DECISION RULE BOUNDS FOR TWO-STAGE STOCHASTIC BILEVEL PROGRAMS

İHSAN YANIKOĞLU* AND DANIEL KUHN†

Abstract. We study two-stage stochastic bilevel programs where the leader chooses a binary here-and-now decision and the follower responds with a continuous wait-and-see decision. Using modern decision rule approximations, we construct lower bounds on an optimistic version and upper bounds on a pessimistic version of the leader's problem. Both bounding problems are equivalent to explicit mixed-integer linear programs that are amenable to efficient numerical solution. The method is illustrated through a facility location problem involving sellers and customers with conflicting preferences.

Key words. bilevel programming, stochastic programming, decision rules

1. Introduction.

We study two-stage stochastic bilevel programming problems, that is, sequential two-player games where the first mover is referred to as the *leader*, and the second mover is termed the *follower*. The leader first chooses a vector of binary *here-and-now* decisions $x \in \mathcal{X} \subseteq \{0,1\}^d$. Next, a vector of uncertain problem parameters $\xi \in \mathbb{R}^k$ is revealed, in response to which the follower selects a vector of continuous *wait-and-see* decisions $y(\xi) \in \mathbb{R}^n$ subject to $Ay(\xi) \le b_x(\xi)$ for some $A \in \mathbb{R}^{m\times n}$ and $b_x(\xi) \in \mathbb{R}^m$. The crux of bilevel programming is that the leader and the follower may have conflicting preferences. Specifically, the follower minimizes $c(\xi)^\top y(\xi)$ for some $c(\xi) \in \mathbb{R}^n$, while the leader minimizes $q^\top x + \mathbb{E}[v(\xi)^\top y(\xi)]$ for some $q \in \mathbb{R}^d$ and $v(\xi) \in \mathbb{R}^n$, where $\mathbb{E}[\cdot]$ denotes the expectation operator with respect to the distribution $\mathbb{P}$ of ξ. We assume henceforth that $v, c \in \mathcal{L}^2_n$ and $b_x \in \mathcal{L}^2_m$ for all $x \in \mathcal{X}$, where $\mathcal{L}^2_r$ for $r \in \mathbb{N}$ denotes the space of all r-dimensional square-integrable functions of ξ, that is, all Borel-measurable functions f from $\mathbb{R}^k$ to $\mathbb{R}^r$ with $\mathbb{E}[\|f(\xi)\|^2] < \infty$.

We emphasize that the leader's decision x affects the follower's feasible set (via the dependence of $b_x(\xi)$ on x), while the follower's decision $y(\xi)$ affects the leader's cost (through the expected wait-and-see cost $\mathbb{E}[v(\xi)^\top y(\xi)]$). The leader's decision x thus has a direct effect on her cost (through $q^\top x$) as well as an indirect effect because x affects the follower's feasible set and therefore impacts $\mathbb{E}[v(\xi)^\top y(\xi)]$.

The two-stage stochastic bilevel program described above is ill-posed if the follower's decision problem admits multiple optimal solutions. Indeed, the leader's cost may depend on the follower's choice among his optimal solutions, whereas the follower's cost is clearly independent of this choice. In the *optimistic* formulation of the bilevel program the follower is assumed to select the optimal solution that is *least* hurtful for the leader.

$$
\begin{array}{lll}
\inf\limits_{x\in\mathcal{X},\, y\in\mathcal{L}^2_n} & q^\top x + \mathbb{E}[v(\xi)^\top y(\xi)] & \\
\text{s.t.} & y(\xi) \in \operatorname*{argmin}\limits_{y'\in\mathbb{R}^n} \left\{ c(\xi)^\top y' : Ay' \le b_x(\xi) \right\} \ \mathbb{P}\text{-a.s.}
\end{array}
\tag{$\mathcal{O}$}
$$

Note that the leader thus optimizes over her own feasible decisions as well as over the follower's optimal decisions. In the *pessimistic* formulation of the two-stage stochastic bilevel program the follower is assumed to select the optimal solution that is *most*

*Industrial Engineering Department, Özyeğin University, İstanbul, Turkey – Corresponding author e-mail: ihsan.yanikoglu@ozyegin.edu.tr

†Risk Analytics and Optimization Chair, EPFL, Lausanne, Switzerland

hurtful for the leader.

$$\begin{array}{lll} \min_{x\in\mathcal{X}} & \sup_{y\in\mathcal{L}^2_n} & q^\top x + \mathbb{E}[v(\xi)^\top y(\xi)] \\ & \text{s.t.} & y(\xi) \in \operatorname*{argmin}_{y'\in\mathbb{R}^n} \left\{ c(\xi)^\top y' : Ay' \le b_x(\xi) \right\} \ \mathbb{P}\text{-a.s.} \end{array} \tag{$\mathcal{P}$}$$

We emphasize that the objective functions of $\mathcal{O}$ and $\mathcal{P}$ are well-defined because the cost function v and the follower's decision y are assumed to be square-integrable.

In this paper we focus on the *pessimistic* version of the two-stage stochastic bilevel program, which is the appropriate formulation if there is no communication between the leader and the follower. In this case the leader has no means to anticipate the follower's decision and should prepare for the worst case.

Example 1. *Assume that ξ is a scalar random variable and that $\mathbb{P}$ denotes the uniform distribution on $[0,1]$. Moreover, assume that the leader selects a single binary decision $x \in \{0,1\}$ and that the follower selects a two-dimensional continuous decision $y(\xi) \in \mathbb{R}^2$. Finally, assume that the follower solves the parametric linear program*

$$\min_{y'\in\mathbb{R}^2} \left\{ y'_2 : -x\xi \le y'_1 \le x\xi,\ x\xi \le y'_2 \le 1 \right\},$$

while the leader minimizes $\mathbb{E}[y_1(\xi)]$, where $y \in \mathcal{L}^2_2$ represents an optimal policy for the follower. Thus, the optimal solution of the optimistic bilevel program $\mathcal{O}$ is $x = 1$. In this case, the follower's optimal policy that is most desirable for the leader is given by $y(\xi) = (-\xi, \xi)$, which results in expected costs of $-\frac{1}{2}$ for the leader and $+\frac{1}{2}$ for the follower. In contrast, the optimal solution of the pessimistic bilevel program $\mathcal{P}$ is $x = 0$. In this case, the unique optimal policy of the follower is given by $y(\xi) = (0,0)$, which results in expected costs of 0 both for the leader and the follower. This example shows that the optimal solutions of $\mathcal{O}$ and $\mathcal{P}$ may be different even in simple cases. The example also highlights that the follower may be better off with a pessimistic leader.

Bilevel programs have found numerous applications in revenue management [8], supply chain management [31], production planning [22], transportation [26], security [32], energy economics [21] and countless other domains. *Deterministic* bilevel programs have been studied in game theory since Stackelberg's celebrated treatise on the strategic interactions between market leaders and followers in 1934 [36], but the development of algorithmic solution procedures has only begun in the 1970s [5]. Most existing algorithms apply only to *optimistic* bilevel programs and assume that the follower's optimization problem is *convex*. In this case the bilevel program can be reformulated as an equivalent single-level program by replacing the follower's subproblem with the coresponding Karush-Kuhn-Tucker optimality conditions. The resulting single-level problem is typically non-convex but can be addressed with global optimization algorithms [15, 27, 37, 39]. For a comprehensive survey of the literature up to 2002 we refer to [10], and for further details on optimality conditions of bilevel programming problems we refer to [9, 11, 12]. The literature on solution methods for *pessimistic* deterministic bilevel programs is much sparser, see, *e.g.* [38, 40] and the references therein.

Despite their rich modeling power, *stochastic* bilevel programs have attracted much less attention—presumably because they are perceived as substantially harder than their (already difficult) deterministic counterparts. *Optimistic* stochastic bilevel models emerged in the context of truss topology optimization [7, 29], traffic planning [1, 28], network design [2] and strategic pricing in electricity markets [16] etc. We note that these optimistic models can also be viewed as special instances of stochastic mathematical programs with equilibrium constraints [14, 30, 33, 41, 42].

Solution algorithms for stochastic bilevel programs typically target the corresponding deterministic equivalents, which constitute deterministic bilevel programs with separate blocks of follower decisions and constraints for each uncertainty realization. If the distribution of the random parameters is continuous, the problems are first discretized via scenario generation techniques, see, *e.g.* [35]. Decomposition algorithms exploiting the block structure of the deterministic equivalents are developed in [20]. To our best knowledge, *pessimistic* stochastic bilevel programs have not yet been studied systematically.

Similarly, most solution methods for stochastic mathematical programs with equilibrium constraints focus on single-stage models with discrete distributions and continuous decision variables only. An entropic approximation scheme for more general mathematical programs with *semi-infinite* equilibrium constraints is discussed in [24].

In this paper we make a first step towards a computationally viable treatment of stochastic bilevel programs. Specifically, we use modern decision rule techniques to show that a *stochastic* bilevel program can be approximated by two simpler *deterministic* bilevel programs that are not significantly harder than the *nominal* version of the original stochastic problem—even if the underlying random parameters follow a continuous distribution. The use of *primal* decision rules will lead to an *upper* bound on the original stochastic bilevel program, whereas the use of *dual* decision rules will yield a *lower* bound. We further demonstrate that the two bounding problems admit exact reformulations as explicit mixed-integer linear programs (MILP) that can be solved with standard software. If *linear* decision rules are used, our bounds depend only on the first two moments and the support of $\mathbb{P}$. Thus, the bounds are robust with respect to model mis-specifications as long as the location and dispersion of $\mathbb{P}$ are correctly estimated.

In the remainder we first delineate the precise model assumptions (Section 2) and then review the primal and dual linear decision rule bounds for single-level stochastic programs (Section 3). Next, we extend these bounds to stochastic bilevel programs (Section 4). After discussing refined bounds based on non-linear decision rules (Section 5) and exact solution procedures for stochastic bilevel programs with discrete distributions (Section 6), we assess the accuracy and computational performance of the decision rule bounds in the context of a facility location problem (Section 7).

Notation. The i-th standard basis vector of a finite-dimensional linear space is denoted by $\mathbf{e}_i$, while $\mathbf{e}$ stands the vector of ones. The dimensions will usually be clear from the context. The trace of any square matrix $M \in \mathbb{R}^{d\times d}$ is denoted by $\mathrm{Tr}(M)$.

2. Preliminary Results and Assumptions.

We first delineate the assumptions underlying the results of this paper.

ASSUMPTION 1 (Linearity). *We assume that* $c(\xi) = C\xi$ *for* $C \in \mathbb{R}^{n\times k}$, $v(\xi) = V\xi$ *for* $V \in \mathbb{R}^{n\times k}$, $b_x(\xi) = B_x\xi$ *for* $B_x \in \mathbb{R}^{m\times k}$ *and* $B_x = \sum_{i=1}^d B_i x_i$ *for* $B_i \in \mathbb{R}^{m\times k}$, $i = 1,\dots,d$. *Moreover, we assume that* $\mathbb{P}$ *is supported on* $\Xi = \{\xi \in \mathbb{R}^k : W\xi \geq h\}$ *for some* $W \in \mathbb{R}^{\ell\times k}$ *and* $h \in \mathbb{R}^\ell$. *Finally, we assume that the first two rows of* W *coincide with* $\mathbf{e}_1$ *and* $-\mathbf{e}_1$*, respectively, while* $h = [1, -1, 0, \dots, 0]^\top$. *The last condition implies that* $\xi_1 = 1$ *for all* $\xi \in \Xi$.

As $\xi_1 = 1$ for all $\xi \in \Xi$, the component ξ_1 constitutes a degenerate random variable that is equal to 1 almost surely under $\mathbb{P}$. Thus, any linear function of ξ can be viewed as an affine function of the non-degenerate random variables $\xi_2, \dots, \xi_k$. This allows us to work with affine functions while benefitting from the notational simplicity of linear functions.

ASSUMPTION 2 (Second-Order Moments). *We assume that* $\mathbb{E}[\|\xi\|^2] < \infty$ *and* *that the second-order moment matrix $\Omega = \mathbb{E}[\xi\xi^\top]$ of ξ is strictly positive definite.*

One can show that Ω is strictly positive definite whenever the non-degenerate random variables $\xi_2, \ldots, \xi_k$ are truly non-degenerate, that is, whenever each of them has a strictly positive variance. This is equivalent to requiring the support of the distribution $\mathbb{P}$ to span $\mathbb{R}^k$; see [23, Proposition 2] for further details. Together with Assumption 1, Assumption 2 further guarantees that the cost function c is square-integrable. Indeed, as ξ has finite second-order moments, we have $\mathbb{E}[\|c(\xi)\|^2] = \mathrm{Tr}(C\Omega C^\top) < \infty$. The square-integrability of v and b_x for any $x \in \mathcal{X}$ can be shown similarly.

