The Use of Decision-dependent Uncertainty Sets in Robust Optimization

Nikolaos H. Lappas and Chrysanthos E. Gounaris*

Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

This paper contemplates the amenability of Robust Optimization to address problems that involve endogenous uncertainty, i.e., uncertainty that is affected by the decision maker's strategy. To that end, we extend generic polyhedral uncertainty sets that are typically considered in RO into sets that depend on the actual decisions. Such dependency allows us not only to capture functional changes in correlations that may be induced by given decisions, but also to alleviate conservatism effects from non-materialized model parameters. Case studies from process network capacity expansion, offshore oilfield development, and clinical trial planning showcase the effect of various levels of uncertainty in the optimal solutions as well as the modeling potential and favorable computational tractability characteristics of the proposed framework, which can be utilized in both short-term and long-term decision making.

Keywords

Robust Optimization, Endogenous Uncertainty, Decision-dependent Uncertainty Sets

Introduction

Uncertainty is inherent in virtually any operation we wish to optimize. Parameters affected by other agents such as market prices and demand, unexpected events such as disruptions, or simply incomplete information about the system, may render solutions of deterministic models suboptimal or even infeasible when parameter realizations deviate from their nominal values. To that end, multiple approaches have been proposed so as to account for uncertainty during the optimization process.

Robust Optimization (RO) (Ben-Tal and Nemirovski, 1999) offers an attractive option, especially for applications where distributional information about the uncertainty is limited and/or where solution feasibility is top priority. RO seeks solutions that remain feasible for any possible uncertainty realization from within a considered uncertainty set. The set captures known correlations among uncertain parameters, but knowledge of detailed probability distributions is not typically required to construct such a set. Multiple types of uncertainty sets (e.g., ellipsoidal, cardinality-constrained budgets, polyhedral) can be used in this context, exploiting correlations among uncertain parameters as a mechanism to control the trade-off between robustness and performance. A common characteristic across traditional uncertainty sets utilized in RO is their "static" nature, i.e., the range of parameter realizations they admit is not affected by the value of the decision variables. This fact leads to two important limitations. Firstly, static sets do not suffice to model settings where one's decisions may render a model parameter physically meaningless (e.g., the yield of a reaction step that was not selected in the optimal flowsheet) or where one's decisions affect the underlying probability distribution from which the parameter realization draws (e.g., deciding to enter a market sooner will give us access to a larger demand). Secondly, the use of a static set may lead to overly conservative solutions in cases where non-materialized parameters contribute in correlations among parameters.

In order to overcome the above challenges, we recently proposed the use of decision-dependent uncertainty sets (DDUS) in the context of process scheduling (Lappas and Gounaris, 2016). In that work, the effects of processing times corresponding to non-materialized tasks were removed dynamically from the uncertainty set. The purpose of this paper is to assess

* *gounaris@cmu.edu*

the concept of using DDUS so to handle endogenous uncertainty via the RO methodology in a number of applications of interest to the Process Systems Engineering (PSE) community.

Robust Optimization

Consider an optimization problem with continuous decision variables, $x \in \mathcal{X}$, binary decision variables, $w \in \mathcal{W}$, and uncertain parameters, $q \in \mathcal{Q}$:

$$\begin{aligned} \min_{x \in \mathcal{X}, w \in \mathcal{W}} \quad & f(x, w, q) \\ \text{subject to} \quad & g_i(x, w, q) \leq 0 \quad \forall i \end{aligned} \tag{1}$$

The fundamental idea behind RO is to guarantee the constraint satisfaction for any parameter realization from within the uncertainty set, and then seek to identify the best feasible solution assessed against the worst possible realization of the uncertainty. The RO problem can thus be cast as the following bi-level problem:

$$\begin{aligned} \min_{x \in \mathcal{X}, w \in \mathcal{W}} \quad & \max_{q \in \mathcal{Q}} \; f(x, w, q) \\ \text{subject to} \quad & g_i(x, w, q) \leq 0 \qquad \forall\, q \in \mathcal{Q} \;\forall\, i \end{aligned} \tag{2}$$

We remark that the problem stated above involves infinitely many constraints and, in order solve it, one shall typically apply standard reformulation techniques from Semi-infinite Programming.

