On Robust Min-Cut and Max-Flow under Uncertainty

Mustafa Ç. Pınar
mustafap@bilkent.edu.tr
Bilkent University
Ankara, Turkey

Deniz Akkaya
Bilkent University
Ankara, Turkey
deniz.akkaya@bilkent.edu.tr

Abstract

The max–flow min–cut theorem links the maximum flow value in a network to the capacity of the minimum s-t cut, with optimal solutions corresponding to binary graph partitions. When arc capacities are uncertain, however, this integrality property may fail. This paper studies robust max–flow and min–cut formulations under uncertainty, with a focus on ellipsoidal uncertainty sets that allow correlated arc capacities. Using a KKT-based analysis, we derive optimality conditions and introduce an induced capacity representation that captures the effect of correlations. Our results provide conditions under which integral solutions remain optimal and highlight how correlation structures can destabilize them.

Keywords

Robust optimization, Min-cut problem, Max-flow problem, Ellipsoidal uncertainty

1 Introduction

Let V be a finite set of nodes and A be a set of directed edges. Consider a directed graph $G = (V, A)$ equipped with a designated source node s, a sink node t, and strictly positive capacities $c_{ij} > 0$ for all arcs $(i, j) \in A$. The maximum flow problem seeks to determine the largest amount of flow that can be sent from the source node to the sink node while respecting the capacity limitations on each arc and ensuring conservation of flow at intermediate nodes.

To establish the standard linear programming (LP) formulation of this problem, we introduce a non-negative decision variable $f_{ij} \geq 0$ representing the amount of flow transmitted along arc $(i, j) \in A$. In addition, we define a scalar variable v that represents the total amount of flow sent from the source to the sink. The objective is to maximize v subject to two fundamental types of constraints. First, flow conservation must hold at every node except the source and sink, meaning that the total incoming flow must equal the total outgoing flow. Second, the flow on each arc cannot exceed its corresponding capacity. Under these constraints, the model determines a feasible routing of flow through the network that achieves the largest possible throughput from s to t.

$$
\begin{aligned}
\max_{\mathbf{f}\in\mathbb{R}^{|A|}, v\in\mathbb{R}} \quad & v \\
\text{s.t.} \quad & \sum_{j\in V} f_{sj} - \sum_{j\in V} f_{js} = v, \quad \text{(source)} \\
& \sum_{j\in V} f_{tj} - \sum_{j\in V} f_{jt} = -v, \quad \text{(sink)} \\
& \sum_{j\in V} f_{ij} - \sum_{j\in V} f_{ji} = 0, \quad \forall i \in V \setminus \{s, t\}, \\
& 0 \leq f_{ij} \leq c_{ij}, \quad \forall (i, j) \in A.
\end{aligned}
$$

To formulate the corresponding dual problem, we associate an unrestricted dual variable y_i with the flow conservation constraint of each node $i \in V$ (including the source s and sink t), and a non-negative dual variable $x_{ij} \geq 0$ with each capacity constraint $f_{ij} \leq c_{ij}$. These dual variables admit a natural interpretation in terms of node potentials and arc selections that define a separating cut in the network.

In particular, the variables $\mathbf{y}$ can be viewed as potential values assigned to nodes, while the variables $\mathbf{x}$ capture the extent to which arcs violate the ordering induced by these potentials. The constraint $y_i - y_j \leq x_{ij}$ ensures that whenever a node i is placed on the source side of a partition and node j lies on the sink side, the corresponding arc (i, j) contributes to the cut through a positive x_{ij} value. The objective then accumulates the capacities of such arcs, effectively measuring the total capacity of the induced cut.

Under this interpretation, the dual problem identifies a cut separating s from t with minimum total capacity. This formulation is equivalent to the classical minimum s-t cut problem, and by strong duality its optimal value coincides with the optimal value of the maximum flow problem. The resulting dual formulation, which identifies a minimum capacity cut in the network [1], is given by:

$$
\begin{aligned}
\min_{\mathbf{y}\in\mathbb{R}^{|V|}, \mathbf{x}\in\mathbb{R}^{|A|}} \quad & \sum_{(i,j)\in A} c_{ij} x_{ij} \\
\text{s.t.} \quad & y_s - y_t = 1, \\
& y_i - y_j \leq x_{ij}, \quad \forall (i, j) \in A, \\
& x_{ij} \geq 0, \quad \forall (i, j) \in A.
\end{aligned}
$$

The classical max-flow min-cut theorem serves as a cornerstone of network optimization, establishing the equivalence between the maximum flow value and the capacity of the minimum s-t cut in a directed graph. While this relationship is well-understood in deterministic settings, uncertainty in arc capacities presents significant challenges for maintaining the *integrality property*, the condition under which an optimal robust solution corresponds to a binary partition of the graph.

