Operations Research
Vol. 61, No. 4, July–August 2013, pp. 941–956
ISSN 0030-364X (print) | ISSN 1526-5463 (online)
http://dx.doi.org/10.1287/opre.2013.1172

Supermodularity and Affine Policies in Dynamic Robust Optimization

Dan A. Iancu
Graduate School of Business, Stanford University, Stanford, California 94305, daniancu@stanford.edu

Mayank Sharma
IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, mxsharma@us.ibm.com

Maxim Sviridenko
Computer Science Department, University of Warwick, Coventry, CV4 7AL, United Kingdom, m.i.sviridenko@warwick.ac.uk

This paper considers a particular class of dynamic robust optimization problems, where a large number of decisions must be made in the first stage, which consequently fix the constraints and cost structure underlying a one-dimensional, linear dynamical system. We seek to bridge two classical paradigms for solving such problems, namely, (1) dynamic programming (DP), and (2) policies parameterized in model uncertainties (also known as decision rules), obtained by solving tractable convex optimization problems.

We show that if the uncertainty sets are integer sublattices of the unit hypercube, the DP value functions are convex and supermodular in the uncertain parameters, and a certain technical condition is satisfied, then decision rules that are *affine* in the uncertain parameters are optimal. We also derive conditions under which such rules can be obtained by optimizing simple (i.e., linear) objective functions over the uncertainty sets. Our results suggest new modeling paradigms for dynamic robust optimization, and our proofs, which bring together ideas from three areas of optimization typically studied separately—robust optimization, combinatorial optimization (the theory of lattice programming and supermodularity), and global optimization (the theory of concave envelopes)—may be of independent interest.

We exemplify our findings in a class of applications concerning the design of flexible production processes, where a retailer seeks to optimally compute a set of strategic decisions (before the start of a selling season), as well as in-season replenishment policies. We show that, when the costs incurred are jointly convex, replenishment policies that depend linearly on the realized demands are optimal. When the costs are also piecewise affine, all the optimal decisions can be found by solving a single linear program of small size (when all decisions are continuous) or a mixed-integer, linear program of the same size (when some strategic decisions are discrete).

Subject classifications: dynamic robust optimization; supermodularity; concave envelopes; lattices; Lovász extension; production planning; inventory management.
Area of review: Optimization.
History: Received March 2012; revisions received September 2012, January 2013; accepted February 2013. Published online in *Articles in Advance* July 19, 2013.

1. Introduction

Dynamic optimization problems under uncertainty have been present in numerous fields of science and engineering, and have elicited interest from diverse research communities, on both a theoretical and a practical level. As a result, many solution approaches have been proposed, with various degrees of generality, tractability, and performance guarantees. One such methodology, which has received renewed interest in recent years because of its ability to provide workable solutions for many real-world problems, is robust optimization and robust control.

The topics of robust optimization and robust control have been studied, under different names, by a variety of academic groups, in operations research (Ben-Tal and Nemirovski 1999, 2002; Ben-Tal et al. 2002; Bertsimas and Sim 2003, 2004), engineering (Bertsekas and Rhodes 1971, Fan et al. 1991, El-Ghaoui et al. 1998, Zhou and Doyle 1998, Dullerud and Paganini 2005), and economics (Hansen and Sargent 2001, 2008), with considerable effort put into justifying the assumptions and general modeling philosophy. As such, the goal of the current paper is not to *motivate* the use of robust (and, more generally, distribution-free) techniques. Rather, we take the modeling approach as a given, and investigate questions of tractability and performance guarantees in the context of a specific class of dynamic optimization problems.

More precisely, we are concerned with models in which a potentially large set of constrained and costly decisions $\mathbf{K}$ must be taken in the first stage, which then critically influence the constraints and cost structure of a linear and one-dimensional system evolving in discrete time, over a finite horizon. Apart from the first-stage decisions $\mathbf{K}$; the system's evolution is also governed by particular actions (or controls) $\mathbf{u}_t$ taken by the decision maker at every time t, and also subject to unknown disturbances $\mathbf{w}_t$. In keeping with the traditional (min–max) robust paradigm, we make the modeling assumption that the uncertain quantities $\mathbf{w}_t$ are only known to lie in a specific uncertainty set $\mathcal{W}_t$. The goal of the decision maker is to compute the first-stage decisions $\mathbf{K}$ and a set of nonanticipative policies $\mathbf{u}_t$ so that the system obeys a set of prespecified constraints robustly (i.e., for any possible realization of the uncertain parameters), while minimizing a worst-case performance measure (see, e.g., Löfberg 2003; Bemporad et al. 2003; Kerrigan and Maciejowski 2003; Ben-Tal et al. 2004, 2005a, and references therein).

Several problems in operations research result in models that fit this description. One such instance, which we use throughout the paper to motivate and exemplify our results, is the following supply chain contracting model, considered in a similar form by Ben-Tal et al. (2005b, 2009).

Problem 1. Consider a retailer selling a single product over a finite planning horizon and facing unknown demands from customers. She is allowed to carry inventory and to backlog unsatisfied demand, and she can renew her inventory in every period by placing replenishment orders.

The retailer faces two types of decisions. Before the start of the selling season, a set of strategic decisions must be made, which fix the structure of the ordering, holding, and backlogging costs, as well as any constraints on order quantities and inventories faced by the retailer during the season.

The goal is to determine, in a centralized fashion, the strategic (i.e., preseason) decisions and the ordering policies that would minimize the overall, worst-case costs for the retailer.

In the model of Ben-Tal et al. (2005b, 2009), the retailer enters a contract with a supplier, whereby the former precommits to a set of orders before the start of the season, which can differ from the actual replenishments during the season. To smoothen the production at the supplier, the contract stipulates penalties for differences between successive precommitments, as well as for deviations of actual orders from precommitments. Here, the first-stage decisions $\mathbf{K}$ are the precommitments, which determine the contractual penalties paid by the retailer.

We note that two key features making such models salient are the *nonlinear* dependency of the cost structure on the strategic decisions, and the potential that at least some strategic decisions may be discrete (e.g., whether to acquire a particular technology, contract with a given vendor, hire more staff, etc.).

The typical approach for solving such problems is via a dynamic programming (DP) formulation (Bertsekas 2001), in which, with a compact notion of the system state $\mathbf{x}_t$, the optimal state-dependent policies $\mathbf{u}_t^*(\mathbf{x}_t)$ and value functions $J_t^*(\mathbf{x}_t)$ are characterized going backward in time. DP is a powerful and flexible technique, enabling the modeling of complex problems, with nonlinear dynamics, incomplete information structures, etc. For certain "simpler" (low-dimensional) problems, the DP approach also allows an exact characterization of the optimal actions; this has lead to numerous celebrated results in operations research, a classic example being the optimality of base stock or (s, S) policies in inventory systems (Scarf 1959, Clark and Scarf 1960, Veinott 1966). Furthermore, the DP approach often entails very useful comparative statics analyses, such as monotonicity results of the optimal policy with respect to particular problem parameters or state variables, monotonicity or convexity of the value functions, etc. (see, e.g., the classical texts Zipkin 2000, Topkis 1998, Heyman and Sobel 1984, Simchi-Levi et al. 2004, and Talluri and van Ryzin 2005 for numerous such examples). We critically remark that such comparative statics results are often possible even for complex problems, where the optimal policies cannot be completely characterized (e.g., Zipkin 2008, Huh and Janakiraman 2010).

The main downside of the DP approach is the well-known "curse of dimensionality," in that the complexity of the underlying Bellman recursions explodes with the number of state variables (Bertsekas 2001), leading to a limited applicability of the methodology in practical settings. In fact, an example of this phenomenon already appears in the model for Problem 1: after the (first-stage) strategic decisions are fixed, the state of the problem consists of the on-hand inventory available at the retailer. As Ben-Tal et al. (2005b, 2009) remark, even though the DP optimal ordering policy might have a simple form (e.g., if the ordering costs were linear, and the holding/backlogging costs were convex, it would be a base-stock policy), the methodology would encounter difficulties, as (i) one may have to discretize the state variable and the actions, and hence produce only an approximate value function; (ii) the dynamic program would have to be solved for any possible choice of strategic decisions; (iii) the value function depending on strategic decisions would, in general, be nonsmooth; and (iv) the DP solution would provide no subdifferential information, leading to the use of zero-order (i.e., gradient-free) methods to solve the resulting first-stage problem, which exhibit notoriously slow convergence. The latter issues would be furthermore exacerbated if some of the strategic decisions were discrete.

An alternative approach is to forgo solving the Bellman recursions (even approximately), and instead focus on particular classes of policies that can be optimized over by solving tractable optimization problems. One of the most popular such approaches is to consider *decision rules* directly parameterized in the observed disturbances, i.e.,

$$\mathbf{u}_t\colon \mathcal{W}_1 \times \mathcal{W}_2 \times \cdots \times \mathcal{W}_{t-1} \to \mathbb{R}^m, \tag{1}$$

where m is the number of control actions at time t. One such example of particular interest has been the class of *affine* decision rules. Originally suggested in the stochastic programming literature (Charnes et al. 1958, Garstka and Wets 1974), these rules have gained tremendous popularity in the robust optimization literature because of their tractability and empirical success (see, e.g., Löfberg 2003; Ben-Tal et al. 2004, 2005a, 2006, 2009; Bemporad et al. 2003; Kerrigan and Maciejowski 2003, 2004; Skaf and Boyd 2010; and Bertsimas et al. 2011a for more references). Recently, they have been reexamined in stochastic settings, with several papers (Shapiro and Nemirovski 2005, Chen et al. 2008, Kuhn et al. 2009) providing tractable methods for determining optimal policy parameters, in the context of both single-stage and multistage linear stochastic programming problems. Several extensions, such as piecewise affine (See and Sim 2010, Goh and Sim 2010) or polynomial decision rules (Ben-Tal et al. 2009, Bertsimas et al. 2011b) have also been recently discussed in the literature.

One central question when restricting attention to a particular subclass of policies (such as affine) is whether this induces large optimality gaps as compared to the DP solution. One such attempt was Bertsimas and Goyal (2010), which considers a two-stage linear optimization problem and shows that affine policies are optimal for a simplex uncertainty set, but can be within a factor of $\mathcal{O}(\sqrt{\dim(\mathcal{W})})$ of the DP optimal objective in general, where $\dim(\mathcal{W})$ is the dimension of the first-stage uncertainty set. Other research efforts have focused on providing tractable dual formulations, which allow a computation of lower or upper bounds, and hence a numerical assessment of the suboptimality level (see Kuhn et al. 2009 for details).

The work that is perhaps closest to ours is Bertsimas et al. (2010), where the authors show that affine decision rules are provably optimal for a considerably simpler setting than Problem 1, namely, one without first-stage (strategic) decisions, with *linear* ordering costs, and with the uncertainty set described by a hypercube. The proofs in the latter paper rely heavily on the problem structure, and cannot be extended to other settings, most importantly to models where the ordering costs depend nonlinearly on the decisions, such as in Problem 1.

However, these (seemingly weak) theoretical results stand in contrast with the considerably stronger empirical observations. In a thorough simulation conducted for an application very similar to Problem 1 Ben-Tal et al. (2009, chap. 14, p. 392) report that affine policies are *optimal* in all 768 instances tested, and Kuhn et al. (2009) find similar results for a related example.

In view of this observation, the goal of the present paper is to enhance the understanding of the modeling assumptions and problem structures that underlie the optimality of affine policies. We seek to do this, in fact, by bridging the strengths of the two approaches suggested above (DP and affine decision rules). Our contributions are as follows.

• We show that if the uncertainty sets are integer sublattices of the unit hypercube, the DP value functions are convex and supermodular in the uncertain parameters, and a certain technical condition is satisfied, then decision rules that are *affine* in the uncertain parameters are optimal. The reason why such conditions are useful is that one can often conduct meaningful comparative statics analyses, even in situations when a DP formulation is computationally challenging. If the optimal value functions and policies happen to match our conditions, then one can forgo numerically solving the DP, and can instead simply focus attention on affine decision rules, which can often be computed by solving particular tractable (convex) mathematical programs.

Our conditions critically rely on the convexity and supermodularity of the objective functions in question, as well as the lattice structure of the uncertainty set $\mathcal{W}$. To the best of our knowledge, these are the first results suggesting that lattice uncertainty sets might play a central role in constructing dynamic robust models, and that they bear a close connection with the optimality of affine forms in the resulting problems. Our proof techniques combine ideas from three areas of optimization typically studied separately—robust optimization, combinatorial optimization (the theory of lattice programming and supermodularity), and global optimization (the theory of concave envelopes)—and may be of independent interest.

• Using these conditions, we reexamine Problem 1, and show that—once the strategic decisions are fixed—affine ordering policies are provably optimal, under *any convex* ordering and inventory costs. Furthermore, the worst-case optimal ordering policy has a natural interpretation in terms of fractional satisfaction of backlogged demands. This generalizes and simplifies the results in Bertsimas et al. (2010), and it enforces the notion that optimal decision rules in robust models can retain a simple form, even as the cost structure of the problem becomes more complex: for instance, when ordering costs are convex, replenishment policies that are affine in historical demands remain optimal, whereas policies parameterized in inventory become considerably more complex (see the discussion in §3.3.1).

