

The impact of the existence of multiple adjustable robust solutions

Frans J. C. T. de Ruiter¹  ·
Ruud C. M. Brekelmans¹ · Dick den Hertog¹

Received: 18 November 2014 / Accepted: 2 January 2016 / Published online: 30 January 2016
© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract In this note we show that multiple solutions exist for the production-inventory example in the seminal paper on adjustable robust optimization in Ben-Tal et al. (Math Program 99(2):351–376, 2004). All these optimal robust solutions have the same worst-case objective value, but the mean objective values differ up to 21.9 % and for individual realizations this difference can be up to 59.4 %. We show via additional experiments that these differences in performance become negligible when using a folding horizon approach. The aim of this paper is to convince users of adjustable robust optimization to check for existence of multiple solutions. Using the production-inventory example and an illustrative toy example we deduce three important implications of the existence of multiple optimal robust solutions. First, if one neglects this existence of multiple solutions, then one can wrongly conclude that the adjustable robust solution does not outperform the nonadjustable robust solution. Second, even when it is a priori known that the adjustable and nonadjustable robust solutions are equivalent on worst-case objective value, they might still differ on the mean objective value. Third, even if it is known that affine decision rules yield (near)

This research was partially supported by the Netherlands Organisation for Scientific Research (NWO) Research Talent Grant 406-14-067.

✉ Frans J. C. T. de Ruiter
f.j.c.t.deruiter@tilburguniversity.edu

Ruud C. M. Brekelmans
r.c.m.brekelmans@tilburguniversity.edu

Dick den Hertog
d.denhertog@tilburguniversity.edu

¹ Department of Econometrics and Operations Research, Tilburg University, Warandelaan 2, 5000 LE Tilburg, The Netherlands

optimal performance in the adjustable robust optimization setting, then still nonlinear decision rules can yield much better mean objective values.

Keywords Adjustable robust optimization · Production-inventory problems · Folding horizon · Multiple solutions

1 Introduction

In [2] the Robust Optimization (RO) methodology is extended to multi-stage problems. The proposed Adjustable Robust Optimization (ARO) techniques appeared to be very effective to solve uncertain multi-stage optimization problems. This first paper on ARO has been cited more than 500 times already, and the ARO methodology has been applied to a wide variety of problems (see e.g. the survey papers [3,6]). Recently, it was shown that (A)RO problems may have multiple optimal solutions, and that not all of these solutions are Pareto robustly optimal [8]. A solution is called Pareto robustly optimal if there is no other robustly feasible solution that has better objective value for at least one scenario, and for all other scenarios in the uncertainty set the objective value is not worse.

In this note we show that the ARO model of the production-inventory problem in [2], which is the seminal work on ARO, also has multiple optimal robust solutions. Although in robust optimization one operates in a distribution-free environment, an often used performance measure is the mean objective value, which is evaluated posteriorly assuming some information on the distribution of the parameters. For the cases considered in [2], we show that among the optimal robust solutions, the difference in mean objective value can be as much as 21.9% and for individual realizations the difference can be up to 59.4%. This underlines the importance of the message in [8] that ARO problems may have multiple optimal robust solutions. In such cases one can often find optimal robust solutions that are much better with respect to the mean objective value than solutions that were initially found.

We also extend the experiments performed in [2] by including a folding horizon approach. In a folding horizon approach the model is re-optimized in each period using the available information at that point of time and only the decisions for the current time are implemented. Using this approach we find that there are still multiple optimal robust solutions, but the differences in mean costs diminish. This is mainly due to the fact that the here-and-now decisions are unique in almost all periods. As a last experiment, we also analyze the model and solutions we found when replacing the worst-case objective by an expected value objective. For the expected value objective we find that, for the seminal production-inventory problem considered here, the solution is unique.

