Mixed-Integer Linear Reformulations of Pessimistic Bilinear Bilevel Infima

Victor Bucarey, Martine Labbé, Gonzalo Muñoz, Martin Schmidt

Abstract. We study pessimistic bilevel optimization with bilinear objectives and linear constraints. Assuming a compact polyhedral leader domain, nonempty follower solution sets, and a finite pessimistic infimum, we derive mixed-integer linear optimization formulations that attain this infimum, even when the original problem has no optimal solution. This setting has applications such as pricing, security, and decision-focused learning. The construction combines LP duality, support selection, and homogeneous witnesses, together with a KKT reformulation for leader-dependent follower constraints. Optimal reformulation solutions yield explicit leader decisions with pessimistic objective values arbitrarily close to the infimum. For rational data, all formulations have polynomial encoding length, also showing that the problem is NP-complete.

1. Introduction

Bilevel optimization problems are hierarchical ones: they consist of an optimization problem, called the upper-level or leader problem, that contains in its constraints the requirement that some variables solve another optimization problem, called the lower-level or follower problem (Dempe et al. 2015; Kleinert et al. 2021; Beck et al. 2027). The leader's decisions enter as parameters in the follower's problem and may affect both its objective function and its constraints.

In the presence of several optimal solutions to the lower-level problem, the way in which a solution is chosen leads to two classic approaches. In the optimistic approach, the solution chosen is the one that favors the upper-level optimization problem. In the pessimistic approach, the leader considers the worst upper-level objective value over all optimal solutions of the lower-level problem (Wiesemann et al. 2013; Liu et al. 2018). In this work, we consider the pessimistic bilevel bilinear optimization problem

$$V := \inf_{x \in X} \sup_{y \in S(x)} x^\top A y + a^\top x + b^\top y, \tag{1}$$

where $X \subseteq \mathbb{R}^{n_x}$ and

$$S(x) := \arg\min\{x^\top B y + c^\top y \,:\, Dy \geq d - Cx\}.$$

The problem data is

$$A, B \in \mathbb{R}^{n_x \times n_y}, \quad C \in \mathbb{R}^{m_y \times n_x}, \quad D \in \mathbb{R}^{m_y \times n_y}, \quad a \in \mathbb{R}^{n_x}, \quad b, c \in \mathbb{R}^{n_y}, \quad d \in \mathbb{R}^{m_y}.$$

We refer to the x and y variables as the leader and follower decisions, respectively. Both objective functions are linear in y for fixed x, whereas their joint dependence on (x, y) can be bilinear.

Problem (1) presents two general challenges from the computational point of view. First, it is NP-hard, since it contains pessimistic linear bilevel optimization as

Date: October 25, 2026.

2020 *Mathematics Subject Classification.* 90C11, 91A65.

Key words and phrases. Pessimistic bilevel optimization, Single-level reformulations, Mixed-integer linear optimization, Attainability.

a special case (Buchheim 2023). Second, even in cases when V is finite, the infimum may not be attained; see, e.g., Lampariello et al. (2019).

The problem structure studied here arises in many applications, such as in decision-focused learning, Stackelberg security games, and in pricing problems. Decision-focused learning trains predictive models to improve downstream decision quality (Elmachtoub and Grigas 2022; Mandi et al. 2024). Pessimistic formulations evaluate the worst realized cost among decisions that are optimal under the predictions (Bucarey et al. 2026; Jiménez et al. 2026). With linear predictors and polyhedral downstream problems, these formulations exhibit the structure considered here. In Stackelberg security games, a defender commits to a randomized protection strategy, after which an attacker chooses a best response (Sinha et al. 2018; Casorrán et al. 2019; Bustamante-Faúndez et al. 2024). Expected payoffs are bilinear in the players' mixed strategies, and pessimistic tie-breaking selects the attacker's best response that is the least favorable to the defender. A similar structure arises in network toll setting (Brotcorne et al. 2001) and deterministic electricity-contract pricing (Jacquet et al. 2024), where a seller sets prices and customers subsequently minimize travel costs or maximize utility. The follower's problem is linear over a price-independent polyhedral feasible region, while the seller's objective contains bilinear price-quantity terms. Under bounded polyhedral price constraints, the pessimistic counterparts of these models are special cases of our setting with $C = 0$. In both security games and pricing problems, optimistic tie-breaking is often justified by arguing that small perturbations of the leader's decision can induce a favorable follower response. Such arguments, however, do not by themselves provide an explicit computational procedure for constructing the required strategies or prices and controlling the resulting loss in the leader's objective. In this article, we propose a constructive approach to this issue, recovering leader decisions whose pessimistic objective values are arbitrarily close to the optimal value, even when this value is not attained.

To state our results, we make the following standing assumption.

Assumption 1.1. The set X is a nonempty polytope, $S(x) \neq \varnothing$ for every $x \in X$, and $V \in \mathbb{R}$ is finite.

Under this assumption, the follower value $\min_{y \in Y(x)} x^\top By + c^\top y$, where $Y(x) := \{y \in \mathbb{R}^{n_y} : Dy \geq d - Cx\}$, is finite and attained for every $x \in X$. The follower feasible sets may be unbounded. Moreover, the pessimistic value of a leader decision,

$$P(x) := \sup_{y \in S(x)} x^\top Ay + a^\top x + b^\top y,$$

may equal $+\infty$ for some $x \in X$.

The main contribution of this short note is the derivation of mixed-integer linear programming (MILP) formulations whose optimal value equals V and is attained, even when the original bilevel problem has no attainable optimal solution. These formulations also allow us to recover leader decisions with pessimistic objective values arbitrarily close to V. More specifically, given an optimal solution of the MILP reformulations, we can construct an explicit infimizing sequence for (1). Finally, the mixed-integer linear formulations are of polynomial size, which, in particular, shows that the decision version of (1) is NP-complete.

1.1. **Related Work.** For a general treatment of pessimistic bilevel optimization, we refer to Wiesemann et al. (2013). The relation to robustness over a decision-dependent set of follower responses is discussed by Goerigk et al. (2025). Reformulations based on LP duality are given by Dempe, Luo and Franke (2018) and Zheng et al. (2018). Moreover, Zeng (2020) shows how to solve the pessimistic problem via solving a suitably chosen optimistic one plus a potentially required single-level problem. In the bilinear setting considered here, directly dualizing the intermediate

maximization problem in this level-reduction approach can introduce cubic terms. The dual stationarity equations permit an equivalent quadratic rewriting; see also Stein and Zemkoho (2026, Section 4.3). Our constructions use LP duality and support selection to obtain MILP formulations that attain the pessimistic infimum and provide an explicit procedure for recovering leader decisions with pessimistic objective values arbitrarily close to it. That optimal solutions need not exist even for compact leader and follower domains is shown by Lampariello et al. (2019).