• Recognizing that, even knowing that affine policies *are* optimal, one could still face the conundrum of solving complex mathematical programs, we provide a set of conditions under which the maximization of a sum of several convex and supermodular functions on a lattice can be replaced with the maximization of a single, *linear* function. With these conditions, we show that, if all the costs in Problem 1 are jointly convex and piecewise affine (with at most m pieces), then *all* the decisions in the problem (strategic and optimal ordering policies) can be obtained by solving *a single* linear program (LP), with $\mathcal{O}(mT^2)$ variables. This explains the empirical results in Ben-Tal et al. (2005b, 2009) and identifies the sole modeling component that renders affine decision rules suboptimal in the latter models. Additionally, if some strategic decisions are discrete, then this LP becomes a mixed-integer linear program (MILP), with the same size.

The rest of the paper is organized as follows. Section 2 contains a precise mathematical description of the two main problems we seek to address. Sections 3 and 4 contain our main results, with the answers to each separate question, and a detailed discussion of their application to Problem 1 stated in the Introduction. Section 5 concludes the paper. The online appendix (available as supplemental material at http://dx.doi.org/10.1287/opre.2013.1172) contains relevant background material on lattice programming and supermodularity (§EC.1), concave envelopes (§EC.2), as well as some of the technical proofs (§EC.3).

1.1. Notation

We use $\bar{\mathbb{R}} \stackrel{\text{def}}{=} \mathbb{R} \cup \{\infty\}$ to denote the set of extended reals. Throughout the text, vector quantities are denoted in bold font. To avoid extra symbols, we use concatenation of vectors in a liberal fashion, i.e., for $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{b} \in \mathbb{R}^k$, we use $(\mathbf{a}, \mathbf{b})$ to denote either the row vector $(a_1, \ldots, a_n, b_1, \ldots, b_k)$ or the column vector $(a_1, \ldots, a_n, b_1, \ldots, b_k)^T$. The meaning should be clear from context. The operators min, max, $\geqslant$ and $\leqslant$ applied to vectors should be interpreted in component-wise fashion.

For a vector $\mathbf{x} \in \mathbb{R}^n$ and a set $S \subseteq \{1, \ldots, n\}$, we use $\mathbf{x}(S) \stackrel{\text{def}}{=} \sum_{j \in S} x_j$, and denote by $\mathbf{x}_S \in \mathbb{R}^n$ the vector with components x_i for $i \in S$ and 0 otherwise. In particular, $\mathbf{1}_S$ is the characteristic vector of the set S, $\mathbf{1}_i$ is the ith unit vector of $\mathbb{R}^n$, and $\mathbf{1} \in \mathbb{R}^n$ is the vector with all components equal to one. We use $\Pi(S)$ to denote the set of all permutations on the elements of S, and $\pi(S)$ or $\sigma(S)$ denote particular such permutations. We let $S^C = \{1, \ldots, n\} \backslash S$ denote the complement of S, and, for any permutation $\pi \in \Pi(S)$, we write $\pi(i)$ for the element of S appearing in the ith position under permutation π, and $\pi^{-1}(i)$ to denote the position of element $i \in S$ under permutation π.

For a set $P \subseteq \mathbb{R}^n$, we use $\text{ext}(P)$ to denote the set of its extreme points, and $\text{conv}(P)$ to denote its convex hull.

2. Problem Statement

As discussed in the Introduction, both the DP formulation and the decision rule approach have well-documented merits. The former is general purpose, and allows very insightful comparative statics analyses, even when the DP approach itself is computationally intractable. For instance, one can check the monotonicity of the optimal policy or value function with respect to particular problem parameters or state variables, or prove the convexity or submodularity/supermodularity of the value function. Such recent examples in the inventory literature are the monotonicity results concerning the optimal ordering policies in single or multiechelon supply chains with positive lead time and lost sales (Zipkin 2008 and Huh and Janakiraman 2010). For more examples, we refer the interested reader to several classical texts on inventory and revenue management: Zipkin (2000), Topkis (1998), Heyman and Sobel (1984), Simchi-Levi et al. (2004), and Talluri and van Ryzin (2005).

In contrast, the decision rule approach does not typically allow such structural results, but instead takes the pragmatic view of focusing on practical decisions, which can be efficiently computed by convex optimization techniques (see, e.g., Ben-Tal et al. 2009, chapter 14).

The goal of the present paper is to provide a link between the two analyses, and to enhance the understanding of the modeling assumptions and problem structures that underlie the optimality of affine decision rules. More precisely, we pose and address two main problems, the first of which is the following.

Problem 2. Consider a one-period game between a decision maker and nature

$$\max_{\mathbf{w} \in \mathscr{W}} \min_{\mathbf{u}(\mathbf{w})} f(\mathbf{w}, \mathbf{u}), \tag{2}$$

where $\mathbf{w}$ denotes an action chosen by nature from an uncertainty set $\mathscr{W} \subseteq \mathbb{R}^n$, $\mathbf{u}$ is a response by the decision maker, allowed to depend on nature's action $\mathbf{w}$, and f is a total cost function. With $\mathbf{u}^*(\mathbf{w})$ denoting the Bellman-optimal policy, we seek conditions on the set $\mathscr{W}$, the policy $\mathbf{u}^*(\mathbf{w})$, and the function $f(\mathbf{w}, \mathbf{u}^*(\mathbf{w}))$ such that there exists an affine policy that is worst-case optimal for the decision maker, i.e.,

$$\exists\, Q \in \mathbb{R}^{m \times n}, \;\; \mathbf{q} \in \mathbb{R}^m \quad \text{such that}$$

$$\max_{\mathbf{w} \in \mathscr{W}} \min_{\mathbf{u}(\mathbf{w})} f(\mathbf{w}, \mathbf{u}) = \max_{\mathbf{w} \in \mathscr{W}} f(\mathbf{w}, Q\mathbf{w} + \mathbf{q}).$$

To understand the question, imagine separating the objective into two components, $f(\mathbf{w}, \mathbf{u}) = h(\mathbf{w}) + J(\mathbf{w}, \mathbf{u})$. Here, h summarizes a sequence of historical costs (all depending on the unknowns $\mathbf{w}$), while J denotes a cost to go (or value function). As such, the outer maximization in (2) can be interpreted as the problem solved by nature at a particular stage in the decision process, whereby the total costs (historical + cost to go) are being maximized. The inner minimization exactly captures the decision maker's problem, of minimizing the cost to go.

We remark that the notion of *worst-case optimal* policies in the previous question is different than that of *Bellman-optimal* policies (Bertsekas 2001). In the spirit of DP, the latter requirement would translate in the policy $\mathbf{u}^*(\mathbf{w})$ being the optimal response by the decision maker for *any* revealed $\mathbf{w} \in \mathscr{W}$, whereas the former notion only requires that $\mathbf{u}(\mathbf{w}) = Q\mathbf{w} + \mathbf{q}$ is an optimal response at points $\mathbf{w}$ that result in the overall worst-case cost (while keeping the cost for all other $\mathbf{w}$ below the worst-case cost). This distinction has been drawn before (Bertsimas et al. 2010), and is one of the key features distinguishing robust (min–max) models from their stochastic counterparts, and allowing the former models to potentially admit optimal policies with simpler structure than those for the latter class. Although one could build a case against worst-case optimal policies by arguing that a rational decision maker should never accept policies that are not Bellman optimal (see, e.g., Epstein and Schneider 2003, Cheridito et al. 2006 for pointers to the literature in economics and risk theory on this topic), we adopt the pragmatic view here that, provided there is no degeneracy in the optimal policies (i.e., there is a unique set of optimal policies in the problem), one can always replicate the true Bellman-optimal policies for a finite-horizon problem through a shrinking horizon approach (Bertsekas 2001), by applying the first-stage decisions and resolving the subproblems of the decision process.

As remarked earlier, an answer to this question would be most useful in conjunction with comparative statics results obtained from a DP formulation: if the optimal value (and policies) matched the conditions in the answer to Problem 2, then one could forgo numerically solving the DP, and could instead simply focus attention on disturbance-affine policies, which could be computable by efficient convex optimization techniques (see, e.g., Löfberg 2003; Ben-Tal et al. 2004, 2005a, 2009; or Skaf and Boyd 2010).

Although answering the above question is certainly very relevant, the results might still remain existential in nature. In other words, even armed with the knowledge that affine policies *are* optimal, one could be faced with the conundrum of solving complex mathematical programs to find such policies. To partially alleviate this issue, we raise the following related problem.

Problem 3. Consider a maximization problem of the form

$$\max_{\mathbf{w}\in\mathcal{W}} \sum_{t\in\mathcal{T}} h_t(\mathbf{w}),$$

where $\mathcal{W} \subseteq \mathbb{R}^n$ denotes an uncertainty set, and $\mathcal{T}$ is a finite index set. Let J^* denote the maximum value in the problem above (assumed finite). We seek conditions on $\mathcal{W}$ and/or h_t such that there exist affine functions $\mathbf{z}_t(\mathbf{w})$, $\forall\, t \in \mathcal{T}$, such that

$$z_t(\mathbf{w}) \geqslant h_t(\mathbf{w}), \quad \forall\, \mathbf{w} \in \mathcal{W}, \forall\, t \in \mathcal{T}$$

$$J^* = \max_{\mathbf{w}\in\mathcal{W}} \sum_{t\in\mathcal{T}} z_t(\mathbf{w}).$$

In words, the latter problem requires that substituting a set of true historical costs h_t with potentially larger (but *affine*) costs z_t results in no change of the worst-case cost. Since one can typically optimize linear functionals efficiently over most uncertainty sets of practical interest (see, e.g., Ben-Tal et al. 2009), an answer to this problem, combined with an answer to Problem 2, might yield tractable and compact mathematical programs for computing worst-case optimal affine policies that depend on disturbances.

We note that conditions involving a linearization of the objectives have been discussed in the recent paper of Gorissen and den Hertog (2013), where the authors show that if the functions h_t are all piecewise affine and convex, then $\sum_{t\in\mathcal{T}} z_t$ exactly corresponds to the Fenchel dual of the function $\sum_{t\in\mathcal{T}} h_t$, which is generally a *strict* upper bound of the latter. By contrast, Problem 3 seeks conditions under which this upper bound yields the same value (when maximized) as the original function.

3. Discussion of Problem 2

We begin by considering Problem 2 in the Introduction. With $\mathbf{u}^*(\mathbf{w}) \in \arg\min_{\mathbf{u}} f(\mathbf{w}, \mathbf{u})$ denoting a Bellman-optimal response by the decision maker, the latter problem can be summarized compactly as finding conditions on $\mathcal{W}$, $\mathbf{u}^*$, and f such that

$$\exists\, Q \in \mathbb{R}^{m\times n},\ \mathbf{q} \in \mathbb{R}^m\colon\ \max_{\mathbf{w}\in\mathcal{W}} f(\mathbf{w}, u^*(\mathbf{w})) = \max_{\mathbf{w}\in\mathcal{W}} f(\mathbf{w}, Q\mathbf{w} + \mathbf{q}).$$

To the best of our knowledge, two partial answers to this question are known in the literature. If $\mathcal{W}$ is a simplex, and $f(\mathbf{w}, Q\mathbf{w} + \mathbf{q})$ is convex in $\mathbf{w}$ for any $Q, \mathbf{q}$, then a worst-case optimal policy can be readily obtained by computing $Q, \mathbf{q}$ so as to match the value of $u^*(\mathbf{w})$ on all the points in $\operatorname{ext}(\mathcal{W})$ (see Bertsimas and Goyal 2010, Ben-Tal et al. 2009, Lemma 14.3.6). This is not a surprising result, since the number of extreme points of the uncertainty set exactly matches the number of policy parameters (i.e., the degrees of freedom in the problem).

A separate instance where the construction is possible is provided in Bertsimas et al. (2010), where $\mathcal{W} = \mathcal{H}_n \overset{\text{def}}{=} [0, 1]^n$ is the unit hypercube in $\mathbb{R}^n$, $u\colon \mathcal{W} \to \mathbb{R}$, and f has the specific form

$$f(\mathbf{w}, u) = a_0 + \mathbf{a}^T\mathbf{w} + c \cdot u + g(b_0 + \mathbf{b}^T\mathbf{w} + u),$$

where a_0, b_0, $c \in \mathbb{R}$, $\mathbf{a}$, $\mathbf{b} \in \mathbb{R}^n$ are arbitrary, and $g\colon \mathbb{R} \to \mathbb{R}$ is any convex function. The proof for the latter result heavily exploits the particular functional form above, and does not lend itself to any extensions or interpretations. In particular, it fails even if one replaces $c \cdot u$ with $c(u)$, for some convex function $c\colon \mathbb{R} \to \mathbb{R}$.

