Technische Universität Braunschweig

Robust Optimization Under Controllable Uncertainty

Eva Ley[1], Rebecca Marx[1], Maximilian Merkert[1], Tim Niemann[1,2], and Sebastian Stiller[1,2]

[1]Technische Universität Braunschweig, Institute for Mathematical Optimization, Universitätsplatz 2, 38106 Braunschweig, Germany

[2]Cluster of Excellence SE^{2}A - Sustainable and Energy-Efficient Aviation, Hermann-Blenk-Str. 42, 38108 Braunschweig, Germany

March 7, 2024

Abstract Applications for optimization with uncertain data in practice often feature a possibility to reduce the uncertainty at a given *query cost*, e.g., by conducting measurements, surveys, or paying a third party in advance to limit the deviations. To model this type of applications we introduce the concept of *optimization problems under controllable uncertainty (OCU)*. For an OCU we assume the uncertain cost parameters to lie in bounded, closed intervals. The optimizer can shrink each of these intervals around a certain value (*hedging point*) possibly reducing it to a single point. Depending on whether the hedging points are known in advance or not, different types of OCU arise. Moreover, the models may differ with respect to when the narrowing down, the underlying optimization, and the revelation of true data take place.

We study two different problem settings - one with known and one with unknown hedging points - in more detail, where we handle the remaining uncertainty by the paradigm of robust optimization. For both settings, we provide bounds on the optimal objective value and a single-level non-linear reformulation. Furthermore, we state assumptions under which the three- respectively four-level problem can be solved as a single-level mixed-integer linear program. We also show that in robust OCU an optimizer might query a parameter solely to reduce the uncertainty for other parameters (*budget deflection*). We give sufficient conditions to avoid this phenomenon.

Keywords: Multi-Level Optimization, Controllable Uncertainty, Mixed-Integer Programming, Robust Optimization, Budgeted Uncertainty Set, Single-Level Reformulation

1 Introduction

1.1 Motivation

Optimization problems in applications often come with uncertainty in the data input. We propose and study a model for optimization under uncertainty in which one can pay to reduce some of the uncertainty while dealing with remaining uncertainty in a robust way.

In *optimization under controllable uncertainty (OCU)*, uncertain parameters are initially only known to lie within bounded, closed intervals that the optimizer can continuously shrink at a cost. We regard optimization problems with uncertain *underlying cost*. The optimizer can continuously shrink each of the initially given intervals for the uncertain parameters at a *query cost*. This may eventually but not necessarily reduce an interval to a single point. We call this point the *hedging point*, as it is the value we get when fully averting uncertainty. However, shrinking of the interval to a single point might not be possible. One major modeling choice in OCU is whether the hedging point is known or unknown. If the hedging point is known, i.e., part of the initial input, this does not mean that we know in advance which value the adversary will choose. However, we know in advance to which value we can force the adversary if we buy full information.

OCU has several applications, both with known and unknown hedging points. Management of currency risks [FWR12] or protection against damage in networks like electric power [BCSW06], supply chain [CS07] or transportation networks [JLSY15, FHEED22] come at some cost and reduce uncertainty in the underlying problems cost. The hedging points are known in these applications. In contrast, revenues are only revealed after some investment is made in research and development portfolio optimization [SCJB10], production planning [JWW98], pharmaceutical clinical trial planning [CM10] or offshore gas-field development [GG04].

In *robust optimization under controllable uncertainty (ROCU)*, we use the worst-case approach of *robust optimization* to deal with the remaining uncertainty in OCU. The possible scenarios for the malign adversary to choose from are restricted to a subset of the Cartesian product of intervals. In the present work, we use the so-called *budgeted uncertainty* set that is widely used in robust optimization [BS03]. The rationale underlying this uncertainty set is that it is over-conservative to protect against a worst case where all uncertain parameters are realized as the maximal value in their interval. Instead, we assume that the sum over all parameters of all relative deviations is limited by a budget parameter.

In ROCU, for elementwise query costs, an optimizer might query a parameter solely to control the uncertainty for other parameters. We call this phenomenon *budget deflection*. We give an example where budget deflection occurs in a problem setting with known hedging point and provide sufficient conditions to avoid it. For another problem setting with unknown hedging point, we show that budget deflection is not possible.

The concept of OCU can be combined with other approaches to deal with remaining uncertainty, e.g., stochastic optimization. In stochastic optimization, the uncertainty realizes according to a random distribution in stochastic optimization, e.g., [BL11]. Using the paradigm of robust optimization, we do not need to assume distributional knowledge.

OCU is an extension of existing models that allow to reduce uncertainty in data input before solving an optimization problem. In *explorable uncertainty*, the optimizer can buy exact information for individual uncertain parameters [Kah91]. Uncertainty for each parameter is either kept or fully erased until the actual optimization problem can be solved exactly in explorable uncertainty. Similarly, in *decision-dependent information discovery (DDID)*, either the exact value is revealed or the full uncertainty for a parameter remains [OP23]. Costs for exact information are not part of the overall objective. Instead, the optimizer has a fixed budget to get some individual uncertain parameters. The possibility to continuously shrink the intervals of uncertain data in OCU extends models in which one can either buy full information or keep the uncertainty as initially given. Furthermore, the query cost for additional information in OCU is an extension of a fixed budget for the reduction of uncertainty. In applications, it might be difficult or unrealistic to completely erase uncertainty in underlying cost and to fix a budget for investments.

In OCU, the optimizer solving the underlying problem and the optimizer choosing the queries is the same entity. In contrast, one could use different objectives for the two decisions of making the queries and solving the underlying problem. Then, we obtain a similar setting as bilevel problems with uncertainty in the follower's data, see [BLS23a] and references therein. Often, one distinguishes whether the uncertainty realizes between the leader's and the follower's decision *(wait-and-see follower)* or after the follower's decision *(here-and-now follower)*. In controllable uncertainty, it is not possible to distinguish beforehand which parts of the underlying problem's data uncertainties realize before and which after the problem is solved as this can be dependent on the queries made. A here-and-now follower who decides before the uncertainty realizes might be turned into a wait-and-see follower if the remaining uncertainty is completely removed.

Structure of the paper Our paper is structured as follows: In the remaining of this section we describe related work. Then, in Section 2, we explain the general concept of optimization with controllable uncertainty for uncertain cost. We describe how one can modify the scenario set at a query cost by reducing the uncertain intervals around the so-called hedging points. In Section 3, we assume that the hedging points are known in advance and part of the input data. For binary queries, we show how one can reformulate the problem to a single-level one, if the underlying optimization problem is given by a linear program (LP). In Section 4, we assume that the hedging points are not initially known and model them as a variable that is chosen in a worst-case fashion. We model this setting in a four-level problem. Further, we investigate a robust optimization approach, and present an equivalent nonlinear single-level formulation.

1.2 Related work

Optimization under controllable uncertainty is closely related to other concepts like bilevel optimization, robust optimization, explorable uncertainty, and decision-dependent information discovery. In this paper we investigate optimization problems with a multilevel structure. Problems with two levels are considered in bilevel optimization. In robust optimization, the uncertainty realizes after the decision has been made. In contrast, in the model of explorable uncertainty, the uncertainty is first handled before the underlying problem is solved exactly. In decision-dependent information discovery these opposite concepts are combined. First some parts of the uncertainty are handled before taking decisions, after which the remaining uncertainty realizes.

In the following, we give a brief overview of the aforementioned concepts. The literature review is far from exhaustive. Recent surveys for further reading are given where available.

Bilevel optimization In a bilevel optimization problem, two optimization problems are nested in a hierarchical order. Two players usually called *leader* and *follower* control disjoint sets of variables, who optimize their own objectives with constraints that both can depend on the other's decisions. Foundations on bilevel programming are explained in the textbook [Dem02] and further advances in bilevel optimization can be found in [DZ21] which includes an extensive bibliography in the last chapter. For mixed-integer bilevel programs, see also the survey [KLLS21] and references therein.

Connections between robust and bilevel optimization, in particular possible reformulations of problems in one setting to the other one and vice versa, are discussed in [GKST23].

Two types of uncertainty that have been considered for bilevel problems are data uncertainty and decision uncertainty, see [BLS23a] and references therein. In data uncertainty, there is an uncertainty about the follower's data that is either realized after the leader's but before the follower's decisions (wait-and-see follower) or after the follower's decision is fixed (here-and-now follower). In decision uncertainty, one or both levels hedge against the other level's decision that for example might not be optimal but only near-optimal due to limited resources. In contrast to controllable uncertainty, neither of the two players can influence the uncertainty in these approaches.

A special type of bilevel optimization problems are min-max problems, i.e., problems in which the two players share the same objective function though optimize in opposite directions. A prominent example are interdiction games. In interdiction games, the upper-level problem interdicts some lower-level elements such that the follower is inhibited as much as possible in pursuing their goal [SS20]. Interdiction games with a monotone Γ-robust follower have been considered in [BLS23b, BLS23c].

An extension of interdiction games are fortification games where a third level is added. In fortification games, some items can be defended before the opponent interdicts some of the remaining items. Binary fortification games can be solved with a decomposition approach [BCSW06]. A generalized solution method is to use a branch-and-cut algorithm with fortification cuts [LLM$^+$23].

Robust optimization In robust optimization, optimization problems with uncertain cost are considered where the scenario is chosen adversarially after the decision of the optimization problem has been fixed, see e.g., the textbook [BEN09].

A widely used scenario set is the so-called *budgeted uncertainty set* introduced in [BS03, BS04] that is restricted in two ways. For each decision variable the uncertain cost is restricted to an interval. Furthermore, there is a budget Γ for the sum of actual increases normalized by the interval sizes. This budgeted uncertainty set is a polytope whose number of vertices grows exponentially with Γ. Robust counterparts of discrete optimization problems with polynomial runtime are still solvable in polynomial time [BS03].

Explorable uncertainty In explorable uncertainty, firstly, the uncertainty is resolved by revealing precise data values at some cost such that, secondly, the underlying problem is solved with certainty to optimality. The seminal work was introduced in [Kah91]. Revealing precise data values at some investment or effort is referred to as a *query*. The goal is to minimize query cost while the underlying problem can still be solved exactly.

