



INFORMS Journal on Computing

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Duality in Two-Stage Adaptive Linear Optimization: Faster Computation and Stronger Bounds

Dimitris Bertsimas, Frans J. C. T. de Ruiter

To cite this article:

Dimitris Bertsimas, Frans J. C. T. de Ruiter (2016) Duality in Two-Stage Adaptive Linear Optimization: Faster Computation and Stronger Bounds. INFORMS Journal on Computing 28(3):500-511. <http://dx.doi.org/10.1287/ijoc.2016.0689>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2016, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Duality in Two-Stage Adaptive Linear Optimization: Faster Computation and Stronger Bounds

Dimitris Bertsimas

Operations Research Center and Sloan School of Management, Massachusetts Institute of Technology, Cambridge,
Massachusetts 02139, dbertsim@mit.edu

Frans J. C. T. de Ruiter

Department of Econometrics and Operations Research, Tilburg University, 5000 LE Tilburg, Netherlands,
f.j.c.t.deruiter@tilburguniversity.edu

In this paper we derive and exploit duality in general two-stage adaptive linear optimization models. The equivalent dualized formulation we derive is again a two-stage adaptive linear optimization model. Therefore, all existing solution approaches for two-stage adaptive models can be used to solve or approximate the dual formulation. The new dualized model differs from the primal formulation in its dimension and uses a different description of the uncertainty set. We show that the optimal primal affine policy can be directly obtained from the optimal affine policy in the dual formulation. We provide empirical evidence that the dualized model in the context of two-stage lot-sizing on a network and two-stage facility location problems solves an order of magnitude faster than the primal formulation with affine policies. We also provide an explanation and associated empirical evidence that offer insight on which characteristics of the dualized formulation make computations faster. Furthermore, the affine policy of the dual formulations can be used to provide stronger lower bounds on the optimality of affine policies.

Keywords: two-stage problems; robust optimization; duality; affine control policies

History: Accepted by Karen Aardal, Area Editor for Design and Analysis of Algorithms; received August 2015; accepted December 2015. Published online May 24, 2016.

1. Introduction

Many applications for decision making under uncertainty can be naturally modeled as two-stage adaptive optimization models. In these models some of the decisions have to be made *here-and-now* before the realization of the uncertain parameter is known. The other decisions are of a *wait-and-see* type, which are chosen after the realization of the uncertain parameter is known. One way of dealing with these problems is via stochastic optimization. These methods assume that a probabilistic description of the realization is known and optimize for expected values. For references on these techniques we refer to Birge and Louveaux (2011) and Kali and Wallace (1994). Stochastic models, especially in a two-stage setting, are known to suffer from the “curse of dimensionality” and are therefore likely not tractable; see e.g., Shapiro and Nemirovski (2005). A different approach is to model these two-stage problems in a robust setting. Robust optimization techniques do not require a probabilistic description of the uncertainty set and have proven to be very useful in a number of practical applications. A selection of applications that use a two-stage robust setting are: unit commitment in the energy sector (Bertsimas et al. 2013, Wang et al. 2013, Zhao and Zeng 2012), emergency supply chain

planning (Ben-Tal et al. 2011), facility location problems (Ardestani-Jaafari and Delage 2014, Atamtürk and Zhang 2007, Gabrel et al. 2014a), capacity expansion of network flows (Ordóñez and Zhao 2007, Yin et al. 2009), and many others; see e.g., the survey papers by Bertsimas et al. (2011) and Gabrel et al. (2014b).

In the last decade or so, there has been a rise in solution techniques tailored to solve two-stage optimization models in a robust setting. One of the first and very popular methods is the use of affine policies for the wait-and-see decisions proposed by Ben-Tal et al. (2004). This method is appealing because it is computationally tractable for problem instances of moderate to large size. Furthermore, the affine policies appear to be near optimal in practical applications (Ardestani-Jaafari and Delage 2014; Ben-Tal et al. 2004, 2005). The use of affine policies is even provably optimal in some special cases (Bertsimas et al. 2010, Iancu et al. 2013). Other methods designed to solve two-stage adaptive optimization models are: approximation by static solutions (Bertsimas and Goyal 2010), finite adaptability (Bertsimas and Caramanis 2010), enumeration of vertices of the uncertainty set (Bertsimas and Goyal 2012), column generation algorithms (Zeng and Zhao 2013), and iterative partitioning of the uncertainty set (Postek and Den Hertog 2016, Bertsimas and Dunning 2016).

In this paper we derive a new dualized formulation of two-stage adaptive linear models that allow for faster computations and stronger bounds. More specifically, the main contributions of this paper can be summarized as follows:

1. We provide a dualized two-stage adaptive model for linear two-stage models with continuous wait-and-see decisions. The new model is derived by consecutively dualizing over the wait-and-see decisions and the uncertain parameters. The new dualized formulations have the same set of feasible (and optimal) here-and-now decisions as the original two-stage models. It has different dimensions, uncertain parameters, wait-and-see decisions, and constraints than the original two-stage adaptive model. Since the model is again a two-stage adaptive model, all existing solution techniques for two-stage adaptive models can be used to solve it.

2. We show that both formulations also have the same set of feasible and optimal here-and-now decisions when we solve the models using the popular method of affine policies. Furthermore, we show how the original affine policy can be obtained instantly from the affine policy in the dualized formulation.

3. We describe an algorithm to strengthen the lower bound method from Hadjiyiannis et al. (2011) to assess the (sub)optimality of affine policies described using both affine policies from the original and the dualized formulation.

4. We provide empirical evidence that the dualized model in the context of two-stage lot-sizing on a network and two-stage facility location problems solves an order of magnitude faster than the primal formulation with affine policies and provides stronger lower bounds. Furthermore, we provide an explanation and associated empirical evidence that offer insight on which characteristics of the dualized formulation make computations faster.

Our dualized formulation can be used for general two-stage adaptive linear models with both continuous and integer here-and-now decisions. However, since we dualize over the second stage variables, the new dualized formulation only works for continuous second stage decisions. Furthermore, to end up with tractable models, our method focuses on polyhedral uncertainty sets.

The rest of this paper is organized as follows. In Section 2, we introduce the two-stage adaptive optimization model and derive the new dualized two-stage model. We explain the use of affine policies in the primal and dual formulation in Section 3. Section 4 gives the computational algorithm to obtain stronger bounds on the optimal value of the fully adaptive model. In Sections 5 and 6, we present our numerical results and show the computational advantage of the dualized formulation. Section 7 gives some concluding remarks.

Notation. Throughout the paper we write vectors and matrices in bold font and scalars in normal font. We use the vector \mathbf{e} to denote the vector of all ones and \mathbf{I} for the identity matrix. The vector $\mathbf{0}$ and matrix \mathbf{O} consist of only zero entries. All inequality signs represent componentwise inequalities.

2. Duality in Two-Stage Adaptive Formulations

We first state the usual two-stage formulation in Section 2.1. The new dualized formulation is given in Section 2.2. We also indicate similarities in structure with the primal formulation and the differences in the two formulations.

2.1. The Primal Formulation

We consider a general two-stage adaptive optimization model with continuous wait-and-see decisions. In the first stage we set the value of the here-and-now decisions \mathbf{x} that have to be decided before the realization of the uncertain parameter is known. The continuous wait-and-see decisions $\mathbf{y} \geq \mathbf{0}$ have to be chosen after the value of the uncertain parameter is revealed. We take a polyhedral description of the uncertainty set of the form

$$\mathcal{U} = \{\boldsymbol{\zeta} \geq \mathbf{0} : \mathbf{D}\boldsymbol{\zeta} \leq \mathbf{d}\}, \quad (1)$$

with $\mathbf{D} \in \mathbb{R}^{p \times L}$ and $\mathbf{d} \in \mathbb{R}^p$. This type of uncertainty sets includes popular sets such as the box-uncertainty and budget uncertainty set (Bertsimas and Sim 2004). The two-stage adaptive optimization problem has a linear objective and a set of linear uncertain constraints. With this general setting we can state the following description of a two-stage adaptive linear optimization model:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \forall \boldsymbol{\zeta} \in \mathcal{U}: \exists \mathbf{y} \geq \mathbf{0}: \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{y} \geq \mathbf{R}\boldsymbol{\zeta} + \mathbf{r} \\ & \mathbf{x} \in \mathcal{X}, \end{aligned} \quad (2)$$

where $\mathcal{X} \subset \mathbb{R}^n$ is a set with additional constraints on the here-and-now decisions (some of the \mathbf{x} variables may be integer). The wait-and-see variable \mathbf{y} has dimension k and we denote the number of constraints in the model by m . The matrix \mathbf{R} is chosen constant in this model, so the model only has uncertainty in the right-hand side. This is mainly done for exposition and all our results can be extended to the case where \mathbf{R} depends on the here-and-now decision \mathbf{x} , for example by taking

$$\mathbf{R}(\mathbf{x}) = \mathbf{R}_0 + \sum_{i=1}^n \mathbf{R}_i x_i$$

for some matrices $\mathbf{R}_0, \mathbf{R}_1, \dots, \mathbf{R}_n$. For our dual derivation to work, we must have the matrix \mathbf{B} to be fixed independent of $\boldsymbol{\zeta}$. Hence, we only consider the case

of *fixed recourse*. Without loss of generality, there is no uncertainty in the objective function and it only includes here-and-now decisions. Objectives including uncertain parameters and wait-and-see decisions can be modelled as an instance of (2) using an epigraph formulation; see Ben-Tal et al. (2009, pp. 10–11). These epigraph formulations are also used in the models of our numerical examples in Sections 5 and 6.

