

Structural Estimation of the Newsvendor Model: An Application to Reserving Operating Room Time

Marcelo Olivares

Columbia Business School, New York, New York 10027, molivares@columbia.edu

Christian Terwiesch

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104,
terwiesch@wharton.upenn.edu

Lydia Cassorla

Department of Anesthesia and Perioperative Care, University of California, San Francisco, California 94143,
cassorla@anesthesia.ucsf.edu

The newsvendor model captures the trade-off faced by a decision maker that needs to place a firm bet prior to the occurrence of a random event. Previous research in operations management has mostly focused on deriving the decision that minimizes the expected mismatch costs. In contrast, we present two methods that estimate the unobservable cost parameters characterizing the mismatch cost function. We present a structural estimation framework that accounts for heterogeneity in the uncertainty faced by the newsvendor as well as in the cost parameters. We develop statistical methods that give consistent estimates of the model primitives, and derive their asymptotic distribution, which is useful to do hypothesis testing. We apply our econometric model to a hospital that balances the costs of reserving too much versus too little operating room capacity to cardiac surgery cases. Our results reveal that the hospital places more emphasis on the tangible costs of having idle capacity than on the costs of schedule overrun and long working hours for the staff. We also extend our structural models to incorporate external information on forecasting biases and mismatch costs reported by the medical literature. Our analysis suggests that overconfidence and incentive conflicts are important drivers of the frequency of schedule overruns observed in our sample.

Key words: newsvendor model; empirical research; health care; structural estimation; operating room reservation

History: Accepted by Linda V. Green, public sector applications; received December 27, 2004. This paper was with the authors 1 year for 4 revisions.

1. Introduction

Many business decisions require that a decision maker takes a firm decision before the occurrence of a random event. Once the uncertainty is resolved, the decision maker observes that her decision was too “large” or too “small,” and incurs costs reflecting the mismatch between her decision and the ex post optimal decision. The newsvendor model captures this trade-off. In operations management, the most frequently analyzed application of this type of decision deals with placing an inventory order in the presence of demand uncertainty.

While previous research related to the newsvendor model has taken the mismatch cost parameters as given and has minimized the expected cost to obtain a cost-minimizing decision, we take a different approach. Following the tradition of structural estimation models in econometrics, we assume that the

decision maker acts rationally and chooses the optimal decision for a cost function that is unobservable for us as researchers. Based on the observed decision making and a set of covariates, we use maximum likelihood estimation to obtain the cost parameters describing the latent cost function. For example, we can use this method to estimate how much value a retailer assigns to a stockout.

In the econometrics community, similar approaches have been taken by Rust (1987) and Berry et al. (1995). Rust (1987) combines a Markov decision process describing a maintenance problem with the empirically observed behavior of the person in charge of managing the maintenance to impute costs of regular cost maintenance and perceived costs of unexpected failures. Berry et al. (1995) use data from the automotive industry to estimate model markups based on a oligopoly model of price competition in a differenti-

ated product market. We contribute to this stream of literature by providing estimable structural models of newsvendor-type decisions, which can be used in a broad variety of managerial applications.

In addition to deriving the estimation procedures for newsvendor-like cost models as well as establishing their econometric properties, we apply our theory to a health-care application. In the setting we studied, the hospital had to reserve a certain amount of operating room (OR) time to specific cardiac procedures. Because the actual procedure time in the OR is random and will—in the best of all cases—vary around the expected procedure time, some procedures will exceed the forecasted durations while others will be completed ahead of schedule. If the hospital reserves too much time to a case, the OR is likely to incur excessive idle time. If, however, the hospital reserves too little time to a case, the hospital is likely to face schedule overruns and decreased service quality. (See Strum et al. 1999 for an application of the newsvendor model in OR management.)

Our econometric models and its health-care application enable us to make the following contributions. First, we extend the long line of newsvendor research by developing a theory that allows for an estimate of the underlying cost function. We present two model specifications, both of which are sufficiently general to capture cost and demand heterogeneity and hence have the potential to be applied to various operations management decisions. Second, for each model, we derive a two-step estimation procedure and establish its key econometric properties, including the asymptotic distribution of the estimators and the associated standard errors required for hypothesis testing. Third, we apply our econometric framework to a health-care setting. We analyze how a hospital balances the costs of reserving too much versus too little OR capacity to individual cardiac surgery cases and demonstrate that our model has significant predictive power for this decision.

2. The Newsvendor Model and Structural Estimation

The newsvendor is a simple and intuitive model and is arguably one of the stepping stones for decision making in operations management. The model is defined as follows. Given a random variable D with distribution $F(\cdot)$, a decision maker (hereon the newsvendor) needs to make a decision Q , before the realization of the random variable D is known. The objective of the newsvendor is to minimize the expected mismatch cost between D and Q . This mismatch cost is assumed to be linear in the amount of the mismatch but typically is not symmetric. If the newsvendor's decision Q exceeds D , the incurred

cost is equal to $C_o(Q - D)^+$, where $(x)^+ = \max\{x, 0\}$. If D exceeds Q , the incurred cost is equal to $C_u(D - Q)^+$. The model parameters C_o and C_u are referred to as the overage and underage cost, respectively, and are assumed to be strictly positive. The optimal decision Q^* solves

$$\min_Q E\{C_o(Q - D)^+ + C_u(D - Q)^+\}. \quad (1)$$

We assume that the problem is unconstrained and therefore Q can take any value in the real domain. Because the objective function defined in (1) is convex in Q , the optimal solution can be characterized by the first-order condition. If the random variable D is continuous, the optimal solution to (1) satisfies


$$F(Q^*) = \frac{C_u}{C_u + C_o} = \frac{1}{1 + \gamma}, \quad (2)$$

where $\gamma = C_o/C_u$ is the ratio between overage and underage costs (see, e.g., Porteus 2002).

In this paper, we extend the long tradition of newsvendor research by developing an econometric framework to impute the cost parameters of a newsvendor based on observed decisions. In our decision problem, a traditional operations research model takes the distribution function F and the cost parameters C_o and C_u as model input, and then characterizes the optimal reservation decision Q^* (see Figure 1, left). Unlike the operations researcher, who is interested in providing a normative theory of how rational agents “should” behave, the econometrician is interested in a descriptive theory of how real-world decision makers actually do behave. There are two econometric approaches toward developing such a theory. In a method known as *reduced form estimation*, researchers collect data on a dependent variable of interest and use regression analysis (or other statistical methods) to explain its variation through a set of explanatory variables. In our decision problem, such models might take the amount of underage, $(D - Q)^+$, as a dependent variable and attempt to explain it through a set of explanatory variables. The outcome of this estimation would be a set of parameters characterizing the marginal impact of an explanatory variable on the amount of overage (see Figure 1, middle). This approach has been the dominant one in empirical research in operations management (e.g., Lieberman and Demeester 1999, Brush and Karnani 1996), in empirical work in health-care operations management (e.g., Milne et al. 1989, Chorba 1976), and is also widely used in the medical community (e.g., Pell et al. 2001, Urbach et al. 2003).

In contrast to reduced form estimation, *structural estimation* first builds a decision model of the situation, similar to the ones used in operations research.

Figure 1 Comparison of Different Approaches to the Newsvendor Problem



When using structural estimation, the econometrician assumes that decision makers already act rationally (and thereby optimally), and then uses observed decision-making behavior (in our case the reservation decision Q) to impute the underlying parameters of the decision model for which this behavior is rational (see Figure 1, right). Structural estimation has actively been used in several fields of economics, including labour economics and industrial organization.¹ With the exception of the work by Cohen et al. (2003), that estimates the cost parameters of a supplier in a semiconductor manufacturing context, structural estimation has had very few applications in the operations management literature.