The work of Stein and Zemkoho (2026) gives a continuous nonlinear single-level reformulation for (1). Their reformulation retains non-attained infima, i.e., if the original problem does not attain its optimal value, neither does the reformulation; see their Example 2 in Section 4.4. Our MILP formulations always attain the pessimistic infimum, provided it is finite, and yield explicit leader decisions with pessimistic objective values arbitrarily close to it.

For $C = 0$, we derive a MILP formulation directly. For general C, we use the KKT reformulation of nonconvex quadratic problems studied by Xia, Vera and Zuluaga (2020) to again obtain a MILP single-level reformulation. The finite bounds required for the MILP formulations are obtained using the basic-solution argument of Buchheim (2023).

1.2. **Additional Notation.** For the ease of notation, we define

$$h(x) := d - Cx, \quad q(x) := B^\top x + c, \quad u(x) := A^\top x + b.$$

For a given leader decision, the follower minimizes $q(x)^\top y$ over $Y(x) = \{y \in \mathbb{R}^{n_y} : Dy \geq h(x)\}$. The leader's objective function is $F(x, y) := a^\top x + u(x)^\top y$ and the pessimistic problem (1) can thus be rewritten as

$$V = \inf_{x \in X} P(x), \quad P(x) = \sup_{y \in S(x)} F(x, y). \tag{2}$$

2. A Single-Level MILP Formulation for the Case of $C = 0$

We first consider the case $C = 0$, in which the follower feasible set is independent of the leader's decision. We denote $Y(x)$ by Y and use the inequality description $X = \{x : Hx \leq k\}$ with $H \in \mathbb{R}^{m_x \times n_x}$ and $k \in \mathbb{R}^{m_x}$.

Let us first describe the main components of our MILP construction. To this end, fix a leader decision x and consider a follower's dual optimal solution λ. Let $I = \{i : \lambda_i > 0\}$. By complementary slackness and strong duality, the set of all optimal follower responses is the face

$$Y_I := \{y : Dy \geq d,\ D_I y = d_I\},$$

where D_I and d_I are the rows indexed by I. Maximizing the leader's objective over this face gives a linear problem whose dual is

$$\min\{d^\top w : D^\top w = u(x),\ w_{I^c} \leq 0,\ w_I \text{ free}\},$$

where I^c is the complement of I. Our formulation involves linear equations and inequalities in the variables (x, y, λ, w) that impose feasibility and optimality. Additionally, we introduce binary variables $s_1, \ldots, s_{m_y}$ which are used to determine I and impose tightness in certain inequalities based on I. Finally, and critically, we introduce a homogeneous copy of all variables $\tilde{x}, \tilde{y}, \tilde{\lambda}, \tilde{w}$ as well as a scaling variable t. These copies and the scaling variable are what allow us to construct arbitrarily near-optimal leader decisions if V is unattainable.

The formulation is given by

$$\min \quad a^\top x + d^\top w \tag{3a}$$

$$\text{s.t.} \quad Hx \leq k, \quad Dy \geq d, \tag{3b}$$

$$D^\top \lambda = B^\top x + c, \quad \lambda \geq 0, \tag{3c}$$
$$D^\top w = A^\top x + b, \tag{3d}$$
$$t \geq 1, \quad H\tilde{x} \leq tk, \quad D\tilde{y} \geq td, \tag{3e}$$
$$D^\top \tilde{\lambda} = B^\top \tilde{x} + tc, \quad \tilde{\lambda} \geq 0, \tag{3f}$$
$$D^\top \tilde{w} = A^\top \tilde{x} + tb, \tag{3g}$$
$$s_i = 0 \implies \begin{cases} \lambda_i = \tilde{\lambda}_i = 0, \\ w_i, \tilde{w}_i \leq 0, \end{cases} \qquad i = 1, \dots, m_y, \tag{3h}$$
$$s_i = 1 \implies \begin{cases} D_i y = d_i, \\ D_i \tilde{y} = t d_i, \\ \tilde{\lambda}_i \geq 1, \end{cases} \qquad i = 1, \dots, m_y, \tag{3i}$$
$$s \in \{0,1\}^{m_y}. \tag{3j}$$

All continuous expressions are linear because t only multiplies fixed data. We use indicator constraints for the support conditions; their big-M linearization is briefly discussed in Section 2.1.

The objective depends on the main variables. We call the variable copies a *strict witness* because, after division by t, its selected dual multipliers are strictly positive. This scaling encodes strict positivity exactly, without the need of a numerical tolerance on the multipliers.

We now formally state the relation between the MILP and the pessimistic problem.

Theorem 2.1. *Suppose that the standing assumptions hold. Then, Problem (3) has an optimal solution and its optimal objective function value equals the pessimistic infimum V.*

Proof. We first show that every feasible objective of (3) is at least V. To this end, for each feasible solution of (3), set

$$U = a^\top x + d^\top w, \quad I = \{i : s_i = 1\}.$$

We normalize the homogeneous witness by setting

$$(\hat{x}, \hat{y}, \hat{\lambda}, \hat{w}) = \frac{1}{t}(\tilde{x}, \tilde{y}, \tilde{\lambda}, \tilde{w}), \quad \hat{U} = a^\top \hat{x} + d^\top \hat{w}. \tag{4}$$

Now consider an arbitrarily chosen $0 < \tau \leq 1$ and consider the convex combination

$$(x_\tau, y_\tau, \lambda_\tau, w_\tau) = (1-\tau)(x, y, \lambda, w) + \tau(\hat{x}, \hat{y}, \hat{\lambda}, \hat{w}). \tag{5}$$

Since the constraints are affine for fixed s, we obtain

$$x_\tau \in X, \quad Dy_\tau \geq d, \quad D_I y_\tau = d_I,$$
$$D^\top \lambda_\tau = q(x_\tau), \quad \lambda_\tau \geq 0, \quad (\lambda_\tau)_{I^c} = 0, \quad (\lambda_\tau)_i \geq \tau/t > 0 \quad (i \in I),$$
$$D^\top w_\tau = u(x_\tau), \quad (w_\tau)_{I^c} \leq 0.$$

With the leader's decision x_τ fixed, the primal-dual gap of any follower decision $y \in Y$ reads

$$q(x_\tau)^\top y - d^\top \lambda_\tau = \sum_{i \in I} (\lambda_\tau)_i (D_i y - d_i) \geq 0. \tag{6}$$

Equality is attained at y_τ. Moreover, since $(\lambda_\tau)_i > 0$ for $i \in I$, equality holds if and only if $D_I y = d_I$. Hence,

$$S(x_\tau) = \{y : Dy \geq d,\ D_I y = d_I\}. \tag{7}$$

Thus, every follower-optimal solution $y \in S(x_\tau)$ satisfies

$$u(x_\tau)^\top y = w_\tau^\top D y = \sum_{i \in I} (w_\tau)_i d_i + \sum_{i \notin I} (w_\tau)_i D_i y \leq d^\top w_\tau,$$

where the last inequality follows from $(w_\tau)_{I^c} \leq 0$. Taking the supremum over all optimal follower responses $y \in S(x_\tau)$ yields

$$V \leq P(x_\tau) \leq a^\top x_\tau + d^\top w_\tau = (1-\tau)U + \tau \hat{U}. \tag{8}$$

Taking the limit $\tau \searrow 0$ proves $V \leq U$.