Assumption 3 (Relatively Complete Recourse). *We assume that $\{y \in \mathbb{R}^n : Ay \leq b_x(\xi)\}$ is non-empty and bounded $\mathbb{P}$-almost surely for every $x \in \mathcal{X}$.*

Assumption 3 is reminiscent of the standard relatively complete recourse condition from two-stage stochastic programming, which requires the follower's feasible set to be non-empty $\mathbb{P}$-almost surely for every fixed decision of the leader. Assumption 3 is stronger in that it also requires boundedness. Note that boundedness holds whenever the recession cone $\{y \in \mathbb{R}^n : Ay \leq 0\}$ reduces to the singleton $\{0\}$ or, equivalently, whenever the positive cone $\{A^\top p : p \in \mathbb{R}^m_+\}$ of $A^\top$ coincides with $\mathbb{R}^n$. An important implication of Assumption 3 is that the follower's subproblem has a compact feasible set and is therefore solvable $\mathbb{P}$-almost surely for every $x \in \mathcal{X}$. This in turn implies that the optimal value of $\mathcal{P}$ exceeds that of $\mathcal{O}$. Note that if the follower's subproblem failed to be $\mathbb{P}$-almost surely solvable for a particular decision $x_\mathcal{P} \in \mathcal{X}$, then the pessimistic leader would be inclined to select $x_\mathcal{P}$ in order to force the supremum over y in $\mathcal{P}$ to $-\infty$, whereas the optimistic leader would select a decision $x_\mathcal{O} \in \mathcal{X}$, for which the infimum over y in $\mathcal{O}$ is finite.

Assumptions 1–3 are assumed to hold throughout the paper. We now derive reformulations of the optimistic and pessimistic bilevel programs $\mathcal{O}$ and $\mathcal{P}$, respectively, that are more amenable to analytical and numerical treatment. For ease of notation, we define

$$z^*(x) = \begin{array}{cl} \displaystyle\min_{y \in \mathcal{L}^2_n} & \mathbb{E}[c(\xi)^\top y(\xi)] \\ \text{s.t.} & Ay(\xi) \leq b_x(\xi) \quad \mathbb{P}\text{-a.s.} \end{array} \tag{2.1}$$

as the follower's expected cost under an optimal policy.

Proposition 2.1 (Elimination of the 'argmin' Operators).

(i) The optimistic bilevel program $\mathcal{O}$ is equivalent to

$$\begin{array}{cl} \displaystyle\inf_{x \in \mathcal{X},\, y \in \mathcal{L}^2_n} & q^\top x + \mathbb{E}[v(\xi)^\top y(\xi)] \\ \text{s.t.} & \mathbb{E}[c(\xi)^\top y(\xi)] \leq z^*(x) \\ & Ay(\xi) \leq b_x(\xi) \quad \mathbb{P}\text{-a.s.} \end{array} \tag{2.2}$$

(ii) The pessimistic bilevel program $\mathcal{P}$ is equivalent to

$$\min_{x \in \mathcal{X}} \begin{array}{cl} \displaystyle\sup_{y \in \mathcal{L}^2_n} & q^\top x + \mathbb{E}[v(\xi)^\top y(\xi)] \\ \text{s.t.} & \mathbb{E}[c(\xi)^\top y(\xi)] \leq z^*(x) \\ & Ay(\xi) \leq b_x(\xi) \quad \mathbb{P}\text{-a.s.} \end{array} \tag{2.3}$$

Proof. As the proofs of assertions (i) and (ii) are similar, we only prove assertion (ii). To this end, we first establish the following equality.

$$\mathbb{E}\Big[\inf_{y \in \mathbb{R}^n} \{c(\xi)^\top y : Ay \leq b_x(\xi)\}\Big] = \inf_{y \in \mathcal{L}^2_n} \Big\{\mathbb{E}[c(\xi)^\top y(\xi)] : Ay(\xi) \leq b_x(\xi) \ \mathbb{P}\text{-a.s.}\Big\} \tag{2.4}$$

It is clear that the left hand side of (2.4) is smaller or equal to the right hand side because the functions $y \in \mathcal{L}_n^2$ are required to square-integrable, a condition which is absent on the left hand side. To prove the converse inequality, note that the linear program inside the expectation on the left hand side of (2.4) is $\mathbb{P}$-almost surely solvable for every $x \in \mathcal{X}$ by virtue of Assumption 3. Thus, there exist finitely many matrices $L_j \in \mathbb{R}^{n \times m}$, one for every basis of the constraint matrix A, such that the basic solutions of the linear program are representable as $L_j b_x(\xi)$ for $j = 1, \ldots, J$. Moreover, for every $x \in \mathcal{X}$ the Borel sets

$$\begin{aligned}\Xi_j(x) = \Big\{ & \xi \in \mathbb{R}^k : AL_j b_x(\xi) \le b_x(\xi), \\ & c(\xi)^\top (L_j b_x(\xi) - L_{j'} b_x(\xi)) \le 0 \ \forall j' = 1, \ldots, j-1 \text{ with } AL_{j'} b_x(\xi) \le b_x(\xi), \\ & c(\xi)^\top (L_j b_x(\xi) - L_{j'} b_x(\xi)) < 0 \ \forall j' = j+1, \ldots, J \text{ with } AL_{j'} b_x(\xi) \le b_x(\xi) \Big\}\end{aligned}$$

form a partition of $\mathbb{R}^k$, and $y^\star(\xi) = L_j b_x(\xi)$ for $\xi \in \Xi_j(x)$ is $\mathbb{P}$-almost surely optimal on the left hand side of (2.4). Moreover, as ξ has finite second moments, we find

$$\begin{aligned}\mathbb{E}[\|y^\star(\xi)\|^2] &= \sum_{j=1}^{J} \mathbb{E}[\|L_j b_x(\xi)\|^2 \mathbf{1}_{\{\xi \in \Xi_j(x)\}}] \le \sum_{j=1}^{J} \mathbb{E}[\|L_j b_x(\xi)\|^2] \\ &= \sum_{j=1}^{J} L_j B_x \mathbb{E}[\xi\xi^\top] B_x^\top L_j^\top = \sum_{j=1}^{J} L_j B_x \Omega B_x^\top L_j^\top < \infty.\end{aligned}$$

Thus, $y^\star \in \mathcal{L}_n^2$ is feasible on the right hand side of (2.4). As the right hand side is no smaller than the left hand side, we conclude that $y^\star$ is also optimal on the right hand side, implying that (2.4) holds indeed as an equality.

The above reasoning shows that problem $\mathcal{P}$ is equivalent to

$$\begin{aligned}\min_{x \in \mathcal{X}} \quad \sup_{y \in \mathcal{L}_n^2} \quad & q^\top x + \mathbb{E}[v(\xi)^\top y(\xi)] \\ \text{s.t.} \quad & y \in \operatorname*{argmin}_{y' \in \mathcal{L}_n^2} \left\{ \mathbb{E}[c(\xi)^\top y'(\xi)] : Ay'(\xi) \le b_x(\xi) \ \mathbb{P}\text{-a.s.} \right\}.\end{aligned}$$

The claim then follows by noting that

$$y \in \operatorname*{argmin}_{y' \in \mathcal{L}_n^2} \left\{ \mathbb{E}[c(\xi)^\top y'(\xi)] : Ay'(\xi) \le b_x(\xi) \ \mathbb{P}\text{-a.s.} \right\}$$

if and only if $Ay(\xi) \le b_x(\xi)$ $\mathbb{P}$-a.s. and $\mathbb{E}[c(\xi)^\top y(\xi)] \le z^*(x)$. ∎

3. Decision Rule Bounds for Stochastic Programs. We first analyze the follower's subproblem (2.1), which constitutes a two-stage stochastic program that involves only wait-and-see decisions. Unfortunately, evaluating $z^*(x)$ for a fixed x is #P-hard even if ξ is uniformly distributed on the unit hypercube in $\mathbb{R}^k$ [13, 19]. Therefore, we can only hope to evaluate $z^*(x)$ approximately. In this paper we will compute upper and lower bounds on $z^*(x)$ by using the primal and dual linear decision rule approximations proposed in [23]. Linear decision rules have first been used to derive tractable upper bound approximations for *adaptive robust optimization problems* [3, 4]. These approximations were later extended to the realm of two- and multi-stage stochastic programming [6, 34]. In this paper we aim to extend these bounds to stochastic bilevel programs.

3.1. Primal Approximation. An upper bound on $z^*(x)$ may be obtained by restricting the set of square-integrable wait-and-see decision $y(\xi)$ in (2.1) to the set of *linear decision rules*, that is, all linear functions of the form $y(\xi) = Y\xi$ for some $Y \in \mathbb{R}^{n\times k}$. The resulting approximate problem involves only a finite number of decision variables (*i.e.*, the entries of the matrix Y) but still infinitely many constraints parameterized by the possible realizations of ξ. Nevertheless, the best linear decision rule can be computed efficiently.

PROPOSITION 3.1. *The stochastic program* (2.1) *is bounded above by*

$$z^u(x) = \min_{Y,\Lambda} \quad \mathrm{Tr}(\Omega C^\top Y) \quad \text{s.t.} \quad AY + \Lambda W = B_x, \quad \Lambda h \geq 0, \Lambda \geq 0. \tag{3.1}$$

Moreover, if (Y^*, Λ^*) *solves the linear program* (3.1)*, then* $y^*(\xi) = Y^*\xi$ *represents an optimal linear decision rule for* (2.1)*.*

Proposition 3.1 is borrowed from [23, § 2.2]. To keep the paper self-contained, we sketch its proof in Appendix A.

3.2. Dual Approximation. By restricting the *primal* decisions of the stochastic program (2.1) to linear decision rules, we obtained an *upper* bound on $z^*(x)$. Similarly, by restricting the *dual* decisions, we can construct an efficiently computable *lower* bound.

PROPOSITION 3.2. *The stochastic program* (2.1) *is bounded below by*

$$z^l(x) = \min_{Y,S} \quad \mathrm{Tr}(\Omega C^\top Y) \quad \text{s.t.} \quad AY + S = B_x, \quad (W - h\mathbf{e}_1^\top)\Omega S^\top \geq 0. \tag{3.2}$$

Proposition 3.1 is borrowed from [23, § 2.3]. To keep the paper self-contained, we sketch its proof in Appendix A.

REMARK 1. *Notice that* $z^u(x) - z^*(x) \geq 0$ *quantifies the loss of optimality incurred by the (primal) linear decision rule approximation. Unfortunately, this optimality gap is as hard to compute as* $z^*(x)$*. However, it can be estimated by* $z^u(x) - z^l(x)$*, which can be computed efficiently by solving the two tractable linear programs* (3.1) *and* (3.2)*.*

4. Decision Rule Bounds for Pessimistic Bilevel Programs. We now demonstrate that the decision rule approximations for stochastic programs can be extended to the realm of bilevel programming. We first show that $\mathcal{P}$ admits an upper and $\mathcal{O}$ a lower decision rule bound. Next, we argue that lower decision rule bounds on $\mathcal{P}$ are generically unbounded, whereas upper decision rule bounds on $\mathcal{O}$ are generically infeasible.

THEOREM 4.1. *The pessimistic version $\mathcal{P}$ of the two-stage stochastic bilevel program is bounded above by the optimal value of the MILP*

$$
\begin{array}{cl}
\min\limits_{x\in\mathcal{X},\,\lambda,\,Y,\,\Lambda,\,\Gamma,\,\beta,\,\gamma} & q^\top x + \mathbf{e}^\top \gamma + \mathrm{Tr}(\Omega C^\top Y) \\
\text{s.t.} & A^\top \Gamma W + \lambda C = V,\ \Gamma h \geq 0,\ \Gamma \geq 0,\ \lambda \geq 0 \\
& AY + \Lambda W = \sum_{i=1}^d B_i \beta_i,\ \Lambda h \geq 0,\ \Lambda \geq 0 \\
& \left.\begin{array}{l}
\beta_i \leq M x_i,\ \beta_i \geq \lambda - M(1-x_i) \\
\gamma_i \leq \mathrm{Tr}(\Omega B_i^\top \Gamma W) + M(1-x_i) \\
\gamma_i \geq \mathrm{Tr}(\Omega B_i^\top \Gamma W) - M(1-x_i) \\
0 \leq \beta_i \leq \lambda, \ \ -Mx_i \leq \gamma_i \leq M x_i
\end{array}\right\} \forall i = 1,\ldots,d,
\end{array}
\tag{$\mathcal{P}^u$}
$$

where $M > 0$ represents a big-M constant.