It is well-established that under a polyhedral budget uncertainty model for arc capacities, the resulting robust optimization problem yields an integral optimal solution [2]. On the other hand, there exist instances where the robust formulation yields highly fractional solutions, depending on the problem structure; see, for example, [3]. Another approach is to approximate axis-parallel ellipsoids by budgeted uncertainty sets and then establish integrality results along the lines of [4].

This paper investigates the conditions under which robust counterparts of these problems admit binary solutions. We first consider ellipsoidal uncertainty models, where robustness is defined over an ellipsoidal region of deviation. Using a KKT-based analysis, we derive necessary conditions for integer optimality. We then study algebraic bounds relating the nominal cost and the robust objective of the min-cut problem, highlighting the trade-offs introduced by robustness. Finally, we extend these results to

INOC '26, Liège (Belgium)

10.48786/inoc.2026.02

correlated uncertainty, corresponding to non-axis-parallel ellipsoidal uncertainty sets, and discuss computational experiments on well-known directed graph benchmark instances.

Our analysis emphasizes the structural impact of uncertainty on the classical dual relationship between flow and cut formulations. In particular, the introduction of an ℓ_2-type robustness term alters the geometry of the feasible region, often leading to fractional optimal solutions even when the nominal problem admits integral ones. By carefully examining the optimality conditions of the resulting convex program, we identify structural properties that either preserve or destroy integrality. These properties reveal how the size of a cut, the distribution of nominal capacities, and the magnitude of uncertainty jointly influence the behavior of optimal solutions.

The results provide both theoretical insight and practical guidance for modeling uncertainty in network flow problems. On the theoretical side, they clarify when the robust min-cut formulation retains the combinatorial structure of the classical problem. From a practical perspective, the derived bounds and optimality conditions help identify network instances where robust optimization may significantly alter the structure of optimal cuts. The computational study illustrates these phenomena and demonstrates how the proposed conditions can be used to anticipate the emergence of fractional solutions in robust network design problems.

2 Robust Min-Cut Problem under Ellipsoidal Uncertainty

In practical applications, inaccuracies in modeling external factors or evaluating feasibility requirements are often unavoidable, making uncertainty in arc capacities c_{ij} a natural modeling consideration. To capture such deviations in a structured way, we represent the possible realizations of capacities through ellipsoidal uncertainty sets, defined as

$$U_E := \left\{ \mathbf{c} \in \mathbb{R}^{|A|} : \|\mathbf{c} - \bar{\mathbf{c}}\|_2 \leq \epsilon \right\},$$

where $\bar{\mathbf{c}} \in \mathbb{R}^{|A|}_{++}$ denotes the nominal capacity vector, and $\epsilon > 0$ is a tolerance parameter controlling the level of robustness. This ellipsoidal uncertainty set represents all capacity realizations whose deviation from the nominal vector remains bounded in the Euclidean norm. In other words, the parameter ϵ limits the magnitude of the aggregate perturbation that may affect the arc capacities.

Under this model, we study a worst-case optimization problem in which the true capacity vector may vary within the uncertainty set. In particular, the capacity of each arc is interpreted as a perturbed value around its nominal level, and the optimization problem seeks solutions that remain feasible and optimal against the most adverse realization within this set. Accordingly, each arc capacity is expressed as

$$\mathbf{c} = \bar{\mathbf{c}} + \mathbf{d}, \quad \text{with} \quad \|\mathbf{d}\|_2 \leq \epsilon.$$

Hence, the robust min-cut problem under this construction can be written as:

$$\begin{aligned}
\min_{\mathbf{y}\in\mathbb{R}^{|V|},\, \mathbf{x}\in\mathbb{R}^{|A|}} \quad & \max_{\substack{\mathbf{d}\in\mathbb{R}^{|A|} \\ \|\mathbf{d}\|_2\leq\epsilon}} \sum_{(i,j)\in A} (\bar{c}_{ij} + d_{ij})\, x_{ij} \\
\text{s.t.} \quad & y_s - y_t = 1, \\
& y_i - y_j \leq x_{ij}, \quad \forall (i,j) \in A, \\
& x_{ij} \geq 0, \quad \forall (i,j) \in A.
\end{aligned}$$