Traditional Uncertainty Sets

One of the most critical elements of applying the RO methodology is the selection of the uncertainty set against which robust feasibility is sought. The shape and size of the set can have a direct impact on the solution conservatism, and extensive literature effort has focused on how to best exploit parameter correlations so as to reduce this conservatism. The most popular uncertainty sets used in RO are the ellipsoidal (Ben-Tal and Nemirovski, 1999) and the cardinality-constrained (Bertsimas and Sim, 2004) ones. In this work, we will focus on general polyhedral uncertainty sets, which generalize the latter. Assuming that the original (deterministic) model for the problem of interest is of linear or mixed-integer linear form, these sets maintain this property also for the reformulated robust counterpart model. A general polyhedral uncertainty set adheres to the following form:

$$\mathcal{Q} = \left\{ \begin{array}{c} q \in \mathbb{R}^{NQ},\ p \in \{0,1\}^{NP} : \\ Hq + Gp \leq d \\ q^L \leq q \leq q^U \end{array} \right\}, \tag{3}$$

where NQ corresponds to the number of continuous uncertain parameters, q, and NP to the number of binary uncertain parameters, p. Matrices H and G store the parameter coefficients, having as many rows as the number of correlations in the set. Vector d stores the right hand side constants. Finally, q^L and q^U are bounds for the admissible realizations of continuous uncertain parameters. We remark that for cases involving binary uncertain parameters ($NP \neq 0$), some solution approaches require certain restrictions on the data, such as total unimodularity of matrix G and an empty matrix H.

Endogenous Uncertainty

In order to properly motivate the use of DDUS, we provide here some background on uncertainty characterization. Uncertain parameters can be classified as *exogenous*, when they are independent of decisions, and *endogenous*, when the decision maker can manipulate their stochastic support. There are various flavors of endogeneity. In certain cases, a decision may render a parameter referenced in a model physically meaningless (e.g., the price of a product under development that did not hit the market). In other cases, a decision may affect the timing of parameter realization (e.g., the period in which we decide to drill an oil well will be the time at which we will learn the true magnitude of our initial well production rate). Finally, a decision may affect the distribution from which an uncertain parameter draws (e.g., the technology we choose will affect the range of possible yields for the process).

Most optimization problems studied in RO literature consider only exogenous uncertainty (e.g., uncertainty in product prices and demand in non-monopolistic markets, weather conditions), which in RO can be modeled using traditional, non decision-dependent uncertainty sets. However, there are still many important problems where uncertainty is subject to the optimizer's decisions (Jonsbråten et al., 1998). In a monopolistic market, for instance, a decision to vastly increase the production will have a negative effect on product prices. Other examples arise in oilfield development planning (Goel and Grossmann, 2004), network capacity expansion (Goel and Grossmann, 2006), network interdiction problems (Peeta et al., 2010), and the planning of clinical trials (Colvin and Maravelias, 2008), to name but a few. These problems are typically handled via two-stage or multi-stage Stochastic Programming approaches. However, the enforcement of nonanticipativity restrictions leads to formulations of excessive size. In order to address the aforementioned challenges, significant efforts have been made regarding better modeling and solution approaches (Apap and Grossmann, 2015). Alternative approaches that model these problems as Markov decision processes have also been proposed.

The above application settings constitute examples where the use of traditional (static) uncertainty sets in RO can prove limited. On the one hand, constant uncertainty set coefficients do not allow functional changes in the applicable correlations they model, i.e., they cannot adapt to changes in underlying distributions. On the other hand, referencing the "realized" value of a parameter that did not materialize in the given solution leads to an expansion of the space of possible realizations of those parameters that did indeed materialize, introducing unnecessary conservatism. This is not to mention that doing so would be unrealistic in practice, since the "realization" of a non-materialized parameter retains no physical meaning.

Decision-dependent Uncertainty Sets

In order to address problems with endogenous uncertainty, we extend the static uncertainty set of Eq. (3) into a set that depends on the discrete decision variables, w, as follows:

$$\mathcal{Q}\left(w\right)=\left\{\begin{array}{c} q\in\mathbb{R}^{NQ},\ p\in\left\{0,1\right\}^{NP}: \\ H\left(v^{q}(w)\circ q\right)+G\left(v^{p}(w)\circ p\right) \\ \leq Dw+d \\ q^{L}\leq q\leq q^{U} \end{array}\right\}, \tag{4}$$

where $v^q(w) \in \{0,1\}^{NQ}$ and $v^p(w) \in \{0,1\}^{NP}$ are problem-specific, binary-valued functions of our discrete decision variables that govern the materialization of each uncertain parameter, q or p, while matrix D contributes decision-dependency to the right-hand-sides. In a DDUS, q^L and q^U retain their definitions as bounds for the admissible realizations of parameters q, but these bounds should now be wide enough to apply under all possible decisions $w \in \mathcal{W}$ for which the corresponding elements of $v^q(w)$ attain the value of one.