Studies investigate the concept of explorable uncertainty on basic combinatorial problems like Selection [GSS11] as well as classical discrete problems like Shortest Path [FMO+03], Minimum Spanning Tree [EHK+08, FMM17], knapsack [GGI+15], and matroids [Mei18, MS19]. The binary query selection revealing an exact value is extended to returning a refined uncertainty interval in [GSS11]. We refer to the survey [EH15] for a good research overview on explorable uncertainty.

The query selection is realized in an online or offline approach. The online query selection is an adaptive model where queries are selected sequentially and for each decision one can use the outcome of all previous value determinations [BHKR05, FMM17]. The offline query model requests a non-adaptive selection simultaneously choosing and revealing as many queries as required to ensure the existence of an exact solution of the underlying optimization problem [MS19, FMO+03].

Decision-Dependent information discovery Decision-dependent revelation of uncertain parameters has mainly been considered in stochastic optimization, see [VGY22] and references therein. Recently, this idea has been combined with robust optimization instead resulting in the problem of decision-dependent information discovery (DDID) [VGY22, PGDT22, OP23]. In DDID, the optimizer has a binary choice in the first step to reveal some exact values, i.e., uncertain values are only either revealed completely or kept uncertain. The chosen values realize in a worst-case manner. Afterwards, the nominal problem is solved in a robust approach for the remaining uncertainties. Thus, DDID can be formulated as a four-level min-max-min-max problem. For general polyhedral uncertainty sets, both exact algorithms and approximations have been proposed. Furthermore, due to a budget instead of a cost for made queries, it can only be beneficial to make additional queries and to exhaust the query budget. As a natural consequence, they assume that it is not possible to make all queries, e.g., [OP23, Assumption 1].

2 Controllable uncertainty for uncertain cost

In general, the concept of controllable uncertainty can be used for both uncertain cost and uncertainty feasible regions. In this paper we solely discuss the case of uncertain cost. For this case we now formalize the concept of controllable uncertainty and point variations in modeling with this concept. In Section 2.3 we summarize additional assumptions to which we restrict the analysis in the rest of this paper. Finally, in Section 2.4 we discuss a peculiar effect of controllable uncertainty with uncertain cost, namely, budget deflection.

2.1 Formalization of the concept

In the following, we introduce the concept of controllable uncertainty. First, we state the optimization problems for which we consider controllable uncertainty. Then, we describe the used model of uncertainty. Afterwards, we introduce queries and explain how they reduce uncertainty. In particular, this includes the definition of the controllable uncertainty set. Finally, we describe the overall objective function of the resulting problem. For a summary of the notation, see Table 1.

Controllable uncertainty is a possibility to model how one can deal with uncertainty in some *underlying (optimization) problem*

$$\min_{y \in Y} f(y).$$

We assume that the underlying problem has non-negative decision variables y chosen from a feasible set $Y \subseteq \mathbb{R}^n_{\geq 0}$. We refer to the indices of vectors like y as *elements* and denote them with e. We assume that the objective function f of the underlying problem is parameterized by non-negative uncertain coefficients $\tilde{c} = (\tilde{c}_1, \dots, \tilde{c}_n)$. These uncertain coefficients $\tilde{c}$ lie within bounded, closed intervals, i.e.,

$$\tilde{c}_e \in [c_e, c_e + d_e] \text{ for all } e \in [n] := \{1, \dots, n\} \text{ with } d \in \mathbb{R}^n_{\geq 0}.$$

Each realization of an uncertain parameter $\tilde{c}_e$ is associated with the corresponding normalized value $u_e \in [0, 1]$ such that $\tilde{c}_e = c_e + u_e d_e$.

A possible choice of u_e as well as the whole vector u is a *realization of the uncertainty* or *scenario*. We assume that the uncertainty set is given as some polyhedron $\mathcal{U} \subseteq [0, 1]^n$. To explicitly denote the dependency on the scenario, we also write $f(u, y)$ for the objective of the underlying problem.

The interval for the choice of u_e is narrowed, from $[0, 1]$ to at most a single point $b_e \in [0, 1]$. This point b_e is the *e-th hedging point*. The set of possible hedging points b is denoted by $\mathcal{B} \subseteq [0, 1]^n$.

The continuous variable x_e expresses how much the size of the interval for u_e is narrowed. We call x_e as well as the vector x as a whole a *query*. The set of possible queries is given by $X \subseteq \mathbb{R}^n_{\geq 0}$. We assume that $\mathbf{0} \in X$ and will refer to $x = \mathbf{0}$ as "making no query". Mostly, we think of X as $\mathbb{R}^n_{\geq 0}$. However, it can include constraints on the query selection. For example, a query for element e might only be allowed if also element e' is queried to at least the same amount, i.e., $x_e \leq x_{e'}$.

The *lower* and *upper query outcome* $\phi^\ell_e(x_e)$ respectively $\phi^u_e(x_e)$ shift the lower respectively upper boundary of the interval for the realization of the uncertainty u_e. We require the functions $\phi^\ell_e \colon \mathbb{R}_{\geq 0} \to [0, 1]$ and $\phi^u_e \colon \mathbb{R}_{\geq 0} \to [0, 1]$ to be monotone and to fulfill $\phi^\ell_e(0) = \phi^u_e(0) = 0$. The lower bound for the realization of the uncertainty u_e raises within the interval $[0, b_e]$ by the fraction of the lower query outcome $\phi^\ell_e(x_e)$. Similarly, the upper query outcome $\phi^u_e(x_e)$ is the fraction by which the upper bound on the realization of the uncertainty u_e is reduced within the interval $[b_e, 1]$. For a visualization, see Figure 1. More precisely, the query x_e narrows the interval $[0, 1]$ for u_e to the interval

$$\left[b_e \phi^\ell_e(x_e), 1 - (1 - b_e)\phi^u_e(x_e)\right].$$

If ϕ^ℓ_e and ϕ^u_e are strictly less than one for all queries x_e, a reduction of the interval for the choice of u_e to a singleton is not possible, see the example of asymptotic behavior below.

Example 2.1 (Query functions)**.** *The query outcome function ϕ_e can be*

(a) $\phi_e(x_e) = \min\{x_e, 1\}$ *for proportional outcome on $[0, 1]$ and constant else or*

(b) $\phi_e(x_e) = \frac{x_e}{x_e + 1}$ *for an asymptotic behavior or*

(c) $\phi_e(x_e) = 0$ *for $x_e < 1$ and $\phi_e(x_e) = 1$ for $x_e \geq 1$ for a binary query outcome. The same outcome is obtained for ϕ_e being the identity and restricting X such that $x_e \in \{0, 1\}$.*

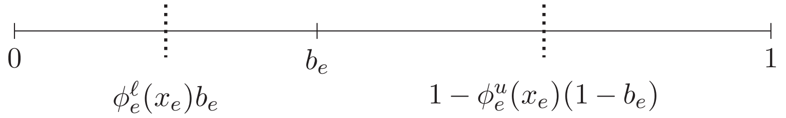

Figure 1: Reduced interval for the choice of u_e

For the sake of convenience, we use the following conventions for the query outcome functions. Whenever ϕ_e^ℓ and ϕ_e^u are equal, we refer to them with one function $\phi_e\colon \mathbb{R}_{\geq 0} \to [0,1]$. Also, we drop the index and only write ϕ^ℓ, ϕ^u or ϕ if the respective query outcome functions $\phi_e^\ell, \phi_e^u, \phi_e$ are equal for all elements $e \in [n]$.

Next, we define the *controllable uncertainty set* $\mathcal{U}(x,b)$ to capture the reduction of the uncertainty set $\mathcal{U}$. We call elements of $\mathcal{U}(x,b)$ the *remaining uncertainty. Adversarial feasibility* is the requirement that the controllable uncertainty set $\mathcal{U}(x,b)$ is non-empty.

Definition 2.2 (Controllable uncertainty set)**.** *For a polyhedral uncertainty set* $\mathcal{U}$*, the* controllable uncertainty set *with respect to query* x *and hedging point* b *is the set*

$$\mathcal{U}(x,b) = \left\{ u \in \mathcal{U} \;\middle|\; \phi_e^\ell(x_e)b_e \leq u_e \leq 1 - \phi_e^u(x_e)(1-b_e)\ \forall e \in [n] \right\}.$$

Observation 2.3. *For no or, respectively, complete reduction of the uncertainty, we have:*

(a) If $\phi_e^\ell(x_e) = \phi_e^u(x_e) = 0$ *for all* $e \in [n]$*, then* $\mathcal{U}(x,b) = \mathcal{U}$*.*

(b) $\mathcal{U}(\mathbf{0},b) = \mathcal{U}$*.*

(c) Let $\overline{x} \in X$ *be a query with* $\phi_e^\ell(\overline{x}_e) = \phi_e^u(\overline{x}_e) = 1$ *for all* $e \in [n]$*. Then* $\mathcal{U}(\overline{x},b) = \{b\} \cap \mathcal{U}$*.*

Proof. The statements follow from the definition of the controllable uncertainty set $\mathcal{U}(x,b)$ and the assumption that $\phi_e^\ell(0) = \phi_e^u(0) = 0$. ☐

If we make a query to narrow an uncertainty interval in our model, a *query cost* $q(x)$ are generated. We assume that the query cost function $q\colon X \to \mathbb{R}_{\geq 0}$ is monotone and that there is no query cost if no query is made, i.e., $q(\mathbf{0}) = 0$.

