

Submitted to
manuscript

The Newsvendor under Demand Ambiguity: Combining Data with Moment and Tail Information

Soroush Saghafian¹, Brian Tomlin²

¹Industrial Eng., School of Computing, Informatics and Decision Systems Engineering, Arizona State University, Tempe, AZ
soroush.saghafian@asu.edu

²Tuck School of Business at Dartmouth, 100 Tuck Hall, Hanover, NH
brian.tomlin@tuck.dartmouth.edu

Operations managers do not typically have full information about the demand distribution. Recognizing this, data-driven approaches have been proposed in which the manager has no information beyond the evolving history of demand observations. In practice, managers often have some partial information about the demand distribution in addition to demand observations. We consider a repeated newsvendor setting, and propose a non-parametric, maximum-entropy based technique, termed SOBME (Second Order Belief Maximum Entropy), which allows the manager to effectively combine demand observations with distributional information in the form of bounds on the moments or tails. In the proposed approach, the decision maker forms a belief about possible demand distributions, and dynamically updates it over time using the available data and the partial distributional information. We derive a closed-form solution for the updating mechanism, and highlight that it generalizes the traditional Bayesian mechanism with an exponential modifier that accommodates partial distributional information. We prove the proposed approach is (weakly) consistent under some technical regularity conditions and we analytically characterize its rate of convergence. We provide an analytical upper bound for the newsvendor's cost of ambiguity, i.e., the extra per-period cost incurred due to ambiguity, under SOBME, and show that it approaches zero quite quickly. Numerical experiments demonstrate that SOBME performs very well. We find that it can be very beneficial to incorporate partial distributional information when deciding stocking quantities, and that information in the form of tighter moment bounds is typically more valuable than information in the form of tighter ambiguity sets. Moreover, unlike pure data-driven approaches, SOBME is fairly robust to the newsvendor quantile. Our results also show that SOBME quickly detects and responds to hidden changes in the unknown true distribution. We also extend our analysis to consider ambiguity aversion, and develop theoretical and numerical results for the ambiguity-averse, repeated newsvendor setting.

Key words: Data-Driven Newsvendor; Maximum Entropy; Distribution Ambiguity, Moment Information, Ambiguity Aversion

History: Last Revision: April 25, 2015

1. Introduction

Randomness in demand and/or supply is a major challenge in sales and operations planning. A common paradigm in the operations literature endows the decision maker with full distributional information about the underlying random variable(s). Full distributional information is, however, a strong assumption and is not typically met in practice. An alternative “pure data-driven” paradigm endows the decision maker only with past observations/realizations of the random variable (the “data”); see, e.g., Azoury (1985), Liyanage and Shanthikumar (2005), Levi et al. (2007, 2010), Huh et al. (2011). Reality often lies somewhere in between these two paradigms: the decision maker typically has some partial distributional information and some data.

In this paper, we propose and analyze a methodology for decision making in settings in which the decision maker does not know the true distribution but has access to some partial distributional information and some limited data. This form of *ambiguity*, a.k.a *Knightian uncertainty*, about the true distribution arises in various settings in operations including new product launches and new process introductions. In such settings, firms devote significant efforts to characterizing demand or yield before the new product or process is launched and then observe a growing time series of realizations, i.e., data, after launch. In new product introductions, forecasting demand and choosing stocking quantities is a common challenge across a wide array of industries (Kahn 2002, AMR 2008), especially “during the commercialization stage (pre-launch preparation and launch) where new product forecasts drive a variety of ... manufacturing decisions,” (Kahn 2002, p. 133). For example, consider the consumer-products company Nabisco (now owned by Mondelez International):

Several statistical techniques are used to forecast sales. However, most of these techniques need a minimum of two to three years of historical information to build a robust model. This presents a big problem when developing forecasts for newly introduced products. Before a new product is rolled out, [the] planning forecast developed by marketing research ... is used for ... production planning ... and inventory deployment. However, when the product is rolled out, actual sales can be substantially less or more than the planned numbers [leading to] lost sales [or] excess inventory. (Amrute 1998, p. 7).

Let us consider the kind of demand-forecast information that firms generate before and after a new product is launched. According to surveys by Lynn et al. (1999) and others, pre-launch forecasts are primarily generated through marketing research and judgmental techniques, with the top four approaches being customer/market research, followed by jury of executive opinion, sales force composite, and look-like analysis (Kahn 2002). Customer/market research encompasses a variety of approaches, including surveys of potential buyers. The jury of executive opinion approach solicits and combines judgments from managers representing various functions, whereas the sales force composite approach solicits and combines judgments from the sales force (Aaker et al. 2010) who are typically better aware of market conditions; see, e.g., Saghafian and Chao (2014), and the references therein. In look-like analysis, the firm identifies “a similar product, sold previously, and extrapolates from that product’s sales” (Fisher and Raman 2010, p. 85). These judgmental approaches, such as sales force composite, can be used to estimate demand uncertainty by having the salespeople “provide a sales estimate range in addition to a single number representing the expected sales level [or] to have them estimate the probability of each potential sales [level]” (Aaker

et al. 2010, p. 747). Due to the subjective nature of these estimation processes, different people will provide different forecasts, i.e., different probability distributions. As such, these pre-launch techniques enable the firm to generate a set of possible demand distributions and a range for the expected sales.

After a product has been launched, the firm can use actual demand observations to update its beliefs about demand. However, market research efforts do not end, at least in part because “the rate of new product introduction is fast and frequent, meaning data series are often short” (Lynn et al. 1999, p.566). Ongoing market research generates forecast information through test sales (Fisher and Raman 2010), simulated test markets (Lynn et al. 1999), and surveys of customer purchase intentions (Aaker et al. 2010). These ongoing post-launch market-research efforts provide valuable demand information separate from the actual demand observations.

Therefore, in planning stocking quantities for a new or recently-released product, the operations function has two separate and evolving sets of information about the unknown demand distribution. One set is the pre- and post-launch market research including the judgmental approaches discussed above. The other set is the actual demand observations over time. This information setting is not captured by the traditional “full-information” paradigm in operations management that endows the decision maker with complete knowledge of the demand distribution. Neither it is captured by the more recent “pure data-driven” approaches that were specifically designed for settings in which the decision maker has access to demand observations but no other information.

We develop an approach tailored to intermediate information settings – such as new product and new process introductions – in which a decision maker operates in an environment of ambiguity characterized by two evolving sources of information about a relevant random variable (e.g., demand): “exogenous research” and observations (the “data”). Based on the above discussion of pre- and post-launch forecasts, we formulate the exogenous research information as a set of possible demand distributions as well as potential upper and/or lower bounds on the moments and/or tails of the demand distribution. Our proposed approach is a second-order, non-parametric, maximum-entropy based approach, which we term SOBME (Second Order Belief Maximum Entropy). It allows the decision maker to effectively combine the above-mentioned sources of information dynamically over time to update its beliefs about the possible distributions governing the demand realizations. In our approach, the decision maker updates its second-order beliefs, i.e., distribution over possible distributions, so as to minimize the relative entropy to its prior belief subject to the most recently observed demand and the current market research information. In doing so, we

follow the core concept of maximum entropy, widely used in estimation and information theory, in which (i) beliefs are updated so that the posterior coincides with the prior as closely as possible, and (ii) only those aspects of beliefs for which new evidence was gained are updated (Jaynes 1957, 1981, Cover and Thomas 1991).

We embed our SOBME approach in a classical operations setting – the repeated newsvendor – and analytically and numerically examine its performance. We focus on a repeated newsvendor setting because that is a common paradigm in the operations literature for exploring approaches that relax the assumption of full demand information. However, SOBME is general and could be used in any setting in which a decision maker does not know the distribution of an underlying random variable but has access to data and to some partial distributional information. Such settings will become increasingly common in the era of big data analytics. Our paper provides both theoretical and managerially-relevant results.

On the theoretical side, we derive a closed-form solution for the SOBME belief-updating mechanism, and show that it contains traditional Bayesian updating – both parametric and nonparametric – as a special case. As we will show, SOBME generalizes the Bayesian updating mechanism by incorporating the partial distributional information through an *exponential modifier*. We explore the asymptotic behavior of SOBME and show that it is weakly consistent. That is, under some technical conditions, the updated distribution converges to the true unknown distribution almost surely in the weak neighborhoods of it. We also analytically establish a rate of convergence for SOBME. Furthermore, we provide a performance guarantee for SOBME by developing an upper bound on the *newsvendor's cost of ambiguity*, i.e., the increase in the expected per-period newsvendor's cost as compared to the full distributional knowledge case. We show that the cost of ambiguity approaches zero quite quickly when SOBME is used. We also explore an ambiguity-averse version of our repeated newsvendor problem, and develop theoretical results that characterize the optimal newsvendor quantity under ambiguity aversion. In doing so, we establish a generalization of the traditional newsvendor quantile result. We also analytically and numerically examine how SOBME performs in this ambiguity-averse setting.

On the managerial side, we numerically establish that SOBME performs well – as measured by the percentage increase in cost as compared to full distributional knowledge – even when the number of demand observations is limited. This is of practical importance because, as noted above, firms often do not have access to a long series of demand observations when planning stocking quantities for recently-introduced products. Given the efforts devoted by companies to generate exogenous

research about the demand distribution, it is important to understand whether such research is valuable, when it is valuable, and what type of research is more valuable. We numerically explore the value of exogenous research by comparing the performance of SOBME with that of pure data-driven approaches, such as Sample Average Approximation (SAA), which rely exclusively on the demand realizations. We show that there is significant value to exogenous research, especially when demand observations are limited or the newsvendor quantile (the desired service level) is high. Unlike SAA, the SOBME approach performs well even at high service levels. We find that research resulting in tighter moment bounds is more valuable than research resulting in a more compact set of possible distributions. As discussed above, “like-products” are one source of exogenous research. Based on a new-product introduction case motivated by a real-world setting (Allon and Van Mieghem 2011), we show how moment bounds (for demand) can be developed from like-products, and numerically show that there is significant benefit to moment bound information even if the number of like-products is quite small. We also numerically show that SOBME is effective at detecting and quickly reacting to changes in the underlying unknown distribution, an important and practical benefit in dynamic environments. We prove that ambiguity aversion leads to lower order quantities. This has important implications for practice if individual managers vary in their ambiguity aversion.

The rest of the paper is organized as follows. The literature is reviewed in §2. The model is presented in §3, and the proposed SOBME approach is developed in §4. A performance guarantee for the repeated newsvendor setting is developed in §5. Numerical results are presented in §6. Ambiguity aversion is considered in §7. §8 concludes the paper.

2. Literature Review

Despite some early and notable exceptions, e.g., Scarf (1958, 1959), the operations literature has typically assumed that the decision maker has full information about the distribution of relevant random variables. A number of approaches have been adopted to relax this strong assumption.

Some papers adopt a parametric approach, in which the distribution comes from a parametric family but the parameter value is unknown. Bayesian approaches have been used in the context of demand uncertainty (Scarf 1959, Azoury 1985, Lariviere and Porteus 1999, Wang and Mersereau 2013) and supply uncertainty (Tomlin 2009, Chen et al. 2010). An “operational statistics” approach for simultaneously estimating an unknown parameter and optimizing an operational planning problem is introduced in Liyanage and Shanthikumar (2005) and further explored in Chu et al. (2008) and Ramamurthy et al. (2012).

The distribution of interest may not come from a parametric family or the decision maker may

not know which family it comes from. With this in mind, a number of different non-parametric approaches have been proposed and analyzed in various inventory planning contexts. Sample Average Approximation (SAA) (see, e.g., Kleywegt et al. 2002, Shapiro 2003, Levi et al. 2007, 2010) is a data-driven approach that solves a sample-based counterpart to the actual problem. For the SAA approach, Levi et al. (2007) and Levi et al. (2010) develop bounds on the number of samples required for the expected cost to come arbitrarily close (with a high confidence) to the expected cost if the true distribution was known. Other data-driven, non-parametric approaches include the bootstrap method (Bookbinder and Lordahl 1989), the CAVE algorithm (Godfrey and Powell 2001), and adaptive inventory methods mainly used for censored demand data (Huh and Rusmevichientong 2009, Huh et al. 2011).