2.2. The New Dualized Formulation

The main contributions of this paper come from Theorem 1, giving a dual formulation of (2).

THEOREM 1. *The here-and-now decision \mathbf{x} is feasible (and optimal) for (2) with nonempty uncertainty set \mathcal{U} as in (1) if and only if \mathbf{x} is feasible (and optimal) for*

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \forall \mathbf{w} \in \mathcal{V}: \exists \boldsymbol{\lambda} \geq \mathbf{0}: \begin{cases} \mathbf{w}^\top (\mathbf{A}\mathbf{x} - \mathbf{r}) - \mathbf{d}^\top \boldsymbol{\lambda} \geq 0, \\ \mathbf{D}^\top \boldsymbol{\lambda} \geq \mathbf{R}^\top \mathbf{w}, \end{cases} \\ & \mathbf{x} \in \mathcal{X}, \end{aligned} \quad (3)$$

where $\mathcal{V} = \{\mathbf{w} \geq \mathbf{0}: \mathbf{B}^\top \mathbf{w} \leq \mathbf{0}, \mathbf{e}^\top \mathbf{w} = 1\}$.

The proof of this theorem is split into two parts. The first part comes from a result known in the literature and the second part is the new contribution leading to the dualized formulation. The result from the literature transforms (2) into a bilinear optimization model by applying duality to the wait-and-see variables. The result from this part is used frequently in the literature, in various settings, to solve two-stage adaptive optimization problems using column generation and Benders decomposition type algorithms (see e.g., Bertsimas et al. 2013, Minoux 2011, Thiele et al. 2009, Zeng and Zhao 2013, and Zhao and Zeng 2012) or to derive an exact solution for special cases (Ordóñez and Zhao 2007). This known result is given in Lemma 1.

LEMMA 1. *The here-and-now decision \mathbf{x} is feasible (and optimal) for (2) if and only if \mathbf{x} is feasible (and optimal) for*

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\boldsymbol{\zeta} \in \mathcal{U}} \max_{\mathbf{w} \geq \mathbf{0}} \{ \mathbf{c}^\top \mathbf{x} + \mathbf{w}^\top (\mathbf{R}\boldsymbol{\zeta} + \mathbf{r} - \mathbf{A}\mathbf{x}) \mid \mathbf{B}^\top \mathbf{w} \leq \mathbf{0} \}. \quad (4)$$

PROOF. For a given $\mathbf{x} \in \mathcal{X}$ we can write (2) as

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\boldsymbol{\zeta} \in \mathcal{U}} \min_{\mathbf{y} \geq \mathbf{0}} \{ \mathbf{c}^\top \mathbf{x} \mid \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{y} \geq \mathbf{R}\boldsymbol{\zeta} + \mathbf{r} \}.$$

The result then follows by dualizing over \mathbf{y} . \square

Note that for every $\boldsymbol{\zeta}$ the variable \mathbf{w} ensures that the problem returns ∞ whenever there exists a $\boldsymbol{\zeta}$ that violates the constraints in the original model (2). The result from Lemma 1 is also used in Kuhn et al. (2011) to assess the suboptimality of affine policies in a two-stage *stochastic* setting. Their bound can also be used in

robust settings, but one has to assign a distribution to the uncertainty set a priori. The authors explain that in that case the quality of the bound depends on the a priori distribution that is chosen. For the rest of the proof we first dualize (4) further to end up with an equivalent two-stage adaptive optimization formulation.

PROOF OF THEOREM 1. Consider, for fixed \mathbf{w} , the inner maximization problem in (4). Dualizing over $\boldsymbol{\zeta}$ gives

$$\begin{aligned} \min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{w} \geq \mathbf{0}} \min_{\boldsymbol{\lambda} \geq \mathbf{0}} \{ & \mathbf{c}^\top \mathbf{x} + \mathbf{w}^\top (\mathbf{r} - \mathbf{A}\mathbf{x}) + \mathbf{d}^\top \boldsymbol{\lambda} \mid \mathbf{D}^\top \boldsymbol{\lambda} \geq \mathbf{R}^\top \mathbf{w}, \\ & \mathbf{B}^\top \mathbf{w} \leq \mathbf{0} \} \\ = \min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbf{w} \in \tilde{\mathcal{V}}} \min_{\boldsymbol{\lambda} \geq \mathbf{0}} \{ & \mathbf{c}^\top \mathbf{x} + \mathbf{w}^\top (\mathbf{r} - \mathbf{A}\mathbf{x}) \\ & + \mathbf{d}^\top \boldsymbol{\lambda} \mid \mathbf{D}^\top \boldsymbol{\lambda} \geq \mathbf{R}^\top \mathbf{w} \}, \end{aligned} \quad (5)$$

where in the last line we introduced $\tilde{\mathcal{V}} = \{\mathbf{w} \geq \mathbf{0}: \mathbf{B}^\top \mathbf{w} \leq \mathbf{0}\}$. Introducing a variable γ we write the model using an epigraph formulation

$$\begin{aligned} \min_{\mathbf{x}, \gamma} \quad & \{ \mathbf{c}^\top \mathbf{x} + \gamma \} \\ \text{s.t.} \quad & \forall \mathbf{w} \in \tilde{\mathcal{V}}: \exists \boldsymbol{\lambda} \geq \mathbf{0}: \begin{cases} \mathbf{w}^\top (\mathbf{r} - \mathbf{A}\mathbf{x}) + \mathbf{d}^\top \boldsymbol{\lambda} \leq \gamma, \\ \mathbf{D}^\top \boldsymbol{\lambda} \geq \mathbf{R}^\top \mathbf{w}, \end{cases} \\ & \mathbf{x} \in \mathcal{X}. \end{aligned} \quad (6)$$

To end up with our final result (3) we have to prove that $\gamma = 0$ for any optimal solution and that we can add the additional restriction $\mathbf{e}^\top \mathbf{w} = 1$ to bound the uncertainty set $\tilde{\mathcal{V}}$ without affecting the set of feasible solutions. From (5) it follows that there has to be an optimal adaptive policy $\boldsymbol{\lambda}^*(\mathbf{w})$ that satisfies

$$\mathbf{d}^\top (\boldsymbol{\lambda}^*(\mathbf{w})) = \min_{\boldsymbol{\lambda} \geq \mathbf{0}} \{ \mathbf{d}^\top \boldsymbol{\lambda} \mid \mathbf{D}^\top \boldsymbol{\lambda} \geq \mathbf{R}^\top \mathbf{w} \}.$$

Note that this policy is not only worst-case optimal, but chooses the best wait-and-see decision $\boldsymbol{\lambda}(\mathbf{w})$ for every scenario \mathbf{w} . For the scenario $\mathbf{0} \in \tilde{\mathcal{V}}$ we have

$$\begin{aligned} \mathbf{d}^\top (\boldsymbol{\lambda}^*(\mathbf{0})) &= \min_{\boldsymbol{\lambda} \geq \mathbf{0}} \{ \mathbf{d}^\top \boldsymbol{\lambda} \mid \mathbf{D}^\top \boldsymbol{\lambda} \geq \mathbf{0} \} \\ &= \min_{\boldsymbol{\lambda} \geq \mathbf{0}} \max_{\boldsymbol{\zeta} \geq \mathbf{0}} \{ \boldsymbol{\lambda}^\top (\mathbf{d} - \mathbf{D}\boldsymbol{\zeta}) \} = 0, \end{aligned}$$

where the last equality holds since $\mathcal{U} = \{\boldsymbol{\zeta} \geq \mathbf{0}: \mathbf{D}\boldsymbol{\zeta} \leq \mathbf{d}\}$ is nonempty. Using this optimal decision for the parameter $\mathbf{0} \in \tilde{\mathcal{V}}$, we see that

$$\gamma \geq \mathbf{0}^\top (\mathbf{r} - \mathbf{A}\mathbf{x}) + \mathbf{d}^\top \boldsymbol{\lambda}^*(\mathbf{0}) = 0. \quad (7)$$