Exogenous parameters can be accommodated in the above scheme by setting the corresponding elements of the materialization indicators, $v^q(w)$ or $v^p(w)$ to the constant value of one. Furthermore, non-decision dependent correlations among exogenous parameters can also be accommodated by setting to zero the corresponding rows of matrix D. Interestingly, the DDUS of Eq. (4) retains the properties of its static precursors of Eq. (3) with regards to the model class of the resulting robust counterpart formulation. We should however remark that, when robustification is performed via the standard dualization approach, an exact linearization (i.e., Glover inequalities) is required to reformulate bilinear terms that occur between variables w and the newly introduced dual variables. Typically, this has mild computational impact and scales well with the number of binary variables referenced in the DDUS.

In the following, we briefly illustrate a number of modeling conveniences that DDUS afford us:

1. A correlation as shown in Eq. (5) would allow for the introduction of decision-dependent bounds for parameter realization and, more generally, allow for the introduction of decision-dependent distributional information in the set.

$$q \leq d_1 w_1 + d_2 w_2 + d_3 w_3 \tag{5}$$

2. A correlation as shown in Eq. (6) would allow for the introduction of a decision-dependent cardinality budget, where the maximum number of parameters that can attain their worst-case realization is limited based on (e.g., investment-related) decisions. Figure 1 illustrates this concept.

$$p_1 + p_2 + p_3 \leq w_1 + w_2 + w_3 \tag{6}$$

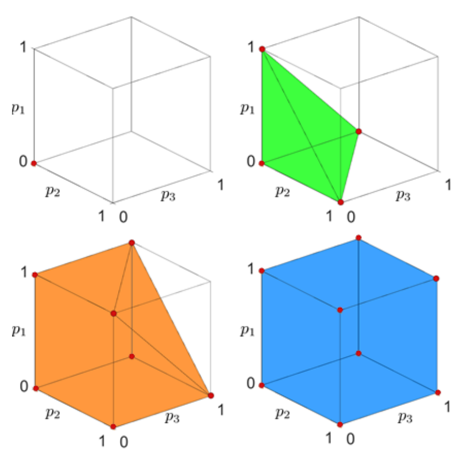

Figure 1. Cardinality budget DDUS for various decisions $w_1+w_2+w_3=n$, where n is 0 (no color), 1 (green), 2 (orange), and 3 (blue). Red dots signify admissible realizations.

3. In the case where a parameter is not materialized as a result of a decision, the contribution of that parameter in the set can be removed by dynamically

projecting the set into a lower dimension. This can be achieved, e.g., with a correlation as shown in Eq. (7). Figure 2 illustrates the difference in the form of the DDUS when parameter q_2 does and does not materialize (as governed by binary variable w_2).

$$w_1 q_1 + w_2 q_2 + w_3 q_3 \leq d_1 w_1 + w_2 d_1 + w_3 d_3 \tag{7}$$

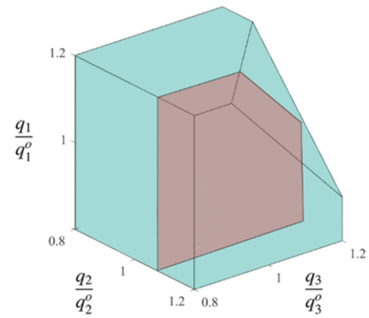

Figure 2. Polyhedral DDUS for the cases of two (magenta) and three (teal) materialized parameters.

Case Studies

Here we present the construction of DDUS for a number of problems that feature endogenous uncertainty and that have been previously studied in the PSE literature using primarily Stochastic Programming methodologies. This now enables us to realistically cast and study these problems as RO problems too. Details about the problem statements, their notation, and exact data for the various benchmark instances can be found in the corresponding cited references. The symbols ξ and ϕ refer to constants that we use to tune the size and shape, respectively, of the polyhedral uncertainty sets we derive. Exact values for those constants are provided later, in the Computational Results section.