In the second part of this note we discuss several important implications for practical ARO. The first implication is that, by ignoring the possibility of multiple solutions, one can incorrectly conclude that the ARO solution is not better than the RO solution, or even incorrectly conclude that ARO is (much) better than RO. The second implication is that even in cases where it is a priori known that RO and ARO are equivalent, i.e., they have the same worst-case optimal objective value, one cannot conclude

that there is no value in using ARO. This is because in many cases there are ARO solutions that give much better solutions for the mean costs. The third implication is that even in cases where affine decision rules are (nearly) optimal, i.e., the optimal robust objective value cannot be improved by using nonlinear decision rule, one cannot conclude that there is no value in using nonlinear decision rules. Such a conclusion might be wrong, since nonlinear decision rules may yield much better solutions for the expected objective value. These implications are illustrated by using both the production-inventory example from [2] and two toy examples.

Our aim is to convince users of ARO that one should always check for the existence of multiple solutions. In many papers on ARO it is not reported that one checked for possible existence of multiple solutions. These papers run the risk that much better solutions could have been found, or even that wrong conclusions have been drawn. For example, researchers who use the same production-inventory example as in the seminal work [2] to test new ARO methods, should be aware of the fact that this problem has many optimal robust solutions with big differences in mean costs.

2 Multiple adjustable robust solutions

To illustrate the implications of multiple adjustable robust solutions we use three problems. The first problem is the production-inventory problem by [2] in its original setting. The second problem is an illustrative toy example where the existence of multiple solutions is more directly visible. The last toy problem we investigate is a two-stage facility location problem. For all models we study both the impact in a folding and in a non-folding horizon approach.

2.1 Production-inventory model by Ben-Tal et al. [2]

We have repeated the experiments for the production-inventory problem by [2]. All solutions are obtained using the commercial solver Gurobi 6.0 [7] programmed in the YALMIP language [10] in MATLAB. All options of Gurobi were left at their default values.

We have found three distinct optimal robust solutions for the original model by [2, pp. 369–370]. All of these solutions are optimal in a robust sense, i.e. they have the same worst-case costs, but costs differ for individual (non worst-case) realizations of the demand. The first solution was obtained by just solving the original model with Gurobi. The average costs of this solution turned out to be much higher than the solution reported by [2]. The second solution is the solution that performs best on the mean costs among all optimal robust solutions. It can be found via the following two-step approach similar to the methods used by [8] to find so-called Pareto robustly optimal solutions:

1. Solve the original model from [2], which gives a solution with minimal worst-case costs.

2. Change the objective into minimizing the costs for the nominal demand trajectory. Furthermore, add a constraint that ensures that the worst-case costs do not exceed the costs found in Step 1.

The solution obtained after step two is the ‘Best’ solution, the one that performs best on the expected objective value among all optimal robust solutions that use linear decision rules, assuming that nominal demand is equal to the expected demand. The third solution is found by changing the objective in the second step into *maximizing* costs for the nominal demand trajectory. This we call the ‘Worst’ solution. Without the two-step approach, and some bad luck, one could have obtained this solution as a ‘First’ solution, i.e. by solving the original problem formulation. The performances of these three optimal robust solutions are given in Table 1. The first column states the uncertainty level, for which we used the same levels as in [2]. If the level of uncertainty is 2.5 %, then this indicates that in each period the realized demand could be up to 2.5 % higher or lower than the nominal demand. The three solutions are all robustly optimal, so they have the same worst-case costs (WC costs). For each of those solutions we have determined the mean costs and the standard deviation. In [2] the mean costs were approximated using 100 simulated demand trajectories drawn from a uniform distribution. The mean costs can also be determined exactly since the objective is linear in the uncertain demand. For the mean costs comparison we assume, as in the original paper, that the mean demand is given by the nominal demand scenario. The standard deviation was derived using the second moment of the uniform distribution, the distribution that was also used in the seminal paper by [2] to sample the scenarios to calculate average costs.

