

How to (and How Not to) Estimate the Salvage Value in the Newsvendor Model*

Gérard P. Cachon
The Wharton School
University of Pennsylvania
Philadelphia PA, 19104
cachon@wharton.upenn.edu
opim.wharton.upenn.edu/~cachon

A. Gürhan Kök
Fuqua School of Business
Duke University
Durham NC, 27708
gurhan.kok@duke.edu
www.fuqua.duke.edu/faculty/alpha/kok.htm

July 25, 2002, revised June 2003, January 2004

Abstract

Despite the importance of the Newsvendor model, little attention has been paid to its robustness with respect to the estimation of one of its inputs, the salvage value. This oversight is inappropriate because we show that the performance of the Newsvendor model is quite sensitive to the particular method chosen to estimate the salvage value. We highlight two reasons for this sensitivity: small errors in the estimate of the salvage value can lead to a significant profit loss and, more interestingly, there exists a dependence between the estimation method used and the resulting data obtained. To explain, the *ex ante* salvage value inputted into the model determines the order quantity which then results in some *ex post* observed salvage value. An estimation method should be evaluated in terms of its equilibrium performance, i.e., the inputted salvage value equals the expected observed salvage value. We evaluate several intuitive estimation methods and find that they perform poorly in equilibrium. We also identify an estimation method that yields the optimal profit. We conclude that the Newsvendor model is robust to errors in the estimation of the salvage value as long as the proper estimation method is utilized.

*The authors would like to thank Awi Federgruen, Marshall Fisher, Ananth Iyer, Anton Kleywegt, Martin Lariviere, Erica Plambeck, Sridhar Seshadri, Jim Smith, Matthew Sobel, Jing-Sheng Song, Christian Terwiesch, Sridhar Tayur, Jan Van Mieghem, Garrett van Ryzin, Eitan Zemel and Paul Zipkin for their helpful comments, as well as the seminar participants at Carnegie Mellon and New York University. This paper is available via the authors' webpages. The previous version of this paper was titled "Heuristic Equilibrium and the Estimation of the Salvage Value in the Newsvendor Model with Clearance Pricing".

The newsvendor model is certainly among the most important models in operations management. It is applied in a wide variety of areas: centralized and decentralized supply chain inventory management (e.g., Shang and Song 2003, Cachon 2003), retail assortment planning (e.g., van Ryzin and Mahajan, 1999), international operations (e.g., Kouvelis and Gutierrez 1997), horizontal competition among firms facing stochastic demand (e.g., Lippman and McCardle, 1995), lead time competition (e.g., Li 1992), outsourcing and subcontracting decisions (e.g., Van Mieghem 1999), product and process redesign (Fisher and Raman 1996 and Lee 1996), and spot markets and inventory control (e.g., Lee and Whang 2002) to name a few. It is taught in most introductory courses in operations management and it is described in detail in most operations management text books.

The newsvendor model is not complicated: the only decision is an order quantity, the purchase cost per unit is c , units are sold during a selling season for a fixed price, p , demand is stochastic during the selling season with a known distribution, sales are bounded by the order quantity and left over inventory is salvaged at the end of the season for a fixed salvage value, v . While there is little uncertainty associated with the cost and price parameters (c and p), and there is a substantial literature on how to choose the demand distribution, there has been surprisingly little research conducted on the remaining input to the model, the salvage value. In many applications of the model the salvage value is *not* a fixed amount received from left over inventory, but rather, it is the output of a clearance pricing decision: if there is little inventory left at the end of the season, the salvage value is likely to be high (a small discount is needed with a popular product in short supply) whereas if there is plenty of inventory at the end of the season, the salvage value is likely to be low (a deep discount is needed with an unpopular product in ample supply). As a result, it is not clear how historical data could be used to estimate the salvage value, nor whether the procedure for estimating the salvage value matters in terms of profits.

This paper studies how a manager should use historical data to estimate the salvage value when applying the newsvendor model to an actual decision. Besides the stature of this model in the operations literature, we offer three reasons why this is an important research topic. First, a small error in the salvage value estimate can translate into a significant profit loss, especially if the salvage value is overestimated. This point is illustrated by the data reported in Table 1 from a simple numerical example: e.g., if demand is gamma distributed,

the gross margin is 25%, the difference between the production cost and the true salvage is 8% of the sales price, and the coefficient of variation is 0.71, then using a salvage value that is 12% higher than the true salvage value results in a profit loss of nearly 60%. In contrast, using an order quantity that is 12% higher than optimal only results in a 1% profit loss. The well known EOQ model also provides an interesting contrast: using an order quantity or setup cost that is 12% greater than the true value increases overall costs by only 0.2% (with either parameter). (See Dobson 1988 for a more detailed discussion of the robustness of the EOQ model to errors in the input parameter estimates).

Our second reason is that there are several intuitively reasonable methods for evaluating the salvage value, but they can result in significantly different estimates, thereby leading to considerable variation in performance. In other words, there is more than one way to look at historical data to derive a salvage value, and, due to sensitivity even to small estimation errors, the estimation technique chosen can have a significant impact on profit.

Our third reason, and most interesting one, is that the estimation of the salvage value cannot be viewed in isolation of the order quantity decision. To explain, consider the following quote that describes the economics of selling fashion ski apparel faced by Sport Obermeyer (Hammond and Raman 1994): "...units left over at the end of the season were sold at a loss that averaged 8% of the wholesale price." That 8% figure was most likely determined with historical data, and it surely depends on how much Sport Obermeyer ordered in the past. For example, had Sport Obermeyer ordered twice as much as they did, then it is reasonable to expect that their historical losses would have been higher (they would have had more parkas to discount). As a result, order quantities determine historical data used to estimate salvage values, which are used to set future order quantities. (We presume here, as with Sport Obermeyer, that the newsvendor model is applied over multiple occasions that are similar in characteristics. For example, clearance pricing decisions are made in the same manner over successive seasons so a set of seasons can be viewed as a set of independent application of the newsvendor model to the same problem, just different realizations of the stochastic variables in each season.)

This relationship between order quantities, data and estimated salvage values raises several questions. For a given salvage value estimation procedure is there a stable equilibrium in terms of inputted salvage values and order quantities? Put into the Sport Obermeyer context,

if they choose order quantities based on their 8% expected loss, then will they observe in their future data 8% losses as well? If there exists this salvage value-order quantity equilibrium, then is it unique, and if unique, then is it optimal? Initial intuition suggests that it should be: data are used to choose actions and the actions result in similar data, so input data and actions are consistent. If not optimal, one would suspect that it is nearly optimal. But in fact, we demonstrate that a substantial profit loss can occur even if there is consistency between the salvage value and the order quantity.

It is important to note that this notion of robustness is different than one based on sampling error alone (as in the data in Table 1, or as in Dobson 1988). We evaluate the performance of each estimation method in the most optimistic scenario that there is no sampling error, i.e., the evaluated salvage value is assumed to be the salvage value that would be estimated with an unlimited sample of data. Nevertheless, because of the dependence between the estimation method and the chosen action, the estimation method can lead to order quantities that deviate substantially from optimal, thereby causing a significant profit loss. To the best of our knowledge, there has been no other work in inventory theory that evaluates robustness of an inventory model in this manner.

Although a cavalier implementation of the newsvendor model is not prudent, this does not mean it is ineffective. We provide an estimation procedure that indeed leads to the optimal solution in equilibrium. Hence, the simple and parsimonious newsvendor model can be effectively applied in situations with clearance pricing.

The next section provides an example to illustrate our ideas and further motivates this research. The subsequent section defines our model, §3 reviews the related literature, §4 identifies the optimal procurement quantity, §5 defines and analyzes several salvage value estimation rules, §6 presents some numerical results and the final section discusses our results.