The application that we present relates to managing capacity in a health-care environment. An important stream of the operations research literature has created a number of tools that directly or indirectly relate to the management of health-care capacity and its utilization (see Green 2004 for an overview). Given that patient demand for health-care services is inherently uncertain, the newsvendor model has found interesting ground for its application. At the strategic level, decisions need to be made with respect to sizing the care capacity. This includes choosing occupancy rates (e.g., Smith-Daniels et al. 1988, Huang 1995, Green and Nguyen 2001), making staffing decisions (e.g., Aiken et al. 2002, Kwak and Lee 1997, Green and Meissner 2002), and choosing the right panel size for physicians (Green et al. 2007). At the tactical level, decisions need to be made with respect to allocating capacity to various demand types (e.g., Green et al. 2003), such as the allocation of OR time to services in

a hospital (Strum et al. 1997). Several of these decisions resemble the newsvendor model and will be discussed more explicitly in §4.

3. Econometric Framework

The first-order condition (2), which defines the optimal decision Q^* , is essential to our imputed cost framework. This equation, commonly referred to as the critical fractile solution, provides a direct relationship between the overage/underage cost ratio and the probability of overestimating D at the optimum. Suppose that we observe a sequence $(D_i, Q_i)_{i=1, \dots, n}$ of realizations of the random variable D and the observed decision Q made prior to each realization. We can then define the fraction of cases in which overage costs were incurred as $\bar{I} = (1/n) \sum_{i=1}^n 1\{D_i \leq Q_i\}$, where $1\{\cdot\}$ denotes the indicator function. \bar{I} provides a crude estimate of the probability of overage of the newsvendor. Assuming that the newsvendor is behaving rationally and that the overage/underage ratio is constant among all observations $i = 1, \dots, n$, we can replace \bar{I} for $F(Q^*)$ in (2) to obtain


$$\frac{1}{1 + \gamma} = \bar{I}, \quad (3)$$

which gives $\gamma = 1/\bar{I} - 1$, or equivalently, $C_o = (1/\bar{I} - 1)C_u$.

While the analysis outlined by Equation (3) is useful as a preliminary data analysis, it suffers from three important problems. First, it does not allow for any statistical tests, which raises the question of whether the cost of overage is larger than the cost of underage (or vice-versa) with statistical significance. Second, this approach ignores the underlying heterogeneity of random component D and the ability of the newsvendor to partially anticipate this heterogeneity. For example, there might be seasonal variation

¹ See Reiss and Wolak (2006) for a review of structural estimation in industrial organization.

Figure 2 General Econometric Framework



in demand, which changes the distribution D across observations. Because the newsvendor can anticipate this seasonal demand change before choosing the quantity Q , she will adjust his quantity accordingly. Third, the cost ratio, $\gamma = C_o/C_u$, could vary across observations. In a retail example, the cost of a lost sale (which is related to the underage cost C_u) might vary with the time of the year or with the margin of the product. Moreover, some factors may affect both the distribution of the random variable and the cost ratio γ at the same time (e.g., price changes affect the demand distribution and the cost of underage faced by a retailer). Thus, a more elaborate model is needed. Below, we develop an estimation framework for the newsvendor problem which incorporates heterogeneity in both the random component, D_i , and the overage/underage cost ratio, γ . Figure 2 illustrates this framework, which is sufficiently general to estimate different applications of newsvendor problems.

The upper right part of Figure 2 accounts for heterogeneity in the random component D_i . Specifically, we assume that the D_i s are given by independent random variables from a common family of distributions $\{F(\cdot; \theta): \theta \in \Theta\}$, where θ is a vector parameter from the parameter space Θ which characterizes each member of the class. The distribution of the duration of procedure i , denoted D_i , is given by $F(\cdot; \theta_i)$. We let this distribution depend on a vector of covariates X_i , which can include different kinds of variables depending on the context. Following common econometric practice (e.g., Bickel and Doksum 2001), we assume the functional form

$$\theta_i = h(X_i, \eta), \tag{4}$$

where η is a vector of parameters to be estimated. Throughout the paper, we will denote covariates (e.g., X_i) as *row vectors* and parameters (e.g., η) as *column vectors*. Thus, the distribution of the random component for observation i , $F(\cdot; \theta_i)$, is characterized by the functional form of the distribution, the function $h(\cdot, \cdot)$, the vector η , and the vector of covariates X_i .

In addition to the ex ante heterogeneity of the random component, the newsvendor might face different trade-offs between overage and underage costs across observations; i.e., the relative cost parameter, γ , might differ on a case to case basis. This is captured in the upper left part of Figure 2. Similar to (4), we let the cost trade-off, γ_i , vary across cases:

$$\gamma_i = g(Z_i, \alpha), \tag{5}$$

where Z_i is a vector of covariates, $g(\cdot)$ is a link function, and α is a vector of parameters to be estimated. Note that the set of explanatory variables for the relative cost parameter underlying Equation (5), Z , may have a nonempty intersection with the set of explanatory variables for the random component, X , as outlined in Equation (4). However, the two sets are not necessarily identical.

Using Equations (4), (5), and (2), we can express the optimal decision Q_i^* as

$$F(Q_i^*; h(X_i, \eta)) = \frac{1}{1 + g(Z_i, \alpha)}. \tag{6}$$

Equation (6) specifies the optimally reserved time, Q_i^* , for each observation $i = 1, \dots, n$ in the data and thereby introduces the newsvendor solution into our estimation framework (see the lower part of Figure 2).

As the number of observations is much higher than the number of parameters that we wish to estimate, it is unlikely to find parameters α and η for which the observed decisions Q_i and the predicted optimal reservation time Q_i^* will exactly match. Thus, as in any econometric estimation, the model needs to account for some unexplained variation of Q_i . In the remainder of this section, we propose two models that account for this unexplained variation in different ways.

In the first model (Model N1), we assume that there are some unobservable (to the researcher) factors that are taken into account by the decision maker when determining the overage/underage ratio. Let ξ_i be an i.i.d. unobservable factor that affects the cost ratio for observation i , and assume that $E(\xi_i|Z_i) = 0$. Given that γ_i is strictly positive and based on Equation (5), we assume the following log-linear specification for the overage/underage cost ratio:

$$\log(\gamma_i) = Z_i\alpha + \xi_i. \quad (7)$$

If we knew γ_i , we could use linear regression to estimate α . Of course, the problem is that we do not know the true γ_i . What we do know is that a rational decision maker will behave according to the critical ratio, which can be rewritten as

$$\gamma_i = \frac{1}{F(Q_i; \theta_i)} - 1. \quad (8)$$

We propose the following two-step procedure to estimate α .

Step 1. Using data from the realizations of D_i , estimate η through maximum likelihood. Use the estimate $\hat{\eta}$ to compute fitted values $\hat{\theta}_i = h(X_i, \hat{\eta})$.

Step 2. Compute the fitted cost ratios $\hat{\gamma}_i = 1/F(Q_i; \hat{\theta}_i) - 1$, and then estimate α in the linear model $\ln(\hat{\gamma}_i) = Z_i\alpha + \xi_i$ through ordinary least squares (OLS). We refer to this procedure as TS-OLS.

In the online appendix (provided in the e-companion),² we show the consistency and asymptotic distribution of the estimator provided by TS-OLS.

The intuition behind TS-OLS is simple. In the first step, it estimates the distribution of D_i as seen by the newsvendor. Then, it computes the cost ratio $\hat{\gamma}_i$ that is consistent with the decision Q_i that was made by the newsvendor. Finally, it uses a regression to describe the variability of the cost ratios through the factors in Z_i . The asymptotic variance of the estimator provided by Step 2 of the OLS regression is adjusted for the estimation error incurred in Step 1.

