  0160-5682/03 $25.00

www.palgrave-journals.com/jors

Two-resource stochastic capacity planning employing a Bayesian methodology

E Stavrulaki[1,*], DKH Fong[2] and DKJ Lin[2]

[1]*Bentley College, Waltham, MA, USA; and* [2]*Penn State University, University Park, PA, USA*

We examine a stochastic capacity-planning problem with two resources that can satisfy demand for two services. One of the resources can only satisfy demand for a specific service, whereas the other resource can provide both services. We formulate the problem of choosing the capacity levels of each resource to maximize expected profits. In addition, we provide analytic, easy-to-interpret optimal solutions, as well as perform a comparative statics analysis. As applying the optimal solutions effectively requires good estimates of the unknown demand parameters, we also examine Bayesian estimates of the demand parameters derived via a class of conjugate priors. We compare the optimal expected profits when demands for the two services follow independent distributions with informative and non-informative priors, and demonstrate that using good informative priors on demand can significantly improve performance.
Journal of the Operational Research Society (2003) **54,** 1198–1208. doi:10.1057/palgrave.jors.2601607

Keywords: applied probability; Bayesian methodology; capacity planning; parameter estimation; single-period inventory models

Introduction

Capacity planning is an inherently difficult problem due to demand uncertainty, particularly in the service sector, because services often cannot be inventoried. As an example, consider the decision of determining how many single and double rooms to build for a hotel expansion project. Once the additional rooms of each type have been built, the hotel has to rely on the existing rooms to satisfy the daily fluctuations in demand. (It is practically impossible to build new rooms on a daily basis to perfectly match demand with capacity.) Recent increased customers' expectations for high service levels, as well as opportunities for substitution among the available resources (eg a customer requesting a single room can stay in a double room) complicate further the problem of planning for the right amount of capacity before demand materializes. Moreover, even when a manager is faced with a well-defined capacity planning problem, for which theoretically optimal solutions can be obtained, applying these solutions in practice requires the estimation of unknown parameters of the demand distribution. For this reason, we examine a stochastic capacity planning problem from both of these perspectives: deriving optimal solutions and estimating the unknown demand parameters.

We consider a problem with two resources and two products where demand for one of the products can be satisfied by either resource. We derive the expected-profit-maximizing capacity levels for general demand distributions and show that they correspond to modified 'critical fractile' ratios. These modified critical fractiles are similar to the ones derived for the classical single resource, single-period newsvendor model, but they also take into account the possibility of satisfying some demand using either of the two resources. We also investigate the impact of changes to the input and demand parameters via a comparative statics analysis, which highlights results that would have been difficult to predict otherwise. Finally, we compare the estimation of the unknown demand parameters via classical and Bayesian approaches.

A common practice when solving the capacity planning problem is to substitute sample estimates for the population parameters in the demand distribution. It is likely, however, that prior information regarding the distribution of the demand parameters is available from various sources, such as industry trends, customer surveys, expert opinions, etc. Ideally, one could use this prior information to derive better estimates for the demand parameters than those obtained from the standard sample estimates, and to generate a distribution of maximum profits, rather than a single expected profit value. A powerful method for encoding prior knowledge through subjective probabilities into the decision process is to use the Bayesian theory. In fact, in some cases (eg the normal demand case to be discussed here), classical estimates are actually Bayesian estimates with respect to certain non-informative priors. Hence, employing Bayesian estimates with informative priors could improve the decision-making process, especially if the noninformative priors are considered unlikely to be true. As modern computing power becomes inexpensive, applying the Bayesian paradigm in practice is not only feasible but also important. In their recent book, for example, French and

Correspondence: E Stavrulaki, 23 Hampshire Road, Peabody, MA 01960, USA.
E-mail: estavrulaki@bentley.edu

Smith[1] cited various real-world applications of Bayesian theory in diverse industries. In addition, many Bayesian applications have recently been developed in the marketing and economics literature (see for example, Fong and Bolton,[2] Pammer *et al*,[3] Fong *et al*,[4] and DeSabo *et al*[5]). In this paper, we demonstrate how we can quantify the value of incorporating prior information in a stochastic capacity planning setting with product substitution.

The remainder of this paper is organized as follows. First we provide a brief literature review. Then we present the model formulation, derive analytically optimal solutions, and perform a comparative statics analysis on the optimal solution. Next, the Bayesian procedures that we employ are discussed. Numerical work, with which we compare the conventional sample estimation approach with a Bayesian approach that assigns informative priors on the demand parameters, is then presented. The final section summarizes our findings and concludes.

Related work

There are two research streams related to this work: single-period capacity planning problems for multiple, substitutable resources (see Sethi and Sethi[6] for a review), and single-period stochastic inventory models for multiple, substitutable products (see Bassok *et al*[7] for a review). The works of Harrison and Van Mieghem,[8] Netessine *et al*[9] (stochastic capacity planning models), Bassok *et al*,[7] and Parlar and Goyal[10] (single-period inventory models) provide standard formulations that are closely related to ours. All of these authors, however, focus on providing conditions for optimality and investigating the properties of the optimal inventory policy, but do not provide easily interpretable analytic solutions, nor perform a comparative statics analysis. In contrast, by focusing on a two-product, two-resource setting, we are able to provide fairly simple equations that the optimal capacities must satisfy, as well as interpret these equations based on the classic single-product newsvendor model. In addition, unlike most of the papers in the aforementioned literature, with the exception of Hill,[11] we consider the aspect of estimating the demand parameters when prior information is available. Hill employs conjugate priors in the context of a standard newsvendor problem and demonstrates the value of Bayesian estimation for a single unknown parameter of three different demand distributions. Similarly, we investigate the effects of parameter estimation on expected profits in the two-resource, capacity-planning setting.

In contrast to the lack of Bayesian applications for single-period inventory models, the application of Bayesian theory to estimate demand parameters of *multiperiod* inventory models has been well studied (see Azouri,[12] for a recent literature review, and Bradford and Sugrue[13] for a two-period application). Most multiperiod inventory models, however, focus on the single-resource, single-product case, and typically obtain optimal Bayesian policies via stochastic dynamic programming formulations. In addition, most of these papers ignore the issue of lost sales, and do not consider the fixed costs associated with acquiring resources. While multiperiod settings may be appropriate in certain business environments, there are cases for which single-period inventory or capacity planning models, such as ours, are more appropriate; consider for example a case for which the planning horizon is too long to allow for repeated updates (as in our hotel expansion example). Thus, our interest is in integrating a Bayesian estimation procedure with a single-period, two-resource, capacity planning model.

In addition to single-period inventory models, revenue management models that focus on optimal stopping rules are somewhat related with this work (for a recent review of revenue management research see McGill and Van Ryzin[14]). This similarity stems from the fact that identifying the optimal booking limit for discount passengers (ie the optimal stopping rule) is based on a formula qualitatively similar to the single-period newsvendor result. However, even though research papers by Bodily and Weatherford,[15] Brumelle *et al*,[16] and Ladany[17] consider two-resource revenue management problems with some notion of substitution (ie allowing customers to upgrade) or Bayesian updating, the assumptions as well as the focus of these papers are significantly different from ours, and thus we do not review them further here; to illustrate, Brumelle *et al* assume that there is a positive probability that a customer denied a discount fare will upgrade to a full fare, whereas we (and capacity planning problems in general) do not account for such a probability, but instead consider the shortage cost of not satisfying all demand, and allow for upgrades only if there is excess capacity.