An alternative non-parametric approach is a robust one in which certain moments of the distribution are known, and the decision maker maximizes (minimizes) the worst-case expected profit (cost) over all distributions with the known attribute; see, e.g., Scarf (1959), Gallego and Moon (1993), Popescu (2007). Because such approaches can be conservative, Perakis and Roels (2008) adopt a regret-based approach in which the decision maker knows the mean. Zhu et al. (2013) consider a newsvendor setting where the demand distribution is only specified by its mean and either its standard deviation or its support, and develop a method that minimizes the ratio of the expected cost to that of full information. Moments, however, might not be known with certainty, and Delage and Ye (2010) develop a robust approach in which the mean and covariance satisfy some known bounds. Wang et al. (2013) devise a likelihood robust optimization (LRO) approach, in which the set of possible distributions are considered to be only those that make the observed data achieve a certain level of likelihood. It is then assumed that the decisions are made based on the worst-case distribution within this set.

A long-established approach when selecting a probability distribution subject to known moments is to choose the distribution that has the maximum entropy (ME). Despite its widespread use in fields such as estimation theory, physics, statistical mechanics, and information theory among others, there are only a few papers in the operations literature that have adopted maximum entropy based approaches. Andersson et al. (2013) examine a non-repeating newsvendor model in which the decision maker only knows the mean and variance of the demand distribution. They compare an ME approach with robust approaches and find that the ME approach performs better on average. Different from Andersson et al. (2013), (i) we examine a repeated newsvendor setting and so the decision maker has access to past data to dynamically update his beliefs; and (ii) we allow for

(evolving) upper/lower bounds on moments and tails instead of assuming a known mean and variance. Maglaras and Eren (2015) apply the ME approach in a revenue management context under a discrete demand assumption and static linear equalities, and establish its asymptotical behavior. Unlike their work, we (i) consider a newsvendor setting, (ii) permit continuously distributed demand, (iii) allow for dynamically changing inequalities (lower and upper bounds on the moments and tails that may change over time), and (iv) build distributions over possible distributions as opposed to the random variable itself, i.e., a second-order approach as opposed to a first-order one.

The concept of maximizing entropy when updating probability distributions in response to new information, under the strong assumption that the first moment is exactly known, has been shown to include Bayesian updating as a special case; see Caticha and Giffin (2006) and Giffin and Caticha (2007). We adopt and develop this approach by (i) creating a mechanism that builds a belief/distribution about the possible distributions of the underlying random parameter (as opposed to directly building a distribution about it); (ii) allowing for partial distributional information in the form of arbitrary upper and/lower bounds on the moments and tails, which may evolve over time (as opposed to the very special case where the first moment is perfectly known and this information is static). Importantly, we also embed it in a sequential decision-making setting (a repetitive newsvendor), provide a performance guarantee for it, and analytically establish its asymptotic behavior including its consistency and rate of convergence. To the best of our knowledge, our proposed method is among the very first in the operations literature that is both data-driven and information-driven and also fully dynamic. The approach provides a method to effectively combine dynamically evolving partial information on moments and tails with data and observations, presenting a widely applicable tool (especially in the current trend of business analytics) for decision making under distributional ambiguity.

3. The Model

We consider a repeated newsvendor problem over periods $t \in \mathcal{T} := \{1, \dots, T\}$, with $T \leq \infty$. Let $h > 0$ and $p > 0$ denote the per-unit holding and shortage penalty costs, respectively. Let $\mathcal{L}(x) = h[x]^+ + p[-x]^+$, where $[x]^+ := \max\{x, 0\}$. Demand realizations are generated as i.i.d. random variables from a true distribution that is unknown to the decision maker (DM). Unless otherwise stated (e.g., §6.3), we assume the true distribution is stationary. The DM, “he”, as a convention hereafter, has two types of evolving information about the demand distribution.

- **Data/Observations.** At the end of period $t = 1, \dots, T$, the DM knows the realization of demand D_i in each of the periods $i = 1, \dots, t$. This information is denoted as \mathcal{F}_t . We define $\mathcal{F}_0 = \emptyset$ to reflect

the fact that there is no demand realization prior to period 1.

- **Exogenous Research.** This information is contained in two sets.

— *Ambiguity Set.* As noted in §1 for the example of new product introductions, pre-launch marketing research generates a set of possible demand distributions based on the judgments of various relevant constituents, e.g., managers, sales people, prospective customers. Suppose the sets of possible demand distributions and associated densities (assuming they exist) are given by the *ambiguity sets* $\mathcal{D} := \{F_D^\omega : \omega \in \Omega\}$ and $\mathcal{D}' := \{f_D^\omega : \omega \in \Omega\}$, respectively, where w is a scalar indicator associated with a distribution, and Ω is an arbitrary, typically uncountable, set. Without loss of generality, we assume all the distributions in \mathcal{D} have a common support denoted by \mathcal{D} , e.g., $\mathcal{D} = [0, \infty)$. We assume that the ambiguity set \mathcal{D} contains the true distribution $F_D^{\omega^*}$, where $\omega^* \in \Omega$ is simply an indicator for the true distribution.

We make no parametric assumptions about the distributions contained in the ambiguity set. The set may contain any parametric and/or non-parametric distributions. The fact that the indicator ω is a scalar is not a limitation. For example, parametric families with more than one parameter can be allowed in our approach. It is simply a matter of associating a scalar ω with an instance of the family. Indeed, our non-parametric approach can allow multiple distributional families to be in the ambiguity; see our numerical study in §6.

— *Moment and Tail Bounds.* As noted in §1, pre- and post-launch marketing research can also provide the DM with additional demand information in the form of bounds on (i) the expected value of demand, and (ii) tail probabilities. At period $t \in \mathcal{T}$, we denote the lower and upper bounds on the expected value of demand by $\underline{\mathbb{E}}_t(D)$ and $\overline{\mathbb{E}}_t(D)$, respectively, with $0 \leq \underline{\mathbb{E}}_t(D) \leq \overline{\mathbb{E}}_t(D)$.¹ We let the tail related information at period t be of the form $Pr(D \geq u) \leq b_t(u)$, for some non-increasing function $b_t(u) : \mathcal{D} \rightarrow [0, 1]$. We denote this moment and tail information by the set $\hat{\mathcal{F}}_t$, and note that allowing it to change over time is a generalization of the static case in which the moment bounds and the tail information are fixed over the horizon. In this latter case, $\hat{\mathcal{F}}_t = \hat{\mathcal{F}}_0$ for all $t \in \mathcal{T}$, where $\hat{\mathcal{F}}_0$ is the initial information before the first time period has occurred. We note that the moment and tail information need to be consistent, i.e., should not contradict each other. Otherwise, there may not be any distribution in the ambiguity set that satisfies them, in which case the problem we discuss below is infeasible. The consistency between the data/observations and the moment/tail bounds can be monitored and statistically tested. Our approach can allow the decision maker to also include endogenously-generated moment bounds based on the data/observations; see §6.3.

¹ For clarity, we only consider the first moment but the approach and results are readily extended to higher moments. Higher moment information, if available, can also improve the convergence rate of our estimation approach.

Because the DM does not know the true demand distribution, his decision in each period will be based on his current belief about the distribution. We will introduce and develop our approach to belief representation and updating in the next section. Before doing so, we specify the sequence of events in each period $t \in \mathcal{T}$:

1. The DM determines an order quantity $q_t \geq 0$ (given his current belief about the demand distribution), orders it, and receives it.
2. Demand is realized, with the realized value denoted by ξ'_t , and the newsvendor cost $\mathcal{L}(q_t - \xi'_t)$ is incurred.
3. New exogenous research, if any, arrives regarding moment and tail bounds, and the information set $\hat{\mathcal{F}}_t$ is specified. Note that $\hat{\mathcal{F}}_t = \hat{\mathcal{F}}_{t-1}$ if no new research arrives. Also, $\hat{\mathcal{F}}_t = \hat{\mathcal{F}}_0$ for all $t \in \mathcal{T}$ if exogenous research is fixed throughout the horizon.
4. The DM updates his belief about the demand distribution for period $t + 1$.

When determining his order quantity at the beginning of period t (step 1), the DM solves:

$$\min_{q_t \geq 0} \mathbb{E}_{D|\mathcal{F}_{t-1}, \hat{\mathcal{F}}_{t-1}} [\mathcal{L}(q_t - D)], \quad (1)$$

where $\mathbb{E}_{D|\mathcal{F}_{t-1}, \hat{\mathcal{F}}_{t-1}}$ denotes that the expectation is taken with respect to the DM's current belief about the demand distribution, and that this current belief (however it was formed) can only avail of the information available at that point in time, i.e., the prior demand realizations \mathcal{F}_{t-1} and the most recent moment and tail information set $\hat{\mathcal{F}}_{t-1}$. We refer to the above optimization problem as the *DM's problem* and to its adopted policy as the *DM's policy*. The values of demand, however, are drawn from a hidden underlying distribution. Therefore, the DM's policy is associated with a *true cost* that depends on the true hidden underlying distribution. This true cost represents the ultimate measure of the policy's performance.

4. A Belief Updating Approach that Combines Data and Moment/Tail Bounds

In this section, we first sketch the outline of our approach to representing and updating beliefs, and contrast it with other existing approaches (§4.1). Next, in §4.2, we formally develop the approach and present the key belief-updating result. Finally, in §4.3, we analytically examine the asymptotic behavior of the approach by investigating its consistency and convergence rate.

4.1. Outline of the Approach

We adopt a distribution-over-distributions approach. At the beginning of period t , the DM's belief about the demand distribution is represented by a distribution over the set of possible demand

distributions contained in the ambiguity set. In essence, the DM assigns a probability that any particular distribution in the ambiguity set is the true distribution. At the end of period t , having observed the value of period- t demand, the DM updates his belief about the probability of each distribution in the ambiguity set; that is, a new distribution (over the distributions in the ambiguity set) is formed. The DM adopts the maximum entropy principle when creating this new distribution: beliefs are updated so that (i) the posterior coincides with the prior as closely as possible, and (ii) only those aspects of beliefs for which new evidence was gained are updated. Furthermore, the updating is done in a manner that ensures conformance to the current moment and tail bound information. These notions are mathematically formalized in §4.2.

We call this a *Second Order Belief Maximum Entropy* (SOBME) approach because the DM's beliefs are represented as a distribution over possible demand distributions rather than being codified directly as a demand distribution (a first-order approach). We note that if the DM is ambiguity neutral, i.e., the DM's period- t objective function is given by (1), then there exists a single first-order demand distribution that would yield an equivalent period- t objective value to that given by a distribution-over-distributions approach. However, there is a question as to how such a first-order demand distribution would be developed. SOBME provides a useful and practically relevant method to dynamically characterize the first-order distribution as a weighted average of the distributions in the ambiguity set. From a practical perspective, our SOBME approach aligns directly with new-product forecasting judgmental methods (see §1) in which different people estimate the probability of each sales level. Such methods yield a set of distributions, one from each person. SOBME provides a method to dynamically update the probability associated with each distribution and, as we will see in subsequent sections, it performs very well. Furthermore, as we will see in §7, a distribution-over-distributions approach is necessary if the DM is ambiguity averse/loving, and SOBME provides a good approach for such settings.

Before we formalize the SOBME approach, we contrast it conceptually with pure data-driven approaches and robust optimization approaches. In pure data-driven approaches, SAA for example, the DM constructs (in each time period) a new demand distribution based exclusively on the history of demand observations. SOBME differs from SAA along two important dimensions: (i) it enables the DM to avail of moment and tail bound information in addition to the demand observations, and (ii) the DM's beliefs are represented as a distribution over distributions rather than as a single distribution. SOBME is also very different from robust optimization approaches in which decisions are made only with respect to the worst-case distribution in the ambiguity set. In SOBME, the

DM assigns weights to all members of the ambiguity set. This helps avoid the tendency of robust optimization methods to make overly conservative decisions. Moreover, our approach updates these weights based on new realizations and evolving moment and tail information.