Model N1 is not the only way to describe the unexplained variation in the observed decision. In

Model N2, we assume that the decision maker behaves approximately rational with some random deviation from the optimal decision. Given that $F(Q|\theta_i)$ is monotone in Q , we can invert Equation (6) to get

$$Q^*(W_i, \alpha, \eta) = F^{-1}\left(\frac{1}{1 + \exp(\alpha'Z_i)}; h(X_i, \eta)\right), \quad (9)$$

where $W_i = [X_i, Z_i]$. We assume that $E(Q_i | W_i) = Q^*(W_i, \alpha, \eta)$. Define the error term $\nu_i = Q_i - Q^*(W_i, \alpha, \eta)$. The model can be written as $Q_i = Q^*(W_i, \alpha, \eta) + \nu_i$, where ν_i is an i.i.d. random variable.

We proceed in a similar way as in the first model, and estimate α through a two-step nonlinear least squares method (TS-NLLS).³ This method can be summarized as follows.

Step 1. Using data from the realizations of D_i , estimate η through maximum likelihood.

Step 2. Use nonlinear least squares to estimate the equation $Q_i = Q^*(W_i, \alpha, \hat{\eta}) + \nu_i$.

We show in the online appendix the consistency and asymptotic distribution of the estimator provided by TS-NLLS.

Model N1 and Model N2 differ in several aspects. The main difference is that they rely on different assumptions to account for the unexplained variability of Q_i . Model N1 assumes that the newsvendor behaves optimally, but has private information regarding the cost ratio. Model N2 assumes that the Z vector describes all the factors that affect the cost ratio, but that the decision maker acts with a “trembling hand” around the optimal decision, i.e., the newsvendor acts optimally in expectation but the actual decision is adjusted by a zero-mean random variable. Which one of these assumptions is more appropriate depends on the context of the application. The two models also differ in the complexity of the estimation method. Both two-step methods, TS-OLS and TS-NLLS, are identical in the first step. In the second step, TS-OLS has a closed-form solution, while TS-NLLS requires the inversion of the distribution, which might have to be done numerically.


4. Application to Operating Room Time Reservation

Operating room (OR) management is a broad and complex problem which involves different levels of decision making. It includes strategic decisions such

³The parameters α and η in Equation (9) could be estimated directly through standard nonlinear least squares methods. We found that this approach may fail under some specifications due to identification problems. For example, suppose that the D_i s are i.i.d. and that the cost ratio is constant, i.e., X_i and Z_i contain only a constant. This implies that $Q^*(W_i, \alpha, \beta)$ is constant—the optimal decision is one and the same for all the observations, but there are multiple pairs (α, η) that yield this decision, which makes the model unidentifiable.

²An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Figure 3 OR Management Decision Time Line



as deciding how much OR capacity to put in place, as well as operational decision such as the scheduling of cases. Given the high complexity of the OR management problem, it is reasonable to decompose the problem into multiple hierarchical decisions. Figure 3 illustrates a time line with five important decisions related to OR management. Two of these decisions ((b) and (c)) involve trade-offs of reserving too much versus too little OR time, and therefore can be modeled as a newsvendor problem.

On a yearly basis, the hospital management determines the OR capacity needed based on future case workload. Dexter et al. (2005a) provide a methodology for making this type of decision using mixed-integer programming. The second stage in Figure 3 determines how much OR time to allocate to a specific block (defined as a set of interchangeable operating suites and personnel). Because the workloads of any given day is uncertain, the optimally allocated time must balance the costs of allocating too much time, which typically translates to idle time for the staff, with the costs of allocating too little time, which typically translates to overtime charges. Strum et al. (1997) develop a newsvendor model to find the optimal time to allocate to each block based on historical workloads. This decision can be revised annually to quarterly. This decision also sets the context in which subsequently case-level time reservations are performed.

At a lower level in the hierarchy, each service needs to decide how much OR time to reserve for any given case. This decision is done for each patient individually and is typically performed during preoperative planning (see Figure 3(c)). Reserving too much OR time to cases will very likely increase idle capacity. Reserving too little OR time will lead to more frequent schedule overruns and overtime hours for the hospital staff.⁴ When deciding how much OR time to

reserve to a case, the decision maker can use information from similar cases that were conducted in the past. This way, forecasts for the case durations can be constructed (Strum et al. 2000a provide goodness-of-fit tests for various duration distributions). Given a forecast of the case duration, the decision maker decides how much OR time to reserve for a specific case. It is important to make the distinction between the “forecasted time” for the case duration and the “reserved time” for the case. The former is a purely statistical concept, while the latter takes into account the overage and underage costs. The decision of how much OR time to reserve for a specific case can be modeled as a newsvendor problem, where the random component (D_i) is the actual duration of cases and the decision variable (Q_i) is the amount of OR time to be reserved. (Weiss 1990 and Charnetski 1984 follow similar approaches to model OR time reservation of individual cases.) We apply our structural estimation methods to this specific decision.

At some point after the OR time has been reserved to a case, the case needs to be scheduled to a specific day and time (see Figure 3(d)). This operational decision will depend on the convenience for the surgeon and patient as well as on the urgency of the case, among other variables. Scheduling the case is a decision that is separate from the previously discussed time reservation decision. Dexter and Traub (2002) provide heuristics for these type of decisions. Even closer to the day of the surgery (Figure 3(e)), OR time is *released*, i.e., scheduled cases are cancelled or moved leaving idle OR time which can be used to schedule other cases, cases can be moved, and emergencies are scheduled (see Dexter et al. 2003).

The decision of how much time to reserve to a given case provides an excellent context to apply

⁴ Predictable work hours are a key driver of employee satisfaction in the health-care industry. For example, Shader et al. (2001) link schedule stability with work satisfaction of nurses as well as

with work stress and employee turnover. Mueller and McCloskey (1990) identified eight dimensions of nursing job satisfaction, of which reliable scheduling is one. Similarly, Stachota et al. (2003) cited hours and schedules as one of the primary reasons for nurses terminating their employment, and Thompson and Brown (2002) identified schedule conflicts as a major driver of nursing turnover.

our structural estimation methods of the newsvendor model. First, there exists significant evidence in the medical and health-care literature that the newsvendor model is used in practice to reserve OR time to balance under and overutilization costs.⁵ Furthermore, these models have been incorporated into decision-support tools to assist OR management in hospitals and to optimize the staffing of these facilities (see <http://www.mda-ltd.com> and Dexter et al. 2001). Second, the estimation of surgical procedure duration has been extensively analyzed in the medical literature (e.g., Strum et al. 2000a, b). These statistical models can be incorporated in the first stage of our two-step procedure described in §3 to fit the distribution of case durations in our data set. Finally, the actual case duration, which is the random component of our model, is fully observed. This feature may not be present in those applications of the newsvendor model where demand is censored by the endogenous stocking quantity.

4.1. Model Specification

We apply our structural estimation method to the decision of how much OR time to reserve to a specific cardiac surgery case. The model input will be the observed reservation decision, Q , and various case characteristics. In addition, we also observe the actual durations of each of the surgery cases (D). Our objective is to estimate the cost ratio γ .

The first stage of our two-step procedure requires fitting the distribution of the duration for each case. Our unit of analysis, indexed by i , is an individual cardiac surgery case (e.g., a triple-bypass surgery for Mr. B conducted by surgeon W). The medical literature related to OR management suggests that the lognormal distribution provides a good statistical fit for the duration of surgery procedures (Strum et al. 2000a). The parameters of this distribution are determined by several case characteristics (Strum et al. 2000b), and can be anticipated by the decision maker. Let X_i denote the factors that describe the duration of case i . Recall that a random variable Y has log-normal distribution with parameters (μ, σ^2) if $\ln(Y)$ is normally distributed with parameters (μ, σ^2) . Following the notation defined in §3, we assume that the distribution of the duration of case i is characterized by the parameter vector $\theta_i = (\mu_i, \sigma_i^2)$. Given historical data of case durations, denoted D_i , and the characteristics of each case, denoted X_i , the actual duration of a case can be written as

$$\log(D_i) = X_i\beta + \varepsilon_i, \quad (10)$$

where the ε_i s are assumed to be i.i.d. normally distributed with mean zero and standard deviation σ .⁶

⁵ We thank two anonymous referees for suggesting references from the medical field.

