Optimization Letters
https://doi.org/10.1007/s11590-026-02331-8

ORIGINAL PAPER

Pricing semi-cooperation in unstable linear–linear bilevel programs via strictly convex tie-breaking and stability contracts

Massimiliano Caramia[1]

Received: 30 December 2025 / Accepted: 20 July 2026
© The Author(s) 2026

Abstract
When the follower problem in a bilevel program admits multiple optimal solutions, the leader's outcome depends on how ties are resolved. This yields the classical optimistic/pessimistic ambiguity and makes the model *unstable* (ill-posed) from a predictive viewpoint. We study linear–linear bilevel programs without additional upper-level coupling constraints involving the follower variables and propose a stabilization mechanism that preserves the follower's linear objective as a primary criterion, while resolving ties within the follower optimal set through a strictly convex quadratic rule. This tie-break selects a unique follower reaction and can be anchored to implement any prescribed follower optimum (e.g., a leader-favorable one). We then interpret stabilization as a semi-cooperation contract: in the presence of a (possibly unknown) secondary preference used by the follower to break ties, the leader may need to compensate the follower to accept the leader-chosen tie-break. We define the price of semi-cooperation as the minimal transfer ensuring acceptance, and we derive explicit formulas and robust bounds under preference uncertainty. A computational illustration highlights stabilization and pricing behavior on randomly generated instances.

Keywords Bilevel optimization · Stability · Optimistic versus pessimistic solutions · Regularization · Stackelberg games · Incentives

✉ Massimiliano Caramia
caramia@dii.uniroma2.it

[1] Dipartimento di Ingegneria dell'Impresa, University of Rome Tor Vergata, Via del Politecnico, 1, 00133 Rome, Italy

Published online: 08 August 2026

1 Introduction and related work

Bilevel optimization models hierarchical decision-making in which a leader chooses a decision x and a follower responds with a decision y by solving an optimization problem parameterized by x. The bilevel program studied can be written as

$$\begin{aligned} \min_{x,y} \quad & F(x,y) \\ \text{s.t.} \quad & x \in X, \\ & y \in \arg\min\{\, f(x,z) : z \in Y(x) \,\}, \end{aligned}$$

where X is the leader feasible set and $Y(x)$ is the follower feasible set induced by the leader's decision. Thus, the leader controls x, while y is not chosen freely by the leader; it must be an optimal response of the follower problem associated with x. These models originate in Stackelberg games [31] and have been widely developed in optimization; see, e.g., the early algorithmic and convexity work in [1], the bilevel linear programming overview in [5], and the surveys and bibliographies [10, 30]. Comprehensive treatments appear in monographs such as [12, 17] and in the MPEC-oriented book [26].

A central issue is *nonuniqueness* of the follower's optimal solutions. When the follower admits multiple optimal responses, the leader's outcome depends on how ties are resolved among follower optima. Two classical conventions are the *optimistic* and *pessimistic* leader, selecting the best and worst follower reactions for the leader, respectively. This ambiguity is a key source of ill-posedness and instability and is explicitly addressed in the bilevel literature, including reformulation and regularity studies [13, 25] and the analysis of ill-posed bilevel programs [32]. Related work investigates solution concepts and selection rules (especially for pessimistic variants and value-function approaches) [11, 14–16]. Robustness to deviations from the predicted follower response has been studied, for example, via limited-observability models [4] and near-optimality robustness [7]. For a broader view on robust and stochastic modeling in bilevel settings, see [2].

The present work is also related to intermediate solution concepts between the optimistic and pessimistic conventions. Such concepts are surveyed, for example, in Section 3.5 of [2], where different ways of relaxing the extreme assumptions on follower behavior are discussed. Our semi-cooperative model has a similar motivation, in the sense that it does not simply assume that the follower always selects the best or worst response for the leader. However, the mechanism studied here is different: the realized follower response is made explicit through a secondary tie-breaking criterion and, when this criterion conflicts with the leader-preferred tie, through a compensating transfer. Thus, the resulting solution lies between the optimistic and pessimistic benchmarks not by imposing an exogenous intermediate attitude, but by pricing the follower's acceptance of a particular optimal tie.

The notion of "stability" is used in the bilevel literature in (at least) two complementary senses: (i) *stability under perturbations* of data/parameters (e.g., continuity/Lipschitz properties of solution maps), and (ii) *well-posedness of the bilevel model* when the follower has multiple optima (i.e., specifying which reaction is realized).

Classical regularization approaches for Stackelberg/bilevel problems include ε-regularizations and vanishing Tikhonov terms [9, 23, 24, 29]. More recently, explicit "stable" formulations of optimistic bilevel optimization have been proposed (see, e.g., [27]), and relaxation-based approaches for pessimistic bilevel variants have been investigated (see, e.g., [6]). Finally, algorithmic procedures that explicitly compute stable solutions for linear–linear bilevel problems have been developed; see [8].

A large fraction of the stability literature proceeds by *perturbing* the follower problem (e.g., adding ε-regularization or Tikhonov terms) and then studying existence, convergence, and sensitivity properties as the perturbation vanishes. In contrast, the present paper separates two roles that are often conflated:

(a) *Selection*: specify an explicit rule for resolving ties *among follower-optimal decisions*, so that the follower's original (primary) objective is not compromised;
(b) *Computation*: approximate that rule through a regularized problem when convenient.

Our contribution combines an explicit tie-breaking rule with an economic interpretation of its enforcement, and provides tractable tools to quantify and manage instability:

- We propose a stabilization mechanism for linear–linear bilevel programs in which the follower may have multiple optima: the follower keeps its original objective as the primary criterion, while ties are resolved by a strictly convex quadratic rule, yielding a unique follower reaction.
- We show that the proposed tie-breaking rule can be tuned to reproduce any preselected follower optimum (in particular, a leader-favorable one) at a given leader decision, thereby making a desired outcome operational.
- For computation, we introduce a regularized follower problem that is easy to solve, and that converges to the exact tie-breaking selection as the regularization vanishes; for any fixed regularization level, the induced follower response varies stably with the leader decision.
- We interpret enforced tie-breaking as a *semi-cooperation contract*: because the follower may have an unmodeled secondary preference, the leader may need to compensate the follower to accept the leader-favorable tie-break. We define the corresponding *price of semi-cooperation* and derive explicit formulas when the follower's secondary preference is known.
- When the follower's secondary preference and/or its importance are uncertain, we introduce a robust acceptance notion and derive computable robust prices and bounds for natural uncertainty descriptions (e.g., scenario-based families and structured classes).
- Finally, we outline a practical single-level solution route for the regularized models based on standard optimality conditions, leading to mixed-integer formulations that can be handled by off-the-shelf solvers.

The idea that coordination between leader and follower may be used to address ill-posedness in bilevel programming is related to the coordination model in [22], where

the leader compensates the follower for collaborative behavior. Our approach differs in both scope and mechanism. We keep the follower's original linear objective as the primary criterion and use compensation only to implement a specified tie-breaking rule within the set of follower-optimal solutions. Thus, the transfer does not buy a violation of follower optimality; it prices the follower's acceptance of one optimal response rather than another. This leads to explicit compensating-variation formulas and robust counterparts for uncertain secondary preferences.

Motivating context. The semi-cooperation mechanism considered in this paper is not intended to model purely adversarial situations, such as interdiction problems between a defender and an attacker, or strictly competitive market interactions between rival firms. In such settings, compensating the follower to select a leader-favorable tie may be unrealistic. Rather, the mechanism is meant for hierarchical environments in which the leader and the follower have distinct objectives but can enter into contractual or regulatory arrangements. Examples include public procurement, regulated logistics, concession contracts, platform–operator interactions, and service-level agreements.

As a simple example, consider a public authority that chooses a tariff, subsidy, or service requirement x, and a contracted operator that then solves a linear operating problem, such as assigning vehicles, routes, or service frequencies. The operator's primary cost may have several optimal operating plans $y \in S(x)$ with the same operating cost. These alternative plans may be equivalent for the operator's primary objective but may differ substantially for the authority, for example in congestion, emissions, reliability, or coverage. If no tie-breaking rule is specified, the leader's predicted outcome is unstable: the same leader decision x may lead to several operator responses with different leader costs. In this context, a contract that compensates the operator for adopting one primary-cost-optimal plan rather than another is realistic: the payment does not ask the operator to violate its original optimality criterion, but only prices the selection of one optimal response within $S(x)$.

Organization of the paper. The rest of the paper is organized as follows. Section 2 introduces the unstable linear–linear bilevel model and formalizes instability exposure. Section 3 develops the strictly convex tie-break selection, its vanishing-regularization approximation, and related stability and computational properties. Section 4 introduces semi-cooperation contracts and the corresponding acceptance prices under an elicited secondary criterion. Sections 5 and 6 propose contract-aware stable solution concepts under elicited and robust acceptance, respectively. Section 7 relates stabilization prices to the leader's willingness to pay. Section 8 presents a numerical illustration. Finally, Section 9 concludes the paper and outlines directions for future research.

2 Unstable linear–linear bilevel programs

We consider a standard *min–min* linear–linear bilevel program. Let $X \subseteq \mathbb{R}^n$ be a nonempty polyhedron. For each $x \in X$, define the follower feasible set

$$Y(x) := \{y \in \mathbb{R}^m : By \le b + Fx\}, \tag{1}$$

where $B \in \mathbb{R}^{q\times m}$, $F \in \mathbb{R}^{q\times n}$ and $b \in \mathbb{R}^q$ are given.

Assumption 1 For all $x \in X$, the follower feasible set $Y(x)$ is nonempty and the linear program

$$\min\{f^\top y : y \in Y(x)\}$$

has at least one optimal solution.
For each $x \in X$, define the follower solution set

$$S(x) := \arg\min\{f^\top y : y \in Y(x)\}, \tag{2}$$

and the follower optimal value function

$$v(x) := \min\{f^\top y : y \in Y(x)\}. \tag{3}$$

The leader objective is affine,

$$C_L(x, y) = c^\top x + d^\top y. \tag{4}$$

We deliberately restrict attention to the case in which there are no additional upper-level coupling constraints involving y. Thus, apart from the follower optimality condition $y \in S(x)$, the leader feasibility constraints depend only on $x \in X$.

The bilevel model is

$$\text{“ } \min_{x\in X} \; C_L(x, y) \quad \text{s.t.} \quad y \in S(x), \text{”} \tag{5}$$

where the quotation marks emphasize that (5) does not yet define a single-valued optimization problem unless the realized element of $S(x)$ is specified. In other words, for a fixed leader decision $x \in X$, the follower is required to choose some element of $S(x)$, but the model has not yet prescribed which one.

If the leader is allowed to select the most favorable element of $S(x)$, then (5) becomes the optimistic formulation

$$\inf_{x\in X} \; \min_{y\in S(x)} \; C_L(x, y).$$

If, instead, the leader evaluates the worst admissible follower tie, one obtains the pessimistic formulation

$$\inf_{x\in X} \; \max_{y\in S(x)} \; C_L(x, y).$$

The instability considered in this paper is precisely the gap between the ambiguous model (5) and any concrete prediction of the realized follower response when $S(x)$ is not a singleton.

Assumption 2 For all $x \in X$, the leader cost $C_L(x, y) = c^\top x + d^\top y$ is bounded above and below on $S(x)$. Equivalently, for every $x \in X$, the quantities $\min_{y \in S(x)} d^\top y$ and $\max_{y \in S(x)} d^\top y$ are finite.