4.2. Formal Development of the SOBME Approach

At the beginning of period $t \in \mathcal{T}$, the DM's belief about the demand distribution is represented as a density function $f_t(\omega)$ over the set of possible demand distributions contained in the ambiguity set $\mathcal{D} := \{F_D^\omega : \omega \in \Omega\}$.² At the end of period t , having observed the period- t demand and knowing the most recent moment and tail information, the DM updates his belief from $f_t(\omega)$ to $f_{t+1}(\omega)$. In what follows, we formally develop the approach for updating $f_t(\omega)$ to $f_{t+1}(\omega)$.

The updating mechanism (at the end of period t) is based on the following three inputs. One input is the joint density (from the beginning of period t) over the demand and the demand distribution indicator; this is given by $f_t(\xi, \omega) = f_t(\omega)f_D^\omega(\xi)$, where $f_D^\omega(\xi)$ is the demand density associated with the indicator ω . The second input is the observed value of the demand in period t , denoted by ξ'_t . The third input is the most recent moment and tail bound information, i.e., $\hat{\mathcal{F}}_t$.

Let $f_{t+1}^o(\xi, \omega)$ denote some joint density (at the end of period t) over the demand *observation* that was realized in period t and the demand distribution indicator random variable. The output of our updating mechanism, defined through the optimization program described below, will be a specific joint density denoted by $f_{t+1}^{*o}(\xi, \omega)$. This will determine the DM's updated belief $f_{t+1}(\omega)$ using $f_{t+1}(\omega) = \int_{\mathcal{D}} f_{t+1}^{*o}(\xi, \omega) d\xi$, i.e., the DM's updated belief $f_{t+1}(\omega)$ is given by the marginal density of $f_{t+1}^{*o}(\xi, \omega)$.

The updating mechanism chooses a joint density $f_{t+1}^o(\xi, \omega)$ [the ‘‘posterior’’] to minimize the Kullback-Leibler (KL) divergence (a.k.a. relative entropy) between $f_{t+1}^o(\xi, \omega)$ and $f_t(\xi, \omega)$ [the ‘‘prior’’] while recognizing that the demand *observation* in period t took on some specific value, ξ'_t , and that the joint density $f_{t+1}^o(\xi, \omega)$ must imply a demand density that conforms to the most recent moment and tail bound information, i.e., $\hat{\mathcal{F}}_t$. The Kullback-Leibler (KL) divergence is formally defined as:

DEFINITION 1 (Kullback-Leibler Divergence). The Kullback-Leibler (KL) divergence or the relative entropy between two joint densities f_{t+1}^o and f_t defined on $\mathcal{D} \times \Omega$ is:

² At the beginning of period $t = 1$, the DM chooses an initial density $f_1(\omega)$ to reflect his starting belief about the probability of each demand distribution and his initial moment/tail bound information. A uniform density over all $\omega \in \Omega$ (representing a maximum initial entropy) would be appropriate if the DM did not believe any particular distribution was more likely to be the true one.

$$d_{KL}(f_{t+1}^o || f_t) = \int_{\mathcal{D}} \int_{\Omega} f_{t+1}^o(\xi, \omega) \log \frac{f_{t+1}^o(\xi, \omega)}{f_t(\xi, \omega)} d\omega d\xi = \mathbb{E}_{f_{t+1}^o} \left[\log \frac{f_{t+1}^o(\xi, \omega)}{f_t(\xi, \omega)} \right].$$

This is a general definition for any two densities, but described in terms of our specific notation.³

Observe that $d_{KL}(f_{t+1}^o || f_t) \geq 0$, and $d_{KL}(f_{t+1}^o || f_t)$ is convex in (f_{t+1}^o, f_t) .

The updating mechanism is given by the following functional optimization program, in which the decision variable is the function $f_{t+1}^o : \mathcal{D} \times \Omega \rightarrow \mathbb{R}_+$, and is chosen to minimize the relative entropy (or equivalently maximize the negative relative entropy).

$$\min_{f_{t+1}^o : \mathcal{D} \times \Omega \rightarrow \mathbb{R}_+} d_{KL}(f_{t+1}^o || f_t) \quad (2)$$

s.t.

$$\mathbb{E}_t(D) \leq \int_{\mathcal{D}} \int_{\Omega} f_{t+1}^o(\xi, \omega) \psi_1(\omega) d\omega d\xi \leq \bar{\mathbb{E}}_t(D) \quad (3)$$

$$\int_{\mathcal{D}} \int_{\Omega} f_{t+1}^o(\xi, \omega) \psi_2(\omega, u) d\omega d\xi \leq b_t(u) \quad \forall u \in \mathcal{D} \quad (4)$$

$$\int_{\Omega} f_{t+1}^o(\xi, \omega) d\omega = \delta(\xi - \xi_t^l) \quad \forall \xi \in \mathcal{D} \quad (5)$$

$$\int_{\mathcal{D}} \int_{\Omega} f_{t+1}^o(\xi, \omega) d\omega d\xi = 1, \quad (6)$$

where $\psi_1(\omega) := \mathbb{E}(D|\omega) = \int_{\mathcal{D}} \xi dF_D^\omega(\xi)$, $\psi_2(\omega, u) := 1 - F_D^\omega(u)$, and $\delta(\cdot)$ is the Dirac delta function.⁴

It should be noted that constraints (4) and (5) are not single constraints: they each represent an infinite number of constraints.

We first discuss the role of constraints (3) and (4). Note that $\int_{\mathcal{D}} f_{t+1}^o(\xi, \omega) d\xi$ gives the marginal density of the demand distribution indicator ω . Therefore, constraint (3) reflects the fact that the DM's updated belief must result in a mean demand that conforms to the most recent upper and lower moment bounds, $\bar{\mathbb{E}}_t(D)$ and $\underline{\mathbb{E}}_t(D)$, respectively. Similarly, constraint (4) reflects the fact the DM's updated belief must result in a demand distribution whose tail probability (at any given u) is bounded above by the most recent tail bound $b_t(u)$.

Next we discuss the role of constraints (5) and (6). Recall that the updating optimization program is selecting a joint density $f_{t+1}^o(\xi, \omega)$ over the demand *observation* that was realized in period t and the demand distribution indicator random variable. The marginal density $\int_{\Omega} f_{t+1}^o(\xi, \omega) d\omega$ in (5) is the density over the possible demand observations. Of course, the realized value of the observation

³ For the more general class of ϕ -divergence, we refer interested readers to Pardo (2006), Ben-Tal et al. (2013), and the references therein. We also note that $d_{KL}(\cdot || \cdot)$ is not a metric in the usual sense as it is not symmetric.

⁴ The Dirac delta function, sometimes called the unit impulse function, is a generalized function (often) defined such that $\delta(x) = \infty$ if $x = 0$ and $\delta(x) = 0$ otherwise, with $\int_{-\infty}^{\infty} \delta(x) dx = 1$ (Kreyszig 1988). It can be defined more rigorously either as a measure or a distribution. It can be viewed as the limit of the Gaussian function: $\delta(x) = \lim_{\sigma \rightarrow 0} \frac{e^{-x^2/2\sigma^2}}{\sqrt{2\pi}\sigma}$. Importantly, the Dirac delta function has the property that for any function $g(\cdot)$, $\int_{-\infty}^{+\infty} g(x)\delta(x-x_0)dx = g(x_0)$.

is already known to be ξ'_t . Therefore, any valid marginal density (valid, in the sense that it reflects what the DM already knows) should place all its weight on $\xi = \xi'_t$ and no weight on $\xi \neq \xi'_t$. However, it also needs to be a legitimate density function. These two requirements are fulfilled by using the Dirac delta function (see footnote 4) and setting $\int_{\Omega} f_{t+1}^o(\xi, \omega) d\omega = \delta(\xi - \xi'_t) \forall \xi \in \mathcal{D}$, which is constraint (5). Finally, constraint (6) is a normalization that ensures $f_{t+1}^o(\xi, \omega)$ is a joint density in the sense that it integrates (over all possible values of both arguments) to 1. We note this constraint is in fact redundant given the property that $\int_{-\infty}^{\infty} \delta(x) dx = 1$ but we include (6) for expositional clarity.

This updating program is an infinite-dimensional optimization program that has a continuously differentiable convex objective function and linear constraints. This allows us to use KKT-like conditions along with the variational principle to characterize the global optimum solution, $f_{t+1}^{*o}(\xi, \omega)$. Let $\bar{\beta}_t$ and $\underline{\beta}_t$ be the Lagrangian multipliers for the upper bound and lower bound in constraint (3), respectively. Let $\gamma_t(u)$ and $\lambda_t(\xi)$ represent the multipliers for (the infinite number of) constraints (4) and (5), respectively. Defining $g_t(\omega) := (\bar{\beta}_t - \underline{\beta}_t)\psi_1(\omega) + \int_{\mathcal{D}} \gamma_t(u)\psi_2(\omega, u)du$, we are able to obtain the following closed-form result.

THEOREM 1 (Second Order Joint Belief Updating). *The solution to the optimization program (2)-(6) is given by:*

$$f_{t+1}^{*o}(\xi, \omega) = \frac{f_t(\xi, \omega)\delta(\xi - \xi'_t)e^{-g_t(\omega)}}{\int_{\Omega} f_t(\xi'_t, \tilde{\omega})e^{-g_t(\tilde{\omega})}d\tilde{\omega}}.$$

PROOF. All proofs are contained in Appendix B.

COROLLARY 1 (Second Order Belief Updating). *The updated density, $f_{t+1}(\omega)$, over the demand distribution indicator ($\omega \in \Omega$) is:*

$$f_{t+1}(\omega) = \frac{f_t(\omega)f_t(\xi'_t|\omega)e^{-g_t(\omega)}}{\int_{\Omega} f_t(\tilde{\omega})f_t(\xi'_t|\tilde{\omega})e^{-g_t(\tilde{\omega})}d\tilde{\omega}},$$

where $f_t(\xi'_t|\omega) = f_D^{\omega}(\xi'_t)$.

REMARK 1 (Bayesian Updating). When the moment and tail bounds, i.e., constraints (3) and (4), are not binding in the updating optimization program, the updating result simplifies to $f_{t+1}(\omega) = \frac{f_t(\omega)f_t(\xi'_t|\omega)}{\int_{\Omega} f_t(\tilde{\omega})f_t(\xi'_t|\tilde{\omega})d\tilde{\omega}}$. That is, the updating result reduces to a traditional Bayesian updating result. This occurs when (i) there are no moment and tail bounds, or (ii) these bounds are sufficiently loose that they do not constrain the DM's updated beliefs, i.e., when they are not informative given his prior belief. In general, however, the updating is performed using an *exponential modifier* that captures the effect of the moment and tail bound information. This is an important result

with widespread applications beyond the motivating example of this paper. It establishes SOBME as a generalization of Bayesian updating for settings in which partial distributional information is available in addition to data/observations.⁵

Recall that the ambiguity set may contain parametric and/or non-parametric distributions. Therefore, in general, SOBME is a non-parametric approach to belief representation and updating. Furthermore, we highlight that the SOBME belief representation and updating result developed above are not limited to a repeated newsvendor setting. They apply to any setting in which a DM has partial distributional information (in the form of an ambiguity set and, possibly, moment and tail bounds) about a random variable of interest, and wishes to dynamically combine this information with an evolving history of observations.

In what follows, we apply the SOBME approach in a repeated newsvendor setting. Considering the sequence of events described in §3, we have the following. At the beginning of period t , the DM determines his order quantity $q_t \geq 0$ to minimize the expected newsvendor cost given his current belief about the demand distribution, which is represented by the density function $f_t(\omega)$ over the set of possible demand distributions contained in the ambiguity set $\mathcal{D} := \{F_D^\omega : \omega \in \Omega\}$. Applying the well-known newsvendor quantile result, the optimal quantity decision under SOBME is given by:

$$q_t^{SOBME} := \inf \left\{ q \in \mathcal{D} : \mathbb{E}_\omega [F_D^\omega(q)] \geq \frac{p}{p+h} \right\}, \quad (7)$$

where $\mathbb{E}_\omega [F_D^\omega(q)] = \int_{\omega \in \Omega} F_D^\omega(q) f_t(\omega) d\omega$. After ordering this quantity, the order is received, demand ξ'_t is realized, the newsvendor cost $\mathcal{L}(q_t^{SOBME} - \xi'_t)$ is incurred, and new exogenous research, if any, arrives regarding moment and tail bounds. At the end of period t , knowing the realized demand and the latest moment and tail bounds, the DM updates his belief from $f_t(\omega)$ to $f_{t+1}(\omega)$ using Corollary 1.

