value at risk (AVaR) do not fall into it.

This paper explores a different approach to modeling multistage stochastic optimization problems which permits the modeler, while maintaining a reasonable form of time consistency, to select any desired objective function for both the top-level problem and the recourse problems 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 familiar single-level optimization forms are simply convenient simplifications that occur when one uses a time-consistent objective. The advantage of the multilevel approach is that the modeler, while having greatly increased freedom in choosing objective functions, can still construct models that have a reasonable form of internal time consistency and are thus defensible from a modeling perspective. The approach can be used in any multistage decision setting with essentially any kind of risk measure, coherent or otherwise.

The drawback of our proposed technique is of course that multilevel optimization models tend to be more computationally 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). Sections 6 and 7 show that the simplest imaginable applications of the proposed modeling approach, while not necessarily leading to fully general bilevel linear programs, do indeed result in $\mathcal{NP}$-hard models. While this finding is natural to expect, its proof is nontrivial. The $\mathcal{NP}$-hardness result also does not necessarily imply that the proposed modeling approach is necessarily fruitless: in the practice of operations research, it is common to solve instances of $\mathcal{NP}$-hard problem classes, especially integer programs, essentially exactly, or to very tight tolerances. 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.

To confirm that $\mathcal{NP}$-hardness is not an insurmountable barrier to using our proposed modeling approach, Section 8 of this paper considers a simple class of three-stage supply-chain stochastic programming models that have been used in earlier work by Collado et al. (2012) and Asamov (2013). We created multilevel versions of these models which are both risk averse and law invariant with respect to the final wealth. We then reformulated them first as mathematical programs with equlibrium constraints (MPEC's) and then into a related form acceptable to standard mixed-integer programming (MIP) solvers. Using the widely-available GuRoBi solver (Gurobi Optimization 2014), we found that the resulting models were readily solvable to global optimality. Thus, the modeling techniques we suggest seem likely to be usable in at least some practial situations.

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, dynamic risk measure systems, and a formal general multistage stochastic problem setting patterned after Ruszczyński and Shapiro (2006a). In this context, it provides definitions of time consistency for systems of risk measures and stochastic optimization models, and proves a formal result relating the two.

Section 3 then demonstrates how multilevel optimization approaches can give the modeler much more flexibility in selecting objective functions than the time-consistent dynamic risk measure framework, while still yielding systems of models that have a reasonable form of time consistency as defined in Section 2. It also provides a formal result demonstrating how the multilevel framework reduces to system of conventional single-level optimization models when used with time-consistent system of risk measures.

Section 4 then specializes the very general model framework of Section 3 to the simplest possible setting in which the drawbacks of time-consistent dynamic risk measures could create difficulties; we call this model class *bilevel risk programming*. Section 5 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 they are $\mathcal{NP}$-hard when using mean-standard-deviation and average-value-at-risk risk measures, respectively. Section 8 then presents our preliminary computational experiments with bilevel risk programming problems, and Section 9 presents some concluding remarks, in particular regarding the kinds of computational modeling tools that would be required to make modeling techniques proposed here accessible to practitioners.

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 $\mathcal{Z}$ of random variables over $(\Omega, \mathcal{F})$, that is, $\mathcal{F}$-measurable functions $Z : \Omega \to \mathbb{R}$. We will take $\mathcal{Z} = \mathcal{L}^p(\Omega, \mathcal{F})$ 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 $\rho : \mathcal{Z} \to \mathbb{R} \cup \{+\infty\}$ to express their preferences between elements of $\mathcal{Z}$. We may think ρ as a function mapping a random variable to a scalar "certainty equivalent".

Definition 1 (SINGLE-STAGE COHERENT RISK MEASURES) *A mapping $\rho : \mathcal{Z} \to \mathbb{R} \cup \{+\infty\}$ is called a* coherent risk measure *(Artzner et al. 1999, Delbaen 2002, Rockafellar et al. 2006) if it has the following properties for any $Z, W \in \mathcal{Z}$:*

Monotonicity. *If $Z \leq W$ (that is, $Z(\omega) \leq W(\omega)$ for all $\omega \in \Omega$), then $\rho(Z) \leq \rho(W)$.*

Convexity. *For $\alpha \in [0,1]$, one has $\rho\big(\alpha Z + (1-\alpha)W\big) \leq \alpha\rho(Z) + (1-\alpha)\rho(W)$.*

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

Translation equivariance. *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. Two common choices are the mean-semideviation

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

and *average value at risk* (Acerbi 2002, Föllmer and Schied 2004, Rockafellar and Uryasev 2000, 2002, Uryasev and Rockafellar 2001), also called *conditional value at risk*,

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

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 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

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

For a discrete random variable, $\operatorname{ess\,sup} Z$ 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* in the sense 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 "recipe" 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 whose value changes 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 subpace 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 essentially a deterministic quantity, and thus Z_1 is deterministic. To compare such sequences of random variables, along with "tail" subsequences of the form $Z_\tau, Z_{\tau+1}, \ldots, Z_T$, where $1 \leq \tau \leq T$, we introduce the following notion:

Definition 2 *Given a probability space $(\Omega, \mathcal{F}, \mathrm{P})$, a filtration $\{\emptyset, \Omega\} = \mathcal{F}_1 \subset \mathcal{F}_2 \subset \cdots \subset \mathcal{F}_T = \mathcal{F}$, and corresponding spaces of random variables $\mathcal{Z}_1, \ldots, \mathcal{Z}_T$, a* dynamic risk measure system *is a sequence of functions $\{\rho_{\tau,T}\}_{\tau=1}^T$ such that $\rho_{\tau,T} : \mathcal{Z}_\tau \times \mathcal{Z}_{\tau+1} \times \cdots \times \mathcal{Z}_T \to \mathcal{Z}_\tau$.*

Each function $\rho_{\tau,T}$ in the risk-measure system converts a sequence of costs respectively measurable at times $\tau, \ldots, T$ to a "partial certainty equivalent" measurable at time τ. To save space throughout the remainder of this paper, we introduce the notation Z_{t_1,t_2} to stand for the concatenation $Z_{t_1}, \ldots, Z_{t_2}$ and $\mathcal{Z}_{t_1,t_2}$ to stand for $\mathcal{Z}_{t_1} \times \cdots \times \mathcal{Z}_{t_2}$.

For the case $\tau = 1$, we have that $\rho_{1,T}(Z_{1,T})$ must be $\mathcal{F}_1$-measurable, and so may be considered as being in $\mathbb{R}$. For all dynamic risk-measure systems considered in the rest of this paper, we make the following assumption, which any reasonably-chosen system of risk measures $\{\rho_{\tau,T}\}_{t=1}^T$ ought to obey:

Assumption 3 (FILTRATION COMPATIBILITY) *For any $\omega \in \Omega$ and $1 \leq \tau \leq t \leq T$, the value of $\rho_{\tau,T}(Z_{\tau,T})(\omega)$ does not depend on the values of $Z_t(\nu)$ for which $\nu \in \Omega$ is distinguishable from ω at time τ.*

2.3 Abstract single-level multistage 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 $X_t \in \mathcal{Z}^{n_t}$, which means we may consider X_t as 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 $\omega \in \Omega$, and that $Z_t \in \mathcal{Z}_t$. 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 $G_t(\omega, x_1, \ldots, x_{t-1}) = G_t(\omega', x_1, \ldots, x_{t-1})$ whenever $\omega, \omega' \in \Omega$ are indistinguishable at time t. We then propose the following very general constraint system:

$$X_t(\omega) \in G_t\big(\omega, X_1(\omega), \ldots, X_{t-1}(\omega)\big) \qquad \forall\, t = 1, \ldots, T,\ \omega \in \Omega. \tag{4}$$

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 over all $\omega \in \Omega$. 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 be influenced by prior decisions in situations that could possibly lead to ω at time t. Note that using the same notation as for Z_{t_1,t_2}, we may condense (4) to $X_t(\omega) \in G_t\big(\omega, X_{1,t-1}(\omega)\big)$ for $t = 1, \ldots, T$ and $\omega \in \Omega$.

Combining (4) with a dynamic risk-measure objective function, 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 $\omega \in \Omega$:

$$\begin{array}{lll} \min & \rho_{1,T}\big(c_1(X_1), \ldots, c_T(X_T)\big) & \\ \text{ST} & X_t(\omega) \in G_t\big(\omega, X_{1,t-1}(\omega)\big) & \forall\, t = 1, \ldots, T,\ \omega \in \Omega \\ & X_t \in \mathcal{Z}_t^{n_t} & \forall\, t = 1, \ldots, T. \end{array} \tag{5}$$

Letting $Z_t = c_t(X_t) \in \mathcal{Z}_t$ for all t, we could also denote the objective function as $\rho_{1,T}(Z_{1,T})$. 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 (5) 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 make the following simplifying assumption:

Assumption 4 (FINITE PROBABILITY SPACE**)** *Ω is a finite set and (without further loss of generality) each "atom" $\omega \in \Omega$ has $\mathrm{P}\big\{\{\omega\}\big\} > 0$.*

This 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 $\mathcal{F}_\tau$ such that for each $E \in \mathcal{E}_\tau$, the only strict subset of E in $\mathcal{F}_\tau$ is $\emptyset$. Each $E \in \mathcal{E}_\tau$ corresponds to a distinguishable scenario in stage τ. Note that any $\mathcal{F}_\tau$-measurable function, including $\rho_{\tau,T}(Z_{\tau,T})$, is necessarily constant across each $E \in \mathcal{E}_\tau$. We will therefore write $\rho_{\tau,T}(Z_{\tau,T})(E)$ for the common value of $\rho_{\tau,T}(Z_{\tau,T})(\omega)$ for all $\omega \in E$. By Assumption 3, this value depends only on the values of $Z_t(\omega)$ for which $\omega \in E$. For any $t = \tau, \dots, T$, we let $Z_t(E)$ denote the function $Z_t : \Omega \to \mathbb{R}$ with its domain restricted to E; the values of this function are precisely those which can influence $\rho_{\tau,T}(Z_{\tau,T})(E)$. We will use a restriction notation similar to $Z_t(E)$ for any random variable.

Now consider the decision maker at time step τ, having taken the decisions $X_{1,\tau-1}$ at the prior time steps, and knowing that elementary event $E \in \mathcal{E}_\tau$ has occurred. Assuming that the risk measure $\rho_{\tau,T}$, which must be constant-valued across E, characterizes our risk preferences at this juncture, we may try to optimize it with respect to the remaining decisions to be made, the only relevant portions of which, by Assumption 3, are constituted by $X_\tau(E), \dots, X_T(E)$, or more compactly $X_{\tau,T}(E)$. Formally, we make the following definition:

Definition 5 *Given a dynamic risk-measure system $\{\rho_{\tau,T}\}_{t=1}^T$ and a problem of the form (5), the corresponding* recourse problem *in state $E \in \mathcal{E}_\tau$ at time $\tau \in \{1, \dots, T\}$ is*

$$\begin{array}{lll} \min & \rho_{\tau,T}\Big(c_\tau\big(X_\tau(E)\big), \dots, c_T\big(X_T(E)\big)\Big)(E) & \\ \mathrm{ST} & X_t(\omega) \in G_t\big(\omega, X_{1,t-1}(\omega)\big) & \forall\, t = \tau, \dots, T,\ \omega \in E \\ & X_t(E) \in \mathcal{Z}_t^{n_t} \mid E & \forall\, t = \tau, \dots, T, \end{array} \tag{6}$$

where $\mathcal{Z}_t^{n_t} \mid E$ denotes the space of functions in Z_t, restricted to the domain E.

