MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 9, No. 3, Summer 2007, pp. 276–290
ISSN 1523-4614 | EISSN 1526-5498 | 07 | 0903 | 0276
DOI 10.1287/msom.1060.0145
© 2007 INFORMS

Implementation of the Newsvendor Model with Clearance Pricing: How to (and How Not to) Estimate a Salvage Value

Gérard P. Cachon
Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104, cachon@wharton.upenn.edu

A. Gürhan Kök
Fuqua School of Business, Duke University, Durham, North Carolina 27708, gurhan.kok@duke.edu

The newsvendor model is designed to decide how much of a product to order when the product is to be sold over a short selling season with stochastic demand and there are no additional opportunities to replenish inventory. There are many practical situations that reasonably conform to those assumptions, but the traditional newsvendor model also assumes a fixed salvage value: all inventory left over at the end of the season is sold off at a fixed per-unit price. The fixed salvage value assumption is questionable when a clearance price is rationally chosen in response to the events observed during the selling season: a deep discount should be taken if there is plenty of inventory remaining at the end of the season, whereas a shallow discount is appropriate for a product with higher than expected demand. This paper solves for the optimal order quantity in the newsvendor model, assuming rational clearance pricing. We then study the performance of the traditional newsvendor model. The key to effective implementation of the traditional newsvendor model is choosing an appropriate fixed salvage value. (We show that an optimal order quantity cannot be generally achieved by merely enhancing the traditional newsvendor model to include a nonlinear salvage value function.) We demonstrate that several intuitive methods for estimating the salvage value can lead to an excessively large order quantity and a substantial profit loss. Even though the traditional model can result in poor performance, the model seems as if it is working correctly: the order quantity chosen is optimal given the salvage value inputted to the model, and the observed salvage value given the chosen order quantity equals the inputted one. We discuss how to estimate a salvage value that leads the traditional newsvendor model to the optimal or near-optimal order quantity. Our results highlight the importance of understanding how a model can interact with its own inputs: when inputs to a model are influenced by the decisions of the model, care is needed to appreciate how that interaction influences the decisions recommended by the model and how the model's inputs should be estimated.

Key words: markdown management; game theory; retailing; inventory; revenue management
History: Received: October 18, 2005; accepted: November 20, 2006.

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, 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 is described in detail in most operations management textbooks.

The newsvendor model is not complicated: an order quantity is the only decision; 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 leftover inventory is salvaged at the end of the season for a fixed salvage value per unit, v. However, in many practical applications of the model the per-unit salvage value is not fixed, but rather depends on a clearance pricing decision: a small discount is needed with a popular product in short supply, whereas a deep discount is needed with an unpopular product in ample supply. In other words, with clearance pricing the salvage value depends on both the amount of inventory ordered and the correlation of demand across time, neither of which are directly accounted for in the newsvendor model.

The focus of this paper is the performance of the newsvendor model in situations with clearance pricing. The key to the successful implementation of the newsvendor model is choosing the correct salvage value to input to the model, i.e., the salvage value that leads the newsvendor model to choose the order quantity that maximizes expected profit, given that the clearance price is chosen optimally to maximize clearance-period revenue. We consider several intuitive methods for estimating the salvage value and evaluate them based on their equilibrium performance. To explain, consider the following quote that describes the economics of selling fashion ski apparel, as faced by Sport Obermeyer (Hammond and Raman 1994, p. 8): "units left over at the end of the season were sold at a loss that averaged 8% of the . . . price." That 8% figure leads to a salvage value, but it depends on how much Sport Obermeyer ordered. Had Sport Obermeyer ordered twice as much as it did, then its historical losses would have been higher (because the company would have had more parkas to discount, causing deeper discounts). Hence, we evaluate the performance of each method when the expected salvage value, given the chosen order quantity, equals the inputted salvage value that leads the newsvendor model to recommend the chosen order quantity. In other words, in equilibrium the input to the newsvendor model is consistent with the observed outcome.

We find that, even in equilibrium, several intuitive salvage-value estimation methods overestimate the correct salvage value, leading to significant profit losses. In fact, when the optimal order quantity is chosen, we find in many situations that those methods estimate the salvage value to be greater than the marginal cost of purchasing the product; the newsvendor model foolishly recommends ordering an unlimited quantity. We provide a simple salvage-value estimation method that yields near-optimal performance when demand is highly correlated across the season (which is likely for products, like apparel, that have variable clearance pricing). We also discuss a more complex method that leads to the optimal salvage value.

Section 1 defines our model, §2 reviews the related literature, §3 identifies the optimal procurement quantity, §4 defines and analyzes several salvage-value estimation rules, §5 presents some numerical results, and §6 discusses our results.

1. The Newsvendor and the Clearance-Pricing Models

In this section we define the two models we study. The first is the traditional newsvendor model, which requires a fixed salvage value as one of its inputs. The second is the clearance-pricing model, which 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, 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 1. Let $F(\cdot)$ be the strictly increasing 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. The clearance-pricing model is nearly identical to the newsvendor model except that at the start of the clearance period, a clearance period price, p_2, is chosen to maximize revenue in the clearance period given the available inventory, $I(q, \xi)$. We assume $p_2 \leq p_1$, which is reasonable in some situations and de facto imposed by many firms. We assume that the inventory left over at the end of the clearance period has zero constant salvage value. Instances with nonzero salvage value w can be transformed to a problem with zero salvage value by subtracting w from p_1 and c, and using $p_2 + w$ in the second-period demand function.