Now, let $t \geq 0$ and $\mathbf{w} \geq \mathbf{0}$. Then we have

$$\begin{aligned} \mathbf{d}^\top (\boldsymbol{\lambda}^*(t\mathbf{w})) &= \min_{\boldsymbol{\lambda} \geq \mathbf{0}} \{ \mathbf{d}^\top \boldsymbol{\lambda} \mid \mathbf{D}^\top \boldsymbol{\lambda} \geq \mathbf{R}^\top (t\mathbf{w}) \} \\ &= \min_{\boldsymbol{\lambda} \geq \mathbf{0}} \{ \mathbf{d}^\top (t\boldsymbol{\lambda}) \mid \mathbf{D}^\top \boldsymbol{\lambda} \geq \mathbf{R}^\top \mathbf{w} \} = \mathbf{d}^\top (t\boldsymbol{\lambda}^*(\mathbf{w})). \end{aligned}$$

Hence, we can impose scalar multiplicity on the adaptive policy $\boldsymbol{\lambda}^*(\mathbf{w})$ without affecting the value of

Table 1 Comparing Dimensions of Uncertainty Parameters, Variables, and Number of Constraints in the Original Two-Stage Adaptive Formulation (2) and Our New Dualized Formulation (3)

	Primal formulation (2)	Dual formulation (3)
No. of uncertain parameters	L	m
No. of wait-and-see decisions	k	p
No. of constraints on variables	m	$L + 1$
No. of constraints on uncertain parameter	p	$k + 1$

$\mathbf{d}^\top(\boldsymbol{\lambda}^*(\mathbf{w}))$. That is, for every $\mathbf{w} \in \tilde{\mathcal{V}}$ and scalar $t \geq 0$ we impose $\boldsymbol{\lambda}^*(t\mathbf{w}) = t\boldsymbol{\lambda}^*(\mathbf{w})$. From (7) we have that $\gamma \geq 0$. Suppose for the sake of contradiction that for an optimal here-and-now decision (\mathbf{x}, γ) we have $\gamma > 0$. Then there exist an $\mathbf{w} \in \tilde{\mathcal{V}}$, $\mathbf{w} \neq \mathbf{0}$ such that

$$\mathbf{w}^\top(\mathbf{r} - \mathbf{A}\mathbf{x}) + \mathbf{d}^\top\boldsymbol{\lambda}^*(\mathbf{w}) = \gamma > 0.$$

Since $\tilde{\mathcal{V}}$ is a cone, we have that $(t\mathbf{w}) \in \tilde{\mathcal{V}}$ for every $t \geq 0$ and $\mathbf{w} \in \tilde{\mathcal{V}}$. Therefore, we have by scalar multiplicity of $\boldsymbol{\lambda}^*(\mathbf{w})$:

$$(t\mathbf{w})^\top(\mathbf{r} - \mathbf{A}\mathbf{x}) + \mathbf{d}^\top\boldsymbol{\lambda}^*(t\mathbf{w}) = t\gamma > \gamma, \quad \text{for all } t > 1.$$

This contradicts the assumption that $\gamma > 0$ is feasible. Hence, we must have $\gamma = 0$. Finally, consider a solution that is feasible for all values in the further restricted uncertainty set

$$\begin{aligned} \mathcal{V} &= \{\mathbf{w} \geq \mathbf{0}: \mathbf{B}^\top\mathbf{w} \leq \mathbf{0}, \|\mathbf{w}\|_1 = 1\} \\ &= \{\mathbf{w} \geq \mathbf{0}: \mathbf{B}^\top\mathbf{w} \leq \mathbf{0}, \mathbf{e}^\top\mathbf{w} = 1\}. \end{aligned}$$

Then, by scalar multiplicity of $\boldsymbol{\lambda}^*(\mathbf{w})$, we can directly construct the other feasible wait-and-see decisions for all other $\mathbf{w} \in \tilde{\mathcal{V}}$ (with $\|\mathbf{w}\|_1 \neq 1$). \square

Any two-stage adaptive optimization model with fixed recourse, continuous wait-and-see decisions and a polyhedral uncertainty set can be readily formulated as an instance of (2). Theorem 1 then directly provides practitioners with the alternative dual formulation (3). Table 1 highlights some differences such as the number of wait-and-see variables, uncertain parameters and constraints in the primal and dual formulation. In our numerical examples in Section 5 and 6 we clarify these differences with explicit values for m, k, L , and p .

3. Solving the Primal and Dual Formulation with Affine Policies

The model (3) is again a two-stage adaptive optimization model with a nonnegative bounded polyhedral uncertainty set and is therefore another instance of (2). Hence, we can directly apply all exact and approximation methods to solve adaptive optimization problems

mentioned in the introduction. We first show the equivalence of the dual formulation with the nonadaptive robust counterpart in the static case. We then continue to show that the optimal solutions of both formulations are the same when we solve the models with affine policies.

3.1. Static Robust Optimization

If we take $\mathbf{B} = \mathbf{O}$, then (2) is the following robust optimization model without wait-and-see decisions:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{c}^\top\mathbf{x} \\ \text{s.t.} \quad & \forall \boldsymbol{\zeta} \in \mathcal{U}: \mathbf{A}\mathbf{x} \geq \mathbf{R}\boldsymbol{\zeta} + \mathbf{r}, \\ & \mathbf{x} \in \mathcal{X}, \end{aligned} \tag{8}$$

where \mathcal{U} is as in (1). This problem is hard to solve in its current form since each constraint has to hold for an infinite number of values for $\boldsymbol{\zeta}$. To reformulate the problem, we can consider the uncertainty constraint-wise (see Ben-Tal et al. 2009), i.e., we only have to look at one row

$$\forall \boldsymbol{\zeta} \in \mathcal{U}: \mathbf{A}_i\mathbf{x} \geq \mathbf{R}_i\boldsymbol{\zeta} + r_i \tag{9}$$

at a time, where $\mathbf{A}_i, \mathbf{R}_i$, and r_i are respectively the i th row of \mathbf{A}, \mathbf{R} , and \mathbf{r} . To make this model tractable we can reformulate each constraint using standard duality techniques to obtain the robust counterpart; see e.g., Ben-Tal et al. (2009).

LEMMA 2 (ROBUST COUNTERPART). *Constraint (9) is satisfied if and only if there exists a $\boldsymbol{\pi}^i \in \mathbb{R}^p$ such that*

$$\begin{aligned} \mathbf{A}_i\mathbf{x} - \boldsymbol{\pi}^{i\top}\mathbf{d} &\geq r_i, \\ \mathbf{D}^\top\boldsymbol{\pi}^i &\geq \mathbf{R}_i, \\ \boldsymbol{\pi}^i &\geq \mathbf{0}. \end{aligned}$$

Note that this dualization approach can also be used for any other polyhedral uncertainty set. For notational convenience we shall use matrix variables for the rest of the section. If we write $\boldsymbol{\Pi} = [\boldsymbol{\pi}^1, \dots, \boldsymbol{\pi}^m]$, then by Lemma 2 we have that (8) is equivalent to

$$\begin{aligned} \min_{\mathbf{x}, \boldsymbol{\Pi}} \quad & \mathbf{c}^\top\mathbf{x}, \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} - \boldsymbol{\Pi}^\top\mathbf{d} \geq \mathbf{r}, \\ & \mathbf{D}^\top\boldsymbol{\Pi} \geq \mathbf{R}^\top, \\ & \mathbf{x} \in \mathcal{X}, \boldsymbol{\Pi} \geq \mathbf{O}. \end{aligned} \tag{10}$$

We can also find a dual formulation for the static model (8) using the dual formulation that is derived in Theorem 1. In that way, we end up with the same dual formulation as in (3), but with the simple uncertainty set

$$\mathcal{V} = \{\mathbf{w} \geq \mathbf{0}: \mathbf{e}^\top\mathbf{w} = 1\}. \tag{11}$$

For these robust models with $\mathbf{B} = \mathbf{O}$ the uncertainty set (11) has only m extreme points $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_m$. As shown in Bertsimas and Goyal (2012, Lemma 1), linear policies are optimal if there are only m extreme points, where m is the size of the uncertainty set. Furthermore, by taking the linear policy $\boldsymbol{\lambda}(\mathbf{w}) = \mathbf{\Pi}\mathbf{w}$ in (3) we end up with the same robust counterpart as (10).