Therefore, we have $\eta = (\beta, \sigma^2)$ and $h(X_i, \beta, \sigma^2) = (X_i\beta, \sigma^2)$. Estimating (β, σ^2) via maximum likelihood is equivalent to estimating β through OLS and σ based on the standard deviation of the regression residuals.

In fitting the log-normal distribution to cardiovascular procedures, we faced one statistical problem. Cardiovascular procedures are much longer than general surgery cases, and therefore are not well fitted by the log-normal distribution. May et al. (2000) suggest adding a third parameter, a location or “shift” parameter that we denote by δ , which fixes the lower bound of the support of D_i . This means that if D_i is the case duration, then $\tilde{D}_i \equiv D_i - \delta$ follows a log-normal distribution. We follow this approach and estimate the shift parameter as

$$\hat{\delta} = (d_{\max}d_{\min} - d_{\text{med}}^2)/(d_{\min} + d_{\max} - 2d_{\text{med}}),$$

where d_{\min} , d_{\max} , and d_{med} are the minimum, the maximum, and the median case duration observed in our cardiac surgery data set, respectively. For our sample, this estimate was 134.75 minutes. Note that for this particular application, the maximum likelihood estimates have closed-form solutions, and therefore numerical optimization routines are not required.

Now, we turn to the second step of the method. Using the log-normality of case duration, we have

$$F(Q_i; \hat{\theta}_i) = \Pr(D_i \leq Q_i) = \Phi\left(\frac{\ln(Q_i - \hat{\delta}) - X_i\hat{\beta}}{\hat{\sigma}}\right), \quad (11)$$

where $\Phi(\cdot)$ denotes the standard normal distribution. Replacing $F(Q_i; \theta_i)$ in Equation (8) with the fitted value (11), we obtain an estimate $\hat{\gamma}_i$ of the cost ratio. Model N1 estimates α via OLS as described in §3.7 Standard errors of the estimator for this specific application are described in the online appendix.

To estimate Model N2, we need to find $Q^*(W, \alpha, \eta)$. Combining Equations (9) and (11) gives

$$\begin{aligned} & Q^*(W_i, \alpha, \hat{\beta}, \hat{\sigma}^2) - \hat{\delta} \\ &= \exp\left\{\hat{\beta}X_i + \hat{\sigma} \cdot \Phi^{-1}\left(\frac{1}{1 + \exp(Z_i\alpha)}\right)\right\}. \end{aligned} \quad (12)$$

Defining $\tilde{Q}_i = Q_i - \delta$ and $\tilde{Q}^*(W_i, \alpha, \hat{\beta}, \hat{\sigma}^2) = Q^*(W_i, \alpha, \hat{\beta}, \hat{\sigma}^2) - \hat{\delta}$, the second step is equivalent to

⁶ To our knowledge, all the previous work that have analyzed the distribution of case durations assume homoskedasticity of the error term, which for the log-normal case implies a constant coefficient of variation across cases.

⁷ Note that our structural model implies that Q_i must be greater than δ for every i because it can never be optimal to reserve less than the minimum possible case duration. In our data set, all of the observed reservation times were above the estimate of the location parameter δ .

estimating the vector α in $\tilde{Q}_i = \tilde{Q}^*(W_i, \alpha, \hat{\beta}, \hat{\sigma}^2) + \nu_i$ via nonlinear least squares. Again, standard errors for the estimates in this application are given in the online appendix.

4.2. Data and Variable Definition

Our empirical analysis is based on 258 cardiac surgery cases. Details of the data collection procedure and the scope of the sample are defined in the online appendix. Our data set includes the date, the time of entry into the OR (*TimeIn*) and the time of departure from the OR, the amount of OR time reserved when the case was booked (Q_i) as well as the actual OR time (D_i), and patient characteristics and procedure characteristics. The actual and reserved times are measured in minutes, while *TimeIn* is measured in hours elapsed starting at 7 A.M.. Patient characteristics include a sex dummy ($SEX = 1$ if male) and *AGE* (in years, normalized to have mean equal to one). Procedure characteristics include the type of the main procedure conducted (see the bottom of Table 1), a dummy to indicate if the procedure was an emergency (*EMERG*), a dummy to indicate if more than one procedure was conducted during the operation (*MPROC*), and information about anesthesia classification (*ASA*). The conventional anesthesia risk-assessment score has six levels: (I) no systemic disease; (II) systemic disease, controlled; (III) systemic disease, symptomatic or uncontrolled; (IV) incapacitated; (V) in extremis, moribund; and (VI) brain death pronounced—organ donor. For cardiovascular procedures, most of the cases fall into categories III or IV (in our data set, no cases were classified I or VI, only one case was classified II, and only two were classified V). Therefore, we defined *ASA* as a dummy variable that is equal to one when the anesthesia

classification was equal to or above IV. We classified procedures into five categories which are defined, following Strum et al. (2000b), based on the current procedural terminology (CPT) codes. Note that each case is classified based on the actual procedures that were conducted. Even though we do not have pre-operation data on the procedures that were planned to be conducted, we learned that there is not much variation between the planned and the actual procedures that are conducted during a surgery. All cases that included a coronary artery bypass procedure were included in the *CABG* category, regardless of additional procedures performed. For these cases, we also coded a measure capturing the number of arteries bypassed (*NBYP*). Aortic valve replacements that included repair of the ascending aorta were classified as *AVR*. Cases of repair of the ascending aorta without *AVR* were included in *OTHER* procedures. Each case was conducted by one of four surgeons, specified by three dummy variables (*S1* through *S3*). Table 1 includes some descriptive statistics for these variables. All dummy variables are coded as binary {0, 1}.

Based on the work by Strum et al. (2000b) and conversations with the hospital management at our research site, we defined covariates for actual duration (X) and cost-ratio parameter (Z) as follows. In X , we included procedure information (dummies for each type of procedure, *NBYP*, *MPROC*, *ASA*, and *EMERG*), patient information (*AGE* and *SEX*), and dummies for surgeons. Even though we would expect postoperative measures to have significant explanatory power for actual duration, we did not include them as covariates in X because this information is not available to the hospital management when reserving the OR. The *NBYP*, *ASA*, and *EMERG* measures are good proxies for case severity, providing valuable information to predict actual procedure duration. Surgeon dummies are included to account for different levels of experience of the surgeons, which can affect actual duration. Previous research in the medical literature (Strum et al. 2003) has shown that predicting duration of multiple procedure cases can be difficult, mainly due to the lack of sufficient historical data for each procedure combination. Given the small size of our data set, we opt to include a single dummy to indicate multiple procedures. This simple approach could be improved if more data on multiple procedure cases were available. In Z , we included the same procedure characteristics (dummies for each type plus *MPROC*, *ASA*, and *EMERG*), dummies for surgeons (*S1*, *S2*, and *S3*) and *TimeIn*. We include procedure characteristics mainly because procedures may use different resources of the hospital, which would affect the overage and underage costs. Surgeon dummies are included to control for potential differences in overtime costs for the surgeons. The