Observation 2.4. *There are the following relations between the query cost* q *and the query outcome function* ϕ*, if* $X = \mathbb{R}^n_{\geq 0}$*:*

(a) If q *is strongly monotone, it is equivalent to use either* q*,* ϕ_e^ℓ *and* ϕ_e^u *or the identity* Id*,* $q^{-1}\phi_e^\ell$ *and* $q^{-1}\phi_e^u$ *as query cost and query function respectively.*

(b) If $\phi = \phi_e^\ell = \phi_e^u$ *is strongly monotone, it is equivalent to use either* q *and* ϕ *or* $\phi^{-1}q$ *and* Id *as query cost and query function respectively.*

Proof. If there is a bijection $g\colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ with $g(0) = 0$, we can get another instance of our problem if we apply g and replace x with $\tilde{x} = g(x)$ in all occurrences. If q or respectively ϕ is strictly monotone, the inverse q^{-1} or ϕ^{-1} exist. If $g = q^{-1}$, then $\phi_e^\ell(\tilde{x}) = \phi_e^\ell(q^{-1}(x))$, $\phi_e^u(\tilde{x}) = \phi_e^u(q^{-1}(x))$ and $q(\tilde{x}) = x$. The same argument applies if $g = \phi^{-1}$. ☐

$\min_{y \in Y} f(y)$	underlying problem
$\tilde{c}_e = c_e + u_e d_e$	uncertain coefficients in $f(y)$
$u \in \mathcal{U} \subseteq [0,1]^n$	realization of the uncertainty
$x \in X \subseteq \mathbb{R}^n_{\geq 0}$	query
$q \colon X \to \mathbb{R}_{\geq 0}$	query cost
$\phi^\ell_e, \phi^u_e \colon \mathbb{R}_{\geq 0} \to [0,1]$	lower, upper query outcome
$b \in \mathcal{B} \subseteq [0,1]^n$	hedging point
$\mathcal{U}(x,b) \subseteq \mathcal{U}$	controllable uncertainty set
$F \colon X \times \mathcal{U} \times Y \to \mathbb{R}$	overall objective function

Table 1: Summary of notation

In optimization under controllable uncertainty, a single optimizer allocates resources to both uncertainty mitigation and the core optimization problem. The overall objective function $F(x,u,y)$ is the sum of the query cost $q(x)$ and the objective function $f(u,y)$ of the underlying problem, i.e.,

$$F(x,u,y) = q(x) + f(u,y).$$

2.2 Variations in modeling with controllable uncertainty

In the following, we outline possibilities to use controllable uncertainty. One can vary when the hedging point is chosen, how and when the remaining uncertainty is dealt with or whether all queries are made at once or successively.

In the following, we assume that all queries are chosen in a single, first step. This is the analogue to offline queries in explorable uncertainty. Successive queries like online queries in explorable uncertainty are not considered here.

For the remaining uncertainty, we use a robust approach, i.e., we consider a worst-case scenario. Thus, we obtain min-max settings. Other approaches for dealing with the remaining uncertainty like stochastic optimization are possible though will not be considered here.

In the two possible problem settings with known hedging points, the uncertainty either realizes before or after the underlying problem is solved. If the underlying problem is solved before the remaining uncertainty realizes, the decisions for both the queries and the underlying problem can be made in the same step. Hence, we obtain a robust optimization problem with a dependent uncertainty set. This setting is not further considered. We consider the other setting in which the remaining uncertainty realizes before the underlying problem is solved in Section 3.

Afterwards, in Section 4, we consider a setting with uncertain hedging points. We will deal with the uncertainty in the hedging points before and with the remaining uncertainty after solving the underlying problem.

2.3 Additional assumptions

For all findings in the remaining of this paper, we add the following two assumptions.

Assumption 2.5. *We assume that the underlying problem has a linear objective function, i.e.,*

$$f(y) = f(u, y) = \tilde{c}^\top y = (c + d \cdot u)^\top y.$$

For example, the underlying problem can be a linear problem (LP) like the diet problem or a discrete problem like shortest path, min cut or TSP. In later sections, we derive some results that only hold for binary problems or problems that can be formulated as an LP.

As uncertainty set, we will only consider the budgeted uncertainty set. This uncertainty set has been introduced in [BS03] and is widely used in robust optimization.

Assumption 2.6. *We assume that $\mathcal{U}$ is the budgeted uncertainty set*

$$\mathcal{U} = \left\{ u \in [0,1]^n \;\middle|\; \sum_{e \in [n]} u_e \leq \Gamma \right\}.$$

In order to have adversarial feasibility, i.e., $\mathcal{U}(x, b) \neq \emptyset$, we derive a necessary condition for the choice of a hedging point.

Observation 2.7 (Necessary condition for adversarial feasibility)**.** $\mathcal{U}(x, b) = \emptyset$ *if*

$$\Gamma - \sum_{e \in [n]} \phi_e^\ell(x_e) b_e < 0.$$

Proof. The statement follows when we combine the lower bounds $\phi_e^\ell(x_e) b_e \leq u_e$ for all $e \in [n]$ from Definition 2.2 and the budget constraint $\sum_{e \in [n]} u_e \leq \Gamma$ from Assumption 2.6. □

Thus, we assume in the following that

$$\mathcal{B} \subseteq \left\{ b \in [0,1]^n \;\middle|\; \Gamma - \sum_{e \in [n]} \phi_e^\ell(x_e) b_e \geq 0 \right\}.$$

2.4 Budget deflection for elementwise query cost

In the following, we define the phenomenon of *budget deflection* for *elementwise query cost*.

Definition 2.8 (Elementwise query cost)**.** *The query cost $q(x)$ is elementwise, if there are functions $q_e \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that*

$$q(x) = \sum_{e \in [n]} q_e(x_e).$$

In controllable uncertainty, a query serves the purpose to gain information or to protect against very unwelcome realizations of the uncertainty at a certain cost. If the query cost is elementwise, one could assume that the query of an element e directly corresponds to controlling the uncertainty u_e. However, in an optimal solution, an element e might be queried though never used in any optimal solution of the underlying problem to control the uncertainty of cost for another element e'. The query for element e decreases the uncertainty budget at a relatively cheap price. Due to the reduced uncertainty budget, the cost for element e' is decreased. We call this phenomenon *budget deflection* and give a more formal definition.

Definition 2.9 (Budget deflection)**.** *Let $x^* \in X = \mathbb{R}^n_{\geq 0}$ be an optimal query for an optimization problem with controllable uncertainty and elementwise query cost. Furthermore, let $Y^*(x^*)$ denote the feasible solutions for the underlying problem that can be optimal when x^* has been fixed.*

We say an instances allows for budget deflection, *if there is an element $e \in [n]$ with positive query cost, i.e., $q_e(x^*_e) > 0$, that is not used in any optimal solution $y^* \in Y^*(x^*)$ of the underlying problem, i.e., $y^*_e = 0$ for all $y^* \in Y^*(x^*)$.*

We will show that budget deflection can occur in the setting with known hedging points described in Section 3. After we provide a small example, we will show how to prevent budget deflection and the impact of this adaption, see Section 3.3. In contrast, in Section 4.3, we will show that there is no budget deflection in the setting with unknown hedging points considered in Section 4.

3 Optimization with known hedging points

We consider the problem with known hedging points (KHP) to determine an optimal query x such that the underlying problem is minimized for the worst-case outcome of the uncertainty u. The values of the hedging point b are input parameters of the problem

$$\min_{x \in X} \max_{u \in \mathcal{U}(x)} \min_{y \in Y} F(x, u, y). \tag{KHP}$$

We use the min-max form to indicate the objective sense for each level. This does not imply that an optimal solution is attained. In fact, an optimal query x might not exist. For example, consider ϕ to be the asymptotic query function suggested in Example 2.1 in combination with zero query cost, i.e., $q = 0$. If there are no restrictions on the possible queries, i.e., $X = \mathbb{R}_{\geq 0}$, for every query x there exists another query x' that results in a smaller objective value. However, as usually done in bilevel optimization, we use "min" in the sense of "minimize" instead of "inf".

Since the hedging points b are input parameters of KHP, we write the uncertainty set only in dependence of the query x, i.e., $\mathcal{U}(x) = \mathcal{U}(x, b)$.

First, we derive bounds on the optimal objective value of KHP in Section 3.1. Afterwards, in Section 3.2, we provide an equivalent single-level reformulation for KHP if the underlying problem is given as a linear program (LP). In Section 3.3, we describe budget deflection, which is a phenomenon that in order to reduce the adversary weight modification for an element e, a different element e' is queried. Finally, after we add additional assumptions on the query cost and query outcome functions, we show how KHP can be formulated as single-level mixed-integer program.

Comparison of KHP with interdiction and fortification games In the following, we compare KHP with interdiction and fortification problems. For a fixed query x, KHP becomes a continuous interdiction problem where interdiction only affects the objective. The uncertainty u is chosen from intervals within a budget of Γ to maximize the minimal outcome of the underlying problem.

Next, we argue how the choice of a query x in KHP resembles the uppermost level of fortification games. In fortification, the uppermost level decides which elements to protect such that they cannot be interdicted. A fortification can be used on the uppermost level to prevent interdiction that realizes in a worst-case manner. In KHP, a query x indicates the reduction of the interval sizes for the weight modification caused by the uncertainty u. For binary query outcome functions, see Example 2.1, KHP is a fortification game with binary fortification and continuous interdiction.

3.1 Bounds for KHP

In the following, we show upper and lower bounds on the optimal objective value of KHP. We fix some particular query and then solve the underlying optimization problem. In general, even when a query is fixed, the resulting bilevel problem cannot be easily solved.

For the first upper bound, we consider that no query is made, i.e., $x = \mathbf{0}$. Recall that if no query is made, there is no query cost. We relax the budget constraint for the uncertainty and use that the upper bounds for the uncertainty are at most one. Then, only the underlying problem with objective $c + d$ remains. Thus, we obtain the following upper bound.

Observation 3.1. *The optimal objective value $F(x^*, u^*, y^*)$ of KHP has the upper bound*

$$F(x^*, u^*, y^*) \leq \min_{y \in Y} (c + d)^\top y.$$

Proof. We use the query $\mathbf{0} \in X$, the definition of $F(x, u, y)$, the assumptions that $q(\mathbf{0}) = 0$, $\mathcal{U}(x) \subseteq [0, 1]^n$ and that c, d, y are non-negative, see Section 2, to obtain

$$\begin{aligned} F(x^*, u^*, y^*) &= \min_{x \in X} \max_{u \in \mathcal{U}(x)} \min_{y \in Y} F(x, u, y) \\ &\leq \max_{u \in \mathcal{U}(\mathbf{0})} \min_{y \in Y} F(\mathbf{0}, u, y) = \max_{u \in \mathcal{U}(0)} \min_{y \in Y} (c + d \cdot u)^\top y \\ &\leq \max_{u \in [0,1]^n} \min_{y \in Y} (c + d \cdot u)^\top y \\ &= \min_{y \in Y} (c + d \cdot \mathbf{1})^\top y = \min_{y \in Y} (c + d)^\top y. \end{aligned}$$

□

For the second upper bound and a lower bound on the optimal objective value of KHP, let $\overline{x}$ denote a query for which the uncertainty set reduces to a singleton. The existence of such a query depends on the allowed queries X and the query outcome functions ϕ_e^ℓ and ϕ_e^u. If such a query $\overline{x}$ exists and we plug this in, only the underlying problem remains. Depending on whether we consider the query cost for $\overline{x}$ or not, we obtain an upper respectively lower bound.