Problem (6) may be denoted $R_\tau^E(X_{1,\tau-1})$. Note that the decision variables of the recourse problem (6) are $X_{\tau,T}(E)$, and the top-level problem (5) may be denoted $R_1^\Omega(X_{1,0})$ or equivalently $R_1^\Omega(\,)$, where we use the convention that X_{t_1,t_2} denotes a zero-length vector if $t_1 > t_2$.

2.5 Law invariance and an intuitive choice of dynamic risk measure system

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

$$\rho_{\tau,T}(Z_{\tau,T})(E) = \mathbb{E}[Z_\tau + \cdots + Z_T \mid E]. \tag{7}$$

That is, the objective function for the event-E recourse problem is the expected value of the sum of the remaining costs, conditioned on the information known at E. For $\tau = 1$, the only elementary event in $\mathcal{E}_1$ is Ω, so this formula reduces to $\rho_{1,T}(Z_{1,T}) = \mathbb{E}[Z_1 + \cdots + Z_T]$. One natural way to attempt to proceed is to generalize this approach to other distribution-based rules for generating risk measures. Specifically, for each elementary event $E \in \mathcal{E}_\tau$ at time τ, we can evaluate the distribution of $Z_\tau + \cdots + Z_T$ conditioned on the event E, and then apply the distribution-based rule to obtain a conditional certainty equivalent. If ρ denotes the general distribution-based recipe for deriving a risk measure, we may write this approach as

$$\rho_{\tau,T}(Z_{\tau,T}) = \rho(Z_\tau + \cdots + Z_T \mid E). \tag{8}$$

For example, if we choose the MSD_γ rule, then we have

$$\rho_{\tau,T}(Z_{\tau,T})(E) = \mathbb{E}\left[\sum_{t=\tau}^{T} Z_t \;\middle|\; E\right] + \gamma\,\mathbb{E}\left[\left[\sum_{t=\tau}^{T} Z_t - \mathbb{E}\left[\sum_{t=\tau}^{T} Z_t \;\middle|\; E\right]\right]_+ \;\middle|\; E\right]. \tag{9}$$

For the AVaR_α rule we obtain

$$\rho_{\tau,T}(Z_{\tau,T})(E) = \frac{1}{\alpha}\int_{1-\alpha}^{1} F^{-1}_{Z_\tau+\cdots+Z_T \mid E}(\nu)\,d\nu, \tag{10}$$

where $F_{Z_\tau+\cdots+Z_T \mid E}$ 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. Next, we define a simple notion of law invariance for dynamic risk-measure systems:

Definition 6 (LAW INVARIANCE) *A dynamic risk-measure system $\rho_{\tau,T}$ is* law invariant *if, given any $\tau \in \{1,\ldots,T\}$, an elementary event $E \in \mathcal{E}_\tau$, and two sequences of random variables $Z_{\tau,T} = (Z_\tau,\ldots,Z_T)$ and $W_{\tau,T} = (W_\tau,\ldots,W_T)$ such that $Z_\tau + \cdots + Z_T$ and $W_\tau + \cdots + W_T$ have the same probability distribution conditional on E, one must have $\rho_{\tau,T}(Z_{\tau,T})(E) = \rho_{\tau,T}(W_{\tau,T})(E)$.*

For $T = 1$ (and thus $E = \Omega$), this definition coincides with the standard single-period notion of law invariance. For modeling situations in which the final time horizon T is fairly close in terms of real elapsed time, objective functions based on the distribution of the final wealth $Z_1 + \cdots + Z_T$, conditioned on any information revealed up to the current time step, seem inherently reasonable, and the kind of law invariance described in the definition thus appears desirable. For costs (or rewards) spread out over very long spans of time, such as decades, it may be argued that the conditional distributions of the sums $Z_\tau + \cdots + Z_T$ might not model all relevant aspects of a decision-maker's preferences; for example, one might be less sensitive to variability of cash flows far in the future than those that are imminent. In this

case, one could argue that a weaker form of law invariance might be more natural, but this paper will focus primarily on the former situation in which the time horizon is not overly distant.

The risk-measure systems given by (8), including the special cases (9) and (10), are clearly law invariant in the sense of Definition 6. 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 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 7 (RISK-MEASURE TIME CONSISTENCY (RUSZCZYŃSKI 2010))
The system of dynamic risk measures $\{\rho_{t,T}\}_{t=1}^{T}$ *is* time consistent *if for each* τ, θ *such that* $1 \leq \tau < \theta \leq T$, *all* $Z_{\tau,\theta-1} \in \mathcal{Z}_{\tau,\theta-1}$ *and all* $Z_{\theta,T}, W_{\theta,T} \in \mathcal{Z}_{\theta,T}$ *with* $\rho_{\theta,T}(Z_{\theta,T}) \leq \rho_{\theta,T}(W_{\theta,T})$, *we have*

$$\rho_{\tau,T}(Z_{\tau,\theta-1}, Z_{\theta,T}) \leq \rho_{\tau,T}(Z_{\tau,\theta-1}, W_{\theta,T}),$$

where all the inequalities are interpreted pointwise.

This definition says that if the $Z_{\theta,T}$ is preferable to $W_{\theta,T}$ 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 "renege" on plans made at earlier stages. To understand this phenomenon, we begin by defining a notion of time consistency for a stochastic program and its recourse models. This definition has a similar inspiration to the notion of time consistency for dynamic risk measures introduced by Boda and Filar (2006), but is far less restricted in its applicability. Before giving the main definition, we first define some notation:

Definition 8 *Given* $1 \leq \tau < \tau' \leq T$, $E \in \mathcal{E}_\tau$, $E' \in \mathcal{E}_{\tau'}$, $E \supseteq E' \neq \emptyset$, *and two sequences of random variables* $X_{\tau,T}(E)$ *and* $X'_{\tau',T}(E')$ *with*

$$\begin{aligned} X_t(E) &\in \mathcal{Z}^{n_t} \mid E && t = \tau, \ldots, T \\ X'_t(E') &\in \mathcal{Z}^{n_t} \mid E' && t = \tau', \ldots, T, \end{aligned}$$

we define $X_{\tau,T}(E) \lhd X'_{\tau',T}(E')$ *to be the random variable sequence*

$$\widetilde{X}_{\tau,T}(E) = \big(\widetilde{X}_\tau(E), \ldots, \widetilde{X}_T(E)\big) \in (\mathcal{Z}^{n_\tau} \mid E) \times \cdots \times (\mathcal{Z}^{n_T} \mid E)$$

given by

$$\widetilde{X}_t(\omega) = \begin{cases} X'_t(\omega), & \text{if } t \geq \tau' \text{ and } \omega \in E', \\ X_t(\omega), & \text{otherwise.} \end{cases} \qquad \forall\, t = \tau, \ldots, T,\ \omega \in E. \tag{11}$$

In short, this substitution notation specifies that the decision vector $X_{\tau,T}(E) \lhd X'_{\tau',T}(E')$ consists of $X_{\tau,T}(E)$, with each of its element corresponding to an element of $X'_{\tau',T}(E')$ replaced by that element.

Definition 9 (MODEL TIME CONSISTENCY) *The system of models R given by* (5) *and* (6) *is called* weakly time consistent *if, for every $\tau, \tau', E, E', X_{1,T}$ with the properties that $1 \leq \tau < \tau' \leq T$, $E \in \mathcal{E}_\tau$, $E' \in \mathcal{E}_{\tau'}$, $E \supseteq E' \neq \emptyset$, and $X_{1,T}$ is feasible for* (5)*, one has for every optimal solution $X^*_{\tau,T}(E)$ of $R^E_\tau\big(X_{1,\tau-1}(E)\big)$ that there exists some optimal solution $X^{**}_{\tau',T}(E')$ of $R^{E'}_{\tau'}\big(X_{1,\tau-1}(E'), X^*_{\tau,\tau'-1}(E')\big)$ such that $X^*_{\tau,T}(E) \lhd X^{**}_{\tau',T}(E')$ remains optimal for $R^E_\tau\big(X_{1,\tau-1}(E)\big)$. The model system is called* strongly time consistent *if the same condition holds for all possible optimal solutions $X^{**}_{\tau',T}(E')$ of $R^{E'}_{\tau'}\big(X_{1,\tau-1}(E'), X^*_{\tau,\tau'-1}(E')\big)$.*

Of particular interest is the case that $\tau = 1$, in which case it follows that $E = \Omega$, meaning that $R^E_\tau\big(X_{1,\tau-1}(E)\big)$ is simply the top-level problem (5), and $R^{E'}_{\tau'}\big(X_{1,\tau-1}(E'), X^*_{\tau,\tau'-1}(E')\big)$ reduces to the recourse problem $R^{E'}_{\tau'}\big(X^*_{1,\tau'-1}(E')\big)$. Essentially, weak time consistency of a model system means that in the context of a recourse problem, one does not have a positive incentive to deviate from plans that appeared optimal from the standpoint of earlier stages. Strong time consistency, which clearly implies weak time consistency, means that no optimal recourse problem solution can invalidate earlier decisions. If a model system fails to conform to weak time consistency, its use is hard to defend in practice: the optimality of decisions at one time will in general be predicated on taking courses of action that could be interpreted as strictly suboptimal at a later time.

We now demonstrate that time consistency of a risk measure system implies time consistency of the associated model system:

Proposition 10 *If the sysem of dynamic risk measures $\{\rho_{\tau,T}\}_{\tau=1}^T$ is time consistent, then the model system R given by* (5) *and* (6) *is strongly time consistent.*

Proof. Consider any τ, τ', E, E', and $X_{1,T}$ with $1 \leq \tau < \tau' \leq T$, $E \in \mathcal{E}_\tau$, $E' \in \mathcal{E}_{\tau'}$, $E \supseteq E' \neq \emptyset$, and $X_{1,T}$ being feasible for (5). Let Q denote the problem $R^E_\tau\big(X_{1,\tau-1}(E)\big)$, and let $X^*_{\tau,T}(E)$ be any optimal solution of Q. Next, let Q' denote the problem $R^{E'}_{\tau'}\big(X_{1,\tau-1}(E'), X^*_{\tau,\tau'-1}(E')\big)$ and suppose that $X^{**}_{\tau',T}(E')$ is some optimal solution of Q'. Furthermore, let $\widetilde{X}_{\tau,T}(E) = X^*_{\tau,T}(E) \lhd X^{**}_{\tau',T}(E')$, where the "$\lhd$" operation is defined as in (11). The proposition will be established if we can show that $\widetilde{X}_{\tau,T}(E)$ is also optimal for Q.

