Robust bilevel optimization: algorithms, complexity and application

Tamás Kis, András Kovács and Csaba Mészáros

HUN-REN Institute for Computer Science and Control (SZTAKI),
1111 Budapest, Kende u. 13-17, Hungary

ARTICLE HISTORY
Compiled February 12, 2025

ABSTRACT
The paper addresses robust bilevel optimization with polyhedral uncertainty in the followers' objective function coefficients. It is assumed that both the leader's and the followers' models are linear, yet, bilinear terms in the followers' objective function are allowed. An efficient algorithm is introduced that finds a bounded-error solution in a finite number of steps. This model captures typical price setting applications, such as network toll setting or electricity tariff optimization for demand response management. The efficiency of the general method is illustrated on demand response management in smart grids. Our computational experiments show that the method solves instances with hundreds of decision variables. The significance of these results is underlined by a proof that the above demand response management problem, and hence, the generic robust bilevel problem as well, are Σ_2^p-complete. Finally, the infinitely robust variant of the problem is discussed, and it is shown to be tractable in polynomial time.

KEYWORDS
Bilevel programming; robust optimization; computational complexity; electricity pricing; demand response management

1. Introduction

Bilevel optimization tackles the challenge of finding the equilibrium in decision problems involving multiple stakeholders. The player called the *leader* makes its choice first, which is observed by its *followers*, who determine their response by optimizing their own objectives. Hence, when looking for its optimal decision, the leader must account for the rational response of the self-interested followers. Accordingly, a critical assumption of classical deterministic bilevel approaches is that the leader has complete knowledge of the decision problems of its followers. Since this assumption is unrealistic in many applications, *robust bilevel optimization* is receiving increasing attention as a promising approach to make use of the available, yet imperfect information about the followers. Theoretical foundations are being elaborated [1] and the first applications are emerging in problems such as hazardous material (hazmat) transportation network design [2], homeland security [3], electric vehicle (EV) charging station design [4], and the scheduling of EV charging [5].

This paper addresses robust bilevel optimization problems where uncertainty occurs in the coefficients of the followers' bilinear objective function in the form of polyhedral

CONTACT András Kovács. Email: andras.kovacs@sztaki.hu

uncertainty sets, and both the leader's and the followers' constraints are linear. It is assumed that uncertainty realizes after the leader has committed to its decision, but before the followers determine their response, i.e., followers solve deterministic problems. This assumption corresponds to so-called *wait-and-see followers* [1]. With this, the *Robust Bilevel Optimization Problem* (RBOP) addressed in this paper can be stated as follows:

$$
\begin{aligned}
&\sup z && (1)\\
&\text{s.t.}\\
&\min\{f(x,y) \mid u \in U,\ y \in \Omega(x,u)\} \geq z && (2)\\
&x \in X && (3)
\end{aligned}
$$

where X is a polytope consisting of all the feasible solutions x of the leader, U is a polyhedral set of uncertain parameters of the follower(s), and $\Omega(x,u)$ is the set of optimal solutions of the follower(s) for fixed x and u, i.e., $y \in \Omega(x,u)$ if and only if y is an optimal solution of the parameterized linear program (LP)

$$\max\{(u-x)^T y \mid y \in Y\} \tag{4}$$

where Y is a polytope. Note that (4) can be a minimization problem as well. Observe that the objective function is linear for fixed u and x. In (2), the minimum is taken over all $u \in U$ and $y \in \Omega(x,u)$, i.e., robustness also incorporates the pessimistic assumption on the choice of the followers' response when $\Omega(x,u)$ is not singleton.

In the objective function (1), we have to use supremum, since our problem contains pessimistic bilevel optimization as a special case when U is a singleton, where the optimum value cannot be attained in general, see e.g., [6]. From now on, $z^\star$ denotes the supremum value.

Throughout this paper, we assume that

$$f(x,y) = (c+x)^T y,$$

where c is a constant vector. Moreover, by assuming that Y is a non-empty polytope, $\Omega(x,u)$ is never empty. The notation applied in the paper is summarized in Table 1. For basic definitions on linear and integer programming, as well as on polyhedral theory, we refer the reader to [7]. The feasible solutions of RBOP (1)-(4) are characterized next.

Definition 1.1. Some (x,y,u) is *robust bilevel feasible* if and only if $x \in X$, $u \in U$, $y \in \Omega(x,u)$ and for any $u' \in U$ and $y' \in \Omega(x,u')$, we have $f(x,y) \leq f(x,y')$.

Definition 1.2. Given an error bound $\varepsilon > 0$, a robust bilevel feasible solution (x,y,u) is *ε-optimal* if $z^\star - f(x,y) \leq \varepsilon(|z^\star| + 1)$.

Main results. For polyhedral U, we propose a novel algorithm that computes solutions for RBOP that approach its supremum arbitrarily well. The approach is illustrated on a demand response management problem in smart electricity grids. The computational results show that the proposed method approaches the supremum with a small error for test instances with up to 15 followers and 15 time periods. Our modeling, algorithmic and computational results are complemented by two theoretical results about the complexity of the problem. On the one hand, we prove that the general problem is

Table 1. Notation.

Problem size and indices	
L	Number of vertices of polyhedron Y
ℓ	Index of vertices of polyhedron Y
K	Number of discrete uncertainty values
k	Index of discrete uncertainty values
M	Application: number of followers
i	Application: follower index
T	Application: number of time periods
t	Application: time period index
$[N]$	set of integers in the interval $[1, N]$ for any integer $N \geq 1$
Input parameters	
X	Leader's feasible region
Y	Followers' feasible region
$\hat{Y}$	Vertices of polyhedron Y
U	Uncertainty set
c	Constant in leader's objective
p_t	Application: wholesale electricity price
$\underline{d}_i$, $\bar{d}_i$	Application: min. and max. total consumption of consumer i
$\underline{y}_{it}$, $\bar{y}_{it}$	Application: min. and max. consumption of consumer i in period t
r^j	Constraint coefficients defining X
v^j	Constraint coefficients defining U
Decision variables and objectives	
x	Leader's variables (application: electricity tariff)
y	Followers' variables (application: electricity consumption)
u	Followers' uncertain parameters (application: consumers' utility)
z	Leader's objective value (application: retailer's profit value)
$f(x, y)$	Leader's objective function (application: retailer's profit function)
$z^\star$	Leader's supremum
Parameters used by the algorithm	
U_δ	Uncertainty set extended by δ
z^δ	Leader's supremum over U_δ
U^D	Discrete uncertainty set built by the algorithm
z^D	Leader's maximum over U^D
$iter$	Iteration count
$N(y)$	Neighboring vertices of y in Y
$\Omega(x, u)$	Follower's set of optimal responses to (x, u)
$Proj_x(\cdot)$	Projection of a set to X
$P(y)$	Set of (x, u) such that $y \in \Omega(x, u)$
α, β	Dual variables
Δ	Characteristic radius (distance between followers' objective values)
θ	Characteristic radius (distance between leader's variables)

Σ_2^p-hard. This complexity result implies that our problem is located outside NP unless the polynomial hierarchy collapses. For basic definitions on the polynomial hierarchy, we refer to [8]. On the other hand, we introduce a special case called the infinitely robust variant, and show that this variant can be solved in polynomial time.

Organization of the paper. The literature on robust bilevel optimization is reviewed in Section 2. Key notions are introduced and core properties of RBOP are established in Section 3. A generic solution method is proposed and its main features are discussed in Section 4. Then, Section 5 illustrates the generic method via an application to demand response management in smart grids. The efficiency of the algorithm is investigated in computational experiments in Section 6. Further theoretical results on the computational complexity of RBOP are presented in Section 7. Finally, conclusions are drawn in Section 8.

2. Literature review

2.1. Robust bilevel optimization

A critical assumption of classical, deterministic bilevel optimization approaches is that the leader is perfectly aware of the followers' decision model and parameters, and hence, it can accurately predict the followers' response to any possible decision [6]. Obviously, this assumption cannot be satisfied in most practical applications. As a possible means to lift this critical assumption, approaches to bilevel optimization under uncertainty, including both robust and stochastic techniques received significant attention in recent years. An excellent review and classification of such robust (and also stochastic) bilevel approaches is presented in [1]. The review highlights that sources of uncertainty in a bilevel context can be significantly richer than in single-level optimization, and categorizes these sources as *data uncertainty* and *decision uncertainty.*

2.2. Robustness against data uncertainty

Approaches to model data uncertainty can be classified further according to the parameters impacted by uncertainty; the assumptions on the uncertainty set, e.g., discrete, box, ellipsoidal or polyhedral sets; as well as the timing of the decisions. For the latter, *wait-and-see followers*, who observe the realization of the uncertain parameters before making their decisions, are differentiated from *here-and-now followers* who must decide before realization, meaning that the follower's sub-problem is in itself a robust optimization problem.

Paper [9] investigates bilevel problems with polynomial leader's objective and interval uncertainty in both the upper- and lower-level linear constraints. The problem is solved via a series of single-level non-convex polynomial relaxations. The approach is illustrated on a few examples with a single leader and a single follower variable. The approach is extended to a more generic problem class with polynomial leader and follower constraint functions in [5], and it is illustrated on an application to scheduling the charging of EVs.

The complexity of the robust bilevel continuous knapsack problem with different forms of uncertainty in the follower's objective is studied in [10]. The problem can be solved in polynomial time in the deterministic case, and the authors show that this result generalizes to the case of discrete and interval uncertainty sets. However, the same problem becomes NP-hard in case of discrete uncorrelated uncertainty (i.e., where the uncertainty set is given as a Cartesian product of finite sets), as well as for polyhedral and ellipsoidal uncertainty. Generic robust linear bilevel problems with uncertainty in the follower's objective are investigated in [11]. It is shown that the robust problem under interval uncertainty can be Σ_2^p-complete even if the deterministic bilevel version is contained in NP and the follower's problem is polynomially solvable. At the same time, the robust problem with discrete uncertainty is located at most one level higher in the polynomial hierarchy than the follower's sub-problem.

A Γ-robust approach to discrete min-max problems with uncertainty in the follower's parameters, including its objective and constraints, is presented and a branch-and-cut algorithm is proposed in [12]. The approach is illustrated on a knapsack interdiction problem.

2.3. Robustness against decision uncertainty

In addition to data uncertainty, bilevel optimization problems may also be exposed to *decision uncertainty* when the leader, even if aware of the exact parameter values, is unable to predict precisely the behavior of its followers. A basic form of decision uncertainty is captured by the *pessimistic bilevel problem*: followers may have multiple optimal responses to the deterministic problem they face, but these bring different benefits for the leader [6]. In such a situation, the leader may want to prepare for receiving the least favorable optimal follower response. Optimality conditions for pessimistic bilevel optimization are derived in [13,14]. An efficient solution approach is proposed for a generic class of pessimistic bilevel problems based on a tight bilevel relaxation and a correction operation in [15]. In [16], it is proven that *independent* pessimistic bilevel problems (i.e., where the followers' feasible region is independent of the leader's decision) take an optimal solution if the followers' objective function is additively separable. On the contrary, if the objective is not separable or the problem is dependent, then problems in general do not admit an optimal solution. Observe that the problem addressed in this paper is independent, but separability for the followers' objective is not ensured. The same paper [16] proves that for independent pessimistic bilevel problems, a sequence of approximations where the followers may return an ε-optimal response instead of the exact optimum converges to the supremum of the original pessimistic problem. An iterative solution scheme motivated by semi-infinite programming techniques is proposed for solving the approximate problem. Yet, the problems investigated differ substantially from those in the current paper: arbitrary non-convex constraints and objectives are allowed, but the typical problem size is only one variable for the leader and one for the follower. The problem of finding the proper balance between the optimistic and pessimistic bilevel cases is called the strong-weak bilevel problem, and it is investigated, e.g., in [17]. An algorithm for constructing stable solutions of linear-linear bilevel problems, also leading to a trade-off between the optimistic and the pessimistic solutions, is introduced in [18].

A relevant source of decision uncertainty can be the inability of the follower to compute an exact optimal response, and accordingly, it may settle for a satisfactory, ε-optimal solution. This phenomenon is known as lower-level near-optimality. Near-optimal robust bilevel optimization problems were investigated in [19], where necessary conditions for the existence of bilevel-feasible solutions are established, and a solution approach based on reformulation to a single-level problem is proposed for the case of convex lower level. In [20], it is shown that, given appropriate conditions, the near-optimal robust multilevel problem remains in the same complexity class as the deterministic variant. The paper [21] calls attention to the practical computational challenge that for problems with non-convex lower levels, where only ε-feasibility can be expected for the non-linear constraints, the resulting ε-feasible solution may be arbitrarily far from the exact optimal solution.

Another possible approach to tackling the limited computational capabilities of the follower is assuming that it chooses its solution strategies from a finite set of methods, known to the leader, which includes heuristics and approximation techniques [22]. After providing generic definitions, the approach is elaborated for the bilevel knapsack problem and heuristics that work with a fixed preference order of the items.

Furthermore, decision uncertainty may stem from the followers' inability to observe the leader's decision precisely, which in turn results in an uncertainty for the leader regarding the follower's response. Bilinear bilevel problems with limited observability are considered in [23], where the leader must prepare for the follower response to any perceived leader action in a polyhedral neighborhood of the true action. A solution method based on reformulation to a single-level problem is proposed, and it is shown that the robust bilevel problem with limited observability belongs to the same complexity class as the original, deterministic bilevel problem.

2.4. Applications of robust bilevel optimization

The aforementioned review [1] emphasizes that, despite the numerous applications of bilevel optimization under uncertainty in the literature, the vast majority of these consider a stochastic setting with known, discrete probability distributions, which allows generating the deterministic equivalent directly. Below, we focus solely on the scarce applications of the robust bilevel approach.

Applications to energy management include [24], where a Γ-robust bilevel approach is taken to compute the optimal bidding strategy of a generator, with the lower level standing for a transmission-constrained economic dispatch problem. Uncertainties arise in the lower level due to rival offers and market demand. The robust bilevel problem is solved by transforming it into a single-level mixed-integer linear program (MILP). EV charging station design is formulated as a robust bilevel problem in [4]. The lower level captures the charging decisions of the EV owners, which is affected by multiple sources of uncertainty, including electricity prices and availability, as well as traffic. Again, the bilevel problem is reformulated into a single-level one, solved in turn by a column-and-constraint generation method.