Observation 3.2. *Let $\overline{x} \in X$ be a query such that no uncertainty is left for this query, i.e., $\phi_e^\ell(\overline{x}_e) = \phi_e^u(\overline{x}_e) = 1$ for all $e \in [n]$. Then, for the optimal objective value $F(x^*, u^*, y^*)$ of KHP, the following lower and upper bounds hold:*

$$\min_{y \in Y} f(b, y) \leq F(x^*, u^*, y^*) \leq q(\overline{x}) + \min_{y \in Y} f(b, y).$$

Proof. We choose $\overline{x} \in X$ and use $\mathcal{U}(\overline{x}) = \{b\}$, see Observation 2.3, to obtain the upper bound:

$$\begin{aligned} F(x^*, u^*, y^*) &= \min_{x \in X} \max_{u \in \mathcal{U}(x)} \min_{y \in Y} q(x) + f(u, y) \\ &\leq q(\overline{x}) + \max_{u \in \mathcal{U}(\overline{x})} \min_{y \in Y} f(u, y) = q(\overline{x}) + \min_{y \in Y} f(b, y). \end{aligned}$$

For the lower bound, we use that query cost $q(x)$ is non-negative and that $\mathcal{U}(\overline{x}) = \{b\}$. If there is no query cost, $\overline{x}$ is optimal for the outer minimization. We have

$$\begin{aligned} F(x^*, u^*, y^*) &= \min_{x\in X} \max_{u\in\mathcal{U}(x)} \min_{y\in Y} f(u,y) + q(x) \\ &\geq \min_{x\in X} \max_{u\in\mathcal{U}(x)} \min_{y\in Y} f(u,y) = \max_{u\in\mathcal{U}(\overline{x})} \min_{y\in Y} f(u,y) = \min_{y\in Y} f(b,y). \end{aligned}$$

□

3.2 Single-level reformulation for KHP

The main result of this section is that KHP can be reformulated as an equivalent single-level non-linear problem (NLP) if the underlying problem is an LP. We call two optimization problems *equivalent*, if they depend on the same parameters and always have the same optimal objective value. In the single-level reformulation, variables u for the realization of the uncertainty are replaced by dual variables for the constraints on the realization of the uncertainty within $\mathcal{U}(x)$. Afterwards, we state conditions on the values of variables in the single-level reformulation that hold for optimal solutions.

Assumption 3.3. *For the remaining of Section 3, we assume that the underlying problem can be solved via a nonempty compact, convex feasible set $Y \subseteq \mathbb{R}^n_{\geq 0}$.*

The underlying problem is for example an LP or a discrete problem given by a totally dual integral (TDI) system. Recall, that the objective of the underlying problem is linear by Assumption 2.5. Discrete problems given by a TDI system can be solved as LPs. TDI systems are known for several optimization problems like Shortest Path, Minimum Spanning Tree, Maximum Flow, and Minimum Cut, e.g., [KV18].

Theorem 3.4 (Single-level NLP)**.** *An optimization problem under controllable uncertainty with known hedging points (KHP)*

$$\min_{x\in X} \max_{u\in\mathcal{U}(x)} \min_{y\in Y} q(x) + c^\top y + (d\cdot u)^\top y \tag{1}$$

is equivalent to the following single-level non-linear problem (NLP)

$$\begin{aligned} \min_{\substack{x\in X,\\ y\in Y,\\ \beta,\theta}} \quad & \Gamma\theta + q(x) + \sum_{e\in[n]} \beta_e + c_e y_e - \phi^\ell_e(x_e) b_e \left(\theta + \beta_e - d_e y_e\right) - \phi^u_e(x_e)\beta_e(1-b_e) \\ \text{s.t.} \quad & \theta + \beta_e - d_e y_e \geq 0 \qquad \forall e\in[n] \\ & \beta,\theta \geq 0. \end{aligned} \tag{2}$$

Proof. First, we argue that we can interchange the innermost minimization and maximization step. The objective function F is linear in both variables u and y if all respective other variables are fixed. Furthermore, the feasible sets for u and y are polytopes by assumption and independent of y, u respectively. Thus, we can apply von Neumann's minimax theorem, see [vN28] or in english e.g., [Sim09, Theorem 2], to obtain

$$\min_{x\in X} \max_{u\in\mathcal{U}(x)} \min_{y\in Y} F(x,u,y) = \min_{x\in X,\, y\in Y} \max_{u\in\mathcal{U}(x)} F(x,u,y). \tag{3}$$

If the underlying problem is an LP, we can also obtain this equality using strong duality twice.

Next, we use strong duality to replace the inner maximization problem by a minimization problem. The inner maximization problem is the LP

$$\begin{aligned} \max_{u} \quad & (d \cdot y)^\top u && (4)\\ \text{s.t.} \quad & u_e \geq \phi_e^\ell(x_e) b_e && \forall e \in [n]\\ & u_e \leq 1 - \phi_e^u(x_e)(1-b_e) && \forall e \in [n]\\ & u^\top \mathbf{1} \leq \Gamma \end{aligned}$$

with dual problem

$$\begin{aligned} \min_{\alpha,\beta,\theta} \quad & \Gamma\theta + \sum_{e\in[n]} \phi_e^\ell(x_e) b_e \alpha_e + (1-\phi_e^u(x_e)(1-b_e))\beta_e && (5)\\ \text{s.t.} \quad & \alpha_e + \beta_e + \theta = d_e y_e && \forall e \in [n]\\ & \alpha \leq 0\\ & \beta, \theta \geq 0. \end{aligned}$$

We substitute $\alpha_e = d_e y_e - \beta_e - \theta \leq 0$ which simplifies Problem (5) to

$$\begin{aligned} \min_{\beta,\theta} \quad & \Gamma\theta + \sum_{e\in[n]} \phi_e^\ell(x_e) b_e (d_e y_e - \beta_e - \theta) + (1-\phi_e^u(x_e)(1-b_e))\beta_e && (6)\\ \text{s.t.} \quad & \theta + \beta_e - d_e y_e \geq 0 && \forall e \in [n]\\ & \beta, \theta \geq 0. \end{aligned}$$

Due to strong duality, Problems (4), (5) and (6) have the same optimal objective value. Thus, we can replace the inner maximization problem in (3) by Problem (6). We obtain the single-level reformulation (2) that is equivalent to KHP. □

For optimal solutions of this equivalent single-level NLP, we make the following observations. The variables y in (1) and (2) are not the same, for example if the underlying problem is given by an TDI system, the values of the variables y in (2) are not necessarily integral in an optimal solution.

Next, we consider elements e that are not used in the underlying problem's optimal solution.

Lemma 3.5. *In an optimal solution* $(x^*, y^*, \beta^*, \theta^*)$ *of Problem* (2)*, if* $y_e^* = 0$ *then*

$$\beta_e^* \left(1 - \phi_e^u(x_e^*)(1-b_e) - \phi_e^\ell(x_e^*) b_e\right) = 0.$$

Proof. If $y_e^* = 0$, the constraint $d_e y_e \leq \theta + \beta_e$ simplifies to $0 \leq \beta_e + \theta$. This is always fulfilled as β and θ are required to be nonnegative. There are no further constraints on β.

For the objective of Problem (2), we have

$$\begin{aligned} & \Gamma\theta + q(x) + \sum_{e\in[n]} \beta_e + c_e y_e + \phi_e^\ell(x_e) b_e (d_e y_e - \beta_e - \theta) - \phi_e^u(x_e)\beta_e(1-b_e)\\ = & \left(\Gamma - \sum_{e\in[n]} \phi_e^\ell(x_e) b_e\right)\theta + q(x)\\ & + \sum_{e\in[n]} c_e y_e + \phi_e^\ell(x_e) b_e d_e y_e + \beta_e \left(1 - \phi_e^u(x_e)(1-b_e) - \phi_e^\ell(x_e) b_e\right). \end{aligned}$$

Variable β_e is included with the factor $\left(1 - \phi_e^u(x_e)(1-b_e) - \phi_e^\ell(x_e)b_e\right)$. This non-negative factor is the size of the remaining interval the adversary chooses u_e from, see Section 2. Thus, in an optimal solution, at least one of the two factors β_e^* and $1 - \phi_e^u(x_e^*)(1-b_e) - \phi_e^\ell(x_e^*)b_e$ is zero. □

In the following lemma, we provide upper bounds for the variables β and θ in an optimal solution of Problem (2) if the underlying problem's decisions y are within the unit cube $[0,1]^n$. For example, let the underlying problem be binary.

Lemma 3.6. *Let $Y \subseteq [0,1]^n$. If $(x^*, y^*, \beta^*, \theta^*)$ is an optimal solution for the single-level reformulation (2), then*

$$\theta^* \leq \max_{e \in [n]} d_e =: D \quad \textit{and} \quad \beta_e^* \leq d_e \ \forall e \in [n].$$

Proof. Consider the constraints $d_e y_e - \beta_e \leq \theta$ for all $e \in [n]$. Since β_e and θ are required to be nonnegative and are included with nonnegative factors in the objective that is minimized, we have $\theta^* \leq \max_e d_e =: D$ and $\beta_e^* \leq d_e$. □

Next, we show upper bounds on elementwise query cost in an optimal solution. The upper bounds depend on whether the corresponding decision variable y_e of the underlying problem has a non-zero value or not. We describe some intuition before we state the lemma and give a proof.