First, we seek to establish that $\widetilde{X}_{\tau,T}(E)$ is feasible for Q, that is,

$$\widetilde{X}_t(\omega) \in G_t\big(\omega, X_{1,\tau-1}(\omega), \widetilde{X}_{\tau,t-1}(\omega)\big) \qquad \forall\, t = \tau, \ldots, T,\ \omega \in E \tag{12}$$

$$\widetilde{X}_t(E) \in \mathcal{Z}_t^{n_t} \mid E \qquad \forall\, t = \tau, \ldots, T. \tag{13}$$

For $t < \tau'$, we have $\widetilde{X}_t(\omega) = X^*_t(\omega)$ for all $\omega \in E$ by construction, and both (12) and (13) hold since $X^*_{\tau,T}(E)$ is feasible for Q. For $t \geq \tau'$, we distinguish between the cases $\omega \in E'$ and $\omega \in E \backslash E'$. In the latter case, we again have $\widetilde{X}_t(\omega) = X^*_t(\omega)$, and so (12) remains true. In the former case, we have $\widetilde{X}_t(\omega) = X^{**}_t(\omega)$, but since $X^{**}_{\tau',T}(E')$ in an optimal solution of Q', we have

$$X^{**}_t(\omega) \in G_t\big(\omega, X_{1,\tau-1}(\omega), X^*_{\tau,\tau'-1}(\omega), X^{**}_{\tau',t-1}(\omega)\big),$$

but in the case $t \geq \tau'$, $\omega \in E'$, this reduces to

$$\begin{aligned}\widetilde{X}_t(\omega) &\in G_t\big(\omega, X_{1,\tau-1}(\omega), \widetilde{X}_{\tau,\tau'-1}(\omega), \widetilde{X}_{\tau,t-1}(\omega)\big)\\ &= G_t\big(\omega, X_{1,\tau-1}(\omega), \widetilde{X}_{\tau,t-1}(\omega)\big),\end{aligned}$$

which is exactly the condition in (12). Therefore, (12) holds in all cases. It remains to verify that (13) holds when $t \geq \tau'$, that is, that $\widetilde{X}_t(E)$ is constant across every elementary event $Y \in \mathcal{E}_t$ that is contained in E. Since $E' \in \mathcal{E}_{\tau'} \subseteq \mathcal{F}_{\tau'} \subseteq \mathcal{F}_t$, the rules of σ-algebras imply that each elementary event $Y \in \mathcal{E}_t$ is either disjoint from E' or contained in E'. In the case that $Y \subseteq E'$, we have $\widetilde{X}_t(\omega) = X^{**}_t(\omega)$ for all $\omega \in Y$, but all these values are must be identical because $X^{**}_t(E')$ is part of an optimal solution of Q', and thus satisfies the constraint $X^{**}_t(E') \in \mathcal{Z}^{n_t}_t(E')$. In the case $Y \cap E' = \emptyset$, we may make a similar argument based on the feasibility of $X^*_t(E)$ for Q. Hence, we have established the feasibility of $\widetilde{X}_{\tau,T}(E)$ for Q.

Next, let $Z^*_t(E) = c_t(X^*_t(E))$ and $\widetilde{Z}_t(E) = c_t(\widetilde{X}_t(E))$ for $t = \tau, \ldots, T$, and observe that for all $\omega \in E'$ we must have

$$\begin{aligned}\rho_{\tau',T}(\widetilde{Z}_{\tau',T})(\omega) &= \rho_{\tau',T}\Big(c_{\tau'}\big(\widetilde{X}_{\tau'}(E)\big), \ldots, c_T\big(\widetilde{X}_T(E)\big)\Big)(\omega)\\ &= \rho_{\tau',T}\Big(c_{\tau'}\big(X^{**}_{\tau'}(E)\big), \ldots, c_T\big(X^{**}_T(E)\big)\Big)(\omega)\\ &\leq \rho_{\tau',T}(Z^*_{\tau',T})(\omega) \qquad (14)\end{aligned}$$

by the optimality of $X^{**}_{\tau',T}(E)$ for Q' (note also that since $\rho_{\tau',T}$ yields a $\mathcal{F}_{\tau'}$-measurable quantity, it is in fact constant across $\omega \in E'$).

Next, we observe that for $\omega \in E \backslash E'$, we have

$$\rho_{\tau',T}(\widetilde{Z}_{\tau',T})(\omega) = \rho_{\tau',T}(Z^*_{\tau',T})(\omega) \qquad (15)$$

because $\widetilde{X}_t(\omega)$ and $X^*_t(\omega)$ and hence $\widetilde{Z}_t(\omega)$ and $Z^*_t(\omega)$ are identical for $t = \tau', \ldots, T$. Combining (14) and (15), we conclude that

$$\rho_{\tau',T}(\widetilde{Z}_{\tau',T})(\omega) \leq \rho_{\tau',T}(Z^*_{\tau',T})(\omega) \qquad \forall\, \omega \in E. \qquad (16)$$

Next, we note that for all $t = \tau, \ldots, \tau' - 1$, we have $\widetilde{Z}_t(\omega) = Z^*_t(\omega)$ for all $\omega \in E$ because $\widetilde{X}_t(\omega) = X^*_t(\omega)$ by construction in such cases.

Thus, we are in the situation that $\widetilde{Z}_{\tau,\tau'-1}(E)$ and $Z^*_{\tau,\tau'-1}(E)$ are identical, after which $\widetilde{Z}_{\tau',T}(E)$ is preferable to $Z^*_{\tau',T}(E)$ in the sense of (16). Because of the assumption that the system $\{\rho_{t,T}\}_{t=1}^T$ is time-consistent, we can then assert that $\rho_{\tau,T}(\widetilde{Z}_{\tau,T})(E) \leq \rho_{\tau,T}(Z^*_{\tau,T})(E)$ (note that both of these quantities are $\mathcal{F}_\tau$-measurable and hence constant across E). This may be seen by extending each of $\widetilde{Z}_t(E)$ and $Z^*_t(E)$, for $t = \tau, \ldots, T$, in any identical $\mathcal{F}_t$-measurable way, to all of $\omega \in \Omega$, rather than just $\omega \in E$, and then applying Definition 7.

Finally, since we have already established that $\widetilde{X}_{\tau,T}(E)$ is feasible for Q, and $X^*_{\tau,T}(E)$ is optimal for the same problem, it follows from $\rho_{\tau,T}(\widetilde{Z}_{\tau,T})(E) \leq \rho_{\tau,T}(Z^*_{\tau,T})(E)$ that $\widetilde{X}_{\tau,T}(E)$ is also optimal for Q. □

Thus, using a time-consistent risk measure system guarantees that the general system of recourse problems given by (5) and (6) will be time consistent.

Unfortunately, intuitive systems of risk measures defined by (8) are not time-consistent, except if ρ is the expected-value and worst-outcome (ess sup) risk-measure. This inconvenient fact can be seen as the consequence of two prior results: first, Ruszczyński (2010, p. 241) shows that translation-invariant time-consistent risk measure systems necessarily have the nested structure

$$\rho_{\tau,T}(Z_\tau, \ldots, Z_T) = Z_\tau + \rho_\tau\Big(Z_{\tau+1} + \rho_{\tau+1}\big(Z_{\tau+2} + \cdots \rho_{T-1}(Z_T)\cdots\big)\Big), \tag{17}$$

where $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$, $t = 1, \ldots, T-1$, are *one-step* risk measures. Furthermore, in order for $\rho_{\tau,T}$ to satisfy conditions analogous to coherency for a single-stage measure, one should further assume that each ρ_t obeys the following axioms generalizing Definition 1:

Definition 11 (ONE-STEP COHERENT RISK MEASURES) *A one-step risk measure function $\rho : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$ is called* coherent *(Ruszczyński 2010) if it has the following properties for any $Z, W \in \mathcal{Z}_{t+1}$, where all inequalities are interpreted pointwise:*

Monotonicity. *If $Z \leq W$, then $\rho_t(Z) \leq \rho_t(W)$.*

Convexity. *For $\alpha \in [0,1]$, one has $\rho_t\big(\alpha Z + (1-\alpha)W\big) \leq \alpha\rho_t(Z) + (1-\alpha)\rho_t(W)$.*

Positive homogeneity. *For $\alpha \geq 0$, then $\rho_t(\alpha Z) = \alpha\rho_t(Z)$.*

Translation equivariance. *For any $U \in \mathcal{Z}_t$, we have $\rho_t(Z + U) = \rho_t(Z) + U$.*

Unfortunately, Shapiro (2012) shows that dynamic risk measures of the form specified by (17) and Definition 11 cannot be law-invariant unless all the one-step risk measures are conditional expectations, or they are all worst-outcome operators, meaning that either $\rho_{1,T}(Z_{1,T}) = \mathbb{E}[Z_1 + \cdots Z_T]$ or $\rho_{1,T} = \text{ess sup}(Z_1 + \cdots Z_T)$, respectively. Kupper and Schachermayer (2009) show a similar result.

Another implication of the result of Shapiro (2012) is that time-consistent dynamic risk measures of the form (17), other than expected value or worst outcome, cannot be simplified into functions of the distribution of the sum of all costs and rewards $Z_1 + \cdots + Z_T$. 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 (17) cannot be simplified to a 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 (17) 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 (5) and (6). 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 property of law invariance with respect to the final wealth, which is desirable in at least some applications. This situation constitutes a significant barrier to wide adoption of such models.

3 Multilevel Optimization Modeling for Multistage Decision-Making

We will now demonstrate that more general way to assure that systems of stochastic programming models are time consistent is to restrict the structure of their constraints, rather than the structure of their objective functions. In particular, we propose to augment the constraints of an arbitrary model system of the form (5)-(6) to include multilevel optimization conditions which directly guarantee weak time consistency. This technique makes it possible to construct time-consistent systems of models without using a time-consistent system of risk measures.

We develop this system of models as follows: first, we select any system of dynamic risk measures $\{\rho_{t,T}\}_{t=1}^T$ that we consider to model 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 (8)-(10). We base this system, which we call $\overline{R}$, on $\{\rho_{t,T}\}_{t=1}^T$ and an arbitrary system of constraints G of the same form as in the previous section. However, we augment these constraints with additional multilevel optimization constraints.

To construct a system of recourse models, we proceed recursively backwards in time starting from time step $T-1$; since the decision model at time step T is deterministic, it is sufficient to start the recursion 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 $R^E_{T-1}(X_{1,T-2})$, which is

$$\begin{array}{lll} \min & \rho_{T-1,T}\Big(c_{T-1}\big(X_{T-1}(E)\big), c_T\big(X_T(E)\big)\Big)(E) & \\ \mathrm{ST} & X_{T-1}(\omega) \in G_{T-1}\big(\omega, X_{1,T-2}(\omega)\big) & \forall\, \omega \in E \\ & X_T(\omega) \in G_T\big(\omega, X_{1,T-1}(\omega)\big) & \forall\, \omega \in E \\ & X_{T-1}(E) \in \mathcal{Z}_{T-1} \mid E. & \end{array} \tag{18}$$

Here, the last constraint just means that $X_{T-1}(\omega)$ should be identical for all $\omega \in E$. We will also denote this problem by $\overline{R}^E_{T-1}(X_{1,T-2})$ and use the notation $\mathcal{X}^*_{T-1,E}(X_{1,T-2})$ 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

$$\begin{aligned} &X_{T-2}(\omega) \in G_{T-2}\big(\omega, X_{1,T-3}(\omega)\big) \quad \forall\, \omega \in E, \\ &X_{T-2}(E) \in \mathcal{Z}_{T-2} \mid E. \end{aligned}$$

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

$$\begin{array}{lll} \min & \rho_{T-2,T}\Big(c_{T-2}\big(X_{T-2}(E)\big), c_{T-1}\big(X_{T-1}(E)\big), c_T\big(X_T(E)\big)\Big)(E) & \\ \text{ST} & X_{T-2}(\omega) \in G_{T-2}\big(\omega, X_{1,T-3}(\omega)\big) & \forall\, \omega \in E \\ & X_{T-2}(E) \in \mathcal{Z}_{T-2} \mid E & \\ & \big(X_{T-1}(D), X_T(D)\big) \in \mathcal{X}^*_{T-1,D}(X_{1,T-2}) & \forall\, D \in \mathcal{E}_{T-1} : D \subseteq E. \end{array} \tag{19}$$

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

$$\begin{array}{lll} \min & \rho_{\tau,T}\Big(c_\tau\big(X_\tau(E)\big), \ldots, c_T\big(X_T(E)\big)\Big)(E) & \\ \text{ST} & X_\tau(\omega) \in G_\tau\big(\omega, X_{1,\tau-1}(\omega)\big) & \forall\, \omega \in E \\ & X_\tau(E) \in \mathcal{Z}_\tau \mid E & \\ & X_{\tau+1,T}(D) \in \mathcal{X}^*_{\tau+1,D}(X_{1,\tau}) & \forall\, D \in \mathcal{E}_{\tau+1} : D \subseteq E, \end{array} \tag{20}$$

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

$$\begin{array}{lll} \min & \rho_{1,T}\big(c_1(X_1), \ldots, c_T(X_T)\big) & \\ \text{ST} & X_1 \in \mathcal{X}_1 & \\ & X_{2,T}(D) \in \mathcal{X}^*_{2,D}(X_1) & \forall\, D \in \mathcal{E}_2. \end{array} \tag{21}$$

Using the notation already defined, this problem may be referred to as $\overline{R}^\Omega_1(\,)$ or equivalently $\overline{R}^\Omega(X_{1,0})$.