In the current paper, we also focus our attention on uncertainty sets $\mathcal{W}$ that are polytopes in $\mathbb{R}^n$. More precisely, with $V = \{1, \ldots, n\}$, we consider any directed graph $G = (V, E)$, where $E \subseteq V^2$ is any set of directed edges, and are interested in uncertainty sets of the form

$$\mathcal{W} = \{\mathbf{w} \in \mathcal{H}_n\colon w_i \geqslant w_j, \forall\, (i, j) \in E\}. \tag{3}$$

It can be shown (see Tawarmalani et al. 2013, and references therein for details) that the polytope $\mathcal{W}$ in (3) has Boolean vertices, since the matrix of constraints is totally unimodular. As such, any $\mathbf{x} \in \operatorname{ext}(\mathcal{W})$ can be represented as $\mathbf{x} = \mathbf{1}_{S_x}$, for some set $S_x \subseteq V$. Furthermore, it can also be checked that the set $\operatorname{ext}(\mathcal{W})$ is a sublattice of[1] $\{0, 1\}^n$ (Topkis 1998),

$$\forall\, \mathbf{x}, \mathbf{y} \in \operatorname{ext}(\mathcal{W})\colon\ \min(\mathbf{x}, \mathbf{y}) = \mathbf{1}_{S_x\cap S_y} \in \operatorname{ext}(\mathcal{W}),$$

$$\max(\mathbf{x}, \mathbf{y}) = \mathbf{1}_{S_x\cup S_y} \in \operatorname{ext}(\mathcal{W}).$$

Among the uncertainty sets typically considered in the modeling literature, the hypercube is one example that fits the description above. Certain hyperrectangles, as well as any simplices or cross-products of simplices could also be reduced to this form via a suitable change of variables[2] (see, e.g., Tawarmalani et al. 2013). For an example of such an uncertainty set and its corresponding graph G, we direct the reader to Figure 1.

Figure 1. Example of a sublattice uncertainty set.

(a) $G = (V, E)$, where $V = \{1, 2, 3\}$ and $E = \{(1, 2), (1, 3)\}$.

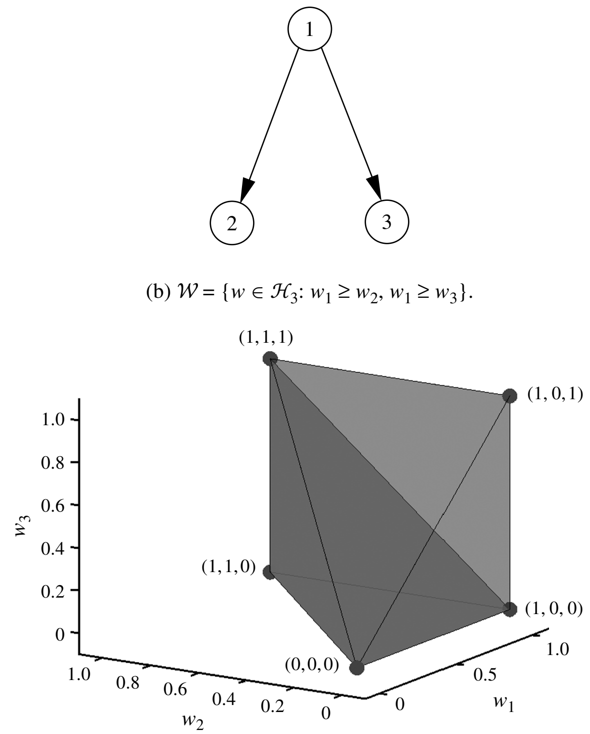

Notes. Figure 1(a) displays the graph of precedence relations, and Figure 1(b) plots the corresponding uncertainty set. Here, $\Pi^{\mathcal{W}} = \{(1,2,3), (1,3,2)\}$, and the two corresponding simplicies are $\Delta_{(1,2,3)} = \text{conv}(\{(0,0,0), (1,0,0), (1,1,0), (1,1,1)\})$ and $\Delta_{(1,3,2)} = \text{conv}(\{(0,0,0), (1,0,0), (1,0,1), (1,1,1)\})$, shown in different shades in (b). Also, $\mathcal{S}_{(0,0,0)} = \mathcal{S}_{(1,0,0)} = \mathcal{S}_{(1,1,1)} = \Pi^{\mathcal{W}}$, while $\mathcal{S}_{(1,0,1)} = \{(1,3,2)\}$, and $\mathcal{S}_{(1,1,0)} = \{(1,2,3)\}$.

For any polytope of the form (3), we define the corresponding set $\Pi^{\mathcal{W}}$ of permutations of V that are consistent with the preorder induced by the lattice $\text{ext}(\mathcal{W})$, i.e.,

$$\Pi^{\mathcal{W}} \stackrel{\text{def}}{=} \{\pi \in \Pi(V)\colon \pi^{-1}(i) \leqslant \pi^{-1}(j), \forall (i,j) \in E\}. \tag{4}$$

In other words, if $(i, j) \in E$, then i must appear before j in any permutation $\pi \in \Pi^{\mathcal{W}}$. Furthermore, with any permutation $\pi \in \Pi(V)$, we also define the simplex Δ_π obtained from vertices of $\mathcal{H}$ in the order determined by the permutation π, i.e.,

$$\Delta_\pi \stackrel{\text{def}}{=} \text{conv}\left(\left\{\mathbf{0} + \sum_{j=1}^{k} \mathbf{1}_{\pi(j)}\colon k = 0, \ldots, n\right\}\right). \tag{5}$$

It can then be checked (see, e.g., Tawarmalani et al. 2013) that any vertex $\mathbf{w} \in \text{ext}(\mathcal{W})$ belongs to several such simplices. More precisely, with $\mathbf{w} = \mathbf{1}_{S_w}$ for a particular $S_w \subseteq V$, we have

$$\mathbf{w} \in \Delta_\pi, \quad \forall \pi \in \mathcal{S}_w \stackrel{\text{def}}{=} \{\pi \in \Pi^{\mathcal{W}}\colon \{\pi(1), \ldots, \pi(|S_w|)\} = S_w\}. \tag{6}$$

In other words, $\mathbf{w}$ is contained in any simplex corresponding to a permutation π that (a) is consistent with the preorder on $\mathcal{W}$, and (b) has the indices in S_w in the first $|S_w|$ positions. An example in included in Figure 1.

Since $\text{ext}(\mathcal{W})$ is a lattice, we can consider functions $f\colon \mathcal{W} \to \mathbb{R}$ that are *supermodular on* $\text{ext}(\mathcal{W})$, i.e.,

$$f(\min(\mathbf{x}, \mathbf{y})) + f(\max(\mathbf{x}, \mathbf{y})) \geqslant f(\mathbf{x}) + f(\mathbf{y}), \quad \forall \mathbf{x}, \mathbf{y} \in \text{ext}(\mathcal{W}).$$

The properties of such functions have been studied extensively in combinatorial optimization and economics (see, e.g., Fujishige 2005, Schrijver 2003, Topkis 1998, for detailed treatments and references). The main results that are relevant for our purposes are summarized in §§EC.1 and EC.2 of the online appendix.

With these definitions, we can now state our first main result, providing a set of sufficient conditions guaranteeing the desired outcome in Problem 2.

Theorem 1. *Consider any optimization problem of the form*

$$\max_{\mathbf{w} \in \mathcal{W}} \min_{\mathbf{u}(\mathbf{w})} f(\mathbf{w}, \mathbf{u}), \tag{7}$$

having a finite optimal value, where $\mathcal{W}$ *is of the form* (3), *and* $f\colon \mathcal{W} \times \mathbb{R}^m \to \bar{\mathbb{R}}$ *is an extended-real function. Let* $\mathbf{u}^*\colon \mathcal{W} \to \mathbb{R}^m$ *denote a Bellman-optimal response of the decision maker, and* $f^*(\mathbf{w}) \stackrel{\text{def}}{=} f(\mathbf{w}, \mathbf{u}^*(\mathbf{w}))$ *be the corresponding optimal cost function. Assume the following conditions are met*:

Assumption 1 (A1). $f^*(\mathbf{w})$ *is convex on* $\mathcal{W}$ *and supermodular in* $\mathbf{w}$ *on* $\text{ext}(\mathcal{W})$.

Assumption 2 (A2). *For* $Q \in \mathbb{R}^{m \times n}$ *and* $\mathbf{q} \in \mathbb{R}^n$, *the function* $f(\mathbf{w}, Q\mathbf{w} + \mathbf{q})$ *is convex in* $(Q, \mathbf{q})$ *for any fixed* $\mathbf{w}$.

Assumption 3 (A3). *There exists* $\hat{\mathbf{w}} \in \arg\max_{\mathbf{w} \in \mathcal{W}} f^*(\mathbf{w}) \cap \text{ext}(\mathcal{W})$ *such that, with* $\mathcal{S}_{\hat{\mathbf{w}}}$ *given by* (6), *the matrices* $\{Q^\pi\}_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}}$ *and vectors* $\{\mathbf{q}^\pi\}_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}}$ *obtained as the solutions to the systems of linear equations*

$$\forall \pi \in \mathcal{S}_{\hat{\mathbf{w}}}\colon Q^\pi \mathbf{w} + \mathbf{q}^\pi = \mathbf{u}^*(\mathbf{w}), \quad \forall \mathbf{w} \in \text{ext}(\Delta_\pi), \tag{8}$$

are such that the function $f(\mathbf{w}, \bar{Q}\mathbf{w} + \bar{\mathbf{q}})$ *is convex in* $\mathbf{w}$ *and supermodular on* $\text{ext}(\mathcal{W})$, *for any* $\bar{Q}$ *and* $\bar{\mathbf{q}}$ *obtained as*

$$\bar{Q} = \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi Q^\pi, \qquad \bar{\mathbf{q}} = \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi \mathbf{q}^\pi, \quad \text{where } \lambda_\pi \geqslant 0, \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi = 1. \tag{9}$$

Then, there exist $\{\lambda_\pi\}_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}}$ *such that, with* $\bar{Q}$ *and* $\bar{\mathbf{q}}$ *given by* (9),

$$\max_{\mathbf{w} \in \mathcal{W}} f^*(\mathbf{w}) = \max_{\mathbf{w} \in \mathcal{W}} f(\mathbf{w}, \bar{Q}\mathbf{w} + \bar{\mathbf{q}}).$$

Before presenting the proof of the theorem, we provide a brief explanation and intuition for the conditions above. A more detailed discussion, together with relevant examples, is included immediately after the proof.

Note first that, since f is taken to be extended real, it can capture constraints on the policy $\mathbf{u}$ by suitable barrier functions, provided that these constraints still yield a finite optimal value, and preserve the Assumptions (A1–A3).

The interpretation and the test for conditions (A1) and (A2) are fairly straightforward. The idea behind (A3) is to consider every simplex Δ_π that contains the maximizer $\hat{\mathbf{w}}$; there are exactly $|\mathcal{S}_{\hat{\mathbf{w}}}|$ such simplices, characterized by (6). For every such simplex, one can compute a corresponding affine decision rule $Q^\pi \mathbf{w} + \mathbf{q}^\pi$ by linearly interpolating the values of the Bellman-optimal response $u^*(\mathbf{w})$ at the extreme points of Δ_π. This is exactly what is expressed in condition (8), and the resulting system is always compatible, since every such matrix-vector pair has exactly m rows, and the $n+1$ variables on each row participate in exactly $n+1$ linearly independent constraints (one for each point in the simplex). Now, the key condition in (A3) considers affine decisions rules obtained as *arbitrary* convex combinations of the rules $Q^\pi \mathbf{w} + \mathbf{q}^\pi$, and requires that the resulting cost function, obtained by using such rules, remains convex and supermodular in $\mathbf{w}$.

3.1. Proof of Theorem 1

In view of these remarks, the strategy behind the proof of Theorem 1 is quite straightforward: we seek to show that, if conditions (A1–A3) are obeyed, then one can find suitable convex coefficients $\{\lambda_\pi\}_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}}$ so that the resulting affine decision rule $\bar{Q}\mathbf{w} + \bar{\mathbf{q}}$ is worst-case optimal. To ensure the latter fact, it suffices to check that the global maximum of the function $f(\mathbf{w}, \bar{Q}\mathbf{w} + \bar{\mathbf{q}})$ is still reached at the point $\hat{\mathbf{w}}$, which is one of the maximizers of $f^*(\mathbf{w})$. Unfortunately, this is not trivial to do, since both functions $f(\mathbf{w}, \bar{Q}\mathbf{w} + \bar{\mathbf{q}})$ and $f^*(\mathbf{w})$ are convex in $\mathbf{w}$ (by (A1–A3)), and it is therefore hard to characterize their global maximizers, apart from stating that they occur at extreme points of the feasible set (Rockafellar 1970).

The first key idea in the proof is to examine the *concave envelopes* of $f(\mathbf{w}, \bar{Q}\mathbf{w} + \bar{\mathbf{q}})$ and $f^*(\mathbf{w})$, instead of the functions themselves. Recall that the concave envelope of a function $f: P \to \mathbb{R}$ on the domain P, which we denote by $\text{conc}_P(f): P \to \mathbb{R}$, is the pointwise smallest concave function that overestimates f on P (Rockafellar 1970) and always satisfies $\arg\max_{\mathbf{x}\in P} f \subseteq \arg\max_{\mathbf{x}\in P} \text{conc}_P(f)$. (The interested reader is referred to §EC.2 of the online appendix for a short overview of background material on concave envelopes, and to the papers Tardella 2008 or Tawarmalani et al. 2013 for other useful references.)

