informs
http://pubsonline.informs.org/journal/opre/

OPERATIONS RESEARCH
Vol. 66, No. 2, March–April 2018, pp. 426–447
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Robust Inventory Management: An Optimal Control Approach

Michael R. Wagner[a]

[a] Michael G. Foster School of Business, University of Washington, Seattle, Washington 98195

Contact: mrwagner@uw.edu, http://orcid.org/0000-0003-2077-3564 (MRW)

Received: May 3, 2016
Revised: December 11, 2016; March 20, 2017; May 14, 2017
Accepted: July 27, 2017
Published Online in Articles in Advance: December 15, 2017

Subject Classifications: inventory/production: uncertainty: deterministic; dynamic programming: optimal control; probability stochastic model applications
Area of Review: Optimization

https://doi.org/10.1287/opre.2017.1669

Abstract. We formulate and solve static and dynamic models of inventory management that lie at the intersection of robust optimization and optimal control theory. Our objective is to minimize cumulative ordering, holding, and shortage costs over a horizon $[0,T]$, where the variable is a nonnegative ordering rate function $q(t) \in \mathscr{L}^2[0,T]$. The demand rate function $d(t)$ is unknown and is only assumed to belong to an uncertainty set $\Omega = \{d(t) \in \mathscr{L}^2[0,T]\colon \mu_a \leqslant (1/T)\int_0^T d(t)\,dt \leqslant \mu_b,\ a \leqslant d(t) \leqslant b,\ \forall\, t \in [0,T]\}$; this set is motivated by the strong law of large numbers for stochastic processes $\lim_{T\to\infty}(1/T)\int_0^T d(t)\,dt = \mu$, where μ is the mean drift. We analyze a static model, where the ordering rate function must be fully specified at time zero, and three dynamic variants, where re-optimizations are allowed during the planning horizon $[0,T]$ at prespecified review epochs. In the dynamic models, at review epoch $\tau \in [0,T]$, the past demand on $[0,\tau)$ is observable. In the first dynamic model, we ignore this information, and define a variant of Ω that is well formed for the remaining planning horizon $[\tau,T]$. In the second model, we define a variant of Ω for $[\tau,T]$ that utilizes the past demand information, though we make a simplifying technical assumption about the consistency of the demand on $[0,\tau)$ and Ω. In the third dynamic model, we remove this assumption, and we remedy the arising complications using the Hilbert Projection Theorem. In all cases we derive optimal *closed-form* ordering rate functions that equal either the bounds a or b, or weighted averages of these bounds $(sb+ha)/(s+h)$ or $(sa+hb)/(s+h)$, where s and h are the shortage and holding costs, respectively. The strategies differ by *when* these four ordering rates are applied, which is determined by an uncertainty-set-dependent partition of the remaining planning horizon. Computational experiments, focused on studying the dynamic variants, supplement the analytical results, and demonstrate that (1) the three variants exhibit comparable performance under well-behaved stochastic demand and (2) the third variant has a significant advantage when demand is seasonal, especially when the review frequency is appropriately selected. Finally, computational comparisons with the omniscient strategy $q(t) = d(t)$, for all $t \in [0,T]$, are encouraging.

Funding: The author gratefully acknowledges the support of a Neal and Jan Dempsey Faculty Fellowship.

Keywords: inventory management • robust optimization • optimal control • stochastic processes

1. Introduction

In recent years robust optimization has emerged as a popular and important approach to decision making under uncertainty, where the model of uncertainty is characterized by set membership rather than stochastic distributions. One of the main reasons for the success of the robust optimization paradigm is that it frequently results in a tractable model, in contrast to stochastic optimization formulations, which can suffer from the curse of dimensionality. In a few cases, closed-form solutions can be derived for robust models, providing additional intuitive understanding of the optimal robust decisions.

In this paper we study a new robust model of inventory management, and we derive new *closed-form* optimal robust solutions. While robust inventory management models have received substantial attention in academic research, our approach differs in a number of ways from existing models. First, while the literature on robust optimization focuses on finite-dimensional models, we propose and study an *infinite-dimensional* robust variant where the ordering variable is a function and the uncertainty set is a set of demand functions; consequently, our model is a natural robust analogue to many inventory models that utilize ideas from optimal control theory and represent demand as a stochastic process (e.g., Brownian motion, renewal process). Second, we structure our uncertainty sets using *strong laws of large numbers for stochastic processes* as motivation, which is in contrast to most of the robust optimization literature that defines uncertainty sets from a structural viewpoint (interval, polyhedral, ellipsoidal, etc.); furthermore, while some recent papers have designed their uncertainty sets using the limit theorems of probability as motivation, they focus on the finite-dimensional central limit theorem (CLT), rather than strong laws of large numbers for stochastic processes. Third, in contrast to most of the literature, we derive *closed-form robust ordering rate functions* for a basic static model and dynamic variants of it, where the ordering rate function can depend on the currently observed inventory position function. Our model overlaps with existing models in the following ways: we consider the control of a single durable product over a finite horizon, with no fixed ordering costs, zero lead times, and backordering allowed; we assume demand is a stochastic process, where the mean (e.g., drift parameter of Brownian motion or mean interarrival time in a renewal process) is known, but any other parameters and relevant distributions (e.g., distribution of interarrival time) are not known. We next provide a literature review to describe more details and better position our contributions.

1.1. Literature Review

The field of robust optimization has burgeoned in recent years and we do not attempt to provide a comprehensive literature review. We point the interested reader to Ben-Tal et al. (2009) for an overview. A similar statement can be made about inventory management, and we suggest (Zipkin 2000) as a primer on this vast field. We focus our review on the papers most related to our work, which concentrates on three streams: (1) robust inventory management, (2) optimal control of stochastic inventory systems, and (3) robust uncertainty sets motivated by the limit theorems of probability.

The foundation of our paper can be found in Bertsimas and Sim (2004), which introduce the notion of a "budget of uncertainty" to reduce the conservatism of robust optimization. Bertsimas and Thiele (2006) apply the ideas in Bertsimas and Sim (2004) to formulate a robust optimization model of inventory control, which can handle fixed costs, capacitated orders and inventory, and network topologies. Bienstock and Özbay (2008) generalize (Bertsimas and Thiele 2006) in multiple directions and also analyze data-driven robust models, focusing on algorithmic issues. Mamani et al. (2016) study a similar inventory problem to that in Bertsimas and Thiele (2006) and Bienstock and Özbay (2008), except that the uncertainty sets are motivated by the CLT, and closed-form solutions, in both static and dynamic contexts, are derived. Other researchers have studied robust inventory management from different perspectives. Chen et al. (2007) study generic robust uncertainty sets allowing for asymmetry, resulting in a second-order cone counterpart and See and Sim (2010) analyze a "factor-based" model of uncertainty, which also results in a second-order cone program. Wagner (2010, 2011) study robust inventory management from the online optimization perspective. More recently, Ardestani-Jaafari and Delage (2016), building upon the work of Gorissen and Hertog (2013), provide approximation approaches for a broader class of robust optimization problems involving sums of piecewise linear functions. Solyalı et al. (2016) propose a new robust formulation of inventory control based on ideas from facility location, which results in polynomial-time solvability when the initial inventory is negative or zero. Our paper differs from this finite-dimensional literature stream in that we study an infinite-dimensional control-theoretic formulation, where the ordering strategy and demand stream are represented by (Lebesgue square-integrable) functions.

All the aforementioned papers deal exclusively with finite dimensional models. Many researchers have instead modeled stochastic inventory management as (nonrobust) stochastic optimal control problems, typically using either (1) Brownian motion or (2) renewal processes to model the evolution of demand. Harrison et al. (1983) and Harrison and Taksar (1983) are some of the first to model storage systems using Brownian motion, and they derive optimal control policies for such systems. Gallego (1990) utilizes Brownian motion to represent cumulative demand in an inventory management system. Jagannathan and Sen (1991) also utilize Brownian motion in an inventory model of blood in a blood bank context. Wein (1992) approximates a make-to-stock production system as a dynamic control problem involving Brownian motion. Asmussen and Perry (1998) focus their study on an operator calculus for matrix-exponential distributions, but also apply their ideas to a (q, Q) inventory system driven by Brownian motion. Liu and Song (2012) study the (S, T) inventory policy under a variety of demand assumptions, including Brownian motion. Wu and Chao (2014) utilize two-dimensional Brownian motion to simultaneously capture correlated cumulative demand and production. Regarding renewal processes, Gallego and van Ryzin (1994) utilize a price-dependent Poisson process to model demand. Rosling (2002) studies different cost-rate models under a renewal process of demand. Plambeck and Ward (2006) and Plambeck (2008) utilize renewal processes to model cumulative demand in assemble-to-order systems. Federgruen and Wang (2015) also utilize a renewal process of demand in an inventory system with shelf-age dependent holding cost and delay-dependent shortage cost. While we do not explicitly utilize Brownian motion or renewal processes in our paper, these stochastic processes motivate the infinite-dimensional nature of our robust model. Furthermore, our work can be considered a robust analogue to these stochastic control-theoretic approaches to inventory management.

We next discuss the design of robust uncertainty sets. Historically, these sets have been designed as either interval-based, polyhedral, ellipsoidal, or, more generally, simply convex. However, recently researchers have attempted to design uncertainty sets that capture salient characteristics of certain limit theorems of probability. The first example of such a work is Bertsimas et al. (2011), which analyzes queuing networks with a robust uncertainty set motivated by the probabilistic law of the iterated logarithm. Next, Bandi and Bertsimas (2012) provide an in-depth study of the use of CLT-style uncertainty sets; these sets have been applied to information theory by Bandi and Bertsimas (2011), option pricing by Bandi and Bertsimas (2014b), auction design by Bandi and Bertsimas (2014a), and queueing theory by Bandi et al. (2015, 2016). More relevant to our paper is Mamani et al. (2016), who studied robust inventory management and whose design of uncertainty sets was also motivated by the CLT. Our paper differs from this literature stream in that our uncertainty sets are motivated by strong laws of large numbers for stochastic processes, rather than the finite-dimensional CLT or the law of the iterated logarithm.

In our paper we consider both static models, where the entire ordering rate function must be specified at time zero, and dynamic models, where the ordering rate function can be updated as uncertain parameters are realized. We therefore also survey the recent work on robust dynamic optimization. Ben-Tal et al. (2004) introduce an "adjustable" robust optimization problem, where some variables can be updated based on realizations of uncertain parameters; these authors prove that the general problem is NP-hard. This difficulty has motivated researchers to study approximations, usually an "affinely adjustable" robust optimization model, where the optimization is over the class of linear policies in the uncertain parameters. Ben-Tal et al. (2005) study such a model in a supply chain context. Chen et al. (2008) utilize second-order cones to improve upon the linear approximations in a generic multiperiod problem. Georghiou et al. (2015) and Bertsimas and Georghiou (2015) study alternative approximations to the general adjustable robust optimization problem. Bertsimas et al. (2010) prove the optimality of affine policies for a general class of multistage robust optimization models where random disturbances are constrained to lie in intervals and are independent. Iancu et al. (2013) continue the study of affine policies, more fully characterizing the problem structures where affine policies are optimal in dynamic robust optimization. Mamani et al. (2016) study a rolling-horizon variant of their basic static model, and show that it exhibits encouraging computational performance with respect to Bertsimas and Thiele (2006) and Bertsimas et al. (2010). Solyalı et al. (2016) also propose a rolling-horizon variant of their static model whose computational performance also compares well with Bertsimas and Thiele (2006), Ben-Tal et al. (2004), See and Sim (2010), and others. In our paper we adopt this rolling-horizon approach to design dynamic strategies, though we define and analyze three separate variants, which depend on whether or not the observed demand stream is consistent with the original robust uncertainty set; this is a consideration that has not been studied in the literature.

1.2. Contributions

Our paper provides a number of contributions to the operations research literature:

1. We are the first to formulate and solve an inventory management problem at the intersection of robust optimization and optimal control theory, where model primitives are functions, not vectors. We analyze a basic static model and three dynamic variants based on a rolling-horizon framework:

    (a) In the first dynamic variant, at a review epoch we define a version of the static uncertainty set that does not depend on the observed demand, and we solve the problem for the remaining horizon.

    (b) In the second dynamic variant, at a review epoch we define an uncertainty set that does depend on the observed demand, and we make a simplifying assumption to ensure the consistency of the observed demand and the structure of original static uncertainty set.

    (c) In the third dynamic variant, we remove the simplifying assumption of consistency, and we resolve the arising complications by applying the Hilbert Projection Theorem via a novel decomposition of projections. If the projected demand stream is not too far from the observed demand stream, we construct an uncertainty set based on the projection; however, if the projection is far, we recommend reparameterizing the uncertainty set based on information learned from the observed demand. To the best of our knowledge, the issue of consistency has not been studied in previous rolling-horizon dynamic robust models (where authors implicitly assume observed demand is consistent, as in our second dynamic variant). We believe this is an important issue to study since inconsistent demand renders a rolling-horizon approach infeasible.

2. We motivate the design of our robust uncertainty sets using strong laws of large numbers for stochastic processes. This design choice continues the recent trend of utilizing the limit theorems of probability to design uncertainty sets. However, we are the first to consider (1) the limit theorems of stochastic processes as well as (2) strong laws of large numbers (the literature has focused on distributional limit theorems, such as the CLT).

3. We derive optimal *closed-form* ordering rate functions for the static problem and all three dynamic variants. All of these optimal strategies order at only four different rates: the lower and upper bounds of the demand rate a and b, and weighted averages of these bounds $(sb + ha)/(s + h)$ and $(sa + hb)/(s + h)$, where s

and h are shortage and holding costs. There is also the possibility of an impulse order to satisfy a backlog in the dynamic cases. The ordering strategies differ by *when* each order rate is applied, which is determined by a partition of the planning horizon:

(a) In the static variant, the planning horizon $[0, T]$ is partitioned into two–three subintervals, based on the value of $\mu_a + \mu_b - (a + b)$, where μ_a and μ_b are lower and upper bounds, respectively, on the average demand rate μ.

(b) In the first dynamic variant, at a review epoch τ, the planning horizon $[\tau, T]$ is again partitioned based on the value of $\mu_a + \mu_b - (a + b)$.

(c) In the second dynamic variant, the planning horizon $[\tau, T]$ is partitioned based on the observed demand and the values of μ_a, μ_b, a, and b.

(d) In the third dynamic variant, the partition of $[\tau, T]$ depends on a projection of the observed demand onto an appropriately defined set, and the values of μ_a, μ_b, a, and b.

4. We complement our theoretical analyses with computational experiments. We first focus on comparing our first and third dynamic strategies (the second is not always feasible) when demand is either (1) a modification of Brownian motion or (2) stochastic with positive trend and seasonality. For the former demand stream, the two dynamic variants exhibit comparable performance, whereas for the latter stream with seasonality, our third dynamic variant has a substantial advantage due to its ability to react to the seasonality. We then study the impact of the review frequency and show that for most scenarios, increasing the frequency reduces the cost, *except* for the third dynamic variant under seasonal demand; in this case, the review frequency should be selected carefully, with respect to the period of the seasonality, as this results in the best performance for the seasonal demand. Finally, we include computational comparisons between our robust strategies and the omniscient strategy $q(t) = d(t)$, for all $t \in [0, T]$, which are encouraging.

While we focus on an inventory management context, we are optimistic that similar approaches, combining robust optimization with optimal control theory and the limit theorems of probability, can be fruitful for other problem contexts.

1.3. Paper Outline

In Section 2, we utilize strong laws of large numbers for stochastic processes to motivate our robust uncertainty sets. In Section 3 we analyze a basic static model, where the robust ordering strategy must be fully characterized at time zero. In Section 4 we introduce a framework for analyzing dynamic rolling-horizon variants of the static problem; in particular, we consider three different *sequences* of uncertainty sets that can be utilized in a dynamic setting. The first sequence defines a new uncertainty set at each review epoch that does not depend on the past demand realization; the dynamic problem for these uncertainty sets is analyzed in Section 4.2. The second sequence of uncertainty sets allows a dependence on the past realization of demand, but assumes a simplifying notion of consistency; this dynamic problem is analyzed in Section 4.3.1. In Section 4.3.2, we remove the assumption of consistency and utilize the Hilbert Projection Theorem to reconcile an inconsistent demand realization with our robust uncertainty sets. Computational experiments, which compare our dynamic strategies under different generators of demand and study the impact of review frequency, are discussed in Section 5. Concluding thoughts are given in Section 6. Certain proofs are presented in the main text, to provide additional insight, but most are provided in the appendix.

2. Robust Uncertainty Set Motivated by Strong Laws for Stochastic Processes

In this paper we consider an infinite-dimensional model indexed by time $t \in [0, T]$, where the cumulative demand up to time t is represented by a function $D(t)$. In this section, we discuss two common stochastic-process models of $D(t)$ used in the literature, Brownian motion and renewal processes, that motivate the definition of our robust uncertainty set. We first discuss Brownian motion, define a corresponding robust uncertainty set, and then argue that the set is also reasonable for representing a renewal process model of demand.