Network interdiction problems are particularly relevant for security applications. In such applications, information asymmetry is inherent, and therefore, robust bilevel approaches are of interest. A Γ-robust bilevel optimization approach is presented to the problem of allocating defense budget to potential targets in [3]. The source of uncertainty is the follower's (attacker's) valuation of the targets, which is unknown to the leader. In [25], a maximum flow interdiction problem is investigated subject to uncertainties related to arc capacities and the resource consumption of arc interdiction, and the performance of several heuristic solution approaches are analyzed. A robust bilevel approach to shortest path network interdiction is introduced in [26], with uncertainties in the leader's arc costs. The problem is solved by reformulating it into a single-level MILP with a second-order cone constraint. In [2], a robust bilevel approach is proposed to hazmat transportation network design, i.e., the problem of selecting arcs from a given network where hazmat transportation should be interdicted to minimize environmental risks. Uncertainty stems from limited knowledge about accident probability. The bilevel model is transformed into a single-level MILP. Further approaches to hazmat network design with uncertainty on arc cost in the followers' problem include [27,28]. A detailed review of network interdiction models and applications is given in [29].

2.5. Positioning of the paper

According to the above classification scheme, the model investigated in this paper is a robust bilevel optimization problem with polyhedral data uncertainty in the objective function of the wait-and-see followers. The model also adopts the pessimistic bilevel assumption to hedge against a potentially unfavorable choice of the follower from the set of multiple optimal responses for a given realization of the uncertainty.

Strongly related contributions have been published recently, which shows the increasing interest towards similar robust bilevel approaches. These include a formal analysis of computational complexity in [11], yet, without arriving at a solution method; as well as algorithms for specific problems, such as the continuous bilevel knapsack problem in [10]. However, to the best of our knowledge, the present paper is the first to propose an efficient solution method for RBOP as defined above.

3. Preliminaries

In this section, we establish some fundamental properties of RBOP.

Observation 3.1. *Since Y is a polytope, the followers' sub-problem (4) admits a finite maximum for any fixed $(x, u) \in X \times U$. Likewise, since both X and Y are polytopes, the overall RBOP (1)-(3) always has a finite supremum.*

Proposition 3.2. *For any $(x, u) \in X \times U$, $\Omega(x, u)$ is a non-empty face of Y.*

Proof. Since Y is a non-empty polytope by assumption, for any fixed x and u the linear program (4) admits an optimal solution, and the optimal solutions constitute a face of Y. □

Corollary 3.3. *For any $(x, u) \in X \times U$, $\min\{f(x, y) \mid y \in \Omega(x, u)\}$ is attained by a vertex of Y.*

Let $\hat{Y} = \{y^\ell \ : \ \ell \in [L]\}$ denote the set of vertices of Y. For any $y \in Y$, let $P(y)$ consist of the pairs $(x, u) \in X \times U$ for which y is an optimal response of the followers to (x, u), i.e., $y \in \Omega(x, u)$.

Proposition 3.4. *$P(y)$ is a polyhedron for any $y \in Y$.*

Proof. Observe that any fixed $y \in Y$ is optimal for some $(x, u) \in X \times U$ if and only if (x, u) satisfies the constraints

$$\begin{aligned}
&(u - x)^T y \geq (u - x)^T y^\ell, \quad \ell \in [L] \\
&x \in X, \ u \in U.
\end{aligned}$$

Clearly, this is a linear system, and the statement follows. □

For non-empty $P(y)$ we define the projection of $P(y)$ to the x variables, that is, let $Proj_x(P(y)) = \{x \in X \mid \exists u \in U \text{ such that } (x, u) \in P(y)\}$. Since the projection of a polyhedron to a linear subspace is a polyhedron, we have

Proposition 3.5. *$Proj_x(P(y))$ is a polyhedron for any $y \in Y$.*

Define $X^{opt}(y) = Proj_x(P(y))$ for any $y \in Y$. Note that $X^{opt}(y)$ may be empty. Furthermore, distinct y and y' in Y may be optimal for the same $x \in X$, that is, $x \in X^{opt}(y) \cap X^{opt}(y')$.

Proposition 3.6. *Suppose X has a non-empty relative interior. Then, for any $\varepsilon > 0$, there exists an ε-optimal robust bilevel feasible solution (x', y', u') such that x' is in the relative interior of X.*

The proof of the proposition is presented in Appendix A.

In the following definition, we assume that U is defined by a system of linear inequalities of the form $\alpha u \leq \alpha_0$, and equations are modelled by two inequalities, i.e., $\alpha^= u \leq \alpha_0^=$ and $-\alpha^= u \leq -\alpha_0^=$.

Definition 3.7. For a given constant $\delta > 0$, let U_δ be the polytope obtained by replacing the right-hand-side by $\alpha_0 + \delta|\alpha_0| + \delta$ of each defining inequality $\alpha u \leq \alpha_0$ of U.

It follows that U_δ is full-dimensional, even if U is not.

Proposition 3.8. *For any given $\varepsilon > 0$, there exists $\delta > 0$ such that the bilevel problem over the extended uncertainty set U_δ admits a solution of value at least $z^\star - \varepsilon(|z^\star|+1)$.*

Proof. First, consider any robust bilevel feasible and ε-optimal solution (x, y, u) that incurs a leader's objective function value of at least $z^\star - \varepsilon(|z^\star| + 1)$ over the original uncertainty set U. By Corollary 3.3, the worst-case response $y \in Y$ for x is a vertex of Y. Observe that the set of vertices $\hat{Y}$ of Y can be partitioned into two disjoint subsets: $\hat{Y}_1(x)$ consisting of the possible responses to x, i.e., $\hat{Y}_1(x) = \{y^\ell \in \hat{Y} \mid \exists u \in U,\ y^\ell \in \Omega(x, u)\}$; and $\hat{Y}_2(x)$ those which are not possible follower responses for x, i.e., $\hat{Y}_2(x) = \{y^\ell \in \hat{Y} \mid \forall u \in U,\ y^\ell \notin \Omega(x, u)\}$. Clearly, for each $y^\ell \in \hat{Y}$, $y^\ell \in \hat{Y}_1(x)$ if and only if $x \in Proj_x(P(y^\ell))$, otherwise $y^\ell \in \hat{Y}_2(x)$, and $x \notin Proj_x(P(y^\ell))$.

Now, for each $y^\ell \in \hat{Y}_2(x)$, we pick a small $\delta_\ell > 0$ and determine the polytope U_{δ_ℓ}. Finally, we define $P_{\delta_\ell}(y^\ell) \subset X \times U_{\delta_\ell}$ analogously to $P(y)$ using the extended set U_{δ_ℓ}. Since $P(y^\ell) \subseteq P_{\delta_\ell}(y^\ell)$, and polytopes are closed sets, there exists $\delta_\ell > 0$ such that $x \notin Proj_x(P_{\delta_\ell}(y^\ell))$. Let $\delta_{\min} = \min_{y^\ell \in \hat{Y}_2(x)} \delta_\ell$. It follows that for all $y^\ell \in \hat{Y}_2(x)$, $x \notin Proj_x(P_{\delta_{\min}}(y^\ell))$.

Finally, for $y^\ell \in \hat{Y}_1(x)$, $P(y^\ell) \subseteq P_{\delta_{\min}}(y^\ell)$, and thus $x \in Proj_x(P_{\delta_{\min}}(y^\ell))$. Hence, for $\delta = \delta_{\min}$, U_δ will do. □

Having made all necessary definitions and established fundamental properties, the robust bilevel problem will be investigated over three different uncertainty sets:

- The *original* problem over the polyhedral uncertainty set U given in the input. Let $z^\star$ denote its supremum.
- Its *extended-uncertainty* variant over the polyhedron $U_\delta \supset U$, where δ is defined according to Proposition 3.8. Its supremum will be denoted by z^δ.
- A *discrete-uncertainty* variant over different, discrete uncertainty sets $U^D \subset U_\delta$. This discrete-uncertainty variant is solved according to the optimistic bilevel assumption, and its optimum is denoted by z^D.

Proposition 3.9. *Assume that the robust bilevel problem (1)-(4), solved over uncertainty sets U_1 and U_2, admits finite suprema z_1 and z_2, respectively. If $U_1 \subseteq U_2$, then $z_1 \geq z_2$.*

Proof. Easily follows from the definitions. □

Corollary 3.10. *For any $\varepsilon > 0$, there exists $\delta > 0$ such that $z^\star$, z^δ, and z^D as defined above, satisfy $z^\star \geq z^\delta \geq z^\star - \varepsilon(|z^\star| + 1)$ from Propositions 3.9 and 3.8, and $z^D \geq z^\delta$ from Proposition 3.9, and consequently, $z^D \geq z^\star - \varepsilon(|z^\star| + 1)$ by transitivity.*

4. Generic solution method

4.1. Definition of the algorithm

This section gives an overview of the algorithm for solving RBOP as defined in (1)-(4). The algorithm addresses finding a feasible solution that approaches the supremum arbitrarily well. For this purpose, it takes as input parameter δ, which defines the extended uncertainty set U_δ as presented in Section 3, and in turn controls the solution quality and the convergence rate of the algorithm. The algorithm is based on solving the discrete-uncertainty variant of the problem over uncertainty set U^D, where U^D is built up iteratively, in such a way that $U^D \subset U_\delta$. The solution of the discrete-uncertainty variant yields a leader's decision x whose performance can be evaluated on the original uncertainty set U. The outcome of that test either proves that a sufficiently good solution is found, in which case the algorithm terminates; or it gives guidance on how U^D should be extended in future iterations. This idea can be elaborated into an algorithm as follows. In the mathematical programs solved in different steps of the algorithm, the decision variables are highlighted with bold font. A detailed illustration of each step on the specific demand response management problem is presented later in Section 5.

(1) Initialize $U^D := \emptyset$, $f_{\text{best}} := -\infty$, $f_{\text{bound}} := \infty$, and $iter := 1$. Compute initial tariff x^1 by solving the high-point relaxation

$$\max f(\mathbf{x}, \mathbf{y}) \text{ subject to } \mathbf{x} \in X,\ \mathbf{y} \in Y$$

if this is efficiently solvable for the specific problem, or take an arbitrary $x^1 \in X$. Go to Step 3.

(2) Solve the discrete-uncertainty variant of the robust bilevel problem over uncertainty set $U^D = \{u^1, \dots, u^K\}$:

$$\text{Maximize } \mathbf{z} \tag{5}$$

subject to

$$\mathbf{z} \le f(\mathbf{x}, \mathbf{y}^k), \qquad k \in [K] \tag{6}$$

$$\mathbf{x} \in X \tag{7}$$

$$\mathbf{y}^k \in \Omega(\mathbf{x}, u^k), \qquad k \in [K]. \tag{8}$$

Here, inequalities (6) state that z is the worst-case leader's objective value achieved by tariff x over the discrete uncertainty values u^k. Constraint (7) expresses that x is feasible for the leader. Finally, constraint (8) states that y^k is an optimal follower response to (x, u^k). Similarly to the common method for transforming deterministic linear bilevel problems into single-level MILPs [30], problem (5)-(8) can be reformulated into a MILP as follows. Constraint (6) contains the bilinear term $x^T y$ in $f(x, y) = (c + x)^T y$, which can be substituted out by exploiting the equality of the primal and dual objectives in the followers' linear problem. In (7), X is a polyhedron, and therefore, this constraint can be formulated directly via a set of linear inequalities. Finally, constraint (8) can be converted into linear inequalities by exploiting the KKT conditions for the followers' linear problem, and linearizing the complementarity constraints by introducing additional binary variables. With this, the reformulation of (5)-(8) into a single-level MILP is complete. The reformulation is illustrated on the specific application in Section 5.3.2.

The solution of this problem defines tariff x^{iter} and solution value z^{iter}. Update $f_{\text{bound}} = z^{iter}$. If $f_{\text{best}} \geq f_{\text{bound}}$ then return $(x_{\text{best}}, y_{\text{best}}, u_{\text{best}})$.

(3) Compute the followers' worst-case response to tariff x^{iter} over the original uncertainty set U.

$$\text{Minimize } f(x^{iter}, \mathbf{y}) \tag{9}$$

subject to

$$\mathbf{u} \in U \tag{10}$$

$$\mathbf{y} \in \Omega(x^{iter}, \mathbf{u}). \tag{11}$$

Observe that finding the worst-case response to the fixed tariff x^{iter} requires *minimizing* the same objective (9) that is originally *maximized* in the RBOP. Like above, constraint (10) stating that u must belong to the polyhedral set U can be encoded into a set of linear inequalities; whereas line (11) expressing that y is an optimal follower response to (x^{iter}, u) can be reformulated into linear inequalities with additional binary variables. Therefore, problem (9)-(11) can be transformed into a MILP.

The solution of the MILP defines y^{iter} and u. If $f_{\text{best}} < f(x^{iter}, y^{iter})$, then update $f_{\text{best}} := f(x^{iter}, y^{iter})$, $x_{\text{best}} := x^{iter}$, $y_{\text{best}} = y^{iter}$ and $u_{\text{best}} = u$. If $f_{\text{best}} \geq f_{\text{bound}}$ then return $(x_{\text{best}}, y_{\text{best}}, u_{\text{best}})$.

(4) Compute the so-called *characteristic utility* u^{iter} for the given tariff x^{iter} and worst-case response y^{iter}, i.e., the utility value that yields y^{iter} as the followers' response to (x', u^{iter}) for every tariff x' in the largest possible Δ-environment of the fixed tariff x^{iter}:

$$\text{Maximize } \boldsymbol{\Delta} \tag{12}$$

subject to

$$\mathbf{u} \in U_\delta \tag{13}$$

$$(\mathbf{u} - x^{iter})^T (y^{iter} - y^\ell) \geq \boldsymbol{\Delta}, \qquad \forall y^\ell \in N(y^{iter}). \tag{14}$$

The objective (12) is maximizing the characteristic radius of the environment of x^{iter} where y^{iter} is the unique response to every (x', u). By constraint (13), the characteristic utility u must belong to the extended uncertainty set U_δ. Inequality (14) states that the given response y^{iter} yields a followers' objective at least Δ higher than any alternative response y^ℓ, where $N(y^{iter}) := \{y^\ell \in \hat{Y} \ : \ y^\ell \text{ is a neighbor of } y^{iter} \text{ in } Y\}$ denotes the set of neighboring vertices of y^{iter} in Y. Add the resulting u^{iter} to U^D.