In this context, a central result used repetitively throughout our analysis is the following characterization for the concave envelope of a function that is convex *and supermodular* on a polytope of the form (3).

Lemma 1. *If $f^*: \mathcal{W} \to \mathbb{R}$ is convex on $\mathcal{W}$ and supermodular on* $\text{ext}(\mathcal{W})$, *then the following results hold*:

1. *The concave envelope of f^* on $\mathcal{W}$ is given by the Lovász extension of f^* restricted to* $\text{ext}(\mathcal{W})$:

$$\text{conc}_{\mathcal{W}}(f^*)(\mathbf{w}) = f^*(\mathbf{0}) + \min_{\pi \in \Pi^{\mathcal{W}}} \sum_{i=1}^{n} \left[f^*\left(\sum_{j=1}^{i} \mathbf{1}_{\pi(j)}\right) - f^*\left(\sum_{j=1}^{i-1} \mathbf{1}_{\pi(j)}\right)\right] w_{\pi(i)}. \quad (10)$$

2. *The inequalities $(\mathbf{g}^\pi)^T \mathbf{w} + g_0 \geqslant f^*(\mathbf{w})$ defining nonvertical facets of $\text{conc}_{\mathcal{W}}(f^*)$ are given by the set $\text{ext}(\mathcal{D}_{f^*,\mathcal{W}}) = \{(\mathbf{g}^\pi, g_0) \in \mathbb{R}^{n+1}: \pi \in \Pi^{\mathcal{W}}\}$, where*

$$g_0 \stackrel{\text{def}}{=} f^*(\mathbf{0}), \quad \mathbf{g}^\pi \stackrel{\text{def}}{=} \sum_{i=1}^{n} \left[f^*\left(\sum_{j=1}^{i} \mathbf{1}_{\pi(j)}\right) - f^*\left(\sum_{j=1}^{i-1} \mathbf{1}_{\pi(j)}\right)\right] \mathbf{1}_{\pi(i)}, \quad \forall \pi \in \Pi^{\mathcal{W}}. \quad (11)$$

3. *The polyhedral subdivision of $\mathcal{W}$ yielding the concave envelope is given by the restricted Kuhn triangulation,*

$$\mathcal{K}^{\mathcal{W}} \stackrel{\text{def}}{=} \{\Delta_\pi: \pi \in \Pi^{\mathcal{W}}\}.$$

The result is essentially Corollary EC.2 in the online appendix, to which we direct the interested reader for more details. This lemma essentially establishes that the concave envelope of a function f^* that is convex and supermodular on an integer sublattice of $\{0,1\}^n$ is determined by the Lovász extension (Lovász 1983). The latter function is *polyhedral* (i.e., piecewise affine), and is obtained by affinely interpolating the function f^* on all the simplicies in the Kuhn triangulation $\mathcal{K}^{\mathcal{W}}$ of the hypercube (see §EC.2.1 of the online appendix). A plot of such a function f and its concave envelope is included in Figure 2.

With this powerful lemma, we can now provide a result that brings us very close to a complete proof of Theorem 1.

Lemma 2. *Suppose $f^*: \mathcal{W} \to \mathbb{R}$ is convex on $\mathcal{W}$ and supermodular on $\text{ext}(\mathcal{W})$. Consider an arbitrary $\hat{\mathbf{w}} \in \text{ext}(\mathcal{W}) \cap \arg\max_{\mathbf{w}\in\mathcal{W}} f^*(\mathbf{w})$, and let $\mathbf{g}^\pi$ be given by* (11). *Then the following results hold*:

1. *For any $\mathbf{w} \in \mathcal{W}$, we have*

$$f^*(\mathbf{w}) \leqslant f^*(\hat{\mathbf{w}}) + (\mathbf{w} - \hat{\mathbf{w}})^T \mathbf{g}^\pi, \quad \forall \pi \in \mathcal{S}_{\hat{\mathbf{w}}}. \quad (12)$$

2. *There exists a set of convex weights $\{\lambda_\pi\}_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}}$ such that $\mathbf{g} = \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi \mathbf{g}^\pi$ satisfies*

$$(\mathbf{w} - \hat{\mathbf{w}})^T \mathbf{g} \leqslant 0, \quad \forall \mathbf{w} \in \mathcal{W}. \quad (13)$$

Proof. The proof is rather technical, and we defer it to §EC.3 of the online appendix. □

For a geometric intuition of these results, we refer to Figure 2. In particular, the first claim simply states that the vectors $\mathbf{g}^\pi$ corresponding to simplicies that contain $\hat{\mathbf{w}}$ are valid supergradients of the function f^* at $\hat{\mathbf{w}}$; this is a direct

Figure 2. A convex and supermodular function (a) and its concave envelope (b).

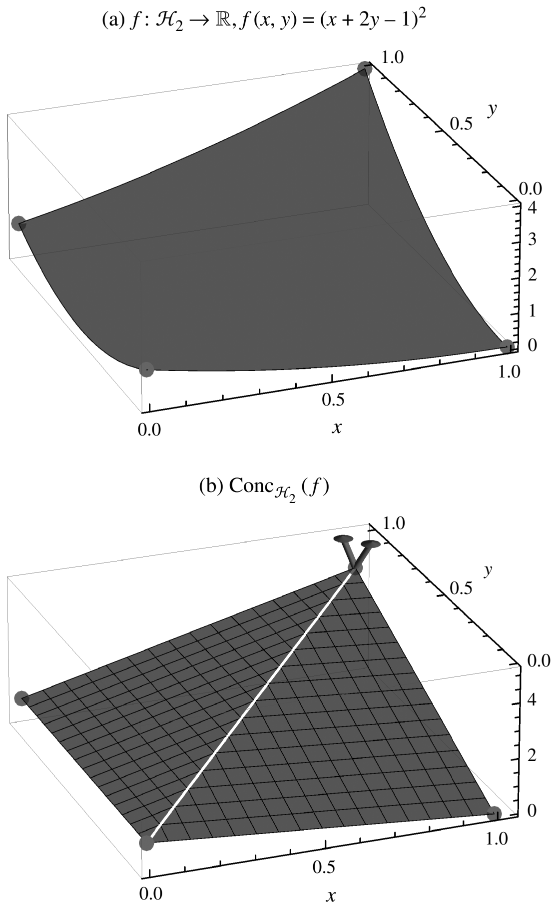

Notes. Here, $\mathscr{W} = \mathcal{H}_2$, $\Pi^{\mathscr{W}} = \{(1, 2), (2, 1)\}$, and $\mathcal{K}^{\mathscr{W}} = \{\Delta_{(1,2)}, \Delta_{(2,1)}\}$, where $\Delta_{(1,2)} = \mathrm{conv}(\{(0, 0), (1, 0), (1, 1)\}$ and $\Delta_{(2,1)} = \mathrm{conv}(\{(0, 0), (0, 1), (1, 1)\})\}$. The plot in Figure 2(b) also shows the two normals of nonvertical facets of $\mathrm{conc}_{\mathscr{W}}(f)$, corresponding to $\mathbf{g}^{(1,2)}$ and $\mathbf{g}^{(2,1)}$.

consequence of Lemma 1, since any such vectors $\mathbf{g}^\pi$ are also supergradients for the concave envelope $\mathrm{conc}_{\mathscr{W}}(f^*)$ at $\hat{\mathbf{w}}$. The second claim states that one can always find a convex combination of the supergradients $\mathbf{g}^\pi$ that yields a supergradient $\mathbf{g}$ that is not a direction of increase for the function f^* when moving in *any* feasible direction away from $\hat{\mathbf{w}}$ (i.e., while remaining in $\mathscr{W}$).

With this lemma, we can now complete the proof of our main result.

Proof of Theorem 1. Consider any $\hat{\mathbf{w}}$ satisfying the requirement (A3). Note that the system of equations in (8) is uniquely defined, since each row of the matrix Q^π and the vector $\mathbf{q}^\pi$ participate in exactly $n+1$ constraints, and the corresponding constraint matrix is nonsingular. Furthermore, from the definition of Δ_π in (5), we have that $\mathbf{0} \in \mathrm{ext}(\Delta_\pi)$, $\forall \pi \in \mathcal{S}_{\hat{\mathbf{w}}}$, so that the system in (8) yields $\mathbf{q}^\pi = \mathbf{u}^*(\mathbf{0})$, $\forall \pi \in \mathcal{S}_{\hat{\mathbf{w}}}$.

By Lemma 2, consider the set of weights $\{\lambda_\pi\}_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}}$, such that $\mathbf{g} = \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi \mathbf{g}^\pi$ satisfies $(\mathbf{w} - \hat{\mathbf{w}})^T \mathbf{g} \leqslant 0$, $\forall \mathbf{w} \in \mathscr{W}$. We claim that the corresponding $\bar{Q} = \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi Q^\pi$, and $\bar{\mathbf{q}} = \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi \mathbf{q}^\pi$ provide the desired affine policy $\bar{Q}\mathbf{w} + \bar{\mathbf{q}}$ such that

$$\max_{\mathbf{w} \in \mathscr{W}} f^*(\mathbf{w}) = \max_{\mathbf{w} \in \mathscr{W}} f(\mathbf{w}, \bar{Q}\mathbf{w} + \bar{\mathbf{q}}).$$

To this end, note that, by (A3), the functions $f(\mathbf{w}, \bar{Q}\mathbf{w} + \bar{\mathbf{q}})$ and $f^\pi(\mathbf{w}) \stackrel{\text{def}}{=} f(\mathbf{w}, Q^\pi \mathbf{w} + \mathbf{q}^\pi)$, $\forall \pi \in \mathcal{S}_{\hat{\mathbf{w}}}$ are convex in $\mathbf{w}$ and supermodular on $\mathrm{ext}(\mathscr{W})$. Also, by construction,

$$\begin{aligned} &\forall \pi \in \mathcal{S}_{\hat{\mathbf{w}}}, \quad f^\pi(\mathbf{w}) = f^*(\mathbf{w}), \quad \forall \mathbf{w} \in \mathrm{ext}(\Delta_\pi), \\ &f(\hat{\mathbf{w}}, \bar{Q}\hat{\mathbf{w}} + \bar{\mathbf{q}}) = f^*(\hat{\mathbf{w}}). \end{aligned} \tag{14}$$

Thus, for any $\pi \in \mathcal{S}_{\hat{\mathbf{w}}}$, the supergradient $\mathbf{g}^\pi$ defined for the function f^* in (11) remains a valid supergradient for f^π at $\mathbf{w} = \hat{\mathbf{w}}$. As such, relation (12) also holds for each function f^π, i.e.,

$$f^\pi(\mathbf{w}) \leqslant f^\pi(\hat{\mathbf{w}}) + (\mathbf{w} - \hat{\mathbf{w}})^T \mathbf{g}^\pi, \quad \forall \pi \in \mathcal{S}_{\hat{\mathbf{w}}}. \tag{15}$$

The following reasoning then concludes our proof

$$\begin{aligned} \forall \mathbf{w} \in \mathscr{W}, \quad f(\mathbf{w}, \bar{Q}\mathbf{w} + \bar{\mathbf{q}}) &\stackrel{[A2]}{\leqslant} \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi f(\mathbf{w}, Q^\pi \mathbf{w} + \mathbf{q}^\pi) \\ &\stackrel{(15)}{\leqslant} \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi [f^\pi(\hat{\mathbf{w}}) + (\mathbf{w} - \hat{\mathbf{w}})^T \mathbf{g}^\pi] \\ &\stackrel{(14)}{=} f(\hat{\mathbf{w}}, \bar{Q}\hat{\mathbf{w}} + \bar{\mathbf{q}}) \\ &\quad + \sum_{\pi \in \mathcal{S}_{\hat{\mathbf{w}}}} \lambda_\pi (\mathbf{w} - \hat{\mathbf{w}})^T \mathbf{g}^\pi \\ &\stackrel{(13)}{\leqslant} f(\hat{\mathbf{w}}, \bar{Q}\hat{\mathbf{w}} + \bar{\mathbf{q}}). \end{aligned}$$

3.2. Examples and Discussion of Existential Conditions

We now proceed to discuss the conditions in Theorem 1, and relevant examples of functions satisfying them. We implicitly assume throughout that the optimal value of the problem in (7) is finite. Condition (A1) can be generally checked by performing suitable comparative statics analyses. For instance, $f^*(\mathbf{w})$ will be convex in $\mathbf{w}$ if $f(\mathbf{w}, \mathbf{u})$ is jointly convex in $(\mathbf{w}, \mathbf{u})$, since partial minimization preserves convexity (Rockafellar 1970). For supermodularity of f^*, more structure is typically needed on $f(\mathbf{w}, \mathbf{u})$.

One such example, which proves instrumental in the analysis of Problem 1, is $f(\mathbf{w}, u) = c(u) + g(b_0 + \mathbf{b}^T \mathbf{w} + u)$. To understand the significance in an inventory setting, the reader can think of $\mathbf{w}$ as a sequence of historical demands, with $b_0 + \mathbf{b}^T \mathbf{w}$ denoting the (affine) dependency of the on-hand inventory on $\mathbf{w}$, u denoting an order quantity with associated ordering cost $c(u)$ (which includes any potential constraints on u), and g the value to go, depending on the inventory position after receiving the order u.