Conversely, we show that the infimum of (3) is at most V. Let $x \in X$ with $P(x) < +\infty$ be given. The follower's primal and dual problems are

$$\min\{q(x)^\top y : Dy \geq d\}, \quad \max\{d^\top \lambda : D^\top \lambda = q(x),\ \lambda \geq 0\}.$$

We choose an optimal primal-dual pair (y, λ) and set

$$I = \{i : \lambda_i > 0\}, \quad s_i = \mathbf{1}_{\{i \in I\}}.$$

Complementary slackness implies $D_I y = d_I$. The calculation in (6), with this primal-dual pair, shows that

$$S(x) = \{y : Dy \geq d,\ D_I y = d_I\}.$$

Since $P(x)$ is finite, the following primal and dual problems have finite optimal values and admit optimal solutions:

$$\max\{u(x)^\top y : Dy \geq d,\ D_I y = d_I\},$$
$$\min\{d^\top w : D^\top w = u(x),\ w_{I^c} \leq 0,\ w_I \text{ free}\}.$$

Let w be a dual optimal solution. By strong duality,

$$a^\top x + d^\top w = P(x). \tag{9}$$

For $I \neq \varnothing$, we choose the finite scalar

$$t = \max\left\{1, \max_{i \in I} \frac{1}{\lambda_i}\right\}.$$

In particular, for $I = \varnothing$, we take $t = 1$. Finally, we set

$$(\tilde{x}, \tilde{y}, \tilde{\lambda}, \tilde{w}) = t(x, y, \lambda, w).$$

By construction, these variables are feasible for Problem (3), and their objective value is $P(x)$ by (9). Hence, the value of Problem (3) is at most $P(x)$ for every finite-cost x. Taking the infimum over x proves the missing inequality of the values.

It remains to show that (3) admits an optimal solution. For fixed s, Problem (3) is an LP. Every feasible such LP is bounded below by V, as shown in the first part of the proof, and thus has an optimal solution. The second part and finiteness of V ensure that at least one binary assignment is feasible. Since there are only finitely many assignments, the best of their optimal values is attained and equals V. □

Remark 2.2. The equality of optimal values does not imply that an optimal main decision solves the pessimistic problem. At a feasible main point, $U = a^\top x + d^\top w$ need not be an upper bound on $P(x)$: selected main multipliers may vanish, so additional follower responses become optimal at x. The pessimistic value can then exceed the cost of the selected face. By the strict witness, this face is the entire follower optimal set at every interpolated x_τ with $\tau > 0$. Thus, (3) attains a limiting value even if the pessimistic problem has no optimal solution.

2.1. **An MILP Reformulation using Big-Ms.** To obtain an ordinary MILP, let M be a sufficiently large constant. Then, the indicator constraints in (3h)–(3i) can be re-written as

$$\begin{aligned} 0 \le \lambda_i \le M s_i, \qquad & s_i \le \tilde{\lambda}_i \le M s_i, \\ w_i \le M s_i, \qquad & \tilde{w}_i \le M s_i, \\ 0 \le D_i y - d_i \le M(1 - s_i), \quad & 0 \le D_i \tilde{y} - t d_i \le M(1 - s_i), \end{aligned} \tag{10}$$

for $i = 1, \dots, m_y$; see, e.g., Bonami et al. (2015) for an overview of optimization problems including indicator constraints. The existence of a finite M preserving an optimal solution follows from Theorem 2.1. For rational data, Buchheim (2023) derives polynomially sized big-M constants that can also be computed in polynomial time, which can also be used in this setting. However, these big-Ms are typically too large to be used in practice and Kleinert et al. (2020) showed for optimistic linear bilevel optimization that computing the tightest big-M value is NP hard. Hence, one typically relies on problem-specific big-M derivations in practice or uses SOS1-type constraints as discussed in, e.g., Kleinert and Schmidt (2023).

2.2. **Recovery of Pessimistic Decisions.** The interpolation in the proof of Theorem 2.1 also gives the following recovery result.

Corollary 2.3. *Let a feasible solution of (3) with leader component x and objective value U be given. Then, for every $0 < \tau \le 1$,*

$$x_\tau = (1 - \tau)x + \tau \hat{x} \in X, \quad P(x_\tau) \le U + \tau(\hat{U} - U). \tag{11}$$

Furthermore, for any target tolerance $\eta > 0$, choose

$$\tau = \begin{cases} 1, & \hat{U} \le U, \\ \min\{1, \eta/(\hat{U} - U)\}, & \hat{U} > U. \end{cases} \tag{12}$$

Then, $P(x_\tau) \le U + \eta$ holds.

Proof. The first claim follows from (8). The choice in (12) ensures $\tau(\hat{U} - U) \le \eta$ and proves the second claim. □

2.3. **Example: A Pessimistic Stackelberg Security Game.** To illustrate the reformulation and the recovery of leader decisions, we consider a simple Stackelberg security game. These are bimatrix games having bilinear objective functions and the players do not affect the feasible region of the other player, i.e., $C = 0$.

In this example, we consider a game in which a defender protects one of two targets with some probability and an attacker chooses a target after observing the defender's mixed strategy. Table 1 gives the payoffs, with the defender's reward first and the attacker's reward second.

TABLE 1. Payoffs of the Stackelberg security game example

	Attack 1	Attack 2
Protect 1	$(10, -10)$	$(-5, 6)$
Protect 2	$(-8, 4)$	$(6, -4)$

We denote by $x \in [0, 1]$ the probability of protecting target 1 and by $y \in [0, 1]$ the probability of attacking it. The attacker's pure-attack rewards are $4 - 14x$ and $10x - 4$, while the defender's are $18x - 8$ and $6 - 11x$. To write this game as a minimization problem, we take the leader's cost to be 6 minus the defender's reward. Omitting a term in the follower objective that is independent of y yields

$$q(x) = 24x - 8, \quad F(x, y) = 11x + (14 - 29x)y, \quad Y = [0, 1].$$

Allowing mixed strategies for the attacker preserves the pessimistic value because, at a tie, the worst affine defender reward is attained at a pure attack. The attacker chooses $y = 1$ for $x < 1/3$, $y = 0$ for $x > 1/3$, and is indifferent at $x = 1/3$. Hence,

$$P(x) = \begin{cases} 14 - 18x, & x < 1/3, \\ 8, & x = 1/3, \\ 11x, & x > 1/3. \end{cases} \tag{13}$$

The infimum $V = 11/3$ is not attained. Equivalently, the defender's pessimistic reward has supremum $6 - V = 7/3$, approached from $x > 1/3$. At the defender strategy $x = 1/3$, where the attacker is indifferent, the defender's actual reward is -2.