As is clear from Table 1, the performances of the three solutions differ significantly. For both the ‘First’ solution and the ‘Worst’ solution we give the mean and maximum *performance gap*. The mean performance gap is just the percentage increase of the mean costs compared to the mean costs of the ‘Best’ solution. The maximum performance gap is the single demand trajectory that results in the largest difference in costs between the ‘Best’ solution and the ‘Worst’ (or ‘First’) solution. To explain how this gap is calculated, we determine the costs for the ‘Worst’ and the ‘Best’ solution, when trajectory \mathbf{d} realizes, by respectively $OPT_W(\mathbf{d})$ and $OPT_B(\mathbf{d})$. These costs are linear in demand \mathbf{d} because the original objective is linear, fixed recourse and we use linear decision rules. The maximum performance gap for the ‘Worst’ solution is given by

$$\max_{\mathbf{d} \in \mathcal{U}} \frac{OPT_W(\mathbf{d}) - OPT_B(\mathbf{d})}{OPT_B(\mathbf{d})},$$

where \mathcal{U} is the box uncertainty set (defined by a set of linear constraints) used in this inventory problem. This is a linear-fractional maximization problem, which can be written as a linear optimization problem using the well-known Charnes-Cooper transformation [5]. The maximum performance gap for the ‘First’ solution is defined and determined analogously. The ‘First’ solution, which is the solution we obtained after solving the original LP problem with our solver, has mean costs of up to 14.5 % higher than the mean costs for the best solution for a 20 % uncertainty level. The ‘Worst’ solution has a performance gap of 21.9 % for the same uncertainty level. If we compare the performance for individual realizations, we see that the costs can

increase up to 39.4 and 59.4 % for the ‘First’ and ‘Worst’ solutions, respectively. For uncertainty levels up to 10 % the mean costs for the ‘Worst’ solution are equal to the worst-case costs, meaning that the worst-case costs are attained in every single scenario. Finally, as reported by [2], only for an uncertainty level of 2.5 % one can find a feasible nonadjustable solution implying that production levels in each period must be determined at the beginning of the planning horizon. The mean costs of 35279 for the nonadjustable solution are only slightly higher than the mean costs for the adjustable ‘Worst’ solution. Note that in the nonadjustable case there is no uncertainty in the objective, hence the mean costs are equal to the worst-case costs.

The mean costs of the solution reported by [2], where no use of a two-step approach was reported, coincides with the performance of our ‘Best solution’. We have tried various settings for our solver to see whether we could also replicate their good result as a ‘First’ solution. We tried both primal/dual simplex methods, interior point methods and a mixture of both in Gurobi. We have also solved the model for each of these options with crossover either enabled or disabled. If the crossover option is enabled, then the solver will push a solution in the optimal facet to a basic solution. None of these alterations led to a solution that was considerably better than our ‘First’ solution depicted in Table 1.

2.2 Folding horizon versus non-folding horizon

One might wonder whether the same differences in mean costs still exist if a so-called folding horizon (FH) is used. In a folding horizon approach the model is re-optimized at each period using the available information at that point of time and only the decisions for the current time are implemented. This is done for each period t starting from the first period until the end of the planning horizon. Using this folding horizon approach we again compared solutions that used the two-step approach in each step (Best FH solution), without a two-step approach (First FH solution) and when the two-step approach was used when maximizing for nominal demand in the second step (Worst FH solution). An exact calculation of the mean costs and the standard deviation is not possible for this experiment. Therefore, we draw 100 demand trajectories independently and uniformly distributed in each period. These trajectories are used to approximate the mean costs and the standard deviation when using the folding horizon approach. Simulations were also used in [2] to approximate the mean costs and the standard deviation for the non-folding horizon approach. The results are depicted in Table 2. We stress that this folding horizon approach was not used in [2]. Clearly, using the two-step approach does not yield significantly better results for the folding horizon approach. Often the resulting costs are the same for both approaches, but for one of our simulated realizations the extra costs incurred when not using the two-step approach is 0.7 %. Even stronger, for each simulated demand trajectory, the costs when using the folding horizon approach (Best FH solution) were at most the costs of the “First FH” solution. Finally, note that the mean costs for the folding horizon solutions are not much lower than the mean costs of the ‘Best’ solution given in Table 1, meaning that there is not much additional gain by re-optimizing in each step as is done in the folding horizon approach. It is at a first glance surprising that the

Table 2 Performance of the best, first obtained and worst optimal robust solutions using the folding horizon approach