For any fixed $x \in X$, the leader faces a set $S(x)$ of follower optima. Define the best- and worst-case leader costs induced by follower tie-breaking as

$$\underline{C}(x) := \min_{y \in S(x)} C_L(x, y), \qquad \overline{C}(x) := \max_{y \in S(x)} C_L(x, y). \tag{6}$$

These quantities are well-defined and finite under Assumption 2. The optimistic and pessimistic bilevel values are then

$$z_{\text{opt}} := \inf_{x \in X} \underline{C}(x), \qquad z_{\text{pess}} := \inf_{x \in X} \overline{C}(x). \tag{7}$$

Definition 1 The bilevel model (5) is *unstable at x* if $S(x)$ is not a singleton, because the realized follower response is not uniquely specified by the lower-level optimality condition alone. The leader-cost instability exposure at x is

$$\Delta_L(x) := \overline{C}(x) - \underline{C}(x) \ \geq 0. \tag{8}$$

If an optimistic solution exists, i.e., if there is $x^O \in \arg\min_{x \in X} \underline{C}(x)$ and $y^O \in \arg\min_{y \in S(x^O)} C_L(x^O, y)$, then

$$\Delta_L(x^O) = C_L(x^O, y^P) - C_L(x^O, y^O) \quad \text{for any } y^P \in \arg\max_{y \in S(x^O)} C_L(x^O, y).$$

Thus $\Delta_L(x^O)$ is the leader's *worst-case loss due to tie-breaking* at the optimistic leader decision.

3 A strictly convex quadratic tie-break and stability

We define a follower selection rule that (i) *preserves* the follower's primary linear optimality and (ii) yields a *unique* reaction for every leader decision. By "preserving the follower's primary linear optimality" we mean that the implemented reaction is always selected from the original follower optimal set $S(x)$. Hence, the selected response y satisfies $f^\top y = v(x) := \min\{f^\top z : z \in Y(x)\}$.

For each $x \in X$, consider the follower LP

$$\min_{y} \quad f^\top y \quad \text{s.t.} \quad By \le b + Fx. \tag{9}$$

Its (Lagrange) dual is

$$\max_{\lambda \ge 0} \quad -(b + Fx)^\top \lambda \quad \text{s.t.} \quad B^\top \lambda = -f. \tag{10}$$

Under Assumption 1, the primal is feasible and has a finite optimal value, hence strong duality holds, and an optimal dual multiplier exists, that is, if we fix $x \in X$, a vector y is optimal for (9) if and only if there exists $\lambda \ge 0$ such that

$$\begin{aligned} By &\le b + Fx, \\ B^\top \lambda &= -f, \qquad \lambda \ge 0, \\ f^\top y &= -(b + Fx)^\top \lambda. \end{aligned} \tag{11}$$

Define the primal–dual optimality set

$$\Omega(x) := \{(y, \lambda) \in \mathbb{R}^m \times \mathbb{R}^q : \ (11) \text{ holds}\}. \tag{12}$$

where the projection of $\Omega(x)$ onto y is exactly $S(x)$.

3.1 Quadratic tie-break over the LP-optimal set

Fix a symmetric positive definite matrix $Q \in \mathbb{R}^{m \times m}$ and a target point $t \in \mathbb{R}^m$. Define

$$q_{Q,t}(y) := \tfrac{1}{2}(y - t)^\top Q(y - t).$$

The tie-break follower problem at leader decision x is

$$\min_{(y,\lambda) \in \Omega(x)} \ q_{Q,t}(y). \tag{13}$$

Equivalently, after projecting out the dual multiplier, it is

$$\min_{y \in S(x)} \ q_{Q,t}(y). \tag{14}$$

Thus, the quadratic criterion is applied only after the follower's linear objective has already been minimized. It selects among the elements of $S(x)$ and therefore does not alter the follower's primary optimal value.

Under Assumption 1, for every $x \in X$ the set $S(x)$ is nonempty. Moreover, since $Y(x) = \{y \in \mathbb{R}^m : By \le b + Fx\}$ is a polyhedron and $S(x) = \{y \in Y(x) : f^\top y = v(x)\} = \{y \in \mathbb{R}^m : By \le b + Fx, \ f^\top y = v(x)\}$, the follower optimal set $S(x)$ is a closed convex polyhedron; geometrically, it is the face

of $Y(x)$ exposed by the linear objective $f^\top y$. Since $q_{Q,t}$ is strongly convex and coercive, the tie-break problem (14) admits a unique optimal solution $y_{Q,t}(x) \in S(x)$. In particular, $f^\top y_{Q,t}(x) = v(x)$, and the map $x \mapsto y_{Q,t}(x)$ is single-valued.

Definition 2 For each $x \in X$, we call $y_{Q,t}(x)$ the *quadratic tie-break reaction*.

For fixed x, the tie-break solution is the (unique) projection of t onto $S(x)$ in the norm induced by Q:

$$y_{Q,t}(x) = \arg\min_{y \in S(x)} \tfrac{1}{2} \|y - t\|_Q^2, \qquad \|z\|_Q := \sqrt{z^\top Q z}.$$

We refer to the choice $t \in S(x)$ as an *anchored* tie-break. In this case, the target point t is not merely a reference point outside the follower optimal set; it is itself a follower optimum, and the tie-break rule implements it exactly.

Corollary 1 *Fix $x \in X$ and let $t \in S(x)$. Then, for every $Q \succ 0$, the quadratic tie-break reaction satisfies $y_{Q,t}(x) = t$. In particular, in the anchored case, the implemented follower response is independent of the choice of Q.*

The matrix Q determines the geometry of the secondary projection onto $S(x)$. In particular, it specifies the relative weights and correlations used to measure deviations from the target point t. Multiplying Q by a positive scalar does not change the selected point, since the objective $q_{Q,t}$ is then rescaled by the same positive factor. Thus, only the relative geometry induced by Q matters.

In practice, several choices are possible. The choice $Q = I$ gives the Euclidean projection onto $S(x)$ and is appropriate when the components of y are already commensurate or have been normalized. A diagonal matrix $Q = \mathrm{diag}(w_1, \ldots, w_m)$ can be used when deviations in different follower variables should be penalized with different intensities, for example because variables have different units, ranges, or economic importance. More general positive definite matrices can encode correlations or covariance information, and may be chosen, for instance, as inverse covariance or inverse scaling matrices when historical follower reactions are available.

The role of Q is different in anchored and non-anchored uses of the mechanism. If $t \in S(x)$, then the selected point is exactly t, independently of the particular choice of $Q \succ 0$. Hence, when the leader intends to implement a prescribed follower optimum by setting t equal to that optimum, Q affects neither the implemented reaction nor the associated compensation. By contrast, when $t \notin S(x)$, or when a fixed target t is used across several leader decisions, Q determines which follower optimum is closest to t in the induced norm and therefore affects the selected tie.

For $\delta > 0$ and fixed $Q \succ 0, t \in \mathbb{R}^m$, define the *regularized follower problem*

$$y_\delta(x) \in \arg\min_{y \in Y(x)} \left(f^\top y + \delta\, q_{Q,t}(y) \right). \tag{15}$$

Since $f^\top y + \delta q_{Q,t}(y)$ is strictly convex and coercive, (15) admits a unique solution for each x.

Theorem 1 *Fix $x \in X$ and $Q \succ 0$, $t \in \mathbb{R}^m$. Let $(\delta_k)_{k\in\mathbb{N}}$ be any sequence such that $\delta_k > 0$ for all k and $\delta_k \downarrow 0$. For each k, let $y_{\delta_k}(x)$ denote the unique optimal solution of (15) with $\delta = \delta_k$. Then*

$$f^\top y_{\delta_k}(x) \to v(x) \qquad \text{and} \qquad y_{\delta_k}(x) \to y_{Q,t}(x),$$

where

$$v(x) = \min\{f^\top y : y \in Y(x)\}.$$

Moreover, for every $\delta > 0$ and every $y \in S(x)$, the unique solution $y_\delta(x)$ of (15) satisfies

$$q_{Q,t}\big(y_\delta(x)\big) \le q_{Q,t}(y). \tag{16}$$

Proof Fix $x \in X$ and abbreviate $S := S(x)$, $q := q_{Q,t}$, and $v := v(x) = \min\{f^\top y : y \in Y(x)\}$. Let $y_\delta := y_\delta(x)$ denote the unique solution of (15) for a fixed $\delta > 0$.

Pick any $\bar{y} \in S$. By optimality of y_δ in (15),

$$f^\top y_\delta + \delta q(y_\delta) \le f^\top \bar{y} + \delta q(\bar{y}) = v + \delta q(\bar{y}).$$

Since $f^\top y_\delta \ge v$, it follows that

$$q(y_\delta) \le q(\bar{y}),$$

and, more generally, $q(y_\delta) \le q(y)$ for every $y \in S$. This proves (16). The same inequality also implies that the family $\{y_\delta : \delta > 0\}$ is bounded, because q is coercive.

Now let $(\delta_k)_k$ be any positive sequence with $\delta_k \downarrow 0$ and set $y_k := y_{\delta_k}(x)$. From the preceding inequality and the optimality relation,

$$0 \le f^\top y_k - v \le \delta_k\big(q(\bar{y}) - q(y_k)\big) \le \delta_k q(\bar{y}),$$

so $f^\top y_k \to v$. Since $\{y_k\}$ is bounded, it has at least one accumulation point. Let y° be any accumulation point. Because $Y(x)$ is closed and $f^\top y_k \to v$, we have $y^\circ \in S$.

Furthermore, for every $y \in S$, inequality (16) gives $q(y_k) \le q(y)$. Passing to the limit along a subsequence converging to y° yields $q(y^\circ) \le q(y)$, $\forall y \in S$. Thus, y° minimizes q over S. Since q is strongly convex and S is convex, this minimizer is unique and equals $y_{Q,t}(x)$. Hence, every accumulation point of $\{y_k\}$ equals $y_{Q,t}(x)$, and therefore the whole sequence converges:

$$y_{\delta_k}(x) \to y_{Q,t}(x).$$

□

Remark 1 It is important to distinguish the exact tie-break rule from the regularized follower problem. The selection rule is (14), namely the minimization of $q_{Q,t}$ over the exact LP-optimal set $S(x)$. Therefore, the implemented reaction $y_{Q,t}(x)$ always satisfies $f^\top y_{Q,t}(x) = v(x)$, and the follower's original linear objective is preserved as a lexicographically primary criterion. By contrast, the auxiliary problem (15) minimizes $f^\top y + \delta q_{Q,t}(y)$ over $Y(x)$. For any fixed $\delta > 0$, this problem is a computational proxy and may select a point with $f^\top y_\delta(x) > v(x)$. Theorem 1 shows that this loss of exact primary optimality vanishes as $\delta \downarrow 0$, and that the limit recovers the exact tie-break reaction $y_{Q,t}(x)$.

One advantage of (15) is *stability under perturbations*. For fixed $\delta > 0$, (15) is a parametric strongly convex quadratic program with polyhedral constraints and an affine right-hand side. Hence, its optimizer $y_\delta(x)$ is uniquely defined, and in regions where the active-set structure is unchanged, it is an affine function of x. Consequently, the solution map is piecewise affine on a polyhedral partition of its domain and is locally Lipschitz on the relative interior of X, denoted by $\mathrm{ri}(X)$, i.e., the interior of X relative to its affine hull. See, e.g., [28] and the general solution-mapping regularity results in [18].

3.2 A KKT-based single-level route (MILP)

For computation, the regularized follower problem admits a KKT characterization that can be embedded in a single-level model.

Let $\delta > 0$ be fixed. The KKT conditions for (15) at a given x are: there exists $\lambda \in \mathbb{R}^q_+$ such that

$$\begin{aligned} By &\le b + Fx, \\ \lambda &\ge 0, \\ f + \delta Q(y - t) + B^\top \lambda &= 0, \\ \lambda_i \left(b_i + F_i x - B_i y\right) &= 0 \qquad i = 1, \ldots, q. \end{aligned} \tag{17}$$

Because the objective is strongly convex, the associated optimizer y is unique. Moreover, (15) is a convex quadratic program with affine constraints. Hence, for every fixed $x \in X$ for which $Y(x)$ is nonempty and the regularized problem has a finite optimal value, the KKT system (17) characterizes the unique optimal solution; see, e.g., [3, 20]. In particular, the stationarity condition can be interpreted as

$$0 \in f + \delta Q(y - t) + N_{Y(x)}(y),$$

where $N_{Y(x)}(y)$ is the normal cone to the polyhedron $Y(x)$ at y. Since $Y(x)$ is polyhedral, this normal cone is generated by the normals of the active affine inequalities, yielding multipliers $\lambda \geq 0$ satisfying (17).