2.1. Brownian Motion Model of Cumulative Demand

Many researchers, as outlined in the introduction, have modeled $D(t)$ as Brownian motion with drift μ and instantaneous volatility σ:

$$D(t) = \mu t + \sigma B(t),$$

where $B(t)$ is a standard Brownian motion process (i.e., $B(t)$ is a normal random variable with zero mean and variance equal to t). While the strong law of large numbers is perhaps best known in the finite-dimensional case, it also applies for Brownian motion, which states that

$$\lim_{t\to\infty} \frac{D(t)}{t} = \mu, \quad \text{almost surely.} \tag{1}$$

However, there are a number of drawbacks to using Brownian motion to model *cumulative* demand: (1) the Brownian motion process can decrease and (2) Brownian motion is nowhere differentiable. In practice, cumulative demand cannot decrease, and the demand rate is typically smoother (especially for high volume products). Therefore, we introduce the instantaneous demand rate

$$d(t) = \frac{dD(t)}{dt}, \tag{2}$$

which we assume exists for all $t \in [0, T]$, is nonnegative, and is bounded from above; in particular, we assume that there exists values $0 \leqslant a \leqslant b < \infty$ such that

$$a \leqslant d(t) \leqslant b, \quad \forall\, t \in [0, T]. \tag{3}$$

Thus, since $a \geqslant 0$, $D(t)$ can only grow, which approximates reality better than Brownian motion, and the existence/boundedness of the derivative introduces a type of smoothness into the model, which is also (arguably) a better model of reality than nondifferentiable Brownian motion. These observations lead us to define a robust uncertainty set for the demand rate function $d(t)$ that (approximately) obeys the strong law in Equation (1), for a finite horizon, and the smoothness constraints in Expression (3).

We now formally define our function-space uncertainty set. To formulate a rigorous model, we begin with the measure space $([0, T], \mathscr{B}([0, T]), m)$, where $\mathscr{B}([0, T])$ is the Borel σ-algebra on $[0, T]$ and m is the standard Lebesgue measure. On this measure space, we focus on the function space $\mathscr{L}^2[0, T]$ of Lebesgue square-integrable functions on $[0, T]$; i.e., $\int_{[0,T]} |f|^2\, dm < \infty$ for all $f \in \mathscr{L}^2[0, T]$. Note that $\mathscr{L}^2[0, T]$ is a complete metric space with inner product $\langle f, g\rangle = \int_{[0,T]} f g\, dm$ and induced norm $\|f\| = \sqrt{\int_{[0,T]} f^2\, dm}$; in other words, $\mathscr{L}^2[0, T]$ is a Hilbert space. Since our ground set is the interval $[0, T]$ on the real line, we utilize the notation $\int_0^T f(t)\, dt$ for $\int_{[0,T]} f\, dm$. We define our robust uncertainty set in terms of the demand rate function $d(t)$ as

$$\Omega = \left\{ d \in \mathscr{L}^2[0, T]\colon \mu_a \leqslant \frac{1}{T} \int_0^T d(t)\, dt \leqslant \mu_b, \right.$$
$$\left. a \leqslant d(t) \leqslant b,\, \forall\, t \in [0, T] \right\}, \tag{4}$$

where we assume $a \leqslant \mu_a \leqslant \mu \leqslant \mu_b \leqslant b$. The set Ω can be interpreted as representing most smooth approximations of nondecreasing Brownian motion sample paths with drift μ over the horizon $[0, T]$, where $(1/T)\int_0^T d(t)\, dt \in [\mu_a, \mu_b]$ is an approximation of the strong law in Equation (1). Furthermore, given the smoothness constraints, this uncertainty set is arguably a better representation of real uncertainty than Brownian motion. We shall also see that a robust model built upon Ω is more tractable, leading to closed-form solutions, which are uncommon in models using Brownian motion.

Note that the range $[\mu_a, \mu_b]$ is an uncertainty set within another uncertainty set (Ω), and our model is implementable without knowing the exact value of μ. Furthermore, even if μ was precisely known, strong laws of large numbers do not necessarily hold under finite durations T. Loosely speaking, if T is large, we can set μ_a and μ_b close to μ; in contrast, if T is small, the interval $[\mu_a, \mu_b]$ must be larger, to allow inevitable deviations from the limit. A straightforward parameterization, to control the degree of conservatism, is also possible: $\mu_a = \mu - \Gamma$ and $\mu_b = \mu + \Gamma$. In Section 3.2, we discuss a regression-based numerical study that provides guidelines for setting the (μ_a, μ_b, a, b) parameters, as a function of the problem's economics, to improve the performance of the robust ordering strategy. Finally, if the horizon T is large enough, and convergence is assumed to (approximately) hold, one can set $\mu_a = \mu_b = \mu$, and still obtain closed-form ordering functions.

2.2. Renewal Process Model of Cumulative Demand

Researchers have also modeled $D(t)$ as a general renewal process, with arbitrary interarrival time distributions. If we denote the mean interarrival time to be $1/\mu$, then the strong law of large numbers for renewal processes is equivalent to Equation (1). The constraints $(1/T)\int_0^T d(t)\, dt \in [\mu_a, \mu_b]$ in the definition of Ω in Equation (4) can alternatively be viewed as an approximation of the strong law for renewal processes.

The smoothness constraints in Equation (3) are slightly more problematic, as a renewal process is essentially a jump process, with a derivative equal to either 0 or infinity. However, we do not aim to precisely capture all characteristics of renewal process samples paths in Ω; rather, we argue that in certain situations, Ω can serve as a good approximation of these sample paths. For instance, if the interarrival time is (stochastically) small with respect to T, then the sample paths of a renewal process can be well approximated by smooth functions.

2.3. Parameterizations

We conclude this section by discussing the values of μ, a, and b. The drift parameter μ could be determined by calculating the average demand rate from historical data. The minimum and maximum rates could similarly be estimated as the extreme values from historical data over some time frame (assuming outliers are appropriately removed); alternatively, a and b could be set at percentiles of historical data, say the 5th and 95th. Also, in Section 3.2, we discuss a more comprehensive approach for jointly selecting the (μ_a, μ_b, a, b) parameters.

In summary, we believe that the uncertainty set Ω in Equation (4) well represents the set of sample paths of Brownian motion that would actually occur in practice. The set Ω can also approximately represent the set of sample paths of general renewal processes. The benefit of using a robust model built upon Ω is the greater tractability and attainable closed-form solutions that are rarely possible in stochastic optimal control. Furthermore, the distribution of the interarrival times is not needed in a robust optimization model.

3. Static Robust Inventory Management

In this section we study a static model, where the ordering rate function $q(t)$ for the entire horizon $[0,T]$ must be specified at time $t=0$. While many inventory applications in practice are dynamic, in the sense that decisions can be adjusted with new information, there is merit in studying static models. First, contractual obligations can require that all orders be specified in advanced; Ardestani-Jaafari and Delage (2016) (see Remark 7 in their section 6) and Mamani et al. (2016) (see the introduction to their section 3) discuss such cases. Furthermore, since robust dynamic models are difficult in general (Ben-Tal et al. 2004 prove that a generic adjustable dynamic robust model is NP-hard), static models are frequently used as subroutines in the design of rolling-horizon dynamic optimization models, as in Mamani et al. (2016) (see their section 4) and Solyalı et al. (2016) (see their section 4); we adopt a similar approach in later sections. The static model considered in this section, for an arbitrary demand uncertainty set Υ, is

$$\begin{aligned}\min\ & \int_0^T (cq(x)+y(x))\,dx\\ \text{s.t. }& I(x)=\int_0^x (q(t)-d(t))\,dt, \quad \forall\, x\in[0,T]\\ & y(x)\geqslant hI(x), \quad \forall\, d\in\Upsilon, \forall\, x\in[0,T]\\ & y(x)\geqslant -sI(x), \quad \forall\, d\in\Upsilon, \forall\, x\in[0,T]\\ & q(x)\geqslant 0, \quad \forall\, x\in[0,T],\end{aligned} \tag{5}$$

where $q\in\mathscr{L}^2[0,T]$ is the ordering rate function, I is the inventory position function, and $y\in\mathscr{L}^2[0,T]$ captures the mismatch cost $\max\{hI(x),-sI(x)\}$, for all $x\in[0,T]$; c is the purchasing cost rate, h is the holding cost rate, and s is the stockout cost rate. We define uncertainty per constraint, which is the more common approach in the literature (e.g., Bertsimas and Sim 2004, Bertsimas and Thiele 2006, Bienstock and Özbay 2008, Ben-Tal et al. 2004, Bertsimas et al. 2010, Mamani et al. 2016, Solyalı et al. 2016, etc.). However, other researchers have focused on determining a single worst-case demand instance per model, as in Gorissen and Hertog (2013) and Ardestani-Jaafari and Delage (2016).

Denote the minimum and maximum cumulative demands on $[0,x]$, for all $x\in[0,T]$, as

$$\underline{D}(x)=\min_{d\in\Upsilon}\int_0^x d(t)\,dt \quad\text{and}\quad \bar{D}(x)=\max_{d\in\Upsilon}\int_0^x d(t)\,dt, \tag{6}$$

respectively. Note that both $\underline{D}(x)$ and $\bar{D}(x)$ are nondecreasing in x. An appealing characteristic of Formulation (5) is that we are able to find an optimality condition that balances the costs associated with the worst-case cumulative demands $\underline{D}(x)$ and $\bar{D}(x)$, for all $x\in[0,T]$, as shown in the following lemma.

Lemma 1. *If $c\leqslant sT$, then the optimal robust ordering rate function q^* satisfies*

1. $\int_0^x q^*(t)\,dt=(s\bar{D}(x)+h\underline{D}(x))/(s+h)$, *for* $x\in[0,T-c/s]$.
2. $q^*(x)=0$, *for* $x\in(T-c/s,T]$.

If $c>sT$, $q^*(x)=0$ for all $x\in[0,T]$.

This lemma is similar in structure to Lemma 2 in Mamani et al. (2016), though, since Formulation (5) is a *continuous* linear program, our result is a condition on a function, rather than a vector, and the proof technique is consequently different as well. In our proof, we guess the solution and prove it is correct using the weak duality of continuous linear programming (as strong duality does not necessarily hold). Furthermore, this lemma is the main driving force of one of our main results, Theorem 1, and after its presentation, we provide an intuitive interpretation of it and Lemma 1 in terms of the newsvendor model for the case where $\Upsilon=\Omega$.

Proof of Lemma 1. We rearrange Formulation (5) with embedded optimization problems:

$$\begin{aligned}\min\ & \int_0^T (cq(x)+y(x))\,dx\\ \text{s.t. }& \frac{y(x)}{h}-\int_0^x q(t)\,dt\geqslant \max_{d\in\Upsilon}\left\{-\int_0^x d(t)\,dt\right\}, \quad \forall\, x\in[0,T]\\ & \frac{y(x)}{s}+\int_0^x q(t)\,dt\geqslant \max_{d\in\Upsilon}\int_0^x d(t)\,dt, \quad \forall\, x\in[0,T]\\ & q(x)\geqslant 0, \quad \forall\, x\in[0,T].\end{aligned} \tag{7}$$

Using Expressions (6), we obtain a continuous linear program with a new auxiliary variable $z\in\mathscr{L}^2[0,T]$:

$$\begin{aligned}\min\ & \int_0^T z(x)\,dx\\ \text{s.t. }& z(x)\geqslant cq(x)+h\int_0^x q(t)\,dt-h\underline{D}(x), \quad \forall\, x\in[0,T]\\ & z(x)\geqslant cq(x)-s\int_0^x q(t)\,dt+s\bar{D}(x), \quad \forall\, x\in[0,T]\\ & q(x)\geqslant 0, \quad \forall\, x\in[0,T].\end{aligned} \tag{8}$$

We first observe that, at optimality,

$$z(x)=\max\left\{cq(x)+h\int_0^x q(t)\,dt-h\underline{D}(x),\ cq(x)-s\int_0^x q(t)\,dt+s\bar{D}(x)\right\}, \quad \forall\, x\in[0,T]. \tag{9}$$

We next assume $c \leqslant sT$ and define $A = T - c/s \in [0,T]$; the case where $c > sT$ is handled at the end of the proof. Consider a feasible solution to Formulation (8) characterized by

$$\int_0^x q(t)\,dt = \frac{s\bar{D}(x) + h\underline{D}(x)}{s+h} \tag{10}$$

for $x \in [0,A]$ and $q(x) = 0$ for $x \in (A,T]$. For $x \in [0,A]$, due to Equation (10), the two arguments of the max operator in Equation (9) are equal. However, for $x \in (A,T]$, the second argument of the max operator in Equation (9) dominates the first; to see this, we notice that

$$\begin{aligned} &cq(x) + h\int_0^x q(t)\,dt - h\underline{D}(x)\\ &\quad\leqslant cq(x) - s\int_0^x q(t)\,dt + s\bar{D}(x)\\ &\Leftrightarrow \int_0^x q(t)\,dt \leqslant \frac{s\bar{D}(x) + h\underline{D}(x)}{s+h}\\ &\Leftrightarrow \int_0^A q(t)\,dt \leqslant \frac{s\bar{D}(x) + h\underline{D}(x)}{s+h}\\ &\Leftrightarrow \frac{s\bar{D}(A) + h\underline{D}(A)}{s+h} \leqslant \frac{s\bar{D}(x) + h\underline{D}(x)}{s+h}, \end{aligned}$$

where the final inequality is due to $(s\bar{D}(x) + h\underline{D}(x))/(s+h)$ being nondecreasing in x.

We therefore assign the auxiliary variable z the value

$$z(x) = \begin{cases} cq(x) + h\left(\dfrac{s\bar{D}(x) + h\underline{D}(x)}{s+h}\right) - h\underline{D}(x), & x \in [0,A]\\ s\bar{D}(x) - s\left(\dfrac{s\bar{D}(A) + h\underline{D}(A)}{s+h}\right), & x \in (A,T], \end{cases} \tag{11}$$

which results in an objective function value of

$$\begin{aligned} &c\left(\frac{s\bar{D}(A) + h\underline{D}(A)}{s+h}\right) + h\int_0^A \left(\frac{s\bar{D}(x) + h\underline{D}(x)}{s+h} - \underline{D}(x)\right)dx\\ &\quad + s\int_A^T \left(\bar{D}(x) - \frac{s\bar{D}(A) + h\underline{D}(A)}{s+h}\right)dx\\ &= h\int_0^A \left(\frac{s\bar{D}(x) + h\underline{D}(x)}{s+h} - \underline{D}(x)\right)dx\\ &\quad + s\int_A^T \bar{D}(x)\,dx. \end{aligned} \tag{12}$$

The dual of Problem (8) is

$$\begin{aligned} \max\ & \int_0^T (s\bar{D}(x)v(x) - h\underline{D}(x)u(x))\,dx,\\ \text{s.t. } & u(x) + v(x) = 1, \quad \forall\, x \in [0,T],\\ & -cu(x) - h\int_x^T u(t)\,dt - cv(x)\\ & \quad + s\int_x^T v(t)\,dt \leqslant 0, \quad \forall\, x \in [0,T],\\ & u(x), v(x) \geqslant 0, \quad \forall\, x \in [0,T], \end{aligned}$$

which, by substituting out the u variable, can be written as

$$\begin{aligned} \max\ & \int_0^T (s\bar{D}(x) + h\underline{D}(x))v(x)\,dx - h\int_0^T \underline{D}(x)\,dx\\ \text{s.t. } & \int_x^T v(t)\,dt \leqslant \frac{c + h(T-x)}{s+h}, \quad \forall\, x \in [0,T]\\ & 0 \leqslant v(x) \leqslant 1, \quad \forall\, x \in [0,T]. \end{aligned}$$

Next, consider the dual feasible solution

$$v(t) = \begin{cases} \dfrac{h}{s+h}, & x \in [0,A]\\ 1, & x \in (A,T], \end{cases}$$

which has a dual objective value of

$$\begin{aligned} &\frac{h}{s+h}\int_0^A (s\bar{D}(x) + h\underline{D}(x))\,dx + \int_A^T (s\bar{D}(x) + h\underline{D}(x))\,dx\\ &\quad - h\int_0^A \underline{D}(x)\,dx - h\int_A^T \underline{D}(x)\,dx\\ &= h\int_0^A \left(\frac{s\bar{D}(x) + h\underline{D}(x)}{s+h} - \underline{D}(x)\right)dx + s\int_A^T \bar{D}(x)\,dx. \end{aligned} \tag{13}$$

Therefore, the primal value (12) and dual value (13) are equal, and the solutions are optimal. We conclude that $\int_0^x q(t)\,dt = (s\bar{D}(x) + h\underline{D}(x))/(s+h)$ for $x \in [0,A]$ and $q(x) = 0$ for $x \in (A,T]$ is an optimal robust ordering solution. If $c > sT$, $A < 0$, and a similar argument shows that $q^*(x) = 0$ for all $x \in [0,T]$. □

Remark 1. Note that Lemma 1 can accommodate continuous discounting with no change in the final result. More precisely, if the objective in Formulation (5) is replaced with $\int_0^T e^{-\delta x}(cq(x) + y(x))\,dx$, for some $\delta > 0$, the results do not change and do not depend on δ (which can be seen by tracing the proof's logic). Therefore, for ease of exposition, we do not discuss discounting in our paper.