Period 2 demand, $D_2(p_2, \xi)$, is a deterministic function of the clearance price and the realization of

Period 1 demand. $D_2(p_2, \xi)$ is nonnegative, differentiable, and decreasing in p_2. Hence, the inverse demand function exists, $p_2(s_2, \xi)$. It is reasonable to assume D_2 depends on ξ if total sales are highly correlated with early season sales, for which there is evidence (Fisher and Raman 1996, Fisher et al. 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 that the profit function is well behaved in q. The second assumption implies that with unlimited inventory the optimal clearance price is lower than p_1, i.e., discounting in the second-period is rational if there is plenty of inventory left over.

We assume $D_2(p_2, \xi)$ is 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 continuous and increasing in ξ. These conditions imply that total demand across the two periods increases in ξ.

A form of $D_2(p_2, \xi)$ that meets our 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, $D_2(p_2, \xi) = x(\xi)\alpha p_2^{-\beta}$ for $\beta > 1$ 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 $\hat{p}_2(\xi) < p_1$.) There is substantial 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 et al. (1995), and Tellis (1988).

2. 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 multiperiod 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 nonlinear 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, or whether it applies in situations with correlated demand. 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 firms 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 et al.(1995), Lippman and McCardle (1995), Schweitzer and Cachon (2000), Van Mieghem and Rudi (2002).

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 et al. (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. In each case except Petruzzi and Dada (2001), 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 traditional newsvendor model.

There are numerous papers that study revenue management, or markdown pricing, or both: e.g., Bitran and Mondschein (1997), Bitran et al. (1997), Brumelle et al. (1990), Federgruen and Heching (1999), Feng and Gallego (1995), Gallego and Van Ryzin (1994), Monahan et al. (2004), and Smith and Achabal (1998). With the exception of Brumelle et al. (1990), these papers assume that 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 economic order quantity (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 four other papers that discuss the consequence of ignoring the interdependence between inputs and actions, albeit in very different settings from ours: Armony and Plambeck (2005) consider demand forecasting in a supply chain in which customers may submit duplicate orders, and Cachon et al. (2005) study assortment planning with consumer search. Balakrishnan et al. (2004) find that iteratively applying the standard EOQ model in a setting where demand depends on the inventory level converges to an equilibrium, which is suboptimal. Cooper et al. (2006) study a similar issue in the context of revenue management. In setting a protection level for the high-fare demand class, the demands for high and low fares must be estimated from data, and those data are influenced by the protection level chosen. They find that the iterative application of the Littlewood rule (which is analogous to the newsvendor problem), combined with an estimation procedure, converges to a point in which the protection level is too low even though it is optimal according to the Littlewood rule and the demand forecasts, given the protection level. They study the behavior of a sequence of controls and estimates, while we study the equilibrium behavior and establish the existence and uniqueness of the equilibrium.

In econometrics, a common challenge is the estimation of a parameter that depends on the decisions made by a firm when the firm is known to possess some knowledge that is relevant to its decision that is not available to the econometrician. There are numerous techniques for handling this endogeneity issue. (Examples include Berry et al. 1995 and Nevo 2001. See Chintagunta et al. 2006 for a review.) In our model, the decision maker has all of the information needed to make a decision, so our challenge is different from the standard endogeneity challenge in econometrics. 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.

3. Optimal Procurement and Clearance Pricing

In this section, we evaluate the optimal decision in the traditional newsvendor model and then the clearance-pricing model. We conclude with a brief explanation for why the clearance-pricing model is more complex (i.e., not equivalent) than the newsvendor model with a nonlinear salvage-value function.

Expected profit in the newsvendor model is

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

where $r_i(q)$ is expected revenue in period i. Period 1 expected revenue is $r_1(q) = p_1(q - I(q))$ and 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 leftover inventory. Leftover inventory for a given demand realization is $I(q, \xi) = (q - \xi)^+$, and expected leftover inventory is given by

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

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:

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

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: First, we establish that clearance-period revenue is concave in the order quantity. Next, we show there exist three threshold functions that partition the regular season demand space into four intervals. The clearance-period revenue function depends on the interval in which the regular season demand realization falls. Finally, we demonstrate that 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 $s_2 \le I(q, \xi)$. When $\xi > q$, there is no inventory left in the clearance period and there is no clearance-period pricing decision. Let $\tilde{R}_2(q, \xi)$, defined over $[0, \infty) \times [0, q]$, be the firm's maximum revenue constrained by available inventory:

$$\tilde{R}_2(q, \xi) = \max_{s_2}(\hat{R}_2(s_2, \xi)\colon s_2 \le I(q, \xi)).$$

From the Maximum Theorem Under Convexity, $\tilde{R}_2(q, \xi)$ is concave in q for fixed ξ because $\hat{R}_2(s_2, \xi)$ is concave in s_2.

The remaining constraint to consider is $p_2 \le p_1$. Let $R_2(q, \xi)$ be the firm's maximum Period 2 revenue, given both the $p_2 \le p_1$ and the $s_2 \le I(q, \xi)$ constraints (and recall that $\hat{p}_2(\xi) < p_1$ for all ξ):

$$R_2(q, \xi) = \min\{p_1 I(q, \xi), \tilde{R}_2(I(q, \xi), \xi)\} \quad \text{for } \xi \le q.$$

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

We now turn to the evaluation of $R_2(q, \xi)$. 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, $I(q,\xi) \geq \hat{s}_2(\xi)$. 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

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

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 satisfies $I(q,\xi) = D_2(p_1,\xi)$, which can be written as

$$q - \xi - D_2(p_1,\xi) = 0. \tag{4}$$

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 $q < \xi$ and $R_2(q,\xi) = 0$. Given that we have established $q \geq \tilde{\xi}(q) \geq \hat{\xi}(q)$ for any q, the second-period revenue is

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

Let $R_i(q)$ be the expected revenue in period i. We have

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

Note that we use upper- and lower-case notation to represent analogous functions in the two models. While $r_1(q) = R_1(q)$, i.e., newsvendor and clearance-pricing 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.