Proof. By Proposition 2.1, the pessimistic bilevel program $\mathcal{P}$ is equivalent to

$$
\begin{array}{lll}
\min\limits_{x\in\mathcal{X}} & \sup\limits_{y\in\mathcal{L}^2_n} & q^\top x + \mathbb{E}[v(\xi)^\top y(\xi)] \\
& \text{s.t.} & \mathbb{E}[c(\xi)^\top y(\xi)] \leq z^*(x) \\
& & Ay(\xi) \leq b_x(\xi) \quad \mathbb{P}\text{-a.s.},
\end{array}
\tag{4.1}
$$

where $z^*(x)$ stands for the optimal value of problem (2.1). Dualizing the inner maximization problem in (4.1) yields via weak duality the following upper bound on (4.1).

$$
\begin{array}{cl}
\inf\limits_{x\in\mathcal{X},\,\lambda,\,p\in\mathcal{L}^2_m} & q^\top x + \mathbb{E}[b_x(\xi)^\top p(\xi)] + \lambda z^*(x) \\
\text{s.t.} & A^\top p(\xi) + \lambda c(\xi) = v(\xi) \quad \mathbb{P}\text{-a.s.} \\
& p(\xi) \geq 0 \quad \mathbb{P}\text{-a.s.}, \quad \lambda \geq 0
\end{array}
\tag{4.2}
$$

As the stochastic program (2.1) is a minimization problem and as λ is non-negative, problem (4.2) can be re-expressed as

$$
\begin{array}{cl}
\inf\limits_{x\in\mathcal{X},\,\lambda,\,p\in\mathcal{L}^2_m,\,y\in\mathcal{L}^2_n} & q^\top x + \mathbb{E}[b_x(\xi)^\top p(\xi)] + \lambda\, \mathbb{E}[c(\xi)^\top y(\xi)] \\
\text{s.t.} & A^\top p(\xi) + \lambda c(\xi) = v(\xi) \quad \mathbb{P}\text{-a.s.} \\
& p(\xi) \geq 0 \quad \mathbb{P}\text{-a.s.}, \quad \lambda \geq 0 \\
& Ay(\xi) \leq b_x(\xi) \quad \mathbb{P}\text{-a.s.}
\end{array}
\tag{4.3}
$$

Note that (4.3) constitutes a non-convex optimization problem because its objective involves a product of λ and $y(\xi)$. The problem can be linearized, however, by applying the variable transformation $y(\xi) \leftarrow \lambda\, y(\xi)$, which results in the following exact reformulation.

$$
\begin{array}{cl}
\inf\limits_{x\in\mathcal{X},\,\lambda,\,p\in\mathcal{L}^2_m,\,y\in\mathcal{L}^2_n} & q^\top x + \mathbb{E}[b_x(\xi)^\top p(\xi)] + \mathbb{E}[c(\xi)^\top y(\xi)] \\
\text{s.t.} & A^\top p(\xi) + \lambda c(\xi) = v(\xi) \quad \mathbb{P}\text{-a.s.} \\
& p(\xi) \geq 0 \quad \mathbb{P}\text{-a.s.}, \quad \lambda \geq 0 \\
& Ay(\xi) \leq \lambda b_x(\xi) \quad \mathbb{P}\text{-a.s.}
\end{array}
\tag{4.4}
$$

The equivalence of (4.3) and (4.4) can be seen as follows. If (x, λ, p, y) is feasible in (4.3), then $(x, \lambda, p, \lambda y)$ is feasible in (4.4) with the same objective value, which implies that (4.4) provides a lower bound on (4.3). Next, consider any (x, λ, p, y) feasible in (4.4). If $\lambda > 0$, then $(x, \lambda, p, y/\lambda)$ is feasible in (4.3) with the same objective value. If $\lambda = 0$, on the other hand, then the last constraint in (4.4) implies that $y(\xi) = 0$ $\mathbb{P}$-almost surely because the recession cone $\{y \in \mathbb{R}^n : Ay \leq 0\}$ of the follower's feasible set coincides with the singleton $\{0\}$ due to Assumption 3. Thus, $(x, 0, p, y')$ is feasible in (4.3) for every $y' \in \mathcal{L}_n^2$ with $Ay'(\xi) \leq b_x(\xi)$ $\mathbb{P}$-almost surely, which exists due to Assumption 3 and an argument familiar from the proof of Proposition 2.1, and the objective value of (x, λ, p, y) in (4.4) coincides with the objective value of $(x, 0, p, y')$ in (4.3). This implies that (4.4) provides also an upper bound on (4.3).

For any fixed x, problem (4.4) can be viewed as a linear two-stage stochastic program with here-and-now decision λ and wait-and-see decisions $p(\xi)$ and $y(\xi)$. In analogy to Proposition 3.1, a conservative approximation of (4.4) is thus obtained by focusing on linear decision rules of the form $y(\xi) = Y\xi$ and $p(\xi) = P\xi$.

$$\begin{array}{cl} \min\limits_{x\in\mathcal{X},\, \lambda,\, P,\, Y,\, \Lambda,\, \Gamma} & q^\top x + \mathrm{Tr}(\Omega B_x^\top P) + \mathrm{Tr}(\Omega C^\top Y) \\ \text{s.t.} & A^\top P + \lambda C = V \\ & P = \Gamma W,\ \Gamma h \geq 0,\ \Gamma \geq 0,\ \lambda \geq 0 \\ & AY + \Lambda W = \lambda B_x,\ \Lambda h \geq 0,\ \Lambda \geq 0 \end{array} \tag{4.5}$$

The derivation of (4.5) parallels the proof of Proposition 3.1. Details are omitted for brevity. Note that the auxiliary variables Λ and Γ can be interpreted as the coefficient matrices of two 'slack'-decision rules. Indeed, the constraints in (4.4) ensure that $\Lambda\xi \geq 0$ and $\Gamma\xi \geq 0$ for all $\xi \in \Xi$. Eliminating the matrix variable P from (4.5) yields

$$\begin{array}{cl} \min\limits_{x\in\mathcal{X},\, \lambda,\, Y,\, \Lambda,\, \Gamma} & q^\top x + \mathrm{Tr}(\Omega B_x^\top \Gamma W) + \mathrm{Tr}(\Omega C^\top Y) \\ s.t. & A^\top \Gamma W + \lambda C = V,\ \Gamma h \geq 0,\ \Gamma \geq 0,\ \lambda \geq 0 \\ & AY + \Lambda W = \lambda B_x,\ \Lambda h \geq 0,\ \Lambda \geq 0. \end{array} \tag{4.6}$$

Unfortunately, the upper bound problem (4.6) still involves bilinear terms in the here-and-now decisions x and the dual variables (λ, Γ). As all components of x constitute binary (0-1) variables and as $B_x = \sum_{i=1}^m B_i x_i$ for some $B_i \in \mathbb{R}^{m\times k}$, $i = 1, \ldots, d$, however, (4.6) can be reformulated as the postulated MILP $\mathcal{P}^u$ by using standard linearization techniques. Specifically, the second term in the objective function of (4.6) can be expressed as $\mathbf{e}^\top \gamma$, where

$$\gamma_i = x_i \,\mathrm{Tr}(\Omega B_i^\top \Gamma W) \quad \Longleftrightarrow \quad \begin{cases} \gamma_i \leq \mathrm{Tr}(\Omega B_i^\top \Gamma W) + M(1 - x_i) \\ \gamma_i \geq \mathrm{Tr}(\Omega B_i^\top \Gamma W) - M(1 - x_i) \\ -Mx_i \leq \gamma_i \leq Mx_i \end{cases}$$

for every $i = 1, \ldots, d$, and $M > 0$ is a suitable big-M constant. Similarly, the product term λB_x in the last constraint of (4.6) can be expressed as $\sum_{i=1}^d B_i\, \beta_i$, where

$$\beta_i = \lambda x_i \quad \Longleftrightarrow \quad 0 \leq \beta_i \leq \lambda,\ \beta_i \leq Mx_i,\ \beta_i \geq \lambda - M(1 - x_i)$$

for every $i = 1, \ldots, d$, and $M > 0$ is a big-M constant. Thus, the claim follows. ■

Note that $\mathcal{P}^u$ can be solved efficiently with powerful MILP solvers such as CPLEX or Gurobi. If the cost functions $c(\xi)$ and $v(\xi)$ coincide (*i.e.*, if $C = V$), the preferences of the leader and the follower are perfectly aligned, in which case the bilevel program $\mathcal{P}$ reduces to the single-level program

$$\left.\begin{array}{cl} \min\limits_{x\in\mathcal{X},\, y\in\mathcal{L}_n^2} & q^\top x + \mathbb{E}[c(\xi)^\top y(\xi)] \\ \text{s.t.} & Ay(\xi) \leq b_x(\xi) \quad \mathbb{P}\text{-a.s.} \end{array}\right\} = \min_{x\in\mathcal{X}} q^\top x + z^\star(x). \tag{4.7}$$

The following corollary demonstrates that $\mathcal{P}^u$ also collapses to the primal linear decision rule approximation of (4.7) whenever $C = V$.

COROLLARY 4.2. *If $C = V$, then $\mathcal{P}^u$ is equivalent to* $\min_{x\in\mathcal{X}} q^\top x + z^u(x)$.

Proof. Set

$$\begin{aligned} \varphi(x,\lambda) \;=\; \min_{Y,\Lambda,\Gamma} \quad & \mathrm{Tr}(\Omega B_x^\top \Gamma W) + \mathrm{Tr}(\Omega C^\top Y) \\ \text{s.t.} \quad & A^\top \Gamma W + \lambda C = C \\ & \Gamma h \geq 0,\ \Gamma \geq 0 \\ & AY + \Lambda W = \lambda B_x,\ \Lambda h \geq 0,\ \Lambda \geq 0, \end{aligned}$$

and note that $\varphi(x,\lambda)$ is convex in λ for any fixed x. From the proof of Theorem 4.1 we know that $\mathcal{P}^u$ is equivalent to (4.6), which in turn is equivalent to $\min_{x\in\mathcal{X},\,\lambda\geq 0} q^\top x + \varphi(x,\lambda)$ whenever $C = V$. For $\theta \in (0,1)$, the substitutions $\Gamma \leftarrow \Gamma/(1-\theta)$, $\Lambda \leftarrow \Lambda/\theta$ and $Y \leftarrow Y/\theta$ yield

$$\varphi(x,\lambda) = (1-\theta)\overline{z}^u(x) + \theta z^u(x),$$

where

$$\overline{z}^u(x) = \min_{\Gamma} \left\{ \mathrm{Tr}(\Omega B_x^\top \Gamma W) : A^\top \Gamma W = C,\ \Gamma h \geq 0,\ \Gamma \geq 0 \right\},$$

while

$$z^u(x) = \min_{Y,\Lambda} \left\{ \mathrm{Tr}(\Omega C^\top Y) : AY + \Lambda W = B_x,\ \Lambda h \geq 0,\ \Lambda \geq 0 \right\} \tag{4.8}$$

denotes the primal linear decision rule bound on the follower's problem (2.1). In the following we will show that $\overline{z}^u(x) \geq z^u(x)$. By weak duality, we find

$$\begin{aligned} \overline{z}^u(x) &\geq \max_{Y,S,\lambda} \left\{ \mathrm{Tr}(\Omega C^\top Y) : AY + S = B_x,\ S\Omega W^\top \geq \lambda h^\top,\ \lambda \geq 0 \right\} \\ &= \max_{Y,S} \left\{ \mathrm{Tr}(\Omega C^\top Y) : AY + S = B_x,\ (W - h\mathbf{e}_1^\top)\Omega S^\top \geq 0 \right\}. \end{aligned} \tag{4.9}$$

Due to the structure of W and h stipulated in Assumption 1, the constraints $S\Omega W^\top \geq \lambda h^\top$ imply that $\lambda = S\Omega\mathbf{e}_1$. The equality in the second line of the above expression thus follows because $S\Omega\mathbf{e}_1 \geq 0$. Indeed, the constraint $(W - h\mathbf{e}_1^\top)\Omega S^\top \geq 0$ is equivalent to the requirement that each row of $S\Omega$ lies within the cone generated by Ξ and thus has a nonnegative first component.