The weak time consistency of the system of models $\overline{R}$ given by (20)-(21) follows essentially immediately from their construction:

Proposition 12 *The system of models* $\overline{R}$ *given by* (18)-(20)-(21) *is weakly time consistent regardless of the choice of* $\{\rho_{\tau,T}\}_{\tau=1}^{T}$.

Proof. Consider any $\tau, \tau', E, E', X_{1,T}$ with $1 \leq \tau < \tau' \leq T$, $E \in \mathcal{E}_\tau$, $E' \in \mathcal{E}_{\tau'}$, $E \supseteq E' \neq \emptyset$, and $X_{1,T}$ being feasible for (21). Suppose that $X^*_{\tau,T}(E)$ is any solution to $\overline{R}_\tau^E(X_{1,\tau-1})$. The proposition will be established if we can show there exists some optimal solution $X^{**}_{\tau',T}(E')$ of $\overline{R}_{\tau'}^{E'}(X_{1,\tau-1}, X^*_{\tau,\tau'-1})$ such that the substituted solution $X^*_{\tau,T}(E) \lhd X^{**}_{\tau',T}(E')$ is optimal for $\overline{R}_\tau^E(X_{1,\tau-1})$.

For $t = \tau+1, \ldots, \tau'$, let E_t be the unique elementary event in $\mathcal{E}_t$ such that one has $E_t \supseteq E'$ (meaning that $E_{\tau'} = E'$). Now, taking $D = E_{\tau+1}$ in last constraint in (20) means that $X^*_{\tau+1,T}(E_{\tau+1})$ is optimal for the problem $\overline{R}_{\tau+1}^{E_{\tau+1}}(X_{1,\tau-1}, X^*_\tau)$. Next, the corresponding constraint in problem $\overline{R}_{\tau+1}^{E_{\tau+1}}(X_{1,\tau-1}, X^*_\tau)$ with $D = E_{\tau+2}$ implies that $X^*_{\tau+2,T}(E_{\tau+2})$ must be optimal for the problem $\overline{R}_{\tau+2}^{E_{\tau+2}}(X_{1,\tau-1}, X^*_{\tau,\tau+1})$. We proceed inductively in this manner for $t = \tau+1, \ldots, \tau'$, obtaining at each step that $X^*_{t,T}(E_t)$ must be optimal for $\overline{R}^{E_t}(X_{1,\tau-1}, X^*_{\tau,t-1})$. Once we reach $t = \tau'$, we have $E_t = E'$, and so we obtain that $X^*_{\tau',T}(E')$ is optimal for $\overline{R}_{\tau'}^{E'}(X_{1,\tau-1}, X^*_{\tau,\tau'-1})$. Therefore, we may choose $X^{**}_{\tau',T}(E') = X^*_{\tau',T}(E')$ and thus obtain that $X^*_{\tau,T}(E) \lhd X^{**}_{\tau',T}(E') = X^*_{\tau,T}(E)$ is optimal for $\overline{R}_\tau^E(X_{1,\tau-1})$. □

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 $X_{2,T}$ by optimally solving the problem (21) are necessarily optimal for any recourse problem that might be encountered. Note that the time consistency is not strong, meaning that 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 Proposition 10. 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 $\left\{\overline{R}_t^E(\,\cdot\,) \mid t = 1, \ldots, T,\ E \in \mathcal{E}_t\right\}$ given by (18)-(20)-(21) is time-consistent in the sense that there should never be any positive incentive to deviate from prior plans. 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 (5)-(6) involve only a single level of optimization (unless there are multilevel optimization constraints embedded in the abstract sets $G_t(\,\cdot\,)$).

We now show that if we apply such a system of models with a time-consistent objective system of objective functions, the multilevel optimization constraints essentially become redundant, and the model system reduces to the more familiar form (5)-(6). First, however, we need a minor definition:

Definition 13 (STRICT MONOTONICITY**)** *Under Assumption 4, a one-step risk measure* $\rho_t : \mathcal{Z}_{t+1} \to \mathcal{Z}_t$ *is* strictly monotonic *if, given an elementary event* $E \in \mathcal{E}_t$, *two* $\mathcal{F}_{t+1}$*-measurable random variables* Z *and* W *with* $Z(\omega) \leq W(\omega)$ *for all* $\omega \in E$ *and* $Z(\nu) < W(\nu)$ *for some* $\nu \in E$, *then one has* $\rho_t(Z)(E) < \rho_t(W)(E)$.

Essentially, this definition says that a strictly monotonic risk measure cannot completely "ignore" an improvement in a particular later outcome: if Z is at least as good as W in all cases and strictly better in at least one possible case, then Z should be strictly preferred to W. The expected-value and mean-semideviation risk measures have this property, but the worst-outcome and AVaR measures do not.

Proposition 14 *Suppose the risk measure system $\{\rho_{t,T}\}_{t=1}^{T}$ is time consistent and monotonic. Then, for every $t = 1, \ldots, T-2$, elementary event $E \in \mathcal{E}_t$, and decision vectors $X_{1,t-1}$, all optimal solutions of $\overline{R}_t^E(X_{1,t-1})$ are also optimal for $R_t^E(X_{1,t-1})$. If the risk measure system is of the form* (17) *with each single-step risk measure ρ_τ strictly monotonic, then the optimal solutions sets of the two models are identical.*

Proof. The proof is by induction on t, starting with $t = T-2$. The base case is immediate, since $\overline{R}_{T-2}^E(X_{1,T-3})$ is defined to be identical to $R_{T-2}^E(X_{1,T-3})$, so their solutions sets are identical regardless of the form of $\{\rho_{t,T}\}_{t=1}^{T}$.

Next, we suppose the claim has been established for $t = \tau + 1$, and seek to prove it for $t = \tau$. Let arbitrary values of $X_{1,\tau-1}$ and an arbitrary choice of $E \in \mathcal{E}_\tau$ be given. Consider any optimal solution $X_{\tau,T}^*$ of $\overline{R}_{T-2}^E(X_{1,\tau-1})$. Considering any elementary event $D \in \mathcal{E}_{\tau+1}$ with $D \subseteq E$, the last constraint in (20) implies that

$$X_{\tau+1,T}^*(D) \in \mathcal{X}_{\tau+1,D}^*(X_{1,\tau-1}, X_\tau^*).$$

Since $\mathcal{X}_{\tau+1,D}^*(X_{1,\tau-1}, X_\tau^*)$ denotes the set of optimal solutions of $\overline{R}_{\tau+1}^D(X_{1,\tau-1}, X_\tau^*)$, the induction hypothesis then implies that $X_{\tau+1,T}^*(D)$ is optimal and thus feasible for $R_{\tau+1}^D(X_{1,\tau-1}, X_\tau^*)$, meaning that

$$X_\theta(\omega) \in G_\theta\big(\omega, X_{1,\tau-1}(\omega), X_{\tau,\theta}^*(\omega)\big) \qquad \forall\, \theta = \tau+1, \ldots, T,\ \omega \in D \tag{22}$$

$$X_\theta(D) \in \mathcal{Z}_\theta^{n_\theta} \mid D \qquad \forall\, \theta = \tau+1, \ldots, T. \tag{23}$$

Now, since each $\omega \in E$ must be a member of some $D \in \mathcal{E}_{\tau+1}$ with $D \subseteq E$, we may deduce from (22) that

$$X_\theta(\omega) \in G_\theta\big(\omega, X_{1,\tau-1}(\omega), X_{\tau,\theta}^*(\omega)\big) \qquad \forall\, \theta = \tau+1, \ldots, T,\ \omega \in E. \tag{24}$$

Furthermore, for any $\theta = \tau+1, \ldots, T$, every elementary event $E_\theta \in \mathcal{E}_\theta$ that is contained in E must be a subset of some $D \in \mathcal{E}_{\tau+1}$ so it follows from (23) that $X_\theta(\omega)$ is constant across $\omega \in E_\theta$ and therefore

$$X_\theta(E) \in \mathcal{Z}_\theta^{n_\theta} \mid E \qquad \forall\, \theta = \tau+1, \ldots, T. \tag{25}$$

Also, the first two constraints in (20) guarantee that we have both $X_\tau(E) \in \mathcal{Z}_\tau^{n_\tau} \mid E$ and $X_\tau(\omega) \in G_\tau(\omega, X_{1,\tau-1})$ for all $\omega \in E$. Combining these observations with (24) and (25) establishes that $X_{\tau,T}^*$ is feasible for $R_\tau^E(X_1, \ldots, X_{\tau-1})$.

Next, consider any feasible solution $\widehat{X}_{\tau,T}(E)$ to $R_\tau^E(X_{1,\tau-1})$. Next, define $\widetilde{X}_{\tau,T}(E)$ by taking $\widetilde{X}_\tau(E) = \widehat{X}_\tau(E)$ and letting $\widetilde{X}_{\tau+1,T}(D)$ be any element of $\mathcal{X}_{\tau+1,D}^*(X_{1,\tau-1}, \widehat{X}_\tau)$ for each

$D \in \mathcal{E}_{\tau+1}$ with $D \subseteq E$. By the induction hypothesis, each $\widetilde{X}_{\tau+1,T}(D)$ must optimal for the corresponding recourse problem $R^D_{\tau+1}(X_{1,\tau-1}, \widetilde{X}_\tau) = R^D_{\tau+1}(X_{1,\tau-1}, \widehat{X}_\tau)$, so we have

$$\rho_{\tau+1,T}\big(c_{\tau+1}(\widetilde{X}_{\tau+1}), \ldots, c_T(\widetilde{X}_T)\big)(D) \leq \rho_{\tau+1,T}\big(c_{\tau+1}(\widehat{X}_{\tau+1}), \ldots, c_T(\widehat{X}_T)\big)(D).$$

Considering that every $\omega \in E$ is a member of some $D \in \mathcal{E}_{\tau+1}$ with $D \subseteq E$, we conclude that

$$\rho_{\tau+1,T}\big(c_{\tau+1}(\widetilde{X}_{\tau+1}), \ldots, c_T(\widetilde{X}_T)\big) \leq \rho_{\tau+1,T}\big(c_{\tau+1}(\widehat{X}_{\tau+1}), \ldots, c_T(\widehat{X}_T)\big)$$

pointwise on E. Next, since $\widetilde{X}_\tau(E)$ and $\widehat{X}_\tau(E)$ are identical by construction, meaning that $c_\tau(\widetilde{X}_\tau)(E) = c_\tau(\widehat{X}_\tau)(E)$, the time consistency of the risk-measure system implies that

$$\rho_{\tau,T}\big(c_\tau(\widetilde{X}_\tau), \ldots, c_T(\widetilde{X}_T)\big)(E) \leq \rho_{\tau,T}\big(c_\tau(\widehat{X}_\tau), \ldots, c_T(\widehat{X}_T)\big)(E),$$

since we are prepending identical costs to the sequences $c_{\tau+1}(\widetilde{X}_{\tau+1})(E), \ldots, c_T(\widetilde{X}_T)(E)$ and $c_{\tau+1}(\widehat{X}_{\tau+1})(E), \ldots, c_T(\widehat{X}_T)(E)$. However, we have constructed $\widetilde{X}_{\tau,T}(E)$ to be feasible for $\overline{R}^E_\tau(X_{1,\tau-1})$, so we must also have

$$\rho_{\tau,T}\big(c_\tau(X^*_\tau), \ldots, c_T(X^*_T)\big)(E) \leq \rho_{\tau,T}\big(c_\tau(\widetilde{X}_\tau), \ldots, c_T(\widetilde{X}_T)\big)(E)$$

by the optimality of $X^*_{\tau,T}(E)$ for $\overline{R}^E_\tau(X_1, \ldots, X_{\tau-1})$. Combining the two preceding inequalities yields

$$\rho_{\tau,T}\big(c_\tau(X^*_\tau), \ldots, c_T(X^*_T)\big)(E) \leq \rho_{\tau,T}\big(c_\tau(\widehat{X}_\tau), \ldots, c_T(\widehat{X}_T)\big)(E).$$

Since the choice of the feasible solution $\widehat{X}_{\tau,T}(E)$ to $R^E_\tau(X_{1,\tau-1})$ was arbitrary, and we have already established that $X^*_{\tau,T}(E)$ is feasible for $R^E_\tau(X_{1,\tau-1})$, we may deduce that $X^*_{\tau,T}(E)$ must be an optimal solution to $R^E_\tau(X_{1,\tau-1})$.