Capacity planning model with substitution

Assumptions and model structure

Consider a monopolist offering two services (products) at fixed prices p_i via two resources. Resource type 2 can only be used to satisfy demand for service 2, whereas resource type 1 can be used to satisfy either type of demand. That is, demand for service type 2 can be upgraded, but demand for service type 1 cannot be downgraded.

Given this structure, the monopolist must choose the capacity level T_i for each resource i, $i=1$, 2, to satisfy demand during a given time horizon (eg the hotel manager decides on the number of single and double rooms available in her hotel within the next 3 years). Note that this decision must be made before the monopolist observes demand. For each unit of type i capacity the firm acquires, it incurs a fixed cost G_i, even if that unit is not used to satisfy demand. The firm also incurs a variable cost v_i, for every unit of type i capacity it uses to satisfy demand.

After the monopolist buys capacity for each resource type, customers arrive randomly during the fixed time horizon and place their orders. Customers must either be served within their respective periods of arrival or be turned away. If the firm is unable to serve a customer by the available capacity, it must compensate that customer. In other words, the firm internalizes the cost of not satisfying demand by paying a shortage cost of c_i, for each unit of unmet demand.

Given the above input parameters (price p_i, fixed G_i and variable v_i capacity costs, and shortage cost c_i) and the firm's capacity choices T_1 and T_2, the *actual* demand values realized during the assumed time horizon determine the actual amount of profit (or loss) for the monopolist. However, since the monopolist will determine the actual profit only at the end of the time horizon, she must decide on the amount of capacities to acquire based on her *expected* profits. Thus, the monopolist's objective is to determine the optimal level of capacity of each resource type T_1 and T_2 so as to maximize expected profits.

To calculate expected profits, the monopolist must have some knowledge of the nature of randomness for the demands. We assume that, irrespective of *how* customers arrive during the given time horizon, the probability distribution function of the *total demand* for each service type is known. Moreover, we assume that the demands D_i, $i = 1, 2$, for the two services are independent and continuous random variables, with probability density functions f_i and cumulative distribution functions F_i, respectively.

We also assume the following with respect to our input parameters: (1) all input parameters (p_i, v_i, c_i, and G_i) are strictly positive, (2) $p_i - v_i > 0$, (3) $v_1 > v_2$, so that it is not optimal to substitute resource 1 for resource 2 whenever there is excess amount of resource 2 available, (4) $p_1 - c_1 > p_2 - c_2 > 0$, so that it is more profitable to satisfy unmet demand of type 1 rather than of type 2, and (5) $p_2 - v_1 > 0$, so that substitution of resource 2 from resource 1 when satisfying demand of type 2 is profitable.

Model formulation

We can now formally express both the net revenue and cost associated with choosing capacities T_1 and T_2. The net revenue function, denoted by $g(T_1, T_2)$, equals

$$g(T_1, T_2) = (p_1 - v_1)\min(D_1, T_1) + (p_2 - v_2)\min(D_2, T_2) + (p_2 - v_1)\min[(D_2 - T_2)^+, (T_1 - D_1)^+]$$

The last term of $g(T_1, T_2)$ describes the net revenue generated from a customer requesting service 2 whose demand is satisfied by resource 1. Note that this substitution is made only when there is excess amount of resource 1 (ie $T_1 > D_1$) and excess demand for service 2 (ie $D_2 > T_2$). Moreover, note that $g(T_1, T_2)$ is a random variable since D_1 and D_2 are also random variables.

The cost function, denoted by $h(T_1, T_2)$, is

$$h(T_1, T_2) = c_1(D_1 - T_1)^+ + c_2(D_2 - T_2 - (T_1 - D_1)^+)^+ + G_1T_1 + G_2T_2$$

The first term of $h(T_1, T_2)$ expresses the shortage cost of not satisfying demand for service 1 from resource 1. The second term of $h(T_1, T_2)$ expresses the shortage cost of not satisfying demand for service 2 given that resource 1 may also be used to satisfy some of that demand. In other words, if $T_1 > D_1$ then there is extra type 1 capacity, equal to $T_1 - D_1$, which can be used to satisfy type 2 demand when $D_2 > T_2$. Thus, the shortage cost associated with type 2 demand is incurred only if the quantity $D_2 - T_2 - (T_1 - D_1)^+$ is non-negative.

The net profit for a given demand realization (D_1, D_2) is $\pi(T_1, T_2) = g(T_1, T_2) - h(T_1, T_2)$, and thus the firm's problem is to

$$\max_{T_1,T_2}\{E_{D_1,D_2}[\pi(T_1,T_2)]\} = \max_{T_1,T_2}\{E_{D_1,D_2}[g(T_1,T_2) - h(T_1,T_2)]\} \quad (1)$$

Note that the term $\min[(D_2 - T_2)^+, (T_1 - D_1)^+]$ in $g(T_1, T_2)$ equals $D_2 - T_2$ if ($D_2 - T_2 < T_1 - D_1$ and $D_2 - T_2 > 0$) and equals $T_1 - D_1$ if ($T_1 - D_1 < D_2 - T_2$ and $T_1 - D_1 > 0$). Similarly, the term $(D_2 - T_2 - (T_1 - D_1)^+)^+$ in $h(T_1, T_2)$ equals $D_2 - T_2$ if ($T_1 - D_1 < 0$ and $D_2 - T_2 > 0$) and equals $D_2 - (T_1 - D_1) - T_2$ if ($T_1 - D_1 > 0$ and $D_2 - (T_1 - D_1) - T_2 > 0$).

Thus, we can express the expected profit $E_{D_1,D_2}[\pi(T_1,T_2)]$ function as follows:

$$\begin{aligned}
&E_{D_1,D_2}[\pi(T_1,T_2)]\\
&= \sum_{i=1}^{2}(p_i - v_i)\left(\int_0^{T_i} x f_i(x)\mathrm{d}x + \int_{T_i}^{\infty} T_i f_i(x)\mathrm{d}x\right)\\
&\quad + (p_2 - v_1)\left(\int_0^{T_1}\left[\int_{T_1+T_2-x}^{\infty}(T_1 - x)f_2(y)\mathrm{d}y\right]f_1(x)\mathrm{d}x\right.\\
&\quad \left. + \int_0^{T_1}\left[\int_{T_2}^{T_1+T_2-x}(y - T_2)f_2(y)\mathrm{d}y\right]f_1(x)\mathrm{d}x\right)\\
&\quad - c_1\int_{T_1}^{\infty}(x - T_1)f_1(x)\mathrm{d}x\\
&\quad - c_2\left(\int_0^{T_1}\left[\int_{T_1+T_2-x}^{\infty}(y + x - T_1 - T_2)f_2(y)\mathrm{d}y\right]f_1(x)\mathrm{d}x\right.\\
&\quad \left. + \int_{T_1}^{\infty}\left[\int_{T_2}^{\infty}(y - T_2)f_2(y)\mathrm{d}y\right]f_1(x)\mathrm{d}x\right) - G_1T_1 - G_2T_2
\end{aligned} \quad (2)$$

Proposition 1 *If the firm can make positive profits, the optimal capacity levels maximizing the expected profit in expression (2) are found by solving the following two equations:*

$$F_1(T_1) = \frac{(p_1 - v_1 + c_1) - G_1 + (p_2 - v_1 + c_2)P[0 < D_1 < T_1 \text{ and } D_2 > T_1 + T_2 - D_1]}{p_1 - v_1 + c_1} \quad (3)$$

$$F_2(T_2) = \frac{(p_2 - v_2 + c_2) - G_2 - (p_2 - v_1 + c_2)P[0 < D_1 < T_1 \text{ and } D_2 < T_1 + T_2 - D_1]}{(p_2 - v_2 + c_2) - (p_2 - v_1 + c_2)P[0 < D_1 < T_1]} \quad (4)$$

with respect to T_1 *and* T_2.