Uncertainty level (%)	Best FH solution		First FH solution		Worst FH solution		Performance gap	
	Mean	Std	Mean	Std	Mean	Std	Mean (%)	Max (%)
2.5	33909	179	33912	179	33912	178	0.0	0.1
5	34057	330	34061	328	34059	328	0.0	0.0
10	34327	676	34350	667	34351	666	0.1	0.7
20	34495	1361	34517	1348	34532	1339	0.1	0.6

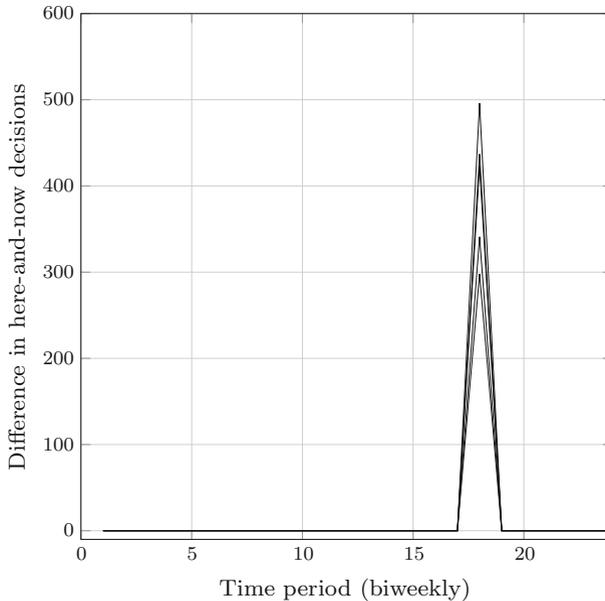


Fig. 1 Here-and-now decisions for factory 2 only differ in period 18 (5 scenarios depicted)

effect of having multiple optimal solutions diminishes when using a folding horizon approach. We found that this is mainly because the first stage decisions are unique for almost all time periods and in all simulated scenarios. The question whether or not the first stage decisions are unique can be answered by fixing the worst case costs in the first step, as in the usual two step approach, and then minimize or maximize the order quantity in the current time period. In this way we get, for each time period t , a lower and upper bound on the feasible first stage decisions. In Fig. 1 we depict the differences between the maximum and the minimum for the 20% uncertainty level for one out of the three factories. The behavior of the solutions depicted was observed for all other cases as well: the vast majority of the first-stage decisions are unique. We only witnessed non-unique optimal here-and-now decisions in time periods 6 and 18, depending on the factory (1, 2 or 3) considered.

Finally, we also investigate what happens if we optimize the expected objective value rather than the worst-case objective value in the non-folding horizon approach. This can be done at comparable computational costs, by replacing the maximization over all realizations in the objective by an objective that solely considers the nominal demand. This expected objective value was also used in [9] to prove optimality of linear decision rules under stochastic and robust settings. The authors did not study the existence of multiple adjustable solutions. We stress that, although we now minimize an expected objective value, we still have a robust problem with ‘hard’ constraints, i.e., the constraints should be satisfied for any realization within the uncertainty set. The main difference with the two step approach is that we do not fix the worst-case objective value as we did in the second step. Arguably, this approach would make more sense in

Table 3 Performance of the linear decision rule that minimizes the mean costs

Uncertainty level (%)	WC	Mean	Std
2.5	35108	33919	178
5	36412	34031	357
10	39040	34311	708
20	44298	35066	1375

problems where the objective is a ‘soft’ criterion as opposed to the constraints which are typically ‘hard’ restrictions. When minimizing the expected objective value, the worst-case objective value is ignored. Hence, in principle, the worst-case costs could be very high. To find the worst-case objective value for a given linear decision rule, a posteriori, one can simply maximize the costs over all possible realizations within the uncertainty set. The results for the optimization problem, with the ‘soft’ expected objective value, but ‘hard’ constraints, are depicted in Table 3. First of all, we note that there is not much difference between the mean costs and the worst-case costs with respect to the ‘Best’ robust solution given earlier in Table 1. There is only a very minor increase in worst-case costs and a very minor decrease in the mean costs. Hence, minimizing the mean costs yields a solution that has very similar costs to the costs of the solution obtained when minimizing worst-case costs. Second, there is no ‘Best’ and ‘Worst’ solution displayed in Table 3. This is because we found that the obtained solution is unique, so there does not exist a linear decision rule, with minimum mean costs, that has a different (neither better nor worse) guarantee on the worst-case objective value.