Lemma 1 does not depend on the structure of Υ. However, letting $\Upsilon = \Omega$ and utilizing the structure of Ω leads to simple intuitive ordering strategies, which we discuss next. Our first theorem leverages Lemma 1 to exactly characterize the optimal robust ordering rate function q^* when we assign $\Upsilon = \Omega$. A representative optimal ordering strategy from this theorem (for the $\mu_a + \mu_b < a + b$ case) is illustrated graphically in Figure 1.

Theorem 1. *Letting $\Upsilon = \Omega$, if $c \leqslant sT$, then the optimal robust ordering rate function $q^*(x) = 0$ for $x \in (T - c/s, T]$ and the following strategy for $x \in [0, T - c/s]$:*

If $\mu_a + \mu_b > a + b$:

$$\begin{cases} (sb + ha)/(s+h), & x \in [0, (b-\mu_a)T/(b-a)],\\ b, & x \in (b-\mu_a)T/(b-a), (\mu_b - a)T/(b-a)],\\ (sa + hb)/(s+h), & x \in ((\mu_b - a)T/(b-a), T]; \end{cases}$$

Figure 1. (Color online) Left: The Optimal Ordering Strategy $q^*(t)$ for the $\mu_a + \mu_b < a + b$ Case of Theorem 1, Where $c = 20$, $s = 2$, $h = 1$, $T = 200$, $\mu_a = 8$, $\mu_b = 12$, $\mu = 10$, $\sigma = 3$, $a = 6$ and $b = 15$. Right: The Cumulative Orders of Theorem 1 Tracking Cumulative Demand, Which Is Generated Using Truncated Normal Random Variables with Mean μ and Standard Deviation σ as the Random (Nonnegative) Instantaneous Growth of Demand

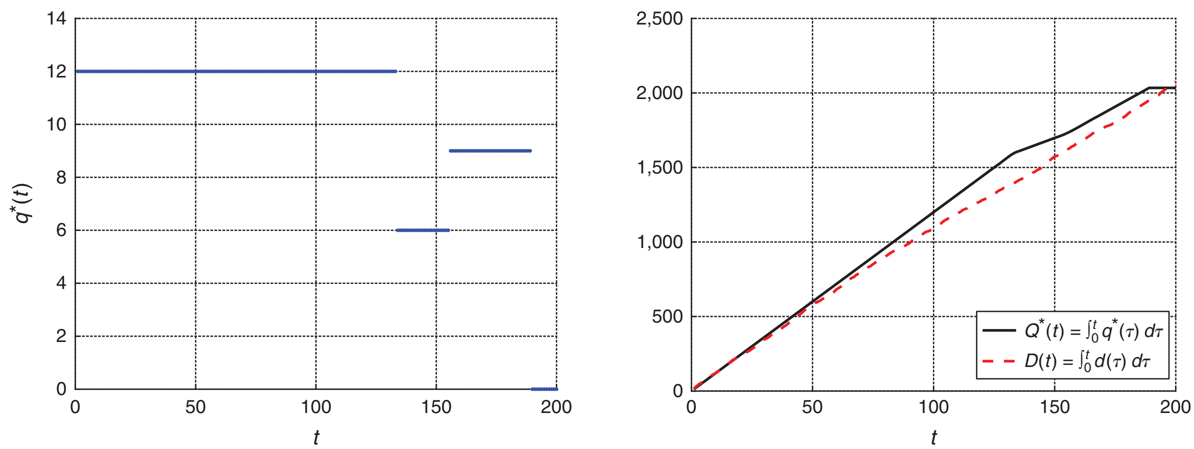

if $\mu_a + \mu_b = a + b$:

$$\begin{cases} (sb+ha)/(s+h), & x \in [0, (b-\mu_a)T/(b-a)], \\ (sa+hb)/(s+h), & x \in ((b-\mu_a)T/(b-a), T]; \end{cases}$$

if $\mu_a + \mu_b < a + b$:

$$\begin{cases} (sb+ha)/(s+h), & x \in [0, (\mu_b - a)T/(b-a)], \\ a, & x \in ((\mu_b - a)T/(b-a), (b-\mu_a)T/(b-a)], \\ (sa+hb)/(s+h), & x \in ((b-\mu_a)T/(b-a), T]. \end{cases}$$

If $c > sT$, $q^*(x) = 0$ *for* $x \in [0, T]$.

Proof of Theorem 1. Given the structure of Ω, we see that

$$\underline{D}(x) = \min_{d\in\Omega} \int_0^x d(t)\,dt = \max\{xa, \mu_a T - (T-x)b\}, \quad \forall x \in [0,T] \qquad (14)$$

and

$$\bar{D}(x) = \max_{d\in\Omega} \int_0^x d(t)\,dt = \min\{xb, \mu_b T - (T-x)a\}, \quad \forall x \in [0,T]. \qquad (15)$$

From Lemma 1, we have that

$$\int_0^x q^*(t)\,dt = \big(s\min\{xb, \mu_b T - (T-x)a\} + h\max\{xa, \mu_a T - (T-x)b\}\big)\cdot(s+h)^{-1}, \qquad (16)$$

for $x \in [0, T - c/s]$ and $q^*(x) = 0$ for $x \in (T - c/s, T]$, assuming that $c \leqslant sT$; if $c > sT$, Lemma 1 implies that $q^*(x) = 0$ for all $x \in [0, T]$. For the remainder of the proof, we assume $c \leqslant sT$ and that Lemma 1 holds for $x \in [0, T]$, and then simply truncate to $[0, T - c/s]$.

The first argument of the max operator in Equation (16) dominates if $x \leqslant (b-\mu_a)T/(b-a)$ and the first argument of the min dominates if $x \leqslant (\mu_b - a)T/(b-a)$. The former threshold for x is strictly less than the latter one iff $\mu_a + \mu_b > a + b$; therefore, we have three cases to consider: (i) $\mu_a + \mu_b > a + b$, (ii) $\mu_a + \mu_b = a + b$, and (iii) $\mu_a + \mu_b < a + b$.

In case (i),

$$\int_0^x q^*(t)\,dt = \begin{cases} \left(\dfrac{sb+ha}{s+h}\right)x, & x \in \left[0, \dfrac{(b-\mu_a)T}{b-a}\right], \\ bx - \dfrac{h(b-\mu_a)T}{s+h}, & x \in \left(\dfrac{(b-\mu_a)T}{b-a}, \dfrac{(\mu_b-a)T}{b-a}\right], \\ \left(\dfrac{sa+hb}{s+h}\right)x + \dfrac{(s\mu_b + h\mu_a - sa - hb)T}{s+h}, & x \in \left(\dfrac{(\mu_b-a)T}{b-a}, T\right], \end{cases}$$

which, taking the derivative with respect to x, results in

$$q^*(x) = \begin{cases} \dfrac{sb+ha}{s+h}, & x \in \left[0, \dfrac{(b-\mu_a)T}{b-a}\right], \\ b, & x \in \left(\dfrac{(b-\mu_a)T}{b-a}, \dfrac{(\mu_b-a)T}{b-a}\right], \\ \dfrac{sa+hb}{s+h}, & x \in \left(\dfrac{(\mu_b-a)T}{b-a}, T\right]. \end{cases}$$

In case (iii), we have that

$$\int_0^x q^*(t)\,dt = \begin{cases} \left(\dfrac{sb+ha}{s+h}\right)x, & x\in\left[0,\dfrac{(\mu_b-a)T}{b-a}\right],\\ ax+\dfrac{s(\mu_b-a)T}{s+h}, & x\in\left(\dfrac{(\mu_b-a)T}{b-a},\dfrac{(b-\mu_a)T}{b-a}\right],\\ \left(\dfrac{sa+hb}{s+h}\right)x+\dfrac{(s\mu_b+h\mu_a-sa-hb)T}{s+h}, & \\ \qquad x\in\left(\dfrac{(b-\mu_a)T}{b-a},T\right], & \end{cases}$$

which implies

$$q^*(x)=\begin{cases}\dfrac{sb+ha}{s+h}, & x\in\left[0,\dfrac{(\mu_b-a)T}{b-a}\right]\\ a, & x\in\left(\dfrac{(\mu_b-a)T}{b-a},\dfrac{(b-\mu_a)T}{b-a}\right]\\ \dfrac{sa+hb}{s+h}, & x\in\left(\dfrac{(b-\mu_a)T}{b-a},T\right].\end{cases}$$

In case (ii), the middle scenario disappears. These solutions are applied for $x\in[T-c/s,T]$ and $q^*(x)=0$ otherwise. □

3.1. Discussion of Theorem 1

To provide a more intuitive discussion, we consider the case where $\mu_a=\mu_b=\mu$ (i.e., T is large enough that convergence approximately holds). Differentiating Equations (14) and (15), we obtain two demand functions in Ω:

$$\underline{d}(t)=\begin{cases} a, & \text{if } t\leqslant\dfrac{(b-\mu)T}{b-a}\\ b, & \text{if } t>\dfrac{(b-\mu)T}{b-a}\end{cases}\quad\text{and}$$

$$\bar{d}(t)=\begin{cases} b, & \text{if } t\leqslant\dfrac{(\mu-a)T}{b-a}\\ a, & \text{if } t>\dfrac{(\mu-a)T}{b-a}.\end{cases}$$

The optimal robust solution in Theorem 1 *balances* the influence of these two extreme demands. To see this, consider demand that is uniformly distributed on the interval $[\underline{D},\bar{D}]$ with CDF F, and consider further the application of the newsvendor model with unit overage cost h and unit underage cost s. The standard newsvendor solution Q^* prescribes

$$F(Q^*)=\frac{s}{s+h}\implies\frac{Q^*-\underline{D}}{\bar{D}-\underline{D}}=\frac{s}{s+h}\implies Q^*=\frac{s\bar{D}+h\underline{D}}{s+h}.$$

Associating $Q^*=\int_0^x q^*(t)\,dt$, $\underline{D}(x)=\underline{D}$, and $\bar{D}(x)=\bar{D}$ for any $x\in[0,T-c/s]$, we obtain the main case of Lemma 1, which drives Theorem 1. Thus, our optimal robust solution balances the influence of these two polar extreme demands, as a function of the mismatch costs s and h in a newsvendor fashion.

The optimal robust ordering strategy depends on the relative sizes of μ and $(a+b)/2$. Since μ already has the interpretation of the mean drift rate of demand, we interpret $(a+b)/2$ as the median. Therefore, there is a notion of skewness that influences the ordering strategy. While linking the relative order of the mean and median might be imprecise in a statistical sense (see von Hippel 2005), it will suffice for our purposes. If $\mu>(a+b)/2$ we shall say our uncertainty set Ω has positive skew; likewise, if $\mu<(a+b)/2$, then we say Ω has negative skew; finally, if $\mu=(a+b)/2$, we say Ω is symmetric.

In all three cases of Theorem 1 (i.e., all skewness possibilities for Ω), there is a partition of the horizon $[0,T]$ into two–three intervals, which we denote as "early," "middle," and "late" (in the second case, there is no middle interval); the specific partitions depend on the level of skewness of Ω. In all three cases, during the early interval, the ordering rate is $(sb+ha)/(s+h)$ and during the late interval, the ordering rate is $(sa+hb)/(s+h)$; both ordering rates are simple weighted averages of the lower and upper bounds of the demand growth rate. The ordering rates also depend on the relative values of s and h: if $s>h$, then the earlier ordering is faster, to avoid the more expensive stockouts, and then the later ordering is slower to avoid inventory holding cost; if $s<h$, the observation is reversed, with slower ordering first, to avoid the more expensive inventory holding cost, and then faster ordering to avoid stockout costs.

We next focus our discussion on nonsymmetric Ω sets. The most prominent difference between the first (positively skewed Ω) and third (negatively skewed Ω) cases in Theorem 1 is the ordering level in the middle interval. If Ω has positive skew, the upper bound on the demand rate b is selected for the ordering rate; the relatively large mean drift of demand, compared to the median, induces the fast ordering in the middle interval so that the orders can "catch up" to demand that is growing fast. In contrast, if Ω has negative skew, the mean drift is relatively small, and the lower bound of the demand growth rate a is selected as the ordering rate, allowing the slow-growing demand to catch up with the supply.

3.2. Guidelines for Selecting (μ_a,μ_b,a,b)

We next provide a numerical study based on linear regression to further explore the issues we addressed above qualitatively for the special case where $\mu_a=\mu_b=\mu$. The economics are as follows: $c=20$, $s=2$, $h=1$. We set the true mean $\mu=10$ and standard deviation $\sigma=3$, and we generate the true demand using truncated normal random variables with mean μ and standard deviation σ as the random (nonnegative) instantaneous growth rate of demand. We consider $\mu_a \in \{8,9\}$, $\mu_b \in \{11,12\}$, $a \in \{6,7\}$, and $b \in \{13,14\}$. These parameters allow us to study asymmetry of μ with respect to both intervals $[\mu_a, \mu_b]$ and $[a,b]$; in addition, the different combinations of the (μ_a, μ_b, a, b) parameters lead to all three of the cases in Theorem 1. For each set of (μ_a, μ_b, a, b) values, we calculate the average percent increase in inventory costs of the robust strategy in Theorem 1 over the omniscient strategy of ordering $q = d$ (which would require perfect knowledge of d in advance), over $n = 10{,}000$ simulation trials; in other words, if Z_i^{robust} and $Z_i^* = c\int_0^T d(t)\,dt$ are the inventory costs of our robust strategy and the omniscient strategy in trial i, respectively, we return the average percent increase

$$PI = \frac{1}{n}\sum_{i=1}^{n}\frac{Z_i^{\text{robust}} - Z_i^*}{Z_i^*};$$

these 16 experiments result in values of PI that ranged from a minimum of 23.1% to a maximum of 76.9%, with a mean (median) of 48.4% (47.0%). Note that our baseline is the omniscient strategy of ordering $q = d$, which is not implementable in practice, and thus provides a very conservative benchmark.

To understand the impact of the (μ_a, μ_b, a, b) parameters precisely, we fitted a multiple-variable linear regression with these parameters as the independent variables and PI as the dependent variable over the 16 observations (experiments), obtaining

$$PI = -5.378 + 0.106\mu_a + 0.223\mu_b + 0.051a + 0.153b.$$

The adjusted-$R^2 = 0.971$, the a coefficient is significant at the 0.01 level, and the other coefficients are significant at the 0.001 level. We observe that the μ_b and b parameters more strongly deteriorate the performance of our robust strategy, and this is due to $s > h$. The robust optimization approach, being naturally conservative, avoids the more costly stockout; by setting the μ_b and b parameters relatively large (with respect to the μ_a and a parameters), we are exacerbating the conservatism, which results in higher costs. Alternatively, if $h > s$, we observe the opposite effect of stronger negative impact from the μ_a and a parameters. In both cases, the μ_a and μ_b parameters have a stronger impact (in terms of regression coefficients) than the respective a and b parameters. Consequently, this analysis provides qualitative guidance on how to set the parameters as a function of the s and h values: if the unit shortage cost is greater than the unit holding cost, set μ_b and b relatively closer to μ than μ_a and a, and vice versa if holding costs are larger than shortage costs, in order to avoid exacerbating the natural conservatism of robust optimization.

3.3. Terminal Costs

In this section, we explore the impact of terminal costs that are incurred at the end of the planning horizon at time $t = T$. In particular, we assume that there is a unit underage cost $F_u > 0$ and a unit overage cost $F_o > 0$ such that the model incurs the terminal cost

$$\max\{F_o I(T), -F_u I(T)\}.$$

These costs suggest a natural modification to the second and third sets of constraints in the model in Equation (5), namely,

$$y(x) \geqslant [h + F_o\delta(x - T)]I(x), \quad \forall\, d \in \Upsilon, \forall\, x \in [0, T]$$

and

$$y(x) \geqslant -[s + F_u\delta(x - T)]I(x), \quad \forall\, d \in \Upsilon, \forall\, x \in [0, T],$$

where δ is the Dirac delta function. It is straightforward to show, by modifying the proofs of Lemma 1 and Theorem 1 in a natural way, the following corollary, where q_F^* is the optimal ordering rate function under terminal costs and q^* is the positive part of the optimal ordering rate function from Theorem 1 (without terminal costs and not restricted to $x \leqslant T - c/s$).

Corollary 1. *Letting* $\Upsilon = \Omega$,
1. *if* $c \leqslant F_u$, *then* $q_F^*(x) = q^*(x)$ *for* $x \in [0, T]$;
2. *if* $F_u < c \leqslant sT + F_u$, *then* $q_F^*(x) = q^*(x)$ *for* $x \in [0, T - (c - F_u)/s]$ *and* $q_F^*(x) = 0$ *for* $x \in (T - (c - F_u)/s, T]$;
3. *if* $sT + F_u < c$, *then* $q_F^*(x) = 0$ *for* $x \in [0, T]$.