Proposition 1 *Fix $\delta > 0$. Consider the bilevel model where the follower uses (15) and the leader minimizes $C_L(x, y)$ (or $C_L(x, y) + T$ if a contract transfer is included, see the next section for details). Then the problem can be formulated as a single-level MPEC using (17). Moreover, using standard big-M linearizations of complementarity, one obtains a mixed-integer linear program (MILP) reformulation.*

The reformulation in Proposition 1 follows the standard KKT-based approach for bilevel programs with convex lower-level problems. The complementarity conditions in (17) can be encoded by Fortuny-Amat–McCarl type big-*M* disjunctions; see [19] and the survey of mixed-integer reformulation techniques in bilevel optimization in [21].

The big-M constants used to linearize the complementarity conditions in (17) must be valid upper bounds on the corresponding primal slacks and dual multipliers. More precisely, if

$$s_i := b_i + F_i x - B_i y \geq 0$$

denotes the slack of the i-th lower-level constraint, then the complementarity condition $\lambda_i s_i = 0$ can be encoded, for example, by a binary variable z_i through

$$0 \leq \lambda_i \leq M_i^{\lambda} z_i, \qquad 0 \leq s_i \leq M_i^{s}(1 - z_i).$$

In the regularized case (15), the lower-level problem is a strongly convex quadratic program with affine constraints. Therefore, the KKT system is necessary and sufficient for optimality. However, the construction of tight valid bounds for the multipliers is more delicate than in a purely linear lower-level problem. In practice, such bounds may be obtained from explicit variable bounds, dual feasibility relations, problem-specific estimates, or auxiliary optimization problems that maximize each slack and multiplier over the relevant feasible KKT region. Poorly chosen big-M constants may lead to weak relaxations or numerical instability. For this reason, the MILP route in Proposition 1 should be understood as a standard computational formulation whose numerical performance depends on the availability of valid and reasonably tight bounds.

This computational route is often more convenient than embedding the exact set $S(x)$ through primal–dual certificates. Here, embedding $S(x)$ means enforcing inside the leader problem that the chosen follower response is optimal for the original lower-level LP, for example by imposing the primal feasibility, dual feasibility, and strong-duality conditions defining $\Omega(x)$. The regularized route instead uses, for each fixed $\delta > 0$, a strongly convex lower-level problem with a unique optimizer, whose KKT conditions define a single-valued approximating reaction. Thus, the regularized route is more convenient in this practical sense; the exact tie-break is then recovered only in the limit as $\delta \downarrow 0$.

The parameter δ controls the trade-off between numerical conditioning and accuracy with respect to the exact lexicographic tie-break. For fixed $\delta > 0$, the regularized lower-level problem is strongly convex and yields a single-valued reaction, but the selected point may satisfy $f^\top y_\delta(x) > v(x)$. By Theorem 1, this loss of exact primary optimality vanishes as $\delta \downarrow 0$. In practice, one may use a continuation strategy: start from a moderate value $\delta_0 > 0$, solve the corresponding MILP, then decrease δ geometrically, for instance $\delta_{\ell+1} = \rho\delta_\ell$ with $\rho \in (0, 1)$, until both the change in the leader objective and, when $v(x)$ is available or can be evaluated, the primary optimality residual $f^\top y_{\delta_\ell}(x) - v(x)$ are below prescribed tolerances. Warm-starting the MILP at each continuation step can reduce the computational burden. Very small values of δ may, however, lead to ill-conditioning and weaker numerical performance, so δ should be decreased only to the accuracy required by the intended approximation of the exact tie-break.

4 Semi-cooperation contracts and the price of semi-cooperation

The tie-break mechanism provides mathematical stabilization. We now attach *economic meaning*: the follower may require compensation to accept a leader-chosen tie-break.

4.1 A contract model

A *semi-cooperation contract* consists of parameters (Q, t, T) where $Q \succ 0$ and $T \geq 0$ is a transfer paid by the leader to the follower (*paid only if the contract is accepted*).

Given a leader decision x, the follower always enforces *primary optimality* by restricting attention to the follower-optimal set $S(x)$. If the follower *accepts* the contract, the implemented action is the quadratic tie-break reaction $y_{Q,t}(x) \in S(x)$.

If the follower *rejects*, it breaks ties within $S(x)$ according to a default (no-contract) secondary criterion $\phi : \mathbb{R}^m \to \mathbb{R}$. The corresponding reservation level, namely the best secondary value achievable by the follower within the original optimal set $S(x)$, is

$$\phi^\star(x) := \inf\{\phi(y) : y \in S(x)\}. \tag{18}$$

Assumption 3 For all $x \in X$, the infimum in (18) is finite and attained.

Under Assumption 3, the set of default secondary minimizers is nonempty and is given by

$$D_\phi(x) := \arg\min\{\phi(y) : y \in S(x)\}. \tag{19}$$

If $D_\phi(x)$ is not a singleton, we adopt a leader-favorable convention within this default set and select

$$y^{\mathrm{def}}(x) \in \arg\min\{C_L(x,y) : y \in D_\phi(x)\}. \tag{20}$$

Thus, the follower first minimizes its secondary disutility ϕ over $S(x)$; only if several such minimizers exist do we use an optimistic tie-break from the leader's perspective. This convention makes the no-contract realized leader cost well defined, but it does not affect the reservation value $\phi^\star(x)$, since all elements of $D_\phi(x)$ have the same secondary value. In particular,

$$\phi(y^{\mathrm{def}}(x)) = \phi^\star(x).$$

The leader aims to minimize its *net* cost

$$J(x;Q,t,T) = \begin{cases} C_L(x, y_{Q,t}(x)) + T, & \text{if accepted,} \\ C_L(x, y^{\mathrm{def}}(x)), & \text{if rejected.} \end{cases} \tag{21}$$

The leader chooses *x* and offers (*Q*, *t*, *T*); the follower observes (*x*, *Q*, *t*, *T*), accepts or rejects, and then the implemented *y* is selected as above.

4.2 Mechanism (i): elicited secondary preference over the follower LP-optimal set

Mechanism (i) is a benchmark where the leader can evaluate (or has elicited) the follower's tie-breaking criterion ϕ and a scale parameter $\epsilon \geq 0$.

Follower acceptance. We interpret ϕ as a *secondary disutility* used only to select among lower-level optima. The scalar ϵ measures how strongly the follower cares about this secondary disutility. If the contract is accepted at *x*, the follower incurs secondary disutility $\epsilon\,\phi(y_{Q,t}(x))$ and receives transfer *T*. If rejected, it incurs the baseline level $\epsilon\,\phi^\star(x)$ (achieved by $y^{\mathrm{def}}(x)$).

Accordingly, the contract is *individually rational* (IR) for the follower at *x* if

$$\epsilon\,\phi\big(y_{Q,t}(x)\big) - T \;\leq\; \epsilon\,\phi^\star(x). \tag{22}$$

Equivalently,

$$T \;\geq\; \epsilon\big(\phi(y_{Q,t}(x)) - \phi^\star(x)\big). \tag{23}$$

In contract-theoretic terms, individual rationality means that the follower must weakly prefer accepting the proposed contract to rejecting it. Rejection gives the follower its outside option: it remains within the original optimal set *S*(*x*) and selects a default tie that minimizes its secondary disutility ϕ. Acceptance, instead, implements the contract-selected tie $y_{Q,t}(x)$ and gives the follower the transfer *T*. Thus, a contract is acceptable only if the transfer compensates the follower for any increase in secondary disutility relative to the no-contract default outcome.

Since $y_{Q,t}(x) \in S(x)$ and $\phi^\star(x) = \min\{\phi(y) : y \in S(x)\}$, we have $\phi\big(y_{Q,t}(x)\big) - \phi^\star(x) \geq 0$. Therefore, the smallest nonnegative transfer satisfying the

individual rationality condition (23) is well defined and is given directly by the following expression.

Definition 3 Fix $x \in X$, $Q \succ 0$, $t \in \mathbb{R}^m$, and an elicited pair (ϵ, ϕ). The *price of semi-cooperation* is

$$\pi_\phi(x; Q, t) := \epsilon\Big(\phi\big(y_{Q,t}(x)\big) - \phi^\star(x)\Big) = \epsilon\Big(\phi\big(y_{Q,t}(x)\big) - \min_{y \in S(x)} \phi(y)\Big). \quad (24)$$

This is a compensating-variation expression: it is the (scaled) increase in the follower's secondary disutility incurred by implementing the contract-selected tie-break action rather than the follower's preferred tie within $S(x)$. In particular, $\pi_\phi(x; Q, t) = 0$ if and only if $y_{Q,t}(x)$ is itself a ϕ-minimizer on $S(x)$.

5 A contract-aware stable solution concept under elicited secondary preference

We now incorporate the semi-cooperation price into the leader's decision problem. Throughout this section, the follower's secondary preference ϕ and its scale $\epsilon \geq 0$ are assumed to be elicited. The contract acceptance price $\pi_\phi(x; Q, t)$ is the quantity defined in (24); it is the minimal transfer that makes the follower weakly prefer the contract-selected tie-break reaction to its no-contract default tie.

Corollary 2 *Fix $x \in X$ and $y \in S(x)$. Let $Q \succ 0$ be arbitrary and set $t := y$. Then $y_{Q,t}(x) = y$ and the minimal individually rational transfer equals*

$$\pi_\phi(x; y) := \pi_\phi(x; Q, y) = \epsilon\big(\phi(y) - \phi^\star(x)\big). \quad (25)$$

5.1 The semi-cooperative (contract-aware) bilevel problem

By Corollary 2, it is without loss of generality to parametrize the leader's choice by (x, y) with $y \in S(x)$ and interpret y as the intended implemented action, enforced through an anchored contract.

Definition 4 Under mechanism (i), define the *semi-cooperative (contract-aware) bilevel problem* as

$$z_{\mathrm{sc}} := \min_{x \in X} \min_{y \in S(x)} \Big\{ C_L(x, y) + \pi_\phi(x; y) \Big\}, \quad (26)$$

where $\pi_\phi(x; y)$ is given by (25). Any pair $(x^{\mathrm{sc}}, y^{\mathrm{sc}})$ attaining (26) is called a *semi-cooperative stable solution*.
Using (4) and (25), (26) admits the equivalent reduced form

$$z_{\mathrm{sc}} = \min_{x \in X} \left\{ c^\top x - \epsilon\, \phi^\star(x) \;+\; \min_{y \in S(x)} \big(d^\top y + \epsilon\, \phi(y) \big) \right\}. \tag{27}$$

Theorem 2 *Suppose that mechanism (i) is used to break ties, with elicited secondary preference ϕ and scale parameter $\epsilon \geq 0$, and that Assumptions 1–3 hold. Then the semi-cooperative value satisfies*

$$z_{\mathrm{opt}} \;\leq\; z_{\mathrm{sc}} \;\leq\; z_{\mathrm{pess}}.$$

Moreover, if $\epsilon = 0$, then $z_{\mathrm{sc}} = z_{\mathrm{opt}}$.

Proof Fix any $x \in X$. For every $y \in S(x)$, the transfer satisfies

$$\pi_\phi(x; y) = \epsilon\big(\phi(y) - \phi^\star(x)\big) \;\geq\; 0,$$

hence

$$C_L(x, y) + \pi_\phi(x; y) \;\geq\; C_L(x, y).$$

Taking the minimum over $y \in S(x)$ yields

$$\min_{y \in S(x)} \big\{ C_L(x, y) + \pi_\phi(x; y) \big\} \;\geq\; \min_{y \in S(x)} C_L(x, y) = \underline{C}(x).$$

Now minimize over $x \in X$ to obtain $z_{\mathrm{sc}} \geq z_{\mathrm{opt}}$.