To complete the proof, we need only show that in the strictly monotonic case, any optimal solution $\widehat{X}_{\tau,T}(E)$ to $R^E_\tau(X_{1,\tau-1})$ must also be optimal for $\overline{R}^E_\tau(X_{1,\tau-1})$. Let $X^*_{\tau,T}(E)$ be any optimal solution to $\overline{R}^E_\tau(X_{1,\tau-1})$. Since we have already established that $X^*_{\tau,T}(E)$ is optimal for $R^E_\tau(X_{1,\tau-1})$, it follows that $\widehat{X}_{\tau,T}(E)$ and $X^*_{\tau,T}(E)$ have the same objective value, the objective functions of the two models being identical. So $\widehat{X}_{\tau,T}(E)$ could only fail to be optimal for $\overline{R}^E_\tau(X_{1,\tau-1})$ if it violated some constraint of that problem, the only possibility being the constraints of the form $\widehat{X}_{\tau+1,T}(D) \in \mathcal{X}^*_{\tau+1,D}(X_{1,\tau-1}, \widehat{X}_\tau)$. Let $\mathcal{D}$ be the collection of elementary events $D \in \mathcal{E}_{\tau+1}$ with $D \subseteq E$ for which this constraint is violated, meaning that there exists an optimal solution $X^{**}_{\tau+1,T}$ to $\overline{R}^D_{\tau+1}(X_{1,\tau-1}, \widehat{X}_{\tau+1})$ with

$$\rho_{\tau+1,T}\big(c_{\tau+1}(X^{**}_{\tau+1}), \ldots, c_T(X^{**}_T)\big)(D) < \rho_{\tau+1,T}\big(c_{\tau+1}(\widehat{X}_{\tau+1}), \ldots, c_T(\widehat{X}_T)\big)(D).$$

We now construct a new solution vector $\widetilde{X}_{\tau,T}(E)$ as follows: we set $\widetilde{X}_\tau(E) = \widehat{X}_\tau(E)$, and for each $D \in \mathcal{E}_{\tau+1}$ with $D \subseteq E$, we set

$$\widetilde{X}_{\tau+1,T}(D) = \begin{cases} X^{**}_{\tau+1,T}(D), & \text{if } D \in \mathcal{D} \\ \widehat{X}_{\tau+1,T}(D), & \text{if } D \notin \mathcal{D}. \end{cases} \tag{26}$$

It is then easily verified that $\widetilde{X}_{\tau,T}(E)$ is feasible for $\overline{R}^E_\tau(X_{1,\tau-1})$. However, the strict monotonicity of the risk measure, the hypothesized nonemptiness of $\mathcal{D}$, and (26) imply that the objective value attained by $\widetilde{X}_{\tau,T}(E)$ is strictly lower than that attained by $\widehat{X}_{\tau,T}(E)$. But this latter value is identical to that attained by $X^*_{\tau,T}(E)$, which was assumed optimal for $\overline{R}^E_\tau(X_{1,\tau-1})$. We thus obtain a contradiction, so it follows that $\widehat{X}_{\tau,T}(E)$ cannot violate any constraints of $\overline{R}^E_t(X_{1,\tau-1})$ and must also be optimal for that problem. □

Thus, in the case of time-consistent risk measure systems, the proposed multilevel system of models essentially reduces to the more familiar setting (5)-(6). When other choices of $\{\rho_{t,T}\}_{t=1}^T$ are made, however, the two frameworks behave differently, with the multilevel framework remaining weakly time consistent without restrictions on the choice of objective functions.

Despite its apparent complexity, a framework of the form (18)-(20)-(21) should really be 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. When one's preferred choice of risk-measure system is time consistent, it reduces to simpler forms that only require single-level optimization.

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 restricting to the most straightforward imaginable special cases that are not immediately equivalent to single-level problems, the modeling framework described in Section 3 leads to an $\mathcal{NP}$-hard problem class. The subclass of models that we consider has the following special form:

- Ω is a finite set, as already assumed.
- $T = 3$; there are three stages.
- We denote the members of $\mathcal{E}_2$ by S, standing for "second-stage scenario". Each second-stage scenario denotes a knowable state of the system at time 2.
- The cost functions $c_t(X_t)$ are all linear; 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 $A_{11}X_1 \leq b_1$; here, A_{11} and b_1 are an $\mathcal{F}_1$-measurable (that is, constant) matrix and vector, respectively.

  - The second-stage constraints are $A_{21}X_1 + A_{22}X_2 \leq b_2$, 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 $A_{31}X_1 + A_{32}X_2 + A_{33}X_3 \leq b_3$, 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 construct the risk-measure system as in (8): we select some law-invariant risk-measure recipe ρ and set

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

Since $T = 3$, the framework (18)-(20)-(21) reduces to a bilevel optimization problem. The top level consists of (21), and the follower problem consists of $|\mathcal{E}_2|$ problems of the form (18). Specifically, we obtain the problem

$$\begin{array}{lll} \min & \rho(c_1^\top X_1 + c_2^\top X_2 + c_3^\top X_3) & \\ \mathrm{ST} & A_{11}X_1(\omega) \leq b_1 & \forall\, \omega \in \Omega \\ & X_1(\omega) = X_1(\omega') & \forall\, \omega, \omega' \in \Omega \\ & \big(X_2(S), X_3(S)\big) \in \mathcal{X}^*_{2,S}(X_1) & \forall\, S \in \mathcal{E}_2, \end{array} \tag{27}$$

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

$$\begin{array}{lll} \min & \rho(c_2^\top X_2 + c_3^\top X_3 \mid S) & \\ \mathrm{ST} & A_{21}X_1(\omega) + A_{22}X_2(\omega) \leq b_2(\omega) & \forall\, \omega \in S \\ & A_{31}X_1(\omega) + A_{32}X_2(\omega) + A_{33}X_3(\omega) \leq b_3(\omega) & \forall\, \omega \in S \\ & X_2(\omega) = X_2(\omega') & \forall\, \omega, \omega' \in S. \end{array} \tag{28}$$

Formally, we define a parameterized class of problems BLRP(ρ) 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 all $\omega \in \Omega$ • 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 (27)-(28).

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 $\text{BLRP}(\rho)$ in which this risk measure parameter, restricted to the rationals, is encoded as part of the problem input. Thus, $\text{BLRP}(\text{MSD}_{(0,1]})$ denotes the class of all $\text{BLRP}(\text{MSD}_\gamma)$ problems, with $\gamma \in (0,1] \cap \mathbb{Q}$ appended to the problem input, and $\text{BLRP}(\text{AVaR}_{(0,1)})$ denotes the class of all $\text{BLRP}(\text{AVaR}_\alpha)$ problems, with $\alpha \in (0,1) \cap \mathbb{Q}$ appended to the problem input.

We will analyze the computational complexity of the problem class $\text{BLRP}(\rho)$ in the cases of $\rho = \text{MSD}_\gamma$, $\gamma \in (0,1)$, and $\rho = \text{AVaR}_\alpha$, $\alpha \in (0,1)$. To do so, we first revisit the 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 (Bard 1991, Ben-Ayed and Blair 1990, Hansen et al. 1992, Jeroslow 1985). This problem class takes the form

$$\begin{array}{lllll} \min & f_1^\top y_1 + f_2^\top y_2 & & & \\ \text{ST} & y_2 \in & \text{Arg min} & \tilde{f}_2^\top y_2 & \\ & & \text{ST} & B_1 y_1 + B_2 y_2 \leq r. & \end{array} \tag{29}$$

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 bilevel linear programming problems corresponding to $\text{BLRP}(\rho)$ are completely general; 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 (29) employ reductions from various combinatorial problems to the special case $\tilde{f}_2 = -f_2$, a problem subclass that we will call *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

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

In principle, the $\mathcal{NP}$-hardness of BOLP may be ascertained 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 15 BOLP *is* $\mathcal{NP}$*-hard.*

Proof. We prove the result 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 $w = (w_1, \ldots, w_n)$:

$$\begin{array}{lll} \min & u_0 + \sum_{i=1}^n w_i & \\ \text{ST} & u_0 \geq \sum_{i=1}^n a_i u_i - \frac{1}{2} \sum_{i=1}^n a_i & \\ & u_0 \geq \frac{1}{2} \sum_{i=1}^n a_i - \sum_{i=1}^n a_i u_i & \\ & 0 \leq u_1, \ldots, u_n \leq 1 & \\ & w \in \quad \text{Arg min} \quad -\sum_{i=1}^n w_i & \\ & \qquad \text{ST} \qquad w_i \leq u_i \qquad & i = 1, \ldots, n \\ & \qquad\qquad\quad w_i \leq 1 - u_i & i = 1, \ldots, n. \end{array} \tag{30}$$

We claim that (30) 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 (30) is equivalent to the (nonconvex) single-level optimization problem

$$\begin{array}{ll} \min & \left|\sum_{i=1}^n a_i u_i - \frac{1}{2} \sum_{i=1}^n a_i\right| + \sum_{i=1}^n \min\{u_i, 1 - u_i\} \\ \text{ST} & 0 \leq u_1, \ldots, u_n \leq 1. \end{array} \tag{31}$$

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 (31) 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 (31) 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 (31) involves the optimization of a continuous function over a compact

set, it must achieve its minimum, and hence the minimum value must be positive. This establishes the claim; it then remains only to observe that problem (30) 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 $a_1, \ldots, a_n$. □

Direct reduction of BOLP or classic bilevel linear programming to problem classes of the form BLRP(ρ) appears to be an intricate task. Thus, we first consider another class of restricted bilevel linear problems and show that it 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

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

We now show that although the leader and follower objectives of BBSBLP(β) may appear very similar — if $\beta > 0$, in particular, then the 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 16 *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

$$g_1 = f_1 \qquad g_2 = \left(1 - \tfrac{2}{1-\beta}\right) f_2 \qquad g_3 = f_2. \tag{32}$$