We conclude this section by providing intuition about this corollary. In Case 1, the unit terminal cost F_u is greater than the unit purchasing cost, and ordering continues until time $t = T$ to minimize the chance of a stockout. In Case 2, ordering stops at time $\tilde{t} = T - (c - F_u)/s$ because it is cheaper to incur a shortage for the remainder of the planning horizon than to purchase new products; in other words, the cumulative shortage cost *per unit* over the interval $[t, T]$, for $t \geqslant \tilde{t}$, is $F_u + s(T - t) \leqslant F_u + s(T - \tilde{t}) = c$; therefore, for $t \geqslant \tilde{t}$, it is cheaper to incur the shortage costs for the remaining horizon, than to procure new units. In Case 3, no ordering ever takes place since it is cheaper to incur a stockout for the entire planning horizon $[0, T]$ than to purchase any units. Finally, we point out that the terminal overage cost F_o does not play a role; this is due to the model ending up with a shortage, as long as $d \in \Omega$ (cf., Equation (11) for the $x \in (A, T]$ case); however, note that if $d \notin \Omega$, the ordering strategy in Theorem 1 could indeed end with positive inventory, incurring the F_o cost.

4. Dynamic Robust Inventory Management

In this section we introduce and analyze numerous dynamic variants of the static model studied in the previous section. These models allow us to study periodic review in a variety of scenarios. Our models are nonanticipatory. In other words, we consider a sequence of reoptimization models in a rolling-horizon framework, rather than one where we can apply a Bellman equation. The reason for this is that we are motivated to derive closed-form solutions, which we found to be intractable in various anticipatory models. Our approach, in a function space, is analogous to that in recent papers in finite-dimensional vector spaces, such as Solyalı et al. (2016) and Mamani et al. (2016).

To motivate our model, suppose that at time $\tau \in [0,T]$ we have observed the past demand stream rate, which we denote as $\hat{d}(t)$, $t \in [0,\tau]$, and we have recorded the past ordering strategy $\hat{q}(t)$, $t \in [0,\tau]$. Consequently, at time τ, we also know the inventory position $I(\tau) = \int_0^\tau (\hat{q}(t) - \hat{d}(t))\,dt$. For now, we only assume that $\hat{d} \in \mathscr{L}^2[0,\tau]$ and that $\hat{d}(t) \geqslant 0$ for all $t \in [0,\tau]$.

In what follows, the *consistency* of the observed demand stream $\hat{d}$ with the uncertainty set Ω will be important. In particular, we define the projection $P_\tau(\Omega)$ as the set of functions on $[0,\tau]$ that can be extended to a function contained in Ω:

$$P_\tau(\Omega) = \left\{ d \in \mathscr{L}^2[0,\tau]\colon \mu_a T - b(T-\tau) \leqslant \int_0^\tau d(x)\,dx \leqslant \mu_b T - a(T-\tau),\ a \leqslant d(x) \leqslant b,\ \forall\, x \in [0,\tau] \right\}. \quad (17)$$

We shall say that $\hat{d}$ is consistent with Ω if $\hat{d} \in P_\tau(\Omega)$. To the best of our knowledge, the issue of consistency has not been studied in previous rolling-horizon dynamic robust models, and the standard assumption is that $\hat{d}$ is indeed consistent; we are the first to study what can be done if there is no consistency between observed realizations and the uncertainty set. Note that it is quite possible that observed demand streams are inconsistent, as the literature typically considers uncertainty sets that do not span the support of the underlying distribution (e.g., in Bandi and Bertsimas 2012, the authors' uncertainty set only covers Γ standard deviations away from the mean for normally distributed uncertainties). Furthermore, the discussion on Ben-Tal et al. (2009, pp. 32–33) recommends choosing an uncertainty set that is smaller than the distributional support. However, an inconsistent demand stream results in problems of well-posedness and implementation of a rolling-horizon strategy. A standard approach in a rolling-horizon framework is to define an uncertainty set for the remaining horizon as the intersection of the observed demand stream and the original uncertainty set; in our context, at time τ, an uncertainty set for $[\tau,T]$ can be defined as $\Omega \cap \hat{d}$. However, if $\hat{d} \notin P_\tau(\Omega)$, then the intersection is empty, and the robust optimization model for $[\tau,T]$ is not well defined. Indeed, in Mamani et al. (2016), this issue is assumed away (see second paragraph of section 4 in Mamani et al. 2016). In this section, we study this issue in depth.

In the following sections we present a basic analysis, for various dynamic models, in terms of a single review epoch τ. We begin, in Section 4.1, by providing generalizations of Lemma 1 that allow for positive and negative initial inventories, respectively, at a generic time τ. In Section 4.2, we present a myopic dynamic model where at time τ we utilize a new uncertainty set for the remaining time horizon $[\tau,T]$ that does not depend on the past demand realization $\hat{d}$. In Sections 4.3.1 and 4.3.2 we define uncertainty sets on $[\tau,T]$ that depend on the demand realization $\hat{d}$. In the former, we assume that $\hat{d}$ is consistent, $\hat{d} \in P_\tau(\Omega)$. In the latter section we consider $\hat{d} \notin P_\tau(\Omega)$, and we find the nearest function $d^* \in P_\tau(\Omega)$ to $\hat{d}$; if d^* is "not too far" from $\hat{d}$, we build an appropriate uncertainty set for $[\tau,T]$ in terms of d^*; if d^* is "too far" from $\hat{d}$, we build a new uncertainty set for $[\tau,T]$ based on the observed demand $\hat{d}$.

4.1. Structural Results

As mentioned above, we consider various robust uncertainty sets for the demand streams defined on $[\tau,T]$. For now, we use Υ as a generic uncertainty set and later specify the full details of various sets. The basic optimization problem used in our dynamic models, for a generic uncertainty set Υ and initial inventory position $I(\tau)$, is

$$\begin{aligned} \min_{q(x)\geqslant 0} \ & \int_\tau^T (cq(x) + y(x))\,dx, \\ \text{s.t.}\ & y(x) \geqslant h\left(I(\tau) + \int_\tau^x (q(t) - d(t))\,dt\right), \quad \forall\, d \in \Upsilon,\ \forall\, x \in [\tau,T], \\ & y(x) \geqslant -s\left(I(\tau) + \int_\tau^x (q(t) - d(t))\,dt\right), \quad \forall\, d \in \Upsilon,\ \forall\, x \in [\tau,T]. \end{aligned} \quad (18)$$

Once Model (18) is solved and an optimal robust ordering rate function $q^*(t)$ is defined on $[\tau,T]$ for a given review epoch τ, we can easily extend the analysis to a countable set of review epochs $\mathscr{T} = \{\tau_1, \tau_2, \ldots\} \subseteq [0,T]$. At review epoch $\tau_i \in \mathscr{T}$, we determine $q^*(t)$ as a function of $I(\tau_i)$ for $t \in [\tau_i, T]$, but only apply it on $[\tau_i, \tau_{i+1})$, and at time τ_{i+1} we solve a new variant of Model (18).

We also generalize Equations (6) for the generic uncertainty set Υ defined on $[\tau,T]$, to determine the minimum and maximum cumulative demands on $[\tau,x]$, for all $x \in [\tau,T]$:

$$\underline{D}_\tau^\Upsilon(x) = \min_{d\in\Upsilon} \int_\tau^x d(t)\,dt \quad \text{and} \quad \bar{D}_\tau^\Upsilon(x) = \max_{d\in\Upsilon} \int_\tau^x d(t)\,dt. \quad (19)$$

We shortly generalize Lemma 1 to accommodate a nonzero initial inventory position $I(\tau)$ for a generic uncertainty set Υ on $[\tau,T]$. We provide two lemmas, one for nonnegative initial inventory $I(\tau) \geqslant 0$ and another for an initial backlog $I(\tau) < 0$. To concisely present the first lemma, it is useful to define a threshold parameter for the case of nonnegative initial inventory.

Definition 1. If $I(\tau) \geqslant 0$ and $(s\bar{D}^{\Upsilon}_{\tau}(T) + h\underline{D}^{\Upsilon}_{\tau}(T))/(s+h) \geqslant I(\tau)$, let

$$\chi^{\Upsilon}_{\tau} = \min x \quad \text{s.t.} \quad \frac{s\bar{D}^{\Upsilon}_{\tau}(x) + h\underline{D}^{\Upsilon}_{\tau}(x)}{s+h} \geqslant I(\tau), \quad \tau \leqslant x \leqslant T.$$

χ^{Υ}_{τ} is well defined and unique since $(s\bar{D}^{\Upsilon}_{\tau}(x) + h\underline{D}^{\Upsilon}_{\tau}(x))/(s+h)$ is nonnegative and nondecreasing in x.

Lemma 2. *If $I(\tau) \geqslant 0$, then the optimal robust ordering rate function q^* satisfies the following:*

1. *If $c \leqslant s(T-\tau)$ and $I(\tau) \leqslant (s\bar{D}^{\Upsilon}_{\tau}(T) + h\underline{D}^{\Upsilon}_{\tau}(T))/(s+h)$, then*
   (a) *if $\chi^{\Upsilon}_{\tau} \leqslant T - c/s$, then*
      i. *$q^*(x) = 0$, for $x \in [\tau, \chi^{\Upsilon}_{\tau})$;*
      ii. *$\int_{\tau}^{x} q^*(t)\,dt = (s\bar{D}^{\Upsilon}_{\tau}(x) + h\underline{D}^{\Upsilon}_{\tau}(x))/(s+h) - I(\tau)$, for $x \in [\chi^{\Upsilon}_{\tau}, T - c/s]$;*
      iii. *$q^*(x) = 0$, for $x \in (T - c/s, T]$.*
   (b) *If $\chi^{\Upsilon}_{\tau} > T - c/s$, then $q^*(x) = 0$ for all $x \in [\tau, T]$.*
2. *If $c > s(T-\tau)$ or $I(\tau) > (s\bar{D}^{\Upsilon}_{\tau}(T) + h\underline{D}^{\Upsilon}_{\tau}(T))/(s+h)$, then $q^*(x) = 0$ for all $x \in [\tau, T]$.*

Proof. The proof is similar to that of Lemma 1, and is presented in the appendix. □

Lemma 3. *If $I(\tau) < 0$, then the optimal robust ordering rate function q^* satisfies the following:*

1. *If $c \leqslant s(T-\tau)$, then*
   (a) *$\int_{\tau}^{x} q(t)\,dt = (s\bar{D}^{\Upsilon}_{\tau}(x) + h\underline{D}^{\Upsilon}_{\tau}(x))/(s+h) - I(\tau)$, for $x \in [\tau, T - c/s]$;*
   (b) *$q^*(x) = 0$, for $x \in (T - c/s, T]$.*
2. *If $c > s(T-\tau)$, then $q^*(x) = 0$ for all $x \in [\tau, T]$.*

Proof. The proof is similar to that of Lemma 2, and is presented in the appendix. □

4.2. Sample-Path-Independent Uncertainty Set

At time τ, we consider an uncertainty set that does not depend on the observed sample demand path $\hat{d}(t)$, $t \in [0, \tau)$. A convenient characteristic of the following approach is that we do not need to concern ourselves with whether or not $\hat{d}$ is consistent with Ω; consequently, this uncertainty set is perhaps the most naive possible, and we later compare its performance with more sophisticated approaches in Section 5.1, demonstrating that in certain situations, this simple approach suffices, whereas in other situations it offers subpar performance. We define the uncertainty set Ω_τ as the set of demand rates on $[\tau, T]$ that fall within the bounds a and b, and have an average drift value within $[\mu_a, \mu_b]$ over the interval $[\tau, T]$:

$$\Omega_\tau = \left\{ d \in \mathcal{L}^2[\tau, T]\colon \mu_a \leqslant \frac{1}{T-\tau}\int_{\tau}^{T} d(x)\,dx \leqslant \mu_b,\ a \leqslant d(x) \leqslant b,\ \forall\, x \in [\tau, T] \right\}.$$

We consider Formulation (18) with $\Upsilon = \Omega_\tau$:

$$\begin{aligned} \min_{q(x) \geqslant 0} \ & \int_{\tau}^{T} (cq(x) + y(x))\,dx \\ \text{s.t.}\ & y(x) \geqslant h\left(I(\tau) + \int_{\tau}^{x} (q(t) - d(t))\,dt\right), \quad \forall\, d \in \Omega_\tau,\ \forall\, x \in [\tau, T], \\ & y(x) \geqslant -s\left(I(\tau) + \int_{\tau}^{x} (q(t) - d(t))\,dt\right), \quad \forall\, d \in \Omega_\tau,\ \forall\, x \in [\tau, T]. \end{aligned} \tag{20}$$

Our subsequent results depend on the sign of the inventory position $I(\tau)$ at time τ, and our proofs utilize Lemmas 2 and 3 with $\Upsilon = \Omega_\tau$. We demonstrate that, under certain conditions, if there is an initial backlog $I(\tau) < 0$, our optimal strategy orders an impulse at time τ, and then applies a variant of the strategy in Theorem 1. Alternatively, if there is an initial nonnegative stock of inventory $I(\tau) \geqslant 0$, then, under certain conditions, the optimal robust strategy is to order nothing on the interval $[\tau, \chi^{\Omega_\tau}_{\tau})$, to allow the stock to deplete, where $\chi^{\Omega_\tau}_{\tau}$ is Definition 1 evaluated with $\Upsilon = \Omega_\tau$, and then a variant of the strategy in Theorem 1 on the interval $[\chi^{\Omega_\tau}_{\tau}, T]$. If the conditions are not met, then it is optimal to order nothing for the entire interval $[\tau, T]$.

The following theorem characterizes the optimal robust ordering rate function q^* for the Ω_τ uncertainty set, where $\delta \in \mathcal{L}^2[\tau, T]$ is the Dirac delta function (i.e., an impulse). Note that inventory control using impulse functions is well established in the literature; see, for example, Harrison et al. (1983) and Harrison and Taksar (1983). This theorem is illustrated in Figure 2.

Theorem 2. *Letting $\theta_1 = ((b-\mu_a)T + (\mu_a - a)\tau)/(b-a)$ and $\theta_2 = ((\mu_b - a)T + (b-\mu_b)\tau)/(b-a)$, we have the following:*

1. *If $I(\tau) < 0$ and $c \leqslant s(T-\tau)$, then the optimal robust ordering rate function $q^*(x) = 0$ for $x \in (T - c/s, T]$ and the following strategy for $x \in [\tau, T - c/s]$:*

If $\mu_a + \mu_b > a + b$:

$$-I(\tau)\delta(x-\tau) + \begin{cases} (sb+ha)/(s+h), & x \in (\tau, \theta_1], \\ b, & x \in (\theta_1, \theta_2], \\ (sa+hb)/(s+h), & x \in (\theta_2, T]; \end{cases}$$

if $\mu_a + \mu_b = a + b$:

$$-I(\tau)\delta(x-\tau) + \begin{cases} (sb+ha)/(s+h), & x \in (\tau, \theta_1], \\ (sa+hb)/(s+h), & x \in (\theta_1, T]; \end{cases}$$

Figure 2. (Color online) Left: The Optimal Ordering Strategy $q^*(t)$ for the $\mu_a + \mu_b > a + b$ Case of Theorem 2, Where $c = 20$, $s = 2$, $h = 1$, $T = 400$, $\tau_1 = 100$, $\tau_2 = 200$, $\tau_3 = 300$, $\mu_a = 8$, $\mu_b = 12$, $\mu = 10$, $\sigma = 3$, $a = 0$ and $b = 15$. Right: The Cumulative Orders of Theorem 2 Tracking Cumulative Demand, Which Is Generated Using Truncated Normal Random Variables with Mean μ and Standard Deviation σ as the Random (Nonnegative) Instantaneous Growth of Demand

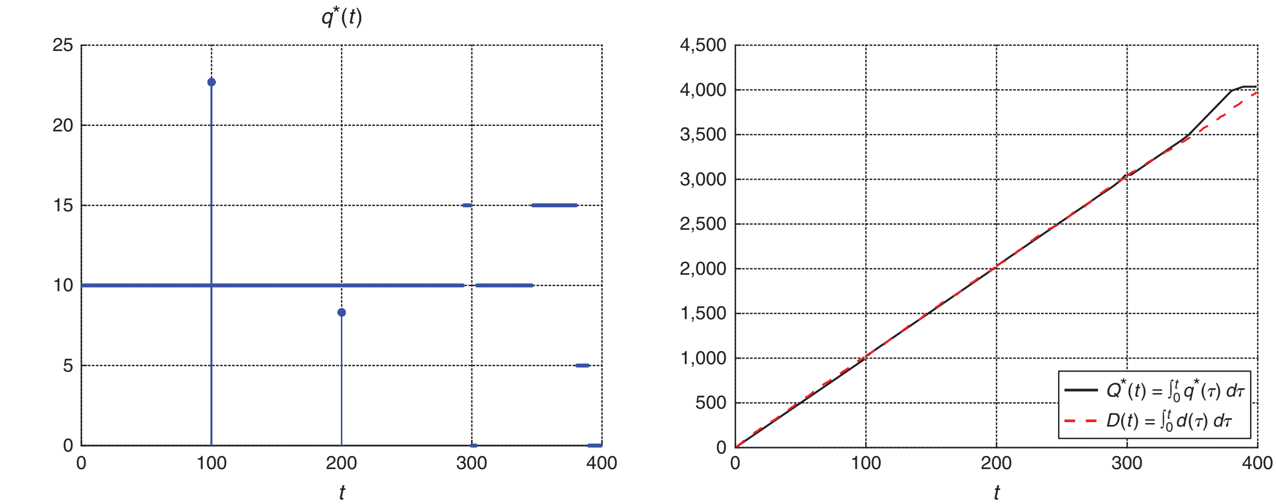

if $\mu_a + \mu_b < a + b$:

$$-I(\tau)\delta(x-\tau) + \begin{cases} (sb+ha)/(s+h), & x \in (\tau, \theta_2], \\ a, & x \in (\theta_2, \theta_1], \\ (sa+hb)/(s+h), & x \in (\theta_1, T]. \end{cases}$$

If $I(\tau) < 0$ *and* $c > s(T-\tau)$, $q^*(x) = 0$ *for all* $x \in [\tau, T]$.