For the upper bound, pick $y^{\mathrm{def}}(x) \in D_\phi(x)$. Then $\phi(y^{\mathrm{def}}(x)) = \phi^\star(x)$, so $\pi_\phi(x; y^{\mathrm{def}}(x)) = 0$ and

$$\min_{y \in S(x)} \big\{ C_L(x, y) + \pi_\phi(x; y) \big\} \;\leq\; C_L\big(x, y^{\mathrm{def}}(x)\big) \;\leq\; \max_{y \in S(x)} C_L(x, y) = \overline{C}(x).$$

Minimizing over $x \in X$ yields $z_{\mathrm{sc}} \leq z_{\mathrm{pess}}$.

Finally, if $\epsilon = 0$, then $\pi_\phi(x; y) = 0$ for all (x, y) and (26) reduces to the optimistic formulation, so $z_{\mathrm{sc}} = z_{\mathrm{opt}}$. □

The following result shows that an anchored contract can implement an attaining solution to the semi-cooperative problem.

Theorem 3 *Let $(x^{\mathrm{sc}}, y^{\mathrm{sc}})$ be a semi-cooperative stable solution. Fix any $Q \succ 0$, set $t := y^{\mathrm{sc}}$, $T := \pi_\phi(x^{\mathrm{sc}}; y^{\mathrm{sc}})$, and offer the contract (Q, t, T) at the leader decision x^{sc}. Then:*

1. *the accepted contract implements the unique tie-break reaction $y_{Q,t}(x^{\mathrm{sc}}) = y^{\mathrm{sc}}$;*
2. *the transfer T makes the contract individually rational for the follower at x^{sc};*
3. *the leader's realized net cost under the accepted contract is*

$$C_L(x^{\mathrm{sc}}, y^{\mathrm{sc}}) + T = C_L(x^{\mathrm{sc}}, y^{\mathrm{sc}}) + \pi_\phi(x^{\mathrm{sc}}; y^{\mathrm{sc}}),$$

which is the objective value attained by $(x^{\mathrm{sc}}, y^{\mathrm{sc}})$ *in (26).*

Proof Since $(x^{\mathrm{sc}}, y^{\mathrm{sc}})$ is a semi-cooperative stable solution, we have $y^{\mathrm{sc}} \in S(x^{\mathrm{sc}})$. By Corollary 1, the anchored choice $t = y^{\mathrm{sc}}$ implies $y_{Q,t}(x^{\mathrm{sc}}) = y^{\mathrm{sc}}$ for every $Q \succ 0$. Uniqueness follows from the strict convexity of $q_{Q,t}$ on the closed convex set $S(x^{\mathrm{sc}})$.

By the definition of the semi-cooperation price, $T = \pi_\phi(x^{\mathrm{sc}}; y^{\mathrm{sc}}) = \epsilon\big(\phi(y^{\mathrm{sc}}) - \phi^\star(x^{\mathrm{sc}})\big)$, and therefore the individual rationality inequality holds with equality: $\epsilon\,\phi(y^{\mathrm{sc}}) - T = \epsilon\,\phi^\star(x^{\mathrm{sc}})$. Thus, the contract is individually rational.

Finally, since the accepted contract implements y^{sc} and pays transfer T, the leader's realized net cost is $C_L(x^{\mathrm{sc}}, y^{\mathrm{sc}}) + T = C_L(x^{\mathrm{sc}}, y^{\mathrm{sc}}) + \pi_\phi(x^{\mathrm{sc}}; y^{\mathrm{sc}})$, which is exactly the value attained by $(x^{\mathrm{sc}}, y^{\mathrm{sc}})$ in (26). □

Remark 2 For fixed x, the inner problem in (27) can be written as

$$\inf_{y \in S(x)} \left\{ d^\top y + \epsilon\big(\phi(y) - \phi^\star(x)\big) \right\}.$$

Thus ϵ controls the trade-off between the leader's preference over the follower ties, represented by $d^\top y$, and the compensation required to move the follower away from its default secondary optimum.

When $\epsilon = 0$, the compensation term vanishes, and the inner problem selects a leader-preferred tie, thus recovering the optimistic convention. Conversely, as ϵ becomes large, any tie satisfying $\phi(y) > \phi^\star(x)$ becomes increasingly expensive to implement. Therefore, for a fixed x, the selected tie is driven toward the default set

$$D_\phi(x) := \arg\min\{\phi(y) : y \in S(x)\}.$$

More precisely, if y_ϵ is a minimizer of the fixed-x inner problem and $\bar{y} \in D_\phi(x)$, then

$$d^\top y_\epsilon + \epsilon\big(\phi(y_\epsilon) - \phi^\star(x)\big) \le d^\top \bar{y}.$$

Since $d^\top y$ is bounded on $S(x)$ under Assumption 2, it follows that

$$\phi(y_\epsilon) - \phi^\star(x) \to 0 \qquad \text{as } \epsilon \to +\infty.$$

Hence, every accumulation point of y_ϵ, when it exists and ϕ is continuous, belongs to $D_\phi(x)$. Among the ties in $D_\phi(x)$, the leader term $d^\top y$ selects the leader-favorable default minimizer.

Thus, the large-ϵ regime should be interpreted as convergence toward the no-contract/default-tie regime, not necessarily toward the pessimistic regime. It may

resemble the pessimistic benchmark only when the follower's default secondary minimizers are unfavorable to the leader.

6 A contract-aware stable solution concept under robust acceptance

We now consider mechanism (ii), in which the leader faces uncertainty about the follower's scale parameter and/or secondary preference. The leader seeks a contract that is accepted *robustly*, i.e., for all admissible follower types, and internalizes the corresponding (worst-case) transfer in the bilevel objective.

6.1 Robust individual rationality and worst-case minimal transfer

Let $\mathcal{E} \subseteq \mathbb{R}_+$ be a nonempty uncertainty set for the scale parameter ϵ and let $\mathcal{P}$ be a nonempty set of admissible secondary preferences $\phi : \mathbb{R}^m \to \mathbb{R}$. For each $\phi \in \mathcal{P}$ and $x \in X$, define the reservation level

$$\phi^\star(x;\phi) := \inf\{\phi(y) : y \in S(x)\}. \tag{28}$$

Assumption 4 For all $x \in X$, $S(x)$ is nonempty and for all $\phi \in \mathcal{P}$ the minimum in (28) is attained and finite. Moreover, for all $x \in X$ and all $y \in S(x)$,

$$\sup_{(\epsilon,\phi)\in\mathcal{E}\times\mathcal{P}} \epsilon\big(\phi(y) - \phi^\star(x;\phi)\big) < \infty.$$

Fix a contract (Q, t, T) with $Q \succ 0$ and consider the tie-break reaction $y_{Q,t}(x)$ at leader decision x. We say that the contract is *robustly individually rational* (robust-IR) at x if, for every $(\epsilon, \phi) \in \mathcal{E} \times \mathcal{P}$,

$$\epsilon\,\phi\big(y_{Q,t}(x)\big) - T \;\le\; \epsilon\,\phi^\star(x;\phi). \tag{29}$$

Equivalently, the robust individual rationality condition (29) holds for every $(\epsilon, \phi) \in \mathcal{E} \times \mathcal{P}$ if and only if

$$T \;\ge\; \sup_{(\epsilon,\phi)\in\mathcal{E}\times\mathcal{P}} \epsilon\Big(\phi\big(y_{Q,t}(x)\big) - \phi^\star(x;\phi)\Big). \tag{30}$$

Definition 5 For fixed $x \in X$ and (Q, t) with $Q \succ 0$, define the *robust minimal transfer* as

$$\pi_{\mathrm{rob}}(x;Q,t) := \inf\{T \ge 0 :\ (29)\ \text{holds}\} = \sup_{(\epsilon,\phi)\in\mathcal{E}\times\mathcal{P}} \epsilon\Big(\phi\big(y_{Q,t}(x)\big) - \phi^\star(x;\phi)\Big). \tag{31}$$

A particularly useful bounded-scale uncertainty set is $\mathcal{E} = [0, \overline{\epsilon}]$, $0 < \overline{\epsilon} < +\infty$, where $\overline{\epsilon}$ is the largest admissible weight attached to the follower's secondary disutility.

Proposition 2 *Assume $\mathcal{E} = [0, \overline{\epsilon}]$, $0 < \overline{\epsilon} < +\infty$. Then, for every $x \in X$ and (Q, t) with $Q \succ 0$,*

$$\pi_{\mathrm{rob}}(x; Q, t) = \overline{\epsilon} \cdot \sup_{\phi \in \mathcal{P}} \Big(\phi\big(y_{Q,t}(x)\big) - \phi^{\star}(x; \phi) \Big). \tag{32}$$

Proof Fix $\phi \in \mathcal{P}$. Since $y_{Q,t}(x) \in S(x)$ and $\phi^{\star}(x; \phi) = \min_{y \in S(x)} \phi(y)$, we have $\phi(y_{Q,t}(x)) - \phi^{\star}(x; \phi) \geq 0$. Therefore, $\epsilon \mapsto \epsilon(\phi(y_{Q,t}(x)) - \phi^{\star}(x; \phi))$ is nondecreasing on $\mathbb{R}_+$, so the supremum over $[0, \overline{\epsilon}]$ is attained at $\overline{\epsilon}$. □

6.2 A tractable robust family and computable prices

A natural way to make mechanism (ii) operational is to restrict the admissible secondary preferences to a *finite scenario set*

$$\mathcal{P} = \{\phi^1, \ldots, \phi^K\},$$

where each $\phi^k : \mathbb{R}^m \to \mathbb{R}$ represents a plausible follower tie-breaking criterion (e.g., learned from data or elicited by negotiation). For each k and $x \in X$, define the scenario reservation level

$$\phi^{k\star}(x) := \min\{\phi^k(y) : y \in S(x)\}.$$

Proposition 3 *Assume that $\mathcal{E} = [0, \overline{\epsilon}]$, $0 < \overline{\epsilon} < +\infty$, and that $\mathcal{P} = \{\phi^1, \ldots, \phi^K\}$. For each $k = 1, \ldots, K$, define $\phi^{k\star}(x) := \min\{\phi^k(u) : u \in S(x)\}$. Then, for every $x \in X$ and every (Q, t) with $Q \succ 0$,*

$$\pi_{\mathrm{rob}}(x; Q, t) = \overline{\epsilon} \cdot \max_{k=1,\ldots,K} \Big(\phi^k\big(y_{Q,t}(x)\big) - \phi^{k\star}(x) \Big).$$

In particular, if an intended follower tie $y \in S(x)$ is implemented by setting $t := y$, then $y_{Q,t}(x) = y$ for every $Q \succ 0$. Hence, the anchored robust price is well defined by $\pi_{\mathrm{rob}}(x; y) := \pi_{\mathrm{rob}}(x; Q, y)$, independently of the choice of $Q \succ 0$, and satisfies

$$\pi_{\mathrm{rob}}(x; y) = \overline{\epsilon} \cdot \max_{k=1,\ldots,K} \Big(\phi^k(y) - \phi^{k\star}(x) \Big).$$

Proof By Proposition 2, since $\mathcal{E} = [0, \overline{\epsilon}]$, we have $\pi_{\mathrm{rob}}(x; Q, t) = \overline{\epsilon} \cdot \sup_{\phi \in \mathcal{P}} \Big(\phi\big(y_{Q,t}(x)\big) - \phi^{\star}(x; \phi) \Big)$. Because $\mathcal{P} = \{\phi^1, \ldots, \phi^K\}$ is finite, the supremum over $\mathcal{P}$ is a maximum over $k = 1, \ldots, K$. Moreover, $\phi^{\star}(x; \phi^k) = \phi^{k\star}(x)$, for $k = 1, \ldots, K$. Therefore,

$\pi_{\mathrm{rob}}(x;Q,t) = \overline{\epsilon} \cdot \max_{k=1,\dots,K}\Big(\phi^k\big(y_{Q,t}(x)\big) - \phi^{k\star}(x)\Big)$. Now consider the anchored case. If $y \in S(x)$ and $t := y$, then Corollary 1 gives $y_{Q,t}(x) = y$ for every $Q \succ 0$. Hence, $\pi_{\mathrm{rob}}(x;Q,y)$ is independent of Q, and it is therefore meaningful to denote this common value by $\pi_{\mathrm{rob}}(x;y)$. Substituting $y_{Q,t}(x) = y$ into the previous formula yields $\pi_{\mathrm{rob}}(x;y) = \overline{\epsilon} \cdot \max_{k=1,\dots,K}\Big(\phi^k(y) - \phi^{k\star}(x)\Big)$, as claimed. □

If $\phi^k(y) = (a^k)^\top y$ is linear, then $\phi^{k\star}(x)$ is the optimal value of the linear program

$$\min_{y \in S(x)} (a^k)^\top y,$$

which can be computed as

$$\phi^{k\star}(x) = \min\{(a^k)^\top y : (y,\lambda) \in \Omega(x)\}.$$

Thus, for fixed $(x,y) \in X \times S(x)$, the robust price $\pi_{\mathrm{rob}}(x;y)$ under finitely many linear scenarios can be evaluated by solving K linear programs of the above form and taking a maximum.