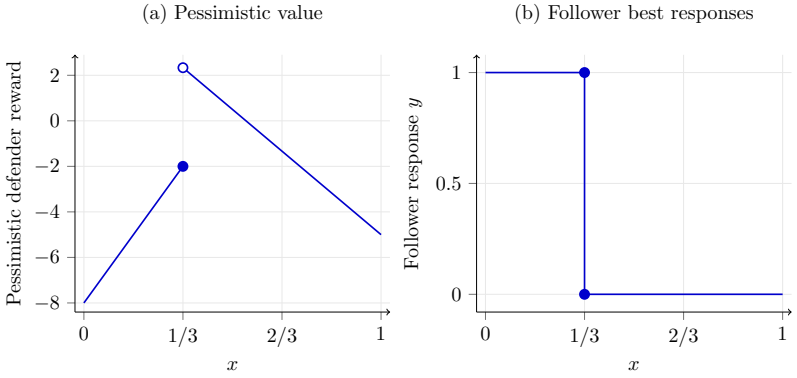


FIGURE 1. Pessimistic defender reward and follower best responses in the security game example. Left: The open circle marks the unattained limiting reward and the solid circle marks the actual reward at indifference. Right: The vertical segment is the full best-response set $S(1/3) = [0, 1]$.

Using reformulation (3), we have that $A = -29$, $a = 11$, $b = 14$, $H = (1, -1)^\top$, $k = (1, 0)^\top$, $B = 24$, $c = -8$, $D = (1, -1)^\top$, and $d = (0, -1)^\top$. The optimal support variables are $(s_1, s_2) = (1, 0)$, an optimal main point and a normalized strict witness are given by

$$\begin{aligned} (x, y, \lambda, w) &= (1/3, 0, (0, 0), (13/3, 0)), \\ (\hat{x}, \hat{y}, \hat{\lambda}, \hat{w}) &= (1/2, 0, (4, 0), (-1/2, 0)), \\ t &= 1. \end{aligned}$$

Corollary 2.3 gives

$$x_\tau = x + (\hat{x} - x)\tau = \tfrac{1}{3} + \tfrac{1}{6}\tau, \quad P(x_\tau) = \tfrac{11}{3} + \tfrac{11}{6}\tau, \quad 0 < \tau \leq 1.$$

These decisions uniquely induce attack 2 and yield defender reward $6 - P(x_\tau) = 7/3 - 11\tau/6$, which converges to $7/3$ for $\tau \searrow 0$. Hence, the finite strict witness allows us to recover pessimistic leader decisions approaching the supremum from the optimal MILP solution.

2.4. **Example: Decision-Focused Learning.** Decision-focused learning trains predictors to improve the decisions obtained from their predictions. Let us consider linear predictors. Let $\bar{q}^{(i)}$ be the realized cost vector of observation $i = 1, \ldots, N_{\text{obs}}$, and let $q^{(i)}(x) = (B^{(i)})^\top x$ be its predicted cost, with the observed features incorporated into $B^{(i)}$ and predictor parameters $x \in X$. For a fixed nonempty

polytope $\bar{Y} = \{y : \bar{D}y \geq \bar{d}\}$, pessimistic regret minimization is the pessimistic bilevel optimization problem

$$\min_{x \in X} \max_{y^{(1)},\ldots,y^{(N_{\text{obs}})}} \frac{1}{N_{\text{obs}}} \sum_{i=1}^{N_{\text{obs}}} (\bar{q}^{(i)})^\top y^{(i)} - \kappa_{\text{DFL}}$$
$$\text{s.t.} \quad y^{(i)} \in \arg\min_{y \in \bar{Y}} (q^{(i)}(x))^\top y, \quad i = 1, \ldots, N_{\text{obs}},$$

where $\kappa_{\text{DFL}} = N_{\text{obs}}^{-1} \sum_i \min_{y \in \bar{Y}} (\bar{q}^{(i)})^\top y$ is the so-called constant perfect-information baseline. The inner minimization selects prediction-optimal decisions, and the maximization evaluates their worst realized cost.

Stacking the decisions into y, this is our model with

$$B = [B^{(1)} \ \cdots \ B^{(N_{\text{obs}})}], \quad b = \frac{1}{N_{\text{obs}}} \operatorname{col}(\bar{q}^{(1)}, \ldots, \bar{q}^{(N_{\text{obs}})}),$$
$$D = \operatorname{diag}(\bar{D}, \ldots, \bar{D}), \qquad d = \operatorname{col}(\bar{d}, \ldots, \bar{d}),$$

and $A = 0$, $a = 0$, $C = 0$, $c = 0$. Hence, Model (3) yields an exact MILP for the optimal pessimistic regret with objective $d^\top w - \kappa_{\text{DFL}}$ and optimal value $V - \kappa_{\text{DFL}}$. This provides a mixed-integer linear alternative to the nonconvex and quadratically constrained quadratic formulation presented by Bucarey et al. (2026).

The connection also explains the role of Theorem 1 by Bucarey et al. (2026), which establishes membership of Simple-Regret in NP by constructing polynomial-size predictors that preserve the downstream optimal faces. Our binary support variables encode these face choices explicitly, while the strict witness certifies their realizability. For rational data, the resulting MILP has polynomial encoding length and provides an optimization counterpart to this certificate argument. NP-completeness follows by combining membership in NP with the known NP-hardness result by Elmachtoub and Grigas (2022).

3. A Single-Level MILP Formulation for the Case of $C \neq 0$

We now extend the construction to general C. Adapting the same formulation as before results in a quadratic objective; we prove that this MIQP reformulation attains the pessimistic infimum. In addition, we establish a recovery result for this formulation. Finally, we use KKT conditions to reformulate the MIQP as an MILP.

3.1. **An MIQP Formulation.**

We use the same variables introduced in the preceding section and define

$$g_C(x, w) = a^\top x + h(x)^\top w = a^\top x + d^\top w - x^\top C^\top w. \tag{14}$$

The formulation for general C reads

$$\begin{aligned}
\min\quad & g_C(x,w) && \text{(15a)}\\
\text{s.t.}\quad & Hx \le k, \quad Dy + Cx \ge d, && \text{(15b)}\\
& D^\top\lambda = B^\top x + c, \quad \lambda \ge 0, && \text{(15c)}\\
& D^\top w = A^\top x + b, && \text{(15d)}\\
& t \ge 1, \quad H\tilde{x} \le tk, \quad D\tilde{y} + C\tilde{x} \ge td, && \text{(15e)}\\
& D^\top\tilde{\lambda} = B^\top\tilde{x} + tc, \quad \tilde{\lambda} \ge 0, && \text{(15f)}\\
& D^\top\tilde{w} = A^\top\tilde{x} + tb, && \text{(15g)}\\
& s_i = 0 \implies \begin{cases} \lambda_i = \tilde{\lambda}_i = 0, \\ w_i, \tilde{w}_i \le 0, \end{cases} \qquad i = 1,\dots,m_y, && \text{(15h)}\\
& s_i = 1 \implies \begin{cases} D_i y + C_i x = d_i, \\ D_i\tilde{y} + C_i\tilde{x} = td_i, \\ \tilde{\lambda}_i \ge 1, \end{cases} \qquad i = 1,\dots,m_y, && \text{(15i)}\\
& s \in \{0,1\}^{m_y}. && \text{(15j)}
\end{aligned}$$

For fixed s, all constraints remain affine, whereas the objective is now quadratic. Hence, the problem is a nonconvex MIQP in indicator form. As before, the binary variables select a follower's face and the homogeneous copy provides a strict witness. The following theorem establishes equality of the optimal values and existence of an optimal MIQP solution.