(5) $iter := iter + 1$. Go to Step 2.

Formal statements about the convergence properties of the algorithm are made and proven in Section 4.3.

4.2. *Graphical illustration*

This section provides a graphical illustration to the above formal description. For each $y \in Y$, $P(Y)$ is a convex polyhedron in $X \times U$. Observe that the followers' objective is invariant to modifications $(x', u') = (x + v, u + v)$, where v is any vector of the same dimension as x and u. Hence, $P(y)$ is an *extrusion* of a polyhedron $X' \subseteq X$ along the diagonal directions of the $X \times U$ space, as shown for one-dimensional X and U spaces in Figure 1. The polyhedra are limited further by constraints defining X (horizontal, lower and upper sides in Figure 1) and those defining U (vertical, left and right sides in Figure 1).

Polyhedra $P(y)$, $y \in \hat{Y}$ define a disjoint partitioning of the $X \times U$ space. If (x, u) is an *internal* point of $P(y)$, then for tariff x and utility u, the unique optimal followers' response is y. On the other hand, if point (x, u) is located on the boundaries of multiple polyhedra $P(y^1)$, $P(y^2)$, ..., $P(y^m)$, then any of the responses $y^1, y^2, ..., y^m$ is an optimal followers' response, and different optimal responses may yield different profits for the leader. This phenomenon, similar to the optimistic versus pessimistic cases in the deterministic problem, brings substantial additional difficulty when solving the robust bilevel problem.

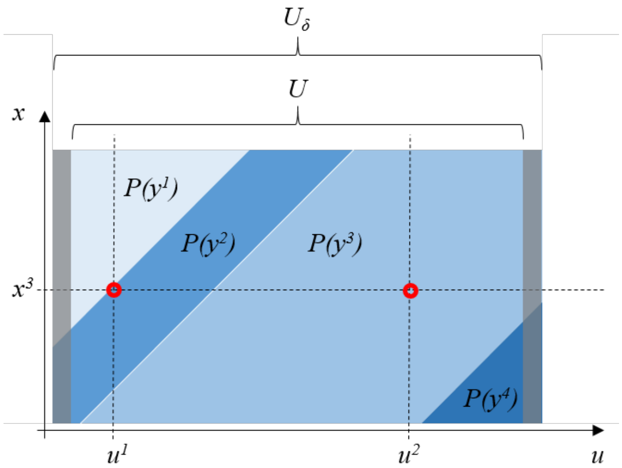

Figure 1. Visualization of the algorithm: solution of the discrete-uncertainty variant over $U^D = \{u^1, u^2\}$.

In the example of Figure 1, after the completion of the second iteration of the algorithm, the discrete uncertainty set contains two elements, $U^D = \{u^1, u^2\}$. Then, in the third iteration, the discrete-uncertainty variant is solved over this set, which results in tariff x^3. Since x^3 is the tariff that maximizes the leader's profit, it is located at an intersection of a vertical line corresponding to u^k (u^1 in the current example) and a boundary of some polyhedra $P(y^\ell)$ ($P(y^1)$ and $P(y^2)$ in the example).

The horizontal line corresponding to x^3 intersects three polyhedra, $P(y^1)$, $P(y^2)$, and $P(y^3)$, which means that the followers may return a response of y^1, y^2, and y^3, depending on the actual value of the uncertain utility u. On the other hand, the followers will never return y^4 as a response to x^3.

Now, let us have a closer look at the solution of the discrete-uncertainty variant over $U^D = \{u^1, u^2\}$. This solution correctly determines y^3 as the only possible followers' response to (x^3, u^2). At the same time, for (x^3, u^1), the two possible responses, y^1 and y^2, yield different profits for the leader, and due to the implicit optimistic assumption, the discrete-uncertainty variant accounts for the most favorable response. Assume this favorable response is y^1, i.e., $f(x^3, y^1) > f(x^3, y^2)$. The possible response y^2 is omitted by the discrete-uncertainty variant.

At this point, two cases must be distinguished. If $f(x^3, y^2) \geq f(x^3, y^3)$, then the algorithm finds y^3 as a worst-case response to x^3 over the original uncertainty set U. Then, the robust solution found and the bound given by the discrete-uncertainty variant coincide, which means that the algorithm terminates with the optimal solution x^3.

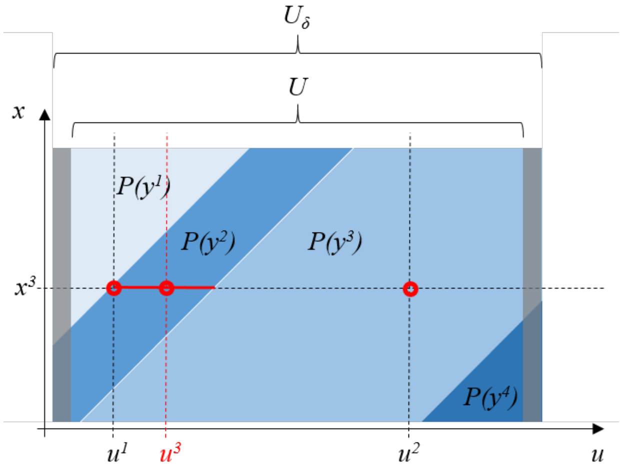

Figure 2. Visualization of the algorithm: if y^2 is the worst-case response to x^3, then the characteristic utility for (x^3, y^2), denoted by u^3, must be added to U^D.

Alternatively, in the unlucky case that $f(x^3, y^2) < f(x^3, y^3)$, the worst-case response to tariff x^3 involves y^2, a response omitted earlier by the discrete-uncertainty variant. Then, to fix this shortcoming, the algorithm computes the characteristic utility for tariff x^3 and response y^2, which corresponds to the mid-point of the intersection of the horizontal line x^3 and polyhedron $P(y^2)$, denoted by u^3 in Figure 2. This is added to collection U^D, and the algorithm continues with $U^D = \{u^1, u^2, u^3\}$ in the next iteration. This ensures that the discrete-uncertainty variant accounts for response y^2 in the largest possible environment of tariff x^3.

An important special case is displayed via another example in Figure 3, where the worst-case response y^1 is taken for a small range of uncertainty values near a vertex of $P(y^1)$ on the boundary of uncertainty set U. This range corresponds to the short red line in the left of Figure 3. The characteristic utility in each subsequent iteration is defined by the midpoint of this short red line. In such a case, subsequent iterations involve slightly modified tariffs (moving downwards in the diagram) and slightly shifted characteristic utilities (moving leftwards). If characteristic utilities were computed over the original uncertainty set U, then this would result in infinite iterations converging to the nearby vertex of $P(y)$. Yet, by computing the characteristic utilities over the extended uncertainty set U_δ, the algorithm arrives in finitely many steps at a characteristic utility $u^3 \in U_\delta \setminus U$. Then, in the next iteration, two cases might occur. Typically, solving the discrete-uncertainty variant will result in a slightly modified tariff x^4 that just avoids y^1 as a worst-case response (see Figure 4); the discrete-uncertainty solution and the worst-case response coincide, which means that tariff x^4 is a close-to-optimal solution. Alternatively, search may continue in completely different regions of the tariff space.

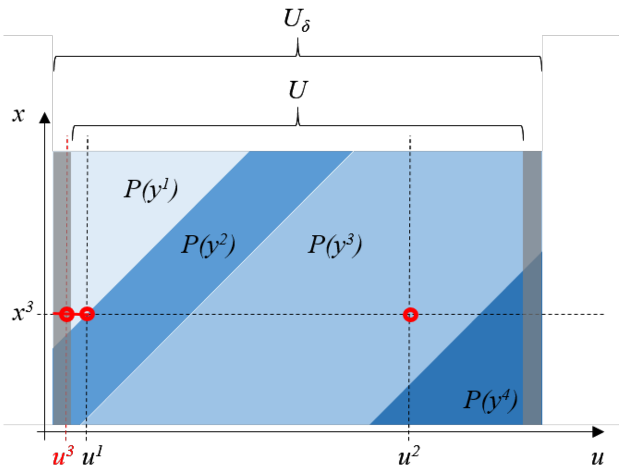

Figure 3. Visualization of the algorithm: an important special case where the worst-case response y^1 is taken for a small range of uncertainty values, near a vertex of $P(y^1)$ on the boundary of uncertainty set U.

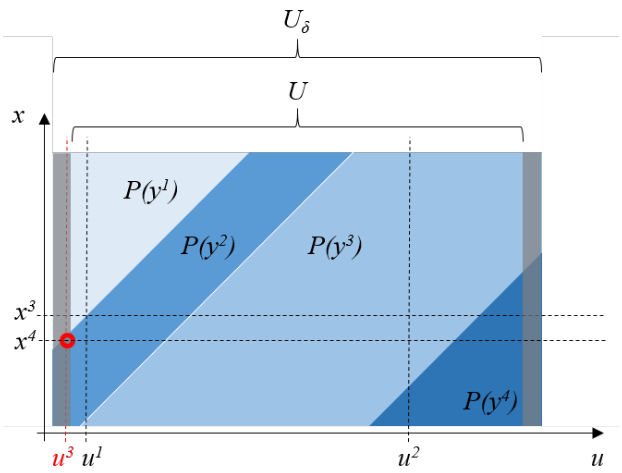

Figure 4. Visualization of the algorithm: computing characteristic utilities over the extended uncertainty set U_δ ensures that iterations arrive at a characteristic utility $u^3 \in U_\delta \setminus U$, and subsequently, to a tariff x^4 that just avoids the unfavorable response y^1.

4.3. Key properties

In this section, key properties of the proposed algorithm are formally proven, including soundness and termination in finitely many steps.

Proposition 4.1. *The solution $(x_{best}, y_{best}, u_{best})$ returned by the algorithm is robust bilevel feasible. Moreover, for any $\varepsilon > 0$, there exists $\delta > 0$ such that $(x_{best}, y_{best}, u_{best})$ is an ε-optimal solution to the robust bilevel problem.*

Proof. $(x_{\text{best}}, y_{\text{best}}, u_{\text{best}})$ is robust bilevel feasible, because x_{best} is derived as a solution of an optimization problem in Step 1 or Step 2 that involves constraint $x \in X$, and

$y_{\text{best}}, u_{\text{best}}$ are obtained in Step 3 as the worst-case response. The leader's objective function value on this solution is f_{best} computed in Step 3.

For a fixed $\varepsilon > 0$, there exists $\delta > 0$ such that $z^\delta \geq z^\star - \varepsilon(|z^\star|+1)$ by Corollary 3.10. With such a parameter δ, $(x_{\text{best}}, y_{\text{best}}, u_{\text{best}})$ is ε-optimal, because

$$f_{\text{best}} \geq f_{\text{bound}} \geq z^\delta \geq z^\star - \varepsilon(|z^\star| + 1),$$

where the first inequality is the termination condition of the algorithm; the second is implied by the fact that f_{bound} is the objective value of the optimistic problem over the discrete uncertainty set $U^D \subset U_\delta$, and the third is ensured by the choice of δ. $\square$

Next we prove that in Step 4, when finding a characteristic utility u^{iter}, it suffices to use only the neighbors of y^{iter}.

Proposition 4.2. *Let u^{iter} be an optimal solution of (12)-(14) of optimum value Δ^{iter}. Then $(u^{iter} - x^{iter})(y^{iter} - y^\ell) \geq \Delta^{iter}$ for all $y^\ell \in \hat{Y}$.*

Proof. To simplify notation, let $w := u^{iter} - x^{iter}$. Consider the cone with apex y^{iter} and generated by the rays $(y^{iter} - y)$ for neighboring vertex $y \in N(y^{iter})$ of y^{iter}. Since this cone encompasses Y, each vertex y^ℓ of Y can be expressed as linear combination of the vectors $(y^{iter} - y)$, $y \in N(y^{iter})$. That is, there exist coefficients $\lambda_j \geq 0$ such that $\sum_{y^j \in N(y^{iter})} \lambda_j \geq 1$, and

$$y^{iter} - y^\ell = \sum_{y^j \in N(y^{iter})} \lambda_j (y^{iter} - y^j).$$

Multiplying both sides by vector w from the left, we obtain

$$w(y^{iter} - y^\ell) = \sum_{y^j \in N(y^{iter})} \lambda_j w(y^{iter} - y^j) \geq \sum_{y^j \in N(y^{iter})} \lambda_j \Delta^{iter} \geq \Delta^{iter}.$$

$\square$

Termination in finitely many steps is ensured by adding a characteristic utility in each iteration that guarantees sufficient progress in each iteration of the search procedure. Recall the set U_δ and Proposition 3.8.

Proposition 4.3. *For every instance of the robust bilevel problem, there exists a constant $\Delta_{\min} > 0$ such that problem (12)-(14) admits an optimal solution with $\Delta \geq \Delta_{\min}$. Moreover, $\Delta_{\min}$ is common over all iterations of the algorithm.*

Proof. We show that for each vertex y^ℓ of Y, and $\Delta_\ell > 0$, there is a vector u^ℓ such that for any $(x, u) \in P(y^\ell)$,

$$(u + u^\ell - x)(y^\ell - y^j) \geq \Delta_\ell, \ \forall j \in [L]. \tag{15}$$

Since Y is a polyhedron, there exists a hyperplane containing y^ℓ with normal vector u^ℓ such that

$$u^\ell(y^\ell - y^j) \geq \Delta_\ell, \ \forall j \in [L].$$

This proves (15). For each $\ell \in [L]$, we choose $\Delta_\ell > 0$ small enough such that there exists u^ℓ satisfying (15) and also $u + u^\ell \in U_\delta$ for any $(x, u) \in P(y^\ell)$, where we exploit that U_δ is full-dimensional. Finally, we define $\Delta_{\min} := \min_{\ell\in[L]} \Delta_\ell$, and the statement is proved. $\square$

Assume that in Step 4 of iteration *iter* of the algorithm, characteristic utility u^{iter} is computed for tariff x^{iter} and corresponding worst-case response y^{iter}. Recall that $N(y^{iter})$ denotes the neighboring vertices of y^{iter} in Y. Then, let us denote by $\hat{\varphi} := (u^{iter} - x^{iter})y^{iter}$ and $\varphi^\ell := (u^{iter} - x^{iter})y^\ell$ the followers' objective function values corresponding to responses y^{iter} and $y^\ell \in N(y^{iter})$, respectively. By constraint (14) and Proposition 4.3 we have for all $y^\ell \in N(y^{iter})$,

$$(\hat{\varphi} - \varphi^\ell) = (u^{iter} - x^{iter})y^{iter} - (u^{iter} - x^{iter})y^\ell \geq \Delta \geq \Delta_{\min} > 0. \tag{16}$$

Since Y is a polytope, there exist upper bounds $\tilde{y}_i$ such that $|y_i| \leq \tilde{y}_i$ for each $y \in Y$ and all coordinates i. In the sequel, let $||\cdot||_{\max}$ denote the maximum norm (also called the infinity norm or L^∞ norm), i.e., the largest absolute value of the components of a vector.