Proof. See Appendix.

Interpreting the optimal solution

By carefully looking at expressions (3) and (4), we can draw an analogy between the optimal two-recourse capacity levels and the optimal single-period newsvendor capacity level. Let us consider, for a moment, resource i, for $i = 1$, 2, in isolation, as in the single-period newsvendor model. Then, the optimal capacity for resource i is derived by setting the critical fractile ratio $(p_i - v_i + c_i - G_i)/(p_i - v_i + c_i)$ equal to $F_i(T_i)$, where $p_i - v_i + c_i - G_i$ is the marginal gain of selling one extra unit of capacity and G_i is the marginal loss of one unused unit of capacity. The optimal solutions to the two-resource, two-service capacity planning problem have similar interpretations. To better understand these interpretations, let us first explain the meaning of two probability expressions: (1) $P[0 < D_1 < T_1$ and $D_2 < T_1 + T_2 - D_1]$ is the probability of the event that there is excess type 1 capacity, which can be used to satisfy *all* of excess type 2 demand (ie no customers are left unserved), whereas (2) $P[0 < D_1 < T_1$ and $D_2 > T_1 + T_2 - D_1]$ is the probability of the event that there is excess type 1 capacity, which can *all* be used to satisfy excess type 2 demand, and there is extra type 2 demand left *unserved*. Keeping these meanings of the above probability expressions in mind, we now return to the two-resource problem, and express the marginal gain of one extra unit of type 1 capacity that can be used to satisfy demand as

$$\begin{aligned} Gain_1 = &(p_1 - v_1 + c_1) - G_1 + (p_2 - v_1 + c_2) \\ &\times P[0 < D_1 < T_1 \text{ and } D_2 > T_1 + T_2 - D_1] \end{aligned}$$

Observe that $(p_2 - v_1 + c_2)$ is multiplied by the probability of having extra type 1 capacity (since $T_1 > D_1$) and extra type 2 demand, which cannot be satisfied by resource 2 alone, or even by all the excess type 1 capacity (since $D_2 > T_1 + T_2 - D_1$). In other words, the marginal gain of one extra unit of capacity for resource 1 equals the corresponding single-resource marginal gain, plus the marginal gain associated with satisfying demand for service 2 from recourse 1, when excess type 1 capacity is surely needed.

The marginal loss of one extra unit of unsold type 1 capacity equals

$$\begin{aligned} Loss_1 = &G_1 - (p_2 - v_1 + c_2) \\ &\times P[0 < D_1 < T_1 \text{ and } D_2 > T_1 + T_2 - D_1] \end{aligned}$$

That is, the fixed cost G_1 is not considered a complete 'loss' if recourse 1 is used to satisfy demand for service 2. Thus, the optimal amount of capacity for resource 1 satisfies $F_1(T_1) = Gain_1/(Gain_1 + Loss_1)$, which has the same intuitive structure as the critical fractile of the single-resource newsvendor model.

Similarly, the marginal gain of one extra unit of type 2 capacity that can be used to satisfy demand equals

$$\begin{aligned} Gain_2 = &(p_2 - v_2 + c_2) - G_2 - (p_2 - v_1 + c_2) \\ &\times P[0 < D_1 < T_1 \quad \text{and} \quad D_2 < T_1 + T_2 - D_1] \end{aligned}$$

In other words, the marginal gain for resource 2 consists of the corresponding single-resource marginal gain, adjusted downwards by the marginal gain associated with satisfying *all* excess type 2 demand by excess *type* 1 capacity (rather than by type 2 capacity). The marginal loss of one extra unit of unsold type 2 capacity equals

$$\begin{aligned} Loss_2 = &G_2 - (p_2 - v_1 + c_2) \\ &\times P[0 < D_1 < T_1 \text{ and } D_2 > T_1 + T_2 - D_1] \end{aligned}$$

In other words, one extra unit of type 2 capacity implies a loss of G_2 only if the firm could *not* have used all of the excess type 1 capacity to satisfy all of the excess type 2 demand. Observe that $Gain_2 + Loss_2 = (p_2 - v_2 + c_2) - (p_2 - v_1 + c_2)\{P[0 < D_1 < T_1$ and $D_2 < T_1 + T_2 - D_1] + P[0 < D_1 < T_1$ and $D_2 > T_1 + T_2 - D_1]\} = (p_2 - v_2 + c_2) - (p_2 - v_1 + c_2)P[0 < D_1 < T_1]$, which equals the denominator of expression (4). Therefore, the optimal amount of capacity for resource 2 satisfies $F_2(T_2) = Gain_2/(Gain_2 + Loss_2)$, which is again similar to the single-resource newsvendor model.

Comparative statics analysis

In the previous subsection we enhanced our understanding of the structure of the optimal policy based on a marginal analysis interpretation of the optimal expressions, (3) and (4), as in the single-resource newsvendor model. In this subsection, as commonly done in the literature (eg Agnihothri *et al*[18]), we further our insights on the behavior of the optimal solution, by performing a comparative statics analysis with respect to the model input and demand parameters. Unfortunately, deriving closed-form solutions from expressions (3) and (4) is not feasible even for fairly simple demand distributions, such as the uniform. Therefore, we perform our analysis numerically. We assume independent normal distributions for the demands of the two service types, because the normal distribution is so often employed in practice, and is also relatively easy to handle mathematically. However (with one exception we report in the section Impact of changes to the demand parameters), our conclusions also hold for other symmetric two-parameter distributions (such as the Student's t-distribution). Table 1 summarizes the parameter values for the base case we considered in our numerical work.