Proposition 1. *Let $f(\mathbf{w}, u) = c(u) + g(b_0 + \mathbf{b}^T\mathbf{w} + u)$, where c, g: $\mathbb{R} \to \bar{\mathbb{R}}$ are arbitrary proper*[3] *convex functions, and $\mathbf{b} \geqslant 0$ or $\mathbf{b} \leqslant 0$. Then, condition* (A1) *is satisfied.*

Proof. If the optimal value in (7) is finite, then f^* must be a real-valued function on $\mathscr{W}$. Since f is jointly convex in $\mathbf{w}$ and u, f^* is convex. Furthermore, note that f^* only depends on $\mathbf{w}$ through $\mathbf{b}^T\mathbf{w}$, i.e., $f^*(\mathbf{w}) = \tilde{f}(\mathbf{b}^T\mathbf{w})$, for some convex $\tilde{f}$. Therefore, since $\mathbf{b} \geqslant 0$ or $\mathbf{b} \leqslant 0$, f^* is supermodular (see Example EC.1 in the online appendix). □

Condition (A2) can also be tested by directly examining the function f. For instance, if f is jointly convex in $\mathbf{w}$ and $\mathbf{u}$, then (A2) is trivially satisfied, as is the case in the example of Proposition 1.

In practice, the most cumbersome condition to test is undoubtedly (A3). Typically, a combination of comparative statics analyses and structural properties on the function f will be needed. We exhibit how such techniques can be used by making reference, again, to the example in Proposition 1.

Proposition 2. *Let $f(\mathbf{w}, u) = c(u) + g(b_0 + \mathbf{b}^T\mathbf{w} + u)$, where c, g: $\mathbb{R} \to \bar{\mathbb{R}}$ are arbitrary proper convex functions, and $\mathbf{b} \geqslant 0$ or $\mathbf{b} \leqslant 0$. Then, condition* (A3) *is satisfied.*

Proof. Let $h(x, y) \overset{\text{def}}{=} c(y) + g(x + y)$. It is shown in Lemma EC.1 of the online appendix that $\arg\min_y h(x, y)$ is decreasing in x, and $x + \arg\min_y h(x, y)$ is increasing in x.

Consider any $\hat{\mathbf{w}} \in \arg\max_{\mathbf{w}\in\mathscr{W}} f^*(\mathbf{w}) \cap \text{ext}(\mathscr{W})$. In this case, the construction in (8) becomes

$$\forall\, \pi \in \mathscr{S}_{\hat{\mathbf{w}}}\colon\ (\mathbf{q}^\pi)^T\mathbf{w} + q_0^\pi = u^*(\mathbf{w}) \equiv y^*(b_0 + \mathbf{b}^T\mathbf{w}),\qquad \forall\, \mathbf{w} \in \text{ext}(\Delta_\pi),$$

for some $y^*(x) \in \arg\min_y h(x, y)$. Note that, since y is scalar here, the affine parametrization is given by a (row) vector $(\mathbf{q}^\pi)^T \in \mathbb{R}^{1\times n}$ and a scalar q_0^π, instead of a matrix $Q^\pi \in \mathbb{R}^{m\times n}$ and a vector $\mathbf{q}^\pi \in \mathbb{R}^m$, respectively. We claim that

$$\mathbf{b} \geqslant 0 \ \Rightarrow\ \mathbf{q}^\pi \leqslant 0 \quad \text{and} \quad \mathbf{b} + \mathbf{q}^\pi \geqslant 0, \quad \forall\, \pi \in \mathscr{S}_{\hat{\mathbf{w}}} \tag{16a}$$

$$\mathbf{b} \leqslant 0 \ \Rightarrow\ \mathbf{q}^\pi \geqslant 0 \quad \text{and} \quad \mathbf{b} + \mathbf{q}^\pi \leqslant 0, \quad \forall\, \pi \in \mathscr{S}_{\hat{\mathbf{w}}}. \tag{16b}$$

We prove the first claim (the second follows analogously). Since $\mathbf{0} \in \text{ext}(\Delta_\pi)$, we have $q_0^\pi = y^*(b_0)$. If $\mathbf{b} \geqslant 0$, then the monotonicity of $y^*(x)$ implies that

$$q_0^\pi + q_i^\pi = y^*(b_0 + b_i) \leqslant y^*(b_0), \quad \forall\, i \in \{1, \ldots, n\},$$

which implies that $\mathbf{q}^\pi \leqslant 0$. Similarly, the monotonicity of $x + y^*(x)$ implies that $\mathbf{b} + \mathbf{q}^\pi \geqslant 0$.

With the previous two claims, it can be readily seen that the functions

$$f^\pi(\mathbf{w}) = c\big((\mathbf{q}^\pi)^T\mathbf{w} + q_0^\pi\big) + g\big(b_0 + q_0^\pi + (\mathbf{b} + \mathbf{q}^\pi)^T\mathbf{w}\big)$$

are convex in $\mathbf{w}$ and supermodular on ext($\mathscr{W}$) (see Example EC.1 in the online appendix), and that the same conclusion holds for affine policies given by arbitrary convex combinations of $(\mathbf{q}^\pi, q_0^\pi)$, hence (A3) must hold. □

In view of Propositions 1 and 2, we have the following example where Theorem 1 readily applies, which will prove essential in the discussion of the two-echelon example of Problem 1.

Lemma 3. *Let $f(\mathbf{w}, u) = h(\mathbf{w}) + c(u) + g(b_0 + \mathbf{b}^T\mathbf{w} + u)$, where h: $[\mathbf{l}, \mathbf{r}] \to \mathbb{R}$ is convex and supermodular on the lattice* ext$([\mathbf{l}, \mathbf{r}])$ *for some $\mathbf{l} \leqslant \mathbf{r} \in \mathbb{R}^n$, and c, g: $\mathbb{R} \to \bar{\mathbb{R}}$ are arbitrary, proper convex functions. Then, if either $\mathbf{b} \geqslant 0$, $\mathbf{b} \leqslant 0$ or h is affine, there exist $\mathbf{q} \in \mathbb{R}^n$, $q_0 \in \mathbb{R}$ such that*

$$\max_{\mathbf{w}\in[\mathbf{l},\mathbf{r}]} f^*(\mathbf{w}) = \max_{\mathbf{w}\in[\mathbf{l},\mathbf{r}]} f(\mathbf{w}, \mathbf{q}^T\mathbf{w} + q_0) \tag{17a}$$

$$\text{sign}(\mathbf{q}) = -\,\text{sign}(\mathbf{b}) \tag{17b}$$

$$\text{sign}(\mathbf{b} + \mathbf{q}) = \text{sign}(\mathbf{b}). \tag{17c}$$

Proof. Assume first that $\mathbf{l} = \mathbf{0}$ and $\mathbf{r} = \mathbf{1}$. (1) When $\mathbf{b} \geqslant 0$ or $\mathbf{b} \leqslant 0$, the results follow directly from Propositions 1 and 2 (note that adding the convex and supermodular function $\tilde{h}$ does not change any of the arguments there). The proofs for the sign relations concerning $\mathbf{q}$ follow from (16a) and (16b), by recognizing that the same inequalities hold for any convex combination of the vectors $\mathbf{q}^\pi$. (2) When h is affine, the case with an arbitrary $\mathbf{b}$ can be transformed, by a suitable linear change of variables for $\mathbf{w}$, to a case with $\mathbf{b} \geqslant 0$ and modified b_0 and affine h.

The case with arbitrary $\mathbf{l} \leqslant \mathbf{r}$ can be reduced to $\mathbf{l} = \mathbf{0}$ and $\mathbf{r} = \mathbf{1}$ by a linear change of variables on $\mathbf{w}$, which does not affect the supermodularity and convexity of the functions in question. □

The latter result directly generalizes that of Bertsimas et al. (2010) in several ways, by allowing the possibility of a nonaffine h, a nonlinear c, and also a nonhypercube uncertainty set (the conclusions hold if the domain of h is $\mathscr{W}$, instead of $[\mathbf{l}, \mathbf{r}]$).

3.3. Application to Problem 1

In this section, we revisit the production planning model discussed in Problem 1 of the Introduction, where the full power of the results introduced in §3 can be used to derive the optimality of ordering policies that are affine in historical demands.

As remarked in the Introduction, a very similar model has been originally considered in Ben-Tal et al. (2005b, 2009); we first describe our model in detail, and then discuss how it relates to that in the other two references.

Let $1, \ldots, T$ denote the finite planning horizon, and introduce the following variables:

• $\mathbf{K} \in \mathbb{R}^m$: the strategic decisions, taken ahead of the selling season, with an associated cost $r(\mathbf{K})$.

• q_t: the realized order quantity from the retailer in period t. The corresponding cost incurred by the retailer is $c_t(q_t, \mathbf{K})$, reflecting the nonlinear dependency on the strategic decisions $\mathbf{K}$. To reflect the possibility of constraints of the form $L_t \leqslant q_t \leqslant U_t$, we take $c_t: \mathbb{R} \times \mathbb{R}^m \to \bar{\mathbb{R}}$.

• I_t: the inventory on the premises of the retailer at the *beginning* of period $t \in \{1, \ldots, T\}$. Let $h_t(I_{t+1}, \mathbf{K})$ denote the holding/backlogging cost incurred at the *end* of period t, also allowed to depend on the strategic decisions. To allow constraints on the inventory of the form $L_t^x \leqslant I_t \leqslant U_t^x$, we take $h_t: \mathbb{R} \times \mathbb{R}^m \to \bar{\mathbb{R}}$.

• d_t: unknown customer demand in period t. We assume that the retailer has very limited information about the demands, so that only bounds are available, $d_t \in \mathcal{D}_t = [\underline{d}_t, \bar{d}_t]$.

The problem of computing strategic and ordering decisions that would minimize the system-level cost in the worst case can then be rewritten as

$$\begin{aligned}\min_{\mathbf{K}} \quad & \Bigg[r(\mathbf{K}) + \min_{q_1} \Bigg[c_1(q_1, \mathbf{K}) + \max_{d_1 \in \mathcal{D}_1} \Bigg[h_1(I_2, \mathbf{K}) + \cdots \\ & + \min_{q_T} \Bigg[c_T(q_T, \mathbf{K}) + \max_{d_T \in \mathcal{D}_T} h_T(I_{T+1}, \mathbf{K}) \Bigg] \ldots \Bigg] \Bigg] \Bigg] \\ \text{s.t.} \quad & I_{t+1} = I_t + q_t - d_t, \quad \forall\, t \in \{1, 2, \ldots, T\}.\end{aligned}$$

By introducing the class of ordering policies that depend on the history of observed demands,

$$q_t: \mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}_{t-1} \to \mathbb{R}, \tag{18}$$

we claim that the theorems of §3 can be used to derive the following structural results.

Theorem 2. *Consider a fixed* $\mathbf{K}$, *and assume the corresponding optimal worst-case cost is finite. If the costs* $c_t(q, \mathbf{K})$ *and* $h_t(I, \mathbf{K})$ *are proper convex, then the following results hold*:

1. *Ordering policies that depend* affinely *on the history of demands are worst-case optimal.*
2. *Any such worst-case optimal order occurring after period* t *is partially satisfying the demands that are still backlogged in period* t.

Before presenting the proof, we discuss the result, and comment on the related literature. The first claim confirms that ordering policies depending affinely on historical demands are (worst-case) optimal, as soon as the preseason (strategic) decisions are fixed, provided that all the costs are proper convex. The second claim provides a structural decomposition of the worst-case optimal affine ordering policies: every such order placed in or after period t can be seen as partially satisfying the demands that are still backlogged in period t, with the free terms (of the affine form) corresponding to safety stock that is built in anticipation for future increased demands. The latter point should become clear after the formal statement and discussion following Lemma 4.

The model is related to that in Ben-Tal et al. (2005b, 2009) in several ways. In the latter model, the vector $\mathbf{K}$ consists of a set of precommitments for orders, $p_1, \ldots, p_T$, one for each period in the selling season. The costs have the specific form

$$\begin{aligned} r(\mathbf{K}) &= \beta_t^- \max(0, p_{t-1} - p_t) + \beta_t^+ \max(0, p_t - p_{t-1}), \\ c_t(q_t, \mathbf{K}) &= \tilde{c}_t \cdot q_t + \alpha_t^- \max(0, p_t - q_t) \\ &\qquad + \alpha_t^+ \max(0, q_t - p_t) + \mathbf{1}_{q_t \in [L_t, U_t]}, \\ h_t(I_{t+1}, \mathbf{K}) &= \max(\tilde{h}_t I_{t+1}, -b_t I_{t+1}). \end{aligned} \tag{19}$$

Here, $\tilde{c}_t$ is the per-unit ordering cost, $\alpha_t^{\pm}$ are the penalties for overordering/underordering (respectively) relative to the precommitments, $\beta_t^{\pm}$ are penalties for differences in pre-commitments for consecutive periods, $\tilde{h}_t$ is the per-unit holding cost, and b_t is the per-unit backlogging cost; $\mathbf{1}_{q_t \in [L_t, U_t]}$ is the indicator function, equal to zero if $q_t \in [L_t, U_t]$, and $+\infty$ otherwise. Such costs are clearly proper convex, and hence fit the conditions of Theorem 2. Note that our model allows more general convex production costs, for instance, reflecting the purchase of units beyond the installed capacity at the supplier, e.g., from a different supplier or an open market, resulting in an extra cost $c_t^{om} \max(0, q_t - K_t)$. More general costs are also possible for holding and backlogging, as well as constraints on the on-hand inventory.