The firm's expected profit, $\Pi(q)$, equals the first-period profit plus the revenue from the second period:

$$\Pi(q) = -cq + R_1(q) + R_2(q).$$

The next step is to identify the optimal quantity.

Theorem 1. *The optimal procurement quantity q^o is the unique solution to the following*:

$$0 = (p_1 - c) - p_1F(q) + \left(\int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} \frac{\partial}{\partial q}(I(q,\xi)p_2(I(q,\xi),\xi))\,dF(\xi) + \int_{\tilde{\xi}(q)}^{q} p_1\,dF(\xi)\right). \tag{6}$$

Proof. Differentiating $\Pi(q)$ yields the right-hand side of (6). Notice that $R_2(q,\xi)$ is a continuous function. Evaluating the second derivatives of $\Pi(q)$ requires the evaluation of $\partial R_2(q,\xi)/\partial q$ at the limits of the integrals (i.e., the break points in its definition). Because it is not necessarily differentiable at those points, $\partial R_2(q,\xi)/\partial q$ at the lower limit of an integral is the right derivative, and that at the upper limit is the left derivative. The second derivative is

$$\begin{aligned} \frac{\partial^2 \Pi(q)}{\partial q^2} = &-p_1 f(q) + \int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} \frac{\partial^2 R_2(q,\xi)}{\partial q^2}\,dF(\xi) \\ &- \hat{\xi}'(q)\frac{\partial R_2(q,\xi)}{\partial q}\bigg|_{\xi=\hat{\xi}(q)} + \tilde{\xi}'(q)\frac{\partial R_2(q,\xi)}{\partial q}\bigg|_{\xi=\tilde{\xi}(q)} \\ &- \tilde{\xi}'(q)p_1 f(\tilde{\xi}(q)) + p_1 f(q). \end{aligned} \tag{7}$$

Terms 1 and 6 cancel each other out. Term 2 is negative, because $R_2(q,\xi)$ is concave in q for all $\xi \in [\hat{\xi}(q), \tilde{\xi}(q)]$. Term 4 is also negative, because $\partial R_2(q,\xi)/\partial q \geq 0$ for any ξ. The sum of Terms 4 and 5 is negative, because $\partial R_2(q,\xi)/\partial q \leq p_1$ for all q and ξ. Hence, $\Pi(q)$ is strictly concave in q. □

It is natural to wonder whether the optimal quantity for the clearance-pricing model can be obtained by using the newsvendor model with a nonlinear salvage-value function. (Porteus 1990 provides the optimal quantity with the latter.) For example, suppose the second-period revenue is

$$r_2(q) = \int_0^q I(q,\xi) v(I(q,\xi))\, dF(\xi),$$

where the salvage value, $v(\cdot)$, is a decreasing function of remaining inventory. Marginal second-period revenue with respect to the order quantity is then

$$\int_0^q \frac{\partial}{\partial q}(I(q,\xi)v(I(q,\xi)))\, dF(\xi). \tag{8}$$

A nonlinear salvage-value function could be constructed so that (8) matches the right-hand side of (6) only if $p_2(I(q,\xi),\xi) = v(I(q,\xi))$ for all ξ, and that is possible only if the second-period price is independent of ξ, i.e., only if there is no correlation in demand. The nonlinear salvage-value model assumes that the salvage revenue depends only on the number of units left over and is independent of how the firm arrived at that number of units. That is a restrictive assumption: For example, the revenue from having 99 units left over must be independent of whether the firm started with 100 units (i.e., first-period demand was one unit) or 10,100 units (first-period demand was 10,001 units). Thus, the newsvendor model with a nonlinear salvage-value function is equivalent to the clearance-pricing model only when demand is independent across time.

4. Salvage-Value Estimation

This section defines and analyzes four methods for estimating the salvage value, v, to be used in the newsvendor model. The first three salvage-value heuristics are intuitively reasonable, but they do not lead the newsvendor model to the optimal quantity. We are interested in the bias they produce (whether they tend to over- or underorder relative to the optimal quantity). In the next section we explore numerically the profit loss associated with each of these methods. The fourth method leads to the correct salvage value (i.e., the salvage value that causes the newsvendor model to output the optimal order quantity), but we illustrate why it is difficult to estimate this salvage value using typically available data. We start with the following definitions to formalize the interaction between the estimation rules and the application of the newsvendor model.

A *salvage-value estimator* yields an estimate of the salvage value from a set of historical data. The *expected salvage value* $v(q)$ is the expected estimate of the salvage value for a given salvage-value estimator and order quantity. A *salvage-value heuristic* is the use of the newsvendor model with a salvage-value estimator. A *heuristic equilibrium* is a pair, $\{v^*, q^*\}$, such that q^* is the traditional newsvendor model's optimal order quantity given an inputted salvage value v^*, and v^* is expected salvage value given an order quantity q^*: $v^* = v(q^*) = v_n(q^*)$. (Given that $v_n(q)$ is increasing, it is sufficient to define an equilibrium merely by the salvage value, v^*, or the order quantity, q^*, but we choose to define an equilibrium as a pair to emphasize the connection between the model's input, v^*, and its recommended action, q^*.)

4.1. Average Salvage-Value Heuristic

The average salvage value is literally the average revenue received per unit of inventory at the start of the clearance period. For a single observation, the average salvage value is

$$v_a(q,\xi) = R_2(q,\xi)/I(q,\xi).$$

For a sample with n observations from similar items for which q units were ordered and units were salvaged, the average salvage-value estimator is

$$\hat{v}_a(q) = \frac{1}{n}\sum_{j=1}^{n}\frac{t_j}{y_j}, \tag{9}$$

where, for the jth observation, y_j is the amount of inventory at the start of the clearance period, and t_j is the total salvage revenue.