Table 1 Sample data

	Input parameters				*Demand parameters*	
Resource	p_i	c_i	G_i	v_i	μ_i	σ_i
1	9	3	2	2	130	22
2	7	2	1	1	150	25

Table 2 Comparative statics

	T_1	T_2	S_{12}	π		T_1	T_2	S_{12}	π
p_1	+	−	+	+	p_2	+	+	−	+
c_1	+	−	+	−	c_2	+	+	−	−
v_1	−	+	−	−	v_2	+	−	+	−
G_1	−	+	−	−	G_2	+	−	+	−
μ_1	+	0	−	+	μ_2	0	+	0	0
σ_1	+	−	+	−	σ_2	+	+	+	−

Note that customer type 1 is willing to pay a higher price for resource type 1 but also expects to be compensated a higher value if her demand cannot be satisfied. Moreover, note that resource type 1 is more expensive as it has higher fixed and variable costs. In Table 2 we report the impact on T_1 and T_2 of changes to the input or demand parameters. We also evaluate and report the impact on the average percentage of type 1 capacity used to satisfy type 2 demand. We refer to this percentage as the *substitution rate* and denote it as S_{12}, where $S_{12} \equiv E[(T_1 - D_1)\mathbf{1}_{(D_1<T_1, D_2>T_1+T_2-D_1)} + (D_2 - T_2)\mathbf{1}_{(D_1<T_1, T_2<D_2<T_1+T_2-D_1)}]/T_1 = (1/T_1)(\int_0^{T_1}(T_1-x)f_1(x)(\int_{T_1+T_2-x}^{\infty} f_2(y)\mathrm{d}y)\mathrm{d}x + \int_0^{T_1} f_1(x)(\int_{T_2}^{T_1+T_2-x}(y-T_2)f_2(y)\mathrm{d}y)\mathrm{d}x)$. The results in Table 2 provide a number of interesting insights, which we discuss next.

Impact of changes to the input parameters. Table 2 illustrates that (with the exception of profits) the directions of change for the optimal decision variables T_1 and T_2 when considering small increases in price or shortage cost are the same (ie other than the profit column, the first two *rows* of Table 2 are the same). Similarly, the directions of change for T_1 and T_2 when considering small increases in the variable or fixed capacity cost are the same (see third and fourth rows of Table 2). Let us first consider increases on the fixed or variable capacity costs. As expected, when the fixed (G_1) or variable (v_1) capacity cost for resource 1 increases, the firm should optimally decrease type 1 capacity and increase type 2 capacity, while also decreasing the substitution rate S_{12}. Analogously, as v_2 and G_2 increase, both T_1 and S_{12} increase while T_2 decreases. In both cases, of course, expected profits decrease.

Turning to changes in price or shortage cost, Table 2 shows that when all other parameters remain the same, and the price p_1 of resource 1 increases (implying that type 1 customers are now more valuable), the firm should optimally increase its type 1 capacity. Interestingly, when p_1 increases, not only should the optimal type 1 capacity increase but type 2 capacity should also optimally decrease. To explain this last, somewhat unexpected result, recall that even though price has increased, the demand distribution remains the same. Therefore, when type 1 capacity increases, we may have excess type 1 capacity more often, which can then be used to satisfy demand for service type 2. To put it differently, since the demand distribution remains constant, random demand realizations would not affect the mean demand, but would now be satisfied from higher levels of type 1 capacity, implying a higher frequency of realizing leftover type 1 capacity. Therefore, as p_1 increases, the firm should not only increase T_1 but at the same time *decrease* T_2, thereby increasing the substitution rate S_{12}. To give one example of the magnitude of these optimal adjustments, we report the impact of a 50% increase in price or variable capacity cost v_1 from the base-case scenario we considered in Table 1: a 50% increase in p_1 implies a 4.6% increase in T_1, a 2.0% decrease in T_2, a 56% increase in profits, and a 35.3% increase in the substitution rate S_{12}, whereas a 50% increase in v_1 implies a 2.8% decrease in T_1, a 4.1% increase in T_2, an 11.3% decrease in profits, and a 34.7% decrease in the substitution rate S_{12}.

The comparative statics results for increases in p_2 are somewhat different from those for increases in p_1, because resource 2 cannot be used to satisfy type 1 demand. From Table 2 we see that as p_2 increases, and service 2 becomes more valuable, both types 1 and 2 capacities should optimally increase to ensure better service for type 2 customers. Interestingly, we also find that as p_2 increases, T_2 increases enough to make the optimal substitution rate S_{12} decrease.

We can now apply arguments similar to the ones we employed for increases in price to explain what happens to the decision variables as shortage cost increases. As the shortage cost c_1 of resource type 1 increases, it becomes more important to the firm to satisfy demand for service type 1, implying, as with increases in p_1, that capacity for resource type 1 should optimally increase, capacity for resource type 2 should decrease (because, as before, there may be extra type 1 capacity to use for type 2 demand), and the substitution percentage should increase. Similarly, as c_2 increases, both T_1 and T_2 increase while S_{12} decreases.

Impact of changes to the demand parameters. As we see from Table 2, changes in the mean demand of one of the resources do not affect the optimal capacity choice for the other resource, whereas changes in the standard deviation of one of the resources affect the capacity choices for *both* resources. Even though the first result is not necessarily counterintuitive, it is somewhat surprising. Before looking at Table 2, for example, we may have expected that when μ_2 increases, both T_1 and T_2 should increase, or that when μ_1

increases, T_2 may decrease. However, we find that an increase in μ_2 only affects T_2, but leaves both T_1 and S_{12} unchanged. In contrast, as σ_2 increases, the optimal values of both capacity types are affected; in fact, when σ_2 increases, the optimal values for all three variables (T_1, T_2, and S_{12}) should optimally increase. Similarly, when the standard deviation σ_1 of type 1 demand increases (while the mean remains the same) both capacities are affected, with T_1 and S_{12} increasing, and T_2 decreasing. To verify whether these results are specific to the normal distribution, we compared them with the corresponding results for t-distributed, independent demands. We found that, as with the normal distribution, increases in μ_2 do *not* affect T_1 and S_{12}, and that changes to the standard deviation of either resource do affect the optimal capacity choices of both recourse types. Unlike the normal distribution case, however, we did observe a slight decrease in T_2 when μ_1 increased. We, therefore, conclude that the impact of changes on the mean demand of one resource is distribution specific, and thus may or may not affect the optimal capacity choice of the other resource.

In summary, our comparative statics analysis provided us with a better understanding of how the optimal capacities should be adjusted when the problem parameters change. One interesting finding, for example, is that when type 1 capacity increases (as a reaction to type 1 demand being more profitable, or type 2 demand being more variable), type 2 capacity should also be adjusted downwards. This finding is of course linked with our independent demands assumption. If demands are positively correlated, for example, then we may see that increases in p_1, or σ_1, may not necessarily imply a decrease in T_2. Our results also suggest the importance of applying accurate estimates for the demand distributions. Particularly, we found that changes in the mean or standard deviation of demand may imply notable differences in expected profits; for example, a 10% increase in the mean of type 1(2) demand from the base level of 130(150) we considered, implies a 5%(3%) increase in profits. Therefore, in the next section we discuss how one could employ an improved estimation method for the demand parameters before calculating the optimal capacity choices.