In the inventory model the decisions are made biweekly. Therefore, it makes sense to use a folding horizon approaches in this case. The impact of multiple adjustable robust solutions on the mean costs is negligible when we re-optimize. However, there might still be value in checking for multiple solutions in (non-)folding horizon approaches for inventory models and related multi-stage optimization models for the following reasons:

1. The non-folding horizon solution can be used as a backup solution in case of failure in hardware or software during the re-optimization steps. This is especially important in more critical multi-stage optimization systems such as power systems.
2. Re-optimization might take too much computation time or might not be possible at all. This happens in multi-stage optimization settings when periods follow up close in time, or when the solutions are implemented in low-end software systems. Examples of low-end computer systems are traffic light systems, that are not designed to solve the more computationally demanding optimization models.

Although for this inventory model the impact of the existence of multiple adjustable robust solutions on the mean costs seems to be negligible, there are other models where there could be a significant impact. This is illustrated by our toy examples in the next section.

Table 4 Comparison of the different nonadjustable and adjustable solutions

	RO_1	RO_2	LDR_1	LDR_2	NDR_1	NDR_2
Here-and-now x	1	0	1	0	1	0
Wait-and-see ^a y	0	0	$-b$	$-\frac{1}{2}b$	$-b^2 - b$	$-\frac{1}{2}b^3$
Profits for scenario (a, b)	a	0	$a + b$	$\frac{1}{2}b$	$a + b^2 + b$	$\frac{1}{2}b^3$
Worst-case profits	0	0	0	0	0	0
Mean profits (with unif. distr.)	$\frac{1}{2}$	0	1	$\frac{1}{4}$	$\frac{4}{3}$	$\frac{1}{8}$

^a Note that for RO_1 and RO_2 the variable y is a here-and-now variable

2.3 Toy examples

Our first illustrative toy example is the following maximization problem:

$$\begin{aligned}
 & \max_{x,y} \min_{a \in [0,1]} ax - y \\
 & \text{subject to} \quad y + b^2 + b \geq 0 \quad \forall b \in [0, 1] \\
 & \quad \quad \quad 0 \leq x \leq 1.
 \end{aligned}
 \tag{toy-1}$$

Let us consider the case where both x and y are nonadjustable. We readily see that the worst-case objective value is 0 and the two solutions, $RO_1 = (1, 0)$ and $RO_2 = (0, 0)$, or any convex combination of these, are worst-case optimal. Without a two-step approach the solver is indifferent between all these optimal robust solutions since they all have optimal worst-case profits. The realized profits as a function of scenario (a, b) are respectively $p_{RO_1}(a, b) = a$ and $p_{RO_2}(a, b) = 0$ and the two-step approach yields solution RO_1 .

Now suppose that y is adjustable and we restrict ourselves to linear decision rules (LDR). Then we find that linear decision rules $y(b) = -b$ or $y(b) = -\frac{1}{2}b$ are optimal in worst-case sense together with any nonadjustable x in $[0, 1]$. For the first solution LDR_1 we take $(x, y) = (1, -b)$ and for the second solution $LDR_2 = (0, -\frac{1}{2}b)$. The profits of these solutions for scenario (a, b) are respectively $p_{LDR_1}(a, b) = a + b$ and $p_{LDR_2}(a, b) = \frac{1}{2}b$. Again, without a two-step approach the solver would be indifferent between these solutions since both have optimal worst-case objective value of 0. The two-step approach yields solution LDR_1 .

Finally, we notice that the so-called *perfect hindsight* solution, where parameters a and b are known before deciding upon x and y , equals $(x, y) = (1, -b^2 - b)$ for any a, b in $[0, 1]$. This perfect hindsight solution can also be obtained in the adjustable robust optimization model by allowing for nonlinear decision rules and setting $NDR_1 = (1, -b^2 - b)$. The profits for this nonlinear decision rule (NDR) are $p_{NDR_1}(a, b) = a + b^2 + b$ for scenario (a, b) . Again, there are many more nonlinear decision rules that are optimal in worst-case sense, but have different mean profits. One example is $NDR_2 = (0, -\frac{1}{2}b^3)$ which yields profit $p_{NDR_2}(a, b) = \frac{1}{2}b^3$. All these results are summarized in Table 4.