We can use the clearance-pricing model to evaluate the expected average salvage value, $v_a(q)$, conditional on having leftover inventory:

$$\begin{aligned} v_a(q) &= E[v_a(q,\xi)] \\ &= \frac{1}{F(q)}\left(\int_0^{\hat{\xi}(q)} \frac{\hat{s}_2(\xi)\hat{p}_2(\xi)}{I(q,\xi)}\, dF(\xi) \right. \\ &\qquad \left. + \int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} p_2(I(q,\xi),\xi)\, dF(\xi) + \int_{\tilde{\xi}(q)}^{q} p_1\, dF(\xi)\right). \end{aligned} \tag{10}$$

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

Figure 1 Average Salvage-Value Heuristic Equilibrium

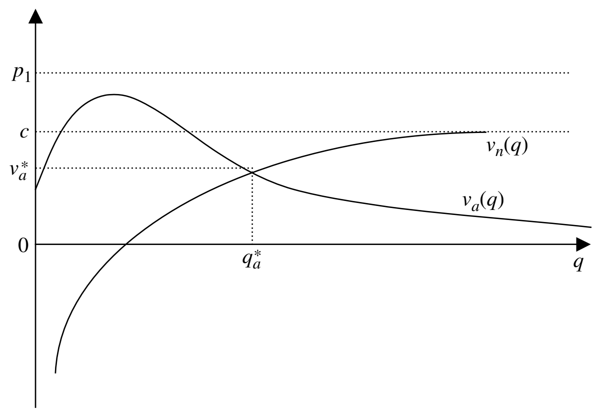

at the clearing price, and the third integral includes outcomes in which all inventory is sold at below the clearing price.

Let $\{v_a^*, q_a^*\}$ be a heuristic equilibrium when the firm uses the average salvage-value heuristic: $\{v_a^*, q_a^*\}$ is an equilibrium if $v_a^* = v_a(q_a^*) = v_n(q_a^*)$. We wish to determine whether an equilibrium exists and, if so, if it is unique and, if so, how it performs, i.e., what the relationship is between q_a^* and q^o and between $\Pi(q_a^*)$ and $\Pi(q^o)$.

To provide some 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. One may think $v_a(q)$ should be decreasing (i.e., the expected salvage value is decreasing in the number of units ordered), 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, but nevertheless inventory must be salvaged, 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 that it holds for all q.)

Theorem 2. *With the average salvage-value heuristic, there exists a unique heuristic 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 $\lim_{q\to\infty} v_n(q) = c$; $v_a(q)$ is a continuous and nonnegative function with $\lim_{q\to\infty} v_a(q) = 0$; 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

$$y_a'(q) - y_n'(q) < 0. \tag{11}$$

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 (11) 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, (11) 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 (6) 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 $z(q^o) > 0$. Because there is a unique q_a^* such that $z(q_a^*) = 0$ and $z'(q_a^*) < 0$, it follows that $q^o < q_a^*$. □

Consider an iterative process that chooses an order quantity based on the current average salvage value, observes a salvage value, updates the average salvage value using the estimator, and so on. Although Theorem 2 guarantees the existence of a unique equilibrium, it does not ensure the convergence of the iterative process to the equilibrium. However, in our numerical experiments, the iterative process always converges to the equilibrium regardless of the initial salvage value.

4.2. Marginal Salvage-Value Heuristic

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, marginal salvage value $v_m(q, \xi)$ is the revenue received from the last unit ordered, i.e., the qth unit, assuming the revenue from that unit is collected in the clearance period.

$$v_m(q,\xi) = \begin{Bmatrix} p_2(I(q,\xi),\xi), & \text{if } s_2(p_2) \geq I(q,\xi) \\ 0, & \text{otherwise.} \end{Bmatrix}$$

From a sample of n observations as described in §4.1, the marginal salvage-value estimator yields

$$\hat{v}_m(q) = \frac{1}{n}\sum_{j=1}^{n}\frac{t_j}{y_j}1\{z_j = y_j\}, \tag{12}$$

where z_j is the number of units sold in the clearance period and $1\{z_j = y_j\}$ is an indicator function equal to one if $z_j = y_j$, and zero otherwise.

From the clearance-pricing model, the expected marginal salvage value, $v_m(q)$, is

$$\begin{aligned} v_m(q) &= E[v_m(q,\xi)] \\ &= \frac{1}{F(q)}\left(\int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} p_2(I(q,\xi),\xi)\,dF(\xi) + \int_{\tilde{\xi}(q)}^{q} p_1\,dF(\xi)\right). \end{aligned} \tag{13}$$

A comparison of (10) with (13) reveals that $v_m(q) = v_a(q)$ when $\hat{\xi}(q) = 0$, i.e., it is expected that the marginal salvage value is identical to the expected average salvage value when it is always optimal to sell all clearance-period inventory, as in the case of constant elasticity demand. Hence, in that situation the marginal salvage value is no different from the average salvage value. However, with exponential demand, $\hat{\xi}(q) > 0$ is possible (i.e., it may be optimal to salvage only a portion of the clearance-period inventory), in which case $v_m(q) < v_a(q)$. While Theorem 3 indicates that the marginal salvage heuristic yields better results in those cases, it nevertheless still does not yield the optimal profit.