Table 1 Summary Statistics

Variables	Mean	Std. dev.	Min.	Max.
<i>NBYP</i>	1.74	1.59	0.00	6.00
<i>AGE</i>	1.00	0.21	0.32	1.42
<i>TimeIn</i>	2.76	3.03	0.13	15.58
<i>Q</i>	338.19	42.13	210.00	467.00
<i>D</i>	366.69	84.35	215.00	742.00
<i>MPROC</i>	0.20			
<i>SEXdum</i>	0.67			
<i>EMERG</i>	0.11			
<i>ASA</i>	0.45			
<i>S1</i>	0.47			
<i>S2</i>	0.23			
<i>S3</i>	0.15			

Procedure	Abbreviation	Mean
Aortic valve replacement	AVR	0.17
Coronary artery bypass graft(s)	CABG	0.62
Mitral valve repair	MV	0.06
Mitral valve replacement	MVR	0.09

Table 2 Estimation Results for Models N1 and N2

Coefficient	Models N1 and N2	Coefficient	Model N1	Model N2
<i>X</i> covariates		<i>Z</i> covariates		
<i>Intercept</i>	5.1757 (0.1169)***	<i>Intercept</i>	0.0065 (0.5475)	-0.4157 (0.4473)
<i>AVR</i>	-0.0796 (0.0854)	<i>TimeIn</i>	-0.0083 (0.0256)	-0.0028 (0.0230)
<i>CABG</i>	-0.1295 (0.1041)	<i>AVR</i>	0.0642 (0.5082)	0.5458 (0.4196)
<i>MV</i>	0.0237 (0.1034)	<i>CABG</i>	-1.0991 (0.5784)*	-0.7279 (0.4977)
<i>MVR</i>	-0.0471 (0.0925)	<i>MV</i>	0.2734 (0.4969)	0.6697 (0.4033)*
<i>NBYP</i>	0.1019 (0.0229)***	<i>MVR</i>	-0.2039 (0.5853)	0.0834 (0.4476)
<i>MPROC</i>	0.3862 (0.0468)***	<i>NBYP</i>	0.6633 (0.0963)***	0.6532 (0.1052)***
<i>SEX</i>	0.0920 (0.0382)**	<i>MPROC</i>	1.7790 (0.1940)***	1.6981 (0.2306)***
<i>AGE</i>	0.0527 (0.0929)	<i>S1</i>	-0.6768 (0.2303)***	-0.6488 (0.2263)***
<i>EMERG</i>	0.0947 (0.0598)	<i>S2</i>	-0.1081 (0.2679)	-0.1031 (0.2561)
<i>ASA</i>	0.0973 (0.0379)**	<i>S3</i>	-0.0054 (0.2805)	0.0166 (0.2608)
<i>S1</i>	-0.1824 (0.0563)***	<i>EMERG</i>	-0.7071 (0.2799)**	-0.7002 (0.2411)***
<i>S2</i>	-0.0859 (0.0588)	<i>ASA</i>	0.3876 (0.1605)**	0.2958 (0.1610)*
<i>S3</i>	-0.0765 (0.0644)			
σ^2	0.0784 (0.0001)***			

Notes. Left (right) of the table shows the estimates for the first (second) step. Standard errors are shown in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% confidence level.

TimeIn covariate allows the cost ratio to vary during the day.

The service we analyzed was allocated one OR per day from 7:30 A.M. to 7 P.M. (11.5 hours). Hence, cardiac surgeons never conducted multiple surgeries in parallel by moving back and forth across ORs. Also, in our setting, a surgeon only very rarely conducted two consecutive surgeries on one day. This specific setting is somewhat unusual for a large hospital. But it provides a perfect empirical setting to apply our estimation framework because it supports the assumption of independence in the time allocation across any two cases. This unique setting clearly limits the generalizability of our cost estimates to other hospitals. In other words, while our estimation method applies to other settings (including other applications of the newsvendor model), our cost estimates do not.

5. Results

Our estimation results are summarized in Table 2. First, consider the variables influencing actual times (variables in *X*; see the left part of Table 2).⁸ As we can see, the variation in case durations can partly be explained by patient characteristics such as *SEX* as well as by variables describing case severity (*ASA* and *NBYP*). Cases with multiple procedures (*MPROC*) tend to take longer. In other words, case durations are not identically distributed. The coefficient of determination (R^2) of the regression is equal to 0.39, which reflects that a significant fraction of the variation in D_i can be predicted through the factors in X_i . The point estimate of σ^2 is approximately 0.08, which

⁸Note that because the first step of the estimation method is the same for the two models, the coefficients for *X* in Models N1 and N2 are identical.


reflects that there still is uncertainty in predicting actual duration after controlling for patient and procedure characteristics.

Second, consider the estimation of the cost-ratio equation (variable *Z*; see the right part of Table 2). As we can see in Table 2, for cases with more than one procedure (*MPROC*), increased emphasis was placed on the costs of OR idle time. The same holds for the number of bypass arteries (*NBYP*) in cardiac bypass surgeries and the anesthesia risk factor (*ASA*). This reflects that complicated cases may use key hospital resources which have high utilization rates, increasing the overage cost for these procedures.

We also observe that emergency cases tend to have a lower cost ratio. Because emergency cases use dedicated resources which are not shared with regular scheduled operations (for example, at nights or during weekends), idle capacity has a lower impact on costs, lowering the cost-ratio parameter. Moreover, because nearly all emergencies are performed within 24 hours of reserving the OR time, the decision maker can use last-minute information such as cancellations of other procedures to schedule emergencies in “time windows” that would not otherwise be utilized, further decreasing overage costs.

After looking at the drivers of the cost parameter γ , we now estimate the numeric value of γ . Toward this task, we compute the predicted values, $\gamma_i = Z_i \hat{\alpha}$, for the observations in our sample. The resulting histogram is shown in Figure 4. The median for estimates of γ for Model N1 (Model N2) is 1.79 (1.56). This is evidence that the hospital indeed emphasizes the costs of OR idle time.

We used the estimate of the cost ratio γ and its associated standard errors to test if γ_i would be larger

Figure 4 Histogram of Cost Ratio $Y_i = Z_i\alpha$ for Models N1 and N2

than one with statistical significance.⁹ For Model N1 (Model N2), we could reject the hypothesis of a cost ratio less than one, $H_0: \gamma_i \leq 1$, at the 95% confidence level for 44% (46%) of the cases, in favor of the alternative hypothesis $H_1: \gamma_i > 1$. Only 23% (21%) of the cases showed a cost ratio significantly less than one for Model N1 (Model N2).

Finally, Figure 4 also shows a substantial heterogeneity in γ . The centered R^2 of the second step regression in Model N1 is equal to 0.40, which reflects that a substantial part of this cost heterogeneity can be explained by the factors in Z . An F -test rejects the null of all Z coefficients being equal to zero (p -value less than 10^{-4}).

We conducted out-of-sample goodness-of-fit tests using Models N1 and N2. The results are described in the online appendix.

While our structural model exhibits good fit to the data, the magnitude and heterogeneity of our cost estimates seem to be at odds with some of the models developed in the OR management literature. Strum et al. (1999) suggest that the relevant costs in allocating OR time are staffing costs, which lead to a cost ratio γ less than one because overtime hours are more expensive than regular hours. In contrast, our results suggest an average cost ratio above one. In the next section, we develop two alternative models that extend the structural models developed in §3. These models provide results which are better aligned with those reported in the OR management literature.

⁹ Asymptotic standard errors for the cost-ratio parameter were obtained via the delta method.