The one feature present in Ben-Tal et al. (2005b), but absent from our model, are cumulative order bounds, of the form

$$\hat{L}_t \leqslant \sum_{k=1}^{t} q_t \leqslant \hat{H}_t, \quad \forall\, t \in \{1, \ldots, T\}.$$

Such constraints have been shown to preclude the optimality of ordering policies that are affine in historical demands, even in the simpler model of Bertsimas et al. (2010). Therefore, the result in Theorem 2 shows that these constraints are, in fact, the *only* modeling component in Ben-Tal et al. (2005b, 2009) that hinders the optimality of affine ordering policies.

The result above also strictly generalizes that of Bertsimas et al. (2010) by allowing arbitrary convex ordering costs c_t. As argued in the Introduction, this is a relevant modeling extension, by allowing the possibility of capturing multiple vendors with different production or distribution technologies.

We also mention some related literature in operations management to which our result might bear some relevance. A particular demand model, which has garnered attention in various operational problems, is the martingale model of forecast evolution (see Hausman 1969, Heath and Jackson 1994, Graves et al. 1998, Chen and Lee 2009, Bray and Mendelson 2012, and references therein), whereby demands in future periods depend on a set of external demand shocks, which are observed in each period. In such models, it is customary to consider so-called generalized order-up-to inventory policies, whereby orders in period t depend in an *affine* fashion on demand signals observed up to period t (see Graves et al. 1998, Chen and Lee 2009, Bray and Mendelson 2012). Typically, the affine forms are considered for simplicity, and, to the best of our knowledge, there are no proofs concerning their optimality in the underlying models. In this sense, if we interpret the disturbances in our model as corresponding to particular demand shocks, our results may provide evidence that affine ordering policies (in historical demand shocks) are provably optimal for particular finite horizon, robust counterparts of the models.

3.3.1. Dynamic Programming Solution. In terms of solution methods, note that Problem 1 can be formulated as a dynamic program (Ben-Tal et al. 2005b, 2009). In particular, for a fixed $\mathbf{K}$, the state space of the problem is one-dimensional, i.e., the inventory I_t, and Bellman recursions can be written to determine the underlying optimal ordering policies $q_t^*(I_t, \mathbf{K})$ and value functions $J_t^*(I_t, \mathbf{K})$,

$$\begin{aligned} J_t(I, \mathbf{K}) &= \min_q \big[c_t(q, \mathbf{K}) + g_t(I+q, \mathbf{K})\big], \\ g_t(y, \mathbf{K}) &\stackrel{\text{def}}{=} \max_{d \in \mathcal{D}_t} \big[h_t(y-d, \mathbf{K}) + J_{t+1}^*(y-d, \mathbf{K})\big], \end{aligned} \tag{20}$$

where $J_{T+1}(I, \mathbf{K})$ can be taken to be 0 or some other *convex* function of I, if salvaging inventory is an option (see Ben-Tal et al. 2005b for details). With this approach, one can derive the following structural properties concerning the optimal policies and value functions.

LEMMA 4. *Consider a fixed* $\mathbf{K}$ *such that the corresponding optimal worst-case cost is finite. Then, the following results hold*:

1. *Any optimal order quantity is nonincreasing in starting inventory, i.e.,* $q_t^*(I_t, \mathbf{K})$ *is nonincreasing in* I_t.
2. *The optimal inventory position after ordering is nondecreasing in starting inventory, i.e.,* $I_t + q_t^*(I_t, \mathbf{K})$ *is nondecreasing in* I_t.
3. *The value functions* $J_t^*(I_t, \mathbf{K})$ *and* $g_t(y, \mathbf{K})$ *are convex in* I_t *and* y, *respectively.*

PROOF. These properties are well known in the literature on inventory management (see Heyman and Sobel 1984, Examples 8–15; Bensoussan et al. 1983, Proposition 3.1; or Topkis 1998, Theorem 3.10.2), and follow by backward induction, and a repeated application of Lemma EC.1 in the online appendix. We omit the complete details because of space considerations. □

When the convex costs c_t are also piecewise affine, the optimal orders follow a *generalized base stock policy*, whereby a different base stock is prescribed for every linear piece in c_t (see Porteus 2002).

In terms of completing the solution of the original problem, once the value function $J_1(I_1, \mathbf{K})$ is available, one can solve the problem $\min_{\mathbf{K}} J_1(I_1, \mathbf{K})$. However, as outlined in Ben-Tal et al. (2005b, 2009), such an approach would encounter several difficulties in practice: (i) one may have to discretize I_t and q_t, and hence only produce an approximate value for J_1; (ii) the DP would have to be solved for any possible choice of $\mathbf{K}$; (iii) $J_1(I_1, \mathbf{K})$ would, in general, be nonsmooth; and (iv) the DP solution would provide no subdifferential information for J_1, leading to the use of zero-order (i.e., gradient-free) methods for solving the resulting first-stage problem, which exhibit notoriously slow convergence. These issues would be further exacerbated if some of the decisions in $\mathbf{K}$ were discrete.

These results are in stark contrast with Theorem 2, which argues that affine ordering policies remain optimal for *arbitrary* convex ordering cost, i.e., the complexity of the policy does not increase with the complexity of the cost function. Furthermore, as we argue in §4, the *exact* solution for the case of piecewise affine costs (such as those considered in Ben-Tal et al. 2005b, 2009) can actually be obtained by solving a single LP, with manageable size.

3.3.2. Proof of Theorem 2. To simplify the notation, let $\mathbf{d}_{[t]} \stackrel{\text{def}}{=} (d_1, \ldots, d_{t-1})$ denote the vector of demands known at the beginning of period t, residing in $\mathcal{D}_{[t]} \stackrel{\text{def}}{=} \mathcal{D}_1 \times \cdots \times \mathcal{D}_{t-1}$. Whenever $\mathbf{K}$ is fixed, we suppress the dependency on $\mathbf{K}$ for all quantities of interest, such as q_t^*, J_t^*, c_t, g_t, etc. The following lemma proves the desired result in Theorem 2.

LEMMA 5. *Consider a fixed* $\mathbf{K}$ *such that the corresponding optimal worst-case cost is finite. For every period* $t \in \{1, \ldots, T\}$, *one can find an* affine *ordering policy* $q_t^{\text{aff}}(\mathbf{d}_{[t]}) = \mathbf{q}_t^T \mathbf{d}_{[t]} + q_{t,0}$ *such that*

$$J_1^*(I_1) = \max_{\mathbf{d}_{[t+1]} \in \mathcal{D}_{[t+1]}} \left[\sum_{k=1}^{t} \big(c_k(q_k^{\text{aff}}) + h_k(I_{k+1}^{\text{aff}})\big) + J_{t+1}^*(I_{t+1}^{\text{aff}})\right], \tag{21}$$

where $I_k^{\text{aff}}(\mathbf{d}_{[k]}) = \mathbf{b}_k^T \mathbf{d}_{[k]} + b_{k,0}$ *denotes the affine dependency of the inventory* I_k *on historical demands, for any* $k \in \{1, \ldots, t\}$. *Furthermore, we also have*

$$\mathbf{b}_t \leqslant \mathbf{0}, \quad \mathbf{q}_t \geqslant \mathbf{0}, \quad \mathbf{q}_t + \mathbf{b}_t \leqslant \mathbf{0}. \tag{22}$$

Let us first interpret the main statements. Equation (21) guarantees that using the affine ordering policies in periods $k \in \{1, \ldots, t\}$ (and then proceeding with the Bellman-optimal decisions in periods $t+1, \ldots, T$) does not increase the overall optimal worst-case cost. As such, it proves the first part of Theorem 2.

Relation (22) confirms the structural decomposition of the ordering policies: if a particular demand d_k no longer appears in the backlog at the beginning of period t (i.e., $\mathbf{b}_t^T \mathbf{1}_k = 0$), then the current ordering policy does not depend on d_k (i.e., $\mathbf{q}_t^T \mathbf{1}_k = 0$). Furthermore, if a fraction $-b_{t,k} \in (0, 1]$ of demand d_k is still backlogged in period t, the order q_t^{aff} will satisfy a fraction $q_{t,k} \in [0, -b_{t,k}]$ of this demand. Put differently, the affine orders decompose the fulfillment of any demand d_k into (a) existing stock in period k and

(b) partial orders in periods $k, \ldots, T$, which is exactly the content of the second part of Theorem 2.

Proof of Lemma 5. The proof is by forward induction on t. At $t = 1$, an optimal *constant* order is available from the DP solution, $q_1^{\text{aff}} = q_1^*(I_1)$. Also, since $I_2 = I_1 + q_1^{\text{aff}} - d_1$, we have $\mathbf{b}_2 \leqslant 0$.

Assuming the induction is true at stages $k \in \{1, \ldots, t-1\}$, consider the problem solved by nature at time $t-1$, given by (21). The cumulative historical costs in stages $1, \ldots, t-1$ are given by

$$\tilde{h}_t(\mathbf{d}_{[t]}) \stackrel{\text{def}}{=} \sum_{k=1}^{t-1} \big(c_k(q_k^{\text{aff}}) + h_k(I_{k+1}^{\text{aff}})\big)$$
$$= \sum_{k=1}^{t-1} \big[c_k(\mathbf{q}_k^T \mathbf{d}_{[k]} + q_{k,0}) + h_k(\mathbf{b}_{k+1}^T \mathbf{d}_{[k+1]} + b_{k+1,0})\big].$$

By the induction hypothesis, $\mathbf{q}_k \geqslant \mathbf{0}$, $\mathbf{b}_k \leqslant \mathbf{0}$, $\forall k \in \{1, \ldots, t-1\}$, and $\mathbf{b}_t \leqslant \mathbf{0}$. Therefore, since c_k and h_k are proper convex, the function $\tilde{h}_t$ is convex and *supermodular* in $\mathbf{d}_{[t]}$ on the lattice $\text{ext}(\mathcal{D}_{[t]})$ (see Example EC.1 in the online appendix). Recalling that J_t^* is derived from the Bellman recursions (20), i.e.,

$$J_t^*(I_t) = \min_q [c_t(q) + g_t(I_t + q)],$$

we obtain that Equation (21) can be rewritten equivalently as

$$J_1^*(I_1) = \max_{\mathbf{d} \in \mathcal{D}_{[t]}} \Big[\tilde{h}_t(\mathbf{d}) + \min_{q_t} [c_t(q_t) + g_t(\mathbf{b}_t^T \mathbf{d} + b_{t,0} + q_t)]\Big]. \tag{23}$$

In this setup, we can directly invoke the result of Lemma 3 to conclude that there exists an *affine* ordering policy $q_t^{\text{aff}}(\mathbf{d}_{[t]}) \stackrel{\text{def}}{=} \mathbf{q}_t^T \mathbf{d}_{[t]} + q_{t,0}$, that is worst-case optimal for problem (23) above. Furthermore, Lemma 3 also states that $\text{sign}(\mathbf{q}_t) = -\text{sign}(\mathbf{b}_t)$ and $\text{sign}(\mathbf{q}_t + \mathbf{b}_t) = \text{sign}(\mathbf{b}_t)$, which completes the proof.

4. Discussion of Problem 3

As suggested in the Introduction, the sole knowledge that affine decision rules are optimal might not necessarily provide a "simple" computational procedure for generating them. An immediate example of this is Problem 1 itself: to find optimal affine ordering policies $q_t^{\text{aff}}(\mathbf{d}_{[t]}) = \mathbf{q}_t^T \mathbf{d}_{[t]} + q_{t,0}$ for any fixed $\mathbf{K}$, we would have to solve the following optimization problem:

$$\min_{\{\mathbf{q}_t, q_{t,0}\}_{t=1}^T} \max_{\mathbf{d}_{[T+1]} \in \mathcal{D}_{[T+1]}} \sum_{t=1}^T \Big[c_t(q_t^{\text{aff}}) + h_t \cdot \Big(I_1 + \sum_{k=1}^t (q_k^{\text{aff}} - d_k)\Big)\Big]. \tag{24}$$

Note that the objective function is seemingly intractable, even when the convex costs c_t and h_t take the piecewise affine form (19) considered in Ben-Tal et al. (2005b, 2009).

With this motivation in mind, we now recall Problem 3 stated in the Introduction, and note that it is exactly geared toward simplifying objectives of the form (24). In particular, if the inner expression in (24) depended *bi-affinely*[4] on the decision variables and the uncertain quantities, then standard techniques in robust optimization could be employed to derive tractable robust counterparts for the problem (see Ben-Tal et al. 2009 for a detailed overview). The following theorem summarizes our main result of this section, providing sufficient conditions that yield the desired outcome.