Theorem 3. *With the marginal salvage-value heuristic, there exists a unique heuristic equilibrium, $\{v_m^*, q_m^*\}$ and $q^o < q_m^*$. If $\hat{R}_2(s_2, \xi)$ is increasing in s_2 for all ξ (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 2, so it is omitted for brevity. □

4.3. Weighted Average Salvage-Value Heuristic

Neither the average salvage value nor the marginal salvage value is weighted to account for the number of units that are salvaged, but the weighted average salvage value is. To be specific, given a sample of observations as described in §4.1, the weighted average salvage estimator is

$$\hat{v}_w(q) = \sum_{j=1}^{n} t_j \Big/ \sum_{j=1}^{n} y_j. \tag{14}$$

The expected weighted average salvage value, $v_w(q)$, expected clearance-period revenue divided by expected clearance-period inventory conditional that there is inventory to liquidate in the clearance period:

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

Because $v_w(q)$ is the ratio of two expectations, while $v_a(q)$ is the expectation of the ratio, the average salvage-value heuristic and the weighted average salvage-value heuristic can yield significantly different results. Given that the average salvage value is too high, we expect the weighted average salvage value to perform better: The lowest observed salvage values tend to occur when inventory is highest.

Analogous to the existence proof for the average salvage-value heuristic equilibrium, it can be shown that there exists a weighted average salvage-value heuristic 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 $q_w^* < q^o$.

4.4. Marginal Revenue Heuristic

The three salvage-value heuristics discussed so far are simple and intuitive, but they do not generate the optimal solution as a unique equilibrium. The marginal revenue heuristic does. Let the input to the traditional newsvendor model be the expected marginal revenue, $v_r(q)$, where

$$\begin{aligned} v_r(q) &= \frac{1}{F(q)}\int_0^q \frac{\partial R_2(q,\xi)}{\partial q}\,dF(\xi) \\ &= \frac{1}{F(q)}\left(\int_{\hat{\xi}(q)}^{\tilde{\xi}(q)}\left(p_2(I(q,\xi),\xi) + I(q,\xi)\frac{\partial p_2(I(q,\xi),\xi)}{\partial q}\right)dF(\xi) + \int_{\tilde{\xi}(q)}^{q} p_1\,dF(\xi)\right). \end{aligned} \quad (15)$$

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

Theorem 4. *With the marginal revenue heuristic, the unique heuristic 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 2 and omitted for brevity. Define $y_r(q) = v_r(q)F(q)$ and $y_n(q) = v_n(q)F(q)$. We have

$$y_r(q) - y_n(q) = \int_0^q \frac{\partial R_2(q,\xi)}{\partial q}\,dF(\xi) - p_1F(q) + (p_1 - c).$$

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, $q_r^* = q^o$. □

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

$$v_r(q) = v_m(q) + \frac{1}{F(q)}\left(\int_{\hat{\xi}(q)}^{\tilde{\xi}(q)} I(q,\xi)\frac{\partial p_2(I(q,\xi),\xi)}{\partial q}\,dF(\xi)\right). \quad (16)$$

The marginal revenue and marginal salvage-value concepts coincide only when the second-period price does not depend on the amount of leftover inventory. However, if clearance-period revenue is concave in the amount of leftover inventory, then marginal revenue is less than the marginal salvage value (which helps to explain why the marginal salvage-value heuristic equilibrium quantity is greater than q^o). In fact, it is possible that if the firm orders the true optimal quantity, q^o, then the marginal salvage value can be greater than the purchase cost $v_m(q^o) > c$ (which renders the newsvendor model infeasible). However, at the optimal solution the marginal revenue is always less than cost, $v_r(q^o) < c$.

The evaluation of $v_r(q)$ from the sample data described in §4.1 is not straightforward. If there is no correlation in demand across time, then a conservative estimate of the marginal revenue is

$$\hat{v}_r(q) = \frac{1}{n}\sum_{j=1}^{n-1}\frac{t_{j+1} - t_j}{y_{j+1} - y_j}. \quad (17)$$

This salvage-value estimator is a piecewise-linear approximation of the concave-increasing salvage revenue curve. It is conservative in the sense that it is biased such that $E[\hat{v}_r(q)] < v_r(q)$, which is prudent, given that the newsvendor model is more sensitive to an overestimation of the salvage value than to an underestimation. The following theorem shows that the heuristic is asymptotically unbiased:

Theorem 5. $E[\hat{v}_r(q)] < v_r(q)$ *for finite n and* $E[\hat{v}_r(q)] \to v_r(q)$ *as* $n \to \infty$.

Proof. Given $(y_j, t_j)_{j=1}^n$, let $\delta_j = |\partial R_2(q,\xi)/\partial q|_{\xi = q - y_j}$. Note that δ_j are not observed, but if they were, the unbiased estimator for $v_r(q)$ would be $(1/n)\sum_{j=1}^n \delta_j$. Define $u_1 = t_1/y_1$ and $u_j = (t_j - t_{j-1})/(y_j - y_{j-1})$ for $j > 1$. Define $l_j = (t_{j+1} - t_j)/(y_{j+1} - y_j)$ for $j < n$ and $l_n = 0$. Clearly, we have $l_j < \delta_j < u_j$ for all j and $l_j = u_{j+1}$ for $j = 1, \ldots, n-1$. Therefore,

$$\frac{1}{n}\sum_{j=1}^{n-1} l_j < \frac{1}{n}\sum_{j=1}^{n}\delta_j < \frac{1}{n}\sum_{j=1}^{n} u_j$$

$$\frac{1}{n}\sum_{j=1}^{n-1} l_j < \frac{1}{n}\sum_{j=1}^{n}\delta_j < \frac{1}{n}\left(\sum_{j=1}^{n-1} l_j + u_1\right)$$

$$\hat{v}_r(q) < \frac{1}{n}\sum_{j=1}^{n}\delta_j < \hat{v}_r(q) + p_1/n.$$

For any sample path, $\hat{v}_r(q) < (1/n)\sum_{j=1}^{n}\delta_j$. Therefore, $E[\hat{v}_r(q)] < E[(1/n)\sum_{j=1}^{n}\delta_j] = v_r(q)$. As $n \to \infty$, $\hat{v}_r(q) \to (1/n)\sum_{j=1}^{n}\delta_j$ and $(1/n)\sum_{j=1}^{n}\delta_j \to v_r(q)$, hence $\hat{v}_r(q) \to v_r(q)$. □

It is relatively simple to estimate marginal revenue with independent demands, because then the clearance-period revenue depends only on the number of units salvaged. With correlated demand, (17) may yield unreasonable estimates: in particular, it may result in a negative estimate for marginal revenue even though $v_r(q)$ is clearly never negative. (Total revenue can decrease in the number of units to be salvaged because a large number of leftover units implies a low demand outcome, which then requires a drastic price cut, thereby cutting total revenue even though there are more units to salvage.)