Case Study I: Capacity Expansion Problem

The first case study is from Goel and Grossmann (2006) and refers to the capacity expansion of a process network, where intermediate product B has to be processed by process 3 so as to produce a high-value product A. Product B itself can be either purchased directly from the market, produced in-house by process 1 using raw material C, or produced in-house by process 2 using raw material D. The above options are not mutually exclusive and can be utilized in conjunction. Moreover, the cost for the installation of necessary equipment has to be accounted, if processes 1 or 2 are to be chosen. The overall objective is to maximize the profit within a 10-period horizon.

The endogenous uncertain parameters involved in this problem are the production yields, θ_{it}, for each unit $i \in \mathcal{I}$ and time period $t \in \mathcal{T}$. The endogeneity stems from the fact that, if a unit is not operated in a given time period, the corresponding production yield does not retain a physical meaning. To that end, the binary-valued function associated with the materialization of each θ_{it} parameter would be $v_{it}^{\theta} = b_{it}$, where b_{it} is the binary decision of operating unit i in time period t. Furthermore, the problem features one additional parameter, the total demand for product A, δ. Since this is an exogenous parameter, the corresponding materialization indicator is $v^{\delta} = 1$.

We define a budget correlation among the materialized uncertain parameters for each unit, which reflects the fact that any process is expected to perform close to its nominal performance, on average across the horizon. In addition, each of the parameters is bounded around their nominal realization values, θ_{it}^0 and δ^0. The nominal values were chosen as the average among the low and high scenario levels reported in the paper. The DDUS we use is as follows:

$$\mathcal{Q}(b) = \left\{ \begin{array}{l} \theta \in \mathbb{R}^{|\mathcal{I}||\mathcal{T}|},\ \delta \in \mathbb{R} : \\ \displaystyle\sum_{t \in \mathcal{T}} b_{it}\theta_{it} \geq (1 - \xi\phi) \sum_{t \in \mathcal{T}} b_{it}\theta_{it}^0 \quad \forall\, i \in \mathcal{I} \\ (1-\xi)\,\theta_{it}^0 \leq \theta_{it} \leq (1+\xi)\,\theta_{it}^0 \quad \forall\, i \in \mathcal{I}\ \forall\, t \in \mathcal{T} \\ (1-\xi)\,\delta^0 \leq \delta \leq (1+\xi)\,\delta^0 \end{array} \right\} \tag{8}$$

Case Study II: Offshore Oil Planning Problem

The second study originates from Goel and Grossmann (2004) (specifically, example 4 in that paper), where an oil company has identified 5 offshore oil reserves and wants to plan its activities so as to develop them. In order to extract oil from each reserve, a well platform has to be installed with the necessary pipelines that transfer the oil to the production platform and, ultimately, to the shore. In addition, the company has at its disposal multiple extraction technologies, which come at given costs and which can achieve different initial production rates. The objective is to maximize the net present value of the oilfield development project over a 15-year planning horizon.

The endogenous uncertain parameters involved in this problem are the initial deliverabilities, $initDeliv_f$, of each field $f \in \mathcal{F}$. These parameters are of endogenous nature because they do not retain a physical meaning if a well is not drilled. The materialization indicator functions in this case are $v_f^{initDeliv} = w_f^A + w_f^B + w_f^C$, where $w_f^{A/B/C}$ are binary variables that decide the type of extraction technology (A, B, or C) used in each case. Note that the model enforces that, for each field $f \in \mathcal{F}$, at most one of these variables can attain the value of one and, hence, the indicators $v_f^{initDeliv}$ are properly restricted to binary evaluations. Finally, the total reserve sizes, $Size_f$, for each field $f \in \mathcal{F}$, are also uncertain. Since these are exogenous uncertain parameters, their materialization indicators are $v_f^{Size} = 1$.