For later use, let Q_t^{SOBME} denote the period- t order quantity random variable. Q_t^{SOBME} takes on a realized value q_t^{SOBME} , specified by (7), given any particular sample path of demand realizations in periods $1, \dots, t-1$. Let

$$q^{\omega*} = \inf \left\{ q \in \mathcal{D} : F_D^{\omega*}(q) \geq \frac{p}{p+h} \right\}$$

denote the optimal order quantity if the DM knows the true distribution. This optimal order

⁵ Giffin and Caticha (2007) also show a related result, but unlike our framework, they only derive it under the very strong assumptions that the first moment is perfectly known and that this is the only information available, i.e., no tail bounds and no upper and/or lower moment bounds. Furthermore, their work does not examine the asymptotic behavior (consistency, rate of convergence, etc.) of the updating mechanism nor does it embed the distribution-formation problem in a decision-making context.

quantity is static, because the true distribution is stationary. In the next section, we show that $Q_t^{SOBME} \rightarrow q^{\omega^*}$ almost surely as $t \rightarrow \infty$.

4.3. Asymptotic Behavior of SOBME: Consistency and Convergence Rate

In this section, we study two important asymptotic properties of SOBME. The first property is consistency. We formalize this concept below, but loosely speaking, this means that the DM eventually places (almost) all his weight on the true demand distribution, with the result that the DM's quantity decision converges (with probability one) to the quantity decision under full distributional information: $Q_t^{SOBME} \rightarrow q^{\omega^*}$ almost surely as $t \rightarrow \infty$. The second property relates to the rate at which the demand distribution built by SOBME converges to the true distribution.

To establish consistency and a rate of convergence, we need to introduce the notion of coherent information for the moment and tail bounds.

DEFINITION 2 (Coherent Information). The sequence of information $\{\hat{\mathcal{F}}_t : t \in \mathcal{T}\}$ is said to be coherent if (i) the sequence of lower (upper) bounds on the first moment, i.e., $\{\underline{\mathbb{E}}_t(D) : t \in \mathcal{T}\}$ ($\{\overline{\mathbb{E}}_t(D) : t \in \mathcal{T}\}$), is non-decreasing (non-increasing), with $\underline{\mathbb{E}}_t(D) \leq \mathbb{E}(D) \leq \overline{\mathbb{E}}_t(D)$ for all $t \in \mathcal{T}$ where the expectation is with respect to the true distribution, and (ii) for any $u \in \mathcal{D}$, the sequence of tail upper bounds $\{b_t(u) : t \in \mathcal{T}\}$ is non-increasing with $b_t(u) \geq F_D^{\omega^*}(u)$.

If the moment and tail bound information is not coherent, then we should not expect an approach that uses such information to necessarily converge to the true distribution. For example, the information might either be “wrong” in some period, i.e., the moment and tail bounds may rule out the true distribution, or the sequence of bounds might be non-monotonic, e.g., periodic, which could prevent convergence even if the information is correct.

Recall that $f_t(\omega)$ is the density function (at the start of period- t) over the demand distribution indicator ω . Let F_t^{SOBME} denote its cumulative distribution, where we use the superscript SOBME to highlight that it is the result of the SOBME updating method. The following result establishes the weak consistency of SOBME by showing that the DM will always, i.e., regardless of the sample path of observations, eventually assign a weight of one to the shrinking KL neighborhoods of ω^* . We refer the reader to Appendix A for formal definitions of weak consistency and KL neighborhoods.

THEOREM 2 (Consistency of SOBME). *Suppose that (a) the sequence of information $\{\hat{\mathcal{F}}_t : t \in \mathcal{T}\}$ is coherent, and (b) the initial distribution F_1^{SOBME} is such that $\int_{K_{\omega^*}^\epsilon} dF_1^{SOBME} > 0$ for every $\epsilon > 0$ (where $K_{\omega^*}^\epsilon$ is a KL neighborhood of ω^*). Then:*

- (i) F_t^{SOBME} is weakly consistent.

(ii) $Q_t^{SOBME} \rightarrow q^*$ almost surely as $t \rightarrow \infty$.

In addition to consistency (showing almost-sure convergence to the true distribution), it is useful to explore the rate/speed of convergence. In a later section, we will examine this and other questions numerically. Here, we develop an almost-sure analytical bound on the distance between the true demand density and the one constructed using SOBME. To this end, for $x \in \mathcal{D}$, let $f_{D,t}^{SOBME}(x) := \int_{\omega \in \Omega} f_t(\omega) f_D^\omega(x) d\omega$ denote the demand density generated by SOBME at the start of period t .

THEOREM 3 (Rate of Convergence). *Suppose that the sequence of information $\{\hat{\mathcal{F}}_t : t \in \mathcal{T}\}$ is coherent and that regularity conditions 3 and 4 (see Appendix A) hold. Then, there exists constants η_1 and η_2 such that for large enough t :*

$$d_H^2(f_D^{\omega^*} || f_{D,t}^{SOBME}) \leq \eta_1 [\epsilon_t^2 + 2e^{-\eta_2 t \epsilon_t^2}]$$

almost surely, where $\epsilon_t = \max\{r_t, y_t^{1/2}\}$, with r_t and y_t defined in Appendix A, and $d_H(f_1 || f_2)$ is the Hellinger distance (defined in Appendix A) between any two probability densities f_1 and f_2 .

The above result establishes $\epsilon_t = \max\{r_t, y_t^{1/2}\}$ as a convergence rate for SOBME. As discussed in Appendix A, this rate of convergence depends on the tightness of the ambiguity set (measured through r_t) and the precision of the initial prior (measured through y_t).

5. Newsvendor's Cost of Ambiguity: A Performance Guarantee

Having characterized the asymptotic behavior of SOBME, we next explore the increase in the average cost per period (during a horizon of $T < \infty$) that the DM experiences due to ambiguity. We refer to this cost as the *newsvendor's cost of ambiguity*. We will derive an analytical upper bound for the newsvendor's cost of ambiguity when SOBME is used. This will enable us to develop a performance guarantee. To that end, let

$$C_T^{SOBME} := \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{f_D^{\omega^*}} [\mathcal{L}(Q_t^{SOBME} - D_t) - \mathcal{L}(q^{\omega^*} - D_t)] \quad (8)$$

denote the random average (per period) newsvendor's cost of ambiguity under SOBME, which depends on the demand realizations up to the beginning of period T , i.e., the realization of vector $\underline{D}_T = (D_i, i = 1, 2, \dots, T-1)$ which is generated based on the true distribution. Taking the expectation with respect to this vector, we denote by

$$c_T^{SOBME} := \mathbb{E}_{\underline{D}_T} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{f_D^{\omega^*}} [\mathcal{L}(Q_t^{SOBME} - D_t) - \mathcal{L}(q^{\omega^*} - D_t)] \right], \quad (9)$$

the expected average per period newsvendor's cost of ambiguity. It should be noted that C_T^{SOBME} is a random variable but c_T^{SOBME} is a number. In (8), Q_t^{SOBME} is a random variable with a realization

that depends on \underline{D}_t : the expectation in (8) is only with respect to D_t , whereas this randomness is removed through the outer expectation in (9).

The following results provide upper bounds for C_T^{SOBME} and c_T^{SOBME} when the support of all the distributions in the ambiguity set can be bounded from above. Bounded demand is not a strong assumption in practice because demand is never infinite.

THEOREM 4 (Newsvendor's Cost of Ambiguity). *Suppose the support \mathcal{D} is $[0, \bar{d}]$ for some $\bar{d} \in \mathbb{R}_+$. Letting $k := 2\sqrt{2\bar{d}} \max\{p, h\}$, the following hold for any T :*

(i)

$$C_T^{SOBME} \leq k \sqrt{\frac{1}{T} \sum_{t=1}^T d_{KL}(f_D^* || f_{D,t}^{SOBME})}, \quad (10)$$

almost surely.

(ii)

$$c_T^{SOBME} \leq k \sqrt{\mathbb{E}_{\underline{D}_T} \left[\frac{1}{T} \sum_{t=1}^T d_{KL}(f_D^* || f_{D,t}^{SOBME}) \right]}. \quad (11)$$

Part (i) of the above result shows that the random average per period KL divergence between the true density and the one built using SOBME nicely bounds, in the almost sure sense, the average per period cost of ambiguity. Part (ii) shows a related result for the expected average per period cost of ambiguity. We next develop a bound for $\mathbb{E}_{\underline{D}_T} \left[\frac{1}{T} \sum_{t=1}^T d_{KL}(f_D^* || f_{D,t}^{SOBME}) \right]$, the expected average per period KL divergence between the true density and the one built using SOBME. Due to the result of Theorem 4 part (ii), this will enable us to reach our end goal of providing a performance guarantee for SOBME.

To bound $\mathbb{E}_{\underline{D}_T} \left[\frac{1}{T} \sum_{t=1}^T d_{KL}(f_D^* || f_{D,t}^{SOBME}) \right]$, we need some general conditions on (i) the (dynamic) information on moment and tail bounds ($\hat{\mathcal{F}}_t$), (ii) the ambiguity set (\mathcal{D}'), (iii) the initial prior used by the DM under SOBME (f_1^{SOBME}), and (iv) the convergence rate. For (i), we use the coherent information condition (Definition 2) to ensure that the DM is not relying on false information. For (ii)-(iv), we use the ‘‘smoothness’’ and ‘‘soundness’’ conditions discussed in Clarke and Barron (1990) for Bayesian updating mechanisms. We briefly list those conditions here for our above-mentioned goal, and refer interested readers to Clarke and Barron (1990) for more detailed discussions. It should be noted that the ‘‘smoothness’’ condition is related to the behavior around ω^* and the ‘‘soundness’’ is related to the speed of weight accumulation around ω^* as more and more observations are made.

CONDITION 1 (Smoothness). *The density $f_D^\omega(\xi)$ and the divergence $d_{KL}(f_D^{\omega^*} || f_D^\omega)$ are both twice continuously differentiable at ω^* . Furthermore, (i) there exists an $\epsilon > 0$ such that:*

$$\mathbb{E} \left[\sup_{\omega \in \Omega: |\omega - \omega^*| < \epsilon} \left| \frac{\partial^2}{\partial \omega^2} \log f_D^\omega(\xi) \right|^2 \right] < \infty,$$

and (ii) $I(\omega)$ and $f_1^{SOBME}(\omega)$ are both finite, positive, and continuous at ω^* , where $I(\cdot)$ is the Fisher information and $f_1^{SOBME}(\cdot)$ is the initial prior density.

CONDITION 2 (Soundness). *The ambiguity set is said to be sound if the convergence of a sequence of ω values is equivalent to the weak convergence of the corresponding demand distributions:*

$$\omega \rightarrow \omega^* \Leftrightarrow F_D^\omega \rightarrow F_D^{\omega^*}.$$

The above conditions allow us to provide an effective bound for the expected value of the average per period KL divergence under our proposed updating mechanism.

THEOREM 5 (Expected Average Per Period KL Divergence). *Suppose the coherent information, smoothness, and soundness conditions hold. Then, for large enough T :*

$$\mathbb{E}_{\mathcal{D}_T} \left[\frac{1}{T} \sum_{t=1}^T d_{KL}(f_D^{\omega^*} || f_{D,t}^{SOBME}) \right] \leq \frac{1}{T} \left[\frac{1}{2} \log \frac{T}{2\pi e} + \frac{1}{2} \log I(\omega^*) + \log \frac{1}{f_1^{SOBME}(\omega^*)} + o(1) \right]. \quad (12)$$

The above result states that the expected average per period KL divergence between the true density and the one built using SOBME is at worst of the form $\frac{1}{T} [\frac{1}{2} \log \frac{T}{2\pi e} + l]$ for some constant l that depends on the initial prior and the Fisher information, which are both also affected by the ‘‘tightness’’ of the ambiguity set, \mathcal{D}' . The important implication is that it allows us to provide the following performance guarantee for our proposed SOBME approach: based on Theorem 4 part (ii), the expected average (per period) newsvendor’s cost of ambiguity under SOBME goes to zero at worst at a fast rate $\sqrt{\frac{\log T}{T}}$ as $T \rightarrow \infty$. In other words, although a DM facing demand ambiguity inevitably incurs an extra per-period cost compared to a DM who knows the true demand distribution, the average extra per-period cost approaches zero quite quickly if SOBME is used. The actual rate depends on the quality of the dynamic information the DM has on moments and tails; Theorem 5 provides an upper bound by considering the worst case, i.e., no information.