1 A motivating example

Suppose a manager wishes to use the newsvendor model to choose an order quantity for an item in advance of its selling season. The manager forecasts that demand follows a gamma distribution with mean 1000 and standard deviation 707. The regular selling price is $p = \$200$ and the purchase cost per unit is $c = \$150$. Left over units at the end of the season are liquidated via a markdown that depends on the number of units left over. To

assess the salvage value the manager looks at historical data of similar items for which the newsvendor model was used to choose an order quantity and units were salvaged. Table 2 provides $n = 10$ observations. In each observation the order quantity was 1400 units. The table provides for the j^{th} observation the amount of inventory left over at the end of the regular season, y_j , (i.e., the inventory at the start of the clearance period), the total revenue earned during the clearance period, t_j , and the revenue earned per unit, $u_j = t_j/y_j$. In some cases the revenue earned per unit in the clearance period is above the purchase cost, \$150, because only a small discount is needed to liquidate the few units in inventory.

How should a manager use these data to estimate a salvage value? Interestingly, we were unable to find any reference that offers a precise method for estimating the salvage value, nor is there any acknowledgment in the literature that there could be multiple methods for evaluating the salvage value or that this estimation must be done carefully. All of the definitions we found read like “the salvage value is the amount earned on each unit left over at the end of the season”. One interpretation of this statement is what we call the *average salvage value*: let \hat{v}_a be the salvage value inputted to the newsvendor model, where \hat{v}_a is the average of the revenue per unit across the observations,

$$\hat{v}_a = \frac{1}{n} \sum_{j=1}^n u_j.$$

For the sample in Table 2, $\hat{v}_a = \$130$. However, if the salvage value is 130 then the order quantity recommended by the newsvendor model is 1254, not 1400 (the order quantity used to collect the data). Furthermore, the model chooses the order quantity 1400 only if the inputted salvage value is \$135. One might suspect that the difference between the observed \$130 and the required \$135 is due to sampling error, but it is not: in these examples the expected salvage value is \$131 when the order quantity is 1400.

The problem with this situation is that our current salvage value-order quantity pair, $\{130, 1400\}$ is not in equilibrium: if the salvage value 130 is inputted to the newsvendor model then the recommended order quantity is not 1400, and if the order quantity 1400 is chosen then the observed clearance period revenue per unit is not \$130. There indeed exists an equilibrium in our example, $\{133, 1346\}$, and it happens to be the unique equilibrium: if 1346 units are ordered then the expected average salvage value is \$133, and if \$133 is the inputted salvage value then 1346 units is the order quantity recommended by the newsvendor

model. But that is where the good news ends: the optimal order quantity is 941 units, and the profit with the equilibrium order quantity is 22% less than the optimal profit. Therefore, even if there is consistency between the inputted salvage value and the observed salvage value, the performance of the model can be grossly suboptimal. In fact, if the manager were to use the average salvage value as the input to the newsvendor model, then there is no hope to reach the optimal order quantity: with an order quantity of 941 units the expected average salvage value is \$157, which is greater than the production cost! Hence, the optimal solution generates data that leads to infeasible inputs to the newsvendor model.¹ We clearly need to consider other methods for estimating the salvage value.

2 The newsvendor and the clearance pricing models

We define in this section the two models we study. We assume a manager uses the newsvendor model to make an order quantity decision and inputs into the model some salvage value estimated with historical data. The newsvendor model correctly represents reality with only one exception, the revenue earned per unit on left over inventory is not a fixed value but rather it is the result of a markdown pricing decision. Therefore, to assess the performance of the newsvendor model we analyze the “clearance pricing model”, a model that is identical to the newsvendor model with the exception that it explicitly incorporates the clearance pricing decision.

In the newsvendor model a firm purchases q units before a single selling season with random demand and pays c per unit. There are no constraints on q (i.e., no capacity constraint), but only a single procurement is feasible. The selling season is divided into two periods. In period 1, called the regular season, the retailer sells each unit for $p_1 > c$. In period 2, the clearance period, the retailer sells all remaining inventory for v per unit, $v < c$. Let $\xi \in (0, \infty]$ be the realization of demand in period one. Let $F(\cdot)$ be the strictly increasing

¹ When teaching the newsvendor model to MBA students we are often asked what should a manager do if the salvage value is greater than the production cost. The standard answer is either that this is not possible (at least on an exam question) or the manager should order more, in which case the salvage value will eventually decrease below the production cost and then the model applies again. But that is incorrect, as this example indicates.

and differentiable distribution function of demand and let $f(\cdot)$ be the density function. The objective in the newsvendor model is to choose an order quantity q to maximize expected profit.

We now describe the clearance pricing model, which is almost identical to the newsvendor model. The difference is that at the start of the clearance period the firm chooses a clearance period price, p_2 , i.e., it is not automatic that all units are salvaged for v per unit. We assume $p_2 \leq p_1$, which is reasonable in some situations and *de facto* imposed by many firms.

Period 2 demand, $D_2(p_2, \xi)$, is a deterministic function of the clearance price and the realization of period one demand. $D_2(p_2, \xi)$ is a non-negative and differentiable function that is decreasing in p_2 . Hence, the inverse demand function exists, $p_2(s_2, \xi)$. The assumption that D_2 depends on ξ is reasonable if total sales are highly correlated with early season sales, for which there is empirical evidence (Fisher and Raman 1996, and Fisher, Rajaram and Raman 2001).

Let s_2 be period 2 sales. We make the following two technical assumptions: period 2 revenue, $s_2 p_2(s_2, \xi)$ is concave in s_2 for all ξ ; and $\hat{p}_2(\xi) < p_1$ for all ξ , where $\hat{p}_2(\xi) = \arg \max(p_2 D_2(p_2, \xi))$, i.e., the period 2 revenue maximizing price is no larger than the period 1 price. (For expositional simplicity, in all references to concavity we mean strict concavity.) The first assumption ensures the firm's profit function is well behaved in q . The second merely implies markdowns are possible: if $\hat{p}_2(\xi) \geq p_1$, then in the clearance period the firm would always choose $p_2 = p_1$ (given the $p_2 \leq p_1$ constraint) because the optimal clearance price is never less than \hat{p}_2 .

$D_2(p_2, \xi)$ is also monotone in ξ for all p_2 . While it is natural to think of $D_2(p_2, \xi)$ as an increasing function of ξ (a product with high regular season demand also has high clearance period demand), we also allow $\partial D_2(p_2, \xi)/\partial \xi = 0$ (i.e., regular season and the clearance period demands are independent) and $\partial D_2(p_2, \xi)/\partial \xi < 0$ (i.e., high regular season demand saturates the market, thereby lowering demand in the clearance period). However, for tractability, $\partial D_2(p_2, \xi)/\partial \xi$ cannot be too negative: we require that $\xi + D_2(\hat{p}_2(\xi), \xi)$ and $\xi + D_2(p_1, \xi)$ are increasing in ξ for all p_2 and ξ . These conditions imply that total demand across the two periods increase in ξ .

One form of $D_2(p_2, \xi)$ that meets these requirements can be constructed by using a multiplicative shock $x(\xi)$ in combination with a commonly used demand function such as the

constant elasticity demand function, $D_2(p_2, \xi) = x(\xi)\alpha p_2^{-\beta}$, or the exponential demand function, $D_2(p_2, \xi) = x(\xi)\alpha e^{-\beta p_2}$ for $1/\beta < p_1$. (The condition on β with exponential demand ensures $\hat{p}_2 < p_1$.) There is substantial empirical evidence to support both demand forms, and both have been observed to fit actual data better than linear demand; see Mulhern and Leone (1991), Hoch, Byung-Do, Montgomery, and Rossi (1995) and Tellis (1988).