Theorem 3. *Consider an optimization problem of the form*

$$\max_{\mathbf{w} \in P} \Big[\mathbf{a}^T \mathbf{w} + \sum_{i \in \mathcal{I}} h_i(\mathbf{w})\Big],$$

having finite optimal value, where $P \subset \mathbb{R}^k$ is any polytope, $\mathbf{a} \in \mathbb{R}^n$ is an arbitrary vector, $\mathcal{I}$ is a finite index set, and h_i: $\mathbb{R}^n \to \mathbb{R}$ are functions satisfying the following properties:

Property 1 (P1). *h_i are concave extendable from* $\text{ext}(P)$, $\forall i \in \mathcal{I}$,

Property 2 (P2). $\text{conc}_P(h_i + h_j) = \text{conc}_P(h_i) + \text{conc}_P(h_j)$, *for any $i \neq j \in \mathcal{I}$.*

Then there exists a set of affine functions $z_i(\mathbf{w})$, $i \in \mathcal{I}$, satisfying $z_i(\mathbf{w}) \geqslant h_i(\mathbf{w})$, $\forall \mathbf{w} \in P$, $\forall i \in \mathcal{I}$, such that

$$\max_{\mathbf{w} \in P} \Big[\mathbf{a}^T \mathbf{w} + \sum_{i \in \mathcal{I}} z_i(\mathbf{w})\Big] = \max_{\mathbf{w} \in P} \Big[\mathbf{a}^T \mathbf{w} + \sum_{i \in \mathcal{I}} h_i(\mathbf{w})\Big].$$

Proof. The proof is slightly technical, so we relegate it to §EC.3 of the online appendix.

Let us discuss the statement of Theorem 3 and relevant examples of functions satisfying the conditions therein. (P1) requires the functions h_i to be concave extendable from $\text{ext}(P)$; by the discussion in §EC.2 of the online appendix, examples of such functions are any convex functions or, when $P = \mathcal{H}_n$, any component-wise convex functions. More generally, concave extendability can be tested using the sufficient condition provided in Lemma EC.2 of the online appendix.

A priori, condition (P2) seems more difficult to test. Note that, by Theorem EC.4 in the online appendix, it can be replaced with any of the following equivalent requirements.

Property 3 (P3). $\text{conc}_P(h_i) + \text{conc}_P(h_j)$ *is concave extendable from vertices, for any $i \neq j \in \mathcal{I}$.*

Property 4 (P4). *For any $i \neq j \in \mathcal{I}$, the linearity domains $\mathcal{R}_{h_i, P} = \{F_k\colon k \in \mathcal{K}\}$ and $\mathcal{R}_{h_j, P} = \{G_l\colon l \in \mathcal{L}\}$ of $\text{conc}_P(h_i)$ and $\text{conc}_P(h_j)$, respectively, are such that $F_k \cap G_l$ has all vertices in $\text{ext}(P)$, $\forall k \in \mathcal{K}$, $\forall l \in \mathcal{L}$.*

The choice of which condition to include should be motivated by what is easier to test in the application of interest. A particularly relevant class of functions satisfying both requirements (P1) and (P2) is the following.

Example 1. Let P be a polytope of the form (3). Then, any functions h_i that are convex and supermodular on $\text{ext}(P)$ satisfy the requirements (P1) and (P2).

The proof for this fact is the subject of Corollary EC.3 of the online appendix. An instance of this, which turns out to be particularly pertinent in the context of Problem 1, is $h_i(\mathbf{w}) = f_i(b_{i,0} + \mathbf{b}_i^T \mathbf{w})$, where $f_i\colon \mathbb{R} \to \mathbb{R}$ are convex functions, and $\mathbf{b}_i \geqslant \mathbf{0}$ or $\mathbf{b}_i \leqslant \mathbf{0}$. A further subclass of the latter is $P = \mathcal{H}_n$ and $\mathbf{b}_i = \mathbf{b} \geqslant 0, \forall\, i \in \mathcal{I}$, which was the object of a central result in Bertsimas et al. (2010) (§4.3 in that paper, and in particular Lemmas 4.8 and 4.9).

We remark that, whereas maximizing convex functions on polytopes is generally NP-hard (the max-cut problem is one such example (Pardalos and Rosen 1986)), maximizing *supermodular* functions on lattices can be done in polynomial time (Fujishige 2005). Therefore, our result does not seem to have direct computational complexity implications. However, as we show in later examples, it does have the merit of drastically simplifying particular computational procedures, particularly when combined with outer minimization problems such as those present in many robust optimization problems.

As another subclass of Example 1, we include the following.

Example 2. Let $P = \mathcal{H}_n$, and $h_i(\mathbf{w}) = \prod_{k\in\mathcal{K}_i} f_k(\mathbf{w})$, where $\mathcal{K}_i$ is a finite index set, and f_k are nonnegative, supermodular, and increasing (decreasing), for all $k \in \mathcal{K}_i$. Then h_i are convex and supermodular.

This result follows directly from Lemma 2.6.4 in Topkis (1998). One particular example in this class are all polynomials in $\mathbf{w}$ with nonnegative coefficients. In this sense, Theorem 3 is useful in deriving a simple (linear-programming based) algorithm for the following problem.

Corollary 1. *Consider a polynomial p of degree d in variables $\mathbf{w} \in \mathbb{R}^n$, such that any monomial of degree at least two has positive coefficients. Then, there is a linear programming formulation of size $\mathcal{O}(n^d)$ for solving the problem* $\max_{\mathbf{w}\in[0,1]^n} p(\mathbf{w})$.

Proof. Note first that the problem is nontrivial because of the presence of potentially negative affine terms. Representing p in the form $p(\mathbf{w}) = \mathbf{a}^T\mathbf{w} + \sum_{i\in\mathcal{I}} h_i(\mathbf{w})$, where each h_i has degree at least two, we can use the result in Theorem 3 to rewrite the problem equivalently as follows:

$$\begin{aligned}\max_{\mathbf{w}\in[0,1]^n} p(\mathbf{w}) = \min_{t,\{\mathbf{z}_i, z_{i,0}\}_{i\in\mathcal{I}}} \quad & t \\ \text{s.t.} \quad & t \geqslant \mathbf{a}^T\mathbf{w} + \sum_{i\in\mathcal{I}} (z_{i,0} + \mathbf{z}_i^T \mathbf{w}), \\ & \qquad \forall\, \mathbf{w} \in [0,1]^n, \quad (*) \\ & h_i(\mathbf{w}) \leqslant z_{i,0} + \mathbf{z}_i^T \mathbf{w}, \\ & \qquad \forall\, \mathbf{w} \in [0,1]^n. \quad (**)\end{aligned}$$

By Theorem 3, the semi-infinite LP on the right-hand side has the same optimal value as the problem on the left. Furthermore, standard techniques in robust optimization can be invoked to reformulate constraints $(*)$ in a tractable fashion (see Ben-Tal et al. 2009 for details), and constraints $(**)$ can be replaced by a finite enumeration over at most 2^d extreme points of the cube (since each monomial term h_i has degree at most d). Therefore, the semi-infinite LP can be rewritten as an LP of size $\mathcal{O}(n^d)$. □

4.1. Application to Problem 1

To exhibit how Theorem 3 can be used in practice, we again revisit Problem 1. More precisely, recall that one had to solve the seemingly intractable optimization problem in (24) in order to find the optimal affine orders q_t^{aff} for any fixed first-stage decisions $\mathbf{K}$, and this was the case even when all the problem costs were piecewise affine.

In this context, the following result is a direct application of Theorem 3.

Theorem 4. *Assume the costs c_t, h_t, and r are jointly convex and piecewise affine, with at most m pieces. Then, the optimal $\mathbf{K}$ and a set of worst-case optimal ordering policies $\{q_t^{\text{aff}}\}_{t\in\{1,\dots,T\}}$ can be computed by solving a single linear program with $\mathcal{O}(m\cdot T^2)$ variables and constraints when all decisions in $\mathbf{K}$ are continuous, or a mixed-integer linear program of the same size when some of the decisions in $\mathbf{K}$ are discrete.*

Proof. Consider first a fixed $\mathbf{K}$. The expression for the inner objective in (24) is

$$\sum_{t=1}^{T} \bigl[c_t(q_t^{\text{aff}}, \mathbf{K}) + h_t(I_{t+1}^{\text{aff}}, \mathbf{K})\bigr],$$

where $I_t^{\text{aff}}(\mathbf{d}_{[t]}) = I_1 + \sum_{k=1}^{t-1}(q_k^{\text{aff}} - d_k) \stackrel{\text{def}}{=} \mathbf{b}_t^T \mathbf{d}_{[t]} + b_{t,0}$ is the expression for the inventory under affine orders. The functions c_t and h_t are convex. Furthermore, by Lemma 4, there exist worst-case optimal affine rules $q_t^{\text{aff}}(\mathbf{d}_{[t]}) = \mathbf{q}_t \mathbf{d}_{[t]} + q_{t,0}$ such that

$$\mathbf{q}_t \geqslant 0, \quad \mathbf{b}_{t+1} \leqslant 0, \quad \forall\, t \in \{1, \dots, T\}.$$

Therefore, $c_t(q_t^{\text{aff}}(\mathbf{d}_{[t]}), \mathbf{K})$ and $h_t(I_{t+1}(\mathbf{d}_{[t+1]}), \mathbf{K})$, as functions of $\mathbf{d}_{[T+1]}$, are convex and *supermodular* on $\text{ext}(\mathcal{D}_{[T+1]})$, and fall directly in the realm of Theorem 3 (see Example 1).

In particular, an application of the latter result implies the existence of a set of *affine* ordering costs $c_t^{\text{aff}}(\mathbf{d}_{[t]}) = \mathbf{c}_t^T \mathbf{d}_{[t]} + c_{t,0}$ and *affine* inventory costs $z_t^{\text{aff}}(\mathbf{d}_{[t+1]}) = \mathbf{z}_t^T \mathbf{d}_{[t]} + z_{t,0}$ such that

$$\begin{aligned}&\max_{\mathbf{d}_{[T+1]}\in\mathcal{D}_{[T+1]}} \sum_{t=1}^{T} \bigl[c_t(q_t^{\text{aff}}, \mathbf{K}) + h_t(I_{t+1}^{\text{aff}}, \mathbf{K})\bigr] \\ &\quad = \max_{\mathbf{d}_{[T+1]}\in\mathcal{D}_{[T+1]}} \sum_{t=1}^{T} (c_t^{\text{aff}} + z_t^{\text{aff}})\end{aligned}$$

$$c_t^{\text{aff}}(\mathbf{d}_{[t+1]}) \geqslant c_t(q_k^{\text{aff}}, \mathbf{K}), \quad \forall \mathbf{d}_{[t]} \in \mathcal{D}_{[t]}, \quad \forall t \in \{1, \ldots, T\}, \quad (*)$$

$$z_t^{\text{aff}}(\mathbf{d}_{[t+1]}) \geqslant h_t(I_{t+1}^{\text{aff}}(\mathbf{d}_{[t+1]}), \mathbf{K}), \quad \forall \mathbf{d}_{[t+1]} \in \mathcal{D}_{[t+1]}, \quad \forall t \in \{1, \ldots, T\}. \quad (**)$$

With this transformation, the objective is a bi-affine function of the uncertainties $\mathbf{d}_{[T+1]}$ and the decision variables $\{\mathbf{c}_t, \mathbf{z}_t\}$. Furthermore, if the costs c_t and h_t are piecewise affine, the constraints $(*)$ and $(**)$ can also be written as bi-affine functions of the uncertainties and decisions. For instance, suppose

$$c_t(q, \mathbf{K}, \mathbf{p}) = \max_{j \in \mathcal{J}_t} \left\{ \boldsymbol{\alpha}_j^T (q, \mathbf{K}) + \beta_j \right\}, \quad \forall t \in \{1, \ldots, T\},$$

for suitably sized vectors $\boldsymbol{\alpha}_j$, $j \in \bigcup_t \mathcal{J}_t$. Then, $(*)$ are equivalent to

$$\mathbf{c}_t^T \mathbf{d}_{[t]} + c_{t,0} \geqslant \boldsymbol{\alpha}_j^T \big(\mathbf{q}_t^T \mathbf{d}_{[t]} + q_{t,0}, \mathbf{K}\big) + \beta_j, \quad \forall t \in \{1, \ldots, T\},$$

which are bi-affine in $\mathbf{d}_{[T+1]}$ and the vector of decision variables $\mathbf{x} \overset{\text{def}}{=} (\mathbf{K}, \mathbf{q}_t, q_{t,0}, \mathbf{c}_t, c_{t,0}, \mathbf{z}_t, z_{t,0})_{t \in \{1,\ldots,T\}}$. As such, the problem of finding the optimal capacity and order precommitments and the worst-case optimal policies can be written as a robust LP (see, e.g., Ben-Tal et al. 2005b, 2009), in which a typical constraint has the form

$$\lambda_0(\mathbf{x}) + \sum_{t=1}^{T} \lambda_t(\mathbf{x}) \cdot d_t \leqslant 0, \quad \forall \mathbf{d} \in \mathcal{D}_{[T+1]},$$

where $\lambda_i(\mathbf{x})$ are affine functions of the decision variables $\mathbf{x}$. It can be shown (see Ben-Tal et al. 2009 for details) that the previous semi-infinite constraint is equivalent to

$$\begin{cases} \lambda_0(\mathbf{x}) + \displaystyle\sum_{t=1}^{T} \left(\lambda_t(\mathbf{x}) \cdot \frac{\underline{d}_t + \bar{d}_t}{2} + \frac{\bar{d}_t - \underline{d}_t}{2} \cdot \xi_t \right) \leqslant 0, \\ -\xi_t \leqslant \lambda_t(\mathbf{x}) \leqslant \xi_t, \quad \forall t \in \{1, \ldots, T\} \end{cases} \quad (25)$$

which are linear constraints in the decision variables $\mathbf{x}$, $\boldsymbol{\xi}$. Therefore, the problem of finding the optimal parameters can be reformulated as an LP with $O(mT^2)$ variables and $O(mT^2)$ constraints, which can be solved very efficiently using commercially available software.

