The Newsvendor under Demand Ambiguity-Aversion: A Data and Information-Driven Approach

Soroush Saghafian, Brian Tomlin

The Newsvendor under Demand Ambiguity-Aversion: A Data and Information-Driven Approach[1]

1 Introduction

Unpredictable demand is a major challenge in sales and operations planning. One paradigm in the operations literature endows the decision maker with complete information about the demand distribution. Full distributional information is a strong assumption but is prevalent in the literature, nevertheless, and is typically referred to as decision making under *uncertainty*. An alternative "data-driven" paradigm endows the decision maker only with demand observations based on past history or simulation (e.g., Azoury 1985, Liyanage and Shanthikumar 2005, Levi et al. 2007, Huh et al. 2011). Reality may lie somewhere in between these two paradigms. The decision maker may face a type of *ambiguity* in which he has access to past observations and to partial - but not full - information about the demand distribution. For instance, firms typically conduct market research prior to a product launch or have access to past sales of similar products, and so production planners have certain beliefs about the demand distribution at launch. Post launch, planners may periodically update their beliefs based on demand observations. Planners may also account for other sources of information - separate from the demand observations - when refining their beliefs. For example, the marketing department may provide insights derived from customer interactions, from analysis of website clickstream data (Huang and Van Mieghem 2014), or other from relevant sources.

With the rapid growth in the type and amount of data available to companies, it is becoming increasingly important for managers to incorporate multiple sources of information in their decisions. Motivated by such data-driven business analytics, we consider a newsvendor-type setting in which a decision maker does not know the demand distribution but has two evolving sources of information about it. One source is the history of observations up to each period (the "data"). The other source is pre- and post-launch research (from marketing, for example) that generates additional information in the form of upper and/or lower bounds on the moments or tails of the distribution. Similar to the first source, this second source may evolve over time as ongoing research generates revised information. We propose a second order, non-parametric, maximum-entropy based approach, which we term SOBME (Second Order Belief Maximum Entropy) that allows the decision maker to effectively combine these two sources of information dynamically over time to update its beliefs about the possible distributions governing the demand realizations. We start by considering this approach in an ambiguity-neutral scenario. We then extend our results and demonstrate how the SOBME approach can be also used when the decision maker (DM) is not ambiguity-neutral. To this latter end, we adopt the general decision-making framework proposed by Klibanoff et al. (2005) and Klibanoff et al. (2009).

[1]This is an extended abstract. The full paper is available upon request.

2 Methodology

The core concept of maximum entropy, widely used in estimation and information theory, updates beliefs so that (a) the posterior coincides with the prior as closely as possible, and (b) only those aspects of beliefs for which new evidence was gained are updated (Jaynes 1981, Cover and Thomas 1991). In the approach we propose, the decision maker updates its beliefs so as to minimize the relatively entropy to its prior belief subject to the observed data and other available information. We derive a closed-formed solution for the belief updating mechanism using an optimization framework, and show that our proposed approach contains traditional Bayesian updating (both parametric and nonparametric) as a special case. In other words, our proposed approach can be viewed as an extension of traditional Bayesian updating that, in addition to observations/data, allows the decision maker to incorporate supplementary information such as moment and tail upper and/or lower bounds. This supplementary information is incorporated in the updating mechanism through an *exponential modifier*. We then show the suitability of our approach for a repeated newsvendor setting under demand ambiguity-aversion.

3 Main Contributions and Key Results

There are many settings in which a decision maker does not have full information about the demand random variable but may have some partial distributional information in the form of bounds on moments or tails. Over time, of course, the decision maker gains access to demand observations. Techniques which rely exclusively on partial distributional information ignore the value of demand observations. In contrast, pure data-driven techniques avail of observations but ignore partial distributional information. With the growing attention to data-driven business analytics, it is essential to overcome this gap and develop approaches that can allow decision makers to effectively and efficiently benefit from all sources of data and information in their decision makings routines.

Under ambiguity-neutrality, we explore the asymptotic behavior of our updating mechanism and show that it is weakly consistent. That is, under some "regularity" conditions, the updated distribution converges to the true unknown distribution almost surely in the weak neighborhoods of it (i.e., distributions that are "close" to the true distribution when closeness of distributions is measured in terms of certain expectation-like distances). We also analytically establish a rate of convergence for the proposed approach, which bounds the distance between the density built using the proposed approach and the true (unknown) density. We numerically compare our approach to (a) the case in which the decision maker knows the true distribution, and (b) the case in which the decision maker is limited to a pure data-driven approach (PDD), in which the decision maker avails of the past history (observations/data) but not additional information, such as moment or tail bounds. PDD approaches are limited, by definition, to only using data and observations and cannot avail of any additional information.

In addition to data, the proposed SOBME approach can accommodate dynamically evolving partial information, which may occur in a variety of applications. We establish that the proposed approach is a natural extension to the traditional Bayesian framework, in which an

exponential modifier is introduced to incorporate the partial distributional information. Our extensive numerical investigation demonstrate that SOBME performs well; not just outperforming a pure data-driven approach but often coming close to the performance of the full information case where the true distribution is completely known. We also find that unlike a pure data derive approach such as Sample Average Approximation (SAA), SOBME's performance does not significantly degrade as the newsvendor quantile increases, making it a suitable candidate for systems with high service level requirements, or for systems in which estimating cost parameters (and hence the optimal newsreader quantile) is subject to estimation errors. We show that SOBME is weakly consistent and analytically characterize its rate of convergence.

Under ambiguity-aversion, we find 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 fractile result. We also demonstrate that a higher level of ambiguity aversion results in a lower ordering quantity. Furthermore, as the level of ambiguity-aversion approaches infinity, the DM only considers the worst-case outcome, i.e., becomes a minimax optimizer, and hence orders nothing so as to avoid any overage cost. Regardless of the DM's ambiguity aversion level, as more and more demand observations are made, 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 also provide much sharper characterization of order quantities in certain cases of interest under demand ambiguity aversion. In particular, under (i) low but positive ambiguity aversion, or (ii) a normality assumption on the disutility distribution, we derive closed-form solutions for the ambiguity-averse newsvendor problem by applying the *certainty equivalence* principle.

References

Azoury, K.S. 1985. Bayes solution to dynamic inventory models under unknown demand distribution. *Management Sci.* **31**(9) 1150–1160.

Cover, P., J. Thomas. 1991. *Elements of Information Theory*. John Wiley and Sons.

Huang, T., J. Van Mieghem. 2014. Clickstream data and inventory management: Model and empirical analysis. *Production Oper. Management* **23**(3) 333–347.

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.

Jaynes, E.T. 1981. On the rationale of maximum-entropy methods. *IEEE Proc.*, vol. 70. 939–952.

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.

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. Let.* **33** 341–348.