2. *If* $I(\tau) \geqslant 0$, $c \leqslant s(T-\tau)$, $I(\tau) \leqslant (s\bar{D}_\tau^{\Omega_\tau}(T) + h\underline{D}_\tau^{\Omega_\tau}(T))/(s+h)$, *and* $\chi_\tau^{\Omega_\tau} \leqslant T - c/s$, *then the optimal robust ordering strategy is the following strategy, applied for* $x \in [\chi_\tau^{\Omega_\tau}, T - c/s]$, *and zero otherwise*:

If $\mu_a + \mu_b > a + b$:

$$\begin{cases} (sb+ha)/(s+h), & x \in (\tau, \theta_1], \\ b, & x \in (\theta_1, \theta_2], \\ (sa+hb)/(s+h), & x \in (\theta_2, T]; \end{cases}$$

if $\mu_a + \mu_b = a + b$:

$$\begin{cases} (sb+ha)/(s+h), & x \in (\tau, \theta_1], \\ (sa+hb)/(s+h), & x \in (\theta_1, T]; \end{cases}$$

if $\mu_a + \mu_b < a + b$:

$$\begin{cases} (sb+ha)/(s+h), & x \in (\tau, \theta_2], \\ a, & x \in (\theta_2, \theta_1], \\ (sa+hb)/(s+h), & x \in (\theta_1, T]. \end{cases}$$

If $I(\tau) \geqslant 0$ *and* $c > s(T-\tau)$ *or* $I(\tau) > (s\bar{D}_\tau^{\Omega_\tau}(T) + h\underline{D}_\tau^{\Omega_\tau}(T))/(s+h)$ *or* $\chi_\tau^{\Omega_\tau} > T - c/s$, *then the optimal robust ordering strategy* $q^*(x) = 0$ *for all* $x \in [\tau, T]$.

Proof. The proof is similar to that of Theorem 1, and is presented in the appendix. □

Note that there are many similarities and differences between the static and dynamic variants described in Theorems 1 and 2 that are worth discusssing. In both cases, the ordering rates (when there is ordering) are the same, either being a demand bound a or b, or one of two weighted averages of the bounds $(sb+ha)/(s+h)$ or $(sa+hb)/(s+h)$. The value of $\mu_a + \mu_b - (a+b)$ also plays the same role in both Theorems 1 and 2; this will not be true for subsequent models where the uncertainty set depends on the observed demand stream $\hat{d}$. Similarly, the partition of $[\tau, T]$ in Theorem 2 is a straightforward generalization of the partition of $[0, T]$ in Theorem 1; again, this will not be true for subsequent models that depend on $\hat{d}$. The main difference between Theorems 1 and 2 is the impact of the observed inventory position $I(\tau)$: if there is a backlog, an impulse order is applied to immediately arrive at the case of zero initial inventory (as in Theorem 1), and if there is nonnegative stock, the optimal robust strategy is to order nothing until a threshold time $\chi_\tau^{\Omega_\tau}$, which allows the inventory to deplete.

4.3. Sample-Path-Dependent Uncertainty Sets

In this section we present two approaches for defining uncertainty sets that incorporate past demand information, as represented by the observed demand stream $\hat{d}(t)$ for $t \in [0, \tau]$. In the first, in Section 4.3.1, we assume that $\hat{d} \in P_\tau(\Omega)$, where the projection $P_\tau(\Omega)$ is the set of functions on $[0, \tau]$ that can be extended to a function in Ω; the set $P_\tau(\Omega)$ is formally defined in Equation (17). In the second, in Section 4.3.2, we assume $\hat{d} \notin P_\tau(\Omega)$ and we find the nearest function $d^* \in P_\tau(\Omega)$ to $\hat{d}$. If d^* is not "too far" from $\hat{d}$, we then define an uncertainty set using d^*; if d^* is too far from $\hat{d}$, we discuss a reparameterized uncertainty set based on $\hat{d}$. In both cases we assume that $\hat{d} \in \mathscr{L}^2[0,\tau]$, and that $\hat{d}$ is nonnegative and bounded.

The literature that studies rolling-horizon robust inventory models (e.g., Mamani et al. 2016, Solyalı et al. 2016) implicitly assumes that $\hat{d}$ is consistent with Ω; thus, our results in Section 4.3.1 can be interpreted as a control-theoretic generalization of these results. However, the assumption of consistency is rather strong, since many uncertainty sets are specifically designed to be strict subsets of the support of the underlying distributions (e.g., see Bandi and Bertsimas 2012 and Ben-Tal et al. 2009), and this issue has, to the best of our knowledge, not been studied. Therefore, the analysis in Section 4.3.2 is new to the literature.

4.3.1. $\hat{d}$ Is Consistent with Ω. In this subsection, we assume $\hat{d} \in P_\tau(\Omega)$. We define an uncertainty set $\Omega_\tau^{\hat{d}}$ that consists of all functions $d \in \mathscr{L}^2[\tau,T]$ that extend $\hat{d}$ to a function contained in the original set Ω:

$$\Omega_\tau^{\hat{d}} = \left\{ d \in \mathscr{L}^2[\tau,T] \colon \mu_a T \leqslant \int_0^\tau \hat{d}(x)\,dx + \int_\tau^T d(x)\,dx \leqslant \mu_b T,\ a \leqslant d(x) \leqslant b,\ \forall\, x \in [\tau,T] \right\}. \quad (21)$$

Since $\hat{d} \in P_\tau(\Omega_\tau)$, the set $\Omega_\tau^{\hat{d}}$ is well defined and nonempty. We consider Formulation (18) with $\Upsilon = \Omega_\tau^{\hat{d}}$:

$$\begin{aligned} \min_{q(x)\geqslant 0} \ & \int_\tau^T (cq(x) + y(x))\,dx \\ \text{s.t.}\ & y(x) \geqslant h\left(I(\tau) + \int_\tau^x (q(t)-d(t))\,dt\right), \quad \forall\, d \in \Omega_\tau^{\hat{d}},\ \forall\, x \in [\tau,T], \\ & y(x) \geqslant -s\left(I(\tau) + \int_\tau^x (q(t)-d(t))\,dt\right), \quad \forall\, d \in \Omega_\tau^{\hat{d}},\ \forall\, x \in [\tau,T]. \end{aligned}$$

The solution to this problem depends on the cumulative demand observed in $\hat{d}$ up to time τ; we next define the relevant metric.

Definition 2. $\hat{D}(\tau) = \int_0^\tau \hat{d}(x)\,dx$.

The following theorem characterizes the optimal robust ordering rate function q^* for the $\Omega_\tau^{\hat{d}}$ uncertainty set, which assumes that $\hat{d} \in P_\tau(\Omega)$.

Theorem 3. *Letting $\hat{\theta}_1 = ((b-\mu_a)T + \hat{D}(\tau) - a\tau)/(b-a)$, $\hat{\theta}_2 = ((\mu_b - a)T + b\tau - \hat{D}(\tau))/(b-a)$, and $\Psi = ((\mu_a+\mu_b)T - (T-\tau)(a+b))/2$, we have the following:*

1. *If $I(\tau) < 0$ and $c \leqslant s(T-\tau)$, then the optimal robust ordering rate function $q^*(x) = 0$ for $x \in (T - c/s, T]$ and the following strategy for $x \in [\tau, T-c/s]$:*

If $\hat{D}(\tau) < \Psi$:

$$-I(\tau)\delta(x-\tau) + \begin{cases} (sb+ha)/(s+h), & x \in (\tau,\hat{\theta}_1], \\ b, & x \in (\hat{\theta}_1,\hat{\theta}_2], \\ (sa+hb)/(s+h), & x \in (\hat{\theta}_2,T]; \end{cases}$$

if $\hat{D}(\tau) = \Psi$:

$$-I(\tau)\delta(x-\tau) + \begin{cases} (sb+ha)/(s+h), & x \in (\tau,\hat{\theta}_1], \\ (sa+hb)/(s+h), & x \in (\hat{\theta}_1,T]; \end{cases}$$

if $\hat{D}(\tau) > \Psi$:

$$-I(\tau)\delta(x-\tau) + \begin{cases} (sb+ha)/(s+h), & x \in (\tau,\hat{\theta}_2], \\ a, & x \in (\hat{\theta}_2,\hat{\theta}_1], \\ (sa+hb)/(s+h), & x \in (\hat{\theta}_1,T]. \end{cases}$$

If $I(\tau) < 0$ and $c > s(T-\tau)$, $q^*(x) = 0$ for all $x \in [\tau,T]$.

2. *If $I(\tau) \geqslant 0$, $c \leqslant s(T-\tau)$, $I(\tau) \leqslant (s\bar{D}_\tau^{\Omega_\tau^{\hat{d}}}(T) + h\underline{D}_\tau^{\Omega_\tau^{\hat{d}}}(T))/(s+h)$, and $\chi_\tau^{\Omega_\tau^{\hat{d}}} \leqslant T - c/s$, then the optimal robust ordering strategy is the following strategy, applied for $x \in [\chi_\tau^{\Omega_\tau^{\hat{d}}}, T - c/s]$, and zero otherwise:*

If $\hat{D}(\tau) < \Psi$:

$$\begin{cases} (sb+ha)/(s+h), & x \in (\tau,\hat{\theta}_1], \\ b, & x \in (\hat{\theta}_1,\hat{\theta}_2], \\ (sa+hb)/(s+h), & x \in (\hat{\theta}_2,T]; \end{cases}$$

if $\hat{D}(\tau) = \Psi$:

$$\begin{cases} (sb+ha)/(s+h), & x \in (\tau,\hat{\theta}_1], \\ (sa+hb)/(s+h), & x \in (\hat{\theta}_1,T]; \end{cases}$$

if $\hat{D}(\tau) > \Psi$:

$$\begin{cases} (sb+ha)/(s+h), & x \in (\tau,\hat{\theta}_2], \\ a, & x \in (\hat{\theta}_2,\hat{\theta}_1], \\ (sa+hb)/(s+h), & x \in (\hat{\theta}_1,T]. \end{cases}$$

If $I(\tau) \geqslant 0$ and $c > s(T-\tau)$ or $I(\tau) > (s\bar{D}_\tau^{\Omega_\tau^{\hat{d}}}(T) + h\underline{D}_\tau^{\Omega_\tau^{\hat{d}}}(T))/(s+h)$ or $\chi_\tau^{\Omega_\tau^{\hat{d}}} > T - c/s$, then the optimal robust ordering strategy $q^*(x) = 0$ for all $x \in [\tau,T]$.

Proof. The proof is similar to that of Theorem 2, and is presented in the appendix. □

We first identify similarities between Theorem 3 and Theorems 1 and 2. The ordering rates are the same for all three theorems, consisting of zero, the bounds a and b, the weighted averages $(sb+ha)/(s+h)$ and $(sa+hb)/(s+h)$, and impulse ordering. Next, we contrast the theorems; we have previously contrasted Theorems 1 and 2, so we focus on the differences between Theorems 3 and 2. The most pronounced difference is the optimal partition of the interval $[\tau,T]$. In Theorem 2, the partition consists of intervals whose breakpoints are weighted averages of τ and T, which only depend on the value of $\mu_a + \mu_b - (a+b)$. In contrast, the optimal partition of Theorem 3 depends on the realized demand stream $\hat{d}$. More prominently, when the results in Theorems 2 and 3 are applied repeatedly for the review epochs $\tau \in \{\tau_1,\tau_2,\ldots,\}$, the partition is *static* under Theorem 2 and *dynamic* under Theorem 3, with respect to the demand stream. In other words, the three cases of Theorem 3 are not fixed for the reoptimization at each review epoch τ_i, as they depend on the updated observation of demand; in contrast, the cases under Theorem 2, for a given set of τ_i, are fixed since the partition only depends on the value of $\mu_a + \mu_b - (a+b)$, which is known a priori.

4.3.2. $\hat{d}$ Is Not Consistent with Ω. In this subsection, we assume $\hat{d} \notin P_\tau(\Omega)$, which precludes the natural definition $\Omega_\tau^{\hat{d}} = \Omega \cap \hat{d}$ that we analyzed in the previous section. The literature on rolling-horizon approaches to robust inventory management use variants of $\Omega_\tau^{\hat{d}}$, which require the consistency of $\hat{d}$ with Ω (e.g., Mamani et al. 2016 and Solyalı et al. 2016 implicitly assume this consistency). Notably, if the observed demand stream is inconsistent (i.e., $\Omega \cap \hat{d} = \varnothing$), then the rolling-horizon strategies from the literature are infeasible, and no ordering strategy can be derived. Here, we address this problem and propose two solution approaches, depending on the degree of consistency violation.

We first find the nearest function $d^* \in P_\tau(\Omega)$ to $\hat{d}$. In particular, we solve the following optimization problem:

$$d^* = \arg\min \|d - \hat{d}\| \quad \text{s.t.} \quad d \in P_\tau(\Omega). \tag{22}$$

Noting that Ω is a closed convex subset of the Hilbert space $\mathcal{L}^2[0,T]$, we can apply the Hilbert Projection Theorem (as presented in Luenberger 1969, p. 51):

Theorem 4. *Let $\mathcal{H}$ be a Hilbert space and M a closed subspace of $\mathcal{H}$. Corresponding to any vector $x \in \mathcal{H}$, there is a unique vector $m_o \in M$ such that $\|x - m_o\| \leqslant \|x - m\|$ for all $m \in M$.*

Letting $\mathcal{H} = \mathcal{L}^2[0,\tau]$, $M = P_\tau(\Omega)$ and $x = \hat{d}$, Theorem 4 implies that Problem (22) has a unique solution $d^* = m_o$, but does not specify how to actually find the solution. In the next lemma, we derive d^*.

Lemma 4. 1. *If $\int_0^\tau \min\{b, \max\{a, \hat{d}(x)\}\}\,dx \in [\mu_a T - b(T-\tau), \mu_b T - a(T-\tau)]$, then*

$$d^*(x) = \min\{b, \max\{a, \hat{d}(x)\}\}, \quad \forall x \in [0,\tau].$$

2. *If $\int_0^\tau \min\{b, \max\{a, \hat{d}(x)\}\}\,dx < \mu_a T - b(T-\tau)$, then*

$$d^*(x) = \min\{b, \min\{b, \max\{a, \hat{d}(x)\}\} + c_1\}, \quad \forall x \in [0,\tau],$$

where $c_1 = \min\{c \in \mathbb{R}^+ \colon \int_0^\tau \min\{b, \min\{b, \max\{a, \hat{d}(x)\}\} + c\}\,dx \geqslant \mu_a T - b(T-\tau)\}$.

3. *If $\int_0^\tau \min\{b, \max\{a, \hat{d}(x)\}\}\,dx > \mu_b T - a(T-\tau)$, then*

$$d^*(x) = \max\{a, \min\{b, \max\{a, \hat{d}(x)\}\} - c_2\}, \quad \forall x \in [0,\tau],$$

where $c_2 = \min\{c \in \mathbb{R}^+ \colon \int_0^\tau \max\{a, \min\{b, \max\{a, \hat{d}(x)\}\} - c\}\,dx \leqslant \mu_b T - a(T-\tau)\}$.

Proof of Lemma 4. We define the set $X = \{d \in \mathcal{L}^2[0,\tau] \colon a \leqslant d(x) \leqslant b,\ \forall x \in [0,\tau]\}$. We decompose the projection onto $P_\tau(\Omega)$ by first projecting $\hat{d}$ onto X, and then projecting the result onto $P_\tau(\Omega)$. This decomposition is valid since the sets are nested, namely, $\mathcal{L}^2[0,\tau] \supset X \supset P_\tau(\Omega)$; see, for example, Berberian (1961, p. 76). Let

$$g^* = \arg\min \|d - \hat{d}\| \quad \text{s.t.} \quad d \in X;$$

it is straightforward to see that $g^*(x) = \min\{b, \max\{a, \hat{d}(x)\}\}$ for $x \in [0,\tau]$.