It appears that estimating $v_r(q)$ in the correlated demand case is quite difficult, with data typically available to a firm because $v_r(q)$ is the incremental amount of revenue if one additional unit is available to salvage and if one additional unit was originally ordered (so that the difference between them, which is the regular-season sales, is held constant). In other words, the firm needs to have an estimate for the change in salvage revenue that would occur while holding regular season demand constant. For every observation of regular-season demand, however, there is generally only one observation of salvage revenue.

A solution for the correlated demand case is to make an assumption regarding the second-period demand function. If second-period demand takes the exponential form, $D_2(p_2, \xi) = x(\xi)\alpha e^{-\beta p_2}$, then (16) simplifies to

$$v_r(q) = v_m(q) - \frac{F(\tilde{\xi}(q)) - F(\hat{\xi}(q))}{\beta F(q)}, \tag{18}$$

and if second-period demand takes the isoelastic form, $D_2(p_2, \xi) = x(\xi)\alpha p_2^{-\beta}$, then (16) simplifies to

$$v_r(q) = v_m(q) - \frac{1}{\beta F(q)}\int_0^{\tilde{\xi}(q)} p_2(I(q, \xi), \xi)\, dF(\xi). \tag{19}$$

Both (18) and (19) are easy to estimate assuming that the elasticity parameter is known (or has been estimated, which is not difficult given the assumed demand form): (18) is equivalent to the evaluation of the marginal salvage value with the one exception that $p_2(I(q, \xi), \xi) - 1/\beta$ is taken as the salvage value instead of just $p_2(I(q, \xi), \xi)$ when there is a markdown that does not clear all inventory, i.e., $\hat{\xi}(q) \le \xi \le \tilde{\xi}(q)$; and (19) is equivalent to the evaluation of the marginal salvage value with the one exception that $p_2(I(q, \xi), \xi)/\beta$ is taken as the salvage value instead of just $p_2(I(q, \xi), \xi)$ when a markdown is made. The estimator $\hat{v}_r(q)$ can be obtained in this case by using $\hat{v}_m(q)$ as defined in Equation (12) and estimates of the additional terms in (18) and (19). Notice that this approach does not require the estimation of the correlation structure $x(\xi)$, whereas finding the optimal solution to the clearance-pricing model does. Of course, if the clearance-period demand model is fully known, then it is not actually necessary to implement the traditional newsvendor model by evaluating $v_r(q)$. Instead, the clearance-period optimal solution, (6), could be evaluated directly. In this case, the two models lead to the same order quantity.

To summarize the average, marginal, and weighted average salvage-value heuristics lead to order quantities greater than the optimal solution, and the marginal revenue heuristic leads to the optimal solution. These results hold when clearance-period demand is independent or correlated with initial period demand. We provide an estimator for the marginal revenue heuristic that works when clearance-period demand is independent of initial period demand, and we suggest an alternative method when demands are correlated across the periods. The alternative method for estimating marginal revenue requires the modeler to assume a form of the clearance-period demand function (e.g., isoelastic or exponential) and an estimate of the elasticity of the demand function (but does not require the modeler to know the correlation between the two periods).

5. 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:

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

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 one and three, with an average of two, so we choose $\{1.2, 2.4\}$ to represent low- and high-demand elasticity.

In each scenario, we set $p_1 = 2$, $\mu = E[D_1] = 1{,}000$; 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 1 presents summary data on the profit performance of the three nonoptimal salvage value heuristics. 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 2 indicates that all three methods order more than the optimal quantity, often by a considerable amount. Table 2 and Figure 2 reveal that the performance of all three methods deteriorates as the variability or critical ratio increases and 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 3 reveals that the weighted salvage value performs quite well when-

Table 1 **Profit Loss,** $(1 - \Pi(q^*)/\Pi(q^o))$

Salvage-value heuristic	σ/μ	Average (%)	Standard deviation (%)	Median (%)	Minimum (%)	Maximum (%)
Average	0.25	4.0	3.4	3.0	0.0	14.3
	0.50	9.3	7.8	7.4	0.1	31.6
	1.00	24.1	17.7	21.2	0.3	63.3
	All	12.5	14.2	6.8	0.0	63.3
Marginal	0.25	3.2	3.2	2.2	0.0	14.3
	0.50	7.6	7.4	5.6	0.0	31.6
	1.00	19.8	17.4	16.1	0.0	63.3
	All	10.2	13.1	5.1	0.0	63.3
Weighted average	0.25	0.5	0.6	0.3	0.0	3.5
	0.50	1.4	2.0	0.6	0.0	11.1
	1.00	5.1	7.7	0.8	0.0	35.0
	All	2.3	5.0	0.5	0.0	35.0

Table 2 **Over Order %,** $(q^*/q^o - 1)$

Salvage-value heuristic	σ/μ	Average (%)	Standard deviation (%)	Median (%)	Minimum (%)	Maximum (%)
Average	0.25	11.4	5.1	11.9	1.4	21.4
	0.50	23.2	10.2	24.7	2.7	43.2
	1.00	48.7	20.5	53.3	6.0	86.6
	All	27.8	20.6	20.8	1.4	86.6
Marginal	0.25	9.9	5.4	10.0	0.5	21.4
	0.50	20.0	10.9	20.6	1.0	43.2
	1.00	42.0	22.0	44.1	2.1	86.6
	All	24.0	19.7	17.6	0.5	86.6
Weighted average	0.25	3.7	2.3	3.4	0.3	10.2
	0.50	7.8	5.4	6.7	0.4	22.9
	1.00	17.1	14.7	10.4	0.5	49.3
	All	9.6	10.7	5.5	0.3	49.3