Next, fix any (Y,Λ) feasible in (4.8) and set $S = \Lambda W$. The objective value of (Y,S) in (4.9) trivially coincides with the objective value of (Y,Λ) in (4.8). Moreover, (Y,S) is feasible in (4.9). To see this, we note that

$$(W - h\mathbf{e}_1^\top)\Omega S^\top = \mathbb{E}\left[(W\xi - h)(\Lambda W\xi)^\top\right] \geq 0,$$

where the equality holds because $\Omega = \mathbb{E}[\xi\xi^\top]$ and $\mathbf{e}_1^\top \xi = 1$ $\mathbb{P}$-a.s., while the inequality holds because $W\xi \geq h$ and $\Lambda W\xi \geq \Lambda h \geq 0$ $\mathbb{P}$-a.s. We conclude that (4.8) has the same objective function as (4.9) but a smaller feasible set (that is, when projected on the Y variables). Thus, the minimum of (4.8) is smaller or equal to the maximum of (4.9). This implies that $\overline{z}^u(x) \geq z^u(x)$, and the infimum of $\varphi(x,\lambda)$ over $\theta \in (0,1)$ is attained for $\theta \uparrow 1$.

For $\theta > 1$, on the other hand, $1-\theta$ is negative, and thus the substitutions $\Gamma \leftarrow \Gamma/(1-\theta)$, $\Lambda \leftarrow \Lambda/\theta$ and $Y \leftarrow Y/\theta$ yield

$$\varphi(x,\lambda) = (1-\theta)\underline{z}^u(x) + \theta z^u(x),$$

where

$$\underline{z}^u(x) = \max_{\Gamma} \left\{ \mathrm{Tr}(\Omega B_x^\top \Gamma W) : A^\top \Gamma W = C,\ \Gamma h \le 0,\ \Gamma \le 0 \right\}.$$

In the following we will show that $z^u(x) \ge \underline{z}^u(x)$. By weak duality, we have

$$\begin{aligned} \underline{z}^u(x) &\le \min_{Y,\, S,\, \lambda} \left\{ \mathrm{Tr}(\Omega C^\top Y) : AY + S = B_x,\ S\Omega W^\top \le \lambda h^\top,\ \lambda \ge 0 \right\} \\ &= \min_{Y,\, S} \left\{ \mathrm{Tr}(\Omega C^\top Y) : AY + S = B_x,\ (W - h\mathbf{e}_1^\top)\Omega S^\top \ge 0 \right\} = z^l(x), \end{aligned}$$

where the first equality holds because $S\Omega W^\top \le \lambda h^\top$ implies $\lambda = S\Omega \mathbf{e}_1 \ge 0$, and the second equality follows from (3.2). Propositions 3.1 and 3.2 then allow us to conclude that $\underline{z}^u(x) \le z^l(x) \le z^u(x)$, and thus the infimum of $\varphi(x, \lambda)$ over $\theta > 1$ is attained for $\theta \downarrow 1$.

The convexity of $\varphi(x, \lambda)$ in λ finally guarantees that the minimum of $\varphi(x, \lambda)$ over $\lambda \ge 0$ is attained at $\lambda = 1$. A direct calculation confirms that $\varphi(x, 1) = z^u(x)$, and thus $\mathcal{P}^u$ is equivalent to $\min_{x \in \mathcal{X},\, \lambda \ge 0} q^\top x + z^u(x)$. ■

As for plain vanilla stochastic programs, decision rule techniques can also be leveraged to obtain *lower* bounds on stochastic bilevel programs.

Theorem 4.3. *The optimistic version $\mathcal{O}$ of the two-stage stochastic bilevel program is bounded below by the optimal value of the MILP*

$$\begin{array}{cl} \displaystyle\min_{x \in \mathcal{X},\, Y,\, S,\, P,\, \beta} & q^\top x + \mathrm{Tr}(\Omega V^\top Y) \\ \text{s.t.} & \mathrm{Tr}(\Omega C^\top Y) \le \mathbf{e}^\top \beta \\ & AY + S = B_x,\ (W - h\mathbf{e}_1^\top)\Omega S^\top \ge 0 \\ & A^\top P = C,\ (W - h\mathbf{e}_1^\top)\Omega P^\top \le 0 \\ & \left. \begin{array}{l} \beta_i \le \mathrm{Tr}(\Omega B_i^\top P) + M(1 - x_i) \\ \beta_i \ge \mathrm{Tr}(\Omega B_i^\top P) - M(1 - x_i) \\ -Mx_i \le \beta_i \le Mx_i \end{array} \right\} \forall i = 1, \ldots, d, \end{array} \qquad (\mathcal{O}^l)$$

where $M > 0$ represents a big-M constant.

Proof. By Proposition 2.1, the optimistic bilevel program $\mathcal{O}$ is equivalent to (2.2), which constitutes an ordinary two-stage stochastic program. Applying the dual linear decision rule approximation of Proposition 3.2 then yields the following lower bound on (2.2).

$$\begin{array}{cl} \displaystyle\min_{x \in \mathcal{X},\, Y,\, S} & q^\top x + \mathrm{Tr}(\Omega V^\top Y) \\ \text{s.t.} & \mathrm{Tr}(\Omega C^\top Y) \le z^*(x) \\ & AY + S = B_x,\ (W - h\mathbf{e}_1^\top)\Omega S^\top \ge 0 \end{array} \qquad (4.10)$$

Note that (4.10) is still intractable as the evaluation of $z^*(x)$ is #P-hard. However, we can further relax (4.10) by replacing $z^*(x)$ with its efficiently computable upper bound $z^u(x)$. In the following, we re-express $z^u(x)$ as

$$\begin{aligned} z^u(x) &= \max_{P,\, \lambda} \left\{ \mathrm{Tr}(\Omega B_x^\top P) : A^\top P = C,\ P\Omega W^\top + \lambda h^\top \le 0,\ \lambda \ge 0 \right\} \\ &= \max_{P} \left\{ \mathrm{Tr}(\Omega B_x^\top P) : A^\top P = C,\ (W - h\mathbf{e}_1^\top)\Omega P^\top \le 0 \right\}. \end{aligned} \qquad (4.11)$$

Here, the first equality follows from strong linear programming duality, which holds because (4.11) constitutes the dual linear decision rule approximation of the stochastic program *dual* to (2.1), which can be shown to be feasible by appealing to Assumption 3. To justify the second equality, we note that the constraints $P\Omega W^\top + h\lambda^\top \leq 0$ and the geometry of Ξ implied by Assumption 1 require that $\lambda = -P\Omega \mathbf{e}_1 \geq 0$ (*cf.* the corresponding argument in the proof of Corollary 4.2). Substituting (4.11) into (4.10) then yields

$$\begin{array}{ll} \min\limits_{x\in\mathcal{X},\, Y,\, S,\, P} & q^\top x + \mathrm{Tr}(\Omega V^\top Y) \\ \text{s.t.} & \mathrm{Tr}(\Omega C^\top Y) \leq \mathrm{Tr}(\Omega B_x^\top P) \\ & AY + S = B_x,\ (W - h\mathbf{e}_1^\top)\Omega S^\top \geq 0 \\ & A^\top P = C,\ (W - h\mathbf{e}_1^\top)\Omega P^\top \leq 0. \end{array} \tag{4.12}$$

Unfortunately, the lower bound problem (4.12) still involves bilinear terms in x and P. As in the proof of Theorem 4.1, however, (4.12) can be reformulated as the explicit MILP $\mathcal{O}^l$ by using similar linearization techniques as in the proof of Theorem 4.1. Details are omitted for brevity. ■

Like the upper bounding problem $\mathcal{P}^u$ from Theorem 4.1, the lower bounding problem $\mathcal{O}^l$ can be solved efficiently with commercial MILP solvers. As in Corollary 4.2, one can further show that $\mathcal{O}^l$ collapses to the dual linear decision rule bound on the stochastic program (4.7) whenever $C = V$.

COROLLARY 4.4. *If* $C = V$*, then* $\mathcal{O}^l$ *is equivalent to* $\min_{x\in\mathcal{X}} q^\top x + z^l(x)$.

Proof. From the proof of Theorem 4.3 we know that $\mathcal{O}^l$ is equivalent to

$$\begin{array}{ll} \min\limits_{x\in\mathcal{X},\, Y,\, S} & q^\top x + \mathrm{Tr}(\Omega C^\top Y) \\ \text{s.t.} & \mathrm{Tr}(\Omega C^\top Y) \leq z^u(x) \\ & AY + S = B_x,\ (W - h\mathbf{e}_1^\top)\Omega S^\top \geq 0, \end{array} \tag{4.13}$$

where we have used that $C = V$. This problem is feasible due to Assumption 3. Therefore, the constraint $\mathrm{Tr}(\Omega C^\top Y) \leq z^u(x)$ cannot be binding at optimality and may be removed without impacting the problem's optimal value. Comparison with (3.2) then indicates that (4.13) is indeed equivalent to $\min_{x\in\mathcal{X}} q^\top x + z^l(x)$, and thus the claim follows. ■

In summary, the upper bound $\mathcal{P}^u$ on $\mathcal{P}$ was constructed by solving the two-stage stochastic program (4.4) in *primal* linear decision rules. Similarly, the lower bound $\mathcal{O}^l$ on $\mathcal{O}$ was obtained by solving the two-stage stochastic program (2.2) in *dual* linear decision rules. It is now tempting to construct a lower bound $\mathcal{P}^l$ on $\mathcal{P}$ by solving (4.4) in *dual* linear decision rules and to construct an upper bound $\mathcal{O}^u$ on $\mathcal{O}$ by solving (2.2) in *primal* linear decision rules. However, $\mathcal{P}^l$ is equivalent to a variant of (2.3) where the inner maximization problem over $y \in \mathcal{L}_n^2$ is solved in primal linear decision rules and $z^*(x)$ is replaced with $z^l(x)$. The inner problem is thus infeasible and evaluates to $-\infty$ unless $z^u(x) = z^l(x)$, that is, unless both primal and dual linear decision rules happen to be optimal. This implies that $\mathcal{P}^l$ generically yields the trivial lower bound $-\infty$. Similarly, $\mathcal{O}^u$ is equivalent to a variant of (2.2) that is solved in primal linear decision rules and where $z^*(x)$ is replaced with $z^l(x)$. This problem is also infeasible unless $z^u(x) = z^l(x)$, and thus $\mathcal{O}^u$ generically provides the trivial upper bound $+\infty$. Note that our inability to obtain meaningful lower bounds on $\mathcal{P}$ and upper bounds on $\mathcal{O}$ is fundamental and not due to the relative inflexibility of *linear* decision rules. Indeed, even more flexible *non-linear* decision rule approximations of the problems (2.3) and (2.2) are necessarily infeasible unless the employed class of decision rules happens to contain the optimal wait-and-see policies. However, as explained in the following section, non-linear decision rules can be used to obtain stronger upper bounds on $\mathcal{P}$ and stronger lower bounds on $\mathcal{O}$.

5. Extension to Non-Linear Decision Rules. Theorems 4.1 and 4.3 imply that

$$\min \mathcal{O}^l \;\leq\; \min \mathcal{O} \;\leq\; \min \mathcal{P} \;\leq\; \min \mathcal{P}^u,$$

where $\min \mathcal{O}^l$ denotes the optimal value of the minimization problem $\mathcal{O}^l$ etc. Recall that $\mathcal{O}^l$ and $\mathcal{P}^u$ constitute efficiently solvable MILPs. We emphasize that any solution of $\mathcal{P}^u$ is feasible in $\mathcal{P}$ and incurs a cost that is bracketed by $\min \mathcal{P}$ and $\min \mathcal{P}^u$ and, *a fortiori*, by the computable bounds $\min \mathcal{O}^l$ and $\min \mathcal{P}^u$. To improve the gap between these bounds, we may use more flexible non-linear decision rules, that is, linear combinations of the components of a non-linear vector-valued lifting operator $L(\xi)$. If there exists a linear retraction operator R with $R(L(\xi)) = \xi$ and if the convex hull of $L(\Xi)$ admits an explicit polyhedral representation, then the best primal and dual decision rules in the class induced by $L(\xi)$ can be computed efficiently [17, 18].