The inner maximization $\max_{\mathbf{c}\in U_E} \mathbf{c}^\top \mathbf{x}$ can be isolated and it yields

$$\max_{\substack{\mathbf{d}\in\mathbb{R}^{|A|} \\ \|\mathbf{d}\|_2\leq\epsilon}} (\bar{\mathbf{c}} + \mathbf{d})^\top \mathbf{x} = \bar{\mathbf{c}}^\top \mathbf{x} + \epsilon \|\mathbf{x}\|_2 = \bar{\mathbf{c}}^\top \mathbf{x} + \epsilon \sqrt{\sum_{(i,j)\in A} x_{ij}^2}.$$

This observation makes the robust min-cut problem a convex (second-order cone) program:

$$\begin{aligned}
\min_{\mathbf{y}\in\mathbb{R}^{|V|},\, \mathbf{x}\in\mathbb{R}^{|A|}} \quad & \sum_{(i,j)\in A} \bar{c}_{ij} x_{ij} + \epsilon \sqrt{\sum_{(i,j)\in A} x_{ij}^2} \\
\text{s.t.} \quad & x_{ij} - y_i + y_j \geq 0, \quad \forall (i,j) \in A, \\
& y_s = 1, \quad y_t = 0, \\
& x_{ij} \geq 0, \quad \forall (i,j) \in A.
\end{aligned}$$

We next highlight several structural properties of this formulation. Since the primal problem is a feasible convex program with linear constraints, strong duality holds. Consequently, the optimal values of the robust min-cut and robust max-flow formulations coincide:

Robust Min-Cut Value = Robust Max-Flow Value.

Despite this dual relationship, the presence of the nonlinear term $\epsilon\|\mathbf{x}\|_2$ in the objective generally destroys the classical integrality property of the min-cut formulation. In contrast to the deterministic setting, the optimal solution $\mathbf{y}^*$ is often fractional, with components satisfying $y_i^* \in [0,1]$ rather than taking binary values.

Furthermore, under uncertainty the arc variables $\mathbf{x}$ may also become fractional even when the node variables $\mathbf{y}$ correspond to an integral cut. This behavior reflects the smoothing effect of the ℓ_2-regularization term, which allows the optimization to distribute mass across multiple arcs in order to reduce the robustness penalty.

Our main objective is therefore to characterize conditions under which the robust min-cut problem admits an optimal solution satisfying $\mathbf{y} \in \{0,1\}^{|V|}$. When such a condition holds, the solution corresponds to a genuine s-t cut in the graph, and the robust objective reduces to the cost of that cut under ellipsoidal perturbations. In this case, the robust cost of a cut can be expressed as

$$C(\mathbf{y}) = \sum_{(i,j)\in\delta^+(S(\mathbf{y}))} \bar{c}_{ij} \; + \; \epsilon\sqrt{|\delta^+(S(\mathbf{y}))|},$$

where $S(\mathbf{y}) := \{i \in V : y_i = 1\}$ and $\delta^+(S) := \{(i,j) \in A : i \in S,\ j \notin S\}$.

3 A Necessary Condition for Integral Optimality

We present a nontrivial necessary condition for an integral cut to be optimal in the robust min-cut convex program with ellipsoidal uncertainty. The condition is derived from the Karush–Kuhn–Tucker (KKT) optimality system and admits a natural combinatorial interpretation: the shifted cost vector on the cut edges must be extendable to a nonnegative s-t flow, that is, to a vector satisfying zero divergence at all internal nodes.