THEOREM 6 (Performance Guarantee). *If $\mathcal{D} = [0, \bar{d}]$ for some $\bar{d} \in \mathbb{R}_+$, then for large enough T under the conditions of Theorem 5:*

$$c_T^{SOBME} \leq k' \sqrt{\frac{\log T}{T}} + o(1), \quad (13)$$

where $k' = 2\bar{d} \max\{p, h\}$, and hence, $c_T^{SOBME} \rightarrow 0$ as $T \rightarrow \infty$ at a rate of $\sqrt{\frac{\log T}{T}}$.

Theorem 6 establishes that the DM does not need to have full distributional information to make suitable stocking decisions: SOBME enables the DM to deal effectively with ambiguity.

6. Numerical Experiments

As discussed in §1, companies reduce ambiguity in new product introductions by devoting significant market-research efforts to gather information about the demand distribution. That is, firms do not rely exclusively on demand observations when forming beliefs about an unknown demand distribution; they can and do generate additional partial information about the demand distribution. In this paper, such information is codified by the ambiguity set and moment and tail bounds, i.e., the exogenous research. It is managerially important to understand the value of such exogenous research in an operational setting (here, a repeated newsvendor). Furthermore, assuming there is value in the exogenous research, it is helpful to examine which type of information (ambiguity set or moment bounds) provides more value, and what factors, e.g., newsvendor quantiles, affect the value. Answers to these questions – which we provide through numerical studies below – can help managers understand when it is worth their effort to obtain exogenous information and what form of information is most valuable.

A natural way to explore the value of exogenous research is to compare our information setting – data and exogenous research – to two benchmark settings: (i) full information, and (ii) a data-only information setting. For this second benchmark, we need to choose an approach tailored to a data-only setting. Various pure data-driven approaches have been developed in the literature as we discussed in §2, among which we consider the Sample Average Approximation (SAA) approach (see, e.g., Kleywegt et al. 2002, Shapiro 2003, Levi et al. 2007, 2010), Scarf’s method (Scarf 1958, Gallego and Moon 1993), and the recently developed likelihood robust optimization (LRO) approach (Wang et al. 2013).

Under SAA, the DM at the beginning of period $t + 1$ sets

$$q_{t+1}^{SAA} := \inf \left\{ q \in \mathcal{D} : \hat{F}_D^t(q) \geq \frac{p}{p+h} \right\}, \quad (14)$$

where $\hat{F}_D^t(q)$ is the empirical c.d.f. after t observations: $\hat{F}_D^t(q) := \frac{1}{t} \sum_{k=1}^t \mathbb{1}\{\xi^k \leq q\}$, with ξ^k denoting the realization in period k . In Scarf’s method, the DM minimizes the worst-case expected cost among all the distributions that have a mean equal to the sample mean and a variance equal to the sample variance. Using Scarf’s result (Scarf 1958), the optimal ordering quantity at the beginning of period $t + 1$ is

$$q_{t+1}^{SCARF} := m_t + \frac{s_t}{2} \left(\sqrt{\frac{p}{h}} - \sqrt{\frac{h}{p}} \right), \quad (15)$$

where m_t and s_t are the sample average and sample standard deviation of the observations made in periods $1, 2, \dots, t$. In the dynamic LRO approach (Wang et al. 2013), the set of possible distributions

in each period t contains only those that make the observed history of data achieve a certain level of likelihood, and the optimal decisions are made based on the worst-case distribution within that set. In our repeated newsvendor setting, we numerically observed that SAA (weakly) outperformed both Scarf’s method and dynamic LRO, and so we adopt the SAA approach as the representative for the data-only setting.⁶ That is, for brevity, we do not present our numerical results for Scarf’s method and LRO, and instead focus on the performance of SAA among the pure data-driven approaches.

6.1. Study Design

In addition to examining the value of exogenous information by comparing the performances of SOBME and SAA, our numerical study explores the performance of SOBME relative to a DM who has full distributional information. It also investigates how the performance of SOBME and the value of exogenous information are influenced by the true distribution, the type of exogenous information (ambiguity set and moment bounds), and the newsvendor quantile $p/(p+h)$.

True Distribution. The true distribution, $F_D^{\omega*}$, in our study is assumed to have a mean of 15. The true distribution is either normal or exponential, with a label “A” denoting the normal case and “B” denoting the exponential case.

Ambiguity Set. We consider two cases for the ambiguity set \mathcal{D} . In Case I, the ambiguity set contains all exponential distributions with a mean in $[10, 20)$ as well as all normal distributions with a mean in $[10, 20]$ and a coefficient of variation (C.V.) of 0.2. In Case II, the ambiguity set includes all those distributions in Case I plus their mixtures (with equal weights of 0.5) of any exponential with mean $[10, 20)$ and any normal with a mean in $[10, 20]$ but a C.V. of 0.2. In other words, Case II represents a larger ambiguity set that contains Case I and mixtures of distributions in Case I.⁷

Therefore, the combination of ambiguity sets and the true distribution leads to four different cases: I-A, I-B, II-A and II-B. In all cases, we assume SOBME starts with a uniform distribution on the ambiguity set, i.e., it starts with the highest information entropy.

Moment Bound Information. For each of these four cases, we considered two possible sets of

⁶ Our results for SAA, LRO and Scarf’s method echo those presented in Wang et al. (2013) which explores the performance of LRO and Scarf’s methods as compared to SAA for a single-shot newsvendor setting with ample data. We also observed that SAA has certain computational advantages over LRO, which is an important benefit in a dynamic multi-period setting such as ours. Scarf’s approach is closer to SAA in terms of computational complexity, but its performance is dominated by SAA.

⁷ In Case I, we consider $\Omega = [10, 30]$, where $\omega \in [10, 20)$ indicates an exponential distribution with mean ω , and $\omega \in [20, 30]$ indicates a normal distribution with a mean $\omega - 10$ and a C.V. of 0.2. The construction of Ω in Case II is done in an analogous manner.

moment bound information, $\{\hat{\mathcal{F}}_t : t \in \mathcal{T}\}$. In one situation, the DM has no moment information beyond those dictated by the ambiguity set (referred to as the “No Bounds” situation). In the other situation (referred to as the “Tight Bounds” situation), the DM has dynamically evolving first moment bound information. In this “Tight Bounds” situation, the initial lower and upper bounds are those dictated by the ambiguity set, i.e., $\underline{\mathbb{E}}_0(D) = 10$ and $\overline{\mathbb{E}}_0(D) = 20$, but these bounds get progressively tighter by 1.5 units in each period until they hit $\underline{\mathbb{E}}_t(D) = 14.5$ and $\overline{\mathbb{E}}_t(D) = 15.5$, after which they remain constant. To be conservative when evaluating SOBME, we assume that the DM does not have tail bound information. Such information, if available, would of course improve the performance.

Newsvendor Quantile. The above combination of two ambiguity sets, two true distributions, and two moment information sets leads to eight cases. For each of these cases, as our base-case scenario, we set the newsvendor quantile to $p/(p+h) = 0.75$, with $h = 1$ and $p = 3$. However, as we will describe later, we also explore the impact of the newsvendor quantile by varying its value in $\{0.50, 0.55, \dots, 0.95\}$.

Runs. Each case is run for 100 periods and the average cost is calculated for the SOBME approach, the full information case, i.e., when the true distribution is known, and the pure data-driven SAA approach. We purposely choose a limited time horizon because we know that the performance of both SOBME and SAA will asymptotically approach that of full information as more and more demand observations become available. In other words, when ample data exists, it is less important which approach is used, as long as it is a consistent approach. As such, we are mainly interested in how well SOBME performs when data is limited. The limited-data setting also better represents the new product introduction we discussed in §1. The average costs are calculated by conducting 50 independent replications with each replication being a sample path generation, i.e., demand realizations in periods 1 to 100. Under each sample path and in each period, the related expected newsvendor costs are calculated directly and based on the closed-form newsvendor overage plus underage costs of the DM’s ordered quantity.

In each period, the closed-form solution of Theorem 1 (and Corollary 1) are used to calculate the updated SOBME belief function and then the newsvendor problem is solved based on the updated beliefs. The belief-updating problem involves calculating the function $g_t(\omega)$ which requires finding Lagrangian multipliers $\overline{\beta}_t$ and $\underline{\beta}_t$. Note that $\gamma_t(u)$ can be set to zero because there is no tail constraint. These two multipliers are calculated efficiently in each period t by employing a numerical search method which takes advantage of the properties of the underlying optimization problem.

The same numerical search can be used for calculating $\gamma_t(u)$ when it is non-zero. We note that the belief-updating optimization problem is independent of the DM's ordering problem, and so the belief-updating aspect of SOBME can be applied for a wide range of decision-making problems in which there is distributional ambiguity. The computational burden in calculating various integrals involved in our method (arising from the continuous first and second-order distributional assumptions, i.e., the sets \mathcal{D} and Ω being continuous) is overcome by using high precision function interpolations and advanced numerical methods for calculating integrals.

6.2. Results

For each of the eight cases in the study design, Figure 1 reports the Average Cost Optimality Gap (%) for both SOBME and SAA, where the optimality gap refers to the percentage difference in cost as compared to the full information case in which the DM knows the true distribution. As can be seen, SOBME performs very well: its optimality gap is less than 1% after 20 periods in six of the eight cases. We next discuss how the exogenous research information (moment bounds and ambiguity set) and the true distribution impact the performance of SOBME relative to full information and to the pure data-driven representative, SAA.

The Effect of Exogenous Information. The impact of the *moment bound information* on SOBME's performance can be explored by comparing the SOBME plots in the left column (No Bounds) with the corresponding SOBME plots in the right column (Tight Bounds). Moment bound information is clearly valuable: SOBME's optimality gap approaches zero much more rapidly when tight moment bound information is available. The impact of the *ambiguity set* on SOBME's performance can be explored by comparing Case I-A with Case II-A and comparing Case I-B with Case II-B. We observe that the performance is quite similar in both comparisons, indicating that SOBME is quite *robust* to the ambiguity set. We tested this observation further (for the case in which the true distribution was normal) by also considering a setting in which the ambiguity set contained only normal distributions. Contracting the ambiguity set did not materially impact SOBME's performance. The fact that SOBME's performance is robust to the ambiguity set is indeed an important benefit because it allows the DM to expand his ambiguity set when he has more doubt about the form of the true distribution.

By comparing the difference between SOBME's performance and that of SAA in each subfigure, we see that the DM benefits significantly from exogenous information in the form of the ambiguity set and moment bounds. This indicates that there is potentially great value in marketing research


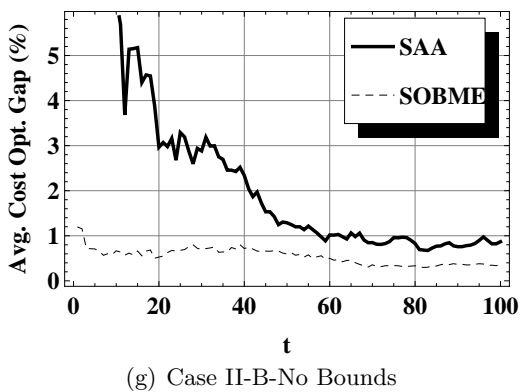

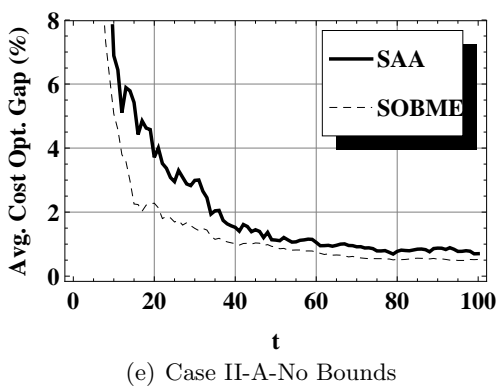



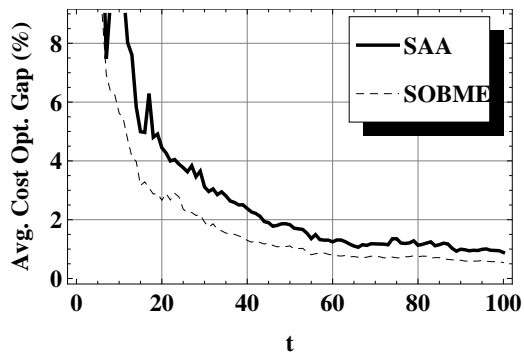


Figure 1 Comparison of SAA and SOBME under various cases.

and related actions that can generate this kind of information. In addition, our results suggest that the information contained in the moment bounds is more valuable than those contained in the ambiguity set. Hence, the DM should prefer a tighter set of moment bounds than a tighter ambiguity set: efforts that result in tighter moment bounds are in general more valuable than those that result in tighter ambiguity sets.