Specifically, assume that the lifting operator $L : \mathbb{R}^k \to \mathbb{R}^{k'}$ and the corresponding inverse retraction operator $R : \mathbb{R}^{k'} \to \mathbb{R}^k$ have the following properties:

(**P1**) L is Lipschitz continuous and satisfies $\mathbf{e}_1^\top L(\xi) = 1$ for all $\xi \in \Xi$;
(**P2**) R is linear;
(**P3**) $R(L(\xi)) = \xi$ for every $\xi \in \mathbb{R}^k$;
(**P4**) the component mappings of L are linearly independent. This means that if any $w \in \mathbb{R}^{k'}$ satisfies $w^\top L(\xi) = 0$ $\mathbb{P}$-almost surely, then we have $w = 0$.

Moreover, suppose that the lifted support $L(\Xi)$ admits a tight polyhedral outer approximation as specified in the following assumption.

Assumption 4 (Lifted Support). *Assume that there exists a polytope $\Xi' = \{\xi' \in \mathbb{R}^{k'} : W'\xi' \geq h'\}$ with $W' \in \mathbb{R}^{\ell' \times k'}$ and $h' \in \mathbb{R}^{\ell'}$ such that $L(\Xi) \subseteq \Xi'$ and $R(\Xi') = \Xi$.*

If Assumptions 1–4 hold, while the lifting operator L and its corresponding retraction operator R satisfy the properties (**P1**)–(**P4**), then the linear decision rule bounds from Section 4 can be improved by the following non-linear decision rule bounds.

Theorem 5.1. *Set $\Omega' = \mathbb{E}[L(\xi)L(\xi)^\top]$, $C' = CR$, $V' = VR$ and $B'_x = B_x R$ for all $x \in \mathcal{X}$. Denote by $\mathcal{LP}^u$ (or by $\mathcal{LO}^l$) the lifted variant of the pessimistic bilevel program $\mathcal{P}^u$ (or the optimistic bilevel program $\mathcal{O}^l$) obtained via the substitutions $C' \leftarrow C$, $V' \leftarrow V$, $B'_x \leftarrow B_x$ for all $x \in \mathcal{X}$, $W' \leftarrow W$, $h' \leftarrow h$, $\Omega' \leftarrow \Omega$. Then, we have*

$$\min \mathcal{O}^l \;\leq\; \min \mathcal{LO}^l \;\leq\; \min \mathcal{O} \;\leq\; \min \mathcal{P} \;\leq\; \min \mathcal{LP}^u \;\leq\; \min \mathcal{P}^u.$$

Proof. Theorem 5.1 generalizes a similar result for two-stage stochastic programs to the realm of bilevel programming; see [17, § 2]. The idea is to reformulate the bilevel programs $\mathcal{O}$ and $\mathcal{P}$ in terms of the lifted random vector $\xi' = L(\xi)$ and to solve the resulting lifted reformulations in decision rules that are linear in ξ'. This is equivalent to solving the original bilevel programs $\mathcal{O}$ and $\mathcal{P}$ in non-linear decision rules representable as linear combinations of the components of L. The proof then follows immediately from Theorems 4.1 and 4.3. Details are omitted for brevity. ∎

Theorem 5.1 guarantees that the non-linear decision rule bounds $\min \mathcal{LO}^l$ and $\min \mathcal{LP}^u$ are at least as strong as the linear decision rule bounds $\min \mathcal{O}^l$ and $\min \mathcal{P}^u$.

While the gaps between $\mathcal{LO}^l$ and $\mathcal{O}$ or between $\mathcal{P}$ and $\mathcal{LP}^u$ can be systematically reduced by optimizing over more expressive classes of non-linear decision rules, the gap between $\mathcal{O}$ and $\mathcal{P}$ is beyond our control. Fortunately, $\mathcal{O}$ and $\mathcal{P}$ share the same optimal value whenever the follower's optimal solution is $\mathbb{P}$-almost unique, which is often the case when $\mathbb{P}$ is absolutely continuous with respect to the Lebesgue measure on Ξ. Several examples of lifting and retraction operators that satisfy the properties (**P1**)–(**P4**) and that admit a lifted support as described in Assumption 4 are given in [17].

Example 2. *Assume that $\mathbb{P}$ represents the uniform distribution on the rectangular support set $\Xi = \{\xi \in \mathbb{R}^k : \underline{\xi} \leq \xi \leq \overline{\xi}\}$ with $\underline{\xi}_1 = 1 = \overline{\xi}_1$, which ensures that $\xi_1 = 1$ for all $\xi \in \Xi$. Thus, the expected value of ξ_i is given by $\mu_i = (\underline{\xi}_1 + \overline{\xi}_i)/2$ for all i. Segregated linear decision rules are induced by the lifting operator $L : \mathbb{R}^k \to \mathbb{R}^{2k-1}$ defined through $L_1(\xi) = \xi_1$, $L_{2i-2}(\xi) = \min\{\xi_i, \mu_i\}$ and $L_{2i-1}(\xi) = \max\{\xi_i - \mu_i, 0\}$ for all $i = 2, \dots, k$. The corresponding retraction operator $R : \mathbb{R}^{2k-1} \to \mathbb{R}^k$ is given by*

$$R = \begin{pmatrix} 1 & & & & & \\ & 1 & 1 & & & \\ & & & \ddots & & \\ & & & & 1 & 1 \end{pmatrix} \in \mathbb{R}^{k \times (2k-1)}.$$

*It is easy to verify that this pair of lifting and retraction operators satisfies the properties (**P1**)–(**P4**). Moreover, a lifted support set obeying Assumption 4 can be defined as*

$$\Xi' = \big\{\xi' \in \mathbb{R}^{2k-1} \; : \; \begin{array}{l} \xi'_1 \leq 1, \; \xi'_i \geq 1, \\ \xi'_{2i-2} \leq \mu_i, \; \xi'_{2i-1} \geq 0, \; \xi'_{2i-2} - \xi'_{2i-1} \geq \underline{\xi}_i \;\; \forall i = 2, \dots, k \big\}. \end{array}$$

For further details on segregated linear decision rules see [17, § 4] or [18, § 4].

We remark that if the follower's decisions are restricted to polynomial decision rules, then the optimistic bilevel program $\mathcal{O}$ can be expressed as a stochastic mathematical program with semi-infinite equilibrium constraints of the type considered in [24]. Such problems can be solved approximately even if Assumption 1 fails to hold by leveraging an entropic approximation together with a constraint randomization.

6. Exact MILP Reformulations for Discrete Distributions. To showcase the efficacy of the decision rule bounds developed in Sections 4 and 5, it is instructive to investigate the optimistic and pessimistic stochastic bilevel programs $\mathcal{O}$ and $\mathcal{P}$, respectively, under an *empirical distribution* of the form

$$\widehat{\mathbb{P}}_N = \frac{1}{N} \sum_{j=1}^{N} \delta_{\widehat{\xi}_j},$$

which assigns equal probabilities to N sample points $\widehat{\xi}_j \in \mathbb{R}^k$, $j = 1, \dots, N$. For $\mathbb{P} = \widehat{\mathbb{P}}_N$, the bilevel programs $\mathcal{O}$ and $\mathcal{P}$ admit exact MILP reformulations. This result is formalized in Theorems 6.1 and 6.2 below, which constitute simple corollaries of

Theorems 4.1 and 4.3, respectively. Thus, their proofs are omitted for brevity.

THEOREM 6.1. *If $\mathbb{P}$ coincides with an empirical distribution $\widehat{\mathbb{P}}_N$, then the pessimistic version $\mathcal{P}$ of the two-stage stochastic bilevel program is equivalent to the MILP*

$$\begin{array}{ll} \min\limits_{x\in\mathcal{X},\,\lambda,\,y,\,p,\,\beta,\,\gamma} & q^\top x + \frac{1}{N}\sum_{j=1}^N\sum_{i=1}^d \gamma_{ij} + \frac{1}{N}\sum_{j=1}^N (C\widehat{\xi}_j)^\top y_j \\ \text{s.t.} & \left.\begin{array}{l} A^\top p_j + \lambda C\widehat{\xi}_j = V\widehat{\xi}_j,\ p_j \geq 0 \\ Ay_j \leq \sum_{i=1}^d \beta_i B_i \widehat{\xi}_j,\ \lambda \geq 0 \\ \beta_i \leq M x_i,\ \beta_i \geq \lambda - M(1-x_i) \\ \gamma_{ij} \leq p_j^\top B_i \widehat{\xi}_j + M(1-x_i) \\ \gamma_{ij} \geq p_j^\top B_i \widehat{\xi}_j - M(1-x_i) \\ 0 \leq \beta_i \leq \lambda,\ \ -Mx_i \leq \gamma_{ij} \leq M x_i \end{array}\right\} \begin{array}{l} \forall i = 1,\ldots,d, \\ \forall j = 1,\ldots,N, \end{array} \end{array} \tag{$\widehat{\mathcal{P}}$}$$

where $M > 0$ represents a big-M constant.

Note that problem $\widehat{\mathcal{P}}$ involves d binary variables as well as $\mathcal{O}(N(d+n+m))$ continuous variables and constraints, that is, the size of $\widehat{\mathcal{P}}$ grows linearly with the number of atoms N of the empirical distribution. In contrast, the pessimistic decision rule problem $\mathcal{P}^u$ involves d binary variables as well as $\mathcal{O}(nk+mk+m\ell+d)$ continuous variables and constraints. If the the atoms of $\widehat{\mathbb{P}}$ are obtained by sampling from a (continuous) distribution $\mathbb{P}$, then a sample size $N \gg \max\{k,\ell\}$ may be needed to guarantee that $\widehat{\mathcal{P}}$ provides a reasonable approximation for $\mathcal{P}$. To be able to offer a similar approximation quality as $\mathcal{P}^u$, problem $\widehat{\mathcal{P}}$ must therefore typically be much larger than $\mathcal{P}^u$. Note also that $\mathcal{P}^u$ is guaranteed to provide an upper bound on $\mathcal{P}$, while $\widehat{\mathcal{P}}$ can systematically overestimate or underestimate $\mathcal{P}$; see Section 7.

If $N = 1$ and $\widehat{\xi}_1 = \mathbb{E}[\xi]$, then $\widehat{\mathcal{P}}$ simplifies to the deterministic nominal counterpart of $\mathcal{P}$, which involves d binary variables as well as $\mathcal{O}(n+m+d)$ continuous variables and constraints. Under the reasonable assumption that the numbers of uncertain parameters k and support constraints ℓ are polynomially bounded in d, n and m, the decision rule problem $\mathcal{P}^u$ is therefore only polynomially larger than its corresponding nominal counterpart.

THEOREM 6.2. *If $\mathbb{P}$ coincides with an empirical distribution $\widehat{\mathbb{P}}_N$, then the optimistic version $\mathcal{O}$ of the two-stage stochastic bilevel program is equivalent to the MILP*

$$\begin{array}{ll} \min\limits_{x\in\mathcal{X},\,y,\,\lambda,\,\beta} & q^\top x + \frac{1}{N}\sum_{j=1}^N (V\widehat{\xi}_j)^\top y_j \\ \text{s.t.} & \sum_{j=1}^N (C\widehat{\xi}_j)^\top y_j + \sum_{j=1}^N\sum_{i=1}^d \beta_{ij} \leq 0 \\ & \left.\begin{array}{l} C\widehat{\xi}_j + A^\top \lambda_j = 0,\ \lambda_j \geq 0 \\ Ay_i \leq \sum_{i=1}^n x_i B_i \widehat{\xi}_j \\ \beta_{ij} \leq \lambda_i^\top B_i \widehat{\xi}_j + M(1-x_i) \\ \beta_{ij} \geq \lambda_i^\top B_i \widehat{\xi}_j - M(1-x_i) \\ -Mx_i \leq \beta_{ij} \leq M x_i \end{array}\right\} \begin{array}{l} \forall i = 1,\ldots,d, \\ \forall j = 1,\ldots,N, \end{array} \end{array} \tag{$\widehat{\mathcal{O}}$}$$

where $M > 0$ represents a big-M constant.

