Investment in Production Resource Flexibility: An Empirical Investigation of Methods for Planning under Uncertainty

Elena Katok,[1] William Tarantino,[2] Terry P. Harrison[1]

[1] *MS&IS Department, Penn State University, University Park, Pennsylvania 16802*

[2] *Center for Army Analysis, Fort Belvoir, Virginia*

Received 16 February 2000; revised 2 May 2002; accepted 2 May 2002

DOI 10.1002/nav.10059

Abstract: We examine several methods for evaluating resource acquisition decisions under uncertainty. Traditional methods may underestimate equipment benefit when part of this benefit comes from decision flexibility. We develop a new, practical method for resource planning under uncertainty, and show that this approach is more accurate than several commonly used methods. We successfully applied our approach to an investment problem faced by a major firm in the aviation information industry. Our recommendations were accepted and resulted in estimated annual savings in excess of $1 million (US). © 2003 Wiley Periodicals, Inc. Naval Research Logistics 50: 105–129, 2003.

Keywords: manufacturing flexibility; stochastic programming; sampling

1. INTRODUCTION

In recent years, many firms have found it increasingly important to invest substantially in technology to maintain a competitive edge. Technological improvements often require superior production methods, and some firms find themselves frequently evaluating opportunities for investments in new production resources. These decisions can easily become crucial to survival in a competitive market place. While essential to the well-being of firms, production investment decisions are extremely difficult because they involve planning under uncertainty. For example, when a new production resource provides manufacturing flexibility, the benefit of this flexibility can be easily underestimated. As Jordan and Graves [17] point out: “[I]n capacity and flexibility planning, investment costs for flexible operations are typically quantified; however, it is less common to quantify the benefits because demand uncertainty is not explicitly considered by the planners. Since flexibility is expensive, this typically results in decisions not to invest in it.”

The benefits of a new production resource can emerge in three ways:

1. Lower cost due to superior performance

Correspondence to: E. Katok

2. Increased capacity
3. Increased decision flexibility.

The first two sources of benefit are fairly intuitive: Cost savings may result if a new resource provides a more efficient production process or introduces a new dedicated process. If a new resource is added to the current production system at a particular stage, capacity at that stage may increase. If that stage previously formed a bottleneck, the throughput of the entire system increases (Goldratt and Cox [13]), potentially yielding cost savings. The third source of benefit comes from increased decision flexibility (Benjaafar, Morin, and Talavage [2]). Decision flexibility is the ability to postpone decisions until more information is obtained. When a new production resource is added to the current system, it can increase decision flexibility by either providing additional capacity where it is needed, or by providing an additional routing for a part. To correctly estimate the impact of a flexible resource, a model must include all three sources of benefit.

With this study we contribute to manufacturing flexibility planning research in three ways:

- We describe a new and better method that accurately accounts for all three sources of flexibility benefit, and that is practical enough to be used for large and complex real-world problems.
- We implement and use the new method to help with a real flexibility-planning decision faced by Jeppesen Sanderson, Inc. (Jeppesen), a manufacturing company, and generate annual savings in excess of $1 million.
- We clearly demonstrate, using the Jeppesen investment problem as a case study, how other commonly used methods can consistently underestimate the benefit of flexibility.

This paper is organized as follows. First, we describe the problem and summarize relevant literature in Section 2. In Section 3, we develop a formal mathematical model for flexibility planning and show analytically that some commonly used methods generally undervalue flexibility. We also describe our new sampling-based optimization algorithm that assesses the benefit of manufacturing flexibility more accurately than existing methods. In Section 4, we describe the flexibility-planning problem faced by Jeppesen and the application of our method to this problem. In Section 5 we demonstrate that alternative methods significantly underestimate the benefit of flexibility at Jeppesen. We discuss the impact of our work at Jeppesen in Section 6 and summarize this work's contributions in Section 7.