Proposition 4.4. *For x^{iter}, u^{iter} and Δ as defined above, assume x' is a tariff such that $||x' - x^{iter}||_{\max} \leq \theta = \Delta/(2\sum_i \tilde{y}_i)$. Then, y^{iter} is also an optimal follower response for x' and u^{iter}. Moreover, if $||x' - x||_{\max} < \theta$, then y^{iter} is the unique optimal response of the followers.*

Proof. For any given i, a change of θ in the single tariff component x_i incurs a change of at most $\theta\tilde{y}_i$ in both of the followers' objective function values $\hat{\varphi}$ and φ^ℓ. If all coefficients are allowed to change at the same time, then both of $\hat{\varphi}$ and φ^ℓ can change by at most $\theta \sum_i \tilde{y}_i = \frac{\Delta}{2}$. Accordingly, for the new tariff x' it holds that

$$(u^{iter} - x')y^{iter} \;\geq\; (u^{iter} - x^{iter})y^{iter} - \frac{\Delta}{2} \;\geq\; (u^{iter} - x^{iter})y^\ell + \frac{\Delta}{2} \;\geq\; (u^{iter} - x')y^\ell.$$

The first and the third inequalities follow from the limited change of the followers' objective upon the variation of x^{iter}, whereas the second inequality is ensured by (16). This means that y^{iter} is also optimal for x' and u^{iter}.

The second part of the proposition, corresponding to strict inequality $||x' - x||_{\max} < \theta$ follows analogously. $\square$

Theorem 4.5. *For any given $\varepsilon > 0$, there exists $\delta > 0$ such that the proposed algorithm terminates in finitely many steps and outputs an ε-optimal solution.*

Proof. By Corollary 3.3, the worst-case response y^{iter} computed in Step 3 of the algorithm is one of the finitely many vertices of Y. Consequently, it suffices to show that each $\hat{y} \in \hat{Y}$ can be received as a worst-case response in finitely many iterations.

An indirect proof is given, assuming that $\hat{y}$ occurs as worst-case response infinitely many times over iterations $iter = 1, 2, \ldots$. We investigate two separate cases according to whether there exist two iterations $iter_1$ and $iter_2$ such that $iter_1 < iter_2$, $y^{iter_1} = y^{iter_2} = \hat{y}$, and $||x^{iter_1} - x^{iter_2}||_{\max} < \theta_{\min} = \frac{\Delta_{\min}}{2\sum_i \tilde{y}_i}$.

In case such iterations $iter_1$ and $iter_2$ exist, then let u^{iter_1} denote the characteristic utility computed for tariff x^{iter_1} and worst-case response $\hat{y}$ in Step 4 of the algorithm. Then, in iteration $iter_2$, we have that $\hat{y}$ is the unique optimal response for x^{iter_2} and

u^{iter_1} as well by Proposition 4.4. This implies that in iteration $iter_2$, Step 2, the value of the discrete-uncertainty solution f_{bound} matches the value of the worst-case response over the original uncertainty set, which is f_{best}, and the algorithm terminates. This contradicts our assumption.

Alternatively, it might happen that the tariff values differ substantially, i.e., $||x^{iter_1} - x^{iter_2}||_{\max} \geq \theta_{\min}$, in any two iterations such that $y^{iter_1} = y^{iter_2} = \hat{y}$. However, the existence of infinitely many tariff values whose distance is greater than the predefined positive constant contradicts the assumption that X is bounded. Hence, the assumption of infinitely many iterations with the same $\hat{y}$ results in a contradiction. Since the worst-case response y^{iter} in Step 3 is always a vertex of Y, and Y has a finite number of vertices, the algorithm terminates in finitely many steps.

Finally, by Proposition 4.1, the solution returned by the algorithm is ε-optimal. □

4.4. Determining the solution quality

Bounding the error of the solution requires calculating an upper bound on the supremum $z^\star$. Observe that f_{bound} computed by the algorithm is inappropriate for this purpose, since it is calculated over the discrete uncertainty set U^D, which may contain uncertainty vectors $u^k \in U_\delta \setminus U$, i.e., which are outside the original uncertainty set U.

To compute a correct upper bound, we depart from the discrete uncertainty set U^D, project all the vectors $u^k \in U^D$ back to U, and then solve the discrete-uncertainty variant (5)-(8) over the projections. Since a finite subset of U is used for solving the optimistic relaxation of RBOP, this leads to an upper bound on $z^\star$ that we denote by UB. Then, $\frac{\text{UB} - f_{\text{best}}}{|\text{UB}|+1}$ is the *relative error* of the best solution.

It remains to sketch how to project the vectors $u^k \in U^D \setminus U$ back to U. Let $\tilde{u}^k \in U$ be a closest vector to u^k in the L^1 norm. Then, $\tilde{u}^k$ is an optimal solution of the optimization problem

$$
\begin{aligned}
&\min \sum_i v_i \\
&\text{s.t.} \\
&u_i^k - \tilde{u}_i^k \leq v_i, \quad \forall i \\
&\tilde{u}_i^k - u_i^k \leq v_i, \quad \forall i \\
&\tilde{u}^k \in U.
\end{aligned}
$$

Note that if $u^k \in U$, then the optimum value is 0, and $\tilde{u}^k = u^k$. One may also use the L^2 norm, which yields a mathematical program with a convex quadratic, separable objective function and the same linear constraints as above.

5. Application to Demand Response Management

5.1. Problem Definition

The proposed generic solution method is illustrated on a specific problem involving demand response management in smart electricity grids. The robust formulation extends the earlier deterministic *Simple Multi-period Energy Tariff Optimization Problem* (SMETOP) [31] with uncertainty as follows.

The leader in the bilevel problem is an electricity retailer who addresses setting a time-of-use electricity tariff for its M consumers. The consumers, who act as M independent followers in the bilevel problem, respond to the electricity tariff by determining their consumption over time in order to maximize their utilities and minimize their electricity costs.

The leader purchases electricity at the wholesale market at given unit prices p_t over the finite time horizon $t = 1, ..., T$. Then, it must set consumer prices x_t subject to regulations agreed a priori expressed in the form arbitrary linear inequalities, which define the polyhedral set X. Each follower i responds to this tariff independently of other followers by determining its consumption y_{it} in such a way that $\sum_{t=1}^{T}\sum_{i=1}^{M}(u_{it} - x_t)y_{it}$ is maximized, which corresponds to maximizing utility and minimizing the cost of electricity. The per period consumption y_{it} must respect the lower bound $\underline{y}_{it}$ and upper bound $\bar{y}_{it}$, and for each $i \in [M]$, total consumption $\sum_{t=1}^{T} y_{it}$ must fall between the lower and upper bounds $\underline{d}_i$ and $\bar{d}_i$. The lower and upper bounds, $\underline{y}_{it}$, $\bar{y}_{it}$, $\underline{d}_i$, and $\bar{d}_i$, are all non-negative for all $i \in [M]$ and $t \in [T]$.

Unlike in the deterministic case, the leader is only partially aware of the followers' parameters. Namely, the leader knows the bounds $\underline{y}_{it}$, $\bar{y}_{it}$, $\underline{d}_i$ and $\bar{d}_i$, but the actual value of the followers' perceived utilities u_{it} is unknown to the leader. Instead, the ensemble of all utility values comes from a given polyhedral uncertainty set U. This also allows that the utilities of different followers can be interrelated. Then, the objective of the leader is maximizing its worst-case profit, $\sum_{t=1}^{T}\sum_{i=1}^{M}(x_t - p_t)y_{it}$, over all possible realizations corresponding to different utility values u_{it}. This robust bilevel problem can be formulated as follows:

$$\sup_{x \in X} z \tag{17}$$

$$\min_{u \in U} \left\{ \sum_{t=1}^{T}\sum_{i=1}^{M} (x_t - p_t)\, y_{it} \;:\; y \in \Omega(x,u) \right\} \geq z, \tag{18}$$

where

$$X = \{x \in \mathbb{R}^T \;:\; r^j x \leq r_0^j, j = 1, \ldots, m_a\},$$
$$U = \{u \in \mathbb{R}^{M \times T} \;:\; v^j u \leq v_0^j, j = 1, \ldots, m_b\},$$

and $\Omega(x, u)$ is the set of optimal solutions of the parametric problem

$$\max \sum_{t=1}^{T}\sum_{i=1}^{M}(u_{it} - x_t)\mathbf{y}_{it} \tag{19}$$

$$\underline{d}_i \leq \sum_{t=1}^{T} \mathbf{y}_{it} \leq \bar{d}_i, \quad i \in [M] \tag{20}$$

$$\underline{y}_{it} \leq \mathbf{y}_{it} \leq \bar{y}_{it}, \quad i \in [M],\ t \in [T]. \tag{21}$$

Let Y be the set of feasible solutions of the linear system (20)-(21).

The above demand response management problem fits into the generic robust bilevel programming framework with electricity tariff x_t as the leader's decision variables, consumption y_{it} as the followers' response, and the consumers' utility u_{it} as the uncertain

Table 2. Sample instance for the demand response management problem.

Parameter	Value
M	1
T	3
p	(1, 1, 100)
X	$[0, 10]^3$
$\bar{y}_1 = \bar{y}_2 = \bar{y}_3$	1
$\underline{y}_1 = \underline{y}_2 = \underline{y}_3$	0
$\underline{d} = \bar{d}$	1
U	$0 \leq u_1, u_2 \leq 10$
	$u_3 = 6$
	$u_1 + u_2 \geq 10$
Vertices of U	$u^A = (10, 0, 6)$
	$u^B = (0, 10, 6)$
	$u^C = (10, 10, 6)$

parameters. A minor mismatch comes from the fact that the leader applies the same tariff x_t to all consumers, whereas the generic framework assumes that the leader's variables x have the same number of dimensions as the utility values u_{it}. The latter would directly correspond to applying consumer-specific tariff variables x_{it} and imposing equality constraints $x_{it} = x_{i't}$, $\forall i, i', t$, which can be handled smoothly by the proposed algorithm. Yet, with a slight abuse of the notation, this paper uses the simple formulation with tariff variables x_t.

5.2. Sample instance

The robust bilevel demand response management problem is illustrated on the small sample instance shown in Table 2, with a single follower and three time periods. Since there is a single follower, the follower index is omitted from the notation. The feasible region for the tariff is the box $X = [0, 10]^3$. The consumer must schedule a single unit of load into any of the three periods. The uncertainty set U for the follower's utilities is the convex hull of the points u^A, u^B, and u^C as displayed in the table. For the sake of simplicity, the third coordinate of the vector is fixed to $u_3 = 6$. We refer to the midpoint of the section $\overline{u^A u^B}$ as $u^{Mid} = \frac{u^A + u^B}{2} = (5, 5, 6)$.

Notice that time periods 1 and 2 are favorable for the leader, since it can realize positive profit by purchasing electricity at a low wholesale price. In contrast, any consumption in the unfavorable period 3 inevitably results in a loss of profit for the leader. Therefore, the leader can maximize its profit by (1) motivating the follower to schedule its load to one of first two time periods, and by (2) increasing the tariff as much as it is allowed by the former requirement.

In case the leader decided for the highest possible tariff, $x^1 = (10, 10, 10)$, then for all vertices of U, the follower would schedule its load into the periods favorable for the leader (periods 1 or 2), resulting in a profit of $10 - 1 = 9$. On the other hand, for utility vector u^{Mid}, the same tariff would result in scheduling the entire load into the last time period with a significant loss, $10 - 100 = -90$. The latter is the leader's objective value for tariff x^1.

Obviously, the leader must decrease the tariff in the favorable periods to attract consumption into those periods. For this purpose, the leader can apply tariff $x^2 = (9 - \varepsilon, 9 - \varepsilon, 10)$, which ensures that either $u_1 - x_1 > u_3 - x_3$ or $u_2 - x_2 > u_3 - x_3$ holds for any $u \in U$. Accordingly, the entire load will be scheduled into the favorable periods 1 or 2, resulting in a profit of $9 - \varepsilon - 1 = 8 - \varepsilon$. The latter tariff x^2 is also the

ε-optimal solution for the sample problem instance.

Finally, observe that the supremum, corresponding to tariff $(9, 9, 10)$ and profit value 8, cannot be attained. The proposed algorithm finds the above ε-optimal solution in two iterations, where the first iteration checks the initial tariff $x^1 = (10, 10, 10)$ and leads to the characteristic utility $u^1 = (4-\delta, 4-\delta, 6+\delta) \in U_\delta$, and the second iteration leads directly to $x^2 = (9 - \varepsilon, 9 - \varepsilon, 10)$.

This sample instance also illustrates that RBOP over a polyhedral uncertainty set U differs essentially from the similar problem over the discrete uncertainty set involving only the vertices of U. Hence, although it may sound like a promising idea to look for the optimal solution by considering only the vertices of U, this idea may easily lead to a sub-optimal solution for the original RBOP.

5.3. Application of the Generic Algorithm

This section presents the application of the generic algorithm outlined in Section 4 to solve the robust bilevel demand response management problem. This requires defining the MILP models corresponding to the abstract problems solved in each step of the algorithm, which is presented in separate subsections below. Like above, decision variables are highlighted with bold font.

5.3.1. Step 1: Initialization

Initial tariff values x are set heuristically by maximizing $\sum_t \mathbf{x}_t$ subject to $\mathbf{x} \in X$.