Note that problem $\widehat{\mathcal{O}}$ displays the same scaling behavior as $\widehat{\mathcal{P}}$ and must therefore typically be much larger than $\mathcal{O}^l$ to be able to offer a similar approximation quality. Moreover, if $N = 1$ and $\widehat{\xi}_1 = \mathbb{E}[\xi]$, then $\widehat{\mathcal{O}}$ simplifies to the deterministic nominal counterpart of $\mathcal{O}$, which suggests that the decision rule problem $\mathcal{O}^l$ is typically only polynomially larger than its corresponding nominal counterpart.

7. Facility Location Problem of a Market Entrant. Consider a market for a homogeneous good that consists of d demand locations. Define $s \in \{0,1\}^d$ through

$s_i = 1$ if location i can accommodate a retail store; $= 0$ otherwise. Company A owns retail stores at a subset of these eligible locations. Thus, we may define $t \in \{0,1\}^d$, $t \leq s$, through $t_i = 1$ if company A owns a store at location i; $= 0$ otherwise. Company B is a market entrant envisaging to build at most r new retail stores. This company will represent the leader in the bilevel program to be formulated. We assume that there can be only one store per location and that new stores can only be built in eligible locations that are not yet occupied by company A. Formally, company B's strategy can be encoded by a vector $x \in \{0,1\}^d$, where $x_i = 1$ if a new store is built at location i; $= 0$ otherwise. The set of feasible strategies is thus given by

$$\mathcal{X} = \{x \in \{0,1\}^d : \|x\|_1 \leq r,\ x + t \leq s\}.$$

We denote by ξ_i the demand at location i. As the good is homogeneous, the customers are inclined to make their purchases at the nearest store irrespective of its owner. We assume, however, that the stores have finite capacity, that is, the store at location i (if it exists) can sell at most b_i units of the good. A fictitious aggregate customer aiming to serve the demand at all locations must thus solve the transportation problem

$$\begin{array}{lll} \min\limits_{y \geq 0} & \sum\limits_{i=1}^{d} \sum\limits_{j=1}^{d} c_{ij}\, y_{ij} & \\ \text{s.t.} & \sum\limits_{i=1}^{d} y_{ij} \geq \xi_j & \forall j = 1, \ldots, d \\ & \sum\limits_{j=1}^{d} y_{ij} \leq b_i(x_i + t_i) & \forall i = 1, \ldots, d, \end{array} \tag{7.1}$$

where $y_{ij} \in \mathbb{R}^{d \times d}$ denotes the amount of the good that is shipped from location i to location j, and c_{ij} stands for the unit transportation costs. The aggregate customer will play the role of the follower in the bilevel program to be formulated. In the sequel we denote by $Y^*(x, \xi)$ the set of minimizers of (7.1) for a given strategy x of company B and a fixed demand pattern ξ. We assume that the largest possible total demand $\max_{\xi \in \Xi} \|\xi\|_1$ is smaller or equal to the capacity $\sum_{i=1}^d b_i t_i$ of all existing stores of company A. Thus, (7.1) is feasible and in fact solvable (*i.e.*, $Y^*(x, \xi) \neq \emptyset$) for every $x \in \mathcal{X}$ and $\xi \in \Xi$.

While the customers make their choices under full knowledge of the realized demand ξ, company B must select a layout of stores $x \in \mathcal{X}$ solely based on knowledge of the demand distribution $\mathbb{P}$. If q_i represents the cost of opening a store at location i and v_i stands for the revenue that company B earns form selling one unit of the good at location i (note that $v_i = 0$ if $t_i = 1$), then company B faces the decision problem

$$\min_{x \in \mathcal{X}}\ q^\top x + \inf_{y \in \mathcal{L}^2_n} \Big/ \sup_{y \in \mathcal{L}^2_n} \left\{ \mathbb{E}[v^\top y(\xi)] : y(\xi) \in Y^*(x, \xi)\ \mathbb{P}\text{-a.s.} \right\}, \tag{7.2}$$

which comes in an optimistic (inf) and a pessimistic (sup) version. The market entrant can use this model to determine how many stores to build and where to site them. As Assumptions 1, 2 and 3 hold, problem (7.2) can be bounded from above and below by the optimal values of the MILPs $\mathcal{P}^u$ and $\mathcal{O}^l$, respectively.

Consider now a particular instance of the facility location problem. Assume that the demand locations correspond to the $d = 15$ most populous North American cities of the SGB128 dataset.[1] The names and populations of these cities as well as all city-to-city road distances are reported in Table 7.1. Assume further that the uncertain demands of different cities are independent and that the demand ξ_i of city i follows a uniform distribution on $[\frac{1}{2}\mu_i, \frac{3}{2}\mu_i]$, where the expected demand μ_i grows affinely with the population from 50 (for the smallest city) to 100 (for the largest city).

All cities are eligible to accommodate retail stores, and the existing company A already operates retail stores in Saint Louis and San Jose with equal capacities $b_i = \frac{3}{4}\sum_{j=1}^d \mu_j$ for $i \in \{5, 9\}$. These two locations were chosen so as to minimize the customers' overall expected transportation costs in the absence of company B. Note that the total capacity of company A suffices to serve all North American customers in every possible demand scenario, and thus the follower's subproblem is always feasible.

Company B can establish at most $r = 3$ new retail stores in different cities not yet occupied by company A. The fixed cost of opening a new store with capacity $b_i = \frac{1}{2}\sum_{j=1}^d \mu_j$ is $q_i = 500$, and the revenue from selling one item amounts to $v_i = 5$ irrespective of the location $i = 1, \ldots, d$. If company B opens up three facilities, then it can serve all North American customers on its own in every possible scenario.

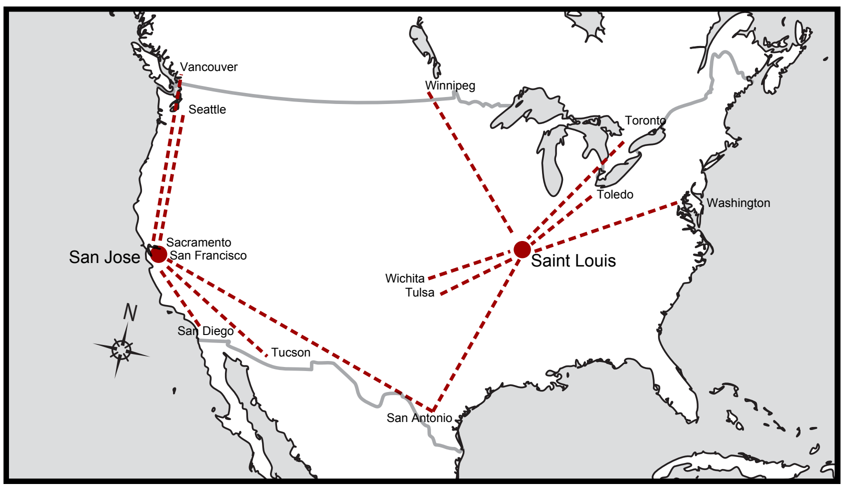

FIG. 7.1. *Assignments of customers to retail stores in the nominal demand scenario and before the market entry of company B. The stores of company A are visualized by red dots.*

The situations before and after the arrival of the market entrant are visualized in Figures 7.1 and 7.2, respectively. Note that customers are always served from the nearest store that has available capacity. The situation shown in Figure 7.2 is based on the assumption that company B solves the conservative upper bound model $\mathcal{P}^u$ to decide where to establish new stores. This model recommends to open up stores in San Francisco, Toledo and Tucson, which enables company B to serve all customers except those based in San Jose, Saint Louis, Tulsa and Wichita. Thus, company B manages to seize 70% of the market.

In order to assess the accuracy and runtime of the proposed approximations more systematically, we solve randomly generated instances of the facility location problem with $d \in \{15, 20, 25, 30, 35\}$ locations. Each location i either represents a demand center (if $s_i = 0$) or a candidate location for a retail store (if $s_i = 1$), where $\|s\|_1 \in$

[1] The SGB128 dataset is available from John Burkardt's web page at `https://people.sc.fsu.edu/~jburkardt/datasets/cities/cities.html`.

Cities	**Abbr.**	Tor.	S.A.	S.D.	S.J.	S.F.	Win.	Was.	Van.	Sea.	Tuc.	Sac.	Tul.	Wic.	S.L.	Tol.	**Population**
Toronto	Tor.	0	1,686	2,581	2,668	480	1,387	480	2,723	2,587	2,207	2,549	1,151	1,198	758	288	2,615,000
San Antonio	S.A.	1,686	0	1,286	1,692	1,746	1,555	1,609	2,281	2,171	867	1,736	531	630	910	1,399	1,409,000
San Diego	S.D.	2,581	1,286	0	460	505	2,017	2,703	1,397	1,265	467	504	1,416	1,410	1,858	2,306	1356000
San Jose	S.J.	2,668	1,692	460	0	48	1,916	2,846	982	840	827	120	1,690	1,637	2,082	2,393	998,537
San Francisco	S.F.	480	1,746	505	48	0	1,561	2,838	950	810	869	87	1,732	1,785	2,073	2,361	837,442
Winnipeg	Win.	1,387	1,555	2,017	1,916	1,561	0	1,916	1,438	1,419	1,945	1,797	1,020	940	1,015	1,114	663,615
Washington	Was.	480	1,609	2,703	2,846	2,838	1,916	0	2,897	2,784	2,278	2,730	1,211	1,279	878	470	658,893
Vancouver	Van.	2,723	2,281	1,397	982	950	1,438	2,897	0	143	1,664	588	2,126	1,952	2,219	2,427	603,500
Seattle	Sea.	2,587	2,171	1,265	840	810	1,419	2,784	143	0	1,529	752	1,990	1,830	2,118	2,292	652,405
Tucson	Tuc.	2,207	867	467	827	869	1,945	2,278	1,664	1,529	0	871	1,057	1,003	1,449	1,919	526,116
Sacramento	Sac.	2,549	1,736	504	120	87	1,797	2,730	588	752	871	0	1,726	1,637	1,971	2,273	479,686
Tulsa	Tul.	1,151	531	1,416	1,690	1,732	1,020	1,211	2,126	1,990	1,057	1,726	0	176	394	864	398,121
Wichita	Wic.	1,198	630	1,410	1,637	1,785	940	1,279	1,952	1,830	1,003	1,637	176	0	636	904	386,552
Saint Louis	S.L.	758	910	1,858	2,082	2,073	1015	878	2,219	2,118	1,449	1,971	394	636	0	470	318,416
Toledo	Tol.	288	1,399	2,306	2,393	2,361	1,114	470	2,427	2,292	1,919	2,273	864	904	470	0	282,313

Table 7.1

Road distances (in miles) and populations of the 15 most populous cities in the SGB128 dataset. The cities are ordered by population size.

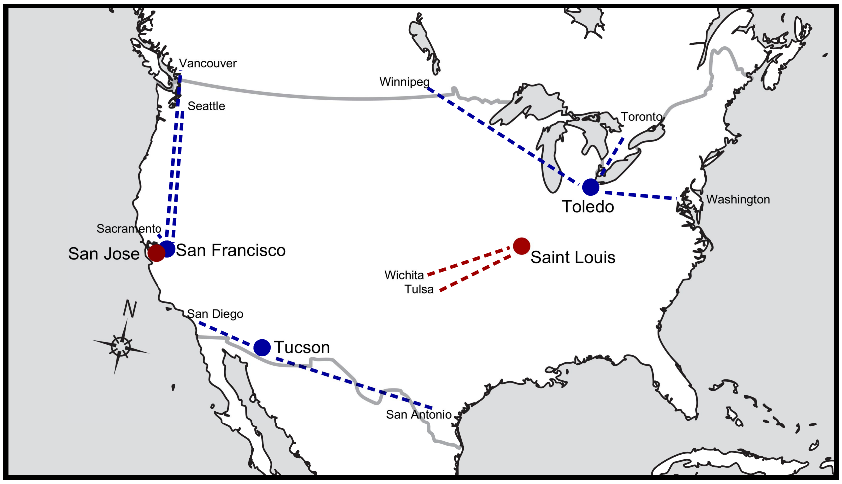

FIG. 7.2. *Assignments of customers to stores in the nominal scenario and after the market entry of company B. The stores of companies A and B are visualized by red and blue dots, respectively.*

$\{5, 10\}$. Candidate location i may already accommodate a store of company A (if $t_i = 1$) or may be available to company B (if $t_i = 0$), where $\|t\|_1 \in \{2, 3, 5\}$. We assume that company B may establish at most $r \in \{2, 3, 5\}$ new stores at a fixed cost $q_i \in \{750, 1000\}$ per store independent of the location i. For every configuration of the problem dimensions we generate 25 independent random instances as follows:

- The coordinates of the d different locations are chosen independently and uniformly at random from $[0, 1]^2$, and the unit transportation cost c_{ij} is set to the Euclidean distance between locations i and j for every $i, j \in \{1, \ldots, d\}$.
- The set of all locations is partitioned uniformly at random into demand centers and candidate locations for retail stores.
- The locations accommodating a store of company A are chosen uniformly at random from all candidate locations.
- The demand ξ_i is uniformly distributed on $[50, 150]$ at demand centers ($s_i = 0$) and vanishes with certainty at candidate locations for retail stores ($s_i = 1$).
- The revenue of selling one unit of the good is $v_i = 5$ at every location.
- Each retail store has capacity $b_i = 150 \cdot (d - \|s\|_1)/\|t\|_1$ irrespective of its location.