In the clearance period the firm sets p_2 to maximize period 2 revenue given the available inventory, $I(q, \xi)$. Without loss of generality, inventory left over at the end of the clearance period has zero salvage value. The firm chooses q before the regular season to maximize expected profit.

3 Literature Review

The literature related to this research can be divided into several broad categories: papers that discuss variations on the newsvendor model; papers on pricing without multiple inventory replenishments; research on multi-period pricing and inventory problems; and research on the robustness of heuristics, especially as applied to inventory models.

A number of papers enrich the newsvendor model along one or more dimensions. Instead of a loss function that is linear in the excess inventory, Porteus (1990) considers a loss function that is quasi-convex in the excess inventory quantity. He provides conditions under which the objective function is well behaved. We demonstrate that in the clearance pricing model the loss function is convex. These models *assume* the non-linear salvage value function is known and accurate, i.e., there is no discussion of how that function could be estimated or how sensitive the performance of the model is to that estimation. Petruzzi and Dada (1999) and Agrawal and Seshadri (2000a) study a newsvendor that chooses both a quantity and a price, but in both cases the newsvendor chooses the regular season price, not the clearance price; they assume a fixed salvage value for inventory remaining at the end of the regular season. In Carr and Lovejoy (2000) the newsvendor also makes multiple decisions, but their newsvendor chooses which customers to serve (each with its own demand distribution) given the newsvendor's fixed capacity. In Dana and Petruzzi (2001) the newsvendor's demand depends on the procurement quantity: more inventory leads to a better fill rate which increases demand. Hence, their model, like ours, has an interdependence between input parameters (the forecasted demand distribution) and the action (quantity). They show

that a unique equilibrium exists if that interdependence is ignored and the procurement quantity in equilibrium is lower than optimal. In contrast, for several of our salvage value estimation methods the firm procures too much in equilibrium and for one of our methods the firm procures the correct amount. In addition, they consider a newsvendor that sets the regular season price and they do not evaluate when the equilibrium leads to a significant loss in profit. The following papers provide other extensions to the newsvendor model that are not closely related to this work: Eeckhoudt, Gollier and Schlesinger (1995), Lippman and McCardle (1995), Schweitzer and Cachon (2000), van Mieghem and Rudi (2001).

Hertz and Schaffir (1960) recognize that the salvage value of clearance inventory depends on the amount of inventory, but then argue that a constant salvage value is an adequate approximation. They do not provide a method for estimating that salvage value.

There are several papers that study a two period version of the newsvendor model with fixed salvage values: Donohue (2000), Fisher and Raman (1996), Fisher, Rajaram and Raman (2001), Kouvelis and Gutierrez (1997), and Petruzzi and Dada (2001). In each case the second period allows a second replenishment, which we do not have. With the exception of Petruzzi and Dada (2001), in each case prices are exogenous. Our model is a special case of Petruzzi and Dada (2001). However, their focus is on a solution procedure for their more complex model whereas our focus is on the robustness of the simple newsvendor model.

There are numerous papers that study revenue management and/or markdown pricing: e.g., Bitran and Mondschein (1997), Bitran, Caldentey and Mondschein (1997), Brumelle et al. (1990), Federgruen and Heching (1997), Feng and Gallego (1995), Gallego and van Ryzin (1994), Monohan, Petruzzi and Zhao (2002) and Smith and Achabal (1998). With the exception of Brumelle et al. (1990), these papers assume demand is independent across time whereas we allow for correlation in demand. Furthermore, their focus is on optimization of a given model without concern for how the model's inputs are determined or whether a simple model can provide an optimal solution.

There are a number of papers that study the robustness of heuristics with inventory models. Dobson (1988) studies the consequence of using incorrect cost parameters due to estimation errors in the classic EOQ model. Lovejoy (1990) shows that myopic optimal policies can be optimal or near-optimal in some dynamic inventory models with parameter adaptive demand processes. Bounds for the (r, Q) inventory policy when a simplifying

heuristic is used to choose the order quantity, Q , are provided by Zheng (1992) and later extended by Axsater (1996), Gallego (1998) and Agrawal and Seshadri (2000b). None of the mentioned papers considers the interaction between actions and data used to estimate input values.

In addition to Dana and Petruzzi (2001), there are two other papers that discuss the consequence of ignoring the interdependence between inputs and actions, albeit in very different settings than ours: Armony and Plambeck (2002) consider demand forecasting in a supply chain in which customers may submit duplicate orders; and Cachon, Terwiesch and Xu (2002) study assortment planning with consumer search.

Finally, the interplay between estimation and controls is a constant theme in stochastic optimal control: estimation modules that produce a unique and consistent input for each realization of the random factor are considered and their existence is assumed. See for example Bertsekas (2000). In contrast to the newsvendor model in this paper, the optimal control models are the best representation of reality and what needs to be estimated (the current state) is well defined. Another difference is that we seek the consistency of inputs and actions at the expectation level rather than for every realization of the random factor.

4 Optimal procurement and clearance pricing

In this section we evaluate the optimal decision in each model, beginning with the newsvendor model. The newsvendor model expected profit is

$$\pi(q) = -cq + r_1(q) + r_2(q)$$

where $r_i(q)$ is expected revenue in period i ,

$$r_1(q) = p_1 \left(q - \int_0^q F(\xi) d\xi \right)$$

Period 2 expected revenue is assumed to be $r_2(q) = vI(q)$, where v is the fixed salvage value per unit and $I(q)$ is expected left over inventory

$$I(q) = \int_0^q I(q, \xi) dF(\xi) = \int_0^q F(\xi) d\xi.$$

where $I(q, \xi) = (q - \xi)^+$. The newsvendor model chooses q to maximize $\pi(q)$,

$$q = F^{-1} \left(\frac{p_1 - c}{p_1 - v} \right) \tag{1}$$

where $F^{-1}(\cdot)$ is the inverse distribution function. From the above we can derive the function $v_n(q)$, which is the salvage value such that q is the optimal quantity with the newsvendor model:

$$v_n(q) = p_1 - \frac{p_1 - c}{F(q)}. \quad (2)$$

Let q^o be the true optimal order quantity (of the clearance pricing model). The newsvendor model recommends q^o as long as $v_n(q^o)$ is the inputted salvage value. In other words, there is nothing that prevents the newsvendor model from finding the optimal quantity. All that we need for that to happen is a method for consistently finding the correct salvage value, $v_n(q^o)$, to input to the model.

In the clearance pricing model there are two decisions, the initial order quantity and a clearance price function that depends on the amount of inventory at the start of the clearance period. We derive the optimal policy in three stages. We first establish that clearance period revenue is concave in the remaining inventory. Next, we show there exists three threshold functions that partition the regular season demand space into four intervals. The clearance period revenue function depends on which interval the regular season demand realization falls. Finally, we demonstrate the expected profit function is concave.

The period 2 price, p_2 , is chosen to maximize revenue after observing period 1 demand, ξ , and the remaining inventory $I(q, \xi)$. Due to the existence of the inverse function, $p_2(s_2, \xi)$, the equivalent decision is to choose the number of units to sell, s_2 , to maximize revenue. Let $\hat{R}_2(s_2, \xi) = s_2 p_2(s_2, \xi)$ be the unconstrained revenue function, which is concave (by our earlier assumption). Let $\hat{s}_2(\xi)$ be the unconstrained optimal period 2 sales quantity:

$$\hat{s}_2(\xi) = \arg \max_{s_2} \hat{R}_2(s_2, \xi).$$

The firm can sell s_2 units only if $s_2 \leq I(q, \xi)$. Therefore let $\tilde{R}_2(I(q, \xi), \xi)$ be the firm's maximum revenue constrained by available inventory:

$$\tilde{R}_2(I(q, \xi), \xi) = \max_{s_2} (\hat{R}_2(s_2, \xi) : s_2 \leq I(q, \xi))$$

From the Maximum Theorem under Convexity (Sundaram 1999, Ch.9), $\tilde{R}_2(I(q, \xi), \xi)$ is concave in $I(q, \xi)$ because $\hat{R}_2(s_2, \xi)$ is concave in s_2 .