2. PROBLEM DESCRIPTION

Flexibility planning has been studied extensively during the last decade, and here we do not attempt to provide an exhaustive survey. For a summary of flexibility categories and measures see Sethi and Sethi [22] or Gupta and Goyal [15]. Recent frameworks for flexibility planning span the spectrum from qualitative and descriptive (Gerwin [12]), to purely theoretical (de Groote [8], van Mieghem [30]), to empirical (Suarez, Cusumano, and Fine [26]), to managerial (Upton [29]).

Our view is that the problem of evaluating an investment in a new production resource in general, and in a flexible resource in particular, consists of two parts. The first part is how to accurately estimate the future benefits the new resource will generate (for example, an uncertain stream of cash flows). The second question is how to properly determine the value of these benefits. In this paper, we focus on the first part: properly estimating the future benefits attributed to an investment. A large decision analysis and real options pricing[1] literature already addresses the second question. Smith and Nau [25] show the circumstances under which the real options and the decision analysis approaches are consistent.

Consider a problem setting where a manufacturing firm has an opportunity to improve its production process by purchasing a new piece of equipment. While the cost of the new resource is known, its real benefit is not. To find the true benefit of the new resource, we need to be able to compare the performance of the resulting production system with and without the new resource.

We have two ways to do this comparison. We can try to model the systems either *as it is actually used* or *as it should be used.* The first method can be achieved with a simulation model, and the second with an optimization model (for example, stochastic programming). Examples of the pure simulation approach include Azzone and Bertele [1], Suresh [27] and Das and Nagendra [7]. Simulation models are conceptually easy to understand and implement, but they can lead to suboptimal results. Ramasesh and Jayakumar [21] take simulation one step further, by using it to generate realizations of uncertain parameters that are then used as data in an optimization model. We will say more about the Ramasesh and Jayakumar approach in Section 3.2. Alternatively, optimization methods such as stochastic programming in theory yield optimal solutions, but real-sized, multistage stochastic programming models with recourse are often intractable. Examples of approaches based on stochastic programming include Sinha and Wei [23], Gupta, Gerchak, and Buzacott [14], and Fine and Freund [11]. Recently much work has been done on developing approximation methods for certain classes of stochastic programming problems (see Birge and Louveaux [3] and Infanger [16]).

Successful applications of stochastic programming include the work of Eppen, Martin, and Schrage [10], who developed a model for General Motors that uses a scenario approach to select the type and level of production capacity. Mulvey, Gould, and Morgan [20] describe an asset-and-liability management system developed for Towers Perrin-Tillinghast that uses stochastic programming to help its clients make major business decisions. Carino et al. [6] describe another asset and liability management system developed for Frank Russell Co. and The Yasuda Fire and Marine Insurance Co., Ltd., to determine the optimal investment strategy.

The new method we describe here is an extension of the approach first described by Katok [18]. It combines optimization with sampling to approximate system performance under uncertainty. The dynamics of the algorithm are consistent with decision-making practices shown to be superior by Benjaafar, Morin, and Talavage [2]. The new method is more intuitive and is easier to implement than stochastic programming, and is more robust and general than pure simulation.

3. MODEL DEVELOPMENT

To determine the benefit of investing in a new production resource, we wish to estimate the additional cash flows that the new resource will generate. To accomplish this, we model the current system (without the new resource) to determine the base cash flows. We then model the

[1] The idea behind the real option pricing approach is to apply finance methods for valuing put and call options to “real” projects. If one can construct a portfolio of financial instruments that exactly replicates the real project’s cash flows in every possible state of the world, the market value of this portfolio is the same as the value of the real project. See Smith and McCardle [24] for a detailed summary of real option pricing approach.

system with the new resource to determine the cash flows from the enhanced system. The difference between the two sets of cash flows can be attributed to the new resource. If the value of these additional cash flows (determined via decision analysis or real option pricing) exceeds the cost of the new resource, the new resource is worth obtaining.