5.3.2. Step 2: Discrete-uncertainty variant

In the second step of the algorithm, the discrete-uncertainty variant with uncertainty set $u^k \in U^D,\ k \in [K]$ is solved. Recall that separate consumption vectors $y^k,\ k \in [K]$ belong to each utility vector u^k, and the lowest profit achieved over the different utility vectors defines the objective value of the solution.

The discrete-uncertainty variant can be converted into a single-level MILP similarly to the well-known technique for the deterministic variant, by using the complementary slackness conditions for the followers' sub-problem. For this purpose, let us first formalize the dual of the followers' sub-problem for a given scenario $k \in [K]$ and given follower $i \in [M]$:

$$\text{Minimize } \bar{d}_i \alpha_i^{k+} - \underline{d}_i \alpha_i^{k-} + \sum_{t=1}^{T} (\bar{y}_{it} \beta_{it}^{k+} - \underline{y}_{it} \beta_{it}^{k-}) \tag{22}$$

subject to

$$\begin{aligned} &\alpha_i^{k+} - \alpha_i^{k-} + \beta_{it}^{k+} - \beta_{it}^{k-} = u_{it}^k - x_t, && t \in [T] \\ &\alpha_i^{k-}, \alpha_i^{k+}, \beta_{it}^{k-}, \beta_{it}^{k+} \geq 0, && t \in [T]. \end{aligned}$$

This formulation uses dual variables α_i^{k+} and α_i^{k-} for the lower and upper bounds in constraint (20), respectively, whereas β_{it}^{k+} and β_{it}^{k-} for the lower and upper bounds in constraint (21) for each scenario $k \in [K]$.

The bilinear terms $\sum_{t=1}^{T} x_t y_{it}$ in the leader's objective function can be reformulated by exploiting strong duality for the followers' linear sub-problem, which implies that the followers' primal and dual objectives are equal:

$$\sum_{t=1}^{T}(u_{it} - \mathbf{x}_t)\mathbf{y}_{it} \;=\; \bar{d}_i\alpha_i^{k+} - \underline{d}_i\alpha_i^{k-} + \sum_{t=1}^{T}(\bar{y}_{it}\beta_{it}^{k+} - \underline{y}_{it}\beta_{it}^{k-}) \quad \forall i \in [M], k \in [K].$$

By rearranging the equation, we get:

$$\sum_{t=1}^{T}\mathbf{x}_t\mathbf{y}_{it} \;=\; \sum_{t=1}^{T}\left(u_{it}^k\mathbf{y}_{it}^k - \bar{y}_{it}\beta_{it}^{k+} + \underline{y}_{it}\beta_{it}^{k-}\right) - \bar{d}_i\alpha_i^{k+} + \underline{d}_i\alpha_i^{k-} \quad \forall i \in [M], k \in [K].$$

This transformation takes us to the following MILP formulation of the discrete-uncertainty variant:

$$\text{Maximize } \mathbf{z} \tag{23}$$

subject to

$$\mathbf{z} \leq \sum_{t=1}^{T}\sum_{i=1}^{m}(u_{it}^k\mathbf{y}_{it}^k - \bar{y}_{it}\beta_{it}^{k+} + \underline{y}_{it}\beta_{it}^{k-} - p_t\mathbf{y}_{it}^k) \;-\; \sum_{i=1}^{m}(\bar{d}_i\alpha_i^{k+} - \underline{d}_i\alpha_i^{k-}), \quad k \in [K] \tag{24}$$

$$\begin{aligned}
&\mathbf{x} \in X && && (25)\\
&0 \leq \alpha_i^{k+} \perp \bar{d}_i - \sum_{t=1}^{T}\mathbf{y}_{it}^k \geq 0, && i \in [M],\; k \in [K] && (26)\\
&0 \leq \alpha_i^{k-} \perp \sum_{t=1}^{T}\mathbf{y}_{it}^k - \underline{d}_i \geq 0, && i \in [M],\; k \in [K] && (27)\\
&0 \leq \beta_{it}^{k+} \perp \bar{y}_{it} - \mathbf{y}_{it}^k \geq 0, && i \in [M],\; t \in [T],\; k \in [K] && (28)\\
&0 \leq \beta_{it}^{k-} \perp \mathbf{y}_{it}^k - \underline{y}_{it} \geq 0, && i \in [M],\; t \in [T],\; k \in [K] && (29)\\
&\alpha_i^{k+} - \alpha_i^{k-} + \beta_{it}^{k+} - \beta_{it}^{k-} = u_{it}^k - \mathbf{x}_t, && i \in [M],\; t \in [T],\; k \in [K]. && (30)
\end{aligned}$$

In the MILP model, the objective (23) is maximizing the worst-case profit z, which is determined by the minimum of the profits realized in the individual scenarios (24). The core of the model is composed of leader's constraints (25) and the complementary slackness conditions for the followers' sub-problem (26)-(30). Here, $0 \leq L \perp R \geq 0$ denotes that $L \geq 0$, $R \geq 0$, and either $L = 0$, or $R = 0$. In an actual implementation, the corresponding constraints can be translated into indicator constraints or classical *big-M* constraints. The optimal solution of this MILP defines tariff x^{iter} in the given iteration.

5.3.3. Step 3: Worst-case followers' response

The problem of finding a worst-case response y^{iter} for the fixed tariff x^{iter} can be captured by the following MILP, again, by applying the KKT reformulation of the followers' sub-problem:

$$\text{Minimize} \sum_{t=1}^{T}\sum_{i=1}^{M}\left(x_t^{iter}-p_t\right)\mathbf{y}_{it} \tag{31}$$

subject to

$$\begin{aligned}
&\mathbf{u}\in U, && && (32)\\
&0\le \alpha_i^+ \perp \bar{d}_i-\sum_{t=1}^{T}\mathbf{y}_{it}\ge 0, && i\in[M] && (33)\\
&0\le \alpha_i^- \perp \sum_{t=1}^{T}\mathbf{y}_{it}-\underline{d}_i\ge 0, && i\in[M] && (34)\\
&0\le \beta_{it}^+ \perp \bar{y}_{it}-\mathbf{y}_{it}\ge 0, && i\in[M],\ t\in[T] && (35)\\
&0\le \beta_{it}^- \perp \mathbf{y}_{it}-\underline{y}_{it}\ge 0, && i\in[M],\ t\in[T] && (36)\\
&\alpha_i^+-\alpha_i^-+\beta_{it}^+-\beta_{it}^-=\mathbf{u}_{it}-x_t^{iter}, && i\in[M],\ t\in[T]. && (37)
\end{aligned}$$

The objective (31) is minimizing the leader's profit. Constraint (32) states that the corresponding utility u must belong to the original uncertainty set U, whereas (33)-(37) encode the optimality conditions for the followers' sub-problem. Variables y in the optimal solution define y^{iter}.

5.3.4. Step 4: Characteristic utility

Upon receiving a worst-case followers' response y^{iter} for tariff x^{iter} worse than what was foreseen based on the solution of the discrete-uncertainty variant, we determine the characteristic utility u^{iter} for (x^{iter}, y^{iter}) to be added to U^D. Before explaining how we find it, we characterize those vectors u such that y^{iter} is an optimal follower response for (x^{iter}, u).

Observation 5.1. *If y is a vertex of Y, then for each $i \in [M]$ one of the following cases holds:*

a) $\underline{d}_i < \sum_{\tau=1}^{T} y_{i\tau} < \bar{d}_i$ *and for each* $t \in [T]$*:* $y_{it} \in \{\underline{y}_{it}, \bar{y}_{it}\}$.

b) $\sum_{\tau=1}^{T} y_{i\tau} \in \{\underline{d}_i, \bar{d}_i\}$, *and for all but at most one* $t \in [T]$, $y_{it} \in \{\underline{y}_{it}, \bar{y}_{it}\}$.

Proposition 5.2. *Given some vector u of appropriate dimensions. y^{iter} is an optimal follower response for (x^{iter}, u) if and only if for each $i \in [M]$, u satisfies the following conditions:*

i) $u_{it} - x_t^{iter} \ge u_{it'} - x_{t'}^{iter}$ *for all* $t \neq t' \in T$ *such that* $\underline{y}_{it} < y_{it}^{iter} \ \wedge\ \bar{y}_{it'} > y_{it'}^{iter}$.

ii) $u_{it} - x_t^{iter} \ge 0$ *for all* $t \in [T]$ *such that* $\underline{y}_{it} < y_{it}^{iter} \ \wedge\ \underline{d}_i < \sum_{\tau=1}^{T} y_{i\tau}^{iter}$.

iii) $u_{it} - x_t^{iter} \le 0$, *for all* $t \in [T]$ *such that* $\bar{y}_{it} > y_{it}^{iter} \ \wedge\ \bar{d}_i > \sum_{\tau=1}^{T} y_{i\tau}^{iter}$.

Proof. The necessity of the conditions is obvious, so we turn to sufficiency. We will prove that no neighboring vertex y of y^{iter} in Y has a smaller cost than y^{iter} provided u satisfies the conditions of the statement. To this end, first we characterize the neighboring vertices of y^{iter}. We can reformulate (19)-(21) as follows:

$$\begin{aligned}
&\max \sum_{i=1}^{M}\sum_{t=1}^{T}(u_{it}-x_t)\mathbf{y}_{it}\\
&\sum_{t=1}^{T}\mathbf{y}_{it}-\mathbf{s}_i=0, \quad i\in[M]\\
&\underline{d}_i \le \mathbf{s}_i \le \bar{d}_i, \quad i\in[M]\\
&\underline{y}_{it} \le \mathbf{y}_{it} \le \bar{y}_{it}, \quad i\in[M],\ t\in[T].
\end{aligned}$$

Observe that this linear program decomposes into M independent linear programs, one for each $i \in M$. Let LP_i denote

$$\begin{aligned}
&\max \sum_{t=1}^{T}(u_{it}-x_t)\mathbf{y}_{it}\\
&\sum_{t=1}^{T}\mathbf{y}_{it}-\mathbf{s}_i=0 \qquad\qquad (38)\\
&\underline{d}_i \le \mathbf{s}_i \le \bar{d}_i\\
&\underline{y}_{it} \le \mathbf{y}_{it} \le \bar{y}_{it}, \quad t\in[T].
\end{aligned}$$

LP_i consists of a single equation (38), and lower and upper bounds on the variables. Therefore, in any basic solution, all but one variables are at lower or upper bounds, and exactly one variable is in the basis. Consequently, if we express the vector $y_i^{iter} = (y_{it}^{iter} \ : \ t \in [T])$ as a basic solution of LP_i, then either $\mathbf{s}_i$, or one of the $\mathbf{y}_{it}$ is basic, and the rest of the variables are non-basic. Each neighbor of y_i^{iter} can be obtained by exchanging a basic and a nonbasic variable, or by swapping a variable from its lower bound to its upper bound or vice versa. Hence, each neighbor of y_i^{iter} in LP_i can be obtained by one of the following transformations:

i) Some $y_{it}^{iter} > \underline{y}_{it}$ is decreased by a positive amount and a distinct $y_{it'}^{iter} < \bar{y}_{it'}$ is increased by the same amount.
ii) Some $y_{it}^{iter} > \underline{y}_{it}$ is decreased by a positive amount.
iii) Some $y_{it}^{iter} < \bar{y}_{it}$ is increased by a positive amount.

Observe that the conditions of the statement ensure that none of the above transformations may decrease the objective value of the solution, hence, the neighbors of y_i^{iter} in LP_i cannot have a smaller objective value than y_i^{iter}. Finally, since each neighboring vertex of y^{iter} in Y can be obtained by one of the above transformations, the statement follows. □

In the remainder of this section, we describe two alternative methods for finding u^{iter}.

5.3.4.1. Method I. The first method is derived directly from the general method of Section 4.3.

$$\text{Maximize } \mathbf{\Delta} \tag{39}$$

subject to

$$\mathbf{u} \in U_\delta \tag{40}$$

$$(\mathbf{u}_{it} - x_t^{iter} - \mathbf{u}_{it'} + x_{t'}^{iter}) \min(y_{it}^{iter} - \underline{y}_{it},\ \bar{y}_{it'} - y_{it'}^{iter}) \geq \mathbf{\Delta}, \quad i \in [M],\ t \neq t' \in [T] : \underline{y}_{it} < y_{it}^{iter} \ \wedge\ \bar{y}_{it'} > y_{it'}^{iter} \tag{41}$$

$$(\mathbf{u}_{it} - x_t^{iter}) \min(y_{it}^{iter} - \underline{y}_{it},\ \sum_{\tau=1}^{T} y_{i\tau}^{iter} - \underline{d}_i) \geq \mathbf{\Delta}, \quad i \in [M],\ t \in [T] : \underline{y}_{it} < y_{it}^{iter} \ \wedge\ \underline{d}_i < \sum_{\tau=1}^{T} y_{i\tau}^{iter} \tag{42}$$

$$-(\mathbf{u}_{it} - x_t^{iter}) \min(\bar{y}_{it} - y_{it}^{iter},\ \bar{d}_i - \sum_{\tau=1}^{T} y_{i\tau}^{iter}) \geq \mathbf{\Delta}, \quad i \in [M],\ t \in [T] : \bar{y}_{it} > y_{it}^{iter} \ \wedge\ \bar{d}_i > \sum_{\tau=1}^{T} y_{i\tau}^{iter}. \tag{43}$$

The objective (39) is maximizing the slack Δ. The characteristic utility u must belong to the extended uncertainty set U_δ (40). Then, inequalities (41)-(43) ensure the proper slack for the different neighbors of y^{iter} in $\hat{Y}$. First, constraint (41) states this requirement for neighbors of y^{iter} received by moving $\min(y_{it}^{iter} - \underline{y}_{it},\ \bar{y}_{it'} - y_{it'}^{iter}) > 0$ load from period t to period t'. Observe that this is the highest possible amount of load that can be moved between the two periods. Similarly, inequality (42) bounds the slack for neighbors obtained by decreasing the load in some period t by the highest possible amount, $\min(y_{it}^{iter} - \underline{y}_{it},\ \sum_{\tau=1}^{T} y_{i\tau}^{iter} - \underline{d}_i) > 0$. Finally, constraint (43) achieves the same for neighbors of y^{iter} received by increasing the load in some period t by the maximum possible amount.