All experiments are run on a 64-bit Windows machine equipped with an Intel Core i7-4700HQ processor and 16 GB of RAM. All MILPs are implemented in MATLAB via the YALMIP [25] interface and solved using CPLEX 12.6.3. All big-M constants are set to 10^6. This value is large enough to guarantee that the big-M constraints are equivalent to exact indicator constraints. On the other hand, the value is small enough to avoid excessive branching. Table 7.2 reports summary statistics of the optimality gaps and CPU times for randomly generated problem instances with different dimensions. The optimality gaps are computed as

$$\Delta = 2 \cdot \left| \frac{\min \mathcal{P}^u - \min \mathcal{O}^l}{\min \mathcal{P}^u + \min \mathcal{O}^l} \right|.$$

We remark that company B may decide not to enter the market at all, that is, $x = 0$ is always feasible in (7.2) and incurs zero cost. Thus, both $\mathcal{P}^u$ and $\mathcal{O}^l$ have non-positive

Problem Dimensions					LDR				SLDR			
d	$\|s\|_1$	$\|t\|_1$	r	q_i	25%Gap	50%Gap	75%Gap	CPU	25%Gap	50%Gap	75%Gap	CPU
15	5	2	3	750	29.2	44.2	58.4	1	0.0	2.5	19.4	4
15	5	3	2	750	0.0	0.0	26.6	2	0.0	0.0	3.7	2
15	5	2	3	1000	24.8	51.4	74.0	1	0.0	9.3	30.8	9
15	5	3	2	1000	0.0	0.0	25.0	1	0.0	0.0	16.2	7
20	5	2	3	750	0.0	0.0	19.7	14	0.0	0.0	16.5	28
20	5	3	2	750	0.0	0.0	42.9	15	0.0	0.0	7.0	17
25	10	5	5	750	24.2	56.9	84.9	36	0.0	11.4	49.8	182
20	5	2	3	1000	0.0	29.8	51.3	51	0.0	5.7	20.8	44
20	5	3	2	1000	0.0	35.7	101.4	18	0.0	0.0	5.7	19
25	10	5	5	1000	3.6	70.1	90.9	112	0.0	10.5	43.3	242
25	5	2	3	750	1.9	22.1	31.9	32	0.0	4.6	19.3	161
25	5	3	2	750	0.0	1.8	19.5	20	0.0	0.6	18.3	128
30	10	5	5	750	21.7	54.1	83.5	180	0.0	12.8	41.3	1259
25	5	2	3	1000	6.3	31.3	41.0	35	0.0	2.6	18.5	133
25	5	3	2	1000	0.0	10.5	59.4	21	0.0	3.9	17.3	112
30	10	5	5	1000	26.8	69.9	90.9	170	0.0	15.1	42.0	1300
30	5	2	3	750	0.5	4.3	28.7	125	0.0	0.0	5.2	526
30	5	3	2	750	0.0	4.2	20.8	187	0.0	0.0	6.4	379
35	10	5	5	750	36.0	58.9	82.9	480	0.0	10.1	20.1	2700
30	5	2	3	1000	3.3	13	24.5	140	0.0	12.1	9.9	540
30	5	3	2	1000	0.0	11.6	37.4	215	0.0	0.0	15.4	335
35	10	5	5	1000	5.3	37.4	54.3	670	0.0	10.5	18.8	2890

TABLE 7.2

Numerical results of the randomized facility location problems. The table reports the 20%, 50% and 75% quantiles of the relative optimality gaps (25%Gap, 50%Gap, 75%Gap) as well as the average CPU times (CPU, measured in seconds) based on 25 independently generated problem instances.

optimal values. Each instance of the facility location problem is solved both in linear decision rules (LDR) and segregated linear decision rules (SLDR); see Example 2.

As the segregated linear decision rules provide more flexibility, they result in significantly smaller optimality gaps than the simple linear decision rules at the expense of higher CPU times. Specifically, the median optimality gaps attained by the segregated linear decision rules range from 0% to about 15%, whereas the median optimality gaps attained by the linear decision rules range from 0% up to about 70%.

Note that the optimality gap can be expressed as $\Delta = \Delta^u + \Delta^l$, where

$$\Delta^u = 2 \cdot \left| \frac{\min \mathcal{P}^u - \min \mathcal{P}}{\min \mathcal{P}^u + \min \mathcal{O}^l} \right| \quad \text{and} \quad \Delta^l = 2 \cdot \left| \frac{\min \mathcal{P} - \min \mathcal{O}^l}{\min \mathcal{P}^u + \min \mathcal{O}^l} \right|.$$

The actual quantity of interest is Δ^u, which measures the suboptimality of the minimizer of $\mathcal{P}^u$ in the original pessimistic bilevel program $\mathcal{P}$. However, unlike its conservative upper bound Δ, the actual optimality gap Δ^u cannot be computed. The optimality gaps reported in Table 7.2 thus represent pessimistic estimates of Δ^u.

Scenario 1						Scenario 2					
i	x_i	y_i	i	x_i	y_i	i	x_i	y_i	i	x_i	y_i
1	**0.616**	**0.642**	14	0.776	0.762	**1**	**0.967**	**0.734**	14	0.501	0.081
2	**0.455**	**0.047**	15	0.204	0.097	**2**	**0.620**	**0.096**	15	0.231	0.744
3	**0.834**	**0.671**	16	0.929	0.966	**3**	**0.055**	**0.813**	16	0.462	0.518
4	**0.803**	**0.038**	17	0.725	0.438	**4**	**0.264**	**0.573**	17	0.009	0.882
5	**0.464**	**0.273**	18	0.422	0.706	**5**	**0.775**	**0.941**	18	0.948	0.454
6	**0.102**	**0.293**	19	0.417	0.826	**6**	**0.826**	**0.624**	19	0.610	0.042
7	**0.683**	**0.270**	20	0.501	0.629	**7**	**0.685**	**0.498**	20	0.359	0.140
8	**0.295**	**0.486**	21	0.397	0.148	**8**	**0.391**	**0.143**	21	0.165	0.802
9	**0.182**	**0.164**	22	0.184	0.108	**9**	**0.015**	**0.073**	22	0.068	0.387
10	**0.336**	**0.690**	23	0.188	0.993	**10**	**0.998**	**0.537**	23	0.462	0.712
11	0.914	0.538	24	0.547	0.993	11	0.280	0.039	24	0.147	0.385
12	0.566	0.445	25	0.901	0.322	12	0.913	0.789	25	0.062	0.940
13	0.509	0.570				13	0.818	0.374			

Table 7.3

Coordinates of the 25 locations in scenarios 1 and 2. The first 10 locations in each scenario represent candidate locations (bold), whereof the first 5 locations are occupied by company A (shaded).

If $\mathbb{P}$ is approximated by a uniform empirical distribution $\widehat{\mathbb{P}}_N$ on N independent samples from $\mathbb{P}$, then the resulting sample average approximations of the bilevel programs $\mathcal{P}$ and $\mathcal{O}$ admit exact MILP reformulations whose sizes scale with N; see Theorems 6.1 and 6.2. While the sample average approximations of classical stochastic programs are known to be optimistically biased, meaning that their optimal values are guaranteed to provide lower confidence bounds on the optimal value of the true stochastic program, the sample average approximations of stochastic bilevel programs can be either optimistically or pessimistically biased. To see this, we consider two scenarios of the market entrant's facility location problem with different cost structures. The coordinates of the demand centers and candidate locations as well as the locations occupied by company A are reported in Table 7.3. In both scenarios, the market entrant may open up at most 5 stores at a cost of 750 each. All existing and tentative new stores have a capacity of 450. The demands at all demand locations follow independent uniform distributions supported on $[50, 150]$. As before, the customers incur costs that reflect the Euclidean shipping distances. The variable costs of the market entrant, however, are fundamentally different in the two scenarios. Specifically, the market entrant earns a revenue of 5 per unit sold in the first scenario and has to cover the customers' shipping costs in the second scenario. Thus, the leader and the follower have conflicting objectives in the first scenario ($C \neq V$) and compatible objectives in the second scenario ($C = V$).

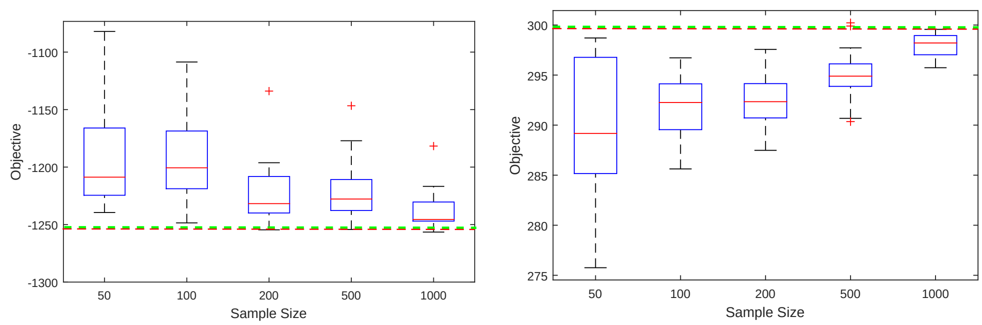

FIG. 7.3. *Expected costs of the market entrant when $C \neq V$ (left) and $V = C$ (right).*

Figure 7.3 visualizes the optimal values of the pessimistic and optimistic stochastic bilevel programs solved in segregated linear decision rules (red and green dashed lines, respectively) as well as the optimal values of 20 independent instances of the sample average approximation $\widehat{\mathcal{P}}$ for increasing sample sizes $N \in \{10, 50, 100, 200, 500, 1000\}$ (boxplots). Note that the decision rule bounds coincide, implying that segregated linear decision rules are optimal in both $\mathcal{O}$ and $\mathcal{P}$. Moreover, the sample average approximations $\widehat{\mathcal{O}}$ and $\widehat{\mathcal{P}}$ share the same optimal values across all instances, that is, there is no difference between optimistic and pessimistic formulations. The two charts of Figure 7.3 correspond to the two outlined scenarios with $C \neq V$ and $C = V$, respectively.

We observe that the sample average approximations display significant variability and are pessimistically biased in scenario 1 and optimistically biased in scenario 2. Thus, the sample average approximations are unreliable, and not even the sign of their bias is known upfront. We emphasize that it is possible to conceive scenarios in which the optimal values of the sample average approximations accumulate between the decision rule bounds already for smaller sample sizes. However, without computing the deterministic decision rule bounds, it would be impossible to recognize the good performance of the sample average approximations in these cases. Finally, while all decision rule bounds in this experiment could be computed in less than 300 seconds, the sample average approximations needed up to 2,000 seconds of runtime for the larger instances displaying a similar accuracy as the decision rule bounds.

8. Conclusion. This paper develops the first decision rule bounds for both pessimistic and optimistic versions of a two-stage stochastic bilevel program, where the leader first chooses a binary here-and-now decision, in response to which the follower selects a continuous wait-and-see decision. Both the upper bound on the pessimistic as well as the lower bound on the optimistic version of the bilevel program coincide with the optimal values of explicit MILPs that can be solved efficiently with commercial solvers. Specifically, the numbers of decision variables and constraints of these MILPs scale only polynomially with the input dimensions. Moreover, both bounding problems reside in the same complexity class as the nominal pessimistic and optimistic bilevel programs. In other words, the MILP reformulations of the bounding problems involve only polynomially more variables and constraints than the MILP reformulations of the corresponding nominal problems. Solving the stochastic bilevel programs approximately via the sample average approximation gives rise to alternative explicit MILPs, which scale, however, with the sample size and thus become quickly computationally excruciating. In contrast to classical stochastic programming, where sample average approximations are optimistically biased, the sample average approximations of stochastic bilevel programs can be both optimistically or pessimistically biased. This uncertainty about the bias underlines the attractiveness of the deterministic bounding methods developed in this paper. Future research should aim to extend the results of this paper to *robust* two-stage bilevel programs. It would be particularly interesting to investigate the ramifications of numerous powerful optimality results for decision rules in the context of robust bilevel programming.