Bayesian updating for the two-resource capacity planning model

So far we focused on maximizing the expected profit function shown in expression (1). We provided analytic solutions for this problem via Equations (3) and (4), which hold for *any* two independent and continuous probability distribution functions. In practice, demands are often assumed to be normal with a given mean and variance. For example, if we denote by $z_i(x)$ the probability density function of a normal distribution with mean μ_i and standard deviation σ_i, that is, $z_i(x) = 1/(\sqrt{2\pi}\sigma_i)\mathrm{e}^{-(1/2)[(x-\mu_i)/\sigma_i]^2}$, $i=1$, 2, and by $Z_i(\cdot)$ the corresponding cumulative distribution function, then the optimal capacity levels are determined by the following relationships:

$$Z_1(T_1) = \frac{(p_1 - v_1 + c_1) - G_1 + (p_2 - v_1 + c_2)\int_0^{T_1}[1 - Z_2(T_1 + T_2 - x)]z_1(x)\mathrm{d}x}{p_1 - v_1 + c_1} \tag{5}$$

$$Z_2(T_2) = \frac{(p_2 - v_2 + c_2) - G_2 - (p_2 - v_1 + c_2)\int_0^{T_1} Z_2(T_1 + T_2 - x)]z_1(x)\mathrm{d}x}{(p_2 - v_2 + c_2) - (p_2 - v_1 + c_2)[Z_1(T_1) - Z_1(0)]} \tag{6}$$

In order to solve these equations we must first estimate the means and variances of the normal distributions. Obviously, the estimation of the unknown parameters is critical when maximizing the expected profit function. Typically, the sample mean $\dot{X}_i$ and the sample variance $S_i^2/(n_i-1)$ (where S_i^2 denotes the sum of squared errors and n_i the sample size) are used to estimate the true demand parameters. However, this traditional estimation approach may not be effective under certain circumstances. For example, traditional estimation would provide only a single value for the maximum expected profit, whereas a Bayesian approach would provide a sample *distribution* of maximum profits so that one can assess the effect of demand estimation on the potential maximum profit. In addition, it is possible that demands are not stationary and have some or all of their parameters varying with time. Are there better estimates for the parameters of the demand distributions that should be used in such cases? In particular, should managers strive to incorporate prior information on the demand parameters when addressing the capacity planning problem? To investigate such questions, we next present how Bayesian theory can be employed for two different demand scenarios.

Normally distributed demands case

Assume that demands are normally distributed with mean μ_i and standard deviation σ_i, that is, that $D_i \sim N(\mu_i, \sigma_i)$, $i=1, 2$. Instead of using the sample mean and sample variance to estimate the population parameters, we could employ Bayesian estimates incorporating existing prior information. A common class of priors for the mean and variance of normal distributions (see Fong and Ord[19]) is

$$\begin{gathered}\mu_i \sim t(2a_i, m_i, \gamma_i b_i/a_i) \text{ and } \sigma_i^2 \sim IG(a_i, b_i),\\ i = 1, 2, \text{ all independent}\end{gathered} \tag{7}$$

where t denotes the Student's t-distribution, IG denotes the inverse gamma distribution (so that $u_i = 1/\sigma_i^2$ is Gamma-distributed with density $f(u_i) = u_i^{a_i-1} b_i^{a_i} \exp(-u_i b_i)/\Gamma(a_i)$, and a_i, b_i, γ_i, and m_i are constants; these constants are specified based on subjective beliefs/expert opinions regarding the demand parameters, and *not necessarily* based on observed data. For example, referring back to our hotel expansion setting, if the hotel manager knows of a change in the population demographics that would increase demand by a certain percentage in the next few years, he would incorporate this information into his beliefs for the prior distribution of the demand parameters, and specify a_i, b_i, γ_i, and m_i accordingly (see the section Numerical results for a specific example of how to determine these four constants). Even though we employ *conjugate* priors in this paper, in practice more realistic, non-conjugate priors can also be handled by using advanced computational approaches such as Markov Chain Monte Carlo methods (see Gilks *et al*[20]).

Note that the prior of (μ_i, σ_i^2) in expression (7) is equivalent to the normal-inverse gamma conjugate prior given by $\mu_i|\sigma_i^2 \sim N(m_i, \sigma_i\sqrt{\gamma_i})$, and $\sigma_i^2 \sim IG(a_i, b_i)$. Thus, the conditional density of μ_i given (σ_i^2, x) is $N((\gamma_i n_i \bar{X}_i + m_i)/(1+\gamma_i n_i), \sigma_i\sqrt{\gamma_i/(1+\gamma_i n_i)})$, and the marginal posterior density of σ_i^2 is $IG(a_i + n_i/2,\ b_i + S_i^2/2 + (\bar{X}_i - m_i)^2/2(\gamma_i + 1/n_i))$, which implies that the Bayesian estimates for μ_i and σ_i^2 are

$$\hat{\mu}_i = \frac{(\gamma_i n_i \bar{X}_i + m_i)}{(1+\gamma_i n_i)} \tag{8}$$

$$\hat{\sigma}_i^2 = \frac{b_i + S_i^2/2 + (\bar{X}_i - m_i)^2/(2(\gamma_i + 1/n_i))}{a_i + n_i/2 - 1} \tag{9}$$

Note that when $b_i = 0$, $a_i = \frac{1}{2}$, and $\gamma_i \to \infty$, then $\hat{\mu}_i \to \bar{X}_i$ and $\hat{\sigma}_i \to S_i^2/(n_i - 1)$. In other words, the traditional sample estimates of the mean and variance are indeed Bayesian estimates with respect to a non-informative prior.

Thus, we are interested in comparing expected profits when using parameter estimates with informative priors ($\hat{\mu}_i$, $\hat{\sigma}_i$), or non-informative priors ($\bar{X}_i$, $S_i^2/(n_i - 1)$). To summarize, assuming that the *true* demand distribution is normal (we refer to this case as the normally distributed demands (NDD) case), we want to compare the following two estimation scenarios when maximizing expected profit:

Scenario NDD-NIP: Substitute simple sample estimates (ie Bayesian estimates with a non-informative prior (NIP)) for the mean and variance of the normal demand distributions in the expected profit expression (2). In other words, substitute $\bar{X}_i$ and $S_i/\sqrt{(n_i - 1)}$ for μ_i and σ_i, respectively, in Equations (5) and (6).

Scenario NDD-IP: Substitute the Bayes estimates of the informative priors (IP) for the mean and variance (see expression (7)) in the expected profit expression (2). In other words, substitute $\hat{\mu}_i$ and $\hat{\sigma}_i$ and for μ_i and σ_i, respectively, in Equations (5) and (6). We present numerical results for these two demand scenarios in the section Numerical results.

t-Distributed demands case

In the previous subsection we assumed that demands follow a *stationary* normal distribution. In certain cases, however, it is possible that the demand distributions are non-stationary. We consider a special non-stationary demands case here, in which the *true* demands are *t*-distributed, $D_i \sim t(2a_i, \mu_i, 2b_i)$, $i = 1, 2$; this case occurs when the means of normally distributed demands remain constant, but the variances are assumed to vary with time and follow an inverse gamma distribution. In other words, we assume that $D_{ij}|\sigma_{ij} \sim N(m_i, \sigma_{ij})$ and $\sigma_{ij}^2 \sim IG(a_i, b_i)$, where the subscript $j = 1, \ldots, n_i$ denotes the current time period.

