

A decision support system for mean–variance analysis in multi-period inventory control



Preetam Basu^{a,*}, Suresh K. Nair^{b,1}

^a Operations Management Group, Indian Institute of Management Calcutta, Diamond Harbor Road, Kolkata 700104, India

^b Operations and Information Management, School of Business, University of Connecticut, Storrs, CT 06269, USA

ARTICLE INFO

Article history:

Received 5 April 2013

Received in revised form 15 September 2013

Accepted 18 September 2013

Available online 27 September 2013

Keywords:

Stochastic dynamic programming

Risk-reward heuristic

Mean–variance analysis

Efficient frontier analysis

Inventory management

ABSTRACT

Traditionally inventory management models have focused on risk-neutral decision making with the objective of maximizing the expected rewards or minimizing costs over a specified time horizon. However, for items marked by high demand volatility such as fashion goods and technology products, this objective needs to be balanced against the risk associated with the decision. Depending on how the product performs vis-à-vis the seller's original forecast, the seller could end up with losses due to either short or surplus supply. Unfortunately, traditional models do not address this issue. Stochastic dynamic programming models have been extensively used for sequential decision making in the context of multi-period inventory management, but in the traditional way where one either minimizes costs or maximizes profits. Risk is implicitly considered by accounting for stock-out costs. Considering risk and reward simultaneously and explicitly in a stochastic dynamic setting is a cumbersome task and often difficult to implement for practical purposes, since dynamic programming is designed to optimize on one variable, not two. In this paper we develop an algorithm, *Variance-Retentive Stochastic Dynamic Programming* that tracks variance as well as expected reward in a stochastic dynamic programming model for inventory control. We use the mean–variance solutions in a heuristic, RiskTracker, to construct efficient frontiers which could be an ideal decision support tool for risk-reward analysis.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Inventory control plays a critical role in day-to-day business operations and is often a crucial differentiator in determining the success or failure of firms. Traditionally inventory control models have focused on risk-neutral decision making where the usual optimization criteria have been either maximizing the sum of discounted rewards or minimizing the sum of accumulated costs over a specified time horizon. Starting from as early as the 1950's operations researchers have studied inventory control models under various economic and market conditions. Arrow et al. [5] and Dvoretzky [20] were the first to analyze a single-period inventory control model under stochastic demand which became popularly known as the newsvendor model. We refer the reader to Khouja [32] for a comprehensive review of the classical newsvendor model and its many extensions. Subsequently the single period stochastic inventory model was extended to multiple periods. The well-known (s,S) policy was proposed where an order is placed to bring the inventory level to S whenever the level fell below s. Significant contributions in this line of inquiry were made by Karlin and Fabens [30], Iglehart [28], Veinott [53] and Sethi and Cheng [47]. Recent papers

by Caro [14], Jain [29], Sana [42–44] and Xu [60] have advanced the extant literature in this field. The focus of the majority of these models is on optimizing the average reward criteria. However, using expected total reward criteria may yield optimal policies that are unacceptable to a risk-sensitive decision maker. For products marked by high demand variability such as fashion goods or technology products, where on the one hand there are risks associated with unsold inventory and on the other potential loss of revenues due to shortages, the variability of the rewards is as important as the expected values. Implicitly considering stock-out costs as a measure of risk is not sufficient in these cases. Instead of identifying one policy to achieve the stated objective, managers are often interested in considering risks more explicitly and obtaining sets of policies at different levels of risk.

There has not been much research done on risk-sensitive inventory management. The limited literature in this domain mainly focuses on single-period models, namely, newsvendor and its various extensions (see reviews by Khouja [32] and Qin et al. [39]) in a single period setting. However, in many managerial scenarios, multiple periods of ordering is involved. Based on demand and available stock-on-hand, firms place orders over multiple periods. Even in the context of fashion supply chains, to take advantage of more accurate demand information, firms often split their orders into an early order and some late orders based on market indicators (Tang et al. [50]). One of the methodologies that has been extensively used in solving these multi-period inventory problems is stochastic dynamic programming. These models provide optimal

* Corresponding author. Tel.: +91 33 2467 8300.

E-mail addresses: preetamb@iimcal.ac.in (P. Basu), suresh.nair@business.uconn.edu (S.K. Nair).