Figure 2 **Average Performance of the Newsvendor Heuristics**

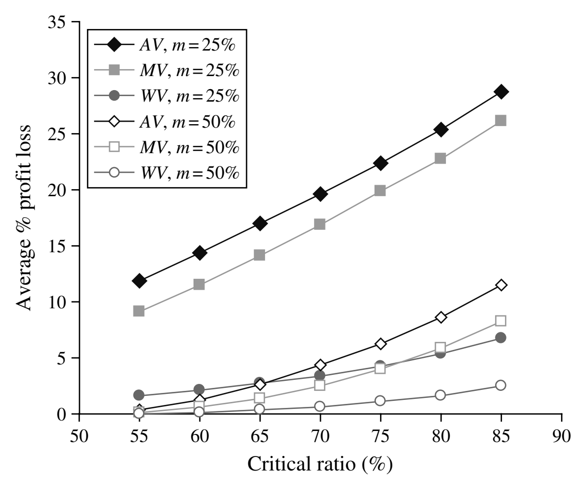

Notes. m = margin, AV = average salvage-value heuristic, MV = marginal salvage-value heuristic, WV = weighted average salvage-value heuristic.

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

			Weighted average salvage value			
σ/μ	$x(\xi)$	$d_2(p_2)$	Average (%)	Standard deviation (%)	Median (%)	Maximum (%)
0.25	ξ	$e^{-\beta p}$	0.29	0.31	0.20	1.35
		$p^{-\beta}$	0.29	0.35	0.16	1.45
	μ	$e^{-\beta p}$	0.67	0.71	0.43	3.02
		$p^{-\beta}$	0.71	0.85	0.41	3.51
0.50	ξ	$e^{-\beta p}$	0.45	0.49	0.31	2.09
		$p^{-\beta}$	0.45	0.53	0.24	2.18
	μ	$e^{-\beta p}$	2.27	2.25	1.64	9.27
		$p^{-\beta}$	2.44	2.78	1.52	11.11
1.00	ξ	$e^{-\beta p}$	0.44	0.49	0.24	1.94
		$p^{-\beta}$	0.42	0.50	0.23	1.98
	μ	$e^{-\beta p}$	9.35	7.55	7.35	28.81
		$p^{-\beta}$	10.29	9.72	7.74	35.01

ever 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. In particular, it is effective in the most relevant condition for a fashion item: demand is positively correlated across seasons. If it is believed that demand is independent across the periods, then the marginal revenue estimator (17) could be applied. We conclude that the method by which the salvage value is estimated does have a significant impact on expected profits.

Because we have established analytically that the average salvage value is too high, we are curious to determine how frequently the average salvage value is greater than purchase cost when the optimal quantity is ordered. In those situations, the traditional newsvendor model recommends ordering an unlimited quantity, which suggests that the chosen order quantity is too low. Table 4 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).

Table 4 **Percentage of Scenarios in Which the Average Salvage Value at the Optimal Order Quantity Is Greater Than Cost**

	Gross margin (%)	
Critical ratio	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

6. Discussion

Our main finding is that the traditional newsvendor model should not be implemented in practice without careful consideration of the estimation of one of its inputs, the salvage value. If the clearance price is independent of the amount of leftover inventory, then the marginal salvage value method is appropriate. For example, if a bakery has a policy of selling day-old products for 50% off no matter the amount of product left over, then the marginal salvage value (the expected revenue on the last unit in the clearance period) leads to the optimal solution. However, if the clearance price depends on the amount of inventory remaining at the end of the regular season (because the firm is following a responsive clearance-pricing strategy), then the marginal salvage value (or worse, the average salvage value) can lead to an order quantity that is significantly greater than optimal, and to a substantial profit loss. Furthermore, there may be no indication that a grossly suboptimal decision is made: the chosen order quantity may be optimal for the inputted salvage value, which is then observed in expectation given the order quantity. Interestingly, given an optimal order quantity, we find that the average salvage value of leftover inventory may actually be larger than marginal cost, in which case a naive manager may conclude, based on the traditional newsvendor model, that the order quantity is too low.

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 either is willing to use the newsvendor model with a nonlinear salvage-value function or is willing to use the more-complex clearance-pricing model. We have shown that the newsvendor model with a nonlinear salvage value function cannot replicate the clearance-pricing model when demand is correlated across seasons. Therefore, the nonlinear salvage-value model is not a viable approach.

To implement the clearance-pricing model, it is necessary to estimate the clearance period demand model: estimate the correlation structure, choose the form of the demand function (exponential, isoelastic, or something else), and then determine their parameters. Conditional that this estimation is done correctly, the clearance-pricing model yields an optimal order quantity. Implementation of the traditional newsvendor model requires fewer assumptions regarding the structure of the clearance-period demand, but does not lead to an optimal solution. Nevertheless, in the important case of correlated demand, the weighted average salvage value appears to generate results that are near optimal. Alternatively, the newsvendor model used with an adjusted marginal salvage value leads to the optimal solution if demand elasticity is known. Thus, we are somewhat agnostic with respect to which model a manager should use in practice. Instead, we focus on the potential pitfalls of using the traditional newsvendor model with an inappropriate method for estimating its salvage-value input.

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 heuristic equilibrium is not a mere intellectual curiosity, but rather is a key construct for understanding the performance of a model. We have demonstrated this in the important newsvendor model.