In the table we use a uniform distribution to calculate the mean profits. For robust optimization one usually assumes to have only very crude information on the distribution function. Nevertheless, if we denote the mean profits of each solution by \bar{P}_{RO_1} , \bar{P}_{RO_2} , \bar{P}_{LDR_1} , \bar{P}_{LDR_2} , \bar{P}_{NDR_1} and \bar{P}_{NDR_2} , then we have

$$\bar{P}_{NDR_1} > \bar{P}_{LDR_1} > \bar{P}_{RO_1} > \bar{P}_{LDR_2} > \bar{P}_{NDR_2} > \bar{P}_{RO_2}$$

for a large class of distribution functions. All these inequalities are valid if (1) not all probability mass of b lies on the extremes, i.e. $P(b = 0 \text{ or } b = 1) \neq 1$ and (2) the mean value of a and b is such that $\mathbb{E}(a) > \frac{1}{2}\mathbb{E}(b)$.

Note that for this toy example, contrary to the model from [2], there could be a significant gain from the two-step method in the folding horizon approach. The variable x has to be chosen in the first step of the optimization. As we have seen, the optimal robust value is indifferent between any x in $[0, 1]$. In the second step we shall always choose $y = -b^2 - b$. However, choosing $x = 0$ instead of $x = 1$ gives us a difference of a in the objective value. The two-step approach combined with the folding horizon approach returns the optimal (folding horizon) solution, which equals NDR_1 .

Similar to our extended experiments for the numerical production-inventory example, we can also replace the worst-case objective by an expected value objective. Again, we find a unique solutions when using an expected value objective to the following optimization model:

$$\begin{aligned} & \max_{x,y} && \mathbb{E}(a)x - \mathbb{E}(y(b)) \\ & \text{subject to} && y(b) + b^2 + b \geq 0 \quad \forall b \in [0, 1] \quad (\text{toy-1-mean}) \\ & && 0 \leq x \leq 1. \end{aligned}$$

Now, if $\mathbb{E}(a) > 0$, then the solver returns the unique optimal $x = 1$. The only optimal (and unique) static and linear decision rules are given by $y(b) = 0$ and $y(b) = -b$, respectively. These are the same solutions as the best decision rules for the optimization problem with worst-case objective value. For the nonlinear decision rule we find that the optimal decision rule is

$$y(b) = -b - b^2 \quad (\text{almost surely}).$$

Our second toy example is a simple facility location problem with two facilities and a set of customers $\{1, \dots, N\}$. The set of customers is such that the unit transportation costs from facility 1 and facility 2 to customer N are both equal to 10. All other customers are (much) closer to both facilities, but unit transportation costs from facility 2 are significantly smaller than from facility 1. This situation is depicted in Fig. 2.

The demand of the customers is uncertain. In the entire network the demand is at most 1, but we do not know at which customers the demand will occur. We model this via the following uncertainty set:

$$\mathcal{U} = \left\{ (d_1, d_2, \dots, d_N) : d_i \geq 0 \ i = 1, \dots, N, \sum_{i=1}^N d_i \leq 1 \right\},$$

Fig. 2 Facility location problem with the most remote customer N at the same distance from both facilities. The two facilities are depicted by *triangles*, the customers by *circles*



where d_i denotes the uncertain demand of customer i . The facility location problem consists of two types of decisions, namely the decision to open facility 1 ($x_1 = 1$) or facility 2 ($x_2 = 1$) and the actual deliveries to the customers from the opened facility. Only one of the facilities may be opened. The total delivery to customer i from facility 1, respectively facility 2, is y_{1i} and y_{2i} and has unit costs c_{1i} and c_{2i} . The goal is to minimize the worst-case transportation costs, which is modeled as:

$$\begin{aligned}
 &\min_{x,y} \sum_{i=1}^N (c_{1i}y_{1i} + c_{2i}y_{2i}) \\
 &\text{subject to } y_{1i} + y_{2i} \geq d_i \quad \forall i = 1, \dots, N \quad \forall (d_1, d_2, \dots, d_N) \in \mathcal{U} \\
 &\quad y_{1i} \leq x_1 \quad \forall i = 1, \dots, N \quad \text{(toy-2)} \\
 &\quad y_{2i} \leq x_2 \quad \forall i = 1, \dots, N \\
 &\quad x_1 + x_2 \leq 1 \\
 &\quad x_1, x_2 \in \{0, 1\}.
 \end{aligned}$$