The remaining constraint to consider is $p_2 \leq p_1$. Let $R_2(I(q, \xi), \xi)$ be the firm's maximum period 2 revenue given both the $p_2 \leq p_1$ and the $s_2 \leq I(q, \xi)$ constraints:

$$R_2(I(q, \xi), \xi) = \min\{p_1 I(q, \xi), \tilde{R}_2(I(q, \xi), \xi)\}$$

The minimum of two concave functions is concave, so $R_2(I(q, \xi), \xi)$ is also concave in $I(q, \xi)$.

We now turn to the evaluation of $R_2(I(q, \xi), \xi)$. There are four relevant cases for the second period revenue maximization problem based on the realization of period 1 demand. In the first case period 2 inventory is larger than the unconstrained optimal selling quantity, $I(q, \xi) \geq \hat{s}_2(\xi)$, which can occur only if $\hat{s}_2(\xi)$ is finite. In this case it is optimal in period 2 to sell $\hat{s}_2(\xi)$ and dispose of the remaining inventory at the end of the clearance period. Define $\hat{\xi}(q) \in [0, q]$ such that $\hat{\xi}(q) = 0$ if $I(q, 0) < \hat{s}_2(0)$, otherwise $\hat{\xi}(q)$ is the set of ξ that satisfy $I(q, \xi) = \hat{s}_2(\xi)$, which can be written as

$$q - \xi - D_2(\hat{p}_2(\xi), \xi) = 0. \quad (3)$$

By assumption $\xi + D_2(\hat{p}_2(\xi), \xi)$ is increasing in ξ , so (3) demonstrates $\hat{\xi}(q)$ is unique.

In the second case there is less inventory than needed to maximize the revenue in the clearance period, $I(q, \xi) < \hat{s}_2(\xi)$. So it is optimal to sell all of the remaining inventory. To do so the firm sets the period 2 price to the clearance price, $p_2(I(q, \xi), \xi)$ as long as the clearance price does not violate the $p_2 \leq p_1$ constraint. Define $\tilde{\xi}(q) \in [0, q]$ such that $\tilde{\xi}(q) = 0$ if $I(q, 0) < D_2(p_1, 0)$, otherwise $\tilde{\xi}(q)$ is the set of ξ that satisfy $I(q, \xi) = D_2(p_1, \xi)$, which can be written as

$$q - \xi - D_2(p_1, \xi) = 0 \quad (4)$$

By assumption $\xi - D_2(p_1, \xi)$ is increasing in ξ , so (4) demonstrates that $\tilde{\xi}(q) > 0$ is unique. Furthermore, if $\hat{\xi}(q) > 0$, then a comparison of (3) with (4) reveals $\tilde{\xi}(q) > \hat{\xi}(q)$ because $D_2(p_1, \xi) < D_2(\hat{p}_2(\xi), \xi)$.

The third case has $\xi > \tilde{\xi}(q)$: the optimal clearance period price is greater than p_1 , but due to the $p_2 \leq p_1$ constraint the firm must settle for $p_2 = p_1$. The fourth case has $\xi \geq q$: there is no inventory left in the clearance period so there is no clearance period pricing decision. Given that we have established $q \geq \tilde{\xi}(q) \geq \hat{\xi}(q)$ for any q , the second period revenue is

$$R_2(I(q, \xi), \xi) = \begin{cases} \hat{s}_2(\xi)p_2(\hat{s}_2, \xi) & 0 \leq \xi \leq \hat{\xi}(q) \\ I(q, \xi)p_2(I(q, \xi), \xi) & \hat{\xi}(q) < \xi \leq \tilde{\xi}(q) \\ I(q, \xi)p_1 & \tilde{\xi}(q) < \xi \leq q \\ 0 & q < \xi \end{cases}. \quad (5)$$

The next step is to identify the optimal quantity. The firm's expected profit, $\Pi(q)$, equals

the first period's profit plus the revenue from the second period:

$$\begin{aligned}\Pi(q) &= -cq + R_1(q) + R_2(q) \\ &= -cq + p_1 \left(q - \int_0^q F(\xi) d\xi \right) + \int_0^q R_2(I(q, \xi), \xi) dF(\xi)\end{aligned}$$

where $R_i(q)$ is expected revenue in period i . Note, we use upper and lower case notation to represent analogous functions in the two models. While $r_1(q) = R_1(q)$, i.e., both models agree in their evaluation of period 1 revenue, the models may disagree in their evaluation of period 2 revenue, i.e., $r_2(q) \neq R_2(q)$ is possible.

Differentiate $\Pi(q)$ ²:

$$\frac{\partial \Pi(q)}{\partial q} = (p_1 - c) - p_1 F(q) + \int_0^q \frac{\partial R_2(I(q, \xi), \xi)}{\partial I} dF(\xi) \quad (6)$$

and

$$\frac{\partial^2 \Pi(q)}{\partial q^2} = \int_0^q \frac{\partial^2 R_2(I(q, \xi), \xi)}{\partial I^2} dF(\xi) + \left(-p_1 + \frac{\partial R_2(0, q)}{\partial I} \right) f(q) \quad (7)$$

(Recall that $\partial I(q, \xi)/\partial q = 1$, $I(q, q) = 0$ and $R_2(I(q, q), q) = 0$.) The marginal second period revenue of an extra unit cannot be greater than p_1 (due to the $p_2 \leq p_1$ constraint), so the second term in (7) is nonpositive. Given that $R_2(I, \xi)$ is concave in $I(q, \xi)$, the integral in (7) is also negative. Hence, $\Pi(q)$ is concave in q . As already defined, q^o is the unique optimal procurement quantity. From (6) and (5), q^o is the unique solution to the following

$$0 = (p_1 - c) - p_1 F(q) + \left(\int_{\tilde{\xi}(q)}^{\tilde{\xi}(q)} \frac{\partial (I(q, \xi) p_2(I(q, \xi), \xi))}{\partial q} dF(\xi) + \int_{\tilde{\xi}(q)}^q p_1 dF(\xi) \right) \quad (8)$$

5 Salvage value estimation

This section defines and analyzes four methods for estimating the salvage value, v .

5.1 Average salvage value

As defined in §1, for one observation the average salvage value, $v_a(q, \xi)$, is second period revenue divided by the number of units left at the end of the first period, assuming there are

² $R_2(I, \xi)$ may not be differentiable w.r.t. I at the break points in its definition. Formally the following integral needs to be written in three separate pieces so that R_2 is differentiable in the ranges of each integral, but we omit that level of detail for brevity.

units left at the end of the regular season to markdown. We can use the clearance pricing model to evaluate the expected average salvage value, $v_a(q) = E[v_a(q, \xi)]$,

$$v_a(q) = \frac{1}{F(q)} \left(\int_0^{\hat{\xi}(q)} \frac{\hat{s}_2(\xi) \hat{p}_2(\xi)}{I(q, \xi)} dF(\xi) + \int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} p_2(I(q, \xi), \xi) dF(\xi) + \int_{\tilde{\xi}(q)}^q p_1 dF(\xi) \right) \quad (9)$$

The first integral includes outcomes in which only a portion of the inventory is liquidated, the second integral includes outcomes in which all inventory is sold at the clearing price and the third integral includes outcomes in which all inventory is sold at below the clearing price.