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

$$B_1x_1 + B_2x_2 \leq r \qquad\qquad x_3 = \left(\tfrac{2}{1-\beta}\right)x_2. \tag{33}$$

Specifically, we may set up such an equivalence by setting

$$C_1 = \begin{bmatrix} B_1 \\ 0 \\ 0 \end{bmatrix} \qquad C_2 = \begin{bmatrix} B_2 \\ -\left(\frac{2}{1-\beta}\right)I \\ \left(\frac{2}{1-\beta}\right)I \end{bmatrix} \qquad C_3 = \begin{bmatrix} 0 \\ I \\ -I \end{bmatrix} \qquad t = \begin{bmatrix} r \\ 0 \\ 0 \end{bmatrix}. \tag{34}$$

Finally, we set $\eta_2 = \zeta$ and $\eta_3 = \left|\frac{2}{1-\beta}\right|\zeta$. Using (32) and that $x_3 = \left(\frac{2}{1-\beta}\right)x_2$ in any feasible solution, the follower problem objective 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 $\|x_3\|_\infty \leq \eta_3$ is equivalent to $\left\|\left(\frac{2}{1-\beta}\right)x_2\right\|_\infty \leq \left|\frac{2}{1-\beta}\right|\zeta$, which is exactly the same as the constraint $\|x_2\|_\infty \leq \eta_2 = \zeta$, so the follower problem may be written

$$\begin{array}{ll} \min & -f_2^\top x_2 \\ \text{ST} & B_1x_1 + B_2x_2 \leq r \\ & \|x_2\|_\infty \leq \zeta. \end{array} \tag{35}$$

Next consider the leader objective, which we may rewrite using (32) 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 (35), 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 that the respective proportionality factors $1 - \frac{2}{1-\beta}$ and $\frac{2}{1-\beta}$ in (32) and (33) were chosen because they are 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 (27)-(28) may be expressed as

$$\begin{array}{ll} \min & c_1^\top x_1 + u_1 + \gamma \sum_{S \in \mathcal{E}_2} \sum_{\omega \in S} \mathrm{P}\{\omega\} \left[c_2(S)^\top x_2(S) + c_3(\omega)^\top x_3(\omega) - u_1 \right]_+ \\ \text{ST} & A_{11} x_1 \leq b_1 \\ & u_1 = \sum_{S \in \mathcal{E}_2} \mathrm{P}\{S\}\, c_2(S)^\top x_2(S) + \sum_{\omega \in \Omega} \mathrm{P}\{\omega\}\, c_3(\omega)^\top x_3(\omega) \\ & \big(x_2(S), x_3(S)\big) \in \widetilde{\mathcal{X}}_S^*(x_1) \qquad\qquad \forall\, S \in \mathcal{E}_2. \end{array} \tag{36}$$

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

$$\begin{array}{ll} \min & c_2(S)^\top x_2(S) + u_2(S) + \gamma \sum_{\omega \in S} \frac{\mathrm{P}\{\omega\}}{\mathrm{P}\{S\}} \left[c_3(\omega)^\top x_3(\omega) - u_2(S) \right]_+ \\ \text{ST} & A_{21}(S) x_1 + A_{22}(S) x_2 \leq b_2(S) \\ & A_{31}(\omega) x_1 + A_{32}(\omega) x_2(S) + A_{33}(\omega) x_3(\omega) \leq b_3(\omega) \qquad\qquad \forall\, \omega \in S \\ & u_2(S) = \sum_{\omega \in S} \big(\frac{\mathrm{P}\{\omega\}}{\mathrm{P}\{S\}} \big) c_3(\omega)^\top x_3(\omega). \end{array} \tag{37}$$

We now show how to construct a subclass of BLRP(MSD_γ) problems that is very similar to BBSBLP(β) for an appropriate choice of β. Consider a three-element probability space with $\Omega = \{\omega_1, \omega_2, \omega_3\}$ partitioned into two 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_γ) problem instance as follows, using the data of the given BBSBLP(β) instance, and as illustrated in Figure 1:

- 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 $\|x_2\|_\infty \leq \eta_2$. 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
    $$C_1 x_1 + C_2 x_2 + C_3 x_3 \leq t \qquad\qquad \|x_3\|_\infty \leq \eta_3.$$
  - For outcome ω_2, the final stage incurs a fixed cost of $K_2 = \eta_3 \|h_3\|_1$.

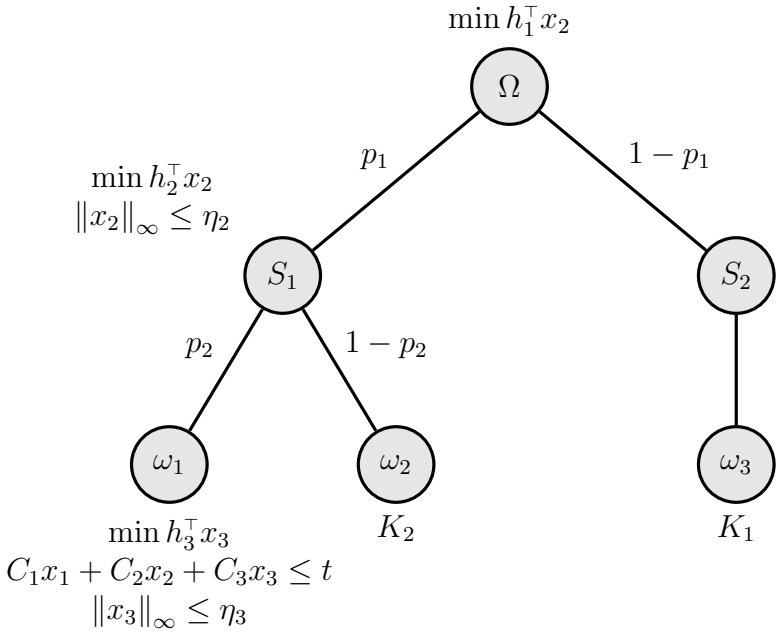

Figure 1: Scenario tree for reduction of BBSBLP(β) instances to BLRP(MSD$_\gamma$) instances.

- Scenario S_2, from which the only possible final-stage consequence is ω_3, incurs a fixed cost of
$$K_1 = \eta_2 \|h_2\|_1 + \left(1 + \tfrac{2p_1p_2}{1-p_1}\right) K_2. \tag{38}$$

Clearly, it is possible configure the input data of a BLRP(MSD$_\gamma$) instance so that the above arrangement is achieved, 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

$$A_{31}(\omega_2) = 0 \quad A_{32}(\omega_2) = 0 \quad A_{33}(\omega_2) = \begin{bmatrix} 1 & 0^\top \\ -1 & 0^\top \end{bmatrix} \quad b_3(\omega_2) = \begin{bmatrix} 1 \\ -1 \end{bmatrix} \quad c_3(\omega_2) = \begin{bmatrix} K_2 \\ 0 \end{bmatrix}.$$

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 (36)-(37) 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:

$$\begin{array}{lll} \min & h_1^\top x_1 + u_1 + \gamma\big(p_1p_2\left[h_2^\top x_2 + h_3^\top x_3 - u_1\right]_+ & \\ & \qquad + p_1(1-p_2)\left[h_2^\top x_2 + K_2 - u_1\right]_+ + (1-p_1)\left[K_1 - u_1\right]_+\big) & \\ \text{ST} & u_1 = p_1p_2(h_2^\top x_2 + h_3^\top x_3) + p_1(1-p_2)(h_2^\top x_2 + K_2) + (1-p_1)K_1 & \\ & (x_2,x_3) \in \quad \text{Arg min} \quad h_2^\top x_2 + u_2 + \gamma\left(p_2\left[h_3^\top x_3 - u_2\right]_+ + (1-p_2)\left[K_2 - u_2\right]_+\right) & \\ & \qquad\qquad \text{ST} \qquad u_2 = p_2 h_3^\top x_3 + (1-p_2)K_2 & \\ & \qquad\qquad\qquad\quad C_1x_1 + C_2x_2 + C_3x_3 \le t & \\ & \qquad\qquad\qquad\quad \|x_2\|_\infty \le \eta_2 & \\ & \qquad\qquad\qquad\quad \|x_3\|_\infty \le \eta_3. & \end{array} \tag{39}$$

Now consider the follower problem of (39). 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

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

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

$$\begin{array}{ll} \min & h_2^\top x_2 + p_2(1-\gamma+\gamma p_2)h_3^\top x_3 \\ \text{ST} & C_1x_1 + C_2x_2 + C_3x_3 \le t \\ & \|x_2\|_\infty \le \eta_2 \\ & \|x_3\|_\infty \le \eta_3. \end{array} \tag{40}$$

We now consider the leader problem in (39). We claim that the choice of K_1 in (38) is sufficiently large for all feasible values of (x_2, x_3) that we have

$$u_1 \ge h_2^\top x_2 + K_2 \ge h_2^\top x_2 + h_3^\top x_3. \tag{41}$$

The second inequality in (41) 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

$$\begin{array}{rl} u_1 = & p_1\big(h_2^\top x_2 + p_2 h_3^\top x_3 + (1-p_2)K_2\big) + (1-p_1)K_1 \\ = & p_1\big(h_2^\top x_2 + p_2 h_3^\top x_3 + (1-p_2)K_2\big) + (1-p_1)\left(\eta_2\|h_2\|_1 + \left(1 + \frac{2p_1p_2}{1-p_1}\right)K_2\right) \\ = & \big[p_1 h_2^\top x_2 + (1-p_1)\eta_2\|h_2\|_1\big] + p_1p_2h_3^\top x_3 + p_1(1-p_2)K_2 + (1-p_1+2p_1p_2)K_2. \end{array} \tag{42}$$

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

$$p_1 h_2^\top x_2 + (1-p_1)\eta_2 \|h_2\|_1 \geq h_2^\top x_2. \tag{43}$$

Next, we observe that

$$h_3^\top x_3 \geq -\|h_3\|_1 \|x_3\|_\infty \geq -\|h_3\|_1 \eta_3 = -K_2. \tag{44}$$

Substituting (43) and (44) into (42), we have

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

establishing (41). It follows immediately that the first two $[\,\cdot\,]_+$ terms in the leader objective of (39) are always zero. Noting that

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

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 (39) may be written, where "$\simeq$" denotes equivalence up to a constant among functions of (x_1, x_2, x_3), as

$$\begin{aligned} & h_1^\top x_1 + u_1 + \gamma(1-p_1)(K_1 - u_1) \\ \simeq\ & h_1^\top x_1 + (1 - \gamma + \gamma p_1)u_1 \\ =\ & h_1^\top x_1 + (1-\gamma+\gamma p_1)\big(p_1 h_2^\top x_2 + p_1p_2 h_3^\top x_3 + p_1(1-p_2)K_2 + (1-p_1)K_1\big) \\ \simeq\ & h_1^\top x_1 + (1-\gamma+\gamma p_1)\big(p_1 h_2^\top x_2 + p_1p_2 h_3^\top x_3\big) \\ =\ & h_1^\top x_1 + p_1(1-\gamma+\gamma p_1)\big(h_2^\top x_2 + p_2 h_3^\top x_3\big). \end{aligned}$$

Combining this form of the leader objective with (40), we may express the entire problem as

$$\begin{array}{lllll} \min & h_1^\top x_1 + p_1(1-\gamma+\gamma p_1)h_2^\top x_2 + p_1p_2(1-\gamma+\gamma p_1)h_3^\top x_3 & & & \\ \text{ST} & (x_2, x_3) \in & \text{Arg min} & h_2^\top x_2 + p_2(1-\gamma+\gamma p_2)h_3^\top x_3 & \\ & & \text{ST} & C_1x_1 + C_2x_2 + C_3x_3 \leq t & \\ & & & \|x_2\|_\infty \leq \eta_2 & \\ & & & \|x_3\|_\infty \leq \eta_3. & \end{array} \tag{45}$$

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 (45) to reduce BBSBLP(β) to BLRP(MSD$_\gamma$) for an appropriate choice of β, thus proving that BLRP(MSD$_\gamma$) is $\mathcal{NP}$-hard.