The Effect of the True Distribution. Comparing the normal cases (I-A and II-A) with the exponential cases (I-B and II-B), we see that SOBME performs better when the true distribution is exponential. In the exponential cases, SOBME's optimality gap is less than 1% after a few periods under both "No Bounds" and "Tight Bounds" situations. That SOBME performs better when the true distribution is exponential is due to fact that the newsvendor cost is much more sensitive to distributional mis-specification when the true distribution is normal than when the true distribution is exponential. To observe this consider the following cases within our numerical setting ($h = 1$ and $p = 3$): (i) the true distribution is normal ($CV=0.2$), but the DM believes with probability one that it is exponential ($CV=1$) (both having a mean of 15), and (ii) the true distribution is exponential ($CV=1$) but the DM believes with probability one it is normal ($CV=0.2$) (both having a mean of 15). In case (i) the DM orders 20.79 units while the true optimal ordering quantity is 17.03. This results in an optimality gap of about 55.15% (in terms of the percentage additional expected cost incurred). In case (ii), the DM orders 17.03 while the true optimal ordering quantity is 20.79. This results in an optimality gap of only about 2.48% (in terms of the percentage additional expected cost incurred). Comparing these two cases shows that a wrong distributional assumption is less consequential when the true distribution is exponential than when it is normal. Hence, having a good estimate of the true distribution is more important when the true distribution is normal than when it is exponential. This fact shows why SOBME performs better when the true distribution is exponential than when it is normal. Interestingly, however, SAA's performance is slightly better when the true distribution is normal than when it is exponential. What accounts for this? There is a force that counteracts the robustness under the exponential distribution. The DM sees a more volatile set of observations under an exponential distribution than under the normal distribution. This significantly influences the rate of convergence of SAA. This second driving force is stronger than the first one for SAA. This is not the case for SOBME, because it is not as reliant on observations as SAA, i.e., SOBME is not using an empirical distribution.

The Effect of the Newsvendor Quantile. To explore the impact of the newsvendor quantile on the performance of SOBME, we focused on the four Case I studies (exponential and normal

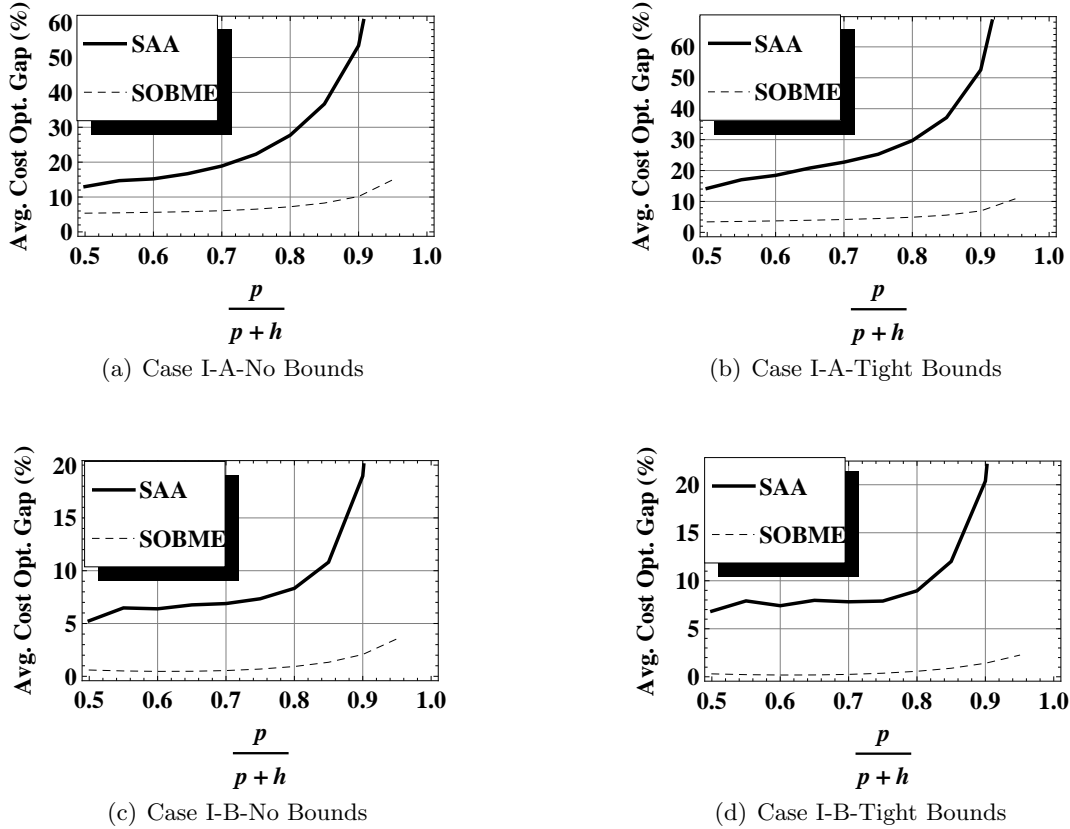


Figure 2 Sensitivity analyses on the newsvendor quantile $p/(p+h)$.

true distributions; No Bounds and Tight Bounds information), and varied the newsvendor quantile $p/(p+h) \in \{0.50, 0.55, \dots, 0.95\}$, while fixing $h = 1$. Figure 2 shows the performance of SOBME and SAA for each of the four cases. The performance of SAA and SOBME are averaged over 50 replications as well as the first 20 periods. We first note that our study confirms what has already been reported in the literature for SAA: its performance significantly degrades as the newsvendor quantile increases (see, e.g., Levi et al. 2010). In contrast, we observe that SOBME’s performance is much more robust: its performance degradation is very low. For instance, in Figure 2 (a), while SAA has an optimality gap of 55% for a newsvendor quantile of 0.9, SOBME has an optimality gap of only 10%. When moment bound information is available, Figure 2 (b) and (d), the difference in performance between SOBME and SAA is even greater. Recall that the difference in performance between SOBME and SAA measures the value of the exogenous information. Thus, our result suggests that this type of information is particularly valuable in settings characterized by high newsvendor quantiles. Moreover, the robustness of SOBME to the newsvendor quantile suggests that SOBME is a very suitable approach if a high service level is required or if inventory cost parameters are subject to estimation error.

In summary, the above study has the following managerial implications: (i) there is significant value in exogenous information (as measured by the comparison of SOBME and SAA), (ii) market research that leads to tighter moment bounds is more valuable than research that leads to a more compact ambiguity set, and (iii) it is especially important to avail of exogenous information (through SOBME) when the newsvendor quantile is high because pure-data driven approaches (SAA) perform much worse at high quantiles.

As mentioned above, the SOBME approach is not limited to a repeated newsvendor setting. It can be used in any multi-period setting in which a DM has partial distributional information (in the form of an ambiguity set and, possibly, moment and tail bounds) about a random variable of interest. A general measure of SOBME's performance is the quality of the distribution generated by SOBME, i.e., how close is it to the true distribution from which the observations are drawn? We investigate this in two different studies in Appendix C, the second of which uses data from a case based on a real-world setting (Allon and Van Mieghem 2011) to show how moment bounds can be developed from "like-products". Among other results, we find that the quality of the SOBME-generated distribution is very good, that partial information in the form of moment bounds is particularly valuable early in the horizon, and that companies can get significant benefit from moment bound information even if the number of like-products is quite limited. We refer the reader to Appendix C for further details.

6.3. Extensions

6.3.1. Endogenous Bounds. We have assumed to this point that the sequence of moment bound information $\{\underline{\mathbb{E}}_t(D), t \in \mathcal{T}\}$ and $\{\overline{\mathbb{E}}_t(D), t \in \mathcal{T}\}$ are derived solely from exogenous research. However, our SOBME framework can also accommodate dynamic bounds that are endogenously generated from the data observations. For instance, using the Central Limit Theorem (CLT) at the end of period t , a $1 - \alpha$ confidence interval for the first moment is given by

$$I_t := (m_t - z_{\alpha/2} \times s_t / \sqrt{t}, m_t + z_{\alpha/2} \times s_t / \sqrt{t}), \quad (16)$$

where z denotes a standard normal random variable, and m_t and s_t are the sample average and sample standard deviation obtained from observations made in periods $1, 2, \dots, t$. In the absence of exogenous bounds, the DM can use the boundaries of I_t as moment bounds. If the DM has exogenous bounds given by some sequences $\{\underline{\mathbb{E}}_t(D), t \in \mathcal{T}\}$ and $\{\overline{\mathbb{E}}_t(D), t \in \mathcal{T}\}$, then he can replace them with the respective boundaries of I_t in a period t when $I_t \subset [\underline{\mathbb{E}}_t(D), \overline{\mathbb{E}}_t(D)]$, i.e., when the endogenous bounds are tighter than the exogenous ones. This ensures that the moment bound

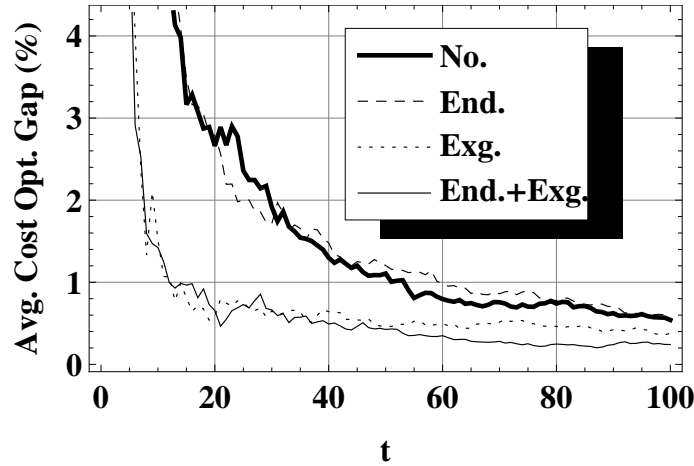


Figure 3 The effect of endogenous bounds on the performance of SOBME (No.: no exogenous or endogenous bounds; End.: endogenous bounds only; Exg.: exogenous bounds only; End.+Exg.: both endogenous and exogenous bounds).

information remains coherent when the exogenous bound information is coherent.

We investigated the value of endogenous bounds by applying this approach (using a 90% confidence interval) to cases I-A-No Bounds and I-A-Tight Bounds. Figure 3 depicts the performance of SOBME for the four possible situations (No./Tight Bounds with/without Endogenous Bounds). As can be seen, endowing the DM with endogenous bounds does not materially improve SOBME’s performance. Irrespective of the endogenous bounds, SOBME is already availing of the data to updates its probabilities across the potential distributions in the ambiguity set. Therefore, the data used to generate the endogenous bounds has also been used (independently) to appropriately “weight” the possible distributions, and those that conform to the endogenous bounds will have received more weight. In other words, SOBME is already using the observations, i.e., “the endogenous data”, in an effective way, and hence using them to create supplementary moment bounds is not of significant help. We also note that the endogenous bounds are typically tight and, therefore, useful, only when they are generated based on ample data, i.e., a large t . However, as discussed earlier, when ample data exists, any estimation approach that is consistent will generate good results and it matters less which one is used. Moreover, it is precisely because they do not have ample data that firms devote efforts to generating exogenous research on the demand distribution.