3.2. Solving the Two-Stage Formulations with Affine Policies

Let us now return to the general case in which $\mathbf{B} \neq \mathbf{O}$, so we do need to take the wait-and-see decisions \mathbf{y} into account. In principle, an optimal policy $\mathbf{y}(\boldsymbol{\zeta})$ in (2) can be any function of the uncertain parameter $\boldsymbol{\zeta}$. However, this results in an intractable model where we would have to optimize over all possible functions. To come up with tractable models, Ben-Tal et al. (2004) suggest restricting the wait-and-see decisions to be affine in $\boldsymbol{\zeta}$:

$$\mathbf{y}(\boldsymbol{\zeta}) = \mathbf{u} + \mathbf{V}\boldsymbol{\zeta},$$

where $\mathbf{u} \in \mathbb{R}^m$ and $\mathbf{V} \in \mathbb{R}^{m \times L}$ are respectively a vector and a matrix of here-and-now variables. Although this restriction might seem very severe, it performs very well in practical applications (Ben-Tal et al. 2004, 2005), and is even provably optimal in some specific cases (Bertsimas et al. 2010, Iancu et al. 2013). With this decision rule, we obtain the following robust counterpart for (2) with affine policies:

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{u}, \mathbf{V}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \forall \boldsymbol{\zeta} \in \mathcal{U}: \begin{cases} \mathbf{A}\mathbf{x} + \mathbf{B}(\mathbf{u} + \mathbf{V}\boldsymbol{\zeta}) \geq \mathbf{R}\boldsymbol{\zeta} + \mathbf{r}, \\ \mathbf{u} + \mathbf{V}\boldsymbol{\zeta} \geq \mathbf{0}, \end{cases} \\ & \mathbf{x} \in \mathcal{X}. \end{aligned} \tag{12}$$

This model does not have wait-and-see variables. Therefore, we can apply Lemma 2 to reformulate each constraint and obtain the robust counterpart. Introducing the auxiliary (matrix) variables $\mathbf{\Pi} \in \mathbb{R}^{p \times m}$ and $\mathbf{\Xi} \in \mathbb{R}^{p \times k}$ we can write down the robust counterpart as

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{u}, \mathbf{V}, \mathbf{\Pi}, \mathbf{\Xi}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} - \mathbf{\Pi}^\top \mathbf{d} \geq \mathbf{r}, \\ & \mathbf{B}\mathbf{V} \geq \mathbf{R} - \mathbf{\Pi}^\top \mathbf{D}, \\ & \mathbf{u} - \mathbf{\Xi}^\top \mathbf{d} \geq \mathbf{0}, \\ & \mathbf{D}^\top \mathbf{\Xi} + \mathbf{V}^\top \geq \mathbf{0}, \\ & \mathbf{\Pi}, \mathbf{\Xi} \geq \mathbf{O}, \\ & \mathbf{x} \in \mathcal{X}. \end{aligned} \tag{13}$$

For the dualized formulation we can also impose linear restrictions, i.e.,

$$\boldsymbol{\lambda}(\mathbf{w}) = \mathbf{Q}\mathbf{w}, \tag{14}$$

where we now introduce here-and-now variables $\mathbf{Q} \in \mathbb{R}^{p \times m}$ to construct the decision rule. Note that we restricted ourselves to *linear* policies in the dual formulation instead of affine policies. However, leaving out the constant term does not restrict the set of feasible and optimal here-and-now decisions as follows from the next proposition.

PROPOSITION 1. *If $(\mathbf{x}, \boldsymbol{\lambda}(\mathbf{w}) = \mathbf{q} + \mathbf{Q}\mathbf{w})$ is feasible for (3), then $(\mathbf{x}, \tilde{\boldsymbol{\lambda}}(\mathbf{w}) = \tilde{\mathbf{Q}}\mathbf{w})$ with $\tilde{\mathbf{Q}} = \mathbf{q}\mathbf{e}^\top + \mathbf{Q}$ is also feasible.*

PROOF. For all $\mathbf{w} \in \mathcal{W}$ we have $\mathbf{e}^\top \mathbf{w} = 1$. Therefore, for all $\mathbf{w} \in \mathcal{W}$ the following relation holds

$$\tilde{\boldsymbol{\lambda}}(\mathbf{w}) = \tilde{\mathbf{Q}}\mathbf{w} = (\mathbf{q}\mathbf{e}^\top + \mathbf{Q})\mathbf{w} = \mathbf{q} + \mathbf{Q}\mathbf{w} = \boldsymbol{\lambda}(\mathbf{w}).$$

Hence, if $\boldsymbol{\lambda}(\mathbf{w})$ is a feasible policy for (15), then so is $\tilde{\boldsymbol{\lambda}}(\mathbf{w})$. \square

Substituting the linear policy (14) in (3), we obtain the following model:

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{Q}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \forall \mathbf{w} \in \mathcal{W}: \begin{cases} \mathbf{w}^\top (\mathbf{A}\mathbf{x} - \mathbf{r}) - \mathbf{d}^\top (\mathbf{Q}\mathbf{w}) \geq 0, \\ \mathbf{D}^\top \mathbf{Q}\mathbf{w} \geq \mathbf{R}^\top \mathbf{w}, \\ \mathbf{Q}\mathbf{w} \geq \mathbf{0}, \end{cases} \\ & \mathbf{x} \in \mathcal{X}. \end{aligned} \tag{15}$$

A robust counterpart for (15) can be derived using standard linear programming dualization as in Lemma 2. With the introduction of the auxiliary variables $\boldsymbol{\epsilon} \in \mathbb{R}^k$, $\boldsymbol{\Lambda} \in \mathbb{R}^{k \times L}$, and $\boldsymbol{\Omega} \in \mathbb{R}^{k \times p}$, the resulting robust counterpart can be written as

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{Q}, \boldsymbol{\epsilon}, \boldsymbol{\Omega}, \boldsymbol{\Lambda}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} + \mathbf{B}\boldsymbol{\epsilon} - \mathbf{Q}^\top \mathbf{d} \geq \mathbf{r}, \\ & \mathbf{B}\boldsymbol{\Lambda} \geq \mathbf{R} - \mathbf{Q}^\top \mathbf{D}, \\ & \mathbf{B}\boldsymbol{\Omega} + \mathbf{Q}^\top \geq \mathbf{0}, \\ & \boldsymbol{\epsilon} \geq \mathbf{0}, \boldsymbol{\Lambda}, \boldsymbol{\Omega} \geq \mathbf{O}, \\ & \mathbf{x} \in \mathcal{X}. \end{aligned} \tag{16}$$

Theorem 2 shows that the primal and dual formulation have the same set of feasible (and optimal) here-and-now decisions.

THEOREM 2. *The solution $(\mathbf{x}, \mathbf{Q}, \boldsymbol{\epsilon}, \boldsymbol{\Omega}, \boldsymbol{\Lambda})$ is feasible for (16) if and only if $(\mathbf{x}, \mathbf{u}, \mathbf{V}, \mathbf{\Pi}, \mathbf{\Xi})$ is feasible for (13), where*

$$\begin{aligned} \mathbf{u} &= \boldsymbol{\epsilon} + \boldsymbol{\Omega}\mathbf{d}, \\ \mathbf{V} &= \boldsymbol{\Lambda} - \boldsymbol{\Omega}\mathbf{D}, \\ \mathbf{\Pi} &= \boldsymbol{\Omega}^\top \mathbf{B}^\top + \mathbf{Q}, \\ \mathbf{\Xi} &= \boldsymbol{\Omega}^\top. \end{aligned}$$

Table 2 Comparing the Number of Affine Constraints and Sign Restrictions in (13) and (16)

	Primal formulation (13)	Dual formulation (16)
No. of affine constraints	$(1+L)(m+k)$	$m(1+L+p)$
No. of sign restrictions	$p(m+k)$	$k(1+L+p)$

The proof is direct and therefore omitted. Theorem 2 is not only useful because it proves equivalence of the primal and dual formulation with affine policies. It also allows us to solve the dual formulation (16) with affine policies and directly obtain the optimal affine policy of the original formulation (or vice versa). Despite this equivalence there may be significant computational benefits from solving two stage problems using the dualized formulation rather than the primal formulation. This can be seen by comparing the two robust counterparts (13) and (16). We compare the number of affine constraints and the number of sign restrictions in Table 2. We use the same parameters as in Table 1 for the number of uncertain parameters (L), the number of wait-and-see decisions (k), the number of affine constraints on the variables (m) and the number of affine constraints in the uncertainty set (p). We observe that the total number of constraints (affine constraints and sign restrictions) is the same in both formulations. However, there is a difference in the break down into the number of affine constraints and the number of sign restrictions. This is important since sign restrictions are much easier to handle by solvers than affine constraints. From Table 2 we see that for a large number of wait-and-see decisions k , relative to the number of constraints in the original model and uncertainty set (m and p), the dual formulation (16) can most likely be solved more efficiently than the primal formulation (13). We observe these computational benefits in our numerical examples in Section 5 and 6 where we present Table 2 with some explicit values for L, k, m , and p .

Finally, we note that the models (12) and (15) can also be solved via cutting plane methods (Mutapcic and Boyd 2009). There have been extensive numerical studies that show that in some cases cutting plane algorithms require slightly less computation time than solving the robust counterpart constructed by Lemma 2 (Fischetti and Monaci 2012, Bertsimas et al. 2016). We have also solved our numerical examples with the cutting plane algorithm described in those papers. As with the reformulation approach, we observe that the dual formulation (15) can be solved an order of magnitude faster than the primal problem. This approach is, however, not elaborated further for two reasons. First, to construct the primal solution from the dual solution by Theorem 2, we need the auxiliary variables that are introduced by the reformulation.