First, consider an element e where the corresponding decision variable y_e is zero in an optimal solution. Then, the cost for the query of e is at most the potential harm the adversary can add when they use their budget elsewhere. A query x_e reduces the adversary's budget by the lower bound for the corresponding realization u_e. This lower bound is $\phi_e^\ell(x_e)b_e$ by definition of the controllable uncertainty set, see Definition 2.2. Recall that u_e is multiplied with d_e in the objective and D denotes the maximal value of $d_{e'}$ for all $e' \in [n]$. Thus, $\phi_e^\ell(x_e)b_e D$ is an upper bound on the query cost $q_e(x_e)$ for element e in an optimal solution.

For an element with non-zero value for the decision variable in the optimal underlying solution, additional cost has to be added to the possible query cost. This cost results from the reduction of the worst remaining realization. More formally, we have the following bounds.

Proposition 3.7. *Let the query cost be elementwise defined, i.e., assume that $q(x) = \sum_{e \in [n]} q_e(x_e)$ with functions $q_e \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Let (x^*, u^*, y^*) be an optimal solution of KHP.*

i) If $y_e^* = 0$, then there is an optimal solution with

$$q_e(x_e^*) \leq \phi_e^\ell(x_e^*)b_e D.$$

If in addition $1 - \phi_e^u(x_e^*)(1-b_e) - \phi_e^\ell(x_e^*) > 0$, i.e., the interval for u_e is not a singleton, the upper bound on $q_e(x_e^*)$ holds for every optimal solution.

ii) If $y_e^* = 1$ and $y \in [0,1]^n$, then

$$q_e(x_e^*) \leq \phi_e^\ell(x_e^*)b_e D + \phi_e^u(x_e^*)d_e(1-b_e).$$

Proof. The objective of the single level reformulation derived in Theorem 3.4 is

$$\Gamma\theta + \sum_{e \in [n]} c_e y_e + q_e(x_e) + \beta_e - \phi_e^\ell(x_e)b_e(\theta + \beta_e - d_e y_e) - \phi_e^u(x_e)\beta_e(1-b_e).$$

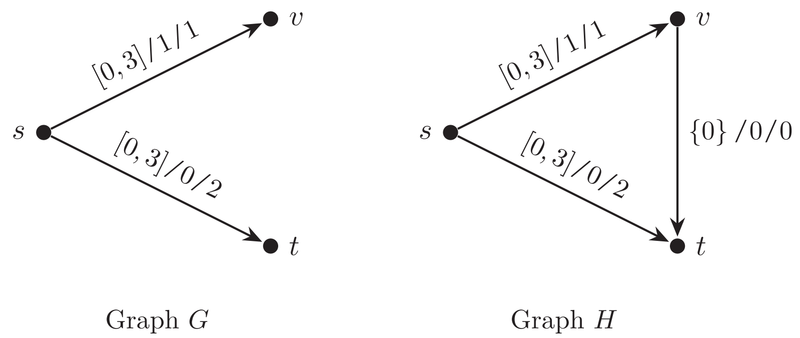

Figure 2: Example for budget deflection: Digraphs G and H with "$[c_e, c_e + d_e]$ / b_e / q_e" along arc.

For $x_e = 0$, the parts of the objective that depend on e simplify to $c_e y_e + \beta_e$. As there are no constraints on x_e and the goal is to minimize, we obtain that in an optimal solution this part is no larger than $c_e y_e + \beta_e$. Thus, we use the upper bound $\theta^* \leq D$ from Lemma 3.6 and get

$$q_e(x_e^*) \leq \phi_e^\ell(x_e^*) b_e \left(D + \beta_e^* - d_e y_e^*\right) + \phi_e^u(x_e^*) \beta_e^* \left(1 - b_e\right).$$

If $y_e^* = 0$, we plug in $y_e^* = 0$ and can set $\beta_e^* = 0$ according to Lemma 3.5 to derive the given upper bound on $q_e(x_e^*)$. If in addition $1 - \phi_e^u(x_e^*)(1 - b_e) - \phi_e^\ell(x_e^*) > 0$, Lemma 3.5 states that $\beta_e^* = 0$ in every optimal solution. Thus, in this case, the upper bound holds in every optimal solution.

If $y_e^* = 1$, we use the upper bound $\beta_e^* \leq d_e$ from Lemma 3.6 and plug in $y_e^* = 1$ to derive the upper bound on the element's query cost. □

3.3 Budget deflection for KHP

In the following, we consider *budget deflection* for KHP. Recall that budget deflection is the phenomenon that one coordinate is queried to control the uncertainty for other coordinates, see Section 2.4.

In this section, we show that budget deflection can occur and give sufficient conditions under which it cannot occur. Further, we show in Theorem 3.10 that under these conditions the optimal objective value does not decrease.

It depends on the application whether budget deflection makes sense or not. We illustrate this by the example of shortest path.

Example 3.8 (Budget deflection for KHP)**.** *Let G be the digraph on the left in Figure 2 and H be the digraph on the right. The two graphs G and H only differ in the existence of an arc between nodes v and t.*

Let the query outcome functions be binary, see Example 2.1. Furthermore, let the query cost be linear, i.e., $q_e(x_e) = q_e x_e$ for all arcs e. Let the adversary have a budget of $\Gamma = 1$.

If no query is made, the controllable budgeted uncertainty set $\mathcal{U}(x)$ is equal to the budgeted uncertainty set $\mathcal{U}$. In graph G, the adversary can spend its entire budget on arc (s, t) that has to be taken such that total cost is three. In graph H, the optimal solution for the adversary is to evenly distribute its budget on both arcs (s, v) and (s, t), resulting in total cost of 1.5.

For both graphs, if at least one of x_{sv}, x_{st} is one, the controllable budgeted uncertainty set is reduced to only contain the hedging point, i.e., $\mathcal{U}(x) = \{b\}$. Then, for both graphs G, H, the cost for arc (s,t) is zero and arc (s,t) is the shortest s-t-path. Thus, the total cost is equal to the according query cost.

In both graphs G and H, the minimal total cost over the queries is one. They are obtained for $x_{sv} = 1$ and $x_{st} = 0$. In both cases, then there is no shortest s-t-path including arc (s,v). Thus, budget deflection occurs.

In graph G, arc (s,v) is never part of an s-t-path. The arc (s,v) could have been removed in some pre-processing for the shortest path problem in graph G. Thus, spending budget on reducing the uncertainty for arc (s,v) as it happens in the optimal solution seems unintended.

In graph H there is an s-t-path that includes arc (s,v). In this case, budget deflection can be a desired aspect of modeling.

If the adversary in KHP represents uncertainty in the worst-case scenario rather than an actual adversarial agent, budget deflection seems unwanted. We can prevent budget deflection if we set the lower query function to zero:

Theorem 3.9. *Let $X = \mathbb{R}^n_{\geq 0}$ and assume that the query cost is elementwise. Let x^* be an optimal solution of KHP with query cost $q_e(x^*_e) > 0$ for some $e \in [n]$ where $\phi^\ell_e = 0$. Then there exists $u \in \mathcal{U}(x^*)$ such that there is an optimal solution y^* to the underlying problem with $y^*_e > 0$.*

Proof. Let x^* be an optimal query. Fix some $e \in [n]$ with positive query cost, i.e., $q(x^*_e) > 0$. Assume for contradiction that for all $u^* \in \mathcal{U}(x^*)$ we have $y^*_e = 0$ for all optimal choices $y^* \in Y$.

First, we argue that without loss of generality, $u^*_e = 0$: Due to $\phi^\ell_e = 0$, the remaining constraints in $\mathcal{U}$ are upper bounds on u such that $u^*_e = 0$ is feasible. The set Y does not depend on u. Thus, it suffices to consider the objective for optimality. The value u_e only occurs in the summand $d_e u_e y_e$. Due to $y^*_e = 0$ by assumption, we have $d_e u_e y_e = 0$ regardless of the value for u_e. Thus, we can set $u^*_e = 0$ without loss of generality.

Let $u^*_e = 0$ and let x' be an alternative query that is equal to x^* in all elements except e where $x'_e = 0$. The query x' is feasible because $X = \mathbb{R}^n_{\geq 0}$. Due to $q_e(x^*_e) > 0$, we obtain $F(x', u^*, y^*) < F(x^*, u^*, y^*)$ which contradicts the optimality of x^*. □

If we set $\phi^\ell_e = 0$, we disable budget deflection. As a corollary of the following theorem, we obtain that the optimal objective value does not decrease if in an instance of KHP budget deflection is disabled while we keep the remaining fixed. More generally, the optimal objective value might increase if query outcome functions are replaced by smaller ones. Formally, we have:

Theorem 3.10. *Let $\tilde{\phi}^\ell_e \leq \phi^\ell_e$ and $\tilde{\phi}^u_e \leq \phi^u_e$ for all $e \in [n]$. Assume that the optimal objective value F^* of an instance of KHP with query outcome functions ϕ^ℓ_e and ϕ^u_e exists. Further assume that the optimal objective value $\tilde{F}^*$ of the same instance except $\tilde{\phi}^\ell_e$ instead of ϕ^ℓ_e and $\tilde{\phi}^u_e$ instead of ϕ^u_e exists. Then, $F^* \leq \tilde{F}^*$.*

Proof. Let $\psi(x,u) := \min_{y \in Y} F(x,u,y)$. Let $\mathcal{U}(x)$ denote the uncertainty set with ϕ^u_e and ϕ^ℓ_e and let $\tilde{\mathcal{U}}(x)$ denote the uncertainty set with $\tilde{\phi}^u_e$ and $\tilde{\phi}^\ell_e$. Then, for all queries x we have $\mathcal{U}(x) \subseteq \tilde{\mathcal{U}}(x)$. Let $\tilde{x}^*$ be an optimal query for the instance with query outcome functions $\tilde{\phi}^u_e$ and $\tilde{\phi}^\ell_e$. The query

$\tilde{x}^*$ is feasible for the instance with ϕ_e^u and ϕ_e^ℓ. Combining, we obtain

$$F^* = \min_{x \in X} \max_{u \in \mathcal{U}(x)} \psi(x, u) \leq \max_{u \in \mathcal{U}(\tilde{x}^*)} \psi(\tilde{x}^*, u) \leq \max_{u \in \tilde{\mathcal{U}}(x_0^*)} \psi(\tilde{x}^*, u) = \tilde{F}^*.$$

□

Corollary 3.11. *Let F^* be the optimal objective value of an instance of KHP. Let F_0^* be the optimal objective value of the same instance except that $\phi^\ell \equiv 0$. Then, $F^* \leq F_0^*$.*

3.4 Single-level MIP for binary problems

KHP can be formulated as a single-level mixed-integer problem for linear query cost q, binary queries and binary underlying optimization problem. We apply McCormick envelopes [McC76] on the single-level problem derived in Theorem 3.4.