When $\mathbf{y}$ is integral, we can define the corresponding cut S in the usual way. In this case, the feasibility constraints link the arc variables $\mathbf{x}$ to the node variables $\mathbf{y}$ and determine which arcs must carry positive values. In particular, these constraints force the variables x_{ij} to be positive exactly on the arcs crossing the cut from S to its complement, while all other arcs may remain at zero. Consequently, the vector $\mathbf{x}$ encodes the structure of the cut and allows the robust objective to be interpreted directly in terms of the set of crossing edges.

$$x_{ij} \geq \max\{0, y_i - y_j\} = \begin{cases} 1, & (i,j) \in \delta^+(S), \\ 0, & \text{otherwise.} \end{cases}$$

Since increasing any x_{ij} only increases the objective value, the optimal choice of $\mathbf{x}$ for a fixed integral $\mathbf{y}$ is

$$x_{ij} = \begin{cases} 1, & (i,j) \in \delta^+(S), \\ 0, & \text{otherwise,} \end{cases} \qquad \forall (i,j) \in A.$$

Let $m := |\delta^+(S)|$. In this case, we have $\|\mathbf{x}\|_2 = \sqrt{m}$.

We proceed by introducing $\mathbf{u} \in \mathbb{R}_+^{|A|}$ as the vector of dual (Lagrange) multipliers associated with the constraints $x_{ij} - y_i + y_j \geq 0$, and $\mathbf{v} \in \mathbb{R}_+^{|A|}$ as the multipliers corresponding to the non-negativity constraints $x_{ij} \geq 0$. These dual variables capture the sensitivity of the objective to perturbations in the corresponding constraints. While it would be possible to define additional multipliers λ_s and λ_t for the equality constraints imposed on the source and sink nodes, such boundary multipliers are not needed explicitly in our subsequent analysis and are therefore omitted for clarity.

Theorem 3.1. *Suppose* $(\mathbf{x}^*, \mathbf{y}^*)$ *is an optimal primal solution and that* $\mathbf{y}^*$ *are integral. Let* $S = \{i \in V : y_i^* = 1\}$ *be the corresponding cut, and* $m = |\delta^+(S)|$. *Then, there exists a vector* $\mathbf{u} \in \mathbb{R}_+^{|A|}$ *satisfying the following properties:*

(i) ***Induced Capacity on Cut Edges:*** *The flow value on cut edges is determined by the robust cost:*

$$u_{ij} = \bar{c}_{ij} + \frac{\epsilon}{\sqrt{m}} \quad \forall (i,j) \in \delta^+(S).$$

(ii) ***Flow Conservation at Internal Nodes:*** *For every node* $k \in V \setminus \{s,t\}$, *the vector* $\mathbf{u}$ *is divergence free:*

$$\sum_{(k,j) \in A} u_{kj} = \sum_{(i,k) \in A} u_{ik}.$$

(iii) ***Capacity Bound on Non-cut Edges:*** *The flow on non-cut edges is bounded by the nominal capacity:*

$$u_{ij} \leq \bar{c}_{ij} \quad \forall (i,j) \notin \delta^+(S).$$

The proof of the preceding theorem relies on a careful examination of the Karush–Kuhn–Tucker (KKT) conditions. In the standard min-cut formulation, the boundary conditions ensure that $\mathbf{x} = 0$ is infeasible. As a result, the norm term $\|\mathbf{x}\|_2$ is evaluated away from the origin, guaranteeing that the Lagrangian function is differentiable throughout the feasible region.

Moreover, the KKT analysis shows that, under suitable assumptions on the feasibility of the dual vector $\mathbf{u}$, an integral primal solution exists. This observation establishes a connection between dual feasibility and primal integrality, and the converse implication provides additional structural insight into the robust min-cut problem, which will be elaborated further in the final version of this work.

4 An Algebraic Bound on Robust Cut Nominal Cost

In this section, we begin by considering the standard maximum flow problem defined with the nominal arc capacities $\bar{\mathbf{c}}$. Let F_{nom} denote the corresponding maximum flow value from the source s to the sink t. Let $\delta^+(S_{\text{nom}})$ represent a nominal minimum cut, that is, a cut achieving the flow value F_{nom}, and let $k := |\delta^+(S_{\text{nom}})|$ denote the number of edges crossing this nominal cut. This setup provides a baseline for comparing the structure and cost of robust cuts under ellipsoidal uncertainty against the deterministic nominal scenario.