Let $\{v_a^*, q_a^*\}$ be an equilibrium when the firm uses $v_a(q)$ as the salvage value input to the newsvendor model: if $\{v_a^*, q_a^*\}$ is an equilibrium then $v_a^* = v_a(q_a^*) = v_n(q_a^*)$. We wish to determine whether an equilibrium exists, if so, is it unique, and if so, how does it perform, i.e., what is the relationship between q_a^* and q^o and the relationship between $\Pi(q_a^*)$ and $\Pi(q^o)$.

To help with intuition, Figure 1 displays $v_n(q)$ and $v_a(q)$ for one example. (The particular parameter values are not important.) As can be seen in the figure, an equilibrium exists and it is unique, i.e., it is the point at which the $v_n(q)$ and $v_a(q)$ functions intersect. Given that $v_n(q)$ is strictly increasing, existence and uniqueness of $\{v_a^*, q_a^*\}$ would be easy to demonstrate if $v_a(q)$ were strictly decreasing. First thought suggests that $v_a(q)$ should be decreasing, i.e., the more units ordered the lower is the expected salvage value. But, as is clear from the figure, that is not necessarily (or even generally) the case because $v_a(q)$ involves a conditional expectation: if q is quite small and we nevertheless must salvage inventory, then the demand realization must have been terribly low.

The next theorem proves uniqueness of $\{v_a^*, q_a^*\}$ by demonstrating that $v_n'(q) > v_a'(q)$ at any equilibrium. (While it appears in the figure that $v_a(q) - v_n(q)$ is decreasing everywhere, which is a sufficient condition for uniqueness, that is a more restrictive condition and it is not clear it holds in general.)

Theorem 1 *With the average salvage value there exists a unique equilibrium, $\{v^*, q_a^*\}$, and $q_a^* > q^o$, i.e., the newsvendor model with the average salvage value input procures too much.*

Proof: Existence is demonstrated geometrically: $v_n(q)$ is a continuous and increasing function with $v_n(0) = -\infty$ and $\lim_{q \rightarrow \infty} v_n(q) = c$; $v_a(q)$ is a continuous and non-negative function with $\lim_{q \rightarrow \infty} v_a(q) = 0$; therefore, there exists at least one q such that $v_n(q) = v_a(q)$.

From the Poincaré-Hopf index theorem (Vives 1999), there is at most one equilibrium if $z'(q) < 0$ for all equilibrium q , where $z(q) = v_a(q) - v_n(q)$. Define the auxiliary functions,

$y_a(q) = v_a(q)F(q)$ and $y_n(q) = v_n(q)F(q)$. Differentiate,

$$z'(q) = \frac{(y'_a(q) - y'_n(q))F(q) - f(q)(y_a(q) - y_n(q))}{F(q)^2}.$$

At an equilibrium $y_a(q) = y_n(q)$, so $z'(q) < 0$ at an equilibrium if

$$y'_a(q) - y'_n(q) < 0. \quad (10)$$

We have $y'_n(q) = p_1 f(q)$,

$$y'_a(q) = - \int_0^{\hat{\xi}(q)} \frac{\hat{s}_2(\xi) \hat{p}_2(\xi)}{I(q, \xi)^2} dF(\xi) + \int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} \frac{\partial p_2(I(q, \xi), \xi)}{\partial I} dF(\xi) + p_1 f(q).$$

Therefore, the condition (10) can be written as

$$- \int_0^{\hat{\xi}(q)} \frac{\hat{s}_2(\xi) \hat{p}_2(\xi)}{I(q, \xi)^2} dF(\xi) + \int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} \frac{\partial p_2(I(q, \xi), \xi)}{\partial I} dF(\xi) < 0,$$

which holds because $\partial p_2(I, \xi)/\partial I < 0$. (Note, (10) does not imply that $z'(q) < 0$ for all q .)

Now demonstrate $q_a^* > q^o$. Differentiate the profit function,

$$\begin{aligned} \Pi'(q) &= \left(\int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} \left(p_2(I(q, \xi), \xi) + I(q, \xi) \frac{\partial p_2(I(q, \xi), \xi)}{\partial I} \right) dF(\xi) + \int_{\tilde{\xi}(q)}^q p_1 dF(\xi) \right) - y_n(q) \\ &= y_a(q) - y_n(q) + \int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} I(q, \xi) \frac{\partial p_2(I(q, \xi), \xi)}{\partial I} dF(\xi) - \int_0^{\hat{\xi}(q)} \frac{\hat{s}_2(\xi) \hat{p}_2(\xi)}{I(q, \xi)} dF(\xi) \end{aligned}$$

The sum of the latter two terms is negative if $\tilde{\xi}(q) > 0$. From (8) it must be that $\tilde{\xi}(q^o) > 0$.

Therefore, $\Pi'(q^o) = 0$ implies $y_a(q^o) - y_n(q^o) > 0$, which implies $z(q^o) > 0$. Since there is a unique q_a^* such that $z(q_a^*) = 0$ and $z'(q_a^*) < 0$, it follows that $q^o < q_a^*$. \square

5.2 Marginal salvage value

Given that the newsvendor model is based on a marginal analysis, one might argue that the marginal salvage value is more appropriate than the average salvage value. To be specific, let $v_m(q, \xi)$ be the revenue received from the last unit ordered, i.e., the q^{th} unit, assuming the revenue from that unit is collected in the clearance period. Let $v_m(q)$ be the expected marginal salvage value:

$$v_m(q) = E[v_m(q, \xi)] = \frac{1}{F(q)} \left(\int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} p_2(I(q, \xi), \xi) dF(\xi) + \int_{\tilde{\xi}(q)}^q p_1 dF(\xi) \right) \quad (11)$$

A comparison of (9) with (11) reveals that $v_m(q) = v_a(q)$ when $\hat{\xi}(q) = 0$, i.e., the marginal salvage value is identical to the average salvage value when it is always optimal to sell

all clearance period inventory, as when there is constant elasticity demand, $D_2(p_2, \xi) = x(\xi)\alpha p_2^{-\beta}$. Hence, in that situation the marginal salvage value is no better than the average salvage value. However with exponential demand, $D_2(p_2, \xi) = x(\xi)\alpha e^{-\beta p_2}$, $\hat{\xi}(q) > 0$ is possible, in which case $v_m(q) < v_a(q)$. While Theorem 2 indicates that the marginal salvage yields better results in those cases, it nevertheless still does not yield the optimal profit.

Theorem 2 *With the marginal salvage value there exists a unique equilibrium, $\{v_m^*, q_m^*\}$ and $q_m^* > q^o$. If $\hat{R}_2(s_2, \xi)$ is increasing in s_2 (so that it is always optimal to liquidate all clearance period inventory) then $q_m^* = q_a^*$, otherwise $q_m^* < q_a^*$.*

Proof: This proof is analogous to Theorem 1, so it is omitted for brevity. \square

5.3 Weighted average salvage value

Neither the average salvage value nor the marginal salvage value are weighted to account for the number of units that are salvaged, but the weighted average salvage value does. To be specific, let $v_w(q)$ be the expected weighted average salvage value, i.e., expected period 2 revenue divided by expected period 2 inventory conditional that there is inventory to liquidate in period 2:

$$v_w(q) = \frac{\frac{1}{F(q)}R_2(q)}{\frac{1}{F(q)}I(q)} = \frac{R_2(q)}{I(q)}$$

where

$$R_2(q) = \int_0^{\hat{\xi}(q)} \hat{s}_2(\xi) \hat{p}_2(\xi) dF(\xi) + \int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} p_2(I(q, \xi), \xi) I(q, \xi) dF(\xi) + \int_{\tilde{\xi}(q)}^q p_1 I(q, \xi) dF(\xi).$$

Because $v_w(q)$ is the ratio of two expectations while $v_a(q)$ is the expectation of the ratio, the methods can yield significantly different results: in the example in §1 the average salvage value is \$130, but the weighted average salvage value is $\$815,812/7,096 = \115 . The weighted average salvage value is closer to the ideal target of \$111, but still does not equal the target. Given that the average salvage value is too high, we can expect the weighted average salvage value to perform better: the lowest observed salvage values tends to occur when inventory is highest. Indeed, in our example the weighted salvage value equilibrium is $\{\$126, 1162\}$, which is still not optimal, but generates a 7% profit loss instead of the 22% profit loss with the average salvage value.