Binary query functions are restrictive compared to our original approach. However, precisely this case is studied in decision-dependent information discovery, e.g., [OP23, PGDT22, VGY22].

Theorem 3.12. *Let the query cost be linear, i.e., $q(x) = q^\top x$ and let the query outcome be binary, see Example 2.1. Furthermore, let the convex hull of the feasible set of the underlying problem be a subset of the unit cube, given by a linear number of linear inequalities. Then, an optimal query exists and KHP can be formulated as an equivalent mixed-integer linear problem with $\mathcal{O}(n)$ variables and constraints.*

Proof. As there are finitely many feasible solutions for a query x, an optimal query x^* exists.

Let $\operatorname{conv}(Y) = \left\{ y \geq 0 \;\middle|\; A^\top y = a \right\}$. Based on Theorem 3.4, in our setting KHP is equivalent to the single-level problem

$$\begin{aligned} \min_{x,y,\beta,\theta} \quad & \Gamma\theta + \sum_{e \in [n]} q_e x_e + c_e y_e + \beta_e + x_e \left(b_e d_e y_e - b_e \theta - \beta_e \right) \\ \text{s.t.} \quad & \theta + \beta_e - d_e y_e \geq 0 \qquad \forall e \in [n] \\ & Ay = a \\ & y, \beta, \theta \geq 0 \\ & x \in \{0,1\}^n . \end{aligned}$$

The only non-linear part are the bilinear summands $x_e \left(b_e d_e y_e - b_e \theta - \beta_e \right)$ in the objective. In the following, we obtain an exact reformulation for these bilinear terms by the McCormick envelopes, since x is binary. First, we deduce upper and lower bounds for the latter factor: Recall that b and d are non-negative, see Section 2. Furthermore, in an optimal solution we have $\theta \leq D$ and $\beta_e \leq d_e$, see Lemma 3.6. Together with y, β and θ being non-negative, in an optimal solution we have

$$b_e d_e y_e - b_e \theta - \beta_e \in [-b_e D - d_e, b_e d_e].$$

Next, we introduce new variables z_e for the bilinear summands. The lower bound

$$z_e \geq \max \left\{ -(b_e D + d_e) x_e, b_e d_e (x_e + y_e - 1) - b_e \theta - \beta_e \right\}$$

suffices as we minimize. In total, we obtain the following mixed-integer linear problem:

$$\begin{aligned}
\min_{x,y,z,\beta,\theta} \quad & \Gamma\theta + \sum_{e\in[n]} q_e x_e + c_e y_e + \beta_e + z_e \\
\text{s.t.} \quad & z_e - b_e d_e x_e - b_e d_e y_e + b_e\theta + \beta_e \geq -b_e d_e && \forall e \in [n] \\
& z_e + (b_e D + d_e) x_e \geq 0 && \forall e \in [n] \\
& \theta + \beta_e - d_e y_e \geq 0 && \forall e \in [n] \\
& Ay = a \\
& y, \beta, \theta \geq 0 \\
& x \in \{0,1\}^n .
\end{aligned}$$

This problem has both $\mathcal{O}(n)$ variables and constraints. □

4 Robustness against uncertain hedging points

We now consider optimization under controllable uncertainty with unknown hedging points (UHP). After we introduce the problem, we compare it to the problem of the previous section and a similar problem from the literature. Then, we provide some bounds on the optimal objective value. Afterwards, we develop an equivalent single-level reformulation of the four-level problem. Finally, we consider whether budget deflection can occur.

In UHP, the hedging points b are chosen adversarially after queries x are made and before the underlying problem's decisions y are fixed. The remaining uncertainty u realizes in a worst-case manner afterwards. In total, we have the following four-level optimization problem

$$\min_{x\in X} \max_{b\in\mathcal{B}} \min_{y\in Y} \max_{u\in\mathcal{U}(x,b)} F(x,u,y). \tag{UHP}$$

Comparison of UHP with KHP and DDID There are two main differences between UHP, considered in this section, and KHP from the previous section. First, in UHP, hedging points b are variables that can change. In contrast, hedging points are fixed values in KHP. Furthermore, in KHP, the uncertainty realizes in-between making queries and solving the underlying problem. However, in UHP, the uncertainty realizes after the decision of the underlying problem is fixed.

UHP is a generalization of decision-dependent information discovery (DDID), e.g., [PGDT22, OP23]. In DDID, there are only binary queries and the objective function does not depend on the query. Binary queries can be modeled in UHP with appropriate query outcome functions ϕ_e^ℓ, ϕ_e^u, see Example 2.1. Furthermore, if a query cost is set to zero in UHP, the objective function does not depend on the query anymore. The remaining settings in UHP and DDID are the same.

4.1 Bounds for UHP

In the following, we show upper and lower bounds on the optimal objective value of UHP. They are based on the boundaries of the intervals $[c_e, c_e + d_e]$ for the uncertain cost coefficients in the underlying problem's objective.

Observation 4.1. *For the optimal objective value $F(x^*, u^*, y^*)$ of UHP holds*

$$\min_{y\in Y} c^\top y \le F(x^*, u^*, y^*) \le \min_{y\in Y}\max_{u\in\mathcal{U}} f(u,y) \le \min_{y\in Y} (c+d)^\top y.$$

Proof. Recall that $F(x,u,y) = q(x) + (c + d\cdot u)^\top y$, that $q(x), c, d$ are non-negative, $\mathcal{U}, \mathcal{B} \subseteq [0,1]^n$, $\mathbf{0} \in X$ and $\mathcal{U}(\mathbf{0}, b) = \mathcal{U}$, see Section 2. Furthermore, the controllable uncertainty set always contains the hedging point b.

If there is no query cost in UHP, an optimal query reduces the upper bounds in the controllable uncertainty set $\mathcal{U}(x,b)$ to a maximal amount. Thus, we have

$$\begin{aligned} F(x^*, u^*, y^*) = \min_{x\in X}\max_{b\in\mathcal{B}}\min_{y\in Y}\max_{u\in\mathcal{U}(x,b)} f(u,y) + q(x) &\ge \min_{x\in X}\max_{b\in\mathcal{B}}\min_{y\in Y}\max_{u\in\mathcal{U}(x,b)} f(u,y) \\ &\ge \max_{b\in\mathcal{B}\cap\mathcal{U}}\min_{y\in Y} f(b,y) \\ &\ge \min_{y\in Y} f(\mathbf{0},y) = \min_{y\in Y} c^\top y. \end{aligned}$$

For the upper bounds, fix $x = \mathbf{0}$ to obtain

$$\begin{aligned} F(x^*, u^*, y^*) &\le \max_{b\in\mathcal{B}}\min_{y\in Y}\max_{u\in\mathcal{U}(\mathbf{0},b)} F(\mathbf{0},u,y) = \max_{b\in\mathcal{B}}\min_{y\in Y}\max_{u\in\mathcal{U}} f(u,y) = \min_{y\in Y}\max_{u\in\mathcal{U}} f(u,y) \\ &\le \min_{y\in Y}\max_{u\in[0,1]^n} (c + d\cdot u)^\top y = \min_{y\in Y} (c+d)^\top y. \end{aligned}$$

□

The robust problem for the underlying problem provides a tighter upper bound. Recall that the uncertainty set $\mathcal{U}$ is the budgeted uncertainty set, see Assumption 2.6. For binary underlying problems, the robust problem can thus be solved by $n+1$ underlying problems [BS03].

4.2 Single-level reformulation for UHP

The main result of this section is an equivalent single-level reformulation for UHP. First, we define some terms to simplify notation and consider the bilevel problem for fixed query and fixed hedging points. Then, we reformulate the inner robust problem as $n+1$ nominal problems for binary underlying problems. Afterwards in Theorem 4.7, for binary underlying problems that can be solved as LP, we obtain a single-level NLP that is equivalent to UHP.

Definition 4.2. *Let*

$$\begin{aligned} \tilde{\Gamma}(x,b) &:= \Gamma - \sum_{e\in[n]} \phi_e^\ell(x_e) b_e, \\ \tilde{c}_e(x_e, b_e) &:= c_e + \phi_e^\ell(x_e) d_e b_e, \\ h_e(x_e, b_e) &:= 1 - \phi_e^u(x_e) + \left(\phi_e^u(x_e) - \phi_e^\ell(x_e)\right) b_e, \\ h_{n+1}(x_{n+1}, b_{n+1}) &:= 0, \\ G_e(x_e, k) &:= c_e + \mathbb{1}_{e<k}(d_e - d_k)\left(1 - \phi_e^u(x_e)\right), \text{ and} \\ g_e(x_e, k) &:= c_e + \mathbb{1}_{e<k}(d_e - d_k)\left(\phi_e^u(x_e) - \phi_e^\ell(x_e)\right) + \phi_e^\ell(x_e) d_e \end{aligned}$$

where $\mathbb{1}_{e<k}$ denotes the indicator whether $e < k$.

By the assumptions on b, c, d, Γ and $\phi_e^\ell(x_e)$, see Section 2, the modified budget $\tilde{\Gamma}(x,b)$ and the modified cost $\tilde{c}_e(x_e, b_e)$ is non-negative.