Motivated from the probability distributions assumed in the referenced paper, we impose budget correlations to restrict the deviations from their nominal values of the sums (across all fields) of initial deliverabilities and total reserve sizes. In addition, each of the oilfield sizes is bounded around their nominal realization values, $Size_f^0$, while the initial deliverabilities are bounded around the nominal realization values yielded by the corresponding technologies that were chosen in each case (denoted as α_f, β_f and γ_f for technologies A, B, and C, respectively). All nominal values were chosen as the medium scenario levels reported in the paper. The DDUS we use is as follows:

$$\mathcal{Q}(w) = \left\{ \begin{array}{l} initDeliv \in \mathbb{R}^{|\mathcal{F}|},\ Size \in \mathbb{R}^{|\mathcal{F}|} : \\ \displaystyle\sum_{f \in \mathcal{F}} \left(w_f^A + w_f^B + w_f^C\right) initDeliv_f \geq \\ \quad (1-\xi\phi) \displaystyle\sum_{f \in \mathcal{F}} \left(\alpha_f w_f^A + \beta_f w_f^B + \gamma_f w_f^C\right) \\ \left(w_f^A + w_f^B + w_f^C\right) initDeliv_f \geq \\ \quad (1-\xi)\left(\alpha_f w_f^A + \beta_f w_f^B + \gamma_f w_f^C\right) \quad \forall\ f \in \mathcal{F} \\ \left(w_f^A + w_f^B + w_f^C\right) initDeliv_f \leq \\ \quad (1+\xi)\left(\alpha_f w_f^A + \beta_f w_f^B + \gamma_f w_f^C\right) \quad \forall\ f \in \mathcal{F} \\ initDeliv_f \geq \\ \quad (1-\xi) \min\{\alpha_f, \beta_f, \gamma_f\} \quad \forall\ f \in \mathcal{F} \\ initDeliv_f \leq \\ \quad (1+\xi) \max\{\alpha_f, \beta_f, \gamma_f\} \quad \forall\ f \in \mathcal{F} \\ (1-\xi)\, Size_f^0 \leq Size_f \leq (1+\xi)\, Size_f^0\ \forall\ f \in \mathcal{F} \end{array} \right\} \tag{9}$$

Case Study III: Clinical Trial Planning Problem

The final case study comes from Colvin and Maravelias (2008) (specifically, example 2 in that paper) and constitutes an R&D portfolio optimization problem. Here, a pharmaceutical company has the opportunity to pursue the development of 5 potential drugs. Before any of them can generate revenue for the company, however, it has to undergo a series of 3 clinical trial phases of given durations. Only those drugs that succeed in all three phases are approved by the regulator and enter the market. There are costs associated with performing each trial, and the company has to plan the use of its limited R&D resources over the next 36-month horizon so as to maximize its portfolio's net present value.

The endogenous uncertain parameters involved in this problem are α_{ij}, indicating the success or not of a clinical trial of drug $i \in \mathcal{I}$ in trial phase $j \in \mathcal{J}$.[1] Endogeneity arises due to the fact that such a parameter has no physical meaning if the corresponding trial is not attempted. The materialization indicator function in this case is $v_{ij}^{\alpha} = \sum_{t \in \mathcal{T}} X_{ijt}$, where X_{ijt} is the binary decision to pursue trial (i, j) in time period $t \in \mathcal{T}$.

According to the data provided in the referenced paper, a trial has a higher probability to succeed in later phases than in earlier ones. To that end, we impose phase-specific correlations to restrict the average number of failures that may occur in the corresponding trials in a manner that is consistent with this observation. The DDUS we use is as follows:

$$\mathcal{Q}(X) = \left\{ \begin{array}{l} \alpha \in \{0,1\}^{|\mathcal{I}||\mathcal{T}|} : \\ \displaystyle\sum_{i \in \mathcal{I}} \sum_{t \in \mathcal{T}} X_{ijt}\alpha_{ij} \geq \\ \quad \displaystyle\sum_{i \in \mathcal{I}} (1-\xi_{ij}) \sum_{t \in \mathcal{T}} X_{ijt} \quad \forall\ j \in \mathcal{J} \end{array} \right\} \tag{10}$$

Computational Results

In order to investigate the trade-off between robustness and performance, we instantiated and solved each case study with 3 levels of uncertainty, namely low (L), medium (M), and high (H). More specifically, for case studies I and II, these levels of uncertainty correspond to (ξ, ϕ) settings of $(0.1, 0.25)$, $(0.2, 0.50)$, and $(0.3, 0.75)$, respectively. For case study III, we chose values $\xi_{ij} = 5(1-\rho_{ij})\xi$, where ξ refers to the three levels above for L/M/H and constants ρ_{ij} correspond to the probability of drug i successfully undergoing trial in phase j (data from the referenced paper).