6.3 A computable upper bound via Lipschitz secondary preferences

Even when $\mathcal{P}$ is infinite, one can derive simple *a priori* bounds on the robust transfer using mild regularity.

Definition 6 For a norm $\|\cdot\|$ on $\mathbb{R}^m$ and $x \in X$, define the diameter of $S(x)$ as

$$\mathrm{diam}_{\|\cdot\|}(S(x)) := \sup\{\|y - y'\| : y, y' \in S(x)\} \in [0, +\infty].$$

Assumption 5 There exist a norm $\|\cdot\|$ on $\mathbb{R}^m$ and a constant $L \geq 0$ such that every $\phi \in \mathcal{P}$ is L-Lipschitz on $S(x)$ for every $x \in X$, i.e.,

$$|\phi(y) - \phi(y')| \leq L\|y - y'\| \qquad \forall x \in X,\ \forall y, y' \in S(x),\ \forall \phi \in \mathcal{P}.$$

Proposition 4 *Assume $\mathcal{E} = [0, \overline{\epsilon}]$ and Assumption 5. Then for every $x \in X$ and every intended tie $y \in S(x)$,*

$$\pi_{\mathrm{rob}}(x;y) \ \leq\ \overline{\epsilon}\, L\, \mathrm{diam}_{\|\cdot\|}(S(x)).$$

Proof Fix $x \in X$, $y \in S(x)$, and $\phi \in \mathcal{P}$. Let $y^\phi \in \arg\min\{\phi(u) : u \in S(x)\}$ so that $\phi^\star(x;\phi) = \phi(y^\phi)$. Then by Lipschitz continuity on $S(x)$,

$$\phi(y) - \phi^\star(x;\phi) = \phi(y) - \phi(y^\phi) \ \leq\ L\|y - y^\phi\| \ \leq\ L\, \mathrm{diam}_{\|\cdot\|}(S(x)).$$

Multiplying by $\bar{\epsilon}$ and taking the supremum over $\phi \in \mathcal{P}$ yields the claim. □

Remark 3 If one chooses $\|\cdot\| = \|\cdot\|_\infty$, then

$$\operatorname{diam}_\infty(S(x)) = \max_{i=1,\dots,m} \left(\max_{y \in S(x)} y_i - \min_{y \in S(x)} y_i \right),$$

and each max/min over $S(x)$ can be computed as a linear program over $\Omega(x)$:

$$\max_{y \in S(x)} y_i = \max\{y_i : (y, \lambda) \in \Omega(x)\}, \qquad \min_{y \in S(x)} y_i = \min\{y_i : (y, \lambda) \in \Omega(x)\}.$$

Hence $\operatorname{diam}_\infty(S(x))$ can be evaluated by solving at most $2m$ linear programs for fixed x.

6.4 Computational remarks: evaluating reactions and prices for fixed leader decision

This subsection clarifies which objects in the proposed stabilization-and-contract framework are efficiently computable for a *fixed* leader decision x.

For given (x, Q, t) with $Q \succ 0$, the exact tie-break reaction $y_{Q,t}(x)$ is the unique minimizer of (14) and can be computed by solving the convex quadratic program (13) over the polyhedron $\Omega(x)$. Equivalently, one can compute $y_{Q,t}(x)$ as the limit of the single strongly convex program (15) for a small $\delta > 0$ (Theorem 1), which only uses the primal constraints $By \le b + Fx$.

In the case $t = y \in S(x)$, Corollary 1 yields $y_{Q,t}(x) = y$ for any $Q \succ 0$. Thus, when the leader intends to enforce a specific $y \in S(x)$, no quadratic program is needed to compute the implemented action.

Under mechanism (i), the price for enforcing an intended tie $y \in S(x)$ is

$$\pi_\phi(x; y) = \epsilon\big(\phi(y) - \phi^\star(x)\big), \qquad \phi^\star(x) = \min\{\phi(u) : u \in S(x)\}.$$

Using the primal–dual certificate set $\Omega(x)$ (Section 3),

$$\phi^\star(x) = \min\{\phi(u) : (u, \lambda) \in \Omega(x)\}.$$

Hence, for fixed x, computing $\phi^\star(x)$ reduces to minimizing ϕ over a polyhedron. In particular:

- if ϕ is linear, this is a linear program;
- if ϕ is convex quadratic, this is a convex quadratic program;
- more generally, if ϕ is convex and efficiently evaluable, this is a standard convex program over a polyhedron.

Under the finite-scenario model $\mathcal{P} = \{\phi^1, \dots, \phi^K\}$ and $\mathcal{E} = [0, \bar{\epsilon}]$ (Section 6.2), the robust price for enforcing $y \in S(x)$ is

$$\pi_{\mathrm{rob}}(x;y) = \overline{\epsilon} \cdot \max_{k=1,\dots,K} \Big(\phi^k(y) - \phi^{k\star}(x) \Big), \qquad \phi^{k\star}(x) = \min\{\phi^k(u) : (u,\lambda) \in \Omega(x)\}.$$

Thus, for fixed x and y, evaluating $\pi_{\mathrm{rob}}(x;y)$ amounts to solving K optimization problems of the above form and taking a maximum.

Under Assumption 5, Proposition 4 yields

$$\pi_{\mathrm{rob}}(x;y) \le \overline{\epsilon}\, L \,\mathrm{diam}_{\|\cdot\|}(S(x)).$$

If one uses the ℓ_∞ norm, Remark 3 shows that $\mathrm{diam}_\infty(S(x))$ can be computed by solving at most $2m$ linear programs over $\Omega(x)$.

Proposition 5 *Fix $x \in X$ and $y \in S(x)$. Let $Q \succ 0$ be arbitrary and set $t := y$. Then $y_{Q,t}(x) = y$, and the robust minimal transfer becomes*

$$\pi_{\mathrm{rob}}(x;y) = \sup_{(\epsilon,\phi)\in\mathcal{E}\times\mathcal{P}} \epsilon \Big(\phi(y) - \phi^\star(x;\phi) \Big). \tag{33}$$

If $\mathcal{E} = [0, \overline{\epsilon}]$, then

$$\pi_{\mathrm{rob}}(x;y) = \overline{\epsilon} \cdot \sup_{\phi\in\mathcal{P}} \Big(\phi(y) - \phi^\star(x;\phi) \Big). \tag{34}$$

Proof Since $t = y \in S(x)$ and $Q \succ 0$, Corollary 1 yields $y_{Q,t}(x) = y$. Substitute into (31) to obtain (33). The specialization (34) follows from Proposition 2. □

6.5 Structured uncertainty sets with computable robust prices

Without structure on $\mathcal{P}$, the supremum in (33) may be hard to evaluate. We now show that for common structured classes, robust prices are computable.

Let $\mathcal{G} \subseteq \mathbb{R}^m$ be a nonempty compact set and define

$$\mathcal{P}_{\mathrm{lin}} := \{\phi_g(\cdot) = g^\top(\cdot) \ : \ g \in \mathcal{G}\}.$$

Then $\phi_g^\star(x) := \min\{g^\top u : u \in S(x)\}$ is finite and attained because $S(x)$ is a polyhedron and g is fixed.

Theorem 4 *Assume $\mathcal{P} = \mathcal{P}_{\mathrm{lin}}$ and $\mathcal{E} = [0, \overline{\epsilon}]$. For any $x \in X$ and any $y \in S(x)$,*

$$\pi_{\mathrm{rob}}(x;y) = \overline{\epsilon} \cdot \sup_{g\in\mathcal{G}} \Big(g^\top y - \min_{u\in S(x)} g^\top u \Big) = \overline{\epsilon} \cdot \sup_{g\in\mathcal{G}} \max_{u\in S(x)} g^\top (y-u). \tag{35}$$

If $\mathcal{G}$ is a compact polytope, the supremum is attained at an extreme point of $\mathcal{G}$.

Proof Plug ϕ_g into (34) and rewrite $g^\top y - \min_{u \in S(x)} g^\top u = \max_{u \in S(x)} g^\top (y - u)$. If $\mathcal{G}$ is a compact polytope, a linear functional attains its maximum on $\mathcal{G}$ at an extreme point. □

Remark 4 If $\mathcal{G}$ is a polytope given by vertices $g^1, \ldots, g^K$, then (35) reduces to a finite maximum:

$$\pi_{\mathrm{rob}}(x; y) = \overline{\epsilon} \cdot \max_{k=1,\ldots,K} \Big((g^k)^\top y - \min_{u \in S(x)} (g^k)^\top u \Big),$$

where each inner minimum is a linear program over $S(x)$. If $\mathcal{G}$ is a norm ball $\{g : \|g\|_* \le L\}$, then $\sup_{g \in \mathcal{G}} g^\top (y - u) = L\|y - u\|$ and (35) becomes a (possibly convex) farthest-point problem over $S(x)$.

6.6 The robust semi-cooperative bilevel problem

Definition 7 Under mechanism (ii), define the *robust semi-cooperative bilevel problem* as

$$z_{\mathrm{rsc}} := \min_{x \in X} \min_{y \in S(x)} \Big\{ C_L(x, y) + \pi_{\mathrm{rob}}(x; y) \Big\}, \tag{36}$$

where $\pi_{\mathrm{rob}}(x; y)$ is given by (33). Any pair $(x^{\mathrm{rsc}}, y^{\mathrm{rsc}})$ attaining (36) is called a *robust semi-cooperative stable solution*.

The contract-aware formulations (26) and (36) remain bilevel optimization problems. Thus, the proposed stabilization and pricing terms should not be interpreted as removing the intrinsic computational difficulty of optimizing over the leader decision and the follower optimality relation. The role of the contract terms is different: they remove ambiguity in the realized follower reaction once an intended tie has been specified. For each fixed x and chosen $y \in S(x)$, the anchored tie-break implements a single follower response, and the corresponding elicited or robust price can be evaluated by the optimization subproblems described above. Hence, the contract layer adds a computable price for stability, while the overall leader problem remains a bilevel/global optimization problem.

The robust counterpart is obtained by the same anchoring argument, with the elicited transfer replaced by the robust minimal transfer.