Lemma 4.3. *Let $x \in X$ be a fixed query and fix a hedging point $b \in \mathcal{B}$. Then, the bilevel problem*

$$\min_{y \in Y} \max_{u \in \mathcal{U}(x,b)} f(u,y) \tag{7}$$

is equivalent to the following LP:

$$\begin{aligned} \min_{\beta,\theta,y} \quad & \tilde{\Gamma}(x,b)\theta + \sum_{e\in[n]} \tilde{c}_e(x_e,b_e)y_e + h_e(x_e,b_e)\beta_e && (8)\\ \text{s.t.} \quad & \theta + \beta_e - d_e y_e \geq 0 \qquad \forall e \in [n] \\ & \beta, \theta \geq 0 \\ & y \in Y. \end{aligned}$$

Proof. In the following, we explicitly formulate the constraints of $\mathcal{U}(x,b)$, see Definition 2.2 and Assumption 2.6, and the objective function $f(u,y)$, see Assumption 2.5. The lower level of Problem (7) is the LP

$$\begin{aligned} \max_{u} \quad & \sum_{e\in[n]} c_e y_e + d_e u_e y_e && (9)\\ \text{s.t.} \quad & u_e \geq \phi_e^\ell(x_e) b_e && \forall e \in [n] \\ & u_e \leq 1 - \phi_e^u(x_e)(1-b_e) && \forall e \in [n] \\ & \sum_{e\in[n]} u_e \leq \Gamma. \end{aligned}$$

The dual problem of Problem (9) is:

$$\begin{aligned} \min_{\alpha,\beta,\theta} \quad & \Gamma\theta + \sum_{e\in[n]} c_e y_e - \phi_e^\ell(x_e) b_e \alpha_e + \left(1 - \phi_e^u(x_e)(1-b_e)\right)\beta_e && (10)\\ \text{s.t.} \quad & -\alpha_e + \theta + \beta_e - d_e y_e = 0 \qquad \forall e \in [n] \\ & \alpha, \beta, \theta \geq 0. \end{aligned}$$

Due to strong duality, Problems (9) and (10) have the same optimal objective value. Thus, we can replace Problem (9) by Problem (10). Furthermore, we substitute $\alpha_e = \theta + \beta_e - d_e y_e \geq 0$, use Definition 4.2 and combine with the minimization of y in Problem (7) to obtain Problem (8). □

The robust counterpart of a binary optimization problem can be effectively optimized via $n + 1$ appropriate nominal optimization problems [BS03, Theorem 3]. These nominal optimization problems only differ in the cost vector. For Problem (7), we obtain the following adaption.

Assumption 4.4. *For the remaining, we assume that the elements are ordered such that entries of the vector d are non-increasing for increasing index and add $d_{n+1} = 0$, i.e.,*

$$d_1 \geq d_2 \geq \cdots \geq d_n \geq d_{n+1} = 0.$$

Theorem 4.5 (Adaption of [BS03])**.** *Let the underlying optimization problem be binary, i.e., $Y \subseteq \{0,1\}^n$. Then, Problem (7) is equivalent to*

$$\min_{k\in[n+1]} \tilde{\Gamma}(x,b)d_k + \min_{y\in Y} \sum_{e\in[n]} \tilde{c}_e(x_e,b_e)y_e + \sum_{j\in[k]} (d_j - d_k)h_j(x_j,b_j)y_j. \tag{11}$$

Proof. The single-level reformulation (8) only differs in the coefficient $h_e(x_e, b_e)$ of the dual variables β_e in the objective function to the problem considered in [BS03, Theorem 3]. The statement follows from the proof given in [BS03]. □

There are several results that reduce the number of subproblems that have to be solved, e.g., [LK14, HRS18, BGK23]. We will not further consider these results.

If the binary underlying problem is furthermore given as an LP, we provide an equivalent single-level LP for the three innermost levels of UHP in the following lemma.

Lemma 4.6. *Let x be a fixed query, assume that the underlying problem be binary, i.e., $Y \subseteq \{0,1\}^n$, fulfilling*

$$\operatorname{conv}(Y) = \left\{ y \;\middle|\; A^\top y = a, y \geq 0 \right\} \text{ with } A \in \mathbb{R}^{m\times n}, a \in \mathbb{R}^m.$$

Let $\mathcal{B}$ be given by a polynomial-sized set of linear inequalities. Then, the remaining problem, i.e.,

$$\max_{b\in\mathcal{B}} \min_{y\in Y} \max_{u\in\mathcal{U}(x,b)} f(u,y) \tag{12}$$

is equivalent to the following LP:

$$\begin{aligned} \max_{R,b,z}\ & R & (13)\\ s.t.\ & -g_e(x_e,k)b_e + \sum_{i\in[m]} A_{i,e} z_i^{(k)} \leq G_e(x_e,k) & \forall k \in [n+1], e\in[n]\\ & R + \sum_{e\in[n]} \phi_e^\ell(x_e)b_e - \sum_{i\in[m]} a_i z_i^{(k)} \leq \Gamma & \forall k\in[n+1]\\ & b \in \mathcal{B}. \end{aligned}$$

Proof. Based on Theorem 4.5, Problem (12) is equivalent to the problem

$$\begin{aligned} \max_{R\in\mathbb{R},\ b\in\mathcal{B}}\ & R \\ \text{s.t.}\ & R - \tilde{\Gamma}(x,b) \leq \psi(x,b,k) \qquad \forall k \in [n+1] \end{aligned}$$

with

$$\psi(x,b,k) := \min_{y\in Y} \sum_{e\in[n]} \tilde{c}_e(x_e,b_e)y_e + \sum_{j\in[k]} (d_j - d_k)h_j(x_j,b_j)y_j. \tag{14}$$

The dual of the minimization problem in (14) is given for every $k \in [n+1]$ by

$$\begin{aligned} \max_{z^{(k)}\in\mathbb{R}^m}\ & \sum_{i\in[m]} a_i z_i^{(k)} \\ \text{s.t.}\ & \sum_{i\in[m]} A_{i,e} z_i^{(k)} \leq \tilde{c}_e(x_e,b_e) + c_e + \mathbb{1}_{e<k}(d_e - d_k)\ h_e(x_e,b_e) \qquad \forall e\in[n]. \end{aligned}$$

Based on strong duality, we can combine this maximization problem with the evaluation of variables R and b to obtain the equivalent LP in Problem (13). □

Finally, we obtain a single-level non-linear problem that is equivalent to UHP.

Theorem 4.7. *Let the underlying problem be binary and solvable as an LP. Furthermore, let the set of hedging points $\mathcal{B}$ be given by a polynomial number of constraints that are linear in b_e. Then, for UHP there exists an equivalent single-level NLP.*

Proof. By Lemma 4.6, the inner three levels of UHP are equivalent to the LP (13). Due to strong duality, we obtain an equivalent single-level problem to UHP, when we combine the dual of Problem (13) with the minimization over queries $x \in X$. □

Corollary 4.8. *For some $A \in \mathbb{R}^{m\times n}, a \in \mathbb{R}^m$, let*

$$\operatorname{conv}(Y) = \left\{ y \;\middle|\; A^\top y = a, y \geq 0 \right\} \text{ and } \mathcal{B} = \left\{ b \in [0,1]^n \;\middle|\; \Gamma - \sum_e \phi_e^\ell(x_e) b_e \geq 0 \right\}.$$

Then, UHP is equivalent to the following single-level bilinear problem

$$\begin{aligned}
\min_{\substack{x,y,\theta,\\ \tilde{\theta},\beta}} \; & \Gamma\left(\theta + \sum_{k\in[n+1]} \tilde{\theta}_k\right) + \sum_{e\in[n]} \beta_e + \sum_{k\in[n+1]} G_e(x_e,k) y_e^{(k)} \\
\text{s.t. } & \beta_e + \phi_e^\ell(x_e)\theta + \sum_{k\in[n+1]} \phi_e^\ell(x_e)\tilde{\theta}_k + g_e(x_e,k) y_e^{(k)} \geq 0 && \forall e \in [n] \\
& -a_e\tilde{\theta}_k + \sum_{i\in[m]} A_{i,e} y_i^{(k)} = 0 && \forall k \in [n+1],\; e \in [n] \\
& \sum_{k\in[n+1]} \tilde{\theta}_k = 1 \\
& \theta,\, \tilde{\theta}_k,\, \beta_e,\, y_e^{(k)} \geq 0 && \forall k \in [n+1],\; e \in [n] \\
& x \in X.
\end{aligned}$$

For binary queries given by a polynomially sized MIP, UHP can be reformulated as a MIP based on linearizing bilinear terms, similarly as done for KHP in Theorem 3.12. Note that under comparable assumptions a similar result to Theorem 4.7 can be obtained in the setting of DDID, see [OP23, Section 4].

4.3 Budget deflection for UHP

In the following, we show that budget deflection does not occur in the setting of UHP. Recall that budget deflection is the phenomenon that a query is made for an element to control the uncertainty of other elements, see Section 2.4. In contrast to Theorem 3.9 for KHP, we do not need any additional assumption on ϕ^ℓ when we consider UHP in the following Theorem.

Theorem 4.9. *Let $X = \mathbb{R}_{\geq 0}$ and let q denote elementwise query cost. For an instance of UHP, let x^* be an optimal solution with $q_e(x_e^*) > 0$ for some $e \in [n]$. Then there exists an optimal b_e^* such that there is an optimal solution y^* with $y_e^* > 0$.*

Proof. Due to elementwise query cost, the overall objective function has the elementwise structure $F(x,u,y) = \sum_{e\in[n]} q_e(x_e) + c_e y_e + d_e u_e y_e$.

We do a proof by contradiction. Let x^* be a fixed optimal query with $q_e(x_e^*) > 0$ and $y_e^* = 0$ for some $e \in [n]$ for all optimal b_e^*. Then, the only summand including u_e in the objective, $u_e b_e y_e$ is already zero. Thus, without loss of generality $u_e^* = \phi_e^\ell(x_e^*) b_e^*$ as there is no other lower bound on u_e. As a consequence, we can have $b_e^* = 0$ without loss of generality by a similar argument. We now set $x_e = 0$ to reduce the objective value as $q_e(x_e^*) > 0$. This does not change the latter decisions for b, y, u and contradicts to the optimality of x^*. □

Corollary 4.10. *In UHP, there is no budget deflection.*

5 Conclusion and outlook

We introduced the concept of controllable uncertainty. OCU models optimization problems with uncertainty in which one can shrink intervals for uncertain parameters at a certain cost. The concept is highly flexible. In particular can be applied to a large variety of robust optimization problems that differ in the number, type and order of levels, the structure of the query function and the parameters that are subject to uncertainty. We distinguish whether hedging points are given as part of the input or chosen by an adversary a posteriori. For the first case of known hedging points, we consider the setting where the uncertainty realizes before the underlying optimization problem is solved. In the latter case of unknown hedging points, we consider the setting where the point in time when uncertainty reveals is influenced by the queries. For both cases, we consider an example problem setting with three and four levels, respectively. In both cases, we were able to simplify the problems to manageable, though still difficult, problem classes. Thereby, we illustrate that optimization models that use the concept of controllable uncertainty can still be accessible to methods which seek global optimality, despite an initially daunting number of levels.