6.3.2. Change Detection. We have assumed to this point that the true demand distribution is stationary throughout the horizon. In practice, however, the distribution may change at certain points in time due to, for example, changes in economic factors, competitor actions, market penetration, and time of year. Some of these changes might be amenable to prediction by managers, and SOBME can accommodate such predications through updated moment bounds. Other changes




Figure 4 Comparison of SAA and SOBME for change detection (Setting Case-I-C) (NB: No Bounds; TB: Tight Bounds).

might not be predictable and completely hidden to the DM, and so it is important to test how well approaches such as SOBME and SAA detect and respond to unpredicted changes.

We numerically examined the ability of SOBME and SAA to detect and respond to such a hidden/unpredicted change in the type of demand distribution. In particular, we use a new setting, setting Case I-C, in which, unknown to the DM, the true distribution changes from normal ($CV=0.2$, $mean=15$) to exponential ($CV=1$, $mean=15$) in the middle of the horizon, i.e., period 50. We evaluate the performance of SOBME (under both “No Bounds” and “Tight Bounds”) and SAA compared to a DM that knows the true distribution throughout the horizon, i.e., both before and after the change. The results are presented in Figure 4. Not surprisingly, both SOBME and SAA are affected by the change at period 50. The change initially affects SOBME slightly more than SAA because SOBME was already getting very close to the true distribution before the change whereas SAA was further from it. Also, we note the spike is lower (and comparable to that for SAA) when SOBME has tight bounds. Very importantly, however, irrespective of the bounds, SOBME recovers much more rapidly than SAA: it quickly recognizes the hidden change and adjusts to it. This reflects the fact (analytically established in the performance guarantee and numerically observed) that SOBME’s cost performance converges quite quickly to that of full information. This characteristic is particularly desirable in dynamic and non-stationary environments (where the underlying true distribution is subject to inevitable changes) as it allows the DM to automatically and quickly adjust its decisions.

7. Repeated Newsvendor under Ambiguity Aversion

We assumed that the DM was ambiguity neutral in the preceding sections. In general decision-making settings, a DM may, in fact, be ambiguity averse, ambiguity neutral, or ambiguity loving;

see, e.g., Klibanoff et al. (2005, 2009). In what follows, we extend our results and demonstrate how the SOBME approach can be also used when the DM is not ambiguity-neutral. We adopt the general framework proposed by Klibanoff et al. (2005), hereafter KMM, for smooth decision making under ambiguity. Tailoring the KMM decision-making framework (see equation (1) of KMM) to a newsvendor problem, we cast the DM's quantity decision at the beginning of period t as

$$\max_{q_t \geq 0} \mathbb{E}_{f_t^{SOBME}(\omega)} \left[\Phi \left(\mathbb{E}_{f_D^\omega} \left[-\mathcal{L}(q_t - D) \right] \right) \right], \quad (17)$$

where $\Phi(\cdot)$ represents the DM's ambiguity attitude, and $f_t^{SOBME}(\omega)$ is the second-order belief built using SOBME. When $\Phi(\cdot)$ is affine, the DM is ambiguity-neutral, which is the case considered in the previous sections. However, when $\Phi(\cdot)$ is concave (convex), the DM is ambiguity-averse (loving).⁸

In the case of an ambiguity-neutral DM, as discussed in §4.1, there is a type of period- t equivalence between second-order and first-order approaches in the sense that the expected value of the DM's objective, see (1), depends on the belief $f_t(\omega)$ and the ambiguity-set densities $f_D^\omega(x)$ only through the demand density $f_{D,t}^{SOBME}(x) := \int_{\omega \in \Omega} f_t(\omega) f_D^\omega(x) d\omega$. Therefore, a first-order approach that somehow arrived at an identical demand density would result in the same period- t objective. Importantly, this theoretical equivalence holds only when the DM is ambiguity neutral. For settings with an ambiguity averse/loving DM, the objective, i.e., (17), depends on $f_t(\omega)$ and $f_D^\omega(x)$ in a manner that cannot be captured through $f_{D,t}^{SOBME}(x) := \int_{\omega \in \Omega} f_t(\omega) f_D^\omega(x) d\omega$. Therefore, it is crucial to adopt a distribution-over-distributions approach. SOBME provides a natural and effective approach for ambiguity averse/loving settings.

In what follows, we first explore some structural properties of the decisions made by an ambiguity-averse newsvendor, i.e., properties of the period- t objective (17), when $\Phi(\cdot)$ is concave. We then explore the effect of the DM's ambiguity aversion on the newsvendor ordering quantity. To this end, we consider the negative exponential function $\Phi(x) = -e^{-\alpha x}$ with parameter $\alpha \geq 0$, which is increasing concave and has a constant *coefficient of ambiguity* $\alpha = -\Phi''(x)/\Phi'(x)$. In this setting, a higher α represents a higher ambiguity aversion. For notational convenience, we use $U^\omega(q) = \mathbb{E}_{f_D^\omega} [\mathcal{L}(q - D)]$ to denote the random variable that represents the expected *disutility* of ordering q units, with its realization depending on $\omega \in \Omega$. The KMM objective (17) is then equivalent to

$$\min_{q_t \geq 0} \mathbb{E}_{f_t^{SOBME}(\omega)} \left[e^{\alpha U^\omega(q_t)} \right] = \min_{q_t \geq 0} \mathcal{M}_{U^\omega(q_t)}(\alpha), \quad (18)$$

where $\mathcal{M}_{U^\omega(q_t)}(\alpha)$ is the moment generating function of $U^\omega(q_t)$ evaluated at α . For technical reasons, we assume that (i) the ambiguity set \mathcal{D} is such that this moment generating function exists

⁸ The internal expectation of the newsvendor cost $\mathcal{L}(\cdot)$ in (17) implies a risk-neutral DM. This could be relaxed but we focus on incorporating ambiguity attitudes and not risk attitudes.

and is differentiable with respect to q , and (ii) all the demand distributions in \mathcal{D} are differentiable and have densities in \mathcal{D}' . Finally, we note that minimizing the moment generating function in (18) is equivalent to minimizing the so-called *entropic* disutility $\frac{1}{\alpha} \log \mathcal{M}_{U^\omega(q_t)}(\alpha)$, and hence we use that form in proofs when more convenient.

LEMMA 1 (Convexity). *The moment generating function $\mathcal{M}_{U^\omega(q)}(\alpha)$ is convex in q ($\forall \alpha \geq 0$), and hence, (18) is a convex program.*

We let $q_t^{KMM}(\alpha)$ denote the optimal solution to (18) on a given sample path. That is, $q_t^{KMM}(\alpha)$ is the period- t optimal ordering quantity of a newsvendor (with a coefficient of ambiguity α) given a particular realization of demands up to and including period $t - 1$. We let $Q_t^{KMM}(\alpha)$ denote the random variable associated with $q_t^{KMM}(\alpha)$ when looking over all possible sample paths through period $t - 1$.

LEMMA 2 (Optimal Ambiguity-Averse Ordering Quantity). *For $\alpha > 0$, let $Z^\omega(q) = U^\omega(q) + \frac{1}{\alpha} \log F_D^\omega(q)$. The optimal ordering quantity $q_t^{KMM}(\alpha)$ is either zero or the solution to the implicit equation*

$$\frac{\mathcal{M}_{Z^\omega(q)}(\alpha)}{\mathcal{M}_{U^\omega(q)}(\alpha)} = \frac{p}{p+h}. \quad (19)$$

The above lemma proves that the optimal ordering quantity, if positive, sets the ratio of the relevant moment generating functions equal to the well-known newsvendor quantile $p/(p+h)$. This can be viewed as a generalization of the classical newsvendor quantile result. Using this result, we now establish how ambiguity aversion influences the optimal ordering quantity (in any given period t), and how the optimal ordering quantity (at any given ambiguity coefficient) behaves as t tends to infinity.

THEOREM 7 (Ambiguity Aversion Effect). *The optimal ordering quantity under ambiguity-aversion satisfies the following:*

- (i) $q_t^{KMM}(\alpha)$ is non-increasing in α .
- (ii) $\lim_{\alpha \rightarrow 0^+} q_t^{KMM}(\alpha) = q_t^{SOBME}$ (defined by (7)).
- (iii) $\lim_{\alpha \rightarrow \infty} q_t^{KMM}(\alpha) = 0$.
- (iv) $\lim_{t \rightarrow \infty} Q_t^{KMM}(\alpha) = q^{\omega^*}$ almost surely when the conditions of Theorem 2 hold.

A higher level of ambiguity aversion, i.e., a higher α , results in a lower ordering quantity. As α approaches zero, the DM acts similarly to an ambiguity-neutral DM, that is, his order quantity approaches q_t^{SOBME} defined in (7). As α approaches ∞ , the DM only considers the worst-case outcome, i.e., becomes a minimax optimizer, and hence orders nothing so as to avoid any overage




Figure 5 Performance of SOBME under the KMM criteria (Case I-A-No Bounds).

cost. Regardless of the DM’s ambiguity aversion level, as more and more demand observations are made ($t \rightarrow \infty$), his ordering quantity converges to that of a DM who completely knows the true demand distribution. This highlights the asymptotical suitability of SOBME under any ambiguity aversion level.

We numerically explore the performance of SOBME under the KMM criterion by considering the “Case I-A-No Bounds” setting introduced in §6.1. Figure 5 presents the average cost optimality gap (%) for three levels of ambiguity aversion: $\alpha = 0.005$, $\alpha = 0.010$, $\alpha = 0.050$. The average cost optimality gap (%) is the percentage extra cost of the ordered quantities of a DM who is facing demand ambiguity and uses the objective (17) compared to a DM who completely knows the true demand distribution. For each period shown in Figure 5, the optimality gap is averaged over 20 independent sample path replications, i.e., demand observations over the horizon. Consistent with Theorem 7 (iv), the average optimality gap goes to zero as t increases. Moreover, as α decreases, the performance under ambiguity aversion tends to that observed in our previous ambiguity-neutral results.

Although Lemma 2 characterizes the period- t optimal order quantity in general, a much sharper characterization can be established in certain cases. In particular, under (i) low but positive ambiguity aversion, or (ii) a normality assumption on the disutility distribution, we can derive closed-form solutions for the ambiguity-averse newsvendor problem by applying the *certainty equivalence* principle. Using this principle drastically simplifies the problem, because we only need to consider the first and second moments of the disutility $U^\omega(q) = \mathbb{E}_{f_D^\omega}[\mathcal{L}(q - D)]$ as opposed to its moment generation function. In what follows, we let $F_{D,t}^{SOBME}(\xi) := \mathbb{E}_{\omega,t}[F_D^\omega(\xi)] = \int_{\omega \in \Omega} f_t(\omega) F_D^\omega(\xi) d\omega$ denote

the demand distribution built by SOBME (used for decision-making in period t) and we let $F_{D,t}^{SOBME(-1)}(\cdot)$ denote its inverse. First, we consider the case in which the DM has a low but positive level of ambiguity aversion.

THEOREM 8 (Low Ambiguity Aversion). *Let $\mu(q_t)$ and $\sigma(q_t)$ denote the mean and standard deviation of $U^\omega(q_t)$ (with respect to ω), respectively. Then*

(i) *The optimal ordering quantity under ambiguity-aversion is*

$$q_t^{KMM}(\alpha) = \operatorname{argmin}_{q_t} \left\{ \mu(q_t) + \frac{\alpha}{2} \sigma^2(q_t) + \frac{\alpha}{2} \mu^2(q_t) + O(\alpha^2) \right\}. \quad (20)$$

(ii) *The solution to (20) is*

$$q_t^{KMM}(\alpha) = F_{D,t}^{SOBME(-1)} \left(\frac{p}{p+h} - \alpha \eta(q_t^{KMM}) - O(\alpha^2) \right), \quad (21)$$

where $\eta(q_t) := \mathbb{E}[F_D^\omega(q_t) U^\omega(q_t)] - \mathbb{E}\left[\frac{p}{p+h} U^\omega(q_t)\right]$.