Analogous to the existence proof for the average salvage value, it can be shown that there exists a weighted average salvage value equilibrium, $\{v_w^*, q_w^*\}$. Among the scenarios

considered in the numerical study discussed in the next section we did not find a scenario in which there existed multiple equilibria with the weighted average salvage value. Nevertheless we are unable to prove (or provide simple conditions for) uniqueness. In all of the scenarios we report in the numerical study we find $q_w^* > q^o$, but it is possible to construct pathological examples in which $q_w^* < q^o$.

5.4 Marginal revenue

The three methods discussed so far are simple and intuitive, but they do not generate the optimal solution as a unique equilibrium. The marginal revenue method does. Let the salvage value be the marginal revenue, $v_r(q)$, where

$$\begin{aligned} v_r(q) &= \frac{1}{F(q)} \int_0^q \frac{\partial R_2(I(q, \xi), \xi)}{\partial I} dF(\xi) \\ &= \frac{1}{F(q)} \left(\int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} \left(p_2(I(q, \xi), \xi) + I(q, \xi) \frac{\partial p_2(I(q, \xi), \xi)}{\partial q} \right) dF(\xi) + \int_{\tilde{\xi}(q)}^q p_1 dF(\xi) \right) \end{aligned} \quad (12)$$

The next theorem demonstrates that the optimal solution is indeed an equilibrium with $v_r(q)$, but, more importantly, the optimal solution is the unique equilibrium.

Theorem 3 *With the marginal revenue, $v_r(q)$, the unique equilibrium is $\{v_r(q^o), q^o\}$, i.e., $v_r(q^o) = v_n(q^o)$.*

Proof: Existence and uniqueness proofs are analogous to Theorem 1 and omitted for brevity. Define $y_r(q) = v_r(q)F(q)$ and $y_n(q) = v_n(q)F(q)$. We have

$$y_r(q) - y_n(q) = \int_0^q \frac{\partial R_2(I(q, \xi), \xi)}{\partial I} dF(\xi) - p_1 F(q) + (p_1 - c) \quad (13)$$

The solution to $y_r(q) - y_n(q) = 0$ is the equilibrium quantity with the marginal revenue method, denoted q_r^* . $y_r(q) - y_n(q)$ is identical to $\Pi'(q)$ given in (6). Hence $q_r^* = q^o$. \square

Interestingly, the marginal revenue is not really a salvage value, i.e., it is not in general the “per unit amount that can be earned on left over inventory”. Note that $v_r(q)$ simplifies to

$$v_r(q) = v_m(q) + \frac{1}{F(q)} \left(\int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} I(q, \xi) \frac{\partial p_2(I(q, \xi), \xi)}{\partial q} dF(\xi) \right) \quad (14)$$

Marginal revenue and marginal salvage value concepts coincide only when clearance period revenue is linear in the amount of left over inventory. But if clearance period revenue is concave in the amount of left over inventory, then the marginal revenue is less than the

marginal salvage value. In fact, it is possible that at q^o the firm discovers that the marginal salvage value is greater than the purchase cost $v_m(q^o) > c$ (which renders the newsvendor model infeasible), but the marginal revenue is always less than cost at the optimal solution, $v_r(q^o) < c$.

While evaluating $v_a(q)$, $v_m(q)$ and $v_w(q)$ from a sample of data requires no explanation, the evaluation of $v_r(q)$ is not straightforward. A conservative estimate is

$$\hat{v}_r(q) = \left(\frac{1}{n-1} \right) \sum_{j=1}^{n-1} \frac{t_{j+1} - t_j}{y_{j+1} - y_j}$$

assuming the data are sorted such that $y_i \leq y_j$ for all $i \leq j$. The above is conservative in the sense that it is biased such that $E[\hat{v}_r(q)] < v_r(q)$, but this is prudent given that the newsvendor model is more sensitive to an overestimation of the salvage value than an underestimation. Using $\hat{v}_r(q)$ with our example we find that the marginal revenue with an order quantity of 1400 is \$81, which clearly suggests the order quantity of \$1400 is too high. (Recall, 1400 is optimal only if the salvage value is \$135.)

6 Numerical Study

This section reports on a numerical study to assess the magnitude of the performance loss from using either $v_a(q)$, $v_m(q)$ or $v_w(q)$ as the salvage value input to the newsvendor model.

We generated 336 scenarios from all combinations of the following parameters:

$$\begin{aligned} m &= (p - c)/p = \{0.25, 0.5\} & d_2(p_2) &= \{\alpha e^{-\beta p_2}, \alpha p_2^{-\beta}\} \\ \sigma/\mu &= \{0.25, 0.5, 1.0\} & \beta &= \{1.2, 2.4\} \\ x(\xi) &= \{\xi, \mu\} & \frac{p-c}{p-v_a^*} &= \{0.55, 0.60, \dots, 0.85\} \end{aligned}$$

where m is the gross margin, σ is the standard deviation of regular season demand, and μ is the mean of regular season demand. We assume that $D_2(p_2, \xi) = x(\xi)d_2(p_2)$, so $x(\xi) = \xi$ means regular season and clearance period demands are positively correlated, whereas $x(\xi) = \mu$ means they are independent. The second period demand function is either exponential, $\alpha e^{-\beta p_2}$, or constant price elasticity, $\alpha p_2^{-\beta}$. Tellis (1988) finds that the β parameter generally ranges between 1 and 3 with an average of 2, so we choose $\{1.2, 2.4\}$ to represent high and low demand elasticity.

In each scenario we set $p_1 = 2$, $\mu = E[D_1] = 1000$ and the regular season demand follows a gamma distribution. In each scenario the α parameter in the second period demand function is chosen such that the average salvage value heuristic equilibrium, $\{v_a^*, q_a^*\}$, yields

the desired critical ratio. Hence, these scenarios could plausibly be observed if a firm were to use the average salvage value.

Table 3 presents summary data on the profit performance of the three non-optimal methods for estimating the salvage value. The average salvage value performs the worst, followed by the marginal salvage value, and the weighted average salvage value performs well on average, but its maximum profit loss can be substantial (35%). Table 4 indicates that all three methods order more than the optimal quantity, often by a considerable amount. Figure 2 reveals that the performance of all three methods deteriorates as the critical ratio is increased or as the gross margin decreases. For example, with a gross margin of 25% and a critical ratio of 75% (which are similar to the parameters faced by Sport Obermeyer, as reported by Fisher and Raman, 1996) the average salvage value’s profit is 22% lower on average than the optimal profit. Table 5 reveals that the weighted salvage value performs quite well whenever the clearance period demand is positively correlated with regular season demand, $x(\xi) = \xi$, but can perform poorly when clearance period demand is independent of regular season demand, $x(\xi) = \mu$, and there is significant demand uncertainty. Thus, while the weighted average salvage value is not optimal, it can be remarkably effective, but, unfortunately, not always. We conclude that the method by which the salvage value is estimated does have a significant impact on expected profits.