From Fig. 2 it is clear that the transportation costs when facility 1 is opened are higher than when facility 2 is opened. The optimal perfect hindsight solution is to open facility 2 and transport exactly the requested demand $y_{2i}(d_i) = d_i$ to each customer. The costs for a particular demand realization (d_1, d_2, \dots, d_N) is then given by

$$\sum_{i=1}^N c_{2i}d_i.$$

The worst-case costs belonging to this solution are

$$\max_{(d_1, d_2, \dots, d_N) \in \mathcal{U}} \sum_{i=1}^N c_{2i}d_i = c_{2N}.$$

In the nonadjustable robust model we decide upon all variables before we know the demand realization d_1, \dots, d_N . The total demand in the network is 1, but all demand could occur at a single customer, so we have to transport one unit to each customer. Therefore, the first constraint in the robust model is equivalent to $y_{1i} + y_{2i} \geq 1$. Since $c_{1i} > c_{2i}$ for all customers $i = 1, \dots, N - 1$, the optimal solution is $x_1 = 0, x_2 = 1$ with $y_{1i} = 0, y_{2i} = 1$ for all $i = 1, \dots, N$ and objective value $\sum_{i=1}^N c_{2i}$. The robust solution vastly overestimates the worst-case costs, but it does open facility 2. In the folding horizon approach, the transportation decisions are re-optimized and we obtain $y_{2i} = d_i$ with costs $\sum_{i=1}^N c_{2i}d_i$, which equal the costs in the perfect hindsight solution.

In the adjustable robust model there are multiple optimal solutions. In the first solution we open facility 1 and transport $y_{1i} = d_i$ to customer $i = 1, \dots, N$. In the second solution we open facility 2 and transport $y_{2i} = d_i$ to each customer. Clearly, we obtain the same worst-case costs $c_{1N} = c_{2N}$ as in the perfect hindsight case. However, the costs when (d_1, d_2, \dots, d_N) realizes equals $\sum_{i=1}^N c_{1i}d_i$ and $\sum_{i=1}^N c_{2i}d_i$ respectively. If the expected demand is $d_i = \frac{1}{N}$ for all $i = 1, \dots, N$, or any other scenario that does not place all probability mass on the demand realization with $d_N = 1$, then the two-step approach picks the solution that opens facility 2.

To conclude, in the first toy-example the here-and-now decisions matter in the folding horizon approach for the costs, but there is no impact of the existence of multiple here-and-now decisions on the choice of the optimal wait-and-see decision in the re-optimization step. In the second toy example we do see an impact: once the wrong facility is opened in the first stage, all demand has to be fulfilled from that location at high expected costs in the re-optimization step.

3 Implications for robust optimization

The inventory-production problem and the toy examples from the previous section allow us to present some important implications. First, if we analyze and compare the mean objective values of arbitrary optimal robust solutions for RO and ARO, then false conclusions can be drawn regarding the added value of ARO over RO. The mean objective value of an *arbitrary* optimal robust solution, obtained by solving the original RO or ARO problem formulations, might very well be much worse than the solution with best mean objective value among all optimal robust solutions. This best performing solution can be obtained by carrying out the two-step approach. In the production-inventory problem with uncertainty level 2.5%, the worst-case objective values of the RO and ARO solution are nearly the same: the difference is only 0.5%. If we compare the RO and ARO solutions on average costs, then the worst ARO solution is also only 0.5% better than the RO solution. The best ARO solution, however, is 3.5% better on average, which could be overlooked if the two-step approach is not carried out. For the 20% uncertainty level, the gap between the average performances of all optimal robust ARO solutions can be as much as 21.9%. The first toy example illustrates that an arbitrary ARO solution is not guaranteed to do better than a RO solution with respect to average performance. For instance, the average performance of robust solution RO_1 is better than the performance of ARO solution LDR_2 . On the other hand, the optimal ARO solution LDR_1 is guaranteed to do better than any