Future research may aim to identify and classify further classes of optimization problems under controllable uncertainty for which levels can be reduced significantly. Possibly this could also discover polynomially solvable ones. Moreover, specialized algorithmic techniques could lead to significant speed-ups to solve problems under controllable uncertainty. An advantage of the OCU approach in particular in contrast to DDID is that OCU could be used in economic applications for quantifying the value of benefits achieved by queries. Depending on the setting, leader queries can serve multiple purposes, namely hedging against highly unwelcome realizations of the uncertainty, quite similar to what is done in fortification games, but also gain access to uncertain information earlier – or, in the interpretation of a two-player game, force the adversary's hand. An important aspect of ROCU models is budget deflection. A third and sometimes unwanted purpose is just to exhaust the uncertainty budget. We discuss modeling subtleties in conjunction with this as a first step to quantitatively differentiate between these purposes. We give fundamental results to study this and similar intriguing phenomena arising for ROCU models.

Acknowledgments

We thank Marc Goerigk for valuable comments on a preliminary version of this paper.

The authors further acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the framework of HyNEAT under grant no. 03SF0670B and by the

Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2163/1 - Sustainable and Energy Efficient Aviation – Project-ID 390881007.

References

[BCSW06] Gerald G. Brown, W. Matthew Carlyle, Javier Salmerón, and R. Kevin Wood. Defending critical infrastructure. *Interfaces*, 36(6):530–544, 2006.

[BEN09] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*, volume 28. Princeton university press, 2009.

[BGK23] Christina Büsing, Timo Gersing, and Arie MCA Koster. A branch and bound algorithm for robust binary optimization with budget uncertainty. *Mathematical Programming Computation*, pages 1–58, 2023.

[BHKR05] Richard Bruce, Michael Hoffmann, Danny Krizanc, and Rajeev Raman. Efficient update strategies for geometric computing with uncertainty. *Theory of Computing Systems*, 38(4):411–423, 2005.

[BL11] John R. Birge and François Louveaux. *Introduction to Stochastic Programming*. Springer, New York, 2 edition, 2011.

[BLS23a] Yasmine Beck, Ivana Ljubic, and Martin Schmidt. A survey on bilevel optimization under uncertainty. *European Journal of Operational Research*, 311(2):401–426, 2023.

[BLS23b] Yasmine Beck, Ivana Ljubić, and Martin Schmidt. Exact methods for discrete γ-robust interdiction problems with an application to the bilevel knapsack problem. *Mathematical Programming Computation*, pages 1867–2957, 2023.

[BLS23c] Yasmine Beck, Ivana Ljubić, and Martin Schmidt. An extension of the Bertsimas & Sim result for discrete, linear, and γ-robust min-max problems, 07 2023. https://optimization-online.org/?p=23532.

[BS03] Dimitris Bertsimas and Melvyn Sim. Robust discrete optimization and network flows. *Mathematical programming*, 98(1):49–71, 2003.

[BS04] Dimitris Bertsimas and Melvyn Sim. The price of robustness. *Operations research*, 52(1):35–53, 2004.

[CM10] Matthew Colvin and Christos T. Maravelias. Modeling methods and a branch and cut algorithm for pharmaceutical clinical trial planning using stochastic programming. *European Journal of Operational Research*, 203(1):205–215, 2010.

[CS07] Richard L. Church and Maria Paola Scaparra. Protecting critical assets: The r-interdiction median problem with fortification. *Geographical Analysis*, 39:129–146, 2007.

[Dem02] Stephan Dempe. *Foundations of Bilevel Programming*. Springer, Dordrecht, 2002.

[DZ21] Stephan Dempe and Alain Zemkoho. *Bilevel Optimization: Advances and Next Challenges*. Springer, Cham, 2021.

[EH15] Thomas Wilhelm Erlebach and Michael Hoffmann. Query-competitive algorithms for computing with uncertainty. *Bull. EATCS*, 116, 2015.

[EHK+08] Thomas Erlebach, Michael Hoffmann, Danny Krizanc, Matúš Mihalák, and Rajeev Raman. Computing minimum spanning trees with uncertainty. *CoRR*, abs/0802.2855, 02 2008.

[FHEED22] Ramy Fakhry, Elkafi Hassini, Mohamed Ezzeldin, and Wael El-Dakhakhni. Tri-level mixed-binary linear programming: Solution approaches and application in defending critical infrastructure. *European Journal of Operational Research*, 298(3):1114–1131, 2022.

[FMM17] Jacob Focke, Nicole Megow, and Julie Meissner. Minimum spanning tree under explorable uncertainty in theory and experiments. In *16th International Symposium on Experimental Algorithms (SEA 2017)*, 08 2017.

[FMO+03] Tomás Feder, Rajeev Motwani, Liadan O'Callaghan, Chris Olston, and Rina Panigrahy. Computing shortest paths with uncertainty. In *Annual Symposium on Theoretical Aspects of Computer Science*, volume 2607, pages 367–378, 02 2003.

[FWR12] Raquel J. Fonseca, Wolfram Wiesemann, and Berc Rustem. Robust international portfolio management. *Computational Management Science*, 9:31–62, 2012.

[GG04] Vikas Goel and Ignacio E. Grossmann. A stochastic programming approach to planning of offshore gas field developments under uncertainty in reserves. *Computers & Chemical Engineering*, 28(8):1409–1429, 2004.

[GGI+15] Marc Goerigk, Manoj Gupta, Jonas Ide, Anita Schöbel, and Sandeep Sen. The robust knapsack problem with queries. *Computers & Operations Research*, 55, 03 2015.

[GKST23] Marc Goerigk, Jannis Kurtz, Martin Schmidt, and Johannes Thürauf. Connections and reformulations between robust and bilevel optimization, 06 2023. https://optimization-online.org/?p=23366.

[GSS11] Manoj Gupta, Yogish Sabharwal, and Sandeep Sen. The update complexity of selection and related problems. *Leibniz International Proceedings in Informatics, LIPIcs*, 13, 08 2011.

[HRS18] Christoph Hansknecht, Alexander Richter, and Sebastian Stiller. Fast robust shortest path computations. In *18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[JLSY15] Jian Gang Jin, Linjun Lu, Lijun Sun, and Jingbo Yin. Optimal allocation of protective resources in urban rail transit networks against intentional attacks. *Transportation Research Part E: Logistics and Transportation Review*, 84:73–87, 2015.

[JWW98] Tore W. Jonsbråten, Roger J.-B. Wets, and David L. Woodruff. A class of stochastic programs with decision dependent random elements. *Annals of Operations Research*, 82:83–106, 1998.

[Kah91] Simon Kahan. A model for data in motion. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, STOC '91, page 265–277, New York, NY, USA, 1991. Association for Computing Machinery.

[KLLS21] Thomas Kleinert, Martine Labbé, Ivana Ljubic, and Martin Schmidt. A survey on mixed-integer programming techniques in bilevel optimization. *EURO Journal on Computational Optimization*, 9:100007, 2021.

[KV18] Bernhard Korte and Jens Vygen. *Combinatorial Optimization*. Springer, Berlin, 6 edition, 2018.

[LK14] Taehan Lee and Changhyun Kwon. A short note on the robust combinatorial optimization problems with cardinality constrained uncertainty. *4OR*, 12:373–378, 2014.

[LLM+23] Markus Leitner, Ivana Ljubic, Michele Monaci, Markus Sinnl, and Kübra Taninmis. An exact method for binary fortification games. *European Journal of Operational Research*, 307(3):1026–1039, 2023.

[McC76] Garth P McCormick. Computability of global solutions to factorable nonconvex programs: Part I—Convex underestimating problems. *Mathematical programming*, 10(1):147–175, 1976.

[Mei18] Julie Meißner. *Uncertainty Exploration: Algorithms, Competitive Analysis, and Computational Experiments*. PhD thesis, Technische Universität Berlin, 2018.

[MS19] Arturo I. Merino and José A. Soto. The minimum cost query problem on matroids with uncertainty areas. In *ICALP*, 2019.

[OP23] Jérémy Omer and Michael Poss. Combinatorial robust optimization with decision-dependent information discovery and polyhedral uncertainty, 09 2023. https://hal.science/hal-04097679v2.

[PGDT22] Rosario Paradiso, Angelos Georghiou, Said Dabia, and Denis Tönissen. Exact and approximate schemes for robust optimization problems with decision dependent information discovery, 08 2022. https://arxiv.org/abs/2208.04115v1.

[SCJB10] Senay Solak, John-Paul B. Clarke, Ellis L. Johnson, and Earl R. Barnes. Optimization of R&D project portfolios under endogenous uncertainty. *European Journal of Operational Research*, 207(1):420–433, 2010.

[Sim09] Stephen Simons. Minimax theorems. In Christodoulos A. Floudas and Panos M. Pardalos, editors, *Encyclopedia of Optimization, Second Edition*, pages 2087–2093. Springer, 2009.

[SS20] J. Cole Smith and Yongjia Song. A survey of network interdiction models and algorithms. *European Journal of Operational Research*, 283(3):797–811, 2020.

[VGY22] Phebe Vayanos, Angelos Georghiou, and Han Yu. Robust optimization with decision-dependent information discovery, August 2022. http://arxiv.org/abs/2004.08490.

[vN28] John von Neumann. Zur Theorie der Gesellschaftsspiele. *Annals of Mathematics Studies*, 100:295–320, 1928.