While our main focus is on the quality of the various performance measures, we were also curious about the frequency at which the estimated salvage value is greater than cost at the optimal order quantity. Recall, in our example in §1 at the optimal order quantity (941 units) the average salvage value is \$157 even though the cost per unit is \$150. Table 6 reveals that this precarious situation for the newsvendor model is actually quite common with a low gross margin (25%) and high critical ratios (70% or higher).

7 Discussion

Our main finding is that the classic newsvendor model should not be implemented in practice without careful consideration of the estimation of one of its inputs, the salvage value. We demonstrate that several intuitive estimation procedures lead to significantly poor performance, especially with products that have low gross margins (e.g., 25%), a high critical ratio (e.g., 75%) and high demand uncertainty (e.g., a coefficient of variation above 0.5).

Furthermore, there can be little evidence to indicate that poor performance is occurring, i.e., data are consistent with actions and actions are consistent with data. We also provide an estimation procedure that does lead to the optimal action in equilibrium, thereby establishing the effectiveness of the newsvendor model in situations with clearance pricing.

This paper is best viewed in terms of work on the robustness of heuristics in other classic inventory models (e.g., Dobson 1988, Gallego 1998, Zheng 1992): the newsvendor model is a simplified version of the clearance pricing model and the question is whether this simplification deteriorates performance. However, in those other papers the issue of input-action dependence does not exist.

One might argue that our results are not necessary if a manager is willing to either directly use the clearance pricing model or is willing to use the newsvendor model with a non-linear salvage value function. We feel, at least in some situations, that there are compelling reasons to favor our approach (i.e., implement the newsvendor model with the salvage value estimated from historical data). To implement the clearance pricing model a manager must have an estimate of the second period demand function, i.e., we cannot just assume the manager knows that function with certainty. (This reminds us of the well-known joke about an engineer, a physicist and an economist stranded on a desert island with only one can of beans for nourishment and the economist announces “let’s assume we have a can opener”.) Where will the manager obtain that function? Is it multiplicative or additive? Does it take a constant elasticity form, an exponential form, or some other form? Surely there is error in the estimation of that function and that error depends on the chosen order quantities. So in practice, the implementation of the clearance pricing model does not avoid the challenge of estimating inputs. The same critique is relevant to the application of the newsvendor model with a non-linear salvage value function: where does that function come from? In addition, neither of those approaches provides the simplicity and elegance of the newsvendor model with a fixed salvage value: it is unlikely the clearance pricing model could be taught effectively in MBA programs, and in complex settings the newsvendor model is a more desirable “building block” (i.e., sub-model) than the clearance pricing model due to its analytical tractability. Finally, we demonstrate that the newsvendor model can be effectively implemented. Therefore its simplicity need not come with any loss of performance and its parsimony can prove to be valuable when it is difficult to specify the parameters or

structure of the dynamics of the clearance period.

Another concern with this research could be whether the concept of an equilibrium is even necessary to evaluate the performance of an estimation method. We believe it is indeed necessary. Any estimation method that works with historical salvage values to estimate future salvage values faces the issue that those historical data depend on previous actions. Hence, there is no way around the dependence between salvage values and actions.

Although we concentrate on the newsvendor model, the issue of input-action dependence is likely to be relevant in many other models in operations management: e.g., an airline's overbooking decision and the observed costs of having customers relinquish their seats are interdependent; demands for a booking class could depend on which fare classes are open (Talluri and van Ryzin 2001); and Lariviere and Porteus (1999) demonstrate that a dependence exists between the chosen inventory policy and the estimate of the demand rate because demand is truncated by available inventory. Although in some cases the input-action dependence is explicitly considered (as is common with optimal control problems), our interest is in the performance of simple heuristics that do not explicitly account for this dependence.

To conclude, we emphasize that a model is not helpful to practitioners if it exists in a vacuum: while in some settings we are forced to make assumptions about the inputs to our models, practicing managers must actually use data to estimate inputs. In our opinion, the concept of input-action dependence and equilibrium is not a mere intellectual curiosity but rather an important construct for understanding the performance of a model. We have demonstrated this in the important newsvendor model.

References

- Agrawal, V., S. Seshadri. 2000a. Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manu. and Service Oper. Mgmt.* **2** 410-413.
- Agrawal, V., S. Seshadri. 2000b. Distribution free bounds for service constrained (Q, r) inventory systems. *Naval Research Logistics.* **47** 635-656.
- Armony, M., E. Plambeck. 2002. The impact of duplicate orders on demand estimation and capacity investment. New York University working paper.

- Axsater, S. 1996. Using the deterministic EOQ formula in stochastic inventory control. *Management Science*. **42**. 830-834.
- Bertsekas, D.P. 2000. *Dynamic Programming and Optimal Control*. Belmont, MA. Athena Scientific.
- Bitran, G. B., S. V. Mondschein. 1997. Periodic pricing of seasonal products in retailing. *Management Sci.* **43** 64-79.
- Bitran, G. B., R. Caldentey, S. V. Mondschein. 1997. Coordinating clearance markdown sales of seasonal products in retail chains. *Oper. Res.* **46** 609-624.
- Brumelle, S.L., J.I. McGill, T.H. Oum, K.Sawaki, M.W. 1990. Tretheway. Allocation of airline seats between stochastically dependent demands. *Trans. Sci.* **24**(3): 183-192.
- Cachon, G. 2003. Supply chain coordination with contracts. *The Handbook of Operations Research and Management Science: Supply Chain Management*. eds. S. Graves, T. de Kok, Kluwer.
- Cachon, G., C. Terwiesch, Yi Xu. 2002. The retail assortment decision in the prescence of consumer search. University of Pennsylvania working paper.
- Carr, S., W. Lovejoy. 2000. The inverse newsvendor problem: choosing an optimal demand portfolio for capacitated resources. *Management Sci.* **46** 912-927.
- Dana Jr., J., N. Petruzzi. 2001. Note: The newsvendor model with endogenous demand. *Management Sci.* **47** 1488-1497.
- Dobson, G. 1988. Sensitivity of the EOQ model to parameter estimates. *Oper. Research.* **36**(4): 570-74.
- Donohue, K. 2000. Efficient supply contracts for fashion goods with forecast updating and two production modes. *Management Sci.* **46**(11): 1397-1411.
- Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. Risk-averse (and prudent) newsboy. *Management Sci.* **41** 786-794.
- Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47** 454-475.
- Feng, Y., G. Gallego. 1995. Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science.* **41**(8): 1371-1391.