Theorem 3.1. *Problem (15) has infimum V and admits an optimal solution with this value.*

Proof. We adapt the proof of Theorem 2.1. For a feasible solution of (15), set $I = \{i : s_i = 1\}$, normalize its witness by t, and interpolate as in (5). By the primal-dual gap argument in the proof of Theorem 2.1, now with right-hand side $h(x_\tau)$, for $0 < \tau \le 1$, we have

$$S(x_\tau) = \{y : Dy + Cx_\tau \ge d,\ D_I y + C_I x_\tau = d_I\}. \tag{16}$$

Using $D^\top w_\tau = u(x_\tau)$ and $(w_\tau)_{I^c} \le 0$ yields

$$V \le P(x_\tau) \le g_C(x_\tau, w_\tau). \tag{17}$$

Continuity of g_C then gives $V \le g_C(x,w)$ as $\tau \searrow 0$.

Conversely, fix $x \in X$ with finite $P(x)$ and an optimal follower primal-dual pair (y,λ). For $I = \{i : \lambda_i > 0\}$, the same face argument and LP duality give an optimal solution w of

$$\min\{h(x)^\top w : D^\top w = u(x),\ w_{I^c} \le 0,\ w_I \text{ free}\},$$

with $g_C(x,w) = P(x)$. Scaling (x,y,λ,w) as in the proof of Theorem 2.1 provides a feasible homogeneous witness for (15). Hence, its infimum equals V.

Finally, every feasible fixed-support problem is a QP over a polyhedron, bounded below by V. It attains its minimum by the Frank–Wolfe existence theorem (Frank and Wolfe 1956). Since V is finite, at least one support is feasible. There are only finitely many supports, so their best minimum is attained and equals V. □

3.2. **Recovery of Pessimistic Decisions.** The interpolation used in the proof again allows us to recover leader decisions for the pessimistic problem. We normalize a feasible homogeneous witness, set $U = g_C(x,w)$, and define

$$\Delta x = \hat{x} - x, \quad \Delta w = \hat{w} - w.$$

The interpolation remains feasible, and its objective value satisfies

$$\begin{aligned} g_C(x+\tau\Delta x, w+\tau\Delta w) &= U + \kappa_1\tau + \kappa_2\tau^2, \\ \kappa_1 &= a^\top \Delta x + (d - Cx)^\top \Delta w - (C\Delta x)^\top w, \\ \kappa_2 &= -(C\Delta x)^\top \Delta w. \end{aligned} \tag{18}$$

By (17),

$$P(x_\tau) \le U + \kappa_1\tau + \kappa_2\tau^2 \le U + \tau K, \quad K = |\kappa_1| + |\kappa_2|, \quad 0 < \tau \le 1. \tag{19}$$

Corollary 3.2. *Let a feasible MIQP solution with objective value U and a tolerance $\eta > 0$ be given. Choose $\tau = 1$ if $K = 0$ and otherwise choose*

$$\tau = \min\{1, \eta/K\}.$$

Then, $P(x_\tau) \le U + \eta$ holds. If the MIQP solution is optimal, i.e., $U = V$, the recovered decision is η-optimal.

Proof. For $0 < \tau \le 1$, combining (17) and (18) gives $P(x_\tau) \le U + \tau K$. The proposed choice makes $\tau K \le \eta$. □

We finally note that for $C = 0$, we have $\kappa_2 = 0$ and $\kappa_1 = \hat{U} - U$, so we obtain the affine bound (11).

3.3. **KKT Linearization of the MIQP.** We now reformulate the MIQP as an MILP with the same optimal value. To this end, we apply the KKT-based reformulation of nonconvex QPs studied by Xia, Vera, and Zuluaga (2020); see, in particular, Section 2.1 and their discussion of valid multiplier bounds. We apply this result to the continuous QP obtained by fixing the support variables.

3.3.1. *Linearizing the Quadratic Objective.* We collect all continuous variables of (15), including the homogeneous witness and t, in a vector ζ and write

$$g_C(x, w) = \phi(\zeta) = \tfrac{1}{2}\zeta^\top Q\zeta + \ell^\top\zeta, \quad Q = Q^\top.$$

The only nonzero blocks of Q are the x, w cross-blocks $Q_{x,w} = -C^\top$ and $Q_{w,x} = -C$. The only nonzero blocks of ℓ are $\ell_x = a$ and $\ell_w = d$.

For a fixed binary assignment of s, we can write the continuous problem as

$$\min\{\phi(\zeta) : G\zeta \le b_*,\ E\zeta = e_*\} \tag{20}$$

for suitable matrices G, E and right-hand side vectors b_*, e_*. The KKT conditions of this problem are

$$\begin{aligned} Q\zeta + \ell + G^\top\mu + E^\top\nu &= 0, \\ G\zeta \le b_*, \quad E\zeta &= e_*, \\ \mu &\ge 0, \\ \mu_j\big((b_*)_j - G_j\zeta\big) &= 0 \quad \text{for all } j. \end{aligned} \tag{21}$$

Here, the equality multipliers ν are unrestricted in sign. Due to all constraints being affine-linear, these conditions are necessary without any further constraint qualification. Consequently, every global minimizer admits KKT multipliers.

Multiplying the stationarity condition by $\zeta^\top$ and using primal feasibility and complementarity, we obtain

$$\zeta^\top Q\zeta + \ell^\top\zeta + b_*^\top\mu + e_*^\top\nu = 0.$$

At every KKT point, we therefore have

$$\phi(\zeta) = \tfrac{1}{2}\big(\ell^\top\zeta - b_*^\top\mu - e_*^\top\nu\big). \tag{22}$$

Thus, we can replace the quadratic objective by the right-hand side of (22) and impose (21) without changing the global minimum. Indeed, a global QP minimizer

admits KKT multipliers. Conversely, every KKT point is feasible for the QP, and the linear objective equals its quadratic objective value. The reformulation therefore requires global minimization over the KKT points since an arbitrary stationary point need not be globally optimal.