We next project g^* onto $P_\tau(\Omega) = \{d \in \mathcal{L}^2[0,\tau] \colon L \leqslant \int_0^\tau d(x)\,dx \leqslant U,\ a \leqslant d(x) \leqslant b,\ \forall x \in [0,\tau]\}$, where $L = \mu_a T - b(T-\tau)$ and $U = \mu_b T - a(T-\tau)$. We have three cases to consider: (1) $\int_0^\tau g^*(x)\,dx \in [L,U]$, (2) $\int_0^\tau g^*(x)\,dx < L$, and (3) $\int_0^\tau g^*(x)\,dx > U$. In Case (1), $d^* = g^*$. In Case (2), we let

$$c_1 = \min_{c \in \mathbb{R}^+} c \quad \text{s.t.} \quad \int_0^\tau \min\{b, g^*(x) + c\}\,dx \geqslant L,$$

where c_1 is the minimum upward shift of $g^*(x)$, with a truncation at the upper bound b, in order to bring the total integral up to L. We claim that $d^*(x) = \min\{b, g^*(x) + c_1\}$ for $x \in [0,\tau]$; to prove this, we write the projection as

$$\begin{aligned} \min_{d \in \mathcal{L}^2[0,\tau]} \quad & \langle d - g^*, d - g^* \rangle \\ \text{s.t.} \quad & \int_0^\tau d(t)\,dt - U \leqslant 0 && : \alpha, \\ & -\int_0^\tau d(t)\,dt + L \leqslant 0 && : \beta, \\ & d(t) - b \leqslant 0,\ t \in [0,\tau] && : \gamma(t), \\ & -d(t) + a \leqslant 0,\ t \in [0,\tau] && : \delta(t), \end{aligned} \tag{23}$$

where the $\alpha, \beta, \gamma(t)$, and $\delta(t)$ are KKT multipliers. Problem (23) is a convex optimization problem; thus the KKT conditions are sufficient. We let $A = \{t \colon d^*(t) = b\}$ and $B = \{t \colon d^*(t) = g^*(t) + c_1\}$, and set $\alpha = 0$, $\beta = 2c_1$, $\delta(t) = 0$ for all $t \in [0,\tau]$, and

$$\gamma(t) = \begin{cases} 2(c_1 - (b - g^*(t)) & t \in A, \\ 0 & t \in B. \end{cases}$$

Clearly d^* is feasible for Problem (23), and $\alpha, \beta, \delta(t) \geqslant 0$ for all $t \in [0,\tau]$; by the definition of d^* and the set A, we also conclude that $\gamma(t) \geqslant 0$ for all $t \in [0,\tau]$. By the definition of c_1, the constraint $-\int_0^\tau d(t)\,dt + L \leqslant 0$ is tight, and for $t \in A$, the constraint $d^*(t) - b \leqslant 0$ is tight; this shows that complementary slackness holds for the positive KKT multipliers.

Finally, we address the vanishing gradient. Since we work in a Hilbert space, we utilize Fréchet derivatives (i.e., the derivative of a functional with respect to its function argument; see Luenberger 1969). The Fréchet derivative of $\langle d-g^*, d-g^*\rangle$, with respect to the function d, is $2(d-g^*)$, the Fréchet derivative of $\int_0^\tau d(t)\,dt$ is $e(t)=1$, $\forall\, t\in[0,\tau]$, and the Fréchet derivative of $d(t)$, namely, evaluating the function d at its argument t, is the Dirac delta function $\delta(x-t)$. The Hilbert space gradient of the Lagrangian, for $t\in[0,\tau]$ and evaluated at $d^*=\min\{b, g^*+c_1\}$, can be written (see Ulbrich 2009, section 2.5.5) as

$$2(d^*(t)-g^*(t))+(\alpha-\beta)e(t)+\int_0^\tau \delta(x-t)(\gamma(x)-\delta(x))\,dx$$
$$=2(d^*(t)-g^*(t))+(\alpha-\beta)e(t)+\gamma(t)-\delta(t),$$

which evaluates to the zero function. Thus, $d^*=\min\{b, g^*+c_1\}$ is optimal. The analysis of Case (3) is similar: we let

$$c_2=\min_{c\in\mathbb{R}^+} c \quad \text{s.t.} \quad \int_0^\tau \max\{a, g^*(x)-c\}\,dx \leqslant U,$$

where c_2 is the minimum downward shift of $g^*(x)$, with a truncation at the lower bound a, in order to bring the total integral down to U; we conclude that $d^*(x)=\max\{a, g^*(x)-c_2\}$ for $x\in[0,\tau]$. □

Note that, since Lemma 4 provides a closed-form solution for d^*, its implementation in practice is very efficient, in the spirit of the closed-form ordering functions derived elsewhere in this paper. Once d^* is determined, we have two alternatives, depending on how close d^* is to $\hat{d}$, which is measured by $\|d^*-\hat{d}\|$. Given a threshold parameter η, we say that d^* is "close enough" to $\hat{d}$ if $\|d^*-\hat{d}\|\leqslant\eta$. In this case, we define an uncertainty set based on d^*, since the violation of consistency is relatively minor; the parameter η can be used to control the conservatism of what is meant by "relatively minor." First, we utilize the uncertainty set $\Omega_\tau^{\hat{d}}$ in Equation (21) with $\hat{d}$ replaced with d^*. Second, we modify Definition 2 for d^*, namely, $\hat{D}(\tau)=\int_0^\tau d^*(x)\,dx$. Finally, Theorem 3 can be applied with the new uncertainty set $\Omega_\tau^{d^*}$ and modified definition of $\hat{D}(\tau)$. Note that, without this step, even very minor consistency violations will result in the failure of the rolling-horizon approach, as the optimization model at time τ is infeasible.

If $\|d^*-\hat{d}\|>\eta$, this is indicative of a relatively major violation of consistency, which suggests that the modeling parameters (μ_a,μ_b,a,b) have been incorrectly chosen, as the uncertainty set is not capturing the observed demand. In this case, it is suggested to revisit these parameters based on the information gleaned from the observed demand stream. In particular, $\hat{\mu}=(1/\tau)\int_0^\tau \hat{d}(t)\,dt$ is the estimated demand mean over the interval $[0,\tau]$, and this can be used to update the prior belief on the true mean μ, which in turn can be used to update the parameters (μ_a,μ_b,a,b); i.e., our model assumes that $a\leqslant\mu_a\leqslant\mu\leqslant\mu_b\leqslant b$, which can be corrected at this step.

5. Computational Studies

In this section we further explore our closed-form strategies in contexts that are difficult to analyze theoretically. In Section 5.1 we compare our dynamic strategies under various demand processes, including those that relax the assumptions of our basic model (by introducing seasonality). In Section 5.2, we examine the impact of different frequencies of review, and demonstrate that more frequent review does not necessarily lead to improved performance when comparing between robust strategies, but more frequent review does improve performance with respect to an omniscient ordering strategy. Finally, in Section 5.3 we study the impact of the ratio c/s and show that, even with relatively large purchase costs, most demand is satisfied.

5.1. A Comparison of Dynamic Strategies

Recall that we have analyzed three variants of a dynamic strategy. In Section 4.2 we considered a variant where at every review epoch τ, a new uncertainty set is constructed that does not depend on past realizations of demand; we denote this strategy the simple (S) strategy. In Section 4.3 we considered two strategies, where at every review epoch τ, a new uncertainty set is constructed that depends on the past demand realizations. In Section 4.3.1 we assumed that the past demand was consistent with the original uncertainty set Ω, which simplifies the analysis; however, in our implementation, we found that this assumption was frequently violated, so we do not consider this strategy in our computational experiments. In Section 4.3.2, we relaxed this assumption, and used the Hilbert Projection Theorem to remedy the inconsistency. We select this latter strategy for study and denote it the complex (C) strategy.

In our experiments we set $T=400$, $\mu=1$, $\sigma=1/3$, $h=1$, $s=2$, $c=20$, $a=0.5$, and $b=1.5$, though our results qualitatively hold for a wide range of parameter values. We set $\mu_a=\mu_b=\mu$, as this results in the best differentiation between the S and C strategies, and best performance for the C strategy, though we observe the same pattern for $\mu_a<\mu<\mu_b$; this provides evidence that, despite the limit of the strong law not holding for finite T, good computational performance can be obtained by setting $\mu_a=\mu_b=\mu$ (which is consistent with the recommendations of the literature that robust uncertainty sets need not exactly coincide with probabilistic behavior). We (initially) consider three review epochs: $\tau_1=100$, $\tau_2=200$, and $\tau_3=300$. We utilize two approaches for generating the demand streams. In the first, which is motivated by Brownian motion and denoted A_1, the instantaneous demand rate $d(t)$ is an

Table 1. Average Costs per Unit Time of the S and C Dynamic Strategies for the A_1 and A_2 Generators of Demand

	Average cost per unit time of strategy S (Section 4.2)	Average cost per unit time of strategy C (Section 4.3.2)
A_1	23.05	23.29
A_2	31.49	27.27

independent truncated normal random variable with mean μ and standard deviation σ: $d(t) = \max\{d, 0\}$, where $d \sim N(\mu, \sigma)$. In the second approach to generate demand, denoted A_2, we relax the assumptions of our modeling to introduce seasonality into the demand stream, and $d(t) = \mu + M \cdot \sin(2\pi t/T))$, where M is a random modulation that is normally distributed with mean $\mu/2$ and standard deviation $\mu/6$ (i.e., to ensure $d(t) \geqslant 0$ with high probability; we truncate, similarly to A_1, in the event that $d(t) < 0$). In Table 1 we present the average costs per unit time for the S and C strategies over 10,000 simulation trials, for the A_1 and A_2 generators of demand.

We see that the S strategy performs quite well, with respect to the C strategy, when the demand stream satisfies the assumptions of the analytical modeling (i.e., A_1 is the demand generator). Thus, when cumulative demand is expected to be well behaved, and in general agreement with the properties of Brownian motion (i.e., expected linear trend with no apparent seasonality), the S strategy is (arguably) satisfactory and there is no need to consider the C strategy. However, when the modeling assumptions are violated (i.e., A_2 is the demand generator), the C strategy is clearly preferred, as it demonstrates a capability to react to information contained in historical demand, resulting in over 15% lower costs. Thus, when demand seasonality is expected, we recommend utilizing the more complex model of Section 4.3.2; in the next subsection, we provide guidance in selecting the review frequency as a function of the seasonality period.

We also point out that the projection operator was indeed invoked in the implementation of the C strategy, under both the A_1 and A_2 demand generators, which provides further motivation for studying the consistency of demand streams and uncertainty sets in rolling-horizon strategies. However, this suggests that there are two potential reasons for the improved performance of C over S, under the A_2 generator: (1) the usage of past demand information and (2) the projection of $\hat{d}$ onto $P_\tau(\Omega)$. To tease out the effect of each, we repeated the experiment with demand modified to always be consistent, and we obtained qualitatively similar improvements of C over S. Thus, we are confident that most of the value comes from using past demand information. However, the resolution of demand inconsistency is required to attain this value of past demand information, as inconsistent demand renders the rolling-horizon approach infeasible.

We conclude this section by presenting, in Figure 3, a visual representation of the S and C strategies under a single sample path of A_2: the C strategy exhibits better performance because it is more responsive in reacting to the seasonality, and is better able to track the cumulative demand toward the end of the planning horizon.

5.2. A Study of the Review Frequency

In this subsection, we explore the impact of the frequency of review epochs. As in the previous subsection, we consider the two generators of demand A_1 and A_2, as well as the same parameters. We consider four frequencies of review: F_2 represents a single review epoch $\tau = T/2$, F_4 represents the review epochs

Figure 3. (Color online) The S Strategy (Left) and the C Strategy (Right) for the Seasonality Generator of Demand A_2, Where $c = 20$, $s = 2$, $h = 1$, $T = 400$, $\tau_1 = 100$, $\tau_2 = 200$, $\tau_3 = 300$, $\mu_a = 8$, $\mu_b = 12$, $\mu = 10$, $\sigma = 3$, $a = 0$ and $b = 15$

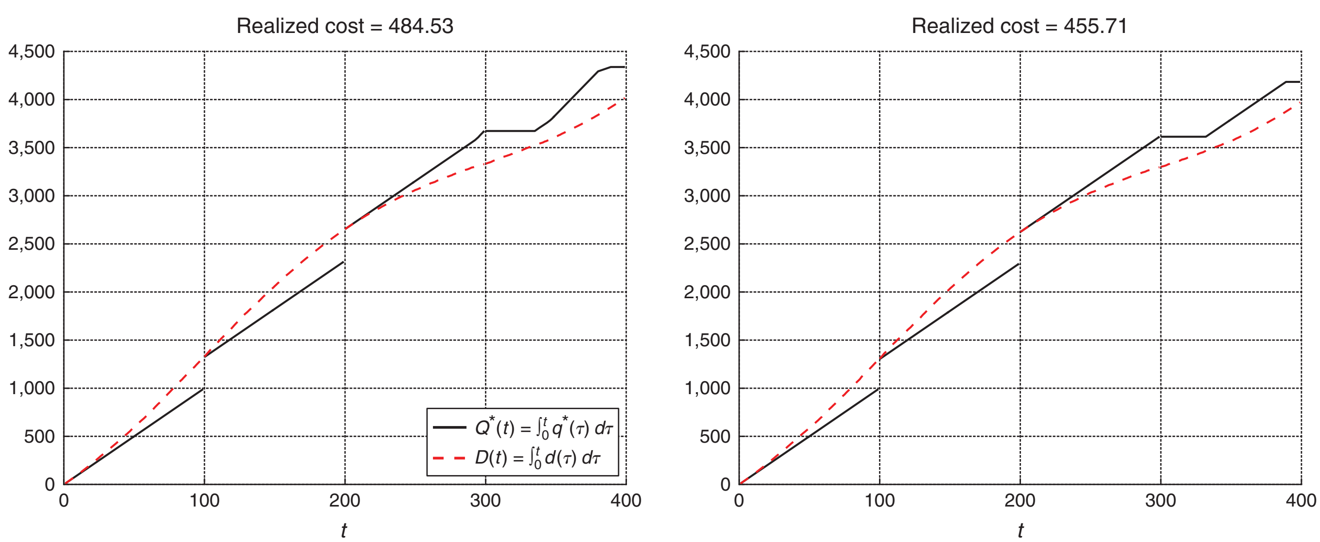

Table 2. Average Costs per Unit Time of the S and C Dynamic Strategies for the A_1 (Left) and A_2 (Right) Generators of Demand, for the Different Review Frequencies F_2, F_4, and F_8

	Average cost of strategy S (Section 4.2)	Average cost of strategy C (Section 4.3.2)		Average cost of strategy S (Section 4.2)	Average cost of strategy C (Section 4.3.2)
F_2	32.94 (64.70%)	33.26 (66.30%)	F_2	58.34 (191.70%)	38.16 (90.80%)
F_4	26.87 (34.35%)	27.21 (36.05%)	F_4	42.31 (111.55%)	31.01 (55.05%)
F_8	23.05 (15.25%)	23.30 (16.50%)	F_8	31.50 (57.50%)	27.27 (36.35%)
F_{16}	21.43 (7.15%)	21.48 (7.40%)	F_{16}	25.35 (26.75%)	24.71 (23.55%)

Table 3. Percent of Average Cumulative Demand Produced by Strategy C for $\rho = c/s$

	$\rho=1$	$\rho=5$	$\rho=10$	$\rho=20$	$\rho=30$	$\rho=50$	$\rho=75$	$\rho=100$
$Q^*(T)/E[D(T)]$	1.0735	1.0640	1.0517	1.0266	1.0014	0.9514	0.9090	0.9091

$\tau \in \{T/4, 2T/4, 3T/4\}$ (the same setup as in the previous subsection), F_8 represents $\tau \in \{iT/8: i \in \mathbb{Z}, 1 \leqslant i \leqslant 7\}$, and F_{16} represents $\tau \in \{iT/16: i \in \mathbb{Z}, 1 \leqslant i \leqslant 15\}$. Our performance metric is the same as in the previous subsection: the average cost per unit time over 10,000 simulation trials. We also report, in parentheses, the percent increase in expected cost with respect to the omniscient strategy of ordering $q = d$ for the entire horizon.

Our first results are presented in the left side of Table 2 for the A_1 generator of demand; unsurprisingly, the more frequent the review epochs, the lower the average cost, for both the S and C strategies. We also observe that the suboptimality, in percentage terms, is attractive, especially for higher frequencies; for instance, for F_{16}, both the S and C strategies return average costs only approximately 7% higher than the omniscient (and unimplementable in practice) strategy of ordering $q(t) = d(t)$ for all $t \in [0, T]$. Finally, the performances of the S and C strategies are comparable for all review frequencies, and thus the simpler S strategy is preferred.