Acknowledgments

The authors would like to thank Awi Federgruen, Marshall Fisher, Ananth Iyer, Anton Kleywegt, Martin Lariviere, Costis Maglaras, 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. The previous version of this paper was titled "Heuristic Equilibrium and the Estimation of the Salvage Value in the Newsvendor Model with Clearance Pricing."

References

Agrawal, V., S. Seshadri. 2000a. Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing Service Oper. Management* **2** 410–413.

Agrawal, V., S. Seshadri. 2000b. Distribution free bounds for service constrained (Q, r) inventory systems. *Naval Res. Logist.* **47** 635–656.

Armony, M., E. Plambeck. 2005. The impact of duplicate orders on demand estimation and capacity investment. *Management Sci.* **51**(10) 1505–1518.

Axsater, S. 1996. Using the deterministic EOQ formula in stochastic inventory control. *Management Sci.* **42** 830–834.

Balakrishnan, A., M. S. Pangburn, E. Stavrulaki. 2004. Stack them high, let 'em fly: Lot-sizing policies when inventories stimulate demand. *Management Sci.* **50**(5) 630–644.

Berry, S., J. Levinsohn, A. Pakes. 1995. Automobile prices in market equilibrium. *Econometrica* **63**(4) 841–890.

Bertsekas, D.P. 2000. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA.

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. Tretheway. 1990. Allocation of airline seats between stochastically dependent demands. *Transportation Sci.* **24**(3) 183–192.

Cachon, G. 2003. Supply chain coordination with contracts. S. Graves, T. de Kok, eds. *The Handbook of Operations Research and Management Science: Supply Chain Management*. Kluwer, Amsterdam, The Netherlands, 229–339.

Cachon, G., C. Terwiesch, Y. Xu. 2005. Retail assortment planning in the presence of consumer search. *Manufacturing Service Oper. Management* **7**(4) 330–346.

Carr, S., W. Lovejoy. 2000. The inverse newsvendor problem: Choosing an optimal demand portfolio for capacitated resources. *Management Sci.* **46** 912–927.

Chintagunta, P., T. Erdem, P. Rossi, M. Wedel. 2006. Structural models in marketing: Review and assessment. *Marketing Sci.* **25**(6) 604–616.

Cooper, W. L., T. Homem-de-Mello, A. J. Kleywegt. 2006. Models of the spiral-down effect in revenue management. *Oper. Res.* **54**(5) 968–987.

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. Res.* **36**(4) 570–574.

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 Sci.* **41**(8) 1371–1391.

Fisher, M., A. Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* **44** 87–99.

Fisher, M., K. Rajaram, R. Raman. 2001. Optimizing inventory replenishment of retail fashion products. *Manufacturing Service Oper. Management* **3**(3) 230–241.

Gallego, G. 1998. New bounds and heuristics for (Q, r) policies. *Management Sci*. **44**(2) 219–233.

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

Hammond, J., A. Raman. 1994. Sport Obermeyer Ltd. Harvard Business School Case 9-695–022, Harvard University, Cambridge, MA.

Hertz, D. B., K. H. Schaffir. 1960. A forecasting method for management of seasonal stylegoods inventories. *Oper. Res.* **8**(1) 45–52.

Hoch, S., B. Kim, A. Montgomery, P. Rossi. 1995. Determinants of store-level price elasticity. *J. Marketing* **13**(1) 7–29.

Kouvelis, P., G. Gutierrez. 1997. 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.

Lee, H. 1996. Effective inventory and service management through product and process redesign. *Oper. Res.* **44**(1) 151–159.

Lee, H., S. Whang. 2002. The impact of the secondary market on the supply chain. *Management Sci.* **48** 719–731.

Li, L. 1992. The role of inventory in delivery-time competition. *Management Sci.* **38** 182–197.

Lippman, S., K. McCardle. 1995. The competitive newsboy. *Oper. Res.* **45** 54–65.

Lovejoy, W. S. 1990. Myopic policies for some inventory models with uncertain demand distributions. *Management Sci.* **36** 724–738.

Monahan, G., N. Petruzzi, W. Zhao. 2004. The dynamic pricing problem from a newsvendor's perspective. *Manufacturing Service Oper. Management* **6**(1) 73–91.

Mulhern, F., R. Leone. 1991. Implicit price bundling of retail products: A multiproduct approach to maximizing store profitability. *J. Marketing* **55** 63–76.

Nevo, A. 2001. Measuring market power in the ready-to-eat cereal industry. *Econometrica* **69**(2) 307–342.

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 Service Oper. Management* **3** 242–263.

Porteus, E. 1990. Stochastic inventory theory. D. Heyman, M. Sobel, eds. *Handbooks in Operations Research and Management Science*, Vol. 2. North-Holland, Amsterdam, The Netherlands, 605–652.

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 Sci.* **49**(5) 618–638.

Smith, S., D. Achabal. 1998. Clearance pricing and inventory policies for retail chains. *Management Sci.* **44** 285–300.

Tellis, G. 1988. The price elasticity of selective demand: A meta analysis of econometric models of sales. *J. Marketing Res.* **25** 331–341.

Van Mieghem, J. 1999. Coordinating investment, production and subcontracting. *Management Sci.* **45**(7) 954–971.

Van Mieghem, J., N. Rudi. 2002. Newsvendor networks: Dynamic inventory management and capacity investments with discretionary activities. *Manufacturing Service Oper. Management* **4** 313–335.

van Ryzin, G., S. Mahajan. 1999. On the relationship between inventory costs and variety benefits in retail assortments. *Management Sci.* **45**(11) 1496–1509.

Vives, X. 1999. *Oligopolistic Competition: Old Ideas and New Tools*. MIT Press, Cambridge, MA.

Zheng, Y. S. 1992. On properties of stochastic inventory systems. *Management Sci.* **38** 87–103.