Theoretically, the proper way to determine the performance of a system under uncertainty is with multistage stochastic programming. The objective function value of this model represents the system's performance. Since such large problems are notoriously difficult to solve to optimality, we develop approximate solutions. In the following sections we develop the stochastic programming formulation (also called the *recourse problem*) of the resource acquisition decision.

3.1. Problem Formulation

We use the stochastic programming notation of Birge and Louveaux [3], where random variables are denoted in bold. Consider a manufacturing firm that produces a set of products $\mathcal{P} = \{p|1, \ldots, P\}$ using a general assembly process. Each $p \in \mathcal{P}$ represents either a finished product or a subassembly. We can specify any type of bill of materials (BOM) structure by defining a set $\mathcal{S}_p$ (successors of p) for each product p to include immediate successors of p in the BOM. We also let $k_{p,j}$ be the number of units of p required to make one unit of j when $j \in \mathcal{S}_p$. If p is an end-item, $\mathcal{S}_p = \emptyset$. Let $\mathcal{R}$ be the set of production resources. Since each $p \in \mathcal{P}$ represents a product at a particular production stage, we assume, without loss of generality, that it needs to be processed only by one resource at each stage, although there may be alternative ways to process a product at a stage. Finally, let us assume that the model can be naturally decomposed into convenient time blocks $\mathcal{J} = \{\tau|\tau = 1, \ldots, J\}$, in such a way that there are not "many" interactions among the time blocks (ideally no interactions at all). Specifically, we assume that inventory cannot be carried across time blocks. We also assume that backorders across time blocks are allowed, but the interpretation of a backorder during the last period of a time block changes to unmet demand, so that there is never any backorder that has to be met in the first period of a time block. Each time block $\tau \in \mathcal{T}$, in turn, consists of time periods $t \in \mathcal{T}_\tau$. Therefore, the model can be decomposed into separate multistage stochastic programs for each time block τ. We identify each period in the model by a pair of indexes, (t, τ), representing the time period t of the time block τ. We introduce the time block notation for convenience, and without loss of generality. If time horizon cannot be broken into time blocks, we simply have a single time block in the problem. Since the Ramasesh and Jayakumar [21] model requires the use of time blocks, we introduce them here, to ensure consistency among models.

Let $\mathbf{d}_p^{t,\tau}$ be the demand for the end items only (in period t of time block τ for product p) and a random variable. When p is an intermediate item, $\mathbf{d}_p^{t,\tau} = 0$. If demand in period t is not filled, a unit backorder cost $\lambda_p^{t,\tau}$ is charged for the period. The processing time for product p on resource r at time t of time block τ is $a_{p,r}^{t,\tau}$, and we assume that processing times do not span multiple periods. If a product does not need to be processed on a particular resource, then $a_{p,r}^{t,\tau} = 0 \ \forall \ t, \tau$. Different resources involve different operating requirements, so let w_r be the cost of one unit of time on resource r. Finally, each resource has $c_r^{t,\tau}$ units of capacity available at time t of time block τ.

Let $\xi^{t,\tau}$ denote the vector of random parameters at time t of time block τ. The elements forming $\xi^{t,\tau}$ are demands $(\mathbf{d}_1^{t,\tau}, \ldots, \mathbf{d}_P^{t,\tau})$.

The decision variables are $\mathbf{x}_{p,r}^{t,\tau}$, representing the number of units of product p processed on resource r at time t of time block τ. These production decisions $\mathbf{x}_{p,r}^{t,\tau}$ are made at the beginning of time period t of time block τ, before the demand $\mathbf{d}_p^{t,\tau}$ for that time period is known. After the production decisions are made, the demand ($\mathbf{d}_p^{t,\tau}$) is revealed. At the end of the period the inventory decisions are made for the next period ($\mathbf{h}_p^{t+1,\tau}$) along with the backordering decisions ($\mathbf{b}_p^{t,\tau}$). So the inventory and the backorder variables are recourse variables that absorb the uncertainty in each period.