Theorem 5 *Let $(x^{\mathrm{rsc}}, y^{\mathrm{rsc}})$ be a robust semi-cooperative stable solution. Fix any $Q \succ 0$, set $t := y^{\mathrm{rsc}}$, $T := \pi_{\mathrm{rob}}(x^{\mathrm{rsc}}; y^{\mathrm{rsc}})$, and offer the contract (Q, t, T) at x^{rsc} . Then:*

1. *the accepted contract implements the unique tie-break reaction $y_{Q,t}(x^{\mathrm{rsc}}) = y^{\mathrm{rsc}}$;*
2. *the transfer T makes the contract robustly individually rational at x^{rsc};*
3. *the leader's realized net cost under the accepted contract is*

$$C_L(x^{\mathrm{rsc}}, y^{\mathrm{rsc}}) + T = C_L(x^{\mathrm{rsc}}, y^{\mathrm{rsc}}) + \pi_{\mathrm{rob}}(x^{\mathrm{rsc}}; y^{\mathrm{rsc}}),$$

which is the objective value attained by $(x^{\mathrm{rsc}}, y^{\mathrm{rsc}})$ *in (36).*

Proof Since $(x^{\mathrm{rsc}}, y^{\mathrm{rsc}})$ is a robust semi-cooperative stable solution, we have $y^{\mathrm{rsc}} \in S(x^{\mathrm{rsc}})$. By Corollary 1, the anchored choice $t = y^{\mathrm{rsc}}$ yields $y_{Q,t}(x^{\mathrm{rsc}}) = y^{\mathrm{rsc}}$ for every $Q \succ 0$. Uniqueness follows from the strict convexity of $q_{Q,t}$ on the closed convex set $S(x^{\mathrm{rsc}})$.

By the definition of the robust minimal transfer, $T = \pi_{\mathrm{rob}}(x^{\mathrm{rsc}}; y^{\mathrm{rsc}}) = \sup_{(\epsilon,\phi)\in\mathcal{E}\times\mathcal{P}} \epsilon\Big(\phi(y^{\mathrm{rsc}}) - \phi^\star(x^{\mathrm{rsc}}; \phi)\Big)$. Therefore, for every admissible type $(\epsilon, \phi) \in \mathcal{E} \times \mathcal{P}$, $T \geq \epsilon\Big(\phi(y^{\mathrm{rsc}}) - \phi^\star(x^{\mathrm{rsc}}; \phi)\Big)$ which is equivalent to $\epsilon\,\phi(y^{\mathrm{rsc}}) - T \leq \epsilon\,\phi^\star(x^{\mathrm{rsc}}; \phi)$. Hence, the contract is robustly individually rational.

Finally, since the accepted contract implements y^{rsc} and pays transfer T, the leader's realized net cost is $C_L(x^{\mathrm{rsc}}, y^{\mathrm{rsc}}) + T = C_L(x^{\mathrm{rsc}}, y^{\mathrm{rsc}}) + \pi_{\mathrm{rob}}(x^{\mathrm{rsc}}; y^{\mathrm{rsc}})$, which is the value attained by $(x^{\mathrm{rsc}}, y^{\mathrm{rsc}})$ in (36). □

6.7 Relation with optimistic and pessimistic benchmarks

The robust semi-cooperative value always dominates the optimistic value. Indeed, for every $x \in X$ and $y \in S(x)$,

$$\pi_{\mathrm{rob}}(x; y) \geq 0,$$

and therefore

$$C_L(x, y) + \pi_{\mathrm{rob}}(x; y) \geq C_L(x, y).$$

Taking the infimum over $y \in S(x)$ and then over $x \in X$ yields

$$z_{\mathrm{opt}} \leq z_{\mathrm{rsc}}.$$

In contrast, the upper bound $z_{\mathrm{rsc}} \leq z_{\mathrm{pess}}$, which holds in the elicited-preference case of Theorem 2, does not necessarily extend to the robust case. The reason is that the elicited upper bound uses a default tie $y^{\mathrm{def}}(x)$ satisfying

$$\phi(y^{\mathrm{def}}(x)) = \phi^\star(x),$$

so that the elicited price of implementing this tie is zero. Under robust acceptance, however, the transfer must compensate all admissible follower types $(\epsilon, \phi) \in \mathcal{E} \times \mathcal{P}$. Unless there exists a single tie $y \in S(x)$ that minimizes every admissible secondary preference $\phi \in \mathcal{P}$, the robust price may be strictly positive for every implementable tie. Consequently, one may have

$$z_{\mathrm{rsc}} > z_{\mathrm{pess}}.$$

This observation gives a practical limit of robust semi-cooperation. A robust contract is meaningful only if the robust transfer required to guarantee acceptance is not too large relative to the leader's exposure to uncontrolled tie-breaking. If the robust price exceeds this exposure, then the leader may prefer not to offer a robust semi-cooperation contract and instead accept the no-contract outcome, or use a less conservative mechanism based on elicited or partially elicited preferences.

Remark 5 Problem (36) is typically more conservative than its elicited-preference analogue, because it prices acceptance against the worst-case admissible follower type. Unlike in the elicited case, the robust semi-cooperative value need not be bounded above by z_{pess}. If the robust price needed to guarantee acceptance exceeds the leader's exposure to uncontrolled tie-breaking, then offering a robust contract may not be economically attractive. In that case, the leader may prefer a no-contract outcome, a less conservative elicited-preference contract, or a contract designed for a restricted uncertainty set. Structured sets (Section 6.5) make this conservatism computable and interpretable.

7 Leader willingness to pay

The instability exposure $\Delta_L(x)$ defined in (8) can also be interpreted as the leader's willingness to pay for controlling tie-breaking at a fixed leader decision x. Indeed, $\Delta_L(x) = \overline{C}(x) - \underline{C}(x)$ is the largest reduction in leader cost that can be obtained at x by moving from a leader-worst follower tie to a leader-best follower tie. Hence, a transfer $\pi \leq \Delta_L(x)$ is affordable in the sense that enforcing a leader-preferred tie costs no more than the leader's worst-case exposure to uncontrolled tie-breaking at the same x.

Theorem 6 *Fix $x \in X$ and let $y^O \in \arg\min_{y \in S(x)} C_L(x, y)$ be a leader-preferred follower tie at x. Suppose that the leader enforces y^O through an anchored contract, namely by choosing the tie-break target $t = y^O$, and let $\pi \geq 0$ be a transfer that guarantees acceptance under the chosen mechanism. If $\pi \leq \Delta_L(x)$, then choosing $T = \pi$ yields acceptance and*

$$C_L(x, y^O) + T \leq \overline{C}(x).$$

Proof Choosing $t = y^O \in S(x)$ implies $y_{Q,t}(x) = y^O$ for every $Q \succ 0$. Since π guarantees acceptance under the chosen mechanism, setting $T = \pi$ yields acceptance and implements y^O. Let $y^P \in \arg\max_{y \in S(x)} C_L(x, y)$. Then $\Delta_L(x) = C_L(x, y^P) - C_L(x, y^O)$. If $\pi \leq \Delta_L(x)$, then $C_L(x, y^O) + T = C_L(x, y^O) + \pi \leq C_L(x, y^P) = \overline{C}(x)$.

□

Remark 6 Theorem 2 shows that pricing acceptance can only worsen the optimistic benchmark and never performs worse than the pessimistic one. Moreover, increasing ϵ makes deviations from the follower's default secondary minimizers increasingly expensive to induce. Thus, the leader's behavior moves toward the no-contract/default-tie regime. When the default secondary minimizer is unfavorable to the leader, this default-tie regime may be close to the pessimistic benchmark, but the two notions are not identical in general.

8 An experimental illustration

This section provides a small computational illustration of the proposed framework. The purpose of the experiments is not to benchmark a new global bilevel solver, but to illustrate three effects: (i) the ambiguity created by multiple follower optima, (ii) the way in which a strictly convex tie-break makes the realized follower response single-valued, and (iii) the economic meaning of the elicited and robust prices of semi-cooperation.

All experiments were implemented in Python 3.12 on a workstation equipped with a 12^{th} Gen Intel(R) Core(TM) i7-1260P CPU (2.10GHz), 32GB RAM, Windows 11.

8.1 Experimental design and random instance generation

We generate $N = 30$ random instances. To make the experiments reproducible, we fix an initial seed s_0 and generate instance j with seed $s_j = s_0 + j$, $j = 1, \ldots, N$. Thus, the master seed (`20240502`) is the single initial seed that reproduces the entire sequence of generated instances, while the instance seed identifies the particular random draw used for each instance.

We consider $n = 2$ leader variables with

$$X = [0, 1]^2$$

and $m = 6$ follower variables written as $y = (u, v)$, where $u \in \mathbb{R}^2$ and $v \in \mathbb{R}^4$. For each leader decision $x \in X$, the follower feasible set is

$$Y(x) = \{(u, v) : \ u = Ax + a_0, \ \ 0 \le v \le h(x)\}, \qquad h(x) := g_0 + Gx.$$

The matrices and vectors are generated independently as follows:

$$A_{ij} \sim U[-0.5, 0.5], \qquad (a_0)_i \sim U[-0.2, 0.2],$$

$$G_{ij} \sim U[0, 1], \qquad (g_0)_i \sim U[0.2, 1].$$

Since $g_0 > 0$ componentwise and $G \ge 0$ componentwise, the upper bound $h(x)$ is positive on X.

The follower objective depends only on u:

$$\min \; f_u^\top u, \qquad f_v = 0.$$

Consequently, once x is fixed, the component $u = Ax + a_0$ is fixed and every $v \in [0, h(x)]$ is lower-level optimal. Hence

$$S(x) = \{(u, v) : \; u = Ax + a_0, \;\; 0 \le v \le h(x)\}.$$

This construction deliberately creates a non-singleton follower optimal set, so that the effect of tie-breaking can be observed directly.

The leader objective is affine:

$$C_L(x, y) = c^\top x + d_u^\top u + d_v^\top v,$$

where

$$c_i, (d_u)_i, (d_v)_i \sim U[-1, 1]$$

independently.

For the elicited secondary preference, we use the quadratic criterion

$$\phi(u, v) = \frac{1}{2}\|v - r\|_2^2,$$

where each component of r is sampled independently from $U[-0.5, 1.5]$. Thus, the default follower tie at a fixed x is the Euclidean projection

$$v^{\text{def}}(x) = \Pi_{[0,h(x)]}(r),$$

and the corresponding default follower response is

$$y^{\text{def}}(x) = \big(Ax + a_0, \; v^{\text{def}}(x)\big).$$

For robust pricing, we consider $K = 6$ quadratic secondary-preference scenarios

$$\phi^k(u, v) = \frac{1}{2}\|v - r^k\|_2^2, \qquad k = 1, \ldots, K,$$

where the vectors $r^1, \ldots, r^K$ are sampled independently, componentwise, from $U[-0.5, 1.5]$. Thus, the robust scenarios use independent centers; they are not all set equal to the elicited center r.

The scale parameters are drawn as

$$\epsilon \sim U[0.1, 1], \qquad \bar{\epsilon} \sim U[0.5, 1.5].$$

8.2 Follower-optimal set and fixed-x quantities

The special structure above makes all fixed-*x* follower quantities explicit. For a given *x*, the follower optimal set is the box

$$S(x) = \{(Ax + a_0, v) : 0 \le v \le h(x)\}.$$

Therefore, for any vector $\alpha \in \mathbb{R}^4$,

$$\min_{v\in[0,h(x)]} \alpha^\top v = \sum_{i:\alpha_i<0} \alpha_i h_i(x), \qquad \max_{v\in[0,h(x)]} \alpha^\top v = \sum_{i:\alpha_i>0} \alpha_i h_i(x).$$

In particular, the optimistic and pessimistic follower ties for the leader are obtained componentwise by setting v_i to either 0 or $h_i(x)$ according to the sign of $(d_v)_i$.