RO solution on the average performance. In our small facility location example we have seen that the robust solution results in a much higher objective value, but that it does open the best facility for folding horizon approaches. The linear decision rule on the other hand results in multiple optimal solutions which could lead to undesirable choices for opening the facilities. The two-step approach results in a solution that opens the cheapest facility, mimicking the solution of perfect hindsight.

Second, one might be inclined to jump to the conclusion that ARO can be safely ignored, when it is a priori known that ARO and RO are equivalent with respect to the *worst-case* objective value. One of the situations that we know where ARO is equivalent to RO is the case of constraint-wise uncertainty see [2, Theorem 2.1]. However, the equivalence is not necessarily true for the mean objective value as well. Therefore, one should not ignore ARO for such problems. This is illustrated by the first toy example: the worst-case objective value is zero for both the RO and ARO solutions, but the mean objective values differ significantly.

Third, even if affine decision rules yield (near) optimal worst-case performance, nonlinear decision rules, such as quadratic decision rules, can yield much better performance on the mean objective value. Most applications of ARO restrict decision rules to affine functions, which is referred to as affinely adjustable robust optimization (AARO) [2]. Affine decision rules are known to perform optimal or nearly optimal in many situations [1,4]. However, once again, this observation is with respect to the worst-case objective value, and not for the mean objective value. This is illustrated by the first toy example. Here, the quadratic decision rule NDR_1 has the same worst-case objective value as any of the other decision rules, but the mean objective value is much better, and, in this particular case, even optimal for each scenario (Bellman optimal).

The encompassing recommendation that follows from these implications is that the two-step approach should always be conducted in any application of robust optimization. The two-step approach enables the optimizer to fully exploit the performance on the mean objective value of the solution, while guaranteeing no deterioration in the worst-case performance. This is especially relevant for ARO, where decision rules can be utilized to enhance the solution's performance in other than worst-case scenarios. We also recommend the use of the two step approach in folding horizon methods, but we do note that the impact of multiple solutions may be less severe.

Acknowledgments The authors thank the associate editor and two anonymous referees for their useful comments.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Ben-Tal, A., El Ghaoui, L., Nemirovski, A.: Robust Optimization. Princeton Series in Applied Mathematics. Princeton University Press, Princeton (2009)

2. Ben-Tal, A., Goryashko, A., Guslitzer, E., Nemirovski, A.: Adjustable robust solutions of uncertain linear programs. *Math. Program.* **99**(2), 351–376 (2004)
3. Bertsimas, D., Brown, D.B., Caramanis, C.: Theory and applications of robust optimization. *SIAM Rev.* **53**(3), 464–501 (2011)
4. Bertsimas, D., Iancu, D.A., Parrilo, P.A.: Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.* **35**(2), 363–394 (2010)
5. Charnes, A., Cooper, W.W.: Programming with linear fractional functionals. *Nav. Res. Logist. Q* **9**(3–4), 181–186 (1962)
6. Gabrel, V., Murat, C., Thiele, A.: Recent advances in robust optimization: an overview. *Eur. J. Oper. Res.* **235**(3), 471–483 (2014)
7. Gurobi Optimization Inc: Gurobi Optimizer Reference Manual (2015). <http://www.gurobi.com>
8. Iancu, D.A., Trichakis, N.: Pareto efficiency in robust optimization. *Manag. Sci.* **60**(1), 130–147 (2013)
9. Kuhn, D., Wiesemann, W., Georghiou, A.: Primal and dual linear decision rules in stochastic and robust optimization. *Math. Program.* **130**(1), 177–209 (2011)
10. Löfberg, J.: Yalmip—a toolbox for modeling and optimization in MATLAB. In: Proceedings of the CACSD Conference. Taipei, Taiwan (2004)