Assume for convenience and without loss of generality that the beginning and ending inventory levels are 0. Also, let $\rho_{t,\tau}$ be the compounded discount rate from the beginning of the planning horizon until period (t, τ); (1)–(5) is a mathematical programming formulation of the stochastic production-planning recourse problem (SPP):

$$\min z_{SPP} = E_\xi\left[\sum_\tau\left(\sum_t \frac{1}{(1+\rho_{t,\tau})}\left(\sum_{pr} w_r a_{p,r}^{t,\tau}\mathbf{x}_{p,r}^{t,\tau} + \sum_p \lambda_p^{t,\tau}\mathbf{b}_p^{t,\tau}\right)\right)\right] \tag{1}$$

subject to

$$\mathbf{h}_p^{t,\tau} + \sum_r \mathbf{x}_{p,r}^{t,\tau} - \sum_{r,j\in\mathbf{S}_p} k_{p,j}\mathbf{x}_{j,r}^{t,\tau} - \mathbf{d}_p^{t,\tau} + \mathbf{b}_p^{t,\tau} - \mathbf{b}_p^{t-1,\tau} - \mathbf{h}_p^{t+1,\tau} = 0 \qquad \forall\ p, t \neq 1, \tau, \tag{2}$$

$$\mathbf{h}_p^{1,\tau} + \sum_r \mathbf{x}_{p,r}^{1,\tau} - \sum_{r,j\in\mathbf{S}_p} k_{p,j}\mathbf{x}_{j,r}^{1,\tau} - \mathbf{d}_p^{1,\tau} + \mathbf{b}_p^{1,\tau} - \mathbf{h}_p^{2,\tau} = 0 \qquad \forall\ p, \tau, \tag{3}$$

$$\sum_p a_{p,r}^{t,\tau}\mathbf{x}_{p,r}^{t,\tau} \leq c_r^{t,\tau} \qquad \forall\ r, t, \tau, \tag{4}$$

$$\mathbf{h}_p^{t,\tau}, \mathbf{x}_{p,r}^{t,\tau}, \mathbf{b}_p^{t,\tau} \geq 0 \qquad \forall\ p, r, t, \tau. \tag{5}$$

Equation (1) is the objective function that minimizes the total expected discounted production and backorder cost, with expectation taken with respect to the random vector ξ. If the planning horizon is sufficiently long, we should include the inventory holding costs as well. Equations (2) and (3) are the set of material balance constraints that ensures that no product is processed until all its predecessors are available. Note that Eq. (3) is for the first period of a time block, where backorder from the previous time block does not have to be met. Equation (4) is the set of capacity constraints.

In practice, more simplistic procedures than stochastic programming are used to determine the value of flexibility, and we review two such procedures in the next section. Sometimes simple simulation-based methods do an adequate job, correctly approximating the benefit of flexible resources; nevertheless, at times, as we will demonstrate, simplistic methods may systematically underestimate the benefit of flexible equipment.

3.2. Alternative Methods

3.2.1. The "Wait and See" Model

If uncertainty can be approximated by a set of scenarios, then one way to determine the value of flexibility is to solve the so-called "wait and see" problem (WS). If we let the individual scenarios correspond to realizations of the random variable ξ, then Eqs. (6)–(10) can define the optimization problem associated with one particular scenario ξ:

$$\min z(\xi) = \sum_{\tau} \sum_{t} \frac{1}{(1+\rho_{t,\tau})} \left(\sum_{pr} w_r a_{pr}^{t,\tau} x_{pr}^{t,\tau} + \sum_{p} \lambda_p^{t,\tau} b_p^{t,\tau} \right) \tag{6}$$

subject to

$$h_p^{t,\tau} + \sum_r x_{p,r}^{t,\tau} - \sum_{r,j \in S_p} k_{p,j} x_{j,r}^{t,\tau} - d_p^{t,\tau} + b_p^{t,\tau} - b_p^{t-1,\tau} - h_p^{t+1,\tau} = 0 \qquad \forall\ p,\ t \neq 1,\ \tau, \tag{7}$$