Acknowledgements. We thank two anonymous referees for their constructive comments and suggestions that led to substantial improvements of the paper. This work was supported by the Scientific and Technological Research Council of Turkey (TÜBİTAK) through the BİDEB 2232 Research Grant (115C100) and by the Swiss National Science Foundation under grant BSCGI0_157733. The authors thank Serkan Buldan for his help with the artwork.

Appendix A. Proofs.

Proof of Proposition 3.1. Here, we just outline the proof idea to keep this paper self-contained. A detailed proof can be found in [23, § 2.2]. When restricting the wait-and-see decision to linear decision rules $y(\xi) = Y\xi$ and using the conventions $b_x(\xi) = B_x\xi$ and $c(\xi) = C\xi$, problem (2.1) reduces to

$$\begin{array}{ll} \min\limits_{Y} & \mathbb{E}[\xi^\top C^\top Y\xi] \\ \text{s.t.} & AY\xi \le B_x\xi \quad \mathbb{P}\text{-a.s.} \end{array} \tag{A.1}$$

Rewriting the inner product in the objective function as a trace, exploiting the cyclicity of the trace and interchanging the trace and expectation operators yields $\mathbb{E}[\xi^\top C^\top Y\xi] = \mathrm{Tr}(\Omega C^\top Y)$. Moreover, as linear functions are continuous, the almost sure constraints in (A.1) are equivalent to the semi-infinite (robust) constraints $AY\xi \le B_x\xi$ for all $\xi \in \Xi$. The linear decision rule problem (A.1) thus reduces to the robust optimization problem

$$\begin{array}{ll} \min\limits_{Y} & \mathrm{Tr}(\Omega C^\top Y) \\ \text{s.t.} & AY\xi \le B_x\xi \quad \forall \xi \in \Xi. \end{array} \tag{A.2}$$

By using the standard machinery of robust linear optimization with polyhedral uncertainty sets [3], the semi-infinite program (A.2) can be reformulated as the tractable linear program (3.1). Thus, the claim follows. ■

Proof of Proposition 3.2. We just outline the proof idea and refer to [23, §§ 2.3–2.4] for a full proof. By introducing a slack variable $s(\xi)$, problem (2.1) can be reformulated as

$$\begin{array}{ll} \min\limits_{y\in\mathcal{L}^2_n,\, s\in\mathcal{L}^2_m} & \mathbb{E}[c(\xi)^\top y(\xi)] \\ \text{s.t.} & \left.\begin{array}{r} Ay(\xi) + s(\xi) - b_x(\xi) = 0 \\ s(\xi) \ge 0 \end{array}\right\} \quad \mathbb{P}\text{-a.s.} \end{array} \tag{A.3}$$

By dualizing the equality constraint, (A.3) reduces to the min-max problem

$$\begin{array}{lll} \min\limits_{y\in\mathcal{L}^2_n,\, s\in\mathcal{L}^2_m} & \sup\limits_{\lambda\in\mathcal{L}^2_m} & \mathbb{E}[c(\xi)^\top y(\xi) + \lambda(\xi)^\top [Ay(\xi) + s(\xi) - b_x(\xi)]] \\ \text{s.t.} & s(\xi) \geq 0 & \mathbb{P}\text{-a.s.} \end{array} \tag{A.4}$$

Instead of maximizing over all square-integrable dual decisions $\lambda(\xi)$, we can alternatively maximize over the subclass of *dual linear decision rules* $\lambda(\xi) = \Lambda\xi$. This leads to a relaxed min-max problem where the inner maximization runs over all matrices $\Lambda \in \mathbb{R}^{m\times k}$. By carrying out this inner maximization explicitly, one can show that the relaxed min-max problem is equivalent to

$$\begin{array}{ll} \min\limits_{y\in\mathcal{L}^2_n,\, s\in\mathcal{L}^2_m} & \mathbb{E}[c(\xi)^\top y(\xi)] \\ \text{s.t.} & \left.\begin{array}{r} \mathbb{E}[[Ay(\xi) + s(\xi) - b_x(\xi)]\xi^\top] = 0 \\ s(\xi) \geq 0 \end{array}\right\} \quad \mathbb{P}\text{-a.s.} \end{array} \tag{A.5}$$

Clearly, any $y(\xi)$ and $s(\xi)$ feasible in (A.3) are also feasible in (A.5), but the converse statement is not necessarily true. Thus, the optimal value of the dual linear decision rule approximation (A.5) provides a *lower* bound on $z^*(x)$. Moreover, (A.5) can be viewed as the *dual* of a linear robust optimization problem with a polyhedral uncertainty set. Under a mild strict feasibility condition and because the second-order moment matrix Ω is strictly positive definite by Assumption 2, one can thus show that (A.5) is equivalent to the tractable linear program (3.2) by invoking a dual variant of the standard robust optimization machinery. However, even if the strict feasibility condition fails to hold, problem (3.2) still provides a lower bound on (A.5); see [23, § 2.4] for details. Thus, the claim follows. ■

REFERENCES

[1] S.M. Alizadeh, P. Marcotte, and G. Savard. Two-stage stochastic bilevel programming over a transportation network. *Transportation Research Part B: Methodological*, 58:92–105, 2013.

[2] J.-A. Audestad, A.A. Gaivoronski, and A. Werner. Extending the stochastic programming framework for the modeling of several decision makers: pricing and competition in the telecommunication sector. *Annals of Operations Research*, 142(1):19–39, 2006.

[3] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009.

[4] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

[5] J. Bracken and J.T. McGill. Mathematical programs with optimization problems in the constraints. *Operations Research*, 21(1):37–44, 1973.

[6] X. Chen, M. Sim, P. Sun, and J. Zhang. A linear decision-based approximation approach to stochastic programming. *Operations Research*, 56(2):344–357, 2008.

[7] S. Christiansen, M. Patriksson, and L. Wynter. Stochastic bilevel programming in structural optimization. *Structural and Multidisciplinary Optimization*, 21(5):361–371, 2001.

[8] J.-P. Côté, P. Marcotte, and G. Savard. A bilevel modelling approach to pricing and fare optimisation in the airline industry. *Journal of Revenue and Pricing Management*, 2(1):23–36, 2003.

[9] S. Dempe. A necessary and a sufficient optimality condition for bilevel programming problems. *Optimization*, 25(4):341–354, 1992.

[10] S. Dempe. *Foundations of Bilevel Programming*. Springer, 2002.

[11] S. Dempe, V. Kalashnikov, and N. Kalashnykova. Optimality conditions for bilevel programming problems. In S. Dempe and V. Kalashnikov, editors, *Optimization with Multivalued Mappings: Theory, Applications, and Algorithms*, pages 3–28. Springer, 2006.

[12] S. Dempe and A.B. Zemkoho. The bilevel programming problem: Reformulations, constraint qualifications and optimality conditions. *Mathematical Programming*, 138(1):447–473, 2013.

[13] M. Dyer and L. Stougie. Computational complexity of stochastic programming problems. *Mathematical Programming*, 106(3):423–432, 2006.

[14] A. Evgrafov and M. Patriksson. On the existence of solutions to stochastic mathematical programs with equilibrium constraints. *Journal of Optimization Theory and Applications*, 121(1):65–76, 2004.

[15] N.P. Faísca, V. Dua, B. Rustem, P.M. Saraiva, and E.N. Pistikopoulos. Parametric global optimisation for bilevel programming. *Journal of Global Optimization*, 38(4):609–623, 2007.

[16] M. Fampa, L.A. Barroso, D. Candal, and L. Simonetti. Bilevel optimization applied to strategic pricing in competitive electricity markets. *Computational Optimization and Applications*, 39(2):121–142, 2008.

[17] A. Georghiou, W. Wiesemann, and D. Kuhn. Generalized decision rule approximations for stochastic programming via liftings. *Mathematical Programming*, 152(1-2):301–338, 2015.

[18] J. Goh and M. Sim. Distributionally robust optimization and its tractable approximations. *Operations Research*, 58(4):902–917, 2010.

[19] G.A. Hanasusanto, D. Kuhn, and W. Wiesemann. A comment on "computational complexity of stochastic programming problems". *Mathematical Programming*, 159(1):557–569, 2016.

[20] C. Henkel. *An Algorithm for the Global Resolution of Linear Stochastic Bilevel Programs*. PhD thesis, Universität Duisburg-Essen, 2014.

[21] X. Hu and D. Ralph. Using EPECs to model bilevel games in restructured electricity markets with locational prices. *Operations Research*, 55(5):809–827, 2007.

[22] R.R. Iyer and I.E. Grossmann. A bilevel decomposition algorithm for long-range planning of process networks. *Industrial & Engineering Chemistry Research*, 37(2):474–481, 1998.

[23] D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2011.

[24] Y. Liu and H. Xu. Entropic approximation for mathematical programs with robust equilibrium constraints. *SIAM Journal on Optimization*, 24(3):933–958, 2014.

[25] J. Löfberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *IEEE International Symposium on Computer Aided Control Systems Design*, pages 284–289, 2004.

[26] A. Migdalas. Bilevel programming in traffic planning: Models, methods and challenge. *Journal of Global Optimization*, 7(4):381–405, 1995.

[27] A. Mitsos, P. Lemonidis, and P.I. Barton. Global solution of bilevel programs with a nonconvex inner program. *Journal of Global Optimization*, 42(4):475–513, 2008.

[28] M. Patriksson. Robust bi-level optimization models in transportation science. *Philosophical Transactions of the Royal Society of London A*, 366(1872):1989–2004, 2008.

[29] M. Patriksson and L. Wynter. Stochastic nonlinear bilevel programming. Technical report, Citeseer, 1997.

[30] M. Patriksson and L. Wynter. Stochastic mathematical programs with equilibrium constraints. *Operations Research Letters*, 25(4):159–167, 1999.

[31] J.-H. Ryu, V. Dua, and E.N. Pistikopoulos. A bilevel programming framework for enterprise-wide process networks under uncertainty. *Computers & Chemical Engineering*, 28(6):1121–1129, 2004.

[32] M.P. Scaparra and R.L. Church. A bilevel mixed-integer program for critical infrastructure protection planning. *Computers & Operations Research*, 35(6):1905–1923, 2008.

[33] A. Shapiro. Stochastic programming with equilibrium constraints. *Journal of Optimization Theory and Applications*, 128(1):221–243, 2006.

[34] A. Shapiro and A. Nemirovski. On complexity of stochastic programming problems. In V. Jeyakumar and A.M. Rubinov, editors, *Continuous Optimization: Current Trends and Applications*, pages 111–144. Springer, 2005.

[35] A. Shapiro and H. Xu. Stochastic mathematical programs with equilibrium constraints, modelling and sample average approximation. *Optimization*, 57(3):395–418, 2008.

[36] J. Tirole and D. Fudenberg. *Game Theory*. MIT Press, 1991.

[37] A. Tsoukalas, B. Rustem, and E.N. Pistikopoulos. A global optimization algorithm for generalized semi-infinite, continuous minimax with coupled constraints and bi-level problems. *Journal of Global Optimization*, 44(2):235–250, 2009.

[38] A. Tsoukalas, W. Wiesemann, and B. Rustem. Global optimisation of pessimistic bi-level problems. In *Lectures on Global Optimization, Fields Communications Series*, pages 215–243. American Mathematical Society, 2009.

[39] H. Tuy, A. Migdalas, and N.T. Hoai-Phuong. A novel approach to bilevel nonlinear programming. *Journal of Global Optimization*, 38(4):527–554, 2007.

[40] W. Wiesemann, A. Tsoukalas, P.-M. Kleniati, and B. Rustem. Pessimistic bilevel optimization. *SIAM Journal on Optimization*, 23(1):353–380, 2013.

[41] H. Xu. An MPCC approach for stochastic Stackelberg-Nash-Cournot equilibrium. *Optimization*, 54(1):27–57, 2005.

[42] H. Xu. An implicit programming approach for a class of stochastic mathematical programs with complementarity constraints. *SIAM Journal on Optimization*, 16(3):670–696, 2006.