- Fisher, M., K. Rajaram, R. Raman. 2001. Optimizing inventory replenishment of retail fashion products. *Manufacturing and Service Operations Management*. **3**(3): 230-241.
- Fisher, M., A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* **44** 87-99.
- Gallego, G. 1998. New bounds and heuristics for (Q, r) policies. *Mgmt. Sci.* **44**(2): 219-233.
- Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*. **40**(8): 999-1020.
- Hammond, J.A. Raman. 1994. Sport Obermeyer Ltd. Havard Business School Case 9-695-022.
- Hertz, D. B., K. H. Schaffir. 1960. A forecasting method for management of seasonal style-goods inventories. *Operations Research*. **8**(1). 45-52.
- Hoch, S., B. Kim, A. Montgomery, P. Rossi. 1995. Determinants of store-level price elasticity. *Journal of Marketing*. **13**(1). 7-29.
- Kouvelis, P., G. Gutierrez. 1994. The newsvendor problem in a global market: optimal centralized and decentralized control policies for a two-market stochastic inventory system. *Management Sci.* **43**(5): 571-585.
- Lariviere, M., E. Porteus. 1999. Stalking information: Bayesian inventory management with unobserved lost sales. *Management Science*. **45**(3). 346-363.
- Lee, H. 1996. Effective inventory and service Management through product and process redesign. *Operations Research*. **44**(1). 151-159.
- Lee, H., S. Whang. 2002. The impact of the secondary market on the supply chain. *Management Science*. **48**. 719-731.
- Li, L. 1992. The Role of Inventory in Delivery-Time Competition. *Management Science*, **38** 182-197.
- Lippman, S., K. McCardle. 1995. The competitive newsboy. *Oper. Research*. **45**. 54-65.
- Lovejoy, W. S. 1990. Myopic policies for some inventory models with uncertain demand distributions. *Management Sci.* **36** 724-738.
- Monohan, G., N. Petruzzi, W. Zhao. 2002. The dynamic pricing problem from a newsvendor's perspective. University of Illinois at Urbana-Champaign working paper.

- Mulhern, F., R. Leone. 1991. Implicit price bundling of retail products: a multiproduct approach to maximizing store profitability. *Journal of Marketing*. **55**. 63-76.
- Petruzzi, N., M. Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Oper. Res.* 47 183-193.
- Petruzzi, N., M. Dada. 2001. Information and inventory recourse for a two-market, price setting retailer. *Manufacturing and Service Oper. Management*. **3** 242-263.
- Porteus, E. 1990. Stochastic inventory theory. *Handbooks in Operations Research and Management Science*. eds., D. Heyman and M. Sobel. Amsterdam, North-Holland. Vol 2.
- Schweitzer, M., G. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution: experimental evidence. *Management Sci.* **46** 404-420.
- Shang, K., J. S. Song. 2003. Newsvendor bounds and heuristic for optimal policies in serial supply chains. *Management Science*. **49**(5). 618-638.
- Smith, S., D. Achabal. 1998. Clearance Pricing and Inventory Policies for Retail Chains. *Management Sci.* **44** 285-300.
- Sundaram, 1999. A First Course in Optimization Theory. Cambridge University Press. Cambridge, UK.
- Tellis, G. 1988. The price elasticity of selective demand: a meta analysis of econometric models of sales. *Journal of Marketing Research*. **25**: 331-341.
- Talluri, K., G. van Ryzin. 2001. Revenue management under a general discrete choice model of consumer behavior. Columbia University working paper.
- Van Mieghem, J. 1999. Coordinating investment, production and subcontracting. *Management Science*. **45**(7). 954-971.
- Van Mieghem, J., N. Rudi. 2002. Newsvendor networks: dynamic inventory management and capacity investments with discretionary activities. *MSOM*. **4** 313-335.
- van Ryzin, G.S. Mahajan. 1999. On the relationship between inventory costs and variety benefits in retail assortments. *Management Science*. 45(11). 1496-1509.
- Vives, X. (1999). *Oligopolistic competition: old ideas and new tools*. Cambridge, MA, MIT Press.
- Zheng, Y. S. 1992. On properties of stochastic inventory systems. *Mgmt Sci.* **38** 87-103.

Table 1: Potential profit losses from errors in the salvage value estimate. Demand follows a gamma distribution with mean 1000, the gross margin and correct salvage value are 1/4th and 2/3rds of the retail price respectively.*

		% loss in profit (relative to the optimal profit) from ...			
Coefficient of variation	δ	Underestimating the salvage value	Overestimating the salvage value	Under ordering	Over ordering
		by $\delta\%$	value by $\delta\%$	by $\delta\%$	by $\delta\%$
1.41	12%	14.2%	302.9%	0.8%	0.9%
0.71	12%	4.4%	59.9%	0.9%	1.0%
0.35	12%	1.7%	18.8%	1.2%	1.5%
1.41	8%	8.0%	38.2%	0.4%	0.4%
0.71	8%	2.4%	8.8%	0.4%	0.4%
0.35	8%	0.9%	3.0%	0.5%	0.6%
1.41	4%	2.6%	5.4%	0.1%	0.1%
0.71	4%	0.7%	1.4%	0.1%	0.1%
0.35	4%	0.3%	0.5%	0.1%	0.2%

* Let q^o be the optimal quantity given the correct salvage value, v , and let q' be the order quantity chosen if the inputed salvage value is $(1+\delta)v$ or $(1-\delta)v$. The profit loss is the difference between expected profit (evaluated with the newsvendor model) with q^o and q' , reported as a percentage of the optimal profit.

Table 2: Ten observations from a newsvendor model. (Demand has a gamma distribution with mean 1000 and standard deviation 707, price is 200, procurement cost is 150 and the order quantity is 1400 units.)

Inventory at the start of the clearance period,	Total clearance period revenue, t_i	Salvage revenue per unit, $u_i = t_i / y_i$
y_i		
106	21,104	200
283	49,989	177
431	63,916	148
561	74,505	133
679	83,234	123
788	90,825	115
893	97,722	109
999	104,279	104
1110	110,919	100
1247	118,688	95
Total	7,096	815,182

Table 3: Profit loss, $(1 - \Pi(q^*) / \Pi(q^o))$

Salvage value	Standard		Median	Minimum	Maximum
	Average	Deviation			
Average	12.5%	14.2%	6.8%	0.0%	63.3%
Marginal	10.2%	13.1%	5.1%	0.0%	63.3%
Weighted average	2.3%	5.0%	0.5%	0.0%	35.0%

Table 4 : Over order %, $(q^* / q^o - 1)$

Salvage value	Standard		Median	Minimum	Maximum
	Average	Deviation			
Average	27.8%	20.6%	20.8%	1.4%	86.6%
Marginal	24.0%	19.7%	17.6%	0.5%	86.6%
Weighted average	9.6%	10.7%	5.5%	0.3%	49.3%

Table 5 : Profit loss, $1 - \Pi(q^*) / \Pi(q^o)$

			Weighted average salvage value		
σ / μ	$x(\xi)$	$d_2(p_2)$	Standard		
			Average	deviation	Maximum
0.25	ξ	$e^{-\beta p}$	0.29%	0.31%	1.35%
		$p^{-\beta}$	0.29%	0.35%	1.45%
	μ	$e^{-\beta p}$	0.67%	0.71%	3.02%
		$p^{-\beta}$	0.71%	0.85%	3.51%
0.50	ξ	$e^{-\beta p}$	0.45%	0.49%	2.09%
		$p^{-\beta}$	0.45%	0.53%	2.18%
	μ	$e^{-\beta p}$	2.27%	2.25%	9.27%
		$p^{-\beta}$	2.44%	2.78%	11.11%
1.00	ξ	$e^{-\beta p}$	0.44%	0.49%	1.94%
		$p^{-\beta}$	0.42%	0.50%	1.98%
	μ	$e^{-\beta p}$	9.35%	7.55%	28.81%
		$p^{-\beta}$	10.29%	9.72%	35.01%

Table 6: Percentage of scenarios in which the average salvage value at the optimal order quantity is greater than cost.

Critical ratio	Gross Margin	
	25%	50%
0.55	0%	0%
0.60	8%	0%
0.65	13%	0%
0.70	29%	0%
0.75	54%	0%
0.80	79%	13%
0.85	100%	50%
Average	40%	9%




Figure 1. Average Salvage Value Heuristic Equilibrium




Figure 2. Average performance of the newsvendor heuristics; m =margin, AV = Average Salvage Value Heuristic, MV = Marginal Salvage Value Heuristic, WV=Weighted Average Salvage Value Heuristic