Proposition 17 *For any $\gamma \in (0,1] \cap \mathbb{Q}$, the problem class* BLRP(MSD$_\gamma$) *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 16. 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 (39); 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 (45). Substituting the above choices of h_1, h_2, and h_3 into the leader objective of (45), 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 (45), 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 can 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

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

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

Corollary 18 *The problem class* BLRP($\text{MSD}_{(0,1]}$)*, with* γ *encoded as part of the problem input, is also* $\mathcal{NP}$*-hard.*

Proof. Consider any instance of the problem class BLRP($\text{MSD}_{1/2}$), which is $\mathcal{NP}$-hard by Proposition 17. 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($\text{MSD}_{(0,1]}$) would imply a polynomial-time algorithm for BLRP($\text{MSD}_{1/2}$). □

7 Complexity of Bilevel AVaR Risk Models

We now consider the complexity of bilevel models using the AVaR_α risk measure instead of the MSD_γ risk measure; the overall analysis technique is similar to the MSD_γ 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, but with different probabilities all based on the parameter α, and the 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 $\|x_2\|_\infty \leq \eta_2$. 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
    $$C_1 x_1 + C_2 x_2 + C_3 x_3 \leq t \qquad\qquad \|x_3\|_\infty \leq \eta_3.$$
  - 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 2. 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:

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

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

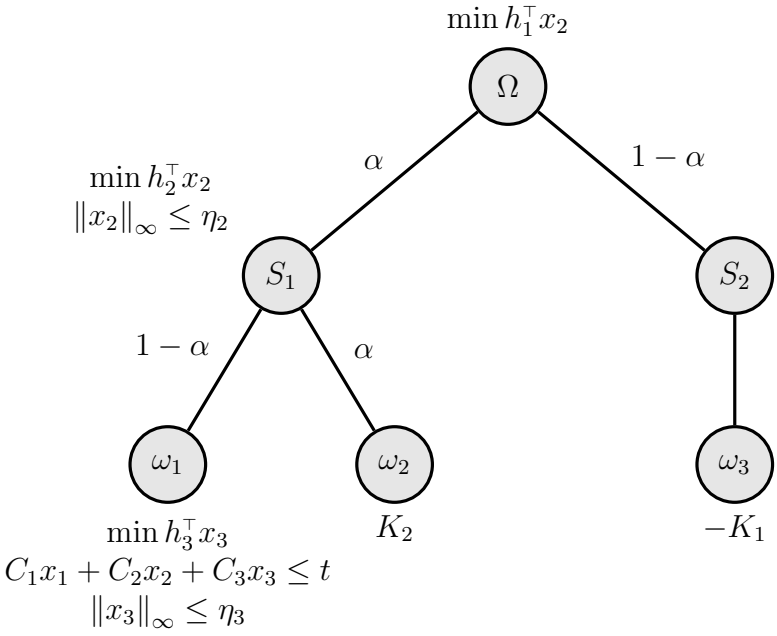

Figure 2: Scenario tree for reduction of BBSBLP(0) instances to BLRP(AVaR_α) instances.

This BLRP(AVaR_α) instance then has the form

$$\begin{array}{lllll} \min & h_1^\top x_1 + \text{AVaR}_\alpha\big(W_2(x_2, x_3)\big) & & & \\ \text{ST} & (x_2, x_3) \in & \text{Arg min} & h_2^\top x_2 + \text{AVaR}_\alpha\big(W_3(x_3)\big) & \\ & & \text{ST} & C_1x_1 + C_2x_2 + C_3x_3 \leq t & \\ & & & \|x_2\|_\infty \leq \eta_2 & \\ & & & \|x_3\|_\infty \leq \eta_3, & \end{array} \tag{46}$$

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 $x_3(\omega_1)$.

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 (46) is 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

$$h_3^\top x_3 \leq \|h_3\|_1 \|x_3\|_\infty \leq \eta_3 \|h_3\|_1 < \eta_3 \|h_3\|_1 + 1 = K_2, \tag{47}$$

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 (46) may be replaced, without any change to 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 $\mathrm{AVaR}_\alpha\big(W_2(x_2, x_3)\big)$ term in the leader objective of (46). We note that

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

for any feasible values of x_2 and x_3. Further, as we have already established in (47) 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 (46), 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

$$\begin{array}{lllll} \min & h_1^\top x_1 + h_2^\top x_2 + (1-\alpha)h_3^\top x_3 & & & \\ \mathrm{ST} & (x_2, x_3) \in & \mathrm{Arg\ min} & h_2^\top x_2 & \\ & & \mathrm{ST} & C_1 x_1 + C_2 x_2 + C_3 x_3 \le t & \\ & & & \left\|x_2\right\|_\infty \le \eta_2 & \\ & & & \left\|x_3\right\|_\infty \le \eta_3. & \end{array} \tag{48}$$

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

Proposition 19 *For any $\alpha \in (0,1) \cap \mathbb{Q}$, the problem class* BLRP(AVaR_α) *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

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

Now construct a BLRP(AVaR_α) instance of the form (46), which by the immediately preceding analysis is equivalent to (48). Substituting the definitions of h_1, h_2, h_3 above into (48), we obtain the equivalent problem

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

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 20 *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 18. □

8 Example Computational Experiments

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 same 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 the 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 our experiments, we used the MSD_γ risk measure with $\gamma = 0.3$, $\gamma = 0.7$, or $\gamma = 0.9$.

For our test datasets, we tried $(P_1, P_2) = (5, 5), (10, 5), (10, 10)$. We also 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.

Having obtained the test data for each instance from the generator of Collado et al. (2012), we used the AMPL modeling language to formulate two different corresponding problems. One of these problems was simply the same model proposed by Collado et al. (2012), in which we use a time-consistent objective of the form (17). This approach applies the mean-semideviation operation twice to all third-stage costs, and results in a moderately large linear programming problem.

The other problem we formulated was equivalent to the $\text{BLRP}(\text{MSD}_\gamma)$ instance corresponding to the problem data. To obtain this problem in AMPL form, we used a series of model transformations:

1. First, using auxiliary variables to model the MSD_γ risk measure, we obtained an equivalent bilevel linear programming problem.
2. Explicitly introducing dual variables for the follower constraints in this bilevel LP, we replaced the follower linear programming problem by an equivalent set of KKT conditions, thus transforming the problem to a mathematical program with equlibrium constraints (MPEC); see for example Luo et al. (1996).
3. We then took each complementarity constraint in the MPEC formulation and transformed it to a special ordered set (SOS) constraint of type 1 (Beale and Tomlin 1970), with only two elements per set. This type of constraint, commonly implemented in commercial mixed-integer programming (MIP) solvers, mandates that at most one variable in the set may be nonzero. While traditionally applied to binary variables, such constraints may also be used for continuous variables.

In this way, we obtained a formulation that could be processed by a standard MIP solver. In the tests we describe below, we used version 5.6.2 of the GuRoBi solver (Gurobi Optimization 2014), which was able to compute globally optimal solutions to the models in reasonably

Table 1: Computational results for example problems with $\gamma = 0.3$.

P_1	P_2	N_1	N_2	Obj=NTC Sol=NTC	Obj=NTC Sol=TC	Δ_{obj}	NTC Elapsed	NTC CPU	TC CPU
5	5	5	5	–821.164	–789.897	4.0%	0.05	0.89	0.004
5	5	5	5	–1454.26	–1429.49	1.7%	0.33	3.32	0.004
5	5	5	5	–498.390	–459.625	8.4%	0.34	4.05	0.004
5	5	5	5	–980.502	–950.32	3.2%	0.06	0.90	0.000
5	5	5	5	–2733.06	–2706.21	1.0%	0.31	3.44	0.004
10	5	5	10	–13097.40	–12989.80	0.8%	0.31	6.27	0.008
10	5	5	10	–1353.64	–1311.23	3.2%	0.28	6.02	0.008
10	5	5	10	–5098.50	–5044.99	1.1%	0.43	8.53	0.008
10	5	5	10	–3839.42	–3766.54	1.9%	0.37	7.88	0.008
10	5	5	10	–4978.18	–4892.19	1.8%	0.38	6.51	0.008
10	10	10	10	–11953.80	–11886.30	0.6%	3.48	99.68	0.020
10	10	10	10	–8777.65	–8692.34	1.0%	6.67	206.69	0.016
10	10	10	10	–5984.11	–5938.5	0.8%	14.25	449.62	0.020
10	10	10	10	–9601.78	–9533.19	0.7%	3.11	91.00	0.020
10	10	10	10	–11862.40	–11765.90	0.8%	2.91	80.94	0.024

short times. We were also able to solve the same models using the CPLEX 12.5.1 solver, but its run times were somewhat longer.

For brevity, we refer to the formulation with the time-consistent choice of objective as "TC", and the multilevel formulation as "NTC", since its choice of objective function is not time consistent; however, both *models* are time consistent in the sense of Definition 9 (strongly in the TC case, weakly in the NTC case).

The results for $\gamma = 0.3$ are summarized in Table 1. The first four columns describe the probem parameters, P_1 being the number of parts, P_2 being the number of products, $N_1 = |\mathcal{E}_2|$ being the number of second-stage scenarios and N_2 being the number of third-stage scenarios for each second-stage scenario, hence $|\Omega| = N_1 N_2$ and $N_2 = |\Omega| \,/\, |\mathcal{E}_2|$. For each different combination of P_1, P_2, N_1, N_2, we randomly generated five different problem instances. The next column, labeled "Obj=NTC, Sol=NTC" is the objective function value attained by the bilevel model, while the following column, labeled "Obj=NTC, Sol=TC" shows the value of the same objective function, computed at the optimal solution of the corresponding single-level model with the nested, time-consistent objective function. This solution is necessarily feasible for the bilevel model, but not optimal for it. The next column, labeled Δ_{obj}, shows the percentage difference between the two objective function values. The final three columns indicate running times in seconds, all on an Ubuntu Linux workstation with two eight-core hyperthreaded 2.2GHz Xeon E5-2660 processors and 1600MHz memory. This system supports up to 32 simultaneous execution threads, which GuRoBi's branch-and-cut MILP algorithm appears to be able to use very efficiently. The "NTC Elapsed" is

Table 2: Computational results for example problems with $\gamma = 0.7$.