Second, initial findings showed that the cutting plane algorithm is much slower for the problems considered Sections 5 and 6. We were only able to solve the smaller instances in reasonable time via cutting planes.

4. Stronger Bounds on the Optimality Gap of Affine Policies

In general, the restriction from fully adaptive policies to affine policies is an approximation of the fully adaptive solution, for both the primal and dual formulation. It is important to provide methods that can efficiently determine bounds on the (sub)optimality of affine policies. Here we extend a method that was first presented in Hadjiyiannis et al. (2011) to provide bounds on the optimality gap of affine policies. We first explain the initial idea from Hadjiyiannis et al. (2011) and then describe the algorithm that provides stronger bounds.

The main idea is to solve the fully adjustable model (2) only for a finite subset of the uncertainty set. Clearly, any optimal solution to this model results in a lower bound since we only guarantee feasibility for a strict subset of the uncertainty region. If we denote the finite subset by $\{\zeta^1, \zeta^2, \dots, \zeta^{\bar{N}}\}$, then we end up with the following equivalent deterministic optimization model

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{y}^1, \dots, \mathbf{y}^{\bar{N}}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{y}^i \geq \mathbf{R}\zeta^i + \mathbf{r}, \quad \forall i = 1, \dots, \bar{N}, \\ & \mathbf{x} \in \mathcal{X}, \mathbf{y}^1, \dots, \mathbf{y}^{\bar{N}} \geq 0. \end{aligned} \tag{17}$$

The crucial question is, of course, which scenarios to include. It is shown by Bertsimas and Goyal (2012) that the lower bound is tight if we include all extreme points of the uncertainty set. This is in practice undoable since there can be a huge number of extreme points, each resulting in extra variables and constraints in (17). Another straightforward way would be to sample \bar{N} scenarios uniformly at random from \mathcal{U} . The model (17) remains tractable for relatively large \bar{N} , but for all our examples we obtain useless bounds, even when the number of random samples \bar{N} is as big as 10^5 . We therefore have to pick the scenarios in a more specific way. To do so, we first introduce the notion of *binding scenarios*.

DEFINITION 1 (BINDING SCENARIOS). Let $f: \mathcal{U} \times \mathcal{X} \rightarrow \mathbb{R}$ be a function of the uncertain parameter $\zeta \in \mathcal{U}$ and here-and-now decision $\mathbf{x} \in \mathcal{X}$. For a given $\mathbf{x} \in \mathcal{X}$ the parameter $\hat{\zeta}$ is called *binding* for the robust constraint

$$f(\zeta, \mathbf{x}) \leq 0, \quad \forall \zeta \in \mathcal{U}$$

if $f(\hat{\zeta}, \mathbf{x}) = 0$.

In the primal formulation with affine policies we only have here-and-now decisions \mathbf{x} , \mathbf{u} and \mathbf{V} . Furthermore, each robust constraint is linear in the here-and-now decision and the uncertain parameter. Therefore, a binding scenario can easily be found for each constraint by solving a small linear optimization model $\hat{\zeta} = \arg \max_{\zeta \in \mathcal{U}} f(\zeta, \mathbf{x})$ and checking whether the maximum is equal to zero (up to a certain precision). The hope is that scenarios that are binding the solution with affine policies are also binding the fully adaptive solution.

The method by Hadjiyiannis et al. (2011) only uses the information derived from the primal formulation with affine policies (2). Using Theorem 2 we can directly construct the optimal affine policy in the dual formulation once the optimal affine policy in the primal formulation is known. Using this other affine policy we can construct another subset of \mathcal{U} consisting of binding scenarios in the dual formulation. The resulting deterministic model of the dual formulation with a finite subset $\{\mathbf{w}^1, \mathbf{w}^2, \dots, \mathbf{w}^{\bar{M}}\}$ is given by

$$\begin{aligned} \min_{\mathbf{x}, \lambda^1, \dots, \lambda^{\bar{M}}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & (\mathbf{w}^j)^\top (\mathbf{A}\mathbf{x} - \mathbf{r}) - \mathbf{d}^\top \lambda^j \geq 0, \quad \forall j = 1, \dots, \bar{M}, \\ & \mathbf{D}^\top \lambda^j \geq \mathbf{R}^\top \mathbf{w}^j, \quad \forall j = 1, \dots, \bar{M}, \\ & \mathbf{x} \in \mathcal{X}, \lambda^1, \dots, \lambda^{\bar{M}} \geq 0. \end{aligned} \tag{18}$$

Combining the constraints from (17) and (18) results in a model that provides a stronger lower bound than the one that only uses the binding scenarios from the primal formulation. We now provide Algorithm 1, which provides the strengthened bound on the optimal value of the fully adaptive model. Step 1 provides us with a feasible solution and an upper bound on the optimal value of the fully adaptive problem. The objective value of the model in step 4 gives us the new lower bound. A binding scenario for each constraints in (12) and (15) can be found directly using the optimal affine policies from steps 1 and 2. We omit here the elaborate description of a more efficient way to finding the set of binding scenarios in step 3 via KKT conditions, which is described in Hadjiyiannis et al. (2011). However, step 3 is not the most time consuming step, because solving the model with affine policies in step 1 takes by far the most time. Finally, we note that we can also solve the dual formulation (16) with affine policies in step 1 and obtain the primal affine policy in step 2 using Theorem 2.

Algorithm 1 (Stronger bounds on optimality of affine policies)

1. Solve (13) to get optimal here-and-now \mathbf{x} , affine policy $\mathbf{y}(\zeta) = \mathbf{u} + \mathbf{V}\zeta$ and auxiliary variables $\mathbf{\Pi}, \mathbf{\Xi}$.
2. Construct the dual affine policy $\lambda(\mathbf{w}) = \mathbf{Q}\mathbf{w}$ using Theorem 2.

3. Find the binding scenarios $\{\zeta^1, \zeta^2, \dots, \zeta^{\bar{N}}\}$ in (12) and $\{\mathbf{w}^1, \mathbf{w}^2, \dots, \mathbf{w}^{\bar{M}}\}$ in (15).
4. Solve the sampled problem with binding scenarios for the primal and dual

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{y}^1, \dots, \mathbf{y}^{\bar{N}}, \lambda^1, \dots, \lambda^{\bar{M}}} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} \quad & \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{y}^i \geq \mathbf{R}\zeta^i + \mathbf{r}, \quad \forall i = 1, \dots, \bar{N}, \\ & \mathbf{y}^1, \dots, \mathbf{y}^{\bar{N}} \geq 0, \\ & (\mathbf{w}^j)^\top (\mathbf{A}\mathbf{x} - \mathbf{r}) - \mathbf{d}^\top \lambda^j \geq 0, \\ & \hspace{15em} \forall j = 1, \dots, \bar{M}, \\ & \mathbf{D}^\top \lambda^j \geq \mathbf{R}^\top \mathbf{w}^j, \quad \forall j = 1, \dots, \bar{M}, \\ & \lambda^1, \dots, \lambda^{\bar{M}} \geq 0, \\ & \mathbf{x} \in \mathcal{X}. \end{aligned}$$

5. Example 1: Lot-Sizing on a Network

In this section we present a natural example in which (15) takes an order of magnitude less time to solve than the primal formulation (12). The new lower bound on the fully adaptive model (2) derived from Algorithm 1 is much stronger than the lower bound from Hadjiyiannis et al. (2011) that only used the binding scenarios from the primal formulation.

5.1. Problem Setting

In lot sizing on a network we have to determine the stock allocation x_i for $i = 1, \dots, N$ stores prior to knowing the realization of the demand at each location. The demand ζ is uncertain and assumed to be in a budget uncertainty set

$$\mathcal{U} = \{\zeta: \mathbf{0} \leq \zeta \leq \hat{\zeta}\mathbf{e}, \mathbf{e}^\top \zeta \leq \Gamma\}.$$

After we observe the realization of the demand we can transport stock y_{ij} from store i to store j at cost t_{ij} in order to meet all demand. The aim is to minimize the worst case storage costs (with unit costs c_i) and the cost arising from shifting the products from one store to another. This network flow model can now be written as a specific instance of the primal problem (2) as follows:

$$\begin{aligned} \min_{\mathbf{x}, \alpha} \quad & \alpha \\ \text{s.t.} \quad & \forall \zeta \in \mathcal{U}: \exists \mathbf{y} \geq \mathbf{0}: \begin{cases} \alpha \geq \sum_{i=1}^N c_i x_i + \sum_{i=1}^N \sum_{j=1}^N t_{ij} y_{ij}, \\ \zeta_i \leq \sum_{j=1}^N y_{ji} - \sum_{j=1}^N y_{ij} + x_i, \\ i = 1, \dots, N, \end{cases} \\ & 0 \leq x_i \leq K_i \quad i = 1, \dots, N, \end{aligned} \tag{19}$$

where the first line in (19) is for the epigraph formulation. The second line contains the balance equations:

we have to shift stock to and from node i such that the initial storage plus the net shift in stock still exceeds the demand at node i . The last constraints restrict the capacity of the stock at each node. Note that this model can be seen as a network flow model with multiple sources and multiple sinks.