The elicited reservation value is also explicit. Here

$$[0, h(x)] := \{w \in \mathbb{R}^4 : 0 \le w_i \le h_i(x),\ i = 1, \ldots, 4\}$$

denotes the four-dimensional box of feasible values for the free follower component v. For a vector $r \in \mathbb{R}^4$, the distance from r to this box is the Euclidean distance from r to the closest feasible vector $w \in [0, h(x)]$, namely

$$\operatorname{dist}(r, [0, h(x)]) := \min_{w\in[0,h(x)]} \|r - w\|_2.$$

Equivalently,

$$\operatorname{dist}^2(r, [0, h(x)]) = \min_{w\in[0,h(x)]} \|r - w\|_2^2 = \|r - \Pi_{[0,h(x)]}(r)\|_2^2,$$

where $\Pi_{[0,h(x)]}(r)$ denotes the Euclidean projection of r onto $[0, h(x)]$. Therefore,

$$\phi^\star(x) = \frac{1}{2} \operatorname{dist}^2\big(r, [0, h(x)]\big),$$

and the default tie is $v^{\mathrm{def}}(x) = \Pi_{[0,h(x)]}(r)$. Similarly, for each robust scenario,

$$\phi^{k\star}(x) = \frac{1}{2} \operatorname{dist}^2\big(r^k, [0, h(x)]\big).$$

For a fixed intended tie $y = (Ax + a_0, v) \in S(x)$, the vector $v \in [0, h(x)]$ denotes the follower response that the leader intends to implement. It is therefore arbitrary within the follower-optimal box and should not be confused with the default tie $v^{\mathrm{def}}(x) = \Pi_{[0,h(x)]}(r)$. The elicited and robust prices are then

$$\pi_\phi(x;y) = \epsilon \left(\frac{1}{2}\|v - r\|_2^2 - \frac{1}{2}\operatorname{dist}^2(r, [0, h(x)]) \right),$$

and

$$\pi_{\mathrm{rob}}(x;y) = \bar{\epsilon} \max_{k=1,\ldots,K} \left(\frac{1}{2}\|v - r^k\|_2^2 - \frac{1}{2}\operatorname{dist}^2(r^k, [0, h(x)]) \right).$$

In the special case where the intended tie is the default no-contract tie, one has $v = v^{\mathrm{def}}(x)$.

8.3 Policies compared

We compare the following policies.

Optimistic benchmark. For each x, the optimistic follower tie minimizes the leader cost over $S(x)$:

$$\underline{C}(x) = \min_{y \in S(x)} C_L(x, y).$$

The optimistic value is

$$z_{\mathrm{opt}} = \inf_{x \in X} \underline{C}(x).$$

Pessimistic benchmark. For each x, the pessimistic follower tie maximizes the leader cost over $S(x)$:

$$\overline{C}(x) = \max_{y \in S(x)} C_L(x, y),$$

and

$$z_{\mathrm{pess}} = \inf_{x \in X} \overline{C}(x).$$

Default no-contract policy. The default policy assumes that no compensation is paid, i.e., $T = 0$, and that the follower breaks ties according to the elicited secondary criterion ϕ. Thus, for each x, the implemented response is

$$y^{\mathrm{def}}(x) = \big(Ax + a_0, \Pi_{[0,h(x)]}(r)\big),$$

and the leader minimizes

$$C_L\big(x, y^{\mathrm{def}}(x)\big)$$

over $x \in X$. The default no-contract policy, therefore, represents the leader's performance when the follower's default secondary tie-breaking rule is anticipated, no compensation is paid, and no semi-cooperation contract is offered.

Stabilized policy with fixed center. We also report a purely stabilizing policy based on a fixed tie-break center t and $Q = I$. The identity matrix is used here as a neutral baseline, corresponding to Euclidean projection after the generated variables have been placed on comparable scales. As discussed in Section 3, alternative diagonal or full positive definite choices of Q could be used to encode different relative weights or covariance structures in the follower variables.

The vector t is sampled once per instance, componentwise from $U[-0.5, 1.5]$, and is kept fixed across all leader decisions. For each x, the implemented follower response is the Euclidean projection of t onto $S(x)$:

$$y_{I,t}(x) = \arg\min_{y \in S(x)} \frac{1}{2}\|y - t\|_2^2.$$

Because $u = Ax + a_0$ is fixed in $S(x)$, this amounts to projecting the v-component of t onto the box $[0, h(x)]$.

This fixed-center stabilized policy is included as a diagnostic baseline showing the effect of imposing an exogenous single-valued tie-breaking rule. It is not meant to represent a rational compensation contract. In particular, if the sampled center is unfavorable for the leader, the leader would not normally pay to enforce it. The contract-aware policy below, by contrast, endogenizes the intended tie and the associated compensation.

Semi-cooperative contract-aware policy. The semi-cooperative value is computed as

$$z_{\mathrm{sc}} = \inf_{x \in X} \inf_{y \in S(x)} \{C_L(x, y) + \pi_\phi(x; y)\}.$$

For a fixed x, this inner problem selects the follower optimum that minimizes the leader cost plus the compensation required to make that tie acceptable to the follower. Hence, this policy differs from the optimistic benchmark whenever the leader-favorable tie is costly to induce under the follower's secondary preference.

8.4 Computation of the reported values

The optimistic and pessimistic values have a particularly simple structure. Since $\underline{C}(x)$ and $\overline{C}(x)$ are affine functions of x in this testbed, their minima over $X = [0, 1]^2$ are attained at the vertices of the leader box. Thus, z_{opt} and z_{pess} are computed by evaluating the four vertices of $\{0, 1\}^2$.

The same vertex argument does not automatically apply to the default, fixed-center stabilized, and semi-cooperative policies. These policies involve projections onto the box $[0, h(x)]$ and therefore lead to piecewise-defined objective functions of the two leader variables x_1 and x_2. For these latter policies, the follower response and compensation terms are first written in closed form, yielding explicit two-dimensional

objective functions over $x \in [0,1]^2$. These functions are then minimized numerically using a reproducible global-search procedure. Specifically, we use SciPy's implementation of differential evolution on the box $[0,1]^2$, with fixed random seeds, tolerance parameters `tol`=10^{-10} and `atol`=10^{-10}, population size parameter `popsize`=20, and maximum number of iterations equal to 1000. The best point returned by differential evolution is then locally refined, and we additionally run local refinements initialized at the four vertices of X. The best objective value among these candidate solutions is reported. Thus, the values reported for z_{sc}, default, and stabilized are not obtained merely by checking the four vertices of X.

For the affordability analysis, we first compute an optimistic leader decision $x^O \in \arg\min_{x \in X} \underline{C}(x)$ and a leader-preferred follower tie $y^O \in \arg\min_{y \in S(x^O)} C_L(x^O, y)$.

The intended tie y^O is implemented through anchoring, i.e., by setting the tie-break target $t := y^O$. By Corollary 1, this gives $y_{Q,t}(x^O) = y^O$ for every $Q \succ 0$. We then compare the elicited and robust prices $\pi_\phi(x^O; y^O)$, $\pi_{\mathrm{rob}}(x^O; y^O)$, with the leader's instability exposure $\Delta_L(x^O) = \overline{C}(x^O) - \underline{C}(x^O)$.

8.5 Leader values and affordability of stabilization

Table 1 reports the mean and standard deviation of the leader values over the $N = 30$ generated instances. Figure 1 complements these summaries with boxplots showing the distribution of the instance-level values.

Consistent with Theorem 2, the semi-cooperative value lies between the optimistic and pessimistic benchmarks in all generated instances. The average optimistic–pessimistic gap is $z_{\mathrm{pess}} - z_{\mathrm{opt}}$, whereas the average loss of the semi-cooperative policy relative to the optimistic benchmark is $z_{\mathrm{sc}} - z_{\mathrm{opt}}$. In the experiments, the ratio $\frac{z_{\mathrm{sc}} - z_{\mathrm{opt}}}{z_{\mathrm{pess}} - z_{\mathrm{opt}}}$ is approximately 28.6% when computed from the reported mean values. Equivalently, the semi-cooperative policy retains about 71.4% of the average improvement of the optimistic benchmark over the pessimistic benchmark, while making the selected tie implementable through an individually rational contract.

Table 2 focuses on buying the optimistic follower tie at an optimistic leader decision. For each instance, we take x^O and y^O as described above, implement y^O by anchoring, and compute the elicited and robust prices. Figure 2 reports the corresponding instance-level distributions.

The elicited price is below the leader's exposure in 96.7% of the instances. In those instances, the leader can enforce the optimistic tie at a cost no higher than the worst-case loss it faces from uncontrolled tie-breaking at the same leader decision. Under robust acceptance, the affordability rate is lower, equal to 60.0%. This is expected, because the robust transfer must guarantee acceptance for all admissible secondary-preference scenarios. In the remaining 40.0% of the instances, the robust transfer required to guarantee acceptance exceeds the leader's instability exposure

Table 1 Leader values on random instances

	z_{opt}	z_{sc}	z_{pess}	Default	Stabilized (fixed t)
mmean ± std	-2.107 ± 0.825	-1.301 ± 0.825	-0.976 ± 0.825	-1.301 ± 0.825	-0.976 ± 0.825

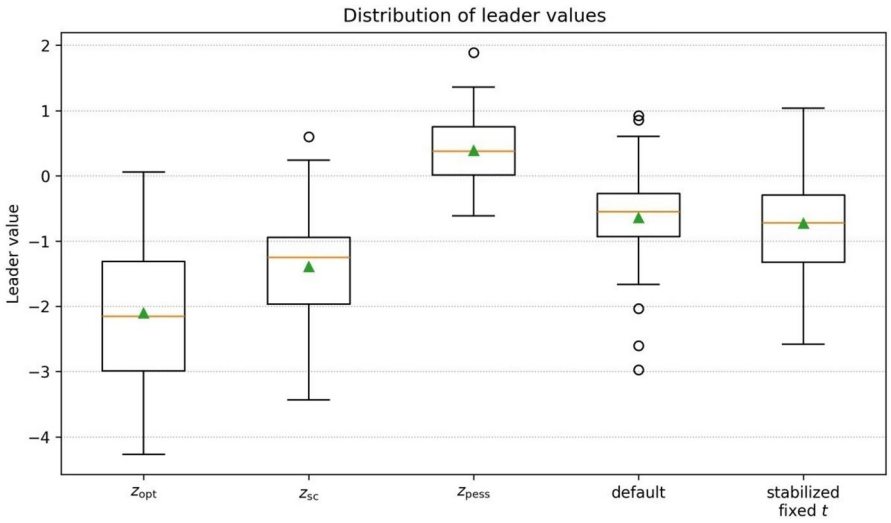


Fig. 1 Distribution of leader values over the $N = 30$ random instances. The boxplots complement the mean–standard deviation summary in Table 1 and show the variability of the optimistic, semi-cooperative, pessimistic, default, and fixed-center stabilized policies

Table 2 Exposure and stabilization prices at an optimistic leader decision x^O

	Exposure $\Delta_L(x^O)$	Elicited price $\pi_\phi(x^O; y^O)$	Robust price $\pi_{\mathrm{rob}}(x^O; y^O)$
Mean ± std	3.083 ± 1.167	1.047 ± 0.775	2.801 ± 1.514

Affordability rates: $\mathbb{P}[\pi_\phi(x^O; y^O) \le \Delta_L(x^O)] = 0.967$, $\mathbb{P}[\pi_{\mathrm{rob}}(x^O; y^O) \le \Delta_L(x^O)] = 0.600$

at x^O. In such cases, enforcing the optimistic tie robustly may not be economically attractive, and the leader may prefer either to accept the uncontrolled/pessimistic outcome or to choose a less favorable but cheaper tie.

8.6 Vanishing-regularization experiment at a fixed leader decision

Finally, we illustrate Theorem 1, which is a statement about the regularized follower problem for a fixed leader decision. For this reason, the last experiment is intentionally a fixed-x follower-level experiment rather than a full bilevel experiment.