Observe that the left hand side of constraints (41)-(43) equals the difference of the followers' objectives in case of followers' responses y^{iter} and some $y^\ell \in N(y^{iter})$, i.e., $(\mathbf{u} - x^{iter})^T(y^{iter} - y^\ell)$, as defined for the generic method in constraint (14). Hence, by Proposition 4.3, the optimal solution value of the LP (39)-(43) is always at least $\Delta_{\min}$ for some universal $\Delta_{\min} > 0$.

5.3.4.2. Method II. In the following mathematical program, we seek a vector u that satisfies the conditions of Proposition 5.2 with a positive slack θ. This ensures that y^{iter} is an optimal follower response not only for (x^{iter}, u), but also for all $x' \in X$ such that $||x' - x^{iter}||_{\max} \leq \theta/2$, see Proposition 5.4.

$$\text{Maximize } \theta \tag{44}$$

subject to

$$\begin{aligned}
\mathbf{u} &\in U_\delta && && (45)\\
\mathbf{u}_{it} - x_t^{iter} &\geq \mathbf{u}_{it'} - x_{t'}^{iter} + \theta, && i \in [M],\ t \neq t' \in [T]:\ \underline{y}_{it} < y_{it}^{iter} \ \wedge\ \bar{y}_{it'} > y_{it'}^{iter} && (46)\\
\mathbf{u}_{it} - x_t^{iter} &\geq \theta, && i \in [M],\ t \in [T]:\ \underline{y}_{it} < y_{it}^{iter} \ \wedge\ \underline{d}_i < \sum_{\tau=1}^{T} y_{i\tau}^{iter} && (47)\\
\mathbf{u}_{it} - x_t^{iter} &\leq -\theta, && i \in [M],\ t \in [T]:\ \bar{y}_{it} > y_{it}^{iter} \ \wedge\ \bar{d}_i > \sum_{\tau=1}^{T} y_{i\tau}^{iter}. && (48)
\end{aligned}$$

The objective (44) maximizes the slack θ. By constraint (45), the characteristic utility u must belong to the extended uncertainty set U_δ. Constraints (46)-(48) correspond to the conditions of Proposition 5.2. Since all variables are continuous, and U_δ is a polytope, this model is an LP. The optimal solution defines the characteristic utility u^{iter} that is added to the discrete uncertainty set U^D.

Proposition 5.3. *There exists a universal constant $\theta_{\min} > 0$ that depends only on δ such that the optimum value of (44)-(48) is always at least $\theta_{\min}$.*

Proof. Let δ_{min} be the minimum distance between a boundary point of U_δ and a boundary point of U. Since $\delta > 0$ by definition, $\delta_{\min} > 0$.

Let $u \in U$ the optimal uncertainty vector obtained when computing y^{iter} in Step 3 of the algorithm. We define u' as follows:

$$u'_{it} = \begin{cases} u_{it} + \delta_{min} & \text{if } \underline{y}_{it} < y_{it}^{iter} \ \wedge\ \underline{d}_i < \sum_{\tau=1}^{T} y_{i\tau}^{iter} \\ u_{it} - \delta_{min} & \text{if } \bar{y}_{it} > y_{it}^{iter} \ \wedge\ \bar{d}_i > \sum_{\tau=1}^{T} y_{i\tau}^{iter} \\ u_{it} & \text{otherwise.} \end{cases}$$

Since y^{iter} is a vertex of Y, the first two cases in the definition of u' cannot occur simultaneously by Observation 5.1, so u' is well-defined. Moreover, $u' \in U_\delta$ and satisfies the constraints (46)-(48) for $\theta = \delta_{\min}$. Therefore, the optimum value of (44)-(48) is always at least $\theta_{\min} := \delta_{\min}$. □

The optimal solution u^{iter} of (44)-(48) has the following property.

Proposition 5.4. *For x^{iter}, u^{iter} and θ as defined above, assume x' is a tariff such that $||x' - x^{iter}||_{\max} \leq \theta/2$. Then $y^{iter} \in \Omega(x', u^{iter})$. Moreover, if $\|x' - x\|_{\max} < \theta/2$, then y^{iter} is the unique optimal response of the followers for (x', u^{iter}).*

Proof. Observe that $x_t^{iter} - \theta/2 \leq x'_t \leq x_t^{iter} + \theta/2$ for all $t \in [T]$. Then (46) implies that $u_{it}^{iter} - x'_t \geq u_{it'}^{iter} - x'_{t'}$ for all $t \neq t' \in [T]$ such that $y_{it}^{iter} > \underline{y}_{it}$ (not at lower bound), and $y_{it'}^{iter} < \bar{y}_{it'}$ (not at upper bound). Moreover, by (47) and (48), $u_{it}^{iter} - x'_t \geq \theta/2$ if y_{it}^{iter} may be decreased, and $u_{it}^{iter} - x'_t \leq -\theta/2$ if y_{it}^{iter} may be increased. By Proposition 5.2, y^{iter} is an optimal follower response for (x', u^{iter}). □

Finally, observe that the two LPs (39)-(43) and (44)-(48) share a common structure, and the only difference between them is the weight of the different neighbors $y^\ell \in N(y^{iter})$ in the definition of the characteristic radii Δ and θ. In Method I, inequalities (41)-(43) are weighted by the difference of the corresponding follower objective values, whereas Method II applies uniform weights in constraints (46)-(48). Yet, the underlying idea can be applied with arbitrary positive weights, and the efficiency of different weight vectors should be investigated in computational experiments. Moreover, it is straightforward to combine different methods by adding more than one characteristic utility in one iteration.

6. Computational Evaluation

6.1. Implementation details

The proposed algorithms were implemented in C++, using the Gurobi 9.5 commercial MILP solver. All logical constraints, including (26)-(29) and (33)-(36) were implemented as indicator constraints. The default algorithms of Gurobi were used, yet, with custom parameter settings for improving numerical stability (see Section 6.4 for details). Also, aggressive cut generation was applied (Cuts=3). All experiments were run with a time limit of 600 seconds, on a personal computer with Intel i7 1.80 GHz CPU and 16 GB RAM.

6.2. Problem Instances

Computational evaluation was carried out on a set of randomly generated problem instances of the demand response management problem. Two different types of instances were investigated: in so-called *independent follower* (IF) instances, the uncertainty sets related to individual followers were independent, i.e., each inequality defining the uncertainty set U contained nonzero coefficients for one follower only; in contrast, in *dependent follower* (DF) instances, the uncertainty sets were interrelated. Instance sizes were varied by selecting both the number of followers, M, and the number of time period, T, from $\{5, 10, 15\}$. Five random instances were generated for each combination of M and T, resulting in 90 instances altogether.

Wholesale prices p_t were drawn from $\mathcal{U}[1, 500]$, where $\mathcal{U}[a, b]$ denotes the discrete uniform distribution over the integers in $[a, b]$. Set X is defined by lower and upper bounds on each individual x_t (two values were drawn from $\mathcal{U}[0, 1000]$, and then the lower (higher) value was used as the lower bound $\underline{x}_t$ (upper bound $\bar{x}_t$) on x_t), and a linear inequality of the form $\sum_{t\in[T]} r_t x_t \leq r_0$, where $r_t,\ t \in [T]$ was taken from $\mathcal{U}[0, 10]$ and r_0 from $\mathcal{U}[\sum_{t\in[T]} r_t \underline{x}_t,\ \sum_{t\in[T]} r_t \bar{x}_t]$.

For generating bounds on the load in individual periods, $\underline{y}_{it}$ and $\bar{y}_{it}$, two random values were drawn from $\mathcal{U}[0, 1000]$, and then the lower (higher) value was used as the lower bound $\underline{y}_{it}$ (upper bound $\bar{y}_{it}$). For bounding the total load of consumer i, two random values were drawn from $\mathcal{U}[\sum_t \underline{y}_{it}, \sum_t \bar{y}_{it}]$, and again, they were used as lower bound $\underline{d}_i$ and upper bound $\bar{d}_i$.

Similarly, random values were generated from $\mathcal{U}[0, 1000]$, and then the lower (higher) value is used as the lower bound $\underline{u}_{it}$ (upper bound $\bar{u}_{it}$) on u_{it}. Additional constraints on U were generated depending on the type of the instance. Dependent follower (DF) instances contained a single linear inequality of the form $\sum_{i\in[M]} \sum_{t\in[T]} v_{it} u_{it} \leq v_0$, where v_{it} was taken from $\mathcal{U}[0, 10]$ and v_0 from $\mathcal{U}[\sum_{i\in[M]} \sum_{t\in[T]} v_{it} \underline{u}_{it},\ \sum_i \sum_t v_{it} \bar{u}_{it}]$. In contrast, independent follower (IF) instances included one separate constraint for each individual follower, i.e., $\forall i : \sum_{t\in[T]} v^i_{it} u_{i,t} \leq v^i_0$, where v^i_{it} was taken from $\mathcal{U}[0, 10]$ and v^i_0 from $\mathcal{U}[\sum_{i\in[M]} \sum_{t\in[T]} v^i_{it} \underline{u}_{it},\ \sum_{i\in[M]} \sum_{t\in[T]} v^i_{it} \bar{u}_{it}]$. The problem instances and the corresponding results are publicly available in the GitHub repository of the project[1].

The robust bilevel problems generated using the above method contain T leader variables to encode the tariff and MT follower variables to capture the consumption, i.e., $15 + 225$ variables for the largest instances with $M = 15$ and $T = 15$. Within the algorithm, the computational challenge lies in solving the discrete-uncertainty variant (23)-(30). When uncertainty set U^D consists of K discrete values, the corresponding single-level MILP contains $3MTK + 2MK + T + 1$ continuous and $2MTK + 2MK$ binary variables, connected by $5MTK + 6MK + K + 1$ constraints. For an instance with $(M, T) = (15, 15)$ in the 10th iteration of the algorithm, this corresponds to 7066 continuous and 4800 binary variables with 12161 constraints.

6.3. Experimental Results

Two versions of the proposed algorithm were investigated and compared in the computational experiments: *Alg-I* computed the characteristic utilities in Step 4 using Method I (refer to Section 5.3.4.1); in contrast, *Alg-II* used Method II (described in Section 5.3.4.2). All other components of the two algorithms were identical. Moreover, to analyze the impact of the tolerance parameter δ used for defining the extended uncertainty set U_δ, both algorithms were run with $\delta = 10^{-2}$ and $\delta = 10^{-3}$, resulting in four different runs on each problem instance.

The results are presented in Tables 3 and 4 for the two algorithms, respectively. Each row in the tables contains aggregated results for a fixed value of parameter δ, the 5 instances of a given type (DF or IF) and a given problem size, as indicated in columns M and T. Column *Terminated* shows the number of instances out of 5 where the solver terminated according to the stopping condition $f_{\text{best}} \geq f_{\text{bound}}$ without hitting the time limit. Then, subsequent columns display the average computation time, the average number of iterations, as well as the average and maximum normalized gaps. Normalized gaps are computed as $\frac{\text{UB}-f_{\text{best}}}{|\text{UB}|+1}$, where f_{best} is the objective value of the best solution found and UB is the upper bound computed according to the technique presented in Section 4.4.

As it can be seen from the tables, algorithm *Alg-II* clearly outperformed *Alg-I*: it terminated on more instances (117 vs. 151 terminations out of 180 runs) while taking less computation time (237 s vs. 112 s on average) and less iterations (14.1 vs. 7.1), as well as achieving lower average gaps (0.34% vs. 0.29%). This difference in performance appears with both investigated values of parameter δ. One possible explanation of this phenomenon is that the uniform weighting of the neighbors when calculating the characteristic radius in Method II leads to more robust performance, whereas the presence of neighbors with very low weights in Method I may lead to slow convergence.

For the more efficient *Alg-II*, the average computation time was 112 s over all instances, also taking into account the instances where the 600 s time limit was hit, whereas it was 18.8 s for the instances where search terminated. This corresponds to 7.1 iterations on average, with 5.7 iterations in case of termination and 14.3 iterations in case of timeout. Timeout occurred in 29 runs on large instances (12 with $\delta = 10^{-2}$ and 17 with $\delta = 10^{-3}$), including 70% of the runs with $(M, T) = (15, 15)$.

The choice of parameter δ clearly impacts the behavior of the algorithm. With a greater δ, the algorithm converges to the supremum in greater steps, i.e., it terminates quicker, but often further away from the supremum. That is, with $\delta = 10^{-2}$, the algorithm terminated in 139 cases with an average gap of 0.41% over all instances.

[1]https://github.com/akovacs-sztaki/Robust-bilevel-optimization

Table 3. Experimental results with algorithm *Alg-I*.

δ	Prob	M	T	Terminated	Avg. time (s)	Avg. iter	Avg. gap	Max. gap
10^{-2}	DF	5	5	5	5.8	8.6	0.04%	0.20%
			10	5	25.0	8.6	0.59%	1.71%
			15	4	185.3	18.6	0.36%	1.33%
		10	5	4	120.5	13.0	0.06%	0.12%
			10	1	502.3	36.2	0.62%	1.30%
			15	2	375.8	19.0	0.35%	1.13%
		15	5	3	243.0	10.2	0.09%	0.16%
			10	4	164.1	14.4	0.30%	1.23%
			15	0	600.0	13.4	0.64%	1.66%
	IF	5	5	5	33.0	9.6	0.25%	0.58%
			10	5	8.9	7.6	0.12%	0.34%
			15	4	134.8	9.0	0.25%	0.58%
		10	5	5	82.9	10.6	0.87%	3.97%
			10	2	385.8	19.8	0.57%	1.03%
			15	3	323.2	16.4	0.81%	1.90%
		15	5	5	24.5	13.0	0.19%	0.60%
			10	1	541.0	20.2	0.77%	1.69%
			15	2	419.7	10.4	0.63%	1.79%
10^{-3}	DF	5	5	5	3.6	8.8	0.03%	0.17%
			10	4	120.5	13.6	0.06%	0.17%
			15	4	214.8	21.2	0.04%	0.13%
		10	5	5	43.7	10.0	0.01%	0.02%
			10	2	365.2	28.4	0.28%	0.95%
			15	2	371.8	19.4	0.20%	0.89%
		15	5	3	246.5	11.4	0.03%	0.06%
			10	4	143.5	14.4	0.19%	0.93%
			15	0	600.0	12.8	0.73%	2.43%
	IF	5	5	5	2.3	5.2	0.02%	0.06%
			10	5	66.9	9.4	0.02%	0.08%
			15	4	226.6	13.6	0.44%	2.11%
		10	5	5	2.9	4.8	0.05%	0.19%
			10	2	368.2	20.0	0.34%	0.70%
			15	1	480.3	16.6	0.78%	2.06%
		15	5	5	3.0	9.2	0.01%	0.06%
			10	0	600.0	21.8	0.51%	1.39%
			15	1	507.5	10.2	0.90%	3.70%

Table 4. Experimental results with algorithm *Alg-II*.