However, under the A_2 generator of demand, we observe qualitatively different results, which are presented in the right side of Table 2. The C strategy significantly outperforms the S strategy for all review frequencies. Furthermore, the largest gap between the two costs is achieved with the *lowest* frequency, which (not coincidentally) is half of the period of the seasonality. Since demand is heaviest in the first half of the period (from the sine function), this single review epoch allows the C strategy to most effectively update the uncertainty set to essentially reduce the expectation of demand in the second half of the period (i.e., to match the negative portion of the sine function). However, when comparing with the omniscient strategy $q = d$, the performances of both strategies improve with increased review frequencies. Thus, if possible, high review frequencies should be selected; however, if this ability is limited under demand seasonality, the review frequency should be carefully selected to take advantage of the periodicity.

5.3. The Impact of Purchase Cost

Our optimal ordering rate functions cease ordering at a certain time that depends on the ratio $\rho = c/s$, in order to minimize cost. Here we study the impact of this ratio on the average percent of cumulative demand filled, for the C strategy (results are comparable for the S strategy). We consider $\rho \in \{1, 5, 10, 20, 30, 50, 75, 100\}$ and evaluate the ratio $Q^*(T)/E[D(T)]$, estimated over 1,000 simulation trials and presented in Table 3, where the experimental setup is identical to those of Figures 2 and 3, except that $c = \rho s$. We observe that small values of ρ allow the robust strategy to satisfy all demand, on average, and only large values of ρ (> 30) sacrifice some demand to reduce cost. However, in the worst case considered ($\rho = 100$), less than 10% of expected demand is lost.

6. Conclusion

In this paper we have analyzed new models of inventory management that uniquely lie at the intersection of robust optimization and optimal control theory. Our uncertainty sets are motivated by strong laws of large numbers for stochastic processes. We analyzed one static and three dynamic models, where in all cases we show that the *closed-form* optimal ordering rate function orders one of four possible rates, namely, the demand rate bounds a and b, and the weighted averages $(sb + ha)/(s + h)$ and $(sa + hb)/(s + h)$. In the dynamic problems, there is also the possibility of an ordering impulse, to satisfy an observed backlog. While the ordering rates come from the same set of four values in all problems, the times at which to apply each rate depend on the partition of the planning horizon, which depend on the problem. In the static case, $[0, T]$ is partitioned based on the value of $\mu_a + \mu_b - (a + b)$; in our simplest dynamic problem, where the uncertainty set does not depend on the past realization of demand, $[\tau, T]$ is also partitioned based on $\mu_a + \mu_b - (a+b)$. When the uncertainty set depends on the past realization of demand, the partition depends on the realization, as well as the values of μ_a, μ_b, a, and b. In our most general dynamic model, we allow for the inconsistency of the realized demand with the structure of a canonical uncertainty set, which we resolve using the Hilbert Projection Theorem; in this case, the partition depends on the projection and the values of μ_a, μ_b, a, and b. Our models can also accommodate discounting: if all costs are discounted at the same continuous rate, our optimal robust ordering strategies do not change. Computational results provide insight into when each dynamic variant should be applied: (1) if demand is well behaved, then any variant suffices, whereas (2) if there is seasonality in demand, the third dynamic variant, with an appropriately selected review frequency, exhibits the strongest performance by far.

We believe the ideas in this paper could be useful for analyzing other new problems defined at the intersection of robust optimization and optimal control theory. In addition, we found the use of probabilistic limit theorems as a guide in forming uncertainty sets to be particularly useful. Similarly, CLT-style uncertainty sets have recently received significant attention in a variety of areas (e.g., auction design, information theory, option pricing, queueing theory; see our literature review), though only in finite dimensions. Combining robust optimization and optimal control ideas to analyze these problems (and others), along with limit-theorem inspired uncertainty sets, could lead to more intuitive, and perhaps closed-form, solutions.

Acknowledgments

The author thanks the associate editor and two reviewers for very insightful comments and suggestions that greatly improved the quality of the paper.

Appendix A. Auxiliary Proofs

Proof of Lemma 2. We rearrange Formulation (18) and apply Expressions (19) to obtain a continuous linear program with a new auxiliary variable $z \in \mathscr{L}^2[\tau, T]$:

$$\begin{aligned}\min\ & \int_\tau^T z(x)\,dx\\ \text{s.t.}\ & z(x) - cq(x) - h\int_\tau^x q(t)\,dt \geqslant h(I(\tau) - \underline{D}_\tau^\Upsilon(x)), \quad \forall x \in [\tau, T],\\ & z(x) - cq(x) + s\int_\tau^x q(t)\,dt \geqslant s(\bar{D}_\tau^\Upsilon(x) - I(\tau)), \quad \forall x \in [\tau, T],\\ & q(x) \geqslant 0, \quad \forall x \in [\tau, T]. \end{aligned} \tag{A.1}$$

We first assume that $c \leqslant s(T-\tau)$, $I(\tau) \leqslant (s\bar{D}_\tau^\Upsilon(T) + h\underline{D}_\tau^\Upsilon(T))/(s+h)$, and $\chi_\tau^\Upsilon \leqslant T - c/s$; these are the necessary conditions for positive ordering. Define $A = T - c/s \in [\tau, T]$ and for clarity we let $\chi = \chi_\tau^\Upsilon$. Consider a feasible solution to Formulation (A.1) characterized by $\int_\tau^x q(t)\,dt = (s\bar{D}_\tau^\Upsilon(x) + h\underline{D}_\tau^\Upsilon(x))/(s+h) - I(\tau)$ for $x \in [\chi, A]$ and $q(x) = 0$ otherwise.

For $x \in [\tau, \chi)$, the auxiliary variable

$$\begin{aligned} z(x) &= \max\{h(I(\tau) - \underline{D}_\tau^\Upsilon(x)), s(\bar{D}_\tau^\Upsilon(x) - I(\tau))\}\\ &= h(I(\tau) - \underline{D}_\tau^\Upsilon(x)), \end{aligned}$$

since $x < \chi$, which implies that $I(\tau) > (s\bar{D}_\tau^\Upsilon(x) + h\underline{D}_\tau^\Upsilon(x))/(s+h)$. For $x \in [\chi, A]$, it is straightforward to see that the first two constraints can both be tight, which implies that the auxiliary variable

$$z(x) = cq(x) + h\left(\frac{s\bar{D}_\tau^\Upsilon(x) + h\underline{D}_\tau^\Upsilon(x)}{s+h} - I(\tau)\right) + h(I(\tau) - \underline{D}_\tau^\Upsilon(x)).$$

Finally, for $x \in (A, T]$, only the second constraint can be tight, since $(s\bar{D}_\tau^\Upsilon(x) + h\underline{D}_\tau^\Upsilon(x))/(s+h) \geqslant (s\bar{D}_\tau^\Upsilon(A) + h\underline{D}_\tau^\Upsilon(A))/(s+h)$, which implies

$$z(x) = s\left(\bar{D}_\tau^\Upsilon(x) - \frac{s\bar{D}_\tau^\Upsilon(A) + h\underline{D}_\tau^\Upsilon(A)}{s+h}\right).$$

The objective function value of (A.1) is

$$\begin{aligned} & h\int_\tau^\chi (I(\tau) - \underline{D}_\tau^\Upsilon(x))\,dx + c\left(\frac{s\bar{D}_\tau^\Upsilon(A) + h\underline{D}_\tau^\Upsilon(A)}{s+h} - I(\tau)\right)\\ &\quad + h\int_\chi^A \left(\frac{s\bar{D}_\tau^\Upsilon(x) + h\underline{D}_\tau^\Upsilon(x)}{s+h} - \underline{D}_\tau^\Upsilon(x)\right) dx\\ &\quad + s\int_A^T \left(\bar{D}_\tau^\Upsilon(x) - \frac{s\bar{D}_\tau^\Upsilon(A) + h\underline{D}_\tau^\Upsilon(A)}{s+h}\right) dx\\ &= h\int_\tau^\chi (I(\tau) - \underline{D}_\tau^\Upsilon(x))\,dx\\ &\quad + h\int_\chi^A \left(\frac{s\bar{D}_\tau^\Upsilon(x) + h\underline{D}_\tau^\Upsilon(x)}{s+h} - \underline{D}_\tau^\Upsilon(x)\right) dx\\ &\quad + s\int_A^T \bar{D}_\tau^\Upsilon(x)\,dx - cI(\tau). \end{aligned} \tag{A.2}$$

The dual of Problem (A.1) is

$$\begin{aligned} \max\ & \left\{\int_\tau^T (h(I(\tau) - \underline{D}_\tau^\Upsilon(x))u(x) + s(\bar{D}_\tau^\Upsilon(x) - I(\tau))v(x))\,dx\right\}\\ \text{s.t.}\ & u(x) + v(x) = 1, \quad \forall x \in [\tau, T],\\ & -cu(x) - h\int_x^T u(t)\,dt - cv(x) + s\int_x^T v(t)\,dt \leqslant 0,\\ & \qquad \forall x \in [\tau, T],\\ & u(x),\ v(x) \geqslant 0, \quad \forall x \in [\tau, T], \end{aligned}$$

which, by substituting out the u variable, can be written as

$$\begin{aligned} \max\ & \left\{\int_\tau^T (s\bar{D}_\tau^\Upsilon(x) + h\underline{D}_\tau^\Upsilon(x) - (s+h)I(\tau))v(x)\,dx\right.\\ & \qquad \left. + h\int_\tau^T (I(\tau) - \underline{D}_\tau^\Upsilon(x))\,dx\right\}\\ \text{s.t.}\ & \int_x^T v(t)dt \leqslant \frac{c + h(T-x)}{s+h}, \quad \forall x \in [\tau, T],\\ & 0 \leqslant v(x) \leqslant 1, \quad \forall x \in [\tau, T]. \end{aligned}$$

Next, consider the dual feasible solution

$$v(t) = \begin{cases} 0, & x \in [\tau, \chi),\\ h/(s+h), & x \in [\chi, A],\\ 1, & x \in (A, T], \end{cases}$$

which has a dual objective value of

$$\begin{aligned}
&\frac{h}{s+h}\int_{\chi}^{A}\big(s\bar{D}_{\tau}^{\Upsilon}(x)+h\underline{D}_{\tau}^{\Upsilon}(x)-(s+h)I(\tau)\big)\,dx\\
&\quad+\int_{A}^{T}\big(s\bar{D}_{\tau}^{\Upsilon}(x)+h\underline{D}_{\tau}^{\Upsilon}(x)-(s+h)I(\tau)\big)\,dx\\
&\quad+h\int_{\tau}^{\chi}(I(\tau)-\underline{D}_{\tau}^{\Upsilon}(x))\,dx\\
&\quad+h\int_{\chi}^{A}(I(\tau)-\underline{D}_{\tau}^{\Upsilon}(x))\,dx+h\int_{A}^{T}(I(\tau)-\underline{D}_{\tau}^{\Upsilon}(x))\,dx\\
&=h\int_{\tau}^{\chi}(I(\tau)-\underline{D}_{\tau}^{\Upsilon}(x))\,dx\\
&\quad+h\int_{\chi}^{A}\left(\frac{s\bar{D}_{\tau}^{\Upsilon}(x)+h\underline{D}_{\tau}^{\Upsilon}(x)}{s+h}-\underline{D}_{\tau}^{\Upsilon}(x)\right)dx\\
&\quad+s\int_{A}^{T}\bar{D}_{\tau}^{\Upsilon}(x)\,dx-cI(\tau). \qquad \text{(A.3)}
\end{aligned}$$

Therefore, the primal value (A.2) and dual value (A.3) are equal, and both solutions are optimal. This concludes the proof of part (1a) of the lemma.

If $\chi_{\tau}^{\Upsilon}>A$, then the proof is modified to exclude the middle case of $[\chi,A]$, which implies that $q^*(x)=0$ for all $x\in[\tau,T]$, which proves part (1b). If $c>s(T-\tau)$, then only the analysis of the $(A,T]$ case is appropriate with $A=\tau$; if $I(\tau)>(s\bar{D}_{\tau}^{\Upsilon}(T)+h\underline{D}_{\tau}^{\Upsilon}(T))/(s+h)$, then only the analysis of the $[\tau,\chi]$ case is appropriate with $\chi=T$. These last observations prove part (2) of the lemma. □

Proof of Lemma 3. The proof is very similar to that of Lemma 2 and we only point out the differences. First, the $x\in[\tau,\chi_{\tau}^{\Upsilon})$ analysis is omitted since the χ_{τ}^{Υ} parameter is only defined for positive inventory positions. Second, the analysis of the $x\in[\chi_{\tau}^{\Upsilon},A]$ case is applicable, except with $\chi_{\tau}^{\Upsilon}=\tau$. The rest of the proof is identical. □

Proof of Theorem 2. We have two cases to consider: (A) $I(\tau)<0$ and (B) $I(\tau)\geqslant 0$. We begin with Case A, where Lemma 3 is applicable with $\Upsilon=\Omega_{\tau}$, and we focus on case (1a) of the lemma (since otherwise it is optimal to order zero): at optimality, for $x\in[\tau,T-c/s]$,

$$\int_{\tau}^{x}q^*(t)\,dt=\frac{s\bar{D}_{\tau}^{\Omega_\tau}(x)+h\underline{D}_{\tau}^{\Omega_\tau}(x)}{s+h}-I(\tau). \qquad \text{(A.4)}$$

Since $\underline{D}_{\tau}^{\Omega_\tau}(\tau)=\bar{D}_{\tau}^{\Omega_\tau}(\tau)=0$, to satisfy Equation (A.4), we require that there is a positive ordering impulse at time τ: $-I(\tau)\delta(x-\tau)$. This impulse order fills the backlogged inventory position $I(\tau)<0$.

Next, given the structure of Ω_τ, we evaluate $\underline{D}_{\tau}^{\Omega_\tau}(x)$ and $\bar{D}_{\tau}^{\Omega_\tau}(x)$, from Equations (19), for $\Upsilon=\Omega_\tau$:

$$\begin{aligned}
\underline{D}_{\tau}^{\Omega_\tau}(x)&=\min_{d\in\Omega_\tau}\int_{\tau}^{x}d(t)\,dt\\
&=\max\{(x-\tau)a,\mu_a(T-\tau)-(T-x)b\},\quad \forall x\in[\tau,T]
\end{aligned}$$

and

$$\begin{aligned}
\bar{D}_{\tau}^{\Omega_\tau}(x)&=\max_{d\in\Omega_\tau}\int_{\tau}^{x}d(t)\,dt\\
&=\min\{(x-\tau)b,\mu_b(T-\tau)-(T-x)a\},\quad \forall x\in[\tau,T].
\end{aligned}$$

We next determine the ordering strategy for $x\in(\tau,T]$. Equation (A.4) can be rewritten as

$$\begin{aligned}
\int_{\tau}^{x}q^*(t)\,dt=\big(&s\min\{(x-\tau)b,\mu_b(T-\tau)-(T-x)a\}\\
&+h\max\{(x-\tau)a,\mu_a(T-\tau)-(T-x)b\}\big)\\
&\cdot(s+h)^{-1}-I(\tau).
\end{aligned}$$

The first term of the max operator dominates when $((b-\mu_a)T+\hat{D}(\tau)-a\tau)/(b-a)\geqslant x$ and the first term of the min operator dominates when $x\leqslant((\mu_b-a)T+b\tau-\hat{D}(\tau))/(b-a)$. The former threshold is strictly less than the latter iff $\hat{D}(\tau)<(\mu_a+\mu_b)T/2-(T-\tau)(a+b)/2$. This implies that we have three cases to consider, which determine the relative order and distinctness of the preceding breakpoints: (i) $\hat{D}(\tau)<(\mu_a+\mu_b)T/2-(T-\tau)(a+b)/2$, (ii) $\hat{D}(\tau)=(\mu_a+\mu_b)T/2-(T-\tau)(a+b)/2$, and (iii) $\hat{D}(\tau)>(\mu_a+\mu_b)T/2-(T-\tau)(a+b)/2$. In each case, we have a partition of $[\tau,T]$ that depends on the value of $\hat{D}(\tau)$; the order and distinctness of the breakpoints depend on the case. We may also omit the inventory position $I(\tau)$ since it is negated by the impulse order discussed above. In case (i) we have that

$$\int_{\tau}^{x}q^*(t)\,dt=\begin{cases}
\left(\dfrac{sb+ha}{s+h}\right)x-\left(\dfrac{sb+ha}{s+h}\right)\tau, & x\in\left(\tau,\dfrac{(b-\mu_a)T+(\mu_a-a)\tau}{b-a}\right],\\[2ex]
bx+\dfrac{h(\mu_a(T-\tau)-bT)-s\tau b}{s+h}, & x\in\left(\dfrac{(b-\mu_a)T+(\mu_a-a)\tau}{b-a},\dfrac{(\mu_b-a)T+(b-\mu_b)\tau}{b-a}\right],\\[2ex]
\left(\dfrac{sa+hb}{s+h}\right)x-\left(\dfrac{sa+hb}{s+h}\right)T+\left(\dfrac{s\mu_b+h\mu_a}{s+h}\right)(T-\tau), & x\in\left(\dfrac{(\mu_b-a)T+(b-\mu_b)\tau}{b-a},T\right],
\end{cases}$$

which, taking derivatives with respect to x and combining with the impulse order at time $x=\tau$, implies

$$q^*(x)=-I(\tau)\delta(x-\tau)+\begin{cases}
\dfrac{sb+ha}{s+h}, & x\in\left(\tau,\dfrac{(b-\mu_a)T+(\mu_a-a)\tau}{b-a}\right],\\[2ex]
b, & x\in\left(\dfrac{(b-\mu_a)T+(\mu_a-a)\tau}{b-a},\dfrac{(\mu_b-a)T+(b-\mu_b)\tau}{b-a}\right],\\[2ex]
\dfrac{sa+hb}{s+h}, & x\in\left(\dfrac{(\mu_b-a)T+(b-\mu_b)\tau}{b-a},T\right].
\end{cases} \qquad \text{(A.5)}$$