3.3.2. *Finite Primal and Dual Bounds.* To obtain a MILP formulation from the KKT system (21), we require the following bounds. First, bounds on the continuous variables of (15) to linearize the indicator constraints and to bound the KKT slacks. Second, bounds on the inequality KKT multipliers μ to linearize the complementarity conditions and the products between them and the binary variables. No bounds on the equality multipliers ν are needed, since they appear neither in complementarity conditions nor in products with binary variables.

Proposition 3.3. *Under the standing assumptions, there exist finite bounds on the continuous variables of* (15) *and on the KKT inequality multipliers such that imposing them preserves an optimal solution of value V. If all problem data and the inequality description of X are rational, such bounds have polynomial binary encoding length and can be computed in polynomial time from the data.*

Proof. For each feasible support assignment in (15), consider the QP KKT system (21). For each complementarity condition, we consider the two alternatives: either the multiplier or the corresponding primal slack is zero. Every choice of these alternatives defines a polyhedron in the KKT variables.

A globally optimal KKT point belongs to at least one of these polyhedra. Minimizing the linear objective over that polyhedron gives the same QP optimum. We can therefore select finite optimal representatives for the finitely many feasible support assignments. In particular, if we let $N = \dim \zeta$, there exists R such that $\zeta \in [-R, R]^N$ preserves the MIQP optimum V. We next provide an estimate for such an R.

For rational data, let L be the input encoding length and p_0, q_0 bound the inequality and equality row counts of any fixed-support QP. These counts are polynomial in L and obtained directly from (15). If we move to standard form, we need to split free variables and add slacks; this gives at most $r = N + 2p_0 + q_0$ equations per linear KKT piece. The coefficients are signed input entries, zeros, or ± 1. Multiplying by the product of all input denominators therefore gives a common integer bound $K \geq 1$ on the absolute coefficients and right-hand sides. This common bound satisfies $\log K = O(L)$ and is computable in polynomial time, without enumerating supports or complementarity choices. Each relevant LP has an optimal basic feasible solution after deleting dependent equations. Buchheim's basic-solution estimate (Buchheim 2023, Lemma 4) bounds its coordinates by $r!K^r$. Undoing the free-variable splits, we may thus take

$$R = r!K^r, \quad \log R = O\big(r \log(r+1) + r \log K\big).$$

Hence, R has polynomial encoding length and is computable in polynomial time. We remark that this bounds optimal representatives, rather than every stationary point.

Once the bounds on ζ are imposed, the support indicators can be linearized as in (10). In particular, for the slacks

$$D_i y + C_i x - d_i, \quad D_i \tilde{y} + C_i \tilde{x} - t d_i,$$

we can use the bounds

$$L_i^C = R(\|D_i\|_1 + \|C_i\|_1) + |d_i|, \quad \tilde{L}_i^C = R(\|D_i\|_1 + \|C_i\|_1 + |d_i|).$$

This yields a bounded MIQP of the form

$$\min \left\{ \tfrac{1}{2} \zeta^\top Q \zeta + \ell^\top \zeta : G\zeta \leq b_0 + Ts,\ E\zeta = e,\ s \in \{0,1\}^{m_y} \right\}. \tag{23}$$

Here all matrices and vectors are fixed. The box $|\zeta_j| \leq R$, the sign restrictions, and the lower bound $t \geq 1$ are included in the linear constraints. The matrix T represents the dependence of the right-hand side on the support variables s. Each conditional equality is enforced by primal feasibility and a conditional inequality.

We then apply the same basic-solution argument to the KKT system of the fixed-support QPs to obtain bounds on the inequality multipliers. The new box and indicator coefficients have polynomial encoding length. If the bounded QP has p inequality and q equality rows, its standard-form KKT pieces have at most $r_2 = N + 2p + q$ equations. As before, clearing denominators and bounding the support-dependent right-hand sides in (23) by $|(b_0)_j| + \sum_i |T_{ji}|$ gives a uniform integer coefficient and right-hand-side bound $K_2 \geq 1$, with $\log K_2$ being polynomial in L. Thus, setting $\mu_j \leq \overline{\mu}_j$ for $\overline{\mu}_j := r_2! K_2^{r_2}$ preserves an optimal KKT certificate. The $\overline{\mu}_j$ have polynomial encoding length and can be computed in polynomial time. □

3.3.3. *An MILP Formulation.* To reformulate (23), for each fixed s, we impose KKT conditions with respect to the continuous vector ζ. The binary variables are held fixed and do not enter the stationarity conditions as optimization variables. We define the primal slacks by

$$\rho = b_0 + Ts - G\zeta.$$

Let $\overline{\mu}_j \geq 0$ be finite bounds preserving an optimal KKT certificate, as given in Section 3.3.2. A valid upper bound on the jth slack is

$$\overline{\rho}_j = |(b_0)_j| + \sum_i |T_{ji}| + R\|G_j\|_1.$$

These bounds have polynomial encoding length and can be computed in polynomial time. To model complementarity, we introduce binary variables u_j. Moreover, we use continuous variables v_{ji} for every pair with $T_{ji} \neq 0$ to represent the products $\mu_j s_i$.

The resulting MILP reads

$$\min \quad \frac{1}{2}\left(\ell^\top \zeta - b_0^\top \mu - e^\top \nu - \sum_{j,i:T_{ji}\neq 0} T_{ji} v_{ji}\right) \tag{24a}$$

$$\text{s.t.} \quad \rho = b_0 + Ts - G\zeta, \quad E\zeta = e, \tag{24b}$$

$$Q\zeta + \ell + G^\top \mu + E^\top \nu = 0, \tag{24c}$$

$$0 \leq \mu_j \leq \overline{\mu}_j u_j, \quad 0 \leq \rho_j \leq \overline{\rho}_j (1 - u_j) \quad \text{for all } j, \tag{24d}$$

$$0 \leq v_{ji} \leq \overline{\mu}_j s_i, \quad v_{ji} \leq \mu_j \quad (T_{ji} \neq 0), \tag{24e}$$

$$v_{ji} \geq \mu_j - \overline{\mu}_j (1 - s_i) \quad (T_{ji} \neq 0), \tag{24f}$$

$$s \in \{0,1\}^{m_y}, \quad u_j \in \{0,1\} \quad \text{for all } j. \tag{24g}$$

The variables ν are unrestricted in sign, and all constraints and the objective are linear. Constraints (24e)–(24f) imply $v_{ji} = \mu_j s_i$. Hence, the objective in (24a) equals the right-hand side of (22) with $b_* = b_0 + Ts$ and $e_* = e$. For p inequality rows, at most pm_y product variables are introduced. Proposition 3.3 therefore ensures that this MILP can be constructed in polynomial time with polynomial encoding length.