5.2. Test Case and Numerical Results

We pick $N \in \{10, 20, 30, \dots, 100\}$ locations uniformly at random from $[0, 10]^2$. Let t_{ij} , the cost to transport one unit of demand from location i to j , be the Euclidean distance and the unit storage cost c_i be equal to 20. The individual maximum demand $\hat{\zeta}$ and the capacity K_i of each store is set to 20 units. The total demand in the network is set to $\Gamma = 20\sqrt{N}$. This is to avoid trivial and unrealistic cases where either all demand can occur at a single store ($\Gamma = 20$) or where the demand in each store is independent ($\Gamma = 20N$). All computations were carried out with Gurobi 6.0.3 (Gurobi Optimization 2015) on an Intel i7-4770 3.40 GHz Windows computer with 8 GB RAM. All modeling was done using the modeling language JuMP (Lubin and Dunning 2015).

We solve both (13) and (15) and depict the average solution times over 10 runs in Table 3, as well as the objective value and the lower bounds. The stock allocation (the here-and-now decision) for an instance with $N = 30$ is depicted in Figure 1. The lower bound from the primal is obtained using the method from Hadjiyiannis et al. (2011). The primal/dual bound is the strengthened bound resulting from Algorithm 1. Solving the model via the new dualized formulation (16) reduces the computation an order of magnitude compared with the original primal formulation (13). For the larger instances we see that the primal formulation is approximately 20 times slower. These results are averaged over 10 runs to avoid random peak performances, but in each individual run we observed the

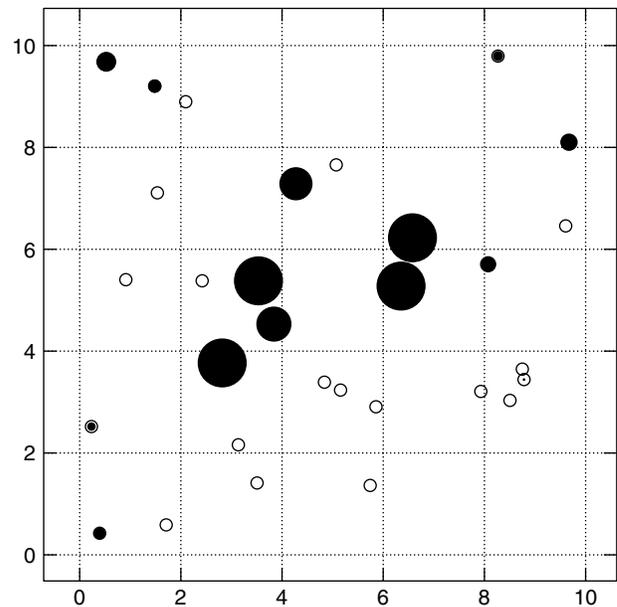


Figure 1 Stock Allocation for an Instance with 30 Stores on the Grid $[0, 10]^2$

Notes. The filled dots have stock, and the larger the dots, the more stock is allocated. The open dots are stores that do not have any stock allocated.

significant decrease in computation time. The strengthened primal/dual bound from Algorithm 1 is much tighter than the primal bound, more than halving the optimality gap for the larger instances.

5.3. Why Is the Dual Formulation Faster?

To understand the significantly faster computation time of the dual formulation displayed in Table 3, we look at the dimensions (number of uncertain parameters, wait-and-see decisions, constraints on variables, and constraints on uncertain parameters) for the case $N = 20$. We give the values of these dimensions in Table 4 using the same format as is in Table 1. We observe that the primal and dual formulation have the same characteristics, except for the number of wait-and-see decisions and the number of constraints on the uncertain parameter in the uncertainty set. Given these values, we can explicitly calculate the number of affine constraints and the number of sign restrictions using the formulas from Table 2. The resulting number of constraints and sign restrictions are given

Table 3 Compare Performance of Primal and Dualized Formulation with Affine Policies for the Lot-Sizing Example

N	Solver time (sec)		Objective value	Lower bound (gap%)	
	Primal	Dual		Primal	Primal/Dual
10	< 0.1	< 0.1	928	797 (14.0%)	824 (11.1%)
20	0.3	0.1	1,353	1,113 (17.7%)	1,190 (12.0%)
30	2.6	0.8	1,670	1,356 (18.8%)	1,465 (12.3%)
40	11.8	2.6	1,947	1,562 (19.8%)	1,728 (11.3%)
50	42.0	7.3	2,188	1,728 (21.0%)	1,934 (11.6%)
60	142.2	20.5	2,421	1,912 (21.0%)	2,160 (10.8%)
70	366.0	41.3	2,598	1,996 (23.2%)	2,312 (11.0%)
80	826.9	88.7	2,781	2,136 (23.2%)	2,495 (10.3%)
90	1,647.1	179.8	2,953	2,252 (23.8%)	2,641 (10.6%)
100	4,026.2	231.0	3,130	2,408 (23.1%)	2,799 (10.6%)

Notes. The percentages in the last columns depict the optimality gap derived from each lower bound compared to the objective value. All results are averaged over 10 runs.

Table 4 Comparing Dimensions of Variables, Uncertainties, and Number of Constraints in the Primal and Dual Formulation for the Lot-Sizing Instance with $N = 20$ Stores

	Primal formulation (2)	Dual formulation (3)
No. of uncertain parameters	20	21
No. of wait-and-see decisions	400	21
No. of constraints on variables	21	21
No. of constraints on uncertain parameter	21	401

in Table 5. We observe that the primal formulation (13) has about 50 times more affine constraints than the dual formulation (16). The dual formulation does have more sign restrictions on its variables, but these are significantly simpler for solvers. To investigate the claim that the number of affine constraints are indeed the cause of the reduction in computation time, we adapt the $N = 20$ instance from the network lot-sizing model (19). From Table 2 we see that increasing p , the number of affine constraints in the uncertainty set \mathcal{U} , leads to an increase of affine constraints in the dual formulation with affine policies. At the same time, the value of p does not affect the number of affine constraints in the primal formulation. To increase p , we add nonredundant constraints of the following type to the polyhedral description of \mathcal{U} :

$$\sum_{i \in S} \zeta_i \leq 20\sqrt{|S|}, \quad (20)$$

where $S \subset \{1, \dots, N\}$ is a random subset of size $\frac{1}{2}N$. The number of constraints p can be increased at will by adding more of these constraints. Note that increasing p also increases the total number of variables and the number of sign constraints, but these grow in more or less the same order of magnitude in both formulations. If we consider the case $N = 20$, then we find that the number of affine constraints in (13) and (16) is equal when the number of constraints in the uncertainty set \mathcal{U} equals $p = 400$. Note that $p \geq 21$, since we need 21 constraints to describe the budget uncertainty set. The case with $p = 21$ is therefore just our original network lot-sizing problem (19). We measure the difference in computation time between the primal and the dual formulation by the quotient

$$\frac{\text{Solver time for (13)}}{\text{Solver time for (16)}}.$$

In Figure 2, we plotted this quotient for each random instance with $p \in \{21, 22, \dots, 1,000\}$ constraints in the uncertainty set. We already know from Table 3 that the dual formulation with affine policies solves the original instance three or more times faster than the primal formulation. If we start adding constraints, the computational advantage progressively decreases and, after a certain point, disappears.

Table 5 Comparing the Number of Affine Constraints and Sign Restrictions in (13) and (16) for the Lot-Sizing Instance with $N = 20$ Stores

	Primal formulation (13)	Dual formulation (16)
No. of affine constraints	8,841	882
No. of sign restrictions	8,841	16,800

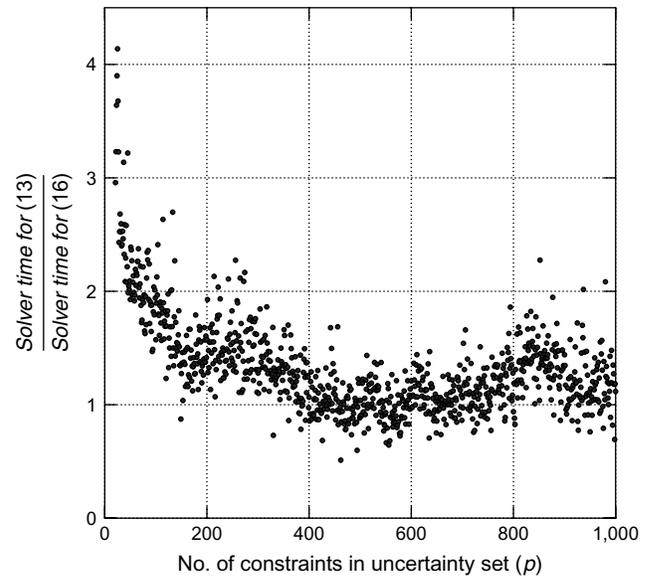


Figure 2 The Computation Time of the Primal Formulation (13) Divided by the Time Needed to Solve the Primal Formulation (13) for $N = 20$ and Various Number of Affine Constraints in the Uncertainty Set p

Notes. For $p = 400$ both (13) and (16) have the same number of affine constraints. Values above the horizontal line at 1 indicate that the dual formulation is solved faster than the primal formulation and vice versa for values smaller than 1.