Given this class of *t*-distributed true demands, and using the same class of conjugate priors that we employed in the previous subsection, the predictive distribution of *future* demand conditional on existing data X for service type i (for any time period $j > n_i$) is $D_i|X \sim t(2a_i + n_i, \tilde{\mu}_i, \tilde{\sigma}_i^2)$, where the location parameter $\tilde{\mu}_i$, and the scale parameter $\tilde{\sigma}_i^2$ are given as follows:

$$\tilde{\mu}_i = \frac{(\gamma_i n_i \bar{X}_i + m_i)}{(1+\gamma_i n_i)} \tag{10}$$

$$\tilde{\sigma}_i^2 = \frac{[(1+\gamma_i n_i)(2b_i + S_i^2) + n_i(\bar{X}_i - m_i)^2][1 + \gamma_i + \gamma_i n_i]}{(2a_i + n_i)(1+\gamma_i n_i)^2} \tag{11}$$

We refer to this case as the *t*-distributed demands (TDD) case, and so we want to compare the following two estimation scenarios when maximizing expected profit:

Scenario TDD-NIP: Use a normal distribution to approximate the distribution of each demand and substitute sample estimates for the mean and variance of the normal demand distributions in expression (2). In other words, substitute $\bar{X}_i$ and $S_i/\sqrt{(n_i - 1)}$ for μ_i and σ_i, respectively, in Equations (5) and (6).

Scenario TDD-IP: Use the predictive distributions $t(2a_i + n_i,\ \tilde{\mu}_i,\ \tilde{\sigma}_i^2)$, $t = 1, 2$. In other words, use this *t*-distribution with location and scale estimates $\tilde{\mu}_i$ and $\tilde{\sigma}_i$ in Equations (3) and (4).

For both the NDD and the TDD cases, we next compare the use of informative and non-informative priors by (i) simulating random samples from the true distributions, (ii) deriving optimal capacity decisions based on these samples, and (iii) comparing expected profits for each scenario.

Numerical results

Our main motivation in numerically comparing the various demand scenarios described in the previous section is to verify and quantify the value of using informative priors on the demand parameters for our capacity planning problem. Our purpose is not to perform a large simulation study, but rather to illustrate that good (not perfect) prior information can be valuable and can improve performance when making capacity decisions under uncertainty. We also give a specific example of how the parameters of the common class of priors introduced in the previous section can be calculated.

Method description

A detailed pseudoalgorithm for the method we use to compare the estimation scenarios is shown in the appendix. We should emphasize that the Bayesian approach we described in the previous section is based on the joint posterior distribution of μ_i and σ_i for both the NDD and the TDD cases. Thus, one can generate samples from the posterior distribution of $\{\mu_i, \sigma_i\}$, and calculate the corresponding optimal values for the capacities $\{T_1, T_2\}$ for each sample; therefore, one can generate a sample distribution for $\{T_1, T_2\}$ and a corresponding sample distribution for maximum expected profits. This approach can provide the decision maker with much more information on profits than when using classical estimation, where only one value of maximum expected profit can be generated. In our numerical results, however, for simplicity, we only consider the posterior mean and variance and evaluate the corresponding expected profit (see the algorithm described in the appendix).

To give a specific numerical example, let us first consider the normally distributed demands case, and let us assume that the *true* mean demand for service 1 is $\mu_1 = 130$ and the *true* variance is $\sigma_1^2 = 22^2 = 484$. Let us consider a decision maker who *believes* that the average demand is $E[\mu_1] = \mu_1 = 130$, give or take 20 (ie $Stdev[\mu_1] = 20$), and that the mean variance is $E[\sigma_1^2] = \sigma_1^2 = 484$, give or take 225 (ie $Stdev[\sigma_1^2] = 225$). In other words, the decision maker has *good* subjective information (because her beliefs regarding the average demand $E[\mu_1]$ and average variance $E[\sigma_1^2]$ are accurate), but *not perfect* information (because she also assigns prior demand distributions around the values of $E[\mu_1]$ and $E[\sigma_1^2]$).

Given these prior beliefs, the values for the parameters of the conjugate priors a_1, b_1, γ_1, and m_1 can be derived as follows. We have assumed that $\sigma_1^2 \sim IG(a_1, b_1)$; thus, the mean $E[\sigma_1^2] = b_1/(a_1 - 1)$ and the standard deviation $Stdev[\sigma_1^2] = (b_1/(a_1 - 1))\sqrt{(a_1 - 2)}$. Solving these two equations for a_1 and b_1 yields $a_1 = 2 + (E[\sigma_1^2]/Stdev[\sigma_1^2])^2$ and $b_1 = (a_1 - 1)E[\sigma_1^2]$. Substituting $E[\sigma_1^2] = 484$ and $Stdev[\sigma_1^2] = 225$ in these expressions we get $a_1 = 6.63$ and $b_1 = 2724$.

To calculate the value of γ_1 as a function of a_1 and b_1, we note that if a random variable is t-distributed with k degrees of freedom, location parameter m, and scale parameter σ^2, $t(k, m, \sigma^2)$, then its standard deviation is $\sqrt{[k/(k-2)]}\sigma$. Thus, since we assume that $\mu_1 \sim t(2a_1, m_1, \gamma_1 b_1/a_1)$, see expression (7), we have that $Stdev[\mu_1] = \sqrt{2a_1/(2a_1 - 2)} \times \sqrt{\gamma_1 b_1/a_1}$ which implies that $\gamma_1 = (Stdev[\mu_1])^2(a_1 - 1)/b_1)$. Therefore, substituting the values of a_1 and b_1 in the expression for γ_1 yields $\gamma_1 = 0.83$. Finally, the location parameter m_1 is set equal to $E[\mu_1]$ so $m_1 = 130$. Summarizing, given the prior beliefs regarding the average mean and average variance, the demand parameters for service type 1, $(a_1, b_1, \gamma_1, m_1) = (6.63, 2724, 0.83, 130)$. Given these parameter values, let us assume that a random sample of size $n_1 = 5$ is drawn from the normal distribution $N(\mu_1, \sigma_1) = N(130, 22)$, which produced the sample estimates $(\bar{X}_1, S_1/\sqrt{(n_1 - 1)}) = (125, 15)$. Then from Equations (8) and (9), the Bayesian estimates are $(\hat{\mu}_1, \hat{\sigma}_1) = (125.97, 19.8)$.

Similarly, if $\mu_2 = 150$, $\sigma_2^2 = 25^2 = 625$, and subjective information suggests that $E[\mu_2] = 150$, $Stdev[\mu_2] = 30$, $E[\sigma_2^2] = 625$, and $Stdev[\sigma_2^2] = 225$, the parameters $(a_2, b_2, \gamma_2, m_2) = (9.72, 5447, 1.44, 150)$, and if a random sample of size $n_2 = 5$ from the normal distribution $N(\mu_2, \sigma_2) = N(150, 25)$ produces the estimates $(\bar{X}_2, S_2/\sqrt{(n_2 - 1)}) = (145, 35)$, then $(\hat{\mu}_2, \hat{\sigma}_2) = (145.61, 26.55)$. Based on these estimates, we compare the two scenarios for the NDD case as summarized in Table 3.

Note that we assumed a small sample size to reflect cases for which there is only a small amount of available demand data, perhaps due to changing business conditions, or due to the services offered being relatively new. Moreover, for both service types, the coefficient of variation implied by the chosen parameter values is low so that the probability of negative demand is very small.