6. A Newsvendor Model with Forecasting Bias

The structural models developed in §3 consistently estimate the overage/underage cost ratio under the assumption that the newsvendor decision Q is based on an unbiased estimate of the distribution of the random component D . However, Dexter et al. (2005b) report systematic biases in OR time forecasts provided by surgical services. There are two main reasons why we believe that the magnitude and variation of the cost estimates of Models N1 and N2 may be driven by systematic forecasting biases. First, we observe that more complicated cases are overrun more often, which is consistent with previous descriptive results related to “overconfidence” developed in other managerial settings (see Kahneman and Tversky 2000). Consider the two subsamples of cases with multiple procedures and single procedures, which correspond to 20% and 80% of the total cases, respectively. The average time required to complete multiple (single) procedure cases is around 430 minutes (350 minutes). Comparing the schedule overruns across these two types of cases, we find that multiple procedure cases go over the schedule 75% of the time, while low-complexity cases go over the scheduled time about 60% of the time. Hence, multiple procedure cases are 1.25 times more likely to go over the scheduled time, which suggests that the decision maker overestimates his ability to perform complex procedures on time and might be able to improve her forecasting capability.

Second, forecasting biases may also arise due to incentive conflicts between the agents participating

in the reservation process. As noted by Dexter et al. (2005b, p. 936), systematic biases can be caused by surgeons “underestimating their case durations to get their cases to ‘fit’ into their allocated OR time.” In our data, we find that emergencies are overrun 54% of the time versus 64% for nonemergency cases. For emergency cases, surgeons can use patient urgency as their “lever” to get the case on schedule soon, reducing the incentive to misreport their estimated duration for the case. This suggests that incentive conflicts may be a cause for the frequent schedule overruns in our sample.

We now extend our econometric framework to account for forecasting biases in the newsvendor’s forecast of D . Continue to assume the same distribution for case duration (log normal with parameters $\mu_i = \beta X_i$ and σ). Instead of assuming a perfectly rational decision maker, we assume that the reservation time is based on a *biased* estimate of the mean distribution. Let D_i^b be the perceived duration of case i , so that $\log D_i^b$ is normal with mean $\mu_i^b = b_i + \mu_i$ and standard deviation σ . For this biased forecast, the reservation decision \tilde{Q}_i is the solution to $\Phi((\ln(\tilde{Q}_i) - (b_i + \mu_i))/\hat{\sigma}) = 1/(1 + \gamma_i)$. Rearranging, we obtain

$$\ln(\tilde{Q}_i) - \mu_i = b_i + \sigma \Phi^{-1}\left(\frac{1}{1 + \gamma_i}\right). \quad (13)$$

The bias term b_i may vary across cases. For example, surgeons may have less incentive to underestimate emergency procedures relative to nonemergency procedures. We can model the bias as

$$b_i = B + \varphi' W_i + \omega_i, \quad (14)$$

where W_i is a vector of observable characteristics (without an intercept), B is a constant, and ω_i denote unobservable factors that affect the bias. We assume that ω_i and ε_i are independent. We also transform W by subtracting its mean, so that B represents the average forecast bias. Under these behavioral assumptions and a cost ratio of the form $\gamma_i = \exp(\alpha' Z_i)$, the actual reservation decision becomes

$$\begin{aligned} \ln(\tilde{Q}_i) - \mu_i \\ = B + \varphi' W_i + \sigma \Phi^{-1}\left(\frac{1}{1 + \exp(\alpha' Z_i)}\right) + \omega_i. \end{aligned} \quad (15)$$

Model (15) can have poor identification of the parameters (B, φ, α) . To see this, suppose that Z_i contains only a constant, so that the cost ratio is constant. Define the critical value $t(\gamma) = \Phi^{-1}(1/(1 + \gamma))$. Note that even if we know σ , we cannot identify $t(\gamma)$ and B separately. In other words, we cannot identify whether the observed reservation times are due to an

underestimation of procedure duration (low B) or due to a high cost ratio γ (low t).¹⁰

To obtain a more informative model, we impose restrictions on the parameters of (15) based on external information. One alternative is to restrict the model to have a constant cost ratio, leading to the following model:

$$\ln(\tilde{Q}_i) - \mu_i = c + \varphi' W_i + \omega_i, \quad (16)$$

where $c = t(\gamma)\sigma + B$. In a study of OR block allocation conducted at two university hospitals, Abouleish et al. (2003) suggest that underage costs are 75% higher than overage costs, so we set $\gamma = 1/1.75 = 4/7$. We refer to (16) as Model B1. In this model, the variance of the adjusted reserved time $\ln(\tilde{Q}_i) - \mu_i$ is explained by heterogeneity in the forecasting bias across cases. Given an estimate of c and a value for γ , we can calculate the implicit average bias B and the coefficient vector φ .

Another alternative is to fix the forecasting bias. Letting $\varphi = 0$ and fixing the average bias at $B = B_o$, model (15) becomes

$$\ln(\tilde{Q}_i) - \mu_i = B_o + \sigma \Phi^{-1}\left(\frac{1}{1 + \exp(\alpha' Z_i)}\right) + \omega_i. \quad (17)$$

We refer to (17) as Model B2. Following Dexter and Ledolter (2005), we define the proportionality bias as $\rho = E(D_i^b)/E(D_i)$. Based on the results reported by Strum et al. (1999) and Dexter et al. (2005b), we fix the value of B_o so that $\rho = 0.8$, i.e., the newsvendor forecast is 20% below the true forecast. Model B2 provides a new estimate of α , which we use to compute a bias-adjusted average cost ratio.

We estimated Models B1 and B2 using a two-step method similar to TS-NLLS. The properties of the estimators and the calculation of the standard errors are detailed in the online appendix. The estimate for the φ coefficients of Model B1 are reported in Table 3. The estimate for the intercept c implies a proportionality bias equal to $\rho = 85\%$ (approximately).¹¹ We also have added the estimates of the corresponding X coefficients in Table 3 for comparison. The results show that longer procedures tend to exhibit a significantly larger bias. Furthermore, the magnitude of the estimates suggest that, while the decision maker seems

¹⁰ However, if case duration is heteroskedastic so that $\text{Var}(\varepsilon_i) = \sigma_i^2$ is not constant across cases, then it is possible to identify t and B separately, and so γ can be identified. In our results, postregression statistical tests do not provide significant evidence of heteroskedasticity. This can be due to the relatively homogenous cases in our sample.

¹¹ This can be calculated by noting that $\rho = \exp(B + \varphi' W_i) \cdot E(\exp(\omega_i))$. Assuming ω_i is normally distributed with mean zero, we obtain $E(\exp(\omega_i)) \approx \exp(\hat{s}^2/2)$, where \hat{s} is the standard deviation of the residuals ω_i in Equation (16).

Table 3 Estimates for Model B1 with Cost Ratio Equal to 4/7

Coefficient	Model B1	X coefficients
<i>c</i>	−0.0906 (0.0116)***	
<i>MPROC</i>	−0.2799 (0.0286)***	0.3862 (0.0468)***
<i>ASA</i>	−0.0632 (0.0250)**	0.0973 (0.0379)**
<i>NBYB</i>	−0.1065 (0.0152)***	0.1019 (0.0229)***
<i>CABG</i>	0.1751 (0.0522)***	−0.1295 (0.1041)
<i>S1</i>	0.1098 (0.0354)***	−0.1824 (0.0563)***
<i>S2</i>	0.0238 (0.0411)	−0.0859 (0.0588)
<i>S3</i>	0.0022 (0.0431)	−0.0765 (0.0644)
<i>EMERG</i>	0.1126 (0.0429)***	0.0947 (0.0598)
<i>Timeln</i>	0.0018 (0.0040)	
<i>R</i> ²	0.4099	

Notes. The two right columns are the coefficient estimates of the actual duration estimation, shown for comparison. Standard errors are shown in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% confidence level.

to be adjusting its forecast for longer procedures, this adjustment is insufficient. To explore this issue further, we looked at the average reserved and actual times of cases with low and high ASA risk factor for each of the four surgeons in our sample, which are illustrated in Figure 5. The figure shows how the adjustment made in the reservation time, while in the right direction, is insufficient to account for all the heterogeneity among surgeons. Similar results were found by Cachon and Schweitzer (2000), who study newsvendor decisions made in a controlled experiment. They find that subjects make an insufficient adjustment in their newsvendor decisions and tend to be biased toward the mean. In contrast to the substantial forecasting bias found for an average case, the positive sign and magnitude of the φ coefficient for *EMERG* suggests that emergency cases exhibit almost no forecasting bias.¹² This suggests that surgeons are intentionally underestimating the duration of nonemergency cases to get them sooner in the schedule.