6. Example 2: Facility Location Problem

The second example we consider is a facility location problem that has also been studied in Ardestani-Jaafari and Delage (2014) and Baron et al. (2011). Similar two-stage adaptive models can be found in Zeng and Zhao (2013). In our results we again observe a significant reduction in computational time required for solving the dualized formulation with affine policies over the primal formulation with affine policies. For this problem, however, the strengthened bounds from Algorithm 1 only slightly improve the bounds obtained from the primal formulation.

6.1. Problem Setting

We consider a facility location problem where we can build factories at candidate sites $s \in \mathcal{S} = \{1, \dots, S\}$, which have to serve customers $c \in \mathcal{C} = \{1, \dots, C\}$ in the area. The uncertain demand for customer c is modelled as $(1 - \zeta_c)\bar{d}_c$, with \bar{d}_c the nominal demand of customer c and ζ_c the uncertain shock in the demand. We take again a budget uncertainty set of the form

$$\mathcal{U} = \{\zeta: \mathbf{0} \leq \zeta \leq \hat{\zeta}\mathbf{e}, \mathbf{e}^\top \zeta \leq \Gamma\},$$

where Γ is our budget parameter.¹ There are two types of decisions in this model. First, there are strategic

¹ In fact, Ardestani-Jaafari and Delage (2014) also consider negative values of the uncertainty parameter. It is not hard to see that these are nonbinding scenarios and we can therefore use this uncertainty set instead.

here-and-now decisions that have to be decided before the demand is known. We have a binary variable x_s to decide whether the facility at site s is opened and a continuous variable p_s to set the capacity level at each opened facility site. Second, there are wait-and-see decisions y_{sc} on the production at facility s , which is transported to customer c . Each unit of demand can generate a revenue of η . There are also several costs incurred for the various strategic and operations decisions. Opening a facility s has a fixed cost f_s and a cost of b_s per unit of capacity installed. The production of one unit at facility s has cost g_s and transporting the goods to customer c bears an additional cost h_{sc} . The goal is to maximize the total profit. This problem can be modelled as a two-stage adaptive optimization model (Ardestani-Jaafari and Delage 2014):

$$\begin{aligned} \max_{t, \mathbf{x}, \mathbf{p}} \quad & \left\{ \alpha - \sum_{s \in \mathcal{S}} (b_s p_s + f_s x_s) \right\} \\ \text{s.t. } \forall \zeta \in \mathcal{U}: \exists \mathbf{y} \geq 0: \quad & \begin{cases} \sum_{s \in \mathcal{S}, c \in \mathcal{C}} (\eta - g_s - h_{sc}) y_{sc} \geq \alpha, \\ \sum_{c \in \mathcal{C}} y_{sc} \leq p_s, \quad \forall s \in \mathcal{S}, \\ \sum_{s \in \mathcal{S}} y_{sc} \leq \bar{d}_c - \zeta_c \bar{d}_c, \quad \forall c \in \mathcal{C}, \end{cases} \\ \mathbf{p} \leq M\mathbf{x}, \mathbf{x} \in \{0, 1\}^N. \end{aligned} \tag{21}$$

Note that we have a maximization objective, but this can easily be turned into a minimization objective by the relation $\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) = -\min_{\mathbf{x} \in \mathcal{X}} (-f(\mathbf{x}))$ before applying Theorem 1.

6.2. Test Case and Numerical Results

We consider the same setting as in Ardestani-Jaafari and Delage (2014), which is based on the setup of an earlier paper on robust facility location planning by Baron et al. (2011). We randomly generate C customers and S sites on a unit square. For the cost parameters we take $f_s = 50,000$, $b_s = 0.1$, $g_s = 0.1$, $\eta = 1$. The nominal demand is drawn uniformly at random from $[17500, 22500]$ and $\hat{\zeta} = 0.15$. The transportation cost t_{ij} is just the Euclidean distance between two points i and j . We take $S = 10$ possible sites and $C \in \{10, 20, 30, 40, 50\}$. The cases with $C = 10$ and $C = 20$ are in Ardestani-Jaafari and Delage (2014) referred to as small and medium instances. For the larger instances the computational time vastly increased and they did not report results on the models with affine policies. We use the same computer and optimization software as mentioned in Section 5.

The results for various numbers of customers C and various percentage levels of uncertainty Γ are given in Table 6. We use the standard notion of budget uncertainty where a budget of 30% means that 30%

Table 6 Numerical Results for Facility Location Problem with Affine Policies

C	Γ%	Solver time (sec)		Objective value	Upper bound (gap%)	
		Primal	Dual		Primal	Combined P/D
10	10	0.2	0.7	30,946	32,233 (3.3%)	32,167 (3.1%)
	30	0.8	1.2	27,894	30,474 (8.0%)	29,835 (6.1%)
	50	1.1	1.3	25,409	28,763 (10.5%)	27,897 (7.9%)
	70	2.0	1.5	23,416	24,895 (5.6%)	24,430 (3.6%)
	90	2.6	0.9	21,889	26,511 (18.3%)	26,353 (17.5%)
20	100	1.9	0.7	21,516	29,136 (28.4%)	26,803 (19.6%)
	10	7.4	3.6	85,895	87,264 (1.3%)	87,264 (1.3%)
	30	10.4	4.2	79,996	82,235 (2.3%)	81,883 (2.0%)
	50	18.0	5.2	75,404	77,060 (1.8%)	76,827 (1.6%)
	70	23.4	5.4	71,872	77,473 (6.4%)	76,854 (5.6%)
30	90	21.2	4.7	69,104	69,874 (0.9%)	69,712 (0.7%)
	100	11.8	1.1	68,226	80,301 (14.7%)	79,810 (14.1%)
	10	55.2	30.3	173,069	174,547 (0.7%)	174,004 (0.5%)
	30	112.5	35.4	163,953	168,422 (2.3%)	166,642 (1.4%)
	50	144.3	35.8	156,451	160,911 (2.3%)	157,913 (0.7%)
40	70	220.1	40.8	150,070	156,881 (3.6%)	153,511 (1.8%)
	90	251.2	31.9	144,873	150,741 (3.4%)	149,310 (2.6%)
	100	111.8	6.4	143,010	164,214 (12.4%)	159,182 (9.5%)
	10	307.4	114.5	243,639	244,628 (0.3%)	244,219 (0.2%)
	30	787.8	220.7	230,556	234,272 (1.3%)	233,557 (1.1%)
50	50	986.2	197.4	219,446	222,396 (1.1%)	221,665 (0.8%)
	70	1,735.4	199.0	209,942	212,479 (1.0%)	211,588 (0.7%)
	90	1,761.8	154.9	202,456	203,607 (0.5%)	203,011 (0.2%)
	100	877.7	25.7	200,044	223,373 (9.7%)	222,408 (9.3%)
	100	10	1,049.0	326.3	341,060	341,951 (0.2%)
30		2,153.2	530.4	323,989	327,184 (0.8%)	325,526 (0.4%)
50		2,766.5	557.1	308,882	312,840 (1.1%)	311,457 (0.7%)
70		4,542.5	536.8	295,599	298,961 (1.0%)	298,129 (0.7%)
90		5,830.9	469.6	284,574	292,716 (2.3%)	291,174 (1.8%)
100	3,582.1	68.2	280,704	304,575 (7.1%)	302,579 (6.5%)	

Notes. The percentages in the last columns depict the optimality gap derived from each upper bound compared to the objective value. All results are averaged over 5 runs.

of the uncertain parameters can be at their extreme value of $\hat{\zeta} = 0.15$. A graph indicating the location and the facilities that are opened for one case is given in Figure 3.

The most striking result is that the dual formulation with affine policies is again solved an order of magnitude faster than the primal formulation with affine policies. This holds especially true for the larger instances and larger values of Γ . We again look at the dimensions of the primal and the dual formulation using Table 1 for its dimensions and Table 2 for the different constraints. For the the case with $C = 50$ customers we present these results in Tables 7 and 8.

Again we see a smaller number of difficult affine constraints in the dual version in exchange for a larger number of easy-to-handle sign restrictions.