Table 1 shows normalized optimal objective values, elucidating the amount of risk premium (difference between nominal deterministic and robust optimal solutions[2]) that has to be paid in order to insure against var-

[1] Unlike the previous case studies, the uncertain parameters in this example are of discrete (binary) nature.

[2] Since all case studies are maximization problems, the robust

ious levels of uncertainty. The table also compares with robust solutions obtained using standard non-DDUS.[3] We observe that DDUS-based RO solutions feature better objective values compared to their non-DDUS counterparts, recovering much of the risk premium paid by the latter (yet without any compromises in the level of insurance against risk). This can be attributed to the fact that, in the case of non-DDUS, the non-materialized parameters (which are not dynamically removed from a non-DDUS) attain their best possible values (since they do not affect the optimal solution), driving the materialized ones (which can negatively impact the optimal solution) to attain worse realizations.

Finally, Table 2 presents formulation sizes and associated solution statistics. We observe that, for these problems, DDUS-based RO did not require excessive sacrifices in terms of computational tractability.

Table 1. Price of Robustness

Case Study	Determ.	DDUS			non-DDUS		
		L	M	H	L	M	H
I	100	90	84	79	71	65	44
II	100	87	79	60	69	63	41
III	100	71	65	54	57	45	36

Table 2. Computational Statistics

	Case Study I		Case Study II		Case Study III	
	Det.	RO	Det.	RO	Det.	RO
geom. avg. CPU time	0.1	1.4	2.2	7.4	0.1	1.4
avg. #Nodes	32	57	144	612	56	182
# Bin. Variables	45	45	125	125	84	84
# Cont. Variables	105	1,120	632	5,487	221	1,865
# Constraints	231	657	435	2,937	335	2,612

Conclusions

In this paper, we studied problems involving endogenous uncertainty, which are abundant in PSE practice. In order to efficiently address these problems with an RO framework, we extended generic polyhedral uncertainty sets into their decision-dependent counterparts, which offer extra modeling flexibility as well as reduce solution conservatism. The application of RO with DDUS was illustrated in three case studies motivated from the stochastic programming literature. Our computational results demonstrated that DDUS can provide considerably less conservative solutions while maintaining the computational tractability advantages of RO. Finally, we remark that while the DDUS introduced in this work are presented in a polyhedral form, this does not preclude the implementation of more involved DDUS of ellipsoidal or conic nature.

solutions always exhibit lower objective values.

[3]The non-DDUS were derived by setting all materialization indicators to the value of one and by replacing the right hand side terms Dw with $\max_{w\in\mathcal{W}:\{v(w)=1\}} Dw$.

Acknowledgments

The authors gratefully acknowledge support from the National Science Foundation (grant CBET–1510787). NHL further acknowledges support from the University of Patras via an Andreas Mentzelopoulos scholarship.

References

Apap, R. M. and Grossmann, I. E. (2015). Models and computational strategies for multistage stochastic programming under endogenous and exogenous uncertainties. *Manuscript in preparation.*

Ben-Tal, A. and Nemirovski, A. (1999). Robust solutions of uncertain linear programs. *Operations research letters*, 25(1):1–13.

Bertsimas, D. and Sim, M. (2004). The price of robustness. *Operations research*, 52(1):35–53.

Colvin, M. and Maravelias, C. T. (2008). A stochastic programming approach for clinical trial planning in new drug development. *Computers & Chemical Engineering*, 32(11):2626–2642.

Goel, V. and Grossmann, I. E. (2004). A stochastic programming approach to planning of offshore gas field developments under uncertainty in reserves. *Computers & chemical engineering*, 28(8):1409–1429.

Goel, V. and Grossmann, I. E. (2006). A class of stochastic programs with decision dependent uncertainty. *Mathematical programming*, 108(2-3):355–394.

Jonsbråten, T. W., Wets, R. J., and Woodruff, D. L. (1998). A class of stochastic programs with decision dependent random elements. *Annals of Operations Research*, 82:83–106.

Lappas, N. H. and Gounaris, C. E. (2016). Multi-stage adjustable robust optimization for process scheduling under uncertainty. *AIChE Journal*, 62(5):1646–1667.

Peeta, S., Salman, F. S., Gunnec, D., and Viswanath, K. (2010). Pre-disaster investment decisions for strengthening a highway network. *Computers & Operations Research*, 37(10):1708–1719.