Theorem 4.1. *Let* $(\mathbf{x}^*, \mathbf{y}^*)$ *be an optimal solution to the Robust Min-Cut Problem where* $\mathbf{y}^*$ *is integral. Define the cut* $S^* = \{i \in V : y_i^* = 1\}$ *and the set of outgoing edges are* $\delta^+(S^*)$. *Let* $|\delta^+(S^*)| = k$. *Then*

$$\sum_{(i,j) \in \delta^+(S^*)} \bar{c}_{ij} \leq F_{\text{nom}} + \epsilon(\sqrt{k} - \sqrt{m}).$$

The proof of this theorem requires a detailed analysis. We begin by establishing strong duality between the robust min-cut and max-flow formulations, proceed to characterize the KKT points of the problem, and finally apply Slater's condition to obtain precise and tight results.

The resulting bound illustrates a fundamental trade-off in robust optimization. The term $\epsilon\sqrt{m}$ represents the robustness penalty incurred by the optimal robust cut S^*, capturing the need to protect against worst-case perturbations across its m edges. Conversely, the term $\epsilon\sqrt{k}$ reflects the robustness advantage available to the adversarial flow, which can exploit uncertainty along the k edges of the nominal minimum cut. When $m < k$, the optimal robust cut contains fewer edges than the nominal min-cut, incurring a smaller robustness penalty; consequently, its nominal cost must also be smaller to satisfy the bound. In the case where $m = k$, the bound reduces to

$$\sum_{(i,j) \in \delta^+(S^*)} \bar{c}_{ij} \leq F_{\text{nom}},$$

which is tight, with equality holding precisely when $S^* = S_{\text{nom}}$. However, if $m > k$, the bound becomes

$$\sum_{(i,j) \in \delta^+(S^*)} \bar{c}_{ij} \leq F_{\text{nom}} + \epsilon(\sqrt{k} - \sqrt{m}) < F_{\text{nom}},$$

which is infeasible since the capacity of any cut must be at least F_{nom}. This reasoning suggests that under ellipsoidal uncertainty, optimal robust cuts tend to have fewer edges than the nominal minimum cuts.

5 Robust Min-Cut under Correlated Ellipsoidal Uncertainty

We extend the ellipsoidal uncertainty model U_E by incorporating a positive semi-definite covariance matrix $\Sigma \in \mathbb{R}^{|A| \times |A|}$ to capture dependencies among arc capacity deviations. Unlike the axis-parallel case, which assumes independent perturbations and relies solely on the Euclidean norm, the covariance matrix allows correlated deviations across arcs. This means that a change in the capacity of one arc may coincide with or influence changes in others, reflecting more realistic systemic uncertainties arising from shared infrastructure, environmental factors, or correlated demand fluctuations. Consequently, the resulting uncertainty set forms a rotated ellipsoid, with its shape and orientation fully determined by the covariance structure encoded in Σ.

The uncertainty set is defined as

$$U_\Sigma = \{\mathbf{c} \in \mathbb{R}_{++}^{A} : \sqrt{(\mathbf{c} - \bar{\mathbf{c}})^\top \Sigma (\mathbf{c} - \bar{\mathbf{c}})} \leq \epsilon\}.$$

For a fixed feasible cut vector $\mathbf{x}$, the inner maximization $\max_{c\in U_\Sigma} \mathbf{c}^\top \mathbf{x}$ yields the robust cost:

$$\begin{aligned}\min_{\mathbf{y}\in\mathbb{R}^{|V|},\, \mathbf{x}\in\mathbb{R}^{|A|}} \quad & \sum_{(i,j)\in A} \bar{c}_{ij}x_{ij} + \epsilon\sqrt{\mathbf{x}\Sigma\mathbf{x}} \\ \text{s.t.} \quad & x_{ij} - y_i + y_j \geq 0, \quad \forall (i,j) \in A, \\ & y_s = 1, \quad y_t = 0, \\ & x_{ij} \geq 0, \quad \forall (i,j) \in A.\end{aligned}$$

Proposition 5.1. *Let $(\mathbf{x}^*, \mathbf{y}^*)$ be an optimal solution where $\mathbf{y}^*$ is integral and let S^* be the corresponding cut. There exists a divergence-free flow $\mathbf{u}$ such that the induced capacity cur on edges $(i,j) \in \delta^+(S^*)$ is given by:*

$$u_{ij} = \bar{c}_{ij} + \epsilon\frac{(\Sigma\mathbf{x}^*)_{ij}}{\sqrt{(\mathbf{x})^\top\Sigma\mathbf{x}^*}}$$

where $(\Sigma\mathbf{x}^*)_{ij} = \sum_{(i',j')\in\delta^+(S^*)} \sigma_{(i,j)(i',j')}$. For non-cut edges $(i,j) \in \delta^+(S^*)$, the flow is bounded by $u_{ij} \leq \bar{c}_{ij} + \epsilon\frac{(\Sigma\mathbf{x}^*)_{ij}}{\sqrt{(\mathbf{x})^\top\Sigma\mathbf{x}^*}} + \gamma_{ij}$, where γ_{ij} is the dual multiplier for non-negativity.