The estimates of Model B2¹³ imply an average bias-adjusted cost ratio of approximately 0.5. This adjusted cost measure is in line with the overage/underage cost ratio reported by Strum et al. (1999).


7. Discussion and Conclusion

As in any single-site research study, one should be cautious to generalize our estimation results to other hospitals; the methods that we developed are generalizable to other settings, but the reported estimation results are specific to our study. Specifically, the trade-offs present in our application depend on how much OR time was allocated to cardiac services at the aggregate level (see Figure 3(a)). The estimated costs are

¹² Recall that the *W* variables are demeaned, so the average bias for emergency cases is $B + \varphi_{EMERG} \cdot (1 - EMERG_i) \approx -0.02$.

¹³ The coefficient estimates of Model B2 are not reported but can be obtained from the authors upon request.

Figure 5 Average Actual Case Duration (*D*) and Reserved Time (*Q*) by Surgeon for Single-Procedure Cases with Low (Top) and High (Bottom) ASA Risk Factor



Note. Numbers above the bars indicate the number of observations on each group.

valid to the cardiac service we analyze under the fixed allocated time of one OR, 11.5 hours per day.

We believe that our results are of substantial interest, both from an academic and a managerial perspective. From an academic perspective, the main contribution of this paper is to provide a general structural model to impute the overage and underage costs in newsvendor-type decisions. Our models are sufficiently general to allow for arbitrary parametric distributions of the random variable, and can accommodate observed heterogeneity in this distribution. The model also allows for observed and unobserved heterogeneity in the overage/underage cost ratio, and therefore can be used to compute different cost estimates for each observation in the sample. We develop methods that give consistent estimates of the parameters of each of the two models, and derive the asymptotic distribution of the estimators which can be used to compute standard errors of the estimates. Therefore, our methodology can be used to conduct hypothesis testing, and is useful for empirical research.

We applied our structural estimation methods to the decision of how much OR time to reserve to a specific surgical case using real data from cardiac surgery. From the perspective of health-care management, our analysis reveals that the hospital under-

Table 4 Comparison of Structural Models

Model	Assumptions				Output
	Cost ratio	Forecast bias	Unobservables	Input	
N1	Varies among cases	None	Cost heterogeneity	Q, D, X, Z	Average cost ratio = 1.7 Factors affecting cost ratio
N2	Varies among cases	None	Deviations from optimal reservation (additive)	Q, D, X, Z	Average cost ratio = 1.6 Factors affecting cost ratio
B1	Constant and equal to 4/7*	Varies among cases	Bias heterogeneity	Q, D, X, W, γ	Average forecasting bias = 15% Variation of bias among cases
B2	Varies among cases	80% of true forecast**	Cost heterogeneity	Q, D, X, Z, B	Average cost ratio = 0.5 Factors affecting cost ratio

*As reported by Abouleish et al. (2003).

**As reported by Dexter et al. (2005b).

lying this study is apparently placing much greater emphasis on OR idle time compared to delays and running over the scheduled time. Specifically, we show that the costs of OR idle time were perceived, in average, as approximately 60% higher than the cost of schedule overrun. It should be emphasized that using such cost parameters is not right or wrong per se: it simply reflects how the hospital balances partly conflicting objectives. It is the role of the hospital administration to evaluate the alignment of these cost estimates with the overall strategic objectives of the hospital.

The structural models we developed are relevant in practice and can be used in different ways. Table 4 summarizes the differences across the four models developed. Consider Models N1 and N2, which assume a fully rational decision maker with no forecasting biases. We found the cost ratio to be approximately 1.6 in this case. This is a helpful piece of information for the hospital management to diagnose if the current system is in line with the expectations and objectives of the hospital. As an example, consider a hospital manager who views the costs of overtime as 75% higher than idle time costs. Further, assume that using Model N1 or N2, she finds that the hospital appears to value overtime costs 50% less than the costs of idle capacity, as we find in our research. This provides statistically significant evidence that there exists a mismatch between the management objectives and the actual behavior of the system. Our analysis is useful to identify potential biases introduced by surgeons, while controlling for the procedure mix performed by each surgeon. A more detailed analysis of the reservation process is in order.

This is where Model B1 can be used. Model B1 helps hospital management to diagnose whether there exists a bias in the forecasting process, and if so, to identify the variables that lead to larger biases. This forecasting bias may arise because of insufficient forecasting capability or due to incentive conflicts. For

example, our hospital manager might find that the hospital systematically underestimates the duration of complex cases, as we did in our research. The system would benefit from an enhanced forecasting system, potentially based on an implementation of the methods developed by Dexter and Ledolter (2005).

Finally, Model B2 provides a refined estimate for the costs of too much versus too little capacity reservation. Unlike Models N1 and N2, Model B2 corrects for the forecasting bias, and hence leads to more realistic estimates of the underlying cost parameters. Based on the adjustment for the forecasting bias, we find the average cost ratio to be 0.5, which is in line with the earlier study by Strum et al. (1999).

One advantage of using structural estimation is that it provides a better understanding of the mechanism by which the different factors affect decisions. In the context of the newsvendor, we can disentangle whether a specific factor affects the observed decision Q_i through the distribution of the random variable D_i or through the overage/underage cost ratio γ_i . This can be helpful for a prescriptive analysis of the system. For example, often it might be easier to adjust factors that affect the cost ratio than changing factors that affect D_i . In addition, disentangling these effects can provide a more robust tool to do prospective analysis when major changes in the system are introduced. For example, we could use our model to measure the economic impact of subsidies of overage and underage costs. In a decentralized supply chain, the imputed parameters could be used to design contracts to coordinate the supply chain and increase efficiency. Most of the contracts suggested in the literature that coordinate newsvendor decisions depend on the overage and underage costs (see Cachon 2003). However, these costs are usually private information of each of the agents negotiating the contract, who might not want to reveal them during the bargaining process. Our structural model can be used to impute these cost parameters from historical data, which can facilitate the specification of such contracts. All of these

contributions easily carry over to the broad range of existing newsvendor applications.

We found the application domain of hospital capacity planning to be particularly well suited for structural estimation methods because hospital operations have been researched both from an analytical and an empirical perspective. We believe that future research could apply our estimation methods to other hospital decisions, such as inventory decisions at blood banks, service-level decisions of trauma surgeons, or resource allocations for elective and emergency procedures. Given the broad range of newsvendor applications in operations management, however, the potential usage of our econometric framework extends to supply chain management, capacity planning, and project management.

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mansci.journal.informs.org/>.

Acknowledgments

The authors thank Gerard Cachon, Justin Ren, Elena Krasnokutskaya, Sergei Savin, Serguei Netessine, and attendants to the Empirical Workshop in Operations Management at The Wharton School (2006) for their helpful ideas and comments. The authors are also grateful to the departmental editor and the review team for a very constructive review process.