If we take a look at the bounds we see they are very close to the objective value, which shows that the use of affine policies is nearly optimal. This observation was also made for the smaller instances in Ardestani-Jaafari and Delage (2014). For $\Gamma = 100\%$, the lower bound is

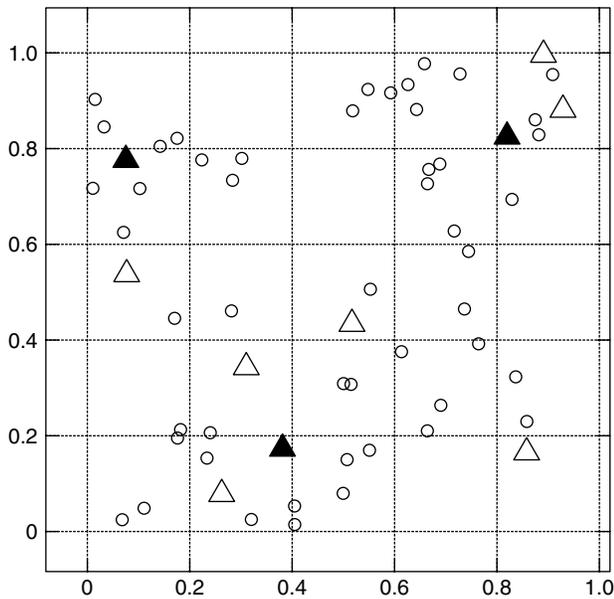


Figure 3 Solution for One Facility Location Instance with $S = 10$ Possible Sites and $C = 50$ Customers on $[0, 1]^2$

Notes. The uncertainty level is set at $\Gamma = 50\%$. Facility locations are indicated by triangles, customers by open circles. The filled triangles are the locations that are picked to be open.

the most far away from the objective value. This is surprising, because for this case (box uncertainty) we know that affine (in fact, static) policies are provably optimal (Ben-Tal et al. 2009, Theorem 14.2.4).

7. Concluding Remarks

In this paper, we have used duality for second-stage decisions and uncertain parameters to derive an equivalent formulation of a primal two-stage adaptive model. The resulting dualized formulation is again a two-stage

Table 7 Comparing Dimensions of Variables, Uncertainties, and Number of Constraints in the Primal and Dual Formulation for the Facility Location Problem (21) with $C = 50$ Customers

	Primal formulation (2)	Dual formulation (3)
No. of uncertain parameters	50	61
No. of wait-and-see decisions	500	51
No. of constraints on variables	61	51
No. of constraints on uncertain parameter	51	501

Table 8 Comparing the Number of Affine Constraints and Sign Restrictions in (13) and (16) for the Facility Location Problem (21) with $C = 50$ Customers

	Primal formulation (13)	Dual formulation (16)
No. of affine constraints	28,661	6,222
No. of sign restrictions	28,661	51,000

adaptive model. We show that optimal affine policies for the primal formulation can be directly constructed from optimal affine policies in the dual formulation. Two examples of lot-sizing and a facility location problem show that the dualized models, when coupled with affine policies, can reduce computational time to solve adaptive problems by an order of magnitude. Furthermore, we provide an algorithm that uses the affine policies in the dual model to strengthen bounds on the optimality gap of affine policies.

Acknowledgments

This research was conducted while the second author was visiting the Operations Research Center at the Massachusetts Institute of Technology and was partly financed by the Netherlands Organisation for Scientific Research (NWO) Research Talent [Grant 406-14-067].

References

- Ardestani-Jaafari A, Delage E (2014) The value of flexibility in robust location-transportation problems *Les Cahiers du GERAD G-2014-83*, GERAD, HEC Montréal.
- Atamtürk A, Zhang M (2007) Two-stage robust network flow and design under demand uncertainty. *Oper. Res.* 55:662–673.
- Baron O, Milner J, Naseraldin H (2011) Facility location: A robust optimization approach. *Production Oper. Management* 20:772–785.
- Ben-Tal A, El Ghaoui L, Nemirovski A (2009) *Robust Optimization*, Princeton Series in Applied Mathematics (Princeton University Press, Princeton, NJ).
- Ben-Tal A, Do Chung B, Mandala SR, Yao T (2011) Robust optimization for emergency logistics planning: Risk mitigation in humanitarian relief supply chains. *Transportation Res. Part B: Methodological* 45:1177–1189.
- Ben-Tal A, Golany B, Nemirovski A, Vial J-R (2005) Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing Service Oper. Management* 7:248–271.
- Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Math. Programming* 99:351–376.
- Bertsimas D, Caramanis C (2010) Finite adaptability in multistage linear optimization. *Automatic Control, IEEE Trans.* 55:2751–2766.
- Bertsimas D, Dunning I (2016) Multistage robust mixed integer optimization with adaptive partitions. *Oper. Res.* Forthcoming.
- Bertsimas D, Goyal V (2010) On the power of robust solutions in two-stage stochastic and adaptive optimization problems. *Math. Oper. Res.* 35:284–305.
- Bertsimas D, Goyal V (2012) On the power and limitations of affine policies in two-stage adaptive optimization. *Math. Programming* 134:491–531.
- Bertsimas D, Sim M (2004) The price of robustness. *Oper. Res.* 52:35–53.
- Bertsimas D, Brown DB, Caramanis C (2011) Theory and applications of robust optimization. *SIAM Rev.* 53:464–501.
- Bertsimas D, Dunning I, Lubin M (2016) Reformulation versus cutting-planes for robust optimization. *Comput. Management Sci.* 13:195–217.
- Bertsimas D, Iancu DA, Parrilo PA (2010) Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.* 35:363–394.
- Bertsimas D, Litvinov E, Sun XA, Zhao J, Zheng T (2013) Adaptive robust optimization for the security constrained unit commitment problem. *Power Systems, IEEE Trans.* 28:52–63.
- Birge JR, Louveaux F (2011) *Introduction to Stochastic Programming* (Springer Science and Business Media, New York).
- Fischetti M, Monaci M (2012) Cutting plane versus compact formulations for uncertain (integer) linear programs. *Math. Programming Comput.* 4:239–273.
- Gabrel V, Murat C, Thiele A (2014b) Recent advances in robust optimization: An overview. *Eur. J. Oper. Res.* 235:471–483.

- Gabrel V, Lacroix M, Murat C, Remli N (2014a) Robust location transportation problems under uncertain demands. *Discrete Appl. Math.* 164:100–111.
- Gurobi Optimization Inc. (2015) Gurobi optimizer reference manual. Accessed July 1, 2015, <http://www.gurobi.com>.
- Hadjiyiannis M, Goulart P, Kuhn D (2011) A scenario approach for estimating the suboptimality of linear decision rules in two-stage robust optimization. *50th IEEE Conf. Decision Control Euro. Control, Orlando, FL*, 7386–7391.
- Iancu DA, Sharma M, Sviridenko M (2013) Supermodularity and affine policies in dynamic robust optimization. *Oper. Res.* 61: 941–956.
- Kali P, Wallace SW (1994) *Stochastic Programming* (Springer, New York).
- Kuhn D, Wiesemann W, Georghiou A (2011) Primal and dual linear decision rules in stochastic and robust optimization. *Math. Programming* 130:177–209.
- Lubin M, Dunning I (2015) Computing in operations research using Julia. *INFORMS J. Comput.* 27:238–248.
- Minoux M (2011) On 2-stage robust LP with RHS uncertainty: Complexity results and applications. *J. Global Optim.* 49: 521–537.
- Mutapcic A, Boyd S (2009) Cutting-set methods for robust convex optimization with pessimizing oracles. *Optim. Methods Software* 24:381–406.
- Ordóñez F, Zhao J (2007) Robust capacity expansion of network flows. *Networks* 50:136–145.
- Postek K, Den Hertog D (2016) Multi-stage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set. *INFORMS J. Comput.* Forthcoming.
- Shapiro A, Nemirovski A (2005) On complexity of stochastic programming problems. Jeyakumar V, Rubinov A, eds. *Continuous Optimization* (Springer, New York), 111–146.
- Thiele A, Terry T, Epelman M (2009) Robust linear optimization with recourse. Technical report TR09-01, Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI.
- Wang Q, Watson J-P, Guan Y (2013) Two-stage robust optimization for NK contingency-constrained unit commitment. *Power Systems, IEEE Trans.* 28:2366–2375.
- Yin Y, Madanat SM, Lu X-Y (2009) Robust improvement schemes for road networks under demand uncertainty. *Eur. J. Oper. Res.* 198:470–479.
- Zeng B, Zhao L (2013) Solving two-stage robust optimization problems using a column-and-constraint generation method. *Oper. Res. Lett.* 41:457–461.
- Zhao L, Zeng B (2012) Robust unit commitment problem with demand response and wind energy. *Power Energy Society General Meeting 2012 IEEE, San Diego*, 1–8.