¹ Tel.: +1 860 486 1727.

sequential decisions where present ordering decisions are taken in consideration of future outcomes. At a specified point in time, which is referred to as a decision epoch, the decision maker observes the state of the system and based on that chooses the order size. This action produces an immediate reward and the system moves to a different state in the next time-period according to a probability distribution. Dynamic programming chooses actions based on reward (whether profit or cost), and does not track the risk associated with the optimal decision, resulting in a point in the risk continuum that corresponds to the maximum expected reward. The contribution of this paper is in developing a novel methodology to track both mean and variance of rewards of a set of ordering policies in the stochastic dynamic programming model. We call our methodology *Variance-Retentive Stochastic Dynamic Programming* (Variance-Retentive SDP). We use the resulting mean–variance solutions in a simple heuristic, which we call RiskTrackr, for creating efficient risk–reward frontiers, similar to those used in portfolio analysis in finance literature (Voros [54]; Markowitz [35,36]; Elton et al. [22]). This is a challenging task from an implementation standpoint, since this requires carrying information on both risk and reward simultaneously for each state, in which standard dynamic programming is not designed to do. The Variance-Retentive SDP algorithm and the RiskTrackr heuristic can be used in practice as a decision support tool for mean–variance analysis in multi-period inventory management systems.

Risk–reward trade-offs are an essential component of inventory decisions. The variability of the possible outcomes often plays an important role in determining the “best” set of ordering decisions. Financial planning models often involve systematic trade-off analysis between an expected return criterion and the variability or the risk associated with the returns. Variance of the outcomes about the expected value is a widely used measure of risk in portfolio theory. Investors use variance to measure the risk of a portfolio of stocks. The basic idea is that variance is a measure of volatility and the more a stock's returns vary from the stock's average return, the more volatile is the stock. Portfolios of financial instruments are chosen to minimize the variance of the returns subject to a level of expected return or vice versa to maximize expected return subject to a level of variance of the return. This paradigm was first introduced by Markowitz's [35] mean–variance analysis for which contribution he was honored with the Nobel Prize in Economics. Mean–Variance analysis has become a standard tool in portfolio management (Fama [23]; Copeland and Weston [19]). In many operational decisions also managers are interested in the mean–variance trade-offs. In inventory control, especially of items marked by high demand volatility, it is extremely important to ascertain the variability associated with a set of policies rather than just the expected return. If the variance of the outcome is large, the chance of deviating from the expected return will also be high. In this paper we define risk as the volatility associated with the outcomes of each of the policies and we measure risk by the variance of the possible outcomes.

In a single period newsvendor model, it is easy to enumerate the mean and the variance of the profit values for different order-sizes. In Fig. 1a we present the mean–variance solutions for a single period newsvendor model on an efficient frontier. The newsvendor optimal solution is obtained by a trade-off between cost of under-stocking and overstocking and maximizes the expected reward. But the optimal solution also has very high risk, as seen by the variance of the solution in the graph. If a decision maker is not comfortable with that level of risk, he may be well served by choosing an alternate solution on the efficient frontier to the left of the newsvendor optimal solution that has lower reward but comes with lower risk. However, traditional newsvendor solutions do not present the decision maker with this choice, since it presents only one “optimal” answer. Similar logic can be applied for the mean–variance solutions for multi-period inventory control models also. However, for multi-period stochastic dynamic programming formulations, enumerating the mean and variance of the profit values for different order sizes becomes an impossible task as state space increases exponentially. In Fig. 1b we present the mean–variance efficient frontier for a multi-period inventory model with five time-periods using Monte Carlo simulation to illustrate the number of sample paths possible and the presence of an efficient frontier. The model that we use for the numerical analysis later in Section 5 has ten time periods and 10^5 possible states, and in each state there are 16 possible actions. This problem has millions of possible policies and even higher number of sample paths. In problems of such dimension it is impossible to enumerate all the possible sample paths and policy tables. Further, dynamic programming models need to be modified completely to track both risk and reward, and this is not an obvious extension of the basic methodology of recursion. In such practical scenarios, we propose the Variance-Retentive SDP algorithm and the RiskTrackr heuristic to construct near-optimal efficient frontiers. The bold curve in Fig. 1b is the efficient frontier obtained by using the RiskTrackr heuristic.

The remainder of the paper is organized as follows: Section 2 provides a literature review. Section 3 develops the analytical model. We present the Variance-Retentive SDP algorithm and the RiskTrackr heuristic for obtaining risk–reward curves in Section 4. Managerial insights are given in Section 5. Finally, we make concluding remarks in Section 6.

2. Literature review

Our main contribution in this paper is in the field of mean–variance analysis in inventory control. Most of the research in this stream considers single-period inventory models. Lau [33] analyzes the classical newsvendor model under two different objectives namely, maximizing the decision-maker's expected utility of total profit and maximizing the probability of a certain level of profit. Chung [17] deduces an algorithm for determining optimal stocking policies for a risk-averse decision maker. Bouakiz and Sobel [12] adopt the utility function approach to characterize




Fig. 1. Mean–variance efficient frontiers.

Download English Version:

<https://daneshyari.com/en/article/552640>

Download Persian Version:

<https://daneshyari.com/article/552640>

[Daneshyari.com](https://daneshyari.com)