References

- Abouleish, A., F. Dexter, R. Epstein, D. Lubarsky, C. Whitten, D. Prough. 2003. Labor costs incurred by anesthesiology groups because of operating rooms not being allocated and cases not being scheduled to maximize operating room efficiency. *Anesthesia Analgesia* **96** 1109–1113.
- Aiken, L. H., S. P. Clarke, D. M. Sloane, J. Sochalski, J. H. Silber. 2002. Hospital nurse staffing and patient mortality, nurse burnout, and job dissatisfaction. *J. Amer. Medical Assoc.* **288** 1987–1993.
- Berry, S., J. Levinson, A. Pakes. 1995. Automobile prices in market equilibrium. *Econometrica* **63**(4) 841–890.
- Bickel, P. J., K. A. Doksum. 2001. *Mathematical Statistics*, 2nd ed. Prentice Hall, Upper Saddle River, NJ.
- Brush, T., A. Karnani. 1996. Impact of plant size and focus on productivity: An empirical study. *Management Sci.* **42**(7) 1065–1081.
- Cachon, G. 2003. Supply chain coordination with contracts. S. Graves, T. de Kok, eds. *Handbooks in Operations Research and Management Science: Supply Chain Management*, Vol. 11. Elsevier, Amsterdam, 229–339.
- Cachon, G., M. Schweitzer. 2000. Decision bias in the newsvendor problem with known demand distribution: Experimental evidence. *Management Sci.* **46**(3) 404–420.
- Charnetski, J. R. 1984. Scheduling operating room surgical procedures with early and late completion penalty costs. *J. Oper. Management* **5** 91–102.
- Chorba, R. W. 1976. Potential avoidability: A statistic for controlling in-patient utilization in acute care hospitals. *Management Sci.* **22**(6) 694–700.
- Cohen, M. A., T. H. Ho, J. Z. Ren, C. Terwiesch. 2003. Measuring imputed cost in the semiconductor equipment supply chain. *Management Sci.* **49**(12) 1653–1670.
- Dexter, F., J. Ledolter. 2005. Bayesian prediction bounds and comparisons of operating room times even for procedures with few or no historical data. *Anesthesiology* **103** 1259–1267.
- Dexter, F., R. Traub. 2002. How to schedule elective surgical cases into specific operating rooms to maximize the efficiency of use of operating room time. *Anesthesia Analgesia* **94** 933–942.
- Dexter, F., R. Epstein, H. M. Marsh. 2001. A statistical analysis of weekday operating room anesthesia group staffing costs at nine independently managed surgical suites. *Anesthesia Analgesia* **92** 1493–1498.
- Dexter, F., J. Ledolter, R. Wachtel. 2005a. Tactical decision making for selective expansion of operating room resources incorporating financial criteria and uncertainty in subspecialties' future workloads. *Anesthesia Analgesia* **100** 1425–1432.
- Dexter, F., R. Traub, A. Macario. 2003. How to release allocated operating room time to increase efficiency: Predicting which surgical service will have the most underutilized operating room time. *Anesthesia Analgesia* **96** 507–512.
- Dexter, F., A. Macario, R. Epstein, J. Ledolter. 2005b. Validity and usefulness of a method to monitor surgical services' average bias in scheduled case durations. *Canadian J. Anesthesia* **52** 935–939.
- Green, L. V. 2004. Capacity planning and management in hospitals. M. L. Brandeau, F. Sainfort, W. P. Pierskalla, eds. *Operations Research and Health Care: A Handbook of Methods and Applications*. Kluwer Academic Publishers, Norwell, MA, 15–42.
- Green, L. V., J. Meissner. 2002. Developing insights for nurse staffing. Working paper, Columbia Business School, New York.
- Green, L. V., V. Nguyen. 2001. Strategies for cutting hospital beds: The impact on patient service. *Health Services Res.* **36** 421–442.
- Green, L. V., S. Savin, M. Murray. 2007. Providing timely access to care: What is the right patient panel size? *Joint Commission J. on Quality and Patient Safety* **33**(4) 211–218.
- Green, L. V., S. Savin, B. Wang. 2003. Managing competing demands in a medical diagnostic facility. Working paper, Columbia Business School, New York.
- Huang, X. 1995. A planning model for requirement of emergency beds. *J. Math. Appl. Medicine Biol.* **12** 345–353.
- Kahneman, D., A. Tversky. 2000. *Choices, Values and Frames*. Cambridge University, Cambridge, UK.
- Kwak, N. K., C. Lee. 1997. A linear programming model for human resource allocation in a health-care organization. *J. Medical Systems* **21** 129–140.
- Lieberman, M. B., L. Demeester. 1999. Inventory reduction and productivity growth: Linkages in the Japanese automotive industry. *Management Sci.* **45**(4) 466–485.
- May, J., D. Strum, L. Vargas. 2000. Fitting the lognormal distribution to surgical procedure times. *Decision Sci.* **31**(1) 129–148.
- Milne, R. G., A. Abebe, B. Torsney. 1989. The impact of teaching on hospital costs: A budgetary approach to non-market institutions. *J. Oper. Res. Soc.* **40**(12) 1089–1098.
- Mueller, C. W., J. C. McCloskey. 1990. Nurses' job satisfaction: A proposed measure. *Nursing Res.* **39**(2) 113–117.
- Pell, J., J. Sirel, A. Marsden, I. Ford, S. Cobbe. 2001. Effect of reducing ambulance response times on deaths from out of hospital cardiac arrest: Cohort study. *British Medical J.* **322** 1385–1388.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford Business Books, Stanford, CA.
- Reiss, P., F. Wolak. 2006. Structural econometric modeling: Rationales and examples from industrial organization. Working paper, Stanford University, Stanford, CA, <http://www.stanford.edu/preiss/makeit.pdf>.

- Rust, J. 1987. Optimal replacement of GMC bus engines: An empirical model of Harold Zucher. *Econometrica* **55**(5) 999–1033.
- Shader, K., M. E. Broome, C. D. Broome, M. E. West, M. Nash. 2001. Factors influencing satisfaction and anticipated turnover for nurses in an academic medical center. *J. Nursing Administration* **31**(4) 210–216.
- Smith-Daniels, V. A., S. B. Schweikhart, D. E. Smith-Daniels. 1988. Capacity management in health care services: Review and future research directions. *Decision Sci.* **19** 889–919.
- Stachota, E., P. Normandin, N. O'Brien, M. Clary, B. Krukow. 2003. Reasons registered nurses leave or change employment status. *J. Nursing Administration* **33**(2) 111–117.
- Strum, D., J. May, L. Vargas. 2000a. Modeling the uncertainty of surgical procedure times. *Anesthesiology* **92**(4) 1160–1167.
- Strum, D., L. Vargas, J. May. 1999. Surgical subspecialty block utilization and capacity planning: A minimal cost analysis model. *Anesthesiology* **90**(4) 1176–1185.
- Strum, D., J. May, A. Sampson, L. Vargas. 2003. Estimating times of surgeries with two component procedures. *Anesthesiology* **98**(1) 232–240.
- Strum, D., A. Sampson, J. May, L. Vargas. 2000b. Surgeon and type of anesthesia predict variability in surgical procedure times. *Anesthesiology* **92** 1454–1466.
- Strum, D., L. Vargas, J. May, G. Bashein. 1997. Surgical suite utilization and capacity planning: A minimal cost analysis model. *J. Medical Systems* **21**(5) 309–322.
- Thompson, T. P., H. N. Brown. 2002. Turnover of licensed nurses in skilled nursing facilities. *Nursing Econom.* **20**(2) 66–69.
- Urbach, D. R., C. M. Bell, P. C. Austin. 2003. Differences in operative mortality between high- and low-volume hospitals in Ontario for 5 major surgical procedures: Estimating the number of lives potentially saved through regionalization. *Canadian Medical Assoc. Res. J.* **168**(11) 1409–1414.
- Weiss, E. N. 1990. Models for determining estimated start times and case ordering in hospital operating rooms. *IIE Trans.* **22** 143–150.