δ	Prob	M	T	Terminated	Avg. time (s)	Avg. iter	Avg. gap	Max. gap
10^{-2}	DF	5	5	5	0.8	3.2	0.07%	0.33%
			10	5	1.6	4.6	0.59%	1.71%
			15	5	1.1	4.6	0.35%	1.33%
		10	5	5	0.8	4.0	0.05%	0.12%
			10	4	124.9	11.2	0.58%	1.22%
			15	4	185.4	9.6	0.31%	1.03%
		15	5	5	8.4	5.2	0.08%	0.18%
			10	5	5.9	5.4	0.31%	1.23%
			15	1	483.5	8.6	0.62%	1.49%
	IF	5	5	5	0.5	2.8	0.24%	0.58%
			10	5	6.8	5.4	0.15%	0.50%
			15	5	22.6	6.8	0.22%	0.48%
		10	5	5	0.3	2.2	0.87%	3.97%
			10	4	147.3	10.2	0.44%	1.09%
			15	4	153.9	8.8	0.83%	1.89%
		15	5	5	2.7	4.2	0.20%	0.60%
			10	3	267.9	11.6	0.58%	1.03%
			15	3	259.7	7.0	0.65%	2.21%
10^{-3}	DF	5	5	5	0.5	3.8	0.03%	0.17%
			10	5	8.5	6.4	0.06%	0.17%
			15	5	8.8	6.2	0.03%	0.13%
		10	5	5	1.2	3.6	0.01%	0.01%
			10	4	137.8	16.2	0.27%	0.95%
			15	3	259.8	12.6	0.20%	0.92%
		15	5	5	15.3	5.4	0.01%	0.02%
			10	5	4.9	5.0	0.19%	0.93%
			15	1	485.3	8.2	0.52%	1.61%
	IF	5	5	5	0.3	2.6	0.02%	0.06%
			10	5	13.4	6.6	0.02%	0.09%
			15	5	30.1	7.8	0.03%	0.07%
		10	5	5	0.3	2.2	0.05%	0.19%
			10	3	242.9	11.4	0.20%	0.41%
			15	3	270.5	10.6	0.45%	1.31%
		15	5	5	1.7	3.8	0.01%	0.06%
			10	3	389.6	15.8	0.38%	1.01%
			15	1	502.1	11.0	0.95%	3.74%

In contrast, with $\delta = 10^{-3}$, it terminated in only 131 cases (8 cases less) with an average gap of 0.22% (0.19% lower). Notably, the average gap over the terminated runs decreased from 0.33% with $\delta = 10^{-2}$ to 0.05% with $\delta = 10^{-3}$, which provides an experimental illustration of Proposition 3.8 that a sufficiently small δ leads to ε-optimality with tighter ε. At the same time, there is no significant difference between the gaps when the algorithm hits the time limit: 0.66% with $\delta = 10^{-2}$ vs. 0.67% with $\delta = 10^{-3}$.

The above gaps stem from three different sources:

(i) stopping the algorithm with a sub-optimal solution due to the time limit;
(ii) the error of the solution due to the extension of the uncertainty set U_δ; and
(iii) the error of the upper bound due to the difference of the pessimistic robust solution and the optimistic upper bound of Section 4.4.

It is rather difficult to separate the three sources of error from each other. Yet, the considerable difference of the gaps in case of termination and timeout (0.05% vs. 0.67% with $\delta = 10^{-3}$) suggests that component (i) is the most significant in case of time-out, whereas the same component is obviously not present in case of termination. By Proposition 3.8, error component (ii) converges to zero as δ is decreased, which is illustrated by the reduction of the total gap from 0.33% with $\delta = 10^{-2}$ to 0.05% with $\delta = 10^{-3}$ over the terminated runs. A special characteristic of error component (iii) is that it persists even if δ is decreased and the algorithm terminates. For a DF instance with $M = 15$ and $T = 10$, which resulted in one of the greatest gaps among the terminated runs (1.23% with $\delta = 10^{-2}$ and 0.93% with $\delta = 10^{-3}$), by observing the difference of the optimistic and the pessimistic follower responses, we could prove a stronger bound that shows that the solutions found are within a 0.005% environment of the supremum. This suggests that component (iii), i.e., the error of the bound is responsible for the few significant gaps even with $\delta = 10^{-3}$ after termination.

DF instances were slightly easier for both algorithms than IF instances: the average gap was 0.25% for DF vs. 0.38% for IF.

As one would expect, almost the entire computation time was taken by solving the large MILPs encoding the discrete-uncertainty variant in Step 2, whereas the compact MILP of Step 3 for determining the worst-case response and the LP of Step 4 for computing the characteristic utility could be solved quickly. Moreover, the computation time for Step 2 increased rapidly with the size of the discrete uncertainty set U^D.

Even when the algorithm hit the time limit, it constructed high-quality solutions: the largest gap encountered by *Alg-II* over all instances was 3.74%. Taking a closer look at the individual iterations, it could be observed that search improved both solutions and bounds compared to the first iteration. Yet, solutions were improved ca. twice as much as bounds, which suggests that developing a better initial solution heuristic is also an intriguing direction for future research.

6.4. Numerical stability

The numerical stability of the computed solutions is a serious issue, which follows directly from the problem definition: the supremum is located in regions of the solution space where certain follower responses, unfavorable for the leader, are “just not optimal”. In order to overcome such numerical issues, the tolerance parameters of the Gurobi MILP solver were chosen to be as strict as possible: primal feasibility tolerance (*FeasibilityTol*), integer feasibility tolerance (*IntFeasTol*), and dual feasibility tolerance (*OptimalityTol*) were all set to 10^{-9}, whereas the Big-M value for feasibility relaxations (*FeasRelaxBigM*) was set to 10^9.

In order to verify the solutions computed by the MILP solver, the feasibility of all solutions were systematically verified. Regarding optimality, the critical step is the solution of the discrete-uncertainty variant in the final iteration. These MILP problems were exported to an LP file and solved using IBM ILOG CPLEX Optimization Studio version 12.6.3 as well, and the solutions from the two MILP solvers were compared. Among the instances solved to optimality by both solvers, the largest relative difference between the reported "optimal" solutions was $6.8 \cdot 10^{-5}$, stemming from slightly different fractional tariff values x_t but identical followers' responses y_{it}^k. For two further instances, the relative difference was $3.8 \cdot 10^{-5}$, yet, with different followers' responses y_{it}^k for a small subset of the indices. For all other instances, the difference was strictly less than 10^{-5}. Observe that these values are considerably worse than what is suggested by the optimality tolerance parameter applied in the MILP solvers. Overall, we believe that such a numerical precision is acceptable in practical applications, especially because the approach addresses finding ε-optimal solutions, but a note should be taken that numerical precision is limited by the MILP solver as well.

7. Computational Complexity

7.1. RBOP is Σ_2^p-Complete

Containment of RBOP in Σ_2^p is trivial, since this class includes all problems that can be expressed in the form $\exists x \forall y P(x, y)$, where P is a boolean predicate on the variables x and y. Then, Σ_2^p-hardness of RBOP is shown by proving the same property for the specific application, the robust bilevel demand response management problem.

Proposition 7.1. *The decision version of the robust bilevel demand response management problem is Σ_2^p-complete.*

Proof. The claim is demonstrated by reduction from the bilevel knapsack problem with interdiction constraints (BKIC), also known as the DeNegre variant of the bilevel knapsack problem, which is known to be Σ_2^p-complete [32,33]. In an instance of BKIC, the two decision makers, the leader and the follower, load a set of R items into their private knapsacks. First, the leader picks some items that fit into its own knapsack ($\sum_{t=1}^R a_t x_t \leq A$), then the follower packs a part of the remaining items ($\sum_{t=1}^R b_t y_t \leq B$, $y_t \leq 1 - x_t \ \forall t$). The objective of the follower is maximizing the total value of the knapsack ($\max \sum_{t=1}^R b_t y_t$), whereas the objective of the hostile leader is to minimize this value ($\min \sum_{t=1}^R b_t y_t$). This definition of BKIC from [32] assumes that for each item, the follower's weight and value, denoted by b_t, match each other. Yet, all our claims apply to the more generic variant where this assumption is omitted, too. In the decision version of the problem, the question asked is whether the leader can pick items in such a way that the follower's knapsack value is at most V.

A corresponding instance of the robust bilevel demand response management problem with one follower and $3R$ time periods is constructed as follows. The crux of the construction is reducing the discrete knapsack problem into a continuous optimization problem. In fact, the first R time periods map the two problems to each other, whereas the remaining $2R$ periods ensure that the optimum of the discrete and the continuous problems coincide. Since the problem involves a single follower, the index of the followers is omitted in the entire proof.

An overview of the constructed energy management problem is presented in Table 5,

Table 5. Parameters of the robust energy management problem, received by reduction from BKIC.

Periods	Core $t = 1, ..., R$	Integrality 1 $t = R+1, ..., 2R$	Integrality 2 $t = 2R+1, ..., 3R$
X	$x_t \in [0, 1]$	$x_t = x_{t-R}$	$x_t = 1 - x_{t-2R}$
U	$u_t \in [-1, 0]$	$u_t = 1 - \varepsilon$	$u_t = 1 - \varepsilon$
Y	$y_t \in [0, 1]$	$y_t \in [0, 1]$	$y_t \in [0, 1]$
p	$p_t = b_t$	$p_t = E$	$p_t = E$

with the definition of the ranges for individual x_t, u_t, and y_t variables, as well as the specification of parameters p_t. Constants used include the big number $E = \sum_{t=1}^{R} b_t + 1$ and the small positive number $\varepsilon = \frac{1}{2RE}$. The feasible set of tariff values, X, is defined by the displayed box constraints and the leader's knapsack constraint $\sum_{t=1}^{R} a_t x_t \leq A$. Likewise, the uncertainty set U is given by the ranges in the table and the follower's knapsack constraint $\sum_{t=1}^{R} b_t(u_t + 1) \leq B$. While $0 \leq y_t \leq 1$ is required for each individual period, there are no tight bounds on $\sum_{t=1}^{R} y_t$. The decision version asks if the leader can achieve a profit of at least $-V - RE$.

First, let us observe that the leader is strongly encouraged to select binary values for x_t, i.e., $x_t = 0$ or $x_t = 1$. The selection of x_t in interval $t \in [R]$ determines the tariff, and hence, indirectly the solution in all other time periods as well. If the leader sets $x_t < \varepsilon$ for some $t \in [R]$, then $u_{R+t} - x_{R+t} > 0$ and $u_{2R+t} - x_{2R+t} < 0$, and the follower's unique response is $y_{R+t} = 1$ and $y_{2R+t} = 0$ in time periods $R + t$, and $2R + t$, respectively. On the other hand, if $x_t > 1 - \varepsilon$, then $u_{R+t} - x_{R+t} < 0$ and $u_{2R+t} - x_{2R+t} > 0$, and the unique response is $y_{R+t} = 0$ and $y_{2R+t} = 1$. However, if $\varepsilon \leq x_t \leq 1-\varepsilon$, then both $u_{R+t}-x_{R+t} \geq 0$ and $u_{2R+t}-x_{2R+t} \geq 0$ and in the worst-case response for the leader, $y_{R+t} = 1$ and $y_{2R+t} = 1$. Hence, if $x_t < \varepsilon$ or $x_t > 1 - \varepsilon$ for all $t \in [R]$, then the leader's profit over the *Integrality* periods is exactly $-RE$, whereas otherwise it is at most $-(R+1)E$. The difference between the profits in the two cases is at least E, which is strictly larger than the profit that can be achieved in the *Core* periods $t = 1, ..., R$, which implies that the leader maximizes its profit by selecting $x_t < \varepsilon$ or $x_t > 1 - \varepsilon$ for all $t \in [R]$.

For the sake of simplicity, it can be assumed that the leader rounds all values $x_t < \varepsilon$ down to $x_t = 0$, and $x_t > 1-\varepsilon$ up to $x_t = 1$. This does not alter the follower's response in the *Core* periods because the knapsack value is modified by strictly less than 1, and does not alter the follower response in the *Integrality* periods either, because the sign of $u_t - x_t$ is preserved. Furthermore, rounding modifies the leader's profit by less than 1, and accordingly, the response to the question asked in the decision variant also remains the same.

Now, assume that the original BKIC instance is a YES-instance, i.e., the leader can select a set of items S such that the follower achieves a profit of at most V. Then, in the demand response management problem, the leader sets $x_t = 1$ for $t \in S$ and $x_t = 0$ for $t \in [R] \setminus S$. For periods $t \in S$, having $x_t = 1$ and $u_t \in [-1, 0]$ implies that $u_t - x_t \leq -1$, and the follower's unique response is $y_t = 0$ in these periods. Hence, these periods do not contribute to the leader's profit. For $t \in [R] \setminus S$, $x_t = 0$, and the follower may set $y_t = 1$ only if $u_t - x_t \geq 0$, that is, the uncertain parameter u_t is 0. This incurs at least $-V$ profit for the leader. Applying binary tariff values also implies that exactly R units of load will be scheduled in the *Integrality* periods, which results in a profit of exactly $-RE$ for the leader. Hence, the total profit over the entire horizon is at least $-V - RE$, and accordingly, the instance of the energy management problem is also a YES-instance.