Theorem 3.4. *Suppose that valid representative bounds are used as specified above. Then, the MILP (24) has an optimal solution with objective value V. Every feasible solution yields main and homogeneous witness variables for which the recovery bound (19) holds.*

Proof. Let a feasible MILP point be given. It satisfies the continuous KKT conditions of (23) for the selected s. By (22), its linear objective value is $\phi(\zeta) = g_C(x, w)$. The

primal variables are feasible for Problem (15), so their objective value is at least V by Theorem 3.1.

Conversely, the primal box preserves a global optimum of (15). Thus, Problem (23) has a global optimum with value V. For its selected s, this point is also globally optimal for the continuous QP and admits KKT multipliers. The multiplier bounds preserve such an optimal certificate. We choose u according to complementarity and set $v_{ji} = \mu_j s_i$. The resulting point is feasible for (24) with objective value V, proving equality and attainment. Finally, the main and witness variables are feasible for (15), so the recovery bound (19) applies. This proves the recovery claim. □

We finalize this subsection with a computational complexity consequence of our MILP reformulation. For the decision version of (1), the input consists of rational, binary-encoded problem data, including H, k, and a rational threshold θ. The question is whether $V \leq \theta$. This need not be equivalent to the existence of $x \in X$ with $P(x) \leq \theta$, since the pessimistic infimum may fail to be attained.

Corollary 3.5. *The decision problem $V \leq \theta$ for Problem (1) is NP-complete.*

Proof. By Theorem 3.4, $V \leq \theta$ holds if and only if (24) has a feasible point with objective at most θ. Fixing the binaries s, u and adding this threshold constraint leaves a rational LP of polynomial encoding length. If feasible, it has a rational solution of polynomial encoding length. The binaries and this solution form a polynomially verifiable certificate, proving NP membership. NP-hardness holds already for the pessimistic linear special case discussed in the introduction. □

3.4. **The Modified Dempe Example.** We consider a modified version of an example in Dempe et al. (2015) in which we allow both endpoints of the follower's feasible interval to depend on the leader's decision. Hence, we consider

$$X = [-1, 1], \quad S(x) = \operatorname*{arg\,min}_{x \leq y \leq x+1} \{-xy\}, \quad F(x, y) = 2xy - x + y.$$

In the notation of the general model, the data are

$$A = 2, \quad B = -1, \quad a = -1, \quad b = 1, \quad c = 0,$$
$$D = \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \quad d = \begin{pmatrix} 0 \\ -1 \end{pmatrix}, \quad C = \begin{bmatrix} -1 \\ 1 \end{bmatrix}.$$

Thus $Dy + Cx \geq d$ is equivalent to $x \leq y \leq x + 1$ and $C \neq 0$. The follower feasible region is nonempty and compact for every $x \in X$. The response map and pessimistic value are given by

$$S(x) = \begin{cases} \{x\}, & x < 0, \\ [0, 1], & x = 0, \\ \{x + 1\}, & x > 0, \end{cases} \quad P(x) = \begin{cases} 2x^2, & -1 \leq x < 0, \\ 1, & x = 0, \\ 2x^2 + 2x + 1, & 0 < x \leq 1. \end{cases} \tag{25}$$

Every pessimistic value is strictly positive, but $P(x) \to 0$ as $x \nearrow 0$. Consequently, $V = 0$ and the infimum is not attained. The global optimistic solution is $(x^o, y^o) = (0, 0)$ with objective value zero; see Figure 2.

For the support $(s_1, s_2) = (1, 0)$, Formulation (15) gives $y = x$, $\lambda_1 = -x \geq 0$, $w_2 \leq 0$, and $w_1 - w_2 = 2x + 1$. Hence, its objective reduces to

$$g_C(x, w) = -x + xw_1 - (1 + x)w_2 = 2x^2 - w_2 \geq 0.$$

The other nonempty feasible support gives the minimum $2x^2 + 2x + 1 \geq 1$ on $[0, 1]$, while the empty support gives value 1 at $x = 0$. Thus, an optimal main point and a

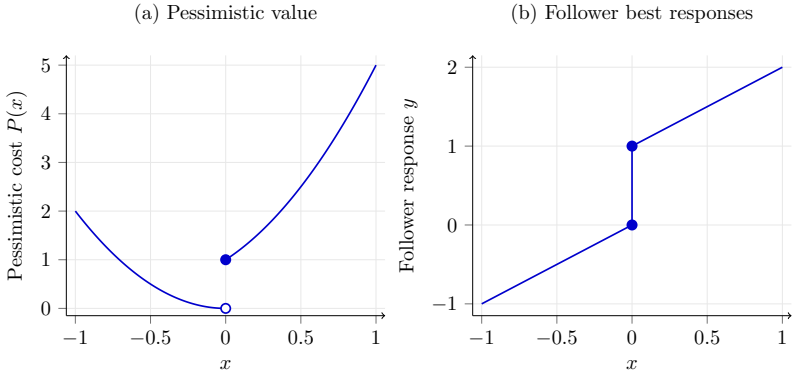


FIGURE 2. Pessimistic cost and follower best responses in the modified Dempe example. Left: The open circle marks the unattained limiting cost and the solid circle marks the actual cost at indifference. Right: The vertical segment is the full best-response set $S(0) = [0, 1]$.

normalized strict witness are

$$(x, y, \lambda, w) = (0, 0, (0, 0), (1, 0)),$$
$$(\hat{x}, \hat{y}, \hat{\lambda}, \hat{w}) = (-1, -1, (1, 0), (-1, 0)),$$
$$t = 1.$$

The main objective is 0, although $P(0) = 1$. This point also extends to an optimal certificate of the KKT MILP (24): use the primal box $R = 3$, the support bound $w_2 \leq 3s_2$, and inequality-multiplier bounds $\overline{\mu}_j = 1$. The objective gradient has only one nonzero component, -1 in the w_2 coordinate. Assigning multiplier 1 to this active support bound and zero to all other inequality and equality rows then verifies stationarity and complementarity.

Interpolating with the witness gives $x_\tau = y_\tau = -\tau$, $\lambda_\tau = (\tau, 0)$, and $w_\tau = (1 - 2\tau, 0)$ for $0 < \tau \leq 1$. The follower response is unique and

$$P(x_\tau) = g_C(x_\tau, w_\tau) = 2\tau^2.$$

Thus, $\kappa_1 = 0$ and $\kappa_2 = 2$ in (18), and the choice $\tau = \min\{1, \sqrt{\eta/2}\}$ yields an η-optimal pessimistic decision for every $\eta > 0$.

Declaration of AI Use. During the preparation of this work, the authors used ChatGPT to assist in developing and checking its content. The authors take full responsibility for the content of this article.

Acknowledgments. Víctor Bucarey and Gonzalo Muñoz were funded by ANID Fondecyt Regular 1261798.

REFERENCES

Beck, Y., I. Ljubić, and M. Schmidt (2027). *Linear and Mixed-Integer Bilevel Optimization. Theory and Algorithms*. Forthcoming. Cambridge University Press. URL: https://www.cambridge.org/core/books/linear-and-mixedinteger-bilevel-optimization/A8BC7630B9204739F1B3D8106798A07F.