Let us now consider an example for the TDD case. For comparison, suppose that the *same* subjective beliefs presented above for the NDD case apply here as well, leading to the same parameter values of $(a_1, b_1, \gamma_1, m_1) = (6.63, 2724, 0.83, 130)$ and $(a_2, b_2, \gamma_2, m_2) = (9.72, 5447, 1.44, 150)$. Moreover, suppose that random samples of sizes $n_1 = n_2 = 5$ were drawn from the true t-distributions $t(2a_i, \mu_i, 2b_i)$, $i = 1, 2$, and produced the *same* sample estimates $(\bar{X}_1, S_1/\sqrt{(n_1 - 1)}, \bar{X}_2, S_2/\sqrt{(n_2 - 1)}) = (125, 15, 145, 35)$. Then from Equations (10) and (11), the Bayes estimates for the demand parameters are $(\tilde{\mu}_1, \tilde{\sigma}_1) = (125.97, 20.13)$ for service 1 and $(\tilde{\mu}_2, \tilde{\sigma}_2) = (145.61, 27.58)$ for service 2. Based on these estimates, we compare the two scenarios for the TDD case as summarized in Table 4.

These numerical examples illustrate that neither of the estimation scenarios (NIP or IP) has perfect information on the demand processes. Both rely on sample data, but in addition, the informative prior estimates incorporate subjective beliefs regarding the prior distribution of the mean and variance. Based on these beliefs, the parameters of the conjugate priors are estimated. Thus, it is important to emphasize that our algorithm for comparing the two estimation scenarios is not designed to favor one *versus* the other. At the same time, however, we chose the subjective beliefs regarding the mean and variance of the unknown

Table 3 Demand scenarios for the NDD case

Scenario	*Estimated distribution*
NDD-NIP	Service 1: $N(125, 15)$ Service 2: $N(145, 35)$
NDD-IP	Service 1: $N(125.97, 19.8)$ Service 2: $N(145.61, 26.55)$

Table 4 Demand scenarios for the TDD case

Scenario	*Estimated distribution*
TDD-NIP	Service 1: $N(125, 15)$ Service 2: $N(145, 35)$
TDD-IP	Service 1: $t(18.25, 125.97, 20.13^2)$ Service 2: $t(24.43, 145.61, 27.58^2)$

parameters to be reasonably close to the true values because we want to know how valuable good information can be. If the decision makers have no specific information regarding their stochastic processes (somewhat unlikely in most business settings), then they might simply apply the classical estimates.

Results

We employ here the same input parameter values we reported in the section Comparative statics analysis (see Table 1). For each demand case (normally distributed and stationary, NDD or t-distributed and non-stationary, TDD) and each estimation scenario (use of informative prior, IP, *versus* use of non-informative prior, NIP) we perform 30 simulations runs. Table 5 summarizes the results for the base case described in Table 1. It reports the average, maximum, and minimum percentage differences when compared with the true expected profit for each scenario. For example, the 'Average' of the NDD-NIP scenario is calculated as the ratio of the difference between the true expected profit and the average expected profit (based on 30 samples of the NDD-NIP case), over the true expected profit. Similarly, the 'max' of the NDD-IP scenario is calculated as the ratio of the difference between the true expected profit and the maximum expected profit (based on 30 samples of the NDD-NIP case), over the true expected profit.

We have also run simulations with larger sample sizes (eg $n_i = 10$ or 15, instead of 5), larger coefficient of variations (eg 25 or 20% instead of about 16%), and various values for parameters such as variable and fixed costs. The results remain qualitatively the same in that the use of informative priors reduces the difference with the true expected profit significantly. Thus, our numerical results suggest that it is desirable to use good prior information, if available.

Table 5 Percentage difference between true profits and expected profits

	NDD-NIP	*NDD-IP*	*TDD-NIP*	*TDD-IP*
Average	1.22%	0.09%	1.97%	0.21%
Max	4.31%	0.36%	5.73%	0.76%
Min	0.01%	0.00%	0.27%	0.02%

Summary and conclusions

We investigate both the modeling and the demand-estimation aspects of a capacity planning problem, in which two available resources provide two services. Demand for one of the services can be satisfied by both types of capacity, whereas demand for the other service can only be provided by a specialized resource (ie downgrades are not allowed). We provide easy-to-interpret, analytic solutions that resemble the familiar critical fractile ratio of the single-resource newsvendor model, but also account for the possibility of substituting one of the resources with the other. We also enhance our understanding of how changes to demand or input parameters impact the optimal capacity choices and profits by performing a Comparative statics analysis. We then turn our attention to understanding the impact of estimating the unknown parameters of the demand distributions. The traditional estimation approach is to use sample mean and sample variance, respectively, to approximate the unknown population mean and variance of a demand distribution. It is often the case, however, that useful prior information is available that could improve the capacity planning process significantly. Thus, we describe how prior information can be used to provide the optimal capacity levels and a distribution of maximum profits based on a Bayesian methodology, and we quantify the value of using such prior information.

In our numerical experiments, we consider the case of independent NDDs and the case of independent TDDs, with the latter case corresponding to non-stationary demands. For each case, we compare expected profits under scenarios that apply Bayesian estimates or standard sample estimates. Our simulations show that the Bayesian estimates yield expected profits that are consistently closer to the true expected profit and thus demonstrate that informative priors can improve expected profits significantly.

One future direction related to this work would be to consider the correlated demands case. In this case, it would be interesting to identify the circumstances under which the benefits of prior information are more pronounced.

Acknowledgments—The authors would like to thank the editor and two anonymous referees for their helpful comments.

Appendix

Proof of Proposition 1. It is easy to show that the objective function in expression (2) is concave in $\{T_1, T_2\}$ (see for example Bassok *et al*[7] and Parlar and Goyal[10]). Thus, if the firm makes positive profits, the first-order conditions are both necessary and sufficient. Taking the partial derivative of the profit expression (2) with respect to T_1 yields

$$\frac{\partial[\text{expression}(2)]}{\partial T_1} = (p_1 - v_1)(1 - F_1(T_1)) - c_1(F_1(T_1) - 1) - G_1 + (p_2 - v_1)\int_0^{T_1}(1 - F_2(T_1 + T_2 - x))f_1(x)\mathrm{d}x + c_2\int_0^{T_1}(1 - F_2(T_1 + T_2 - x))f_1(x)\mathrm{d}x$$

Thus, setting the partial with respect to T_1 equal to zero we get that

$$(p_1 - v_1 + c_1 - G_1) - (p_1 - v_1 + c_1)F_1(T_1) + (p_2 - v_1 + c_2)\int_0^{T_1}(1 - F_2(T_1 + T_2 - x))f_1(x)\mathrm{d}x = 0$$

which implies Equation (3). Similarly, taking the partial derivative of the profit expression (2) with respect to T_2 yields

$$\frac{\partial[\text{expression}(2)]}{\partial T_2} = (p_2 - v_2)(1 - F_2(T_2)) - G_2 - (p_2 - v_1)\int_0^{T_1}[F_2(T_1 + T_2 - x) - F_2(T_2)]f_1(x)\mathrm{d}x + c_2\left\{\int_0^{T_1}[1 - F_2(T_1 + T_2 - x)]f_1(x)\mathrm{d}x + \int_{T_1}^{\infty}[1 - F_2(T_2)]f_1(x)\mathrm{d}x\right\}$$