Conversely, assume that the demand response management problem instance is a YES-instance with profit at least $-V - RE$. Then, by the above, for all $t \in [R]$, the leader must select $x_t = 0$ or $x_t = 1$, and the profit over the *Integrality* periods is exactly $-RE$. This solution corresponds to a feasible solution for the bilevel knapsack problem as well, where $x_t = 1$ encodes that the leader selects item t. The fact that the leader collects a profit of at least $-V$ over the *Core* periods implies that follower's knapsack (set of items with $u_t = 0$) has a value of at most V, and the bilevel knapsack problem is also a YES-instance. □

An additional consequence of the above proposition is that, since in general no polynomial-size MILP formulation exists for a Σ_2^p-hard problem, the number of iterations of the proposed algorithm is super-polynomial in the worst case. Also, these findings are in line with those of Buchheim at al. [11], who showed that the complexity of robust linear bilevel problems with uncertainty in the follower's objective depends on the structure of the uncertainty set: with interval uncertainty, the linear bilevel problem may become Σ_2^p-hard, whereas with discrete uncertainty, the robust problem remains in NP.

7.2. The Infinitely Robust Variant is Solvable in Polynomial Time

In the infinitely robust variant of RBOP, for any $x \in X$ and $y \in \hat{Y}$, there exists a $u \in U$ such that y is an optimal follower response for (x, u), that is, $\Omega(x, u) = \{y\}$. Therefore, the leader must choose x in such a way that it maximizes the minimum profit which can be achieved over all the possible follower responses. Hence, the infinitely robust bilevel optimization problem can be formulated as follows:

$$\max z \tag{49}$$
$$\text{s.t}$$
$$(c + x)^T y \geq z, \quad \forall y \in \hat{Y} \tag{50}$$
$$x \in X. \tag{51}$$

Since X is a polyhedron, this is a linear program. Note that a polytope may have an exponential number of vertices in the number of facets (consider, e.g., the n-dimensional cube in $\mathbb{R}^n$), the number of inequalities in (50) may be exponential in the size of Y specified by a system of inequalities. Suppose only a subset of the constraints in (50) is included in the linear program, and let (x', z') be an optimal solution of the relaxation. We can decide if (x', z') satisfies all the constraints in (50) by solving the following optimization problem parametrized by x':

$$\min\{(c + x')^T y \ : \ y \in Y\}. \tag{52}$$

Denote the optimal vertex solution of (52) by $y(x')$. If the optimum value of (52) is less than z', then the inequality in (50) corresponding to $y(x')$ is violated by (x', z'). Therefore, we add the inequality $(c + x)y(x') \geq z$ to the LP relaxation. Since (52) is a linear program, we can solve the *separation problem* for (50) in polynomial time by any polynomial time algorithm for linear programming [34,35]. Consequently, by the equivalence of separation and optimization [36], we can solve (49)-(51) in polynomial time by the ellipsoid method of Khachiyan [34]. Thus, we have proved the following result:

Theorem 7.2. *The infinitely robust variant of RBOP can be solved in polynomial time.*

It is easy to see that if $X \subseteq U$, then the RBOP is infinitely robust.

8. Conclusions

This paper introduced a robust bilevel optimization framework for tackling bilevel problems with polyhedral uncertainty in the coefficients of the followers' objective function. The framework assumes that both the leader's and the followers' constraints are linear, whereas the followers' objective is bilinear. This robust bilevel problem is, in general, Σ_2^p-hard, i.e., it is located one level higher in the polynomial hierarchy than its deterministic counterpart.

An efficient solution method was proposed, based on the idea of iteratively building a discrete set of so-called characteristic utilities to map the relevant areas of the original continuous, polyhedral set. It was formally proved that the algorithm terminates in finitely many steps with an ε-optimal solution for any given $\varepsilon > 0$ for a problem that, in general, does not admit a maximum but only a finite supremum.

The approach was illustrated on a demand response management problem in smart grids with uncertainty in the consumers' utilities. The proposed algorithm solved instances of relevant sizes, with up to 15 followers and 15 time units, to proven ε-optimality, whereas it found high-quality robust solutions and proved strong bounds for the remaining open instances as well. Notably, this is significantly larger than the instances typically investigated in earlier contributions in the field of robust bilevel optimization.

There are a number of intriguing directions for future research. A key to the efficiency of the proposed method was the relatively low number of discrete utility values that could properly represent the original polyhedral uncertainty set. It is a natural question whether this observation generalizes to other applications. Regarding the improvement of the algorithm, the adaptive adjustment of parameter δ can ensure further performance guarantees, namely, achieving ε-optimality for an a priori given $\varepsilon > 0$. Also, the problem size tractable by the exact method may fall short of that desirable in practical applications, and hence, the development of effective heuristics is essential. Finally, it will be interesting to generalize the approach to broader classes of problems, e.g., with polyhedral uncertainty in the followers' constraint coefficients or integer variables for the leader.

Acknowledgments

This research was supported by the TKP2021-NKTA-01 NRDIO grant on "Research on cooperative production and logistics systems to support a competitive and sustainable economy" and by the European Union within the framework of the National Laboratory for Autonomous Systems (RRF-2.3.1-21-2022-00002).

References

[1] Beck Y, Ljubić I, Schmidt M. A survey on bilevel optimization under uncertainty. European Journal of Operational Research. 2023;311(2):401–426.

[2] Wang J, Bai X, Liu Y. Globalized robust bilevel optimization model for hazmat transport network design considering reliability. Reliability Engineering & System Safety. 2023;239.
[3] Nikoofal ME, Zhuang J. Robust allocation of a defensive budget considering an attacker's private information. Risk Analysis. 2012;32(5):930–943.
[4] Zeng B, Dong H, Sioshansi R, et al. Bilevel robust optimization of electric vehicle charging stations with distributed energy resources. IEEE Transactions on Industry Applications. 2020;56(5):5836–5847.
[5] Chuong TD, Yu X, Eberhard A, et al. Convergences for robust bilevel polynomial programmes with applications. Optimization Methods & Software. 2023;38(5):975–1008.
[6] Dempe S. Foundations of bilevel programming. Springer Science & Business Media; 2002.
[7] Schrijver A. Theory of linear and integer programming. John Wiley & Sons; 1998.
[8] Garey MR, Johnson DS. Computers and intractability: A guide to the theory of NP-completeness. W. H. Freeman and Company; 1979.
[9] Chuong TD, Jeyakumar V. Finding robust global optimal values of bilevel polynomial programs with uncertain linear constraints. Journal of Optimization Theory and Applications. 2017;173(2):683–703.
[10] Buchheim C, Henke D. The robust bilevel continuous knapsack problem with uncertain coefficients in the follower's objective. Journal of Global Optimization. 2022;83:803–824.
[11] Buchheim C, Henke D, Hommelsheim F. On the complexity of robust bilevel optimization with uncertain follower's objective. Operations Research Letters. 2021;49(5):703–707.
[12] Beck Y, Ljubić I, Schmidt M. Exact methods for discrete Γ-robust interdiction problems with an application to the bilevel knapsack problem. Mathematical Programming Computation. 2023;15(4):733–782.
[13] Dempe S, Mordukhovich BS, Zemkoho AB. Necessary optimality conditions in pessimistic bilevel programming. Optimization. 2014;63(4):505–533.
[14] Dempe S, Mordukhovich BS, Zemkoho AB. Two-level value function approach to non-smooth optimistic and pessimistic bilevel programs. Optimization. 2019;68(2-3):433–455.
[15] Zeng B. A practical scheme to compute the pessimistic bilevel optimization problem. INFORMS Journal on Computing. 2020;32(4):1128–1142.
[16] Wiesemann W, Tsoukalas A, Kleniati PM, et al. Pessimistic bilevel optimization. SIAM Journal on Optimization. 2013;23(1):353–380.
[17] Zheng Y, Wan Z, Jia S, et al. A new method for strong-weak linear bilevel programming problem. Journal of Industrial and Management Optimization. 2015;11(2):529–547.
[18] Caramia M. An algorithm to find stable solutions in linear-linear bilevel problems. OPSEARCH. 2024;61(2):972–988.
[19] Besançon M, Anjos MF, Brotcorne L. Near-optimal robust bilevel optimization; 2019. Unpublished manuscript, https://hal.science/hal-02414848v2.
[20] Besançon M, Anjos MF, Brotcorne L. Complexity of near-optimal robust versions of multilevel optimization problems. Optimization Letters. 2021;15(8):2597–2610.
[21] Beck Y, Bienstock D, Schmidt M, et al. On a computationally ill-behaved bilevel problem with a continuous and nonconvex lower level. Journal of Optimization Theory and Applications. 2023;198(1):428–447.
[22] Zare MH, Prokopyev OA, Sauré D. On bilevel optimization with inexact follower. Decision Analysis. 2020;17:74–95.
[23] Beck Y, Schmidt M. A robust approach for modeling limited observability in bilevel optimization. Operations Research Letters. 2021;49(5):752–758.
[24] Haghighat H. Strategic offering under uncertainty in power markets. International Journal of Electrical Power & Energy Systems. 2014;63:1070–1077.
[25] Chauhan D, Unnikrishnan A, Boyles SD, et al. Robust maximum flow network interdiction considering uncertainties in arc capacity and resource consumption. Annals of Operations Research. 2024;335(2):689–725.
[26] Punla-Green SZ, Mitchell JE, Gearhart JL, et al. Shortest path network interdiction with asymmetric uncertainty. Networks. 2024;83(3):605–623.
[27] Sun L, Karwan MH, Kwon C. Robust hazmat network design problems considering risk

uncertainty. Transportation Science. 2015;50(4):1188–1203.
[28] Xin C, Qingge L, Wang J, et al. Robust optimization for the hazardous materials transportation network design problem. Journal of Combinatorial Optimization. 2015; 30(2):320–334.
[29] Smith JC, Song Y. A survey of network interdiction models and algorithms. European Journal of Operational Research. 2020;283(3):797–811.
[30] Kleinert T, Labbé M, Ljubić I, et al. A survey on mixed-integer programming techniques in bilevel optimization. EURO Journal on Computational Optimization. 2021;9:100007.
[31] Kovács A. On the computational complexity of tariff optimization for demand response management. IEEE Transactions on Power Systems. 2018;33(3):3204–3206.
[32] Caprara A, Carvalho M, Lodi A, et al. A complexity and approximability study of the bilevel knapsack problem. In: Integer Programming and Combinatorial Optimization; (LNTCS; Vol. 7801). Springer Berlin Heidelberg; 2013. p. 98–109.
[33] DeNegre S. Interdiction and discrete bilevel linear programming [dissertation]. Lehigh University; 2011.
[34] Khachiyan LG. Polynomial algorithms in linear programming. USSR Computational Mathematics and Mathematical Physics. 1980;20(1):53–72.
[35] Karmarkar N. A new polynomial-time algorithm for linear programming. In: Proceedings of the sixteenth annual ACM symposium on Theory of computing; 1984. p. 302–311.
[36] Grötschel M, Lovász L, Schrijver A. Geometric algorithms and combinatorial optimization, 2nd edition. Springer Science & Business Media; 2012.

Appendix A. Proof of Proposition 3.6

Proof. Suppose $(x^\star, y^\star, u^\star)$ is an $\frac{\varepsilon}{2}$-optimal robust bilevel feasible solution. If $x^\star$ is in the relative interior of X, we are done. Otherwise, $x^\star$ is on the boundary of X. By Corollary 3.3, we may assume $y^\star \in \hat{Y}$.

We call $y' \in Y$ *bad* for some $x' \in X$ if $x' \in X^{opt}(y')$ and $z^\star - f(x', y') > \varepsilon|z^\star| + \varepsilon$. We partition $\hat{Y}$ into two subsets: $\hat{Y}^{bad}$ containing all the vertices of Y which are bad for $x^\star$, and the rest $\hat{Y}^{good} := \hat{Y} \setminus \hat{Y}^{bad}$. Clearly, $y^\star \in \hat{Y}^{good}$. Moreover, for each $y^b \in \hat{Y}^{bad}$ and $u \in U$, there exists $y^g \in \hat{Y}^{good}$ such that

$$(u - x^\star)^T y^g > (u - x^\star)^T y^b,$$

otherwise y^b would be an optimal follower answer for $(x^\star, u)$, and then $z^\star - f(x^\star, y^b) > \varepsilon|z^\star| + \varepsilon$. This implies $f(x^\star, y^\star) > f(x^\star, y^b)$, hence, $(x^\star, y^\star, u^\star)$ is not a robust bilevel feasible solution, a contradiction. We argue there exists $\Delta > 0$ such that for each $y^b \in \hat{Y}^{bad}$ and $u \in U$, there is $y^g \in \hat{Y}^{good}$ such that

$$(u - x^\star)^T y^g \geq (u - x^\star)^T y^b + \Delta.$$

To determine Δ, for each $y^b \in \hat{Y}^{bad}$, consider the optimization problem

$$value(y^b) := \min_{u \in U} \max_{y^g \in \hat{Y}^{good}} (u - x^\star)^T y^g - uy^b. \tag{A1}$$

Inside the minimum, we have a piecewise-linear, convex, affine function of u. So the minimum is attained at some point in U. Since y^b is bad for $x^\star$, the minimum value in (A1) is strictly greater than $-(x^\star)^T y^b$. Therefore, $\Delta(y^b) := value(y^b) + (x^\star)^T y^b > 0$. By the choice of $\Delta(y^b)$, for any $u \in U$, there is $y^g \in \hat{Y}^{good}$ such that $(u - x^\star)^T y^g \geq$

$(u-x^\star)^T y^b + \Delta(y^b)$. Define $\Delta := \min\{\frac{\varepsilon}{2}|z^\star|, \min_{y^b \in \hat{Y}^{bad}} \Delta(y^b)\}$. Then, we can choose a non-zero vector v such that $x' := x^\star + v$ is in the relative interior of X, and $|v^T y^g| + |v^T y^b| < \Delta$ for any pair of $(y^g, y^b) \in \hat{Y}^{good} \times \hat{Y}^{bad}$. It follows that any $y^g \in \hat{Y}^{good}$ is good for x', namely,

$$(c+x')^T y^g = (c+x^\star)^T y^g + v^T y^g \geq z^\star - \frac{\varepsilon}{2}|z^\star| - \Delta - \varepsilon > z^\star - \varepsilon|z^\star| - \varepsilon.$$

Moreover, for any $u \in U$, and $y^b \in \hat{Y}^{bad}$, we can find $y^g \in \hat{Y}^{good}$ such that

$$(u-x')^T y^g > (u-x')^T y^b,$$

hence, y^b cannot be the follower's optimal answer for (x', u). Therefore, for any u, there exists $y^g \in \hat{Y}^{good}$ such that (x', y^g, u) is an ε-optimal robust bilevel solution. □