In case (iii), we have In case (iii), we have

$$\int_{\tau}^{x}q^*(t)\,dt=\begin{cases}
\left(\dfrac{sb+ha}{s+h}\right)x-\left(\dfrac{sb+ha}{s+h}\right)\tau, & x\in\left(\tau,\dfrac{(\mu_b-a)T+(b-\mu_b)\tau}{b-a}\right],\\[2ex]
ax+\dfrac{-h\tau a+s(\mu_b(T-\tau)-aT)}{s+h}, & x\in\left(\dfrac{(\mu_b-a)T+(b-\mu_b)\tau}{b-a},\dfrac{(b-\mu_a)T+(\mu_a-a)\tau}{b-a}\right],\\[2ex]
\left(\dfrac{sa+hb}{s+h}\right)x-\left(\dfrac{sa+hb}{s+h}\right)T+\left(\dfrac{s\mu_b+h\mu_a}{s+h}\right)(T-\tau), & x\in\left(\dfrac{(b-\mu_a)T+(\mu_a-a)\tau}{b-a},T\right],
\end{cases}$$

which implies

$$q^*(x) = -I(\tau)\delta(x-\tau) + \begin{cases} \dfrac{sb+ha}{s+h}, & x \in \left(\tau, \dfrac{(\mu_b-a)T+(b-\mu_b)\tau}{b-a}\right], \\ a, & x \in \left(\dfrac{(\mu_b-a)T+(b-\mu_b)\tau}{b-a}, \dfrac{(b-\mu_a)T+(\mu_a-a)\tau}{b-a}\right], \\ \dfrac{sa+hb}{s+h}, & x \in \left(\dfrac{(b-\mu_a)T+(\mu_a-a)\tau}{b-a}, T\right]. \end{cases} \quad \text{(A.6)}$$

In case (ii), the middle interval disappears since the thresholds $((b-\mu_a)T+(\mu_a-a)\tau)/(b-a)$ and $((\mu_b-a)T+(b-\mu_b)\tau)/(b-a)$ are equal.

In Case B we apply Lemma 2 with $\Upsilon = \Omega_\tau$. In particular, there is only nonzero ordering when $c \leqslant s(T-\tau)$, $I(\tau) \leqslant (s\bar{D}_\tau^{\Omega_\tau}(T) + h\underline{D}_\tau^{\Omega_\tau}(T))/(s+h)$, and $\chi_\tau^{\Omega_\tau} \leqslant T - c/s$; this nonzero ordering only occurs on the interval $x \in [\chi_\tau^{\Omega_\tau}, T-c/s]$ and the structure is identical to that derived above for Case A, with the impulse order omitted. □

Proof of Theorem 3. The proof follows the same logic as the proof of Theorem 2; the differences are mainly technical. We have two cases to consider: (A) $I(\tau) < 0$ and (B) $I(\tau) \geqslant 0$. We begin with Case A, where Lemma 3 is applicable with $\Upsilon = \Omega_\tau^{\hat{d}}$, and we focus on case (1a) of the lemma (since otherwise it is optimal to order zero): at optimality, for $x \in [\tau, T-c/s]$,

$$\int_\tau^x q^*(t)\,dt = \frac{s\bar{D}_\tau^{\Omega_\tau^{\hat{d}}}(x) + h\underline{D}_\tau^{\Omega_\tau^{\hat{d}}}(x)}{s+h} - I(\tau). \quad \text{(A.7)}$$

Since $\underline{D}_\tau^{\Omega_\tau^{\hat{d}}}(\tau) = \bar{D}_\tau^{\Omega_\tau^{\hat{d}}}(\tau) = 0$, to satisfy Equation (A.7), we require that there is a positive ordering impulse at time τ: $-I(\tau)\delta(x-\tau)$. This impulse order fills the backlogged inventory position $I(\tau) < 0$.

Next, given the structure of $\Omega_\tau^{\hat{d}}$, we evaluate $\underline{D}_\tau^{\Omega_\tau^{\hat{d}}}(x)$ and $\bar{D}_\tau^{\Omega_\tau^{\hat{d}}}(x)$, from Equations (19), for $\Upsilon = \Omega_\tau^{\hat{d}}$:

$$\underline{D}_\tau^{\Omega_\tau^{\hat{d}}}(x) = \min_{d \in \Omega_\tau^{\hat{d}}} \int_\tau^x d(t)\,dt = \max\{(x-\tau)a, \mu_a T - (T-x)b - \hat{D}(\tau)\}, \ \forall\, x \in [\tau, T]$$

and

$$\bar{D}_\tau^{\Omega_\tau^{\hat{d}}}(x) = \max_{d \in \Omega_\tau^{\hat{d}}} \int_\tau^x d(t)\,dt = \min\{(x-\tau)b, \mu_b T - (T-x)a - \hat{D}(\tau)\}, \ \forall\, x \in [\tau, T].$$

We next determine the ordering strategy for $x \in (\tau, T]$. Equation (A.7) can be rewritten as

$$\int_\tau^x q^*(t)\,dt = (s\min\{(x-\tau)b, \mu_b T - (T-x)a - \hat{D}(\tau)\} + h\max\{(x-\tau)a, \mu_a T - (T-x)b - \hat{D}(\tau)\}) \cdot (s+h)^{-1} - I(\tau).$$

The first term of the max operator dominates when $((b-\mu_a)T + \hat{D}(\tau) - a\tau)/(b-a) \geqslant x$ and the first term of the min operator dominates when $x \leqslant ((\mu_b-a)T + b\tau - \hat{D}(\tau))/(b-a)$. The former threshold is strictly less than the latter iff $\hat{D}(\tau) < (\mu_a+\mu_b)T/2 - (T-\tau)(a+b)/2$. This implies that we have three cases to consider, which determine the relative order and distinctness of the preceding breakpoints: (i) $\hat{D}(\tau) < (\mu_a+\mu_b)T/2 - (T-\tau)(a+b)/2$, (ii) $\hat{D}(\tau) = (\mu_a+\mu_b)T/2 - (T-\tau)\cdot(a+b)/2$, and (iii) $\hat{D}(\tau) > (\mu_a+\mu_b)T/2 - (T-\tau)(a+b)/2$. In each case, we have a partition of $[\tau, T]$ that depends on the value of $\hat{D}(\tau)$; the order and distinctness of the breakpoints depend on the case. We may also omit the inventory position $I(\tau)$ since it is negated by the impulse order discussed above. In case (i) we have that

$$\int_\tau^x q^*(t)\,dt = \begin{cases} \left(\dfrac{sb+ha}{s+h}\right)x - \left(\dfrac{sb+ha}{s+h}\right)\tau, \\ \quad x \in \left(\tau, \dfrac{(b-\mu_a)T + \hat{D}(\tau) - a\tau}{b-a}\right] \\ bx + \dfrac{h(\mu_a T - \hat{D}(\tau) - bT) - s\tau b}{s+h}, \\ \quad x \in \left(\dfrac{(b-\mu_a)T + \hat{D}(\tau) - a\tau}{b-a}, \dfrac{(\mu_b-a)T + b\tau - \hat{D}(\tau)}{b-a}\right] \\ \left(\dfrac{sa+hb}{s+h}\right)x - \left(\dfrac{sa+hb}{s+h}\right)T + \left(\dfrac{s\mu_b + h\mu_a}{s+h}\right)T - \hat{D}(\tau), \\ \quad x \in \left(\dfrac{(\mu_b-a)T + b\tau - \hat{D}(\tau)}{b-a}, T\right]. \end{cases}$$

Taking derivatives with respect to x and combining with the impulse order at time $x = \tau$, implies

$$q^*(x) = -I(\tau)\delta(x-\tau) + \begin{cases} \dfrac{sb+ha}{s+h}, & x \in \left(\tau, \dfrac{(b-\mu_a)T+\hat{D}(\tau)-a\tau}{b-a}\right], \\ b, & x \in \left(\dfrac{(b-\mu_a)T+\hat{D}(\tau)-a\tau}{b-a}, \dfrac{(\mu_b-a)T+b\tau-\hat{D}(\tau)}{b-a}\right], \\ \dfrac{sa+hb}{s+h}, & x \in \left(\dfrac{(\mu_b-a)T+b\tau-\hat{D}(\tau)}{b-a}, T\right]. \end{cases} \quad \text{(A.8)}$$

Note that the first and third intervals, $(\tau, ((b-\mu_a)T + \hat{D}(\tau) - a\tau)/(b-a)]$ and $(((\mu_b-a)T + b\tau - \hat{D}(\tau))/(b-a), T]$, respectively, have nonnegative length since $\hat{d} \in P_\tau(\Omega)$ implies that

$$\mu_a T - b(T-\tau) \leqslant \hat{D}(\tau) \leqslant \mu_b T - a(T-\tau);$$

these two inequalities can be rearranged to show the intervals are well defined. The middle interval is well defined because we are in case (i).

In case (iii), a similar analysis gives In case (iii), a similar analysis gives

$$q^*(x) = -I(\tau)\delta(x-\tau) + \begin{cases} \dfrac{sb+ha}{s+h}, & x \in \left(\tau, \dfrac{(\mu_b-a)T+b\tau-\hat{D}(\tau)}{b-a}\right], \\ a, & x \in \left(\dfrac{(\mu_b-a)T+b\tau-\hat{D}(\tau)}{b-a}, ds\,\dfrac{(b-\mu_a)T+\hat{D}(\tau)-a\tau}{b-a}\right], \\ \dfrac{sa+hb}{s+h}, & x \in \left(\dfrac{(b-\mu_a)T+\hat{D}(\tau)-a\tau}{b-a}, T\right]. \end{cases} \quad \text{(A.9)}$$

In case (ii), the middle interval disappears since the thresholds $((b-\mu_a)T + \hat{D}(\tau) - a\tau)/(b-a)$ and $((\mu_b-a)T + b\tau - \hat{D}(\tau))/(b-a)$ are equal.

In Case B we apply Lemma 2 with $\Upsilon = \Omega_\tau^{\hat{d}}$. In particular, there is only nonzero ordering when $c \leqslant s(T-\tau)$, $I(\tau) \leqslant (s\bar{D}_\tau^{\Omega_\tau^{\hat{d}}}(T) + h\underline{D}_\tau^{\Omega_\tau^{\hat{d}}}(T))/(s+h)$, and $\chi_\tau^{\Omega_\tau^{\hat{d}}} \leqslant T - c/s$; this nonzero ordering only occurs on the interval $x \in [\chi_\tau^{\Omega_\tau^{\hat{d}}}, T - c/s]$ and the structure is identical to that derived above for Case A, with the impulse order omitted. □

References

Ardestani-Jaafari A, Delage E (2016) Robust optimization of sums of piecewise linear functions with application to inventory problems. *Oper. Res.* 64(2):474–494.

Asmussen S, Perry D (1998) An operational calculus for matrix-exponential distributions, with applications to a Brownian (q, q) inventory model. *Math. Oper. Res.* 23(1):166–176.

Bandi C, Bertsimas D (2011) Network information theory via robust optimization. Submitted for publication.

Bandi C, Bertsimas D (2012) Tractable stochastic analysis in high dimensions via robust optimization. *Math. Programming* 134(1):23–70.

Bandi C, Bertsimas D (2014a) Optimal design for multi-item auctions: A robust optimization approach. *Math. Oper. Res.*39(4): 1012–1038.

Bandi C, Bertsimas D (2014b) Robust option pricing. *Eur. J. Oper. Res.* 239(3):842–853.

Bandi C, Bertsimas D, Youssef N (2015) Robust queueing theory. *Oper. Res.* 63(3):676–700.

Bandi C, Bertsimas D, Youssef N (2016) Robust transient multi-server queues and feedforward networks. Submitted for publication.

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

Ben-Tal A, Golany B, Nemirovski A (2005) Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing and 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.

Berberian S (1961) *Introduction to Hilbert Space* (Oxford University Press, Oxford, UK).

Bertsimas D, Georghiou A (2015) Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Oper. Res.* 63(3):610–627.

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

Bertsimas D, Thiele A (2006) A robust optimization approach to inventory theory. *Oper. Res.* 54(1):150–168.

Bertsimas D, Gamarnik D, Rikun A (2011) Performance analysis of queueing networks via robust optimization. *Oper. Res.* 59(2): 455–466.

Bertsimas D, Iancu D, Parrilo P (2010) Optimality of affine policies in multi-stage robust optimization. *Math. Oper. Res.* 35(2):363–394.

Bienstock D, Özbay N (2008) Computing robust basestock levels. *Discrete Optim.* 5(2):389–414.

Chen X, Sim M, Sun P (2007) A robust optimization perspective on stochastic programming. *Oper. Res.* 55(6):1058–1071.

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

Federgruen A, Wang M (2015) Inventory models with shelf-age and delay-dependent inventory costs. *Oper. Res.* 63(3):701–715.

Gallego G (1990) Scheduling the production of several items with random demands in a single facility. *Management Sci.* 36(12): 1579–1592.

Gallego G, van Ryzin G (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* 40(8):999–1020.

Georghiou A, Wiesemann W, Kuhn D (2015) Generalized decision rule approximations for stochastic programming via liftings. *Math. Programming* 152(1–2):301–338.

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

Harrison J, Taksar M (1983) Instantaneous control of Brownian motion. *Math. Oper. Res.* 8(3):439–453.

Harrison J, Sellke T, Taylor A (1983) Impulse control of Brownian motion. *Math. Oper. Res.* 8(3):454–466.

Iancu D, Sharma M, Sviridenko M (2013) Supermodularity and affine policies in dynamic robust optimization. *Oper. Res.* 61(4): 941–956.

Jagannathan R, Sen T (1991) Storing crossmatched blood: A perishable inventory model with prior allocation. *Management Sci.* 37(3):251–266.

Liu F, Song J-S (2012) Good and bad news about the (s, t) policy. *Manufacturing and Service Oper. Management* 14(1):42–49.

Luenberger D (1969) *Optimization by Vector Space Methods* (John Wiley & Sons, New York).

Mamani H, Nassiri S, Wagner M (2016) Closed-form solutions for robust inventory management. *Management Sci.* 63(5): 1625–1643.

Plambeck E, Ward A (2006) Optimal control of a high-volume assemble-to-order system. *Math. Oper. Res.* 31(3):453–477.

Plambeck EL (2008) Asymptotically optimal control for an assemble-to-order system with capacitated component production and fixed transport costs. *Oper. Res.* 56(5):1158–1171.

Rosling K (2002) Inventory cost rate functions with nonlinear shortage costs. *Oper. Res.* 50(6):1007–1017.

See C, Sim M (2010) Robust approximation to multi-period inventory management. *Oper. Res.* 58(3):583–594.

Solyalı O, Cordeau J, Laporte G (2016) The impact of modeling on robust inventory management under demand uncertainty. *Management Sci.* 62(4):1188–1201.

Ulbrich M (2009) Optimization methods in Banach spaces. Hinze M, Pinnau R, Ulbrich M, Ulbrich S, eds. *Optimization with PDE Constraints* (Springer, Dordrecht, Netherlands), 97–156.

von Hippel P (2005) Mean, median, and skew: Correcting a textbook rule. *J. Statist. Ed.* 13(2). http://www.amstat.org/publications/jse/v13n2/vonhippel.html.

Wagner M (2010) Fully distribution-free profit maximization: The inventory management case. *Math. Oper. Res.* 35(4):728–741.

Wagner M (2011) Online lot-sizing problems with ordering, holding and shortage costs. *Oper. Res. Lett.* 39(2):144–149.

Wein L (1992) Dynamic scheduling of a multiclass make-to-stock queue. *Oper. Res.* 40(4):724–735.

Wu J, Chao X (2014) Optimal control of a Brownian production/inventory system with average cost criterion. *Math. Oper. Res.* 39(1):163–189.

Zipkin P (2000) *Foundations of Inventory Management* (McGraw-Hill/Irwin, New York).

Michael R. Wagner is an assistant professor of operations management at the Michael G. Foster School of Business at the University of Washington. His research interests are in decision making under uncertainty (stochastic, online, and robust optimization, as well as hybrids thereof). In particular, he investigates the value of information in a variety of application domains, such as supply chain management, inventory control, resource allocation, and logistics.