When some of the decisions in $\mathbf{K}$ are discrete, the reformulation above remains unchanged, and the LP becomes a mixed-integer linear program of the same size. □

5. Conclusions

In this paper, we strive to bridge two well-established paradigms for solving a particular class of dynamic robust problems, in which a large number of first-stage decisions must be made, which govern the constraints and cost structure of a simple (linear, one-dimensional) dynamical system. The first is dynamic programming—a methodology with very general scope, which allows insightful comparative statics analyses, but suffers from the curse of dimensionality, which limits its use in practice. The second involves the use of decision rules, i.e., policies parameterized in model uncertainties, which are typically obtained by restricting attention to particular functional forms and solving tractable convex optimization problems. The main downside of the latter approach is the lack of control over the degree of suboptimality of the resulting decisions.

We focus our analysis on the popular class of *affine* decision rules, and discuss sufficient conditions on the value functions of the dynamic program and the uncertainty sets, which ensure their optimality. We exemplify our findings in a class of applications concerning a retailer's strategic decisions and replenishment policies, where we show that all the optimal decisions can be found by solving a single linear (or mixed-integer) program of small size. From a theoretical standpoint, our results emphasize the interplay between the convexity and supermodularity of the value functions, and the lattice structure of the uncertainty sets, suggesting new modeling paradigms for dynamic robust optimization.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/opre.2013.1172.

Endnotes

1. We could also state these results in terms of $\mathcal{W}$ itself being a sublattice of $\mathcal{H}_n$. However, the distinction will turn out to be somewhat irrelevant, since the convexity of all the objectives will dictate that only the structure of the extreme points of $\mathcal{W}$ matters.
2. For a simplex, if $\mathcal{W}^\Gamma = \{\mathbf{w} \geqslant \mathbf{0}\colon \sum_{i=1}^n w_i \leqslant \Gamma\}$, then, with the change of variables $y_k \overset{\text{def}}{=} (\sum_{i=1}^k w_i)/\Gamma$, $\forall k \in \{1, \ldots, n\}$, the corresponding uncertainty set in the $\mathbf{y}$ variables is $\mathcal{W}_y = \{\mathbf{y} \in [0,1]^n\colon 0 \leqslant y_1 \leqslant y_2 \leqslant \cdots \leqslant y_n \leqslant 1\}$.
3. A function f is said to be *proper* if $f(\mathbf{x}) < +\infty$ for at least one $\mathbf{x}$, and $f(\mathbf{x}) > -\infty$ for every $\mathbf{x}$ (Rockafellar 1970).
4. That is, it would be affine in one set of variables when the other set is fixed.

Acknowledgments

The authors thank the associate editor and three anonymous referees for their insightful comments and suggestions, which considerably improved the quality of the paper. Part of this research was conducted while the first author was a Goldstine Fellow in the Department of Business Analytics and Mathematical Sciences of the IBM T. J. Watson Research Center, whose financial support is gratefully acknowledged. The third author's research was partially supported by Engineering and Physical Sciences Research Council (EPSRC) [Grant EP/J021814/1], FP7 Marie Curie Career Integration Grant, and the Royal Society Wolfson Merit Award.

References

Bemporad A, Borrelli F, Morari M (2003) Min–max control of constrained uncertain discrete-time linear systems. *IEEE Trans. Automatic Control* 48(9):1600–1606.

Ben-Tal A, Nemirovski A (1999) Robust solutions of uncertain linear programs. *Oper. Res. Lett.* 25:1–13.

Ben-Tal A, Nemirovski A (2002) Robust optimization—Methodology and applications. *Math. Programming* 92(3):453–480.

Ben-Tal A, Boyd S, Nemirovski A (2005a) Control of uncertainty-affected discrete time linear systems via convex programming. Working paper, Georgia Institute of Technology, Atlanta.

Ben-Tal A, Boyd S, Nemirovski A (2006) Extending scope of robust optimization: Comprehensive robust counterparts of uncertain problems. *Math. Programming* 107(1):63–89.

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

Ben-Tal A, Nemirovski A, Roos C (2002) Robust solutions of uncertain quadratic and conic-quadratic problems. *SIAM J. Optim.* 13(2): 535–560.

Ben-Tal A, Golany B, Nemirovski A, Vial J-P (2005b) Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing Service Oper. Management* 7(3):248–271.

Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Math. Programming* 99(2):351–376.

Bensoussan A, Crourty M, Proth JM (1983) *Mathematical Theory of Production Planning* (Elsevier Science Publishing Company, New York).

Bertsekas DP (2001) *Dynamic Programming and Optimal Control* (Athena Scientific, Belmont, MA).

Bertsekas DP, Rhodes IB (1971) On the minmax reachability of target tubes. *Automatica* 7:233–247.

Bertsimas D, Goyal V (2010) On the power and limitations of affine policies in two-stage adaptive optimization. *Math. Programming* 134(2):1–41.

Bertsimas D, Sim M (2003) Robust discrete optimization and network flows. *Math. Programming* 98(1–3):49–71.

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

Bertsimas D, Caramanis C, Brown DB (2011a) Theory and applications of robust optimization. *SIAM Rev.* 53(3):464–501.

Bertsimas D, Iancu DA, Parrilo PA (2010) Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.* 35(2):363–394.

Bertsimas D, Iancu DA, Parrilo PA (2011b) A hierarchy of near-optimal policies for multistage adaptive optimization. *IEEE Trans. Automatic Control* 56(12):2809–2824.

Bray RL, Mendelson H (2012) Deriving supply chain metrics from financial statements. *Management Sci.* 58(5):860–875.

Charnes A, Cooper WW, Symonds GH (1958) Cost horizons and certainty equivalents: An approach to stochastic programming of heating oil. *Management Sci.* 4(3):235–263.

Chen L, Lee HL (2009) Information sharing and order variability control under a generalized demand model. *Management Sci.* 55(5):781–797.

Chen X, Sim M, Sun P, Zhang J (2008) A linear decision-based approximation approach to stochastic programming. *Oper. Res.* 56(2): 344–357.

Cheridito P, Delbaen F, Kupper M (2006) Dynamic monetary risk measures for bounded discrete-time processes. *Electronic J. Probab.* 11:57–106.

Clark AJ, Scarf H (1960) Optimal policies for a multi-echelon inventory problem. *Management Sci.* 50(12):475–490.

Dullerud GE, Paganini F (2005) *A Course in Robust Control Theory* (Springer, New York).

El-Ghaoui L, Oustry F, Lebret H (1998) Robust solutions to uncertain semidefinite programs. *SIAM J. Optim.* 9(1):33–52.

Epstein LG, Schneider M (2003) Recursive multiple-priors. *J. Econom. Theory* 113(1):1–31.

Fan MKH, Tits A, Doyle J (1991) Robustness in the presence of mixed parametric uncertainty and unmodeled dynamics. *IEEE Trans. Automatic Control* 36(1):25–38.

Fujishige S (2005) *Submodular Functions and Optimization*, 2nd ed. Annals of Discrete Mathematics, Vol. 58 (Elsevier B.V., Amsterdam).

Garstka SJ, Wets RJB (1974) On decision rules in stochastic programming. *Math. Programming* 7(1):117–143.

Goh J, Sim M (2010) Distributionally robust optimization and its tractable approximations. *Oper. Res.* 58(4):902–917.

Gorissen BL, den Hertog D (2013) Robust counterparts of inequalities containing sums of maxima of linear functions. *Eur. J. Oper. Res.* 227(1):30–43.

Graves SC, Kletter DB, Hetzel WB (1998) A dynamic model for requirements planning with application to supply chain optimization. *Oper. Res.* 46(3):35–49.

Hansen LP, Sargent TJ (2001) Robust control and model uncertainty. *Amer. Econom. Rev.* 91(2):60–66.

Hansen LP, Sargent TJ (2008) *Robustness* (Princeton University Press, Princeton, NJ).

Hausman WH (1969) Sequential decision problems: A model to exploit existing forecasters. *Management Sci.* 16(2):93–111.

Heath DC, Jackson PL (1994) Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE Trans.* 26(3):17–30.

Heyman DP, Sobel MJ (1984) *Stochastic Models in Operations Research* (Dover Publications, Mineola, NY).

Huh W-T, Janakiraman G (2010) On the optimal policy structure in serial inventory systems with lost sales. *Oper. Res.* 58(2):486–491.

Kerrigan EC, Maciejowski JM (2003) On robust optimization and the optimal control of constrained linear systems with bounded state disturbances. *Proc.* 2003 *Eur. Control Conf., Cambridge, UK*.

Kerrigan EC, Maciejowski JM (2004) Properties of a new parameterization for the control of constrained systems with disturbances. *Proc. 2004 Amer. Control Conf.* 5:4669–4674.

Kuhn D, Wiesemann W, Georghiou A (2009) Primal and dual linear decision rules in stochastic and robust optimization. *Math. Programming* 130(1):1–33.

Löfberg J (2003) Approximations of closed-loop minimax MPC. *Proc. 42nd IEEE Conf. Decision and Control*, Vol. 2 (IEEE, Piscataway, NJ), 1438–1442.

Lovász L (1983) Submodular functions and convexity. *Mathematical Programming—The State of the Art* (Springer, Berlin), 235–257.

Pardalos PM, Rosen JB (1986) Methods for global concave minimization: A bibliographic survey. *SIAM Rev.* 28(3):367–379.

Porteus EL (2002) *Foundations of Stochastic Inventory Management* (Stanford Business Books, Stanford, CA).

Rockafellar T (1970) *Convex Analysis* (Princeton University Press, Princeton, NJ).

Scarf H (1959) The optimality of (s, S) policies in the dynamic inventory problem. Arrow KJ, Karlin S, Suppes P, eds. *Mathematical Methods in the Social Sciences* (Stanford University Press, Stanford, CA), 196–202.

Schrijver A (2003) *Combinatorial Optimization: Polyhedra and Efficiency*, 1st ed. (Springer-Verlag, Berlin).

See CT, Sim M (2010) Robust approximation to multiperiod inventory management. *Oper. Res.* 58(3):583–594.

Shapiro A, Nemirovski A (2005) On complexity of stochastic programming problems. Jeyakumar V, Rubinov A, eds. *Continuous Optimization: Current Trends and Modern Applications* (Springer, New York), 111–146.

Simchi-Levi D, Chen X, Bramel J (2004) *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics and Supply Chain Management*, 2nd ed. (Springer).

Skaf J, Boyd S (2010) Design of affine controllers via convex optimization. *IEEE Trans. Automatic Control* 55(11):2476–2487.

Talluri KT, van Ryzin GJ (2005) *The Theory and Practice of Revenue Management* (Springer Science and Business Media, Inc., New York).

Tardella F (2008) Existence and sum decomposition of vertex polyhedral convex envelopes. *Optim. Lett.* 2:363–375.

Tawarmalani M, Richard JP, Xiong C (2013) Explicit convex and concave envelopes through polyhedral subdivisions. *Math. Programming* 138(1–2):531–577.

Topkis DM (1998) *Supermodularity and Complementarity* (Princeton University Press, Princeton, NJ).

Veinott AF Jr (1966) The status of mathematical inventory theory. *Management Sci.* 12(11):745–777.

Zhou K, Doyle JC (1998) *Essentials of Robust Control* (Prentice Hall, Upper Saddle River, NJ).

Zipkin P (2000) *Foundations of Inventory Management* (McGraw-Hill, Boston).

Zipkin P (2008) On the structure of lost-sales inventory models. *Oper. Res.* 56(4):937–944.

Dan A. Iancu is an assistant professor of operations, information, and technology at the Stanford Graduate School of Business. His research interests include dynamic optimization under uncertainty and risk, with applications in inventory and supply chain management, revenue management, and problems at the interface of operations and finance.

Mayank Sharma is a research staff member at the IBM T. J. Watson Research Center in Yorktown Heights, New York. His research interests include the modeling and analysis of complex systems, and the design of efficient algorithms for such systems, with a focus on the areas of stochastic optimization and probabilistic inference.

Maxim Sviridenko is a professor in the Department of Computer Science at the University of Warwick, United Kingdom. His current research interests include mathematical aspects of operations research, theoretical computer science, optimization, and probability theory.