Similar to the results obtained for the axis-parallel uncertainty model, the proof of this proposition relies on a careful examination of the Karush-Kuhn-Tucker (KKT) conditions. By analyzing the Lagrangian system under the generalized covariance-based uncertainty, we can extend the structural insights from the simpler axis-parallel case. The key relationships between the dual variables, divergence-free flows, and induced capacities on cut edges remain largely intact, with only minor adjustments required to account for the correlations encoded in Σ. As a result, the main structural properties of integral optimal solutions, and necessary conditions carry over to the correlated setting essentially without loss.

Moreover, the algebraic bounds that relate the nominal minimum cut to the robust optimal cut are preserved under this generalization. For any feasible integral cut vector $\mathbf{x}^*$, we define the total cut variance as

$$V(\mathbf{x}^*) = (\mathbf{x}^*)^\top\Sigma\mathbf{x}^*,$$

which measures the cumulative impact of correlated perturbations across the cut edges. Similarly, for a nominal minimum cut S_{nom} with vector $\mathbf{x}^*_{\text{nom}}$, we define

$$V_{\text{nom}} = (\mathbf{x}^*_{\text{nom}})^\top\Sigma\mathbf{x}^*_{\text{nom}},$$

which quantifies the variance of the nominal cut under the same correlation structure. These variance measures play a crucial role in extending the previously derived upper bounds on the nominal cost of robust cuts and provide insight into how correlations influence the trade-off between robustness and cut capacity in the generalized uncertainty setting.

Theorem 5.2. *Let $(\mathbf{x}^*, \mathbf{y}^*)$ be an optimal solution to the Robust Min-Cut Problem with Generalized Uncertainty where $\mathbf{y}^*$ is integral. Define the cut $S^* = \{i \in V : y_i^* = 1\}$ and the set of outgoing edges are $\delta^+(S^*)$. Then*

$$\sum_{(i,j)\in\delta^+(S^*)} \bar{c}_{ij} \leq F_{\text{nom}} + \epsilon(\sqrt{V_{\text{nom}}} - \sqrt{V(\mathbf{x}^*)}).$$

The proof is quite similar to axis-parallel case, but still can be extended to the generalized case.

6 Conclusion

In this paper we studied the robust minimum cut problem under ellipsoidal uncertainty. While the classical min-cut formulation enjoys the integrality property, introducing robustness through a convex uncertainty set alters the geometry of the problem and may lead to fractional optimal solutions.

Using a KKT-based analysis, we derived a necessary condition characterizing when an integral cut can remain optimal in the robust formulation. We further established algebraic bounds relating the nominal capacity of robust optimal cuts to the value of the nominal minimum cut, revealing a structural trade-off between nominal cost and robustness. These results also lead to conditions that restrict the size of candidate robust cuts.

We then extended the analysis to a generalized uncertainty model with correlated perturbations represented through a covariance matrix. The main structural insights obtained in the axis-parallel case continue to hold under this more general formulation.

Future work will focus on validating the proposed results through computational experiments on large-scale network instances.

References

[1] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. 1993. *Network Flows: Theory, Algorithm, and Applications.* Prentice Hall, Englewood Cliffs, NJ.
[2] Dimitri Bertsimas and Melvyn Sim. 2003. Robust discrete optimization and network flows. *Mathematical Programming* 98 (2003), 49–71.
[3] Christina Büsing, Timo Gersing, and Arie M. C. A. Koster. 2025. Recycling valid inequalities for robust combinatorial optimization with budgeted uncertainty. *Mathematical Programming* 210, 1 (2025), 97–146. doi:10.1007/s10107-024-02135-7
[4] Michael Poss. 2018. Robust combinatorial optimization with knapsack uncertainty. *Discrete Optimization* 27 (2018), 88–102. doi:10.1016/j.disopt.2017.09.004