Bonami, P., A. Lodi, A. Tramontani, and S. Wiese (2015). "On mathematical programming with indicator constraints." In: *Mathematical Programming* 151.1, pp. 191–223. DOI: 10.1007/s10107-015-0891-4.

Brotcorne, L., M. Labbé, P. Marcotte, and G. Savard (2001). "A bilevel model for toll optimization on a multicommodity transportation network." In: *Transportation Science* 35.4, pp. 345–358.

Bucarey, V., S. Calderón, G. Muñoz, and F. Semet (2026). "Decision-focused predictions via pessimistic bilevel optimization: complexity and algorithms." In: *Computational Optimization and Applications*, pp. 1–31.

Buchheim, C. (2023). "Bilevel linear optimization belongs to NP and admits polynomial-size KKT-based reformulations." In: *Operations Research Letters* 51.6, pp. 618–622. DOI: 10.1016/j.orl.2023.10.006.

Bustamante-Faúndez, P., V. Bucarey, M. Labbé, V. Marianov, and F. Ordoñez (2024). "Playing stackelberg security games in perfect formulations." In: *Omega* 126, p. 103068.

Casorrán, C., B. Fortz, M. Labbé, and F. Ordóñez (2019). "A study of general and security Stackelberg game formulations." In: *European Journal of Operational Research* 278.3, pp. 855–868.

Dempe, S., V. Kalashnikov, G. A. Pérez-Valdés, and N. Kalashnykova (2015). *Bilevel Programming Problems. Theory, Algorithms and Applications to Energy Networks*. Energy Systems. Berlin, Heidelberg: Springer. DOI: 10.1007/978-3-662-45827-3.

Dempe, S., G. Luo, and S. Franke (2018). "Pessimistic Bilevel Linear Optimization." In: *Journal of Nepal Mathematical Society* 1.1, pp. 1–10. URL: https://www.journal.nms.org.np/JNMS-Online-Articles/Vol_1.Issue_1.Year-2018/JNMS-2018-013_Dempe_Etal.pdf.

Elmachtoub, A. N. and P. Grigas (2022). "Smart "predict, then optimize"." In: *Management Science* 68.1, pp. 9–26.

Frank, M. and P. Wolfe (1956). "An algorithm for quadratic programming." In: *Naval Research Logistics Quarterly* 3.1–2, pp. 95–110. DOI: 10.1002/nav.3800030109.

Goerigk, M., J. Kurtz, M. Schmidt, and J. Thürauf (2025). "Connections between Robust and Bilevel Optimization." In: *Open Journal of Mathematical Optimization* 6, 2, pp. 1–17. DOI: 10.5802/ojmo.38.

Jacquet, Q., W. van Ackooij, C. Alasseur, and S. Gaubert (2024). "Quadratic regularization of bilevel pricing problems and application to electricity retail markets." In: *European Journal of Operational Research* 313.3, pp. 841–857.

Jiménez, D., B. Pagnoncelli, and H. Yaman (2026). "Pessimistic bilevel optimization approach for decision–focused learning." In: *Mathematical Programming Computation*, pp. 1–23.

Kleinert, T., M. Labbé, I. Ljubić, and M. Schmidt (2021). "A Survey on Mixed-Integer Programming Techniques in Bilevel Optimization." In: *EURO Journal on Computational Optimization* 9, 100007. DOI: 10.1016/j.ejco.2021.100007.

Kleinert, T., M. Labbé, F. Plein, and M. Schmidt (2020). "Technical Note—There's No Free Lunch: On the Hardness of Choosing a Correct Big-M in Bilevel Optimization." In: *Operations Research* 68.6, pp. 1716–1721. DOI: 10.1287/opre.2019.1944.

Kleinert, T. and M. Schmidt (2023). "Why there is no need to use a big-M in linear bilevel optimization: A computational study of two ready-to-use approaches." In: *Computational Management Science* 20.1, p. 3. DOI: 10.1007/s10287-023-00435-5.

Lampariello, L., S. Sagratella, and O. Stein (2019). "The Standard Pessimistic Bilevel Problem." In: *SIAM Journal on Optimization* 29.2, pp. 1634–1656. DOI: 10.1137/18M119759X.

Liu, J., Y. Fan, Z. Chen, and Y. Zheng (2018). "Pessimistic Bilevel Optimization: A Survey." In: *International Journal of Computational Intelligence Systems* 11.1, pp. 725–736. DOI: 10.2991/ijcis.11.1.56.

Mandi, J., J. Kotary, S. Berden, M. Mulamba, V. Bucarey, T. Guns, and F. Fioretto (2024). "Decision-focused learning: Foundations, state of the art, benchmark and future opportunities." In: *Journal of Artificial Intelligence Research* 80, pp. 1623–1701.

Sinha, A., F. Fang, B. An, C. Kiekintveld, and M. Tambe (2018). "Stackelberg security games: Looking beyond a decade of success." In: IJCAI.

Stein, O. and A. Zemkoho (2026). *A true single-level reformulation for pessimistic bilevel optimization*. Preprint. URL: https://arxiv.org/html/2609.38014v1.

Wiesemann, W., A. Tsoukalas, P.-M. Kleniati, and B. Rustem (2013). "Pessimistic Bilevel Optimization." In: *SIAM Journal on Optimization* 23.1, pp. 353–380. DOI: 10.1137/120864015.

Xia, W., J. C. Vera, and L. F. Zuluaga (2020). "Globally Solving Nonconvex Quadratic Programs via Linear Integer Programming Techniques." In: *INFORMS Journal on Computing* 32.1, pp. 40–56. DOI: 10.1287/ijoc.2018.0883.

Zeng, B. (2020). "A Practical Scheme to Compute the Pessimistic Bilevel Optimization Problem." In: *INFORMS Journal on Computing* 32.4, pp. 1128–1142. DOI: 10.1287/ijoc.2019.0927.

Zheng, Y., G. Zhang, Z. Zhang, and J. Lu (2018). "A reducibility method for the weak linear bilevel programming problems and a case study in principal-agent." In: *Information Sciences* 454–455, pp. 46–58. DOI: 10.1016/j.ins.2018.04.045.

(V. Bucarey) Universidad de O'Higgins, Institute of Engineering Sciences, Lib. Bdo. O'Higgins 611, Rancagua, Chile
Email address: victor.bucarey@uoh.cl

(M. Labbé) Université Libre de Bruxelles, Département d'Informatique, Bruxelles, Belgium
Email address: martine.labbe@ulb.be

(G. Muñoz) Universidad de Chile, Department of Industrial Engineering, Beauchef 851, Santiago, Chile
Email address: gonzalo.m@uchile.cl

(M. Schmidt) Trier University, Department of Mathematics, Universitätsring 15, 54296 Trier, Germany
Email address: martin.schmidt@uni-trier.de