We generate bounded polytopes $Y = \{y : Ay \le b\} \subset \mathbb{R}^4$ using random halfspaces together with box bounds, and use the degenerate linear objective $f = (1, 0, 0, 0)$. This creates a lower-level LP with multiple optima. For each instance, we draw a random positive definite matrix Q and a target point t, compute the exact tie-break minimizer $y_{Q,t} \in \arg\min_{y \in S} \frac{1}{2}(y - t)^\top Q(y - t)$, and compare it with the regularized solutions $y_\delta \in \arg\min_{y \in Y} \left\{ f^\top y + \delta \frac{1}{2}(y - t)^\top Q(y - t) \right\}$ for $\delta \in \{10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\}$. Table 3 reports the Euclidean error $\|y_\delta - y_{Q,t}\|_2$ over 30 random instances.

The errors decrease as δ becomes smaller, in line with Theorem 1. The convergence is not necessarily monotone instance by instance, because active-set changes and numerical tolerances may affect the computed solutions for very small δ.

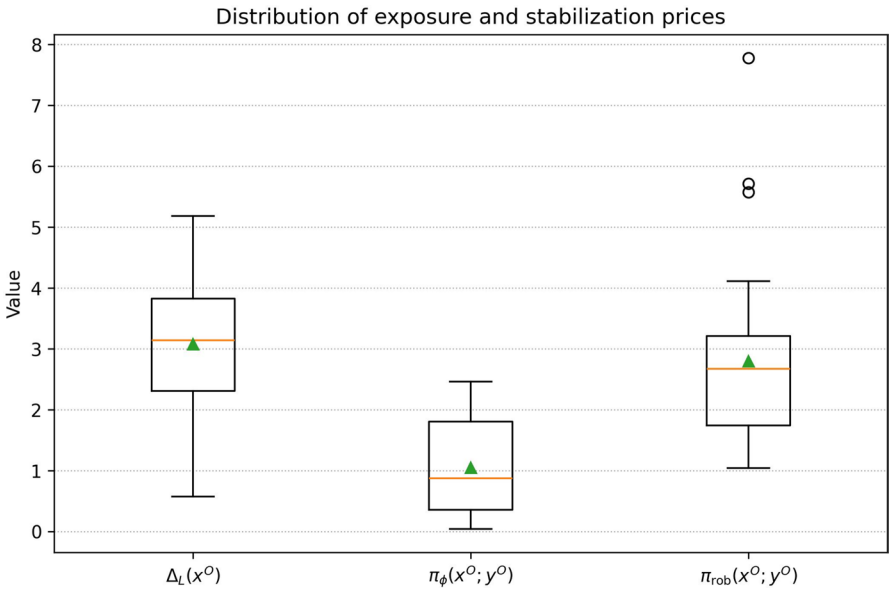


Fig. 2 Distribution of instability exposure and stabilization prices at an optimistic leader decision x^O over the $N = 30$ random instances. The figure compares the leader exposure $\Delta_L(x^O)$ with the elicited price $\pi_\phi(x^O; y^O)$ and the robust price $\pi_{\text{rob}}(x^O; y^O)$

Table 3 Empirical convergence of $y_\delta \to y_{Q,t}$ as $\delta \downarrow 0$.

δ	10^{-1}	10^{-2}	10^{-3}	10^{-4}
mean $\|y_\delta - y_{Q,t}\|_2$	$4.30 \cdot 10^{-1}$	$7.40 \cdot 10^{-2}$	$4.87 \cdot 10^{-2}$	$3.02 \cdot 10^{-2}$
$\|y_\delta - y_{Q,t}\|_2 < 10^{-6}$	0.267	0.333	0.367	0.367

We also monitored solution times in these experiments. Very small values of δ may worsen numerical conditioning because the quadratic term becomes weak relative to the linear follower objective. Consequently, decreasing δ improves the approximation of the exact lexicographic tie-break but can increase computational effort or make the solver more sensitive to tolerances, although in our small fixed-x experiments this effect was not pronounced. This supports the continuation strategy described in Section 3.2: decrease δ only until the primary optimality residual and the change in the implemented response are below the desired tolerances.

Overall, the experiments support the two main messages of the paper. First, strictly convex tie-breaking resolves lower-level ambiguity and makes the follower reaction single-valued without changing the follower's primary LP optimum in the exact tie-break model. Second, semi-cooperation prices provide a quantitative way to decide whether enforcing a leader-preferred tie is economically worthwhile, with robust prices reflecting the additional conservatism induced by uncertainty about the follower's secondary preference.

9 Conclusions

Unstable bilevel models naturally arise when the follower problem admits multiple optimal solutions: the leader's realized outcome then depends on an implicit tie-breaking rule, which leads to optimistic/pessimistic ambiguity and undermines predictive robustness.

Focusing on linear–linear bilevel programs, we proposed a stabilization mechanism that preserves the follower's primary linear optimal value while selecting a unique follower response through a strictly convex quadratic tie-break applied over the follower optimal set. We complemented this stabilization with an economic layer by interpreting enforced tie-breaking as a semi-cooperation contract. Under an elicited secondary preference, the price of semi-cooperation admits a transparent compensating-variation form over the follower optimal set. We also introduced a robust counterpart and showed that, for structured uncertainty classes, robust prices admit explicit reformulations and computable bounds.

From a computational perspective, we connected the exact tie-break to a vanishing regularization scheme and analyzed the corresponding regularized reaction map: for fixed regularization, the reaction is single-valued, Lipschitz, and piecewise affine over polyhedral regions, and it converges to the exact tie-break as the regularization vanishes. This also enables a practical single-level solution route via KKT-based MILP reformulation. Experiments have been presented to illustrate the theoretical framework further.

Future work may address several natural extensions of the present framework. First, a wider experimental campaign on different classes of linear bilevel instances would be useful to assess the computational behavior of the proposed stabilization and pricing mechanisms. Second, richer uncertainty models for secondary preferences could be studied, with particular attention to classes that preserve tractability of robust prices. Third, sharper perturbation and sensitivity results for exact lexicographic selections could be developed in broader convex bilevel settings. Finally, an important direction is the extension of the proposed contract-based stabilization approach to bilevel programs with upper-level coupling constraints involving the follower variable. In such models, tie-breaking affects not only the leader's objective value but also upper-level feasibility, so the present mechanism would have to be complemented by feasibility-aware selection and compensation rules.

Acknowledgements None.

Funding Open access funding provided by Università degli Studi di Roma Tor Vergata within the CRUI-CARE Agreement.

Data availability Data will be made available on reasonable request.

Declarations

Conflict of interest None.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License,

which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

1. Bard, J.F.: Convex two-level optimization. Math. Program. **40**, 15–27 (1988)
2. Beck, Y., Ljubić, I., Schmidt, M.: A survey on bilevel optimization under uncertainty. Eur. J. Oper. Res. **311**(2), 401–426 (2023)
3. Beck, Y., Ljubić, I., Schmidt, M.: Linear and Mixed-Integer Bilevel Optimization: Theory and Algorithms. Cambridge University Press (2026)
4. Beck, Y., Schmidt, M.: A robust approach for modeling limited observability in bilevel optimization. Oper. Res. Lett. **49**(5), 752–758 (2021)
5. Ben-Ayed, O.: Bilevel linear programming. Comput. Oper. Res. **20**(5), 485–501 (1993)
6. Benchouk, I., Jolaoso, L., Nachi, K., Zemkoho, A.B.: Relaxation methods for pessimistic bilevel optimization, (2025). arXiv:2412.11416 arXiv preprint
7. Besançon, M., Anjos, M.F., Brotcorne, L.: Robust bilevel optimization for near-optimal lower-level solutions. J. Global Optim. **90**, 813–842 (2024)
8. Caramia, M.: An algorithm to find stable solutions in linear-linear bilevel problems. Opsearch **61**(2), 972–988 (2024)
9. Caruso, F., Lignola, M.B., Morgan, J.: Regularization and approximation methods in Stackelberg games and bilevel optimization. In: Dempe, S., Zemkoho, A. (eds.) Bilevel Optimization, Springer Optimization and Its Applications, vol. 161. Springer, Cham (2020)
10. Colson, B., Marcotte, P., Savard, G.: Bilevel programming: a survey. 4OR **3**, 87–107 (2005)
11. Corpus, C., Camacho-Vallejo, J.-F.: A study on the impact of selecting the follower's reaction in solving semi-vectorial bilevel problems. Optim. Lett. **19**, 1347–1373 (2025)
12. Dempe, S.: Foundations of Bilevel Programming. Springer (2002)
13. Dempe, S., Gadhi, N., Zemkoho, A.B.: Bilevel programming: reformulations, constraint qualifications and optimality conditions. Math. Program. **138**, 447–473 (2013)
14. Dempe, S., Mordukhovich, B.S., Zemkoho, A.B.: Sensitivity analysis for two-level value functions with applications to bilevel programming. SIAM J. Optim. **22**(4), 1309–1343 (2012)
15. Dempe, S., Mordukhovich, B.S., Zemkoho, A.B.: Necessary optimality conditions in pessimistic bilevel programming. Optimization **63**(4), 505–533 (2014)
16. Dempe, S., Mordukhovich, B.S., Zemkoho, A.B.: Two-level value function approach to nonsmooth optimistic and pessimistic bilevel programs. Optimization **68**(2–3), 433–455 (2019)
17. Dempe, S., Kalashnikov, V., Pérez-Valdés, G.A., Kalashnykova, N.: Bilevel Programming Problems: Theory Algorithms and Applications to Energy Networks. Springer (2015)
18. Dontchev, A.L., Rockafellar, R.T.: Implicit Functions and Solution Mappings: A View from Variational Analysis. Springer (2009). https://doi.org/10.1007/978-0-387-87821-8
19. Fortuny-Amat, J., McCarl, B.: A representation and economic interpretation of a two-level programming problem. J. Oper. Res. Soc. **32**(9), 783–792 (1981). https://doi.org/10.1057/jors.1981.156
20. Hiriart-Urruty, J.B., Lemaréchal, C.: Convex Analysis and Minimization Algorithms I: Fundamentals. Springer (2013)
21. Kleinert, T., Labbé, M., Ljubić, I., Schmidt, M.: A survey on mixed-integer programming techniques in bilevel optimization. EURO J. Comput. Optim. **9**, 100007 (2021). https://doi.org/10.1016/j.ejco.2021.100007
22. Jia, S., Zheng, Y.: A new coordination model for ill-posed bilevel programming problem. Appl. Mech. Mater. **411–414**, 1943–1947 (2013). https://doi.org/10.4028/www.scientific.net/AMM.411-414.1943
23. Lignola, M.B., Morgan, J.: Stability of regularized bilevel programming problems. J. Optim. Theory Appl. **93**(3), 575–596 (1997)

24. Loridan,P., Morgan,J.: ε-regularized two-level optimization problems: Approximation and existence results. In: Dolecki, S. (eds) Optimization. *Lecture Notes in Mathematics*, vol. 1405, Springer, Berlin, Heidelberg, (1989)
25. Outrata, J.V.: Necessary optimality conditions for Stackelberg problems. J. Optim. Theory Appl. **76**, 305–320 (1993)
26. Outrata, J.V., Kočvara, M., Zowe, J.: Nonsmooth Approach to Optimization Problems with Equilibrium Constraints: Theory Applications and Numerical Results. Springer (1998)
27. Royset, J.O.: On Stability in Optimistic Bilevel Optimization, (2024). arXiv:2408.13323 arXiv preprint
28. Tøndel, P., Johansen, T.A., Bemporad, A.: An algorithm for multi-parametric quadratic programming and explicit MPC solutions. Automatica **39**(3), 489–497 (2003). https://doi.org/10.1016/S0005-1098(02)00250-9
29. White, D.J., Anandalingam, G.: A penalty function approach for solving bi-level linear programs. J. Global Optim. **3**, 397–419 (1993)
30. Vicente, L.N., Calamai, P.H.: Bilevel and multilevel programming: a bibliography review. J. Global Optim. **5**, 291–306 (1994)
31. von Stackelberg, H.: Marktform und Gleichgewicht. Springer (1934)
32. Zemkoho, A.B.: Solving ill-posed bilevel programs. Set-Valued Var. Anal. **24**(3), 423–448 (2016)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.