Observing that the last term within the braces of the previous equation, $\int_{T_1}^{\infty}[1 - F_2(T_2)]f_1(x)\mathrm{d}x$, equals $(1 - F_2(T_2)) - \int_0^{T_1}(1 - F_2(T_2))f_1(x)\mathrm{d}x$, we can simplify the first derivative with respect to T_2 as

$$\frac{\partial[\text{expression}(2)]}{\partial T_2} = (p_2 - v_2 + c_2)(1 - F_2(T_2)) - G_2 - (p_2 - v_1 + c_2)\int_0^{T_1}[F_2(T_1 + T_2 - x) - F_2(T_2)]f_1(x)\mathrm{d}x$$

and setting the partial with respect to T_2 equal to zero we get that

$$(p_2 - v_2 + c_2)(1 - F_2(T_2)) - G_2 + (p_2 - v_1 + c_2)F_2(T_2)\int_0^{T_1}f_1(x)\mathrm{d}x - (p_2 - v_1 + c_2)\int_0^{T_1}F_2(T_1 + T_2 - x)f_1(x)\mathrm{d}x = 0$$

which implies Equation (4). □

Algorithm for calculating our numerical results

(1) Choose a demand scenario R to be either the NDD or the TDD (ie $R \in \{NDD, TDD\}$). Choose values for the true μ_i and σ_i, $i = 1, 2$. Denote the true demand distribution for service i by $f_i^R(.)$, $i = 1, 2$ (in other words, $f_i^R(.)$ has mean μ_i and variance σ_i^2). For every demand scenario R, repeat steps (2) and (3).

(2) Choose an estimation scenario $S(R)$ for a given true demand distribution R, where $S(R) \in \{NIP, IP\}$. For every estimation scenario $S(R)$, repeat steps (3a)–(3f) several times.

(3a) Generate a random sample of size n_i following $f_i^R(.)$, $i = 1, 2$.

(3b) Derive the estimated demand distributions $f_i^{E(S(R))}(.)$, $i = 1, 2$, for approach $S(R)$, denoted for simplicity by $f_i^E(.)$ (in other words, the parameters of $f_i^E(.)$ are $(\hat{\mu}_i, \hat{\sigma}_i)$ or $(\tilde{\mu}_i, \tilde{\sigma}_i)$ depending on $S(R)$).

(3c) Derive optimal capacity levels, denoted by $\hat{T}_1$ and $\hat{T}_2$, that maximize expected profit for scenario $S(R)$ given the estimated demand distributions $f_i^E(.)$

$$\pi(\hat{T}_1, \hat{T}_2, f_1^E(.), f_2^E(.)) = \max_{T_1, T_2}\left\{E_{f_1^E(.), f_2^E(.)}[\pi(T_1, T_2)]\right\}.$$

(3d) Evaluate the true expected profit $\pi(\hat{T}_1, \hat{T}_2, f_1^R(.), f_2^R(.))$ given capacity levels $\hat{T}_1$ and $\hat{T}_2$ and the true demand distributions $f_i^R(.)$, $i = 1, 2$.

(3e) Derive the true optimal capacity levels T_1^R and T_2^R that maximize expected profit for scenario $S(R)$ given the true demand distributions $f_i^R(.)$:

$$\pi(T_1^R, T_2^R, f_1^R(.), f_2^R(.)) = \max_{T_1, T_2}\left\{E_{f_1^R(.), f_2^R(.)}[\pi(T_1, T_2)]\right\}.$$

(3f) Evaluate the difference between the true expected profit for approximately optimal capacities, $\pi(\hat{T}_1, \hat{T}_2, f_1^R(.), f_2^R(.))$, and the true expected profit, $\pi(T_1^R, T_2^R, f_1^R(.), f_2^R(.))$, given the optimal capacities. When iteration limit is reached, go to step (2).

References

1 French S and Smith JQ (eds) (1997). *The Practice of Bayesian Analysis*, J Wiley and Sons: New York.

2 Fong DKH and Bolton GE (1997). Analyzing ultimatum bargaining: a bayesian approach to the comparison of two potency curves under shape constraints. *J Bus Econ Stat* **15**: 335–344.

3 Pammer SE, Fong DKH and Arnold SF (2000). Forecasting the penetration of a new product: a Bayesian approach. *J Bus Econ Stat* **18**: 428–435.

4 Fong DKH, Pammer SE, Arnold SF and Bolton GE (2002). Reanalyzing ultimatum bargaining: comparing nondecreasing curves without shape constraints. *J Bus Econ Stat* **20**: 423–430.

5 DeSabo WS, Kim Y and Fong DKH (1999). A bayesian multidimensional scaling procedure for the spatial analysis of revealed choice data. *J Econom* **89**: 79–108.

6 Sethi AK and Sethi SP (1990). Flexibility in manufacturing: A Survey. *Int J Flexible Manuf* **2**(4): 289–328.
7 Bassok Y, Anupindi R and Akella R (1999). Single-period multiproduct inventory models with substitution. *Opns Res* **47**(4): 632–642.
8 Harrison JM and Van Mieghem JA (1999). *Multi-resource investment strategies: operational hedging under demand uncertainty*. Working Paper, Stanford University.
9 Netessine S, Dobson G and Shumsky R (1999). *Capacity investment decisions for flexible resources under uncertainty*, Working Paper, University of Rochester.
10 Parlar M and Goyal SK (1984). Optimal ordering decisions for two substitutable products with stochastic demands. *Opsearch* **21**(1): 1–15.
11 Hill RM (1997). Applying bayesian methodology with a uniform prior to the single period inventory model. *Eur J Opl Res* **98**: 555–562.
12 Azouri KS (1985). Bayes solution to dynamic inventory model under unknown demand distribution. *Mngt Sci* **31**(9): 1150–1160.
13 Bradford JW and Sugrue PK (1990). A bayesian approach to the two-period style-goods inventory problem with single replenishment and heterogeneous Poisson demands. *J Opl Res Soc* **41**(3): 211–218.
14 McGill JI and Van Ryzin GJ (1999). Revenue management: research overview and prospects. *Transportation Science* **33**(2): 233–256.
15 Bodily SE and Weatherford LR (1995). Perishable-asset revenue management: generic and multiple-price yield management with diversion. *Omega* **23**(3): 173–185.
16 Brumelle SL *et al* (1990). Allocation of airline seats between stochastically dependent demands. *Transport Sci* **24**(3): 183–192.
17 Ladany SP (1977). Bayesian dynamic operating rules for optimal hotel reservations. *Z Opns Res* **21**: 165–176.
18 Agnihothri S, Lee JS and Kim J (2002). Lot sizing with random yields and tardiness costs. *Comput Opns Res* **27**: 437–459.
19 Fong DKH and Ord KJ (1993). Estimating moments of the effective lead time. *J Opl Res Soc* **44**(3): 247–252.
20 Gilks WR, Richardson S and Spiegelhalten DJ (eds) (1996). *Markov Chain Monte Carlo in Practice*, Chapman & Hall: London.

Received August 2001;
accepted April 2003 after one revision