P_1	P_2	N_1	N_2	Obj=NTC Sol=NTC	Obj=NTC Sol=TC	Δ_{obj}	NTC Elapsed	NTC CPU	TC CPU
5	5	5	5	–714.803	–643.076	11.2%	0.37	7.12	0.004
5	5	5	5	–1342.34	–1274.82	5.3%	0.32	3.39	0.004
5	5	5	5	–325.178	–234.128	38.9%	0.06	1.04	0.004
5	5	5	5	–886.559	–806.171	10.0%	0.31	3.16	0.004
5	5	5	5	–2643.91	–2579.49	2.5%	0.30	3.74	0.004
10	5	5	10	–12643.6	–12381.4	2.1%	0.37	7.11	0.008
10	5	5	10	–1189.21	–1097.29	8.4%	0.30	5.66	0.008
10	5	5	10	–4734.64	–4595.77	3.0%	0.43	8.86	0.008
10	5	5	10	–3504.99	–3317.26	5.7%	0.42	8.86	0.008
10	5	5	10	–4674.50	–4485.90	4.2%	0.36	7.07	0.008
10	10	10	10	–11671.4	–11502.5	1.5%	8.47	259.27	0.016
10	10	10	10	–8398.79	–8234.36	2.0%	7.61	234.80	0.016
10	10	10	10	–5785.69	–5667.47	2.1%	16.51	519.17	0.020
10	10	10	10	–9351.43	–9168.16	2.0%	3.29	100.84	0.020
10	10	10	10	–11456.1	–11230.9	2.0%	3.13	89.09	0.020

the total elapsed time to solve the NTC model on this system, whereas the "NTC CPU" column is the total number of CPU seconds reported, that is, the total CPU time summed over all threads. Note that for the more difficult problems, GuRoBi was able to obtains parallel speedups of over 31.5. The last column is the reported CPU time for solving the TC model, which is simply a linear program. For these problems, parallel speedups were minimal, and the elapsed time and CPU time were virtually identical, so only CPU times are shown. While the NTC models take considerably longer to solve than the TC models, most can still be solved in under a second of elapsed time and a few seconds of total CPU time. The 100-scenario problems are the most demanding, but the most difficult still takes under 17 seconds of elapsed time.

Generally speaking, the difference between the NTC and TC solution values is fairly small, but is as high as 8.4% for one of the smaller models, meaning that if one's objective is to minimize the $\text{MSD}_{0.3}$ risk measure applied to the sum of all costs, one can do about 8% better by using the bilevel model. Curiously, the difference between the two models seems to decrease as one increases the number of scenarios.

Tables 2 and 3 show computational results for the same models, but with $\gamma = 0.7$ and $\gamma = 0.9$, respectively. While the run times are largely similar, the difference between the NTC and TC solution values is much more pronounced, with one Δ_{obj} value close to 40% for $\gamma = 0.7$ and one over 100% for $\gamma = 0.9$. These results accord with intuition, since the effects of applying a risk measure in a nested, recursive way should be magnified with increasing risk aversion. Specifically, as compared to simple risk measure, a nested time-consistent risk

Table 3: Computational results for example problems with $\gamma = 0.9$.

P_1	P_2	N_1	N_2	Obj=NTC Sol=NTC	Obj=NTC Sol=TC	Δ_{obj}	NTC Elapsed	NTC CPU	TC CPU
5	5	5	5	–659.865	–567.346	16.3%	0.51	10.02	0.004
5	5	5	5	–1286.40	–1192.27	7.9%	0.10	1.84	0.004
5	5	5	5	–219.852	–105.659	108.1%	0.50	9.39	0.004
5	5	5	5	–839.831	–730.047	15.0%	0.32	3.61	0.004
5	5	5	5	–2599.82	–2508.45	3.6%	0.34	3.96	0.004
10	5	5	10	–12416.7	–12072.1	2.9%	0.37	7.77	0.008
10	5	5	10	–1112.78	–986.016	12.9%	0.30	6.55	0.008
10	5	5	10	–4595.11	–4399.01	4.5%	0.50	10.64	0.008
10	5	5	10	–3337.78	–3085.05	8.2%	0.38	8.26	0.008
10	5	5	10	–4522.83	–4289.57	5.4%	0.36	6.96	0.008
10	10	10	10	–11530.6	–11305.5	2.0%	8.50	263.57	0.020
10	10	10	10	–8218.75	–7993.42	2.8%	9.19	284.41	0.016
10	10	10	10	–5690.80	–5539.66	2.7%	16.61	523.43	0.024
10	10	10	10	–9225.67	–8980.46	2.7%	10.16	318.62	0.024
10	10	10	10	–11259.0	–10970.0	2.6%	3.22	91.96	0.024

measure increases risk aversion and sensitivity to "tail" events as costs and rewards become more distant in time; the higher the level of risk aversion in the one-step risk-measures making up the formula (17), the more pronounced this effect becomes.

For all values of γ, the tendency of the objective values of the two models to become more similar as the number of scenarios increases remains present; the reasons for this phenomenon are unclear at present.

9 Modeling Considerations and Concluding Remarks

In conclusion, while we have shown that the class of multilevel models suggested in this paper is $\mathcal{NP}$-hard for two of the most common risk measure families, we have also shown that at least one class of realistic instances are easily solvable to global optimality with standard modern workstation hardware and standard optimization solvers, even without trying to exploit their special structural properties. Thus, using multilevel models seems a possible path to approach the application of risk measures to multistage decision problems, avoiding the complexities of nested time-consistent objective functions. Of course, more work needs to be done to investigate the use of these models, including examining different classes of applications, applying different risk measures than MSD_γ, attempting to solve larger models, and possible development of specialized solution algorithms.

The modeling approach we suggest has the advantage of giving the modeler essentially complete freedom in the choice of objective functions for the top-level problem and every

recourse problem without violating weak time consistency, thus removing one of the key barriers to applying coherent and related risk measures to multistage stochastic decision problems. However, it has its own "barriers to entry". One drawback is that the modeler has to become comfortable with thinking about a multistage decision problem as a multilevel optimization, even if there is only one decision maker. This philosophy, which may be difficult to grasp initially, is intrinsic to the modeling approach and cannot be avoided. The approach's other modeling drawbacks, however, have to do with the availability of appropriate modeling tools, and could likely be overcome with appropriate advances in modeling software. To put ourselves in a position to perform the computational tests described in Section 8, we had to follow a long series of model manipulations, first creating a multilevel model, then converting it to an MPEC, and finally converting the MPEC complementarity constraints into SOS1 constraints in a form understood by our modeling/solution environment. Each of these steps is tedious and error-prone, and all could probably be automated. Furthermore, the final form of the resulting AMPL model is complicated and opaque.

The final step of converting a linear (or perhaps quadratic) MPEC model into an SOS1 form suitable for a MIP solver is likely the easiest to automate. AMPL already supports a `complements` keyword for expressing complementarity constraints, and other modeling languages have similar facilities. However, AMPL currently only communicates `complements` constraints intelligibly to nonlinear solvers such as LOQO. Such solvers are able to process complementarity constraints (Benson et al. 2006), but generally guarantee only locally optimal solutions. In earlier computational work, we used LOQO to compute locally solutions to the same models described in Section 8: while the AMPL files were simpler, global optimality was not guaranteed, and solution times were longer than with GuRoBi. Unfortunately, converting the complementarity constraints to SOS1 form proved to be awkward, because AMPL uses a somewhat opaque "suffix" syntax to create SOS constraints, and the user must assign a unique index number to every SOS constraint, requiring the AMPL file to contain complex indexing calculations which are not an intrinsic part of the model. As a first step towards making multilevel modeling more accessible, it seems that it should be possible to have modeling languages automatically translate complentarity constructs such as AMPL's `complements` into a form intelligible to MIP solvers such as GuRoBi and CPLEX.

As an aside, AMPL supports an "implies" syntax which, if one introduces auxiliary binary variables, allows for clearer expression of complementarity constraints than through SOS groupings. This syntax uses the operator "`==>`". For example, the complementarity condition $\lambda \geq 0 \perp \langle a, x \rangle \leq b$ could be expressed in AMPL as

```
var lambda >= 0;
var flag binary;
subject to feasibility : sum {i in 1..n} a[i]*x[i] <= b;
subject to complementarity :
   flag = 0 ==> lambda = 0
            else sum {i in 1..n} a[i]*x[i] = b;
```

While resulting in a more human-readable model than the SOS1 approach, this technique is currently understood only by the CPLEX solver and cannot be used with GuRoBi. For each complementarity condition, it also requires an auxiliary binary variable (here, `flag`) that in some sense it should not be necessary for the modeler to explicitly create. Furthermore, we found that models transformed in this way solved slower than their SOS1 counterparts. Ideally, the modeling language should be able to translate explicit complementarity constraints into whatever form is most appropriate for the solver being used.

While the abilit to automate the transformation of complementarity constraints into solver-appropriate equivalent forms would be helpful, it would be even better if translation of multilevel optimization models into complementarity form could also be automated. In the experiments of Section 8, we had to take each constraint of the recourse model, create an explicit dual variable for it in AMPL, and then formulate the correct KKT complementarity condition. This process is tedious and error-prone, and another candidate for automation. However, it appears to be more challenging to implement than translation of already-formulated complementarity conditions into solver-appropriate forms.

Standard coherent risk measures such as AVaR_α and MSD_γ also present some modeling challenges using present technology. To use such risk measures in an AMPL model, one must explicitly introduce auxiliary variables and constraints, obscuring the basic structure of the model. It would be better to have modeling tools that would allow one to specify the application's constraints and objective function elements, and then call on a library of standard risk measures to apply to the objective. The necessary model transformations, including the addition of any auxiliary variables and constraints, as well as transformation into a multilevel form in the cases of non-time-consistent risk measure systems, could then in principle be automated. Ideally, such facilities should be integrated into a modeling environment tailored to the efficient expression of stochastic programming problems; see for example Colombo et al. (2009) and Valente et al. (2009).

With an entire layered suite of tools as described above, multistage risk-averse optimization problems could be modeled without the lengthy sequence of manual transformations described in Section 8. However, creation of such tools would clearly be a major, multi-year effort.

Finally, we make additional concluding observations. First, the observed phenomenon of the difference between the multilevel and nested-objective models decreasing with the number of scenarios bears further investigation. If it occurs for other classes of applications, it needs to be better understood, and might prove helpful in devising specialized algorithms for solving large-scale problems.

Investigation of the use of entropic risk measures of the form suggested by Kupper and Schachermayer (2009) also bears further investigation. These risk measures are time consistent but not scale invariant. If they can gain acceptance as primary modeling techniques, they would permit simpler alternatives to the multilevel models proposed here. If not, they still might have some applicability as time-consistent approximations to other formulations.

Although the $\mathcal{NP}$-hardness results we have obtained cover only two specific 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 linear programming problem instances in which the law-invariant use of the MSD_γ or AVaR_α risk measures in the form of (9) or (10) 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 phenomenon underscores that when using "natural" conditional measures of risk like (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.

Asamov, Tsvetan. 2013. Time-consistent approximations of risk-averse multistage stochastic optimization problems. Ph.D. thesis, Rutgers University.

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

Beale, E. M. L., John A. Tomlin. 1970. Special facilities in a general mathematical programming system for non-convex problems using ordered sets of variables. J. Lawrence, ed., *Proceedings of the Fifth International Conference on Operational Research*. Tavistock Publications, London, 447–454.

Ben-Ayed, O., C. E. Blair. 1990. Computational difficulties of bilevel linear programming. *Oper. Res.* **38** 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** 155–182.

Boda, K., J. A. Filar. 2006. Time consistent dynamic risk measures. *Math. Methods Oper. Res.* **63** 169–186.

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.

Colombo, M., A. Grothey, J. Hogg, K. Woodsend, J. Gondzio. 2009. A structure-conveying modelling language for mathematical and stochastic programming. *Math. Program. Comput.* **1** 223–247.

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

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

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

Gurobi Optimization, Inc. 2014. Gurobi optimizer reference manual. URL `http://www.gurobi.com`.

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

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

Kupper, M., W. Schachermayer. 2009. Representation results for law invariant time consistent functions. *Math. Financ. Econ.* **2** 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** 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** 51–74.

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

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

Ruszczyński, A., A. Shapiro. 2006b. Optimization of convex risk functions. *Math. Oper. Res.* **31** 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. *Oper. Res. Let.* **40** 436–439.

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.

Valente, C., G. Mitra, M. Sadki, R. Fourer. 2009. Extending algebraic modelling languages for stochastic programming. *INFORMS J. Comput.* **21** 107–122.