We see that when the DM has a low but positive level of ambiguity aversion, the optimal ordering quantity is a perturbed version of the traditional critical quantile newsvendor solution. This perturbation, is approximately (with an error of order α^2 denoted by $O(\alpha^2)$) *linear* in the coefficient of ambiguity, α , with a slope of $-\eta(q_t^{KMM})$.

Next, we consider a case where $U^\omega(q)$ has an approximately Normal distribution based on the DM's second-order belief. See Appendix D for an example illustrating the validity of this approximation.

REMARK 2 (Normal Approximation). Suppose at some period t , $U^\omega(q_t)$ is approximately Normal with mean $\mu(q_t)$ and standard deviation $\sigma(q_t)$. Then

(i) The newsvendor problem under ambiguity aversion is equivalent to the mean-variance optimization $\min_{q_t} \mu(q_t) + \frac{\alpha}{2} \sigma^2(q_t)$. That is,

$$q_t^{KMM}(\alpha) \simeq \operatorname{argmin}_{q_t} \left\{ \mu(q_t) + \frac{\alpha}{2} \sigma^2(q_t) \right\}. \quad (22)$$

(ii) The solution to (22) is

$$q_t^{KMM}(\alpha) \simeq F_{D,t}^{SOBME(-1)} \left(\frac{p}{p+h} - \alpha \operatorname{Cov}(F_D^\omega(q_t^{KMM}), U^\omega(q_t^{KMM})) \right). \quad (23)$$

As in the case of low ambiguity aversion, the optimal ordering quantity is again a perturbed version of the traditional critical quantile solution, with the linear perturbation now having a slope of $-\operatorname{Cov}(F_D^\omega(q_t^{KMM}), U^\omega(q_t^{KMM}))$.

8. Conclusion

There are many multi-period settings in which a decision maker does not have full distributional information about a random variable of interest but may have some partial distributional information in the form of bounds on moments or tails. Over time, the decision maker gains access to demand observations. Techniques which rely exclusively on partial distributional information ignore the value of observations. In contrast, pure data-driven techniques avail of observations but ignore partial distributional information. With the growing attention to business analytics, it is essential to bridge this gap by developing approaches that allow decision makers to benefit from both data and partial distributional information.

In this paper, using the maximum entropy principle, we developed an approach (SOBME) that allows a decision maker to effectively use partial distributional information and data/observations. Moreover, SOBME can accommodate dynamically-evolving partial information. We established that SOBME is a natural generalization of traditional Bayesian updating; an exponential modifier is introduced to incorporate the partial distributional information. This is an important result, with applications far beyond the newsvendor focus of this paper. We proved that SOBME is weakly consistent and characterized its rate of convergence. Applying SOBME in a repeated newsvendor setting, we introduced and analyzed the notion of newsvendor's cost of ambiguity. We provided an analytical bound for this cost when the DM uses SOBME. More broadly, our results indicate that information-theoretical approaches, such as the one we developed in this paper, provide a natural way for characterizing an upper bound for the newsvendor's cost of ambiguity.

Our numerical investigation demonstrated that SOBME performs well, not just outperforming a pure data-driven approach but often coming close to the performance of full distributional information. We also found that unlike a pure data-driven approach such as SAA, SOBME's performance does not significantly degrade as the newsvendor quantile increases, making it a suitable candidate for settings with high service level requirements. By comparing our approach to pure data-driven approaches, we also generated insights into the value of partial distribution information. Our results indicate that it can be very beneficial to incorporate such information when deciding about the stocking quantities. Also, we found that information in the form of tighter moment bounds is typically more valuable than information in the form of tighter ambiguity sets.

Finally, we examined an ambiguity-averse version of our repeated newsvendor problem, and showed how SOBME can be applied in that setting. We established theoretical results that characterize the optimal newsvendor quantity under ambiguity aversion, and demonstrated that SOBME

is well-suited to an ambiguity-averse setting.

In closing, we emphasize that the SOBME belief representation and updating approach can be applied to any setting in which a decision maker has access to observations/data and some partial distributional information. We hope to apply and test SOBME in various other settings in future work.

Acknowledgment: The authors are grateful to *Retsef Levi* for his invaluable suggestions, comments, and discussions which helped to improve this paper. The authors also thank Prof. *Chung Piaw Teo*, the associate editor, and two anonymous referees for their constructive suggestions that improved the content and exposition of this paper.

References

- Aaker, D.A., V. Kumar, G.S. Dayand, R.P. Leone. 2010. *Marketing Research*. John Wiley & Sons.
- Allon, G., J.A. Van Mieghem. 2011. The mexico china dual sourcing strategy simulation. *Kellogg Case Number 1-212-500* .
- AMR. 2008. Innovation in consumer products: New products from concept to launch. AMR Research 2008 Study.
- Amrute, S. 1998. Forecasting new products with limited history: Nabisco's experience. *J. Bus. Forcast. Meth. Syst.* **17** 7–11.
- Andersson, J., K. Jörnsten, S.L. Nonås, L. Sandal, J. Ubøe. 2013. A maximum entropy approach to the newsvendor problem with partial information. *Eur. J. Oper. Res.* **228** 190–200.
- Azoury, K.S. 1985. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Sci.* **31**(9) 1150–1160.
- Barron, A., M.J. Schervish, L. Wasserman. 1999. The consistency of posterior distributions in nonparametric problems. *Ann. of Stat.* **27**(2) 536–561.
- Ben-Tal, A., D. den Hertog, A. De Waegenaere, B. Melenberg, G. Rennen. 2013. Robust solutions of optimization problems affected by uncertain probabilities. *Management Sci.* **59**(2) 341–357.
- Bookbinder, J.H, A.E. Lordahl. 1989. Estimation of inventory reorder level using the bootstrap statistical procedure. *IIE Trans.* **21** 302–312.
- Boyd, S. P., L. Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press, Cambridge.
- Caticha, A., A. Giffin. 2006. Updating probabilities. *Bayesian Inference and Maximum Entropy Methods in Sci. and Eng.*, vol. 872. 31–42.
- Chen, M., Y. Xia, X. Wang. 2010. Managing supply uncertainties through bayesian information update. *IEEE Trans. on Automation Sci. and Eng.* **7**(1) 24–36.
- Chu, L.Y., J.G. Shanthikumar, Z.-J. Shen. 2008. Solving operational statistics via a bayesian analysis. *Oper. Res. Lett.* **36**(1) 110–116.
- Clarke, B., A. Barron. 1990. Information theoretic asymptotics of Bayes methods. *IEEE Trans. Info. Theo.* **36**(3) 453–471.
- Cover, P., J. Thomas. 1991. *Elements of Information Theory*. John Wiley and Sons.
- Delage, E., Y. Ye. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* **58**(3) 595–612.
- Diaconis, P., D. Freedman. 1986. On the consistency of Bayes estimates. *Ann. of Stat.* **14**(1) 1–26.
- Fisher, M., A. Raman. 2010. *The New Science of Retailing: How Analytics are Transforming the Supply Chain and Improving Performance*. Harvard Bus. Press.

- Gallego, G., I. Moon. 1993. The distribution free newsboy problem: Review and extensions. *J. Oper. Res. Soc.* **44**(8) 825–834.
- Giffin, A., A. Caticha. 2007. Updating probabilities with data and moments. *Amer. Inst. of Phys. Conf. Proc.*, vol. 954. 74–84.
- Godfrey, G.A., W.B. Powell. 2001. An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Sci.* **47**(8) 1101–1112.
- Huh, W.T., R. Levi, P. Rusmevichientong, J.B. Orlin. 2011. Adaptive data-driven inventory control with censored demand based on kaplan-meier estimator. *Oper. Res.* **59**(4) 929–941.
- Huh, W.T., P. Rusmevichientong. 2009. A nonparametric asymptotic analysis of inventory planning with censored demand. *Math. Oper. Res.* **34**(1) 103–123.
- Jaynes, E.T. 1957. Information theory and statistical mechanics. *Physical Rev.* **106**(4) 620.
- Jaynes, E.T. 1981. On the rationale of maximum-entropy methods. *IEEE Proc.*, vol. 70. 939–952.
- Kahn, K.B. 2002. An exploratory investigation of new product forecasting practices. *J. of Prod. Innov. Management* **19**(2) 133–143.
- Kleywegt, A. J., A. Shapiro, T. Homem-De-Mello. 2002. The sample average approximation method for stochastic discrete optimization. *SIAM J. Opt.* **12**(479–502).
- Klibanoff, P., M. Marinacci, S. Mukerji. 2005. A smooth model of decision making under ambiguity. *Econometrica* **73**(6) 1849–1892.
- Klibanoff, P., M. Marinacci, S. Mukerji. 2009. Recursive smooth ambiguity preferences. *J. Econ. Theo.* **144** 930–976.
- Kreyszig, Erwin. 1988. *Advanced Engineering Mathematics*. Wiley, New York.
- Lariviere, M.A., E.L. Porteus. 1999. Stalking information: Bayesian inventory management with unobserved lost sales. *Management Sci.* **45**(3) 346–363.
- Levi, R., G. Perakis, J. Uichanco. 2010. The data-driven newsvendor problem: New bounds and insights. Working Paper, MIT.
- Levi, R., R.O. Roundy, D.B. Shmoys. 2007. Provably near-optimal sampling-based policies for stochastic inventory control models. *Math. Oper. Res.* **32**(4) 821–839.
- Liyanage, L.H., J.G. Shanthikumar. 2005. A practical inventory control policy using operational statistics. *Oper. Res. Lett.* **33** 341–348.
- Luenberger, D.G. 1994. *Optimization by Vector Space Methods, 1969*. Wiley, New York.
- Lynn, G.S., S.P. Schnaars, R.B. Skov. 1999. A survey of new product forecasting practices in industrial high technology and low technology businesses. *Indust. Market. Management* **28**(6) 565–571.
- Maglaras, C., S. Eren. 2015. A maximum entropy joint demand estimation and capacity control policy. *Production Oper. Management* **24**(3) 438–450.
- Pardo, L. 2006. *Statistical Inference Based on Divergence Measures*. Chapman and Hall/CRC, Boca Raton, FL.
- Perakis, G., G. Roels. 2008. Regret in the newsvendor model with partial information. *Oper. Res.* **56**(1) 188–203.
- Popescu, I. 2007. Robust mean-covariance solutions for stochastic optimization. *Oper. Res.* **55**(1) 98–112.
- Ramamurthy, V., George Shanthikumar, Z.-J. Shen. 2012. Inventory policy with parametric demand: Operational statistics, linear correction, and regression. *Production Oper. Management* **21**(2) 291–308.
- Saghafian, S., X. Chao. 2014. The impact of operational decisions on the optimal design of salesforce incentives. *Naval Res. Log.* **61**(4) 320–340.
- Sarf, H. 1958. A min-max solution of an inventory problem. *Studies in the Mathematical Theory of Inventory and Production* 201–209.
- Sarf, H. 1959. Bayes solution to the statistical inventory problem. *Ann. Math. Stat.* **30**(2) 490–508.
- Schwartz, L. 1965. On Bayse procedures. *Z. Wahrsch. Verw. Gebiete* **4** 10–26.

- Shapiro, A. 2003. *Monte Carlo Sampling Methods*. A. Ruszczyński, A. Shapiro, eds. *Stochastic Programming, Vol. 10 of Handbooks in Operations Research and Management Science*. Elsevier, Amsterdam, The Netherlands.
- Shen, X., L. Wasserman. 2001. Rates of convergence of posterior distributions. *Ann. of Stat.* **29**(3) 687–714.
- Tomlin, B. 2009. Impact of supply learning when suppliers are unreliable. *Manufacturing Service Oper. Management* **11**(2) 192–209.
- Wang, Z., P.W. Glynn, Y. Ye. 2013. Likelihood robust optimization for data-driven problems. Working Paper, University of Minnesota.
- Wang, Z., A.J. Mersereau. 2013. Bayesian inventory management with potential change-points in demand. Working Paper, University of North Carolina.
- Zhu, Z., J. Zhang, Y. Ye. 2013. Newsvendor optimization with limited distribution information. *Opt. Meth. and Soft.* **28** 640–667.