$$h_p^{1,\tau} + \sum_r x_{p,r}^{1,\tau} - \sum_{r,j \in S_p} k_{p,j} x_{j,r}^{1,\tau} - d_p^{1,\tau} + b_p^{1,\tau} - h_p^{2,\tau} = 0 \qquad \forall\ p,\ \tau, \tag{8}$$

$$\sum_p a_{p,r}^{t,\tau} x_{p,r}^{t,\tau} \leq c_r^{t,\tau} \qquad \forall\ r,\ t,\ \tau, \tag{9}$$

$$h_p^{t,\tau},\ x_{p,r}^{t,\tau},\ b_p^{t,\tau} \geq 0 \qquad \forall\ p,\ r,\ t,\ \tau. \tag{10}$$

Here all variables and parameters indexed by t and τ represent quantities in period t of time-block τ.

Denote an optimal solution to (6)–(10) as $x^*(\xi)$ (since the x variables uniquely determine the h and the b variables), and the corresponding objective function value as $z(\xi)$. We can then compute $z_{WS} = E_{\xi} z(\xi)$, as the expected value of objective function values of deterministic subproblems corresponding to realizations of the random variables in all scenarios. This solution is known in the literature as the "wait and see" solution (Birge and Louveaux [3]).

Computing $z_{WS} = E_{\xi} z(\xi)$ exactly is unlikely to be practical because the number of scenarios can be extremely large. If this is the case, we must approximate $z_{WS} = E_{\xi} z(\xi)$ with a sample-mean estimate of z_{WS}. This is the approach we take for the empirical comparisons discussed in Section 4.

3.2.2. The Aggregate Model

A natural method to simplify computations of the optimal value of the objective function for the deterministic production planning problem and establish a base line on the benefit of new equipment is to consider an aggregate formulation (APP). This method is especially convenient when random variables can be naturally separated into several time blocks, with not many interactions among the time blocks. This is the Ramasesh and Jayakumar [21] approach. Eppen, Martin, and Schrage [10] use a similar approach, aggregating their capacity planning model developed for General Motors into five yearly time blocks.

In the aggregate formulation, the planning horizon consists of time blocks, $\mathcal{J} = \{\tau | \tau = 1, \ldots, J\}$. We aggregate the products into end-items. In this case the set $\mathcal{P}$ of products includes end-items only, and the decision variables $\mathbf{x}_{p,r}^{\tau}$ represent the number of units of the end-item p processed on resource r during the time-block τ. We can measure the per unit requirement of resource r by product p in time block τ, $A_{p,r}^{\tau} = \sum_{t \in \mathcal{T}_\tau} \sum_q a_{q,r}^{t,\tau}$ where q was in the BOM for p in the SPP model. The demand for product p is now the aggregate demand for the time block, $\mathbf{D}_p^{\tau} = \sum_{t \in \mathcal{T}_\tau} \mathbf{d}_p^t$. The capacity of resource r is the aggregate capacity for the time block, $C_r^{\tau} = \sum_{t \in \mathcal{T}_\tau} c_r^t$. If capacity is insufficient to fill current time-block demand, the product is backordered, and $\mathbf{B}_p^{\tau}$ is the total backorder of product p for time block τ. Note that the nature of backordering can be different in APP than in SPP, since in APP backordering represents the unmet demand, while in SPP backorders can be filled in subsequent periods. If we allow some demand at the end of a time block to remain unmet in SPP, that unmet demand has the same meaning as $\mathbf{B}_p^\tau$ in APP. The backorder cost $\lambda_p = \lambda_p^{t,\tau}$ when t is the last period in a time block τ. Since no inventory is carried across time blocks, we do not need the inventory variables. The discount factor ρ_τ in the aggregate model is the single period discount factor $\rho_{t,\tau}$, compounded over the time block τ, $\rho_\tau = (\Pi_{t\in\mathcal{T}_\tau}(1 + \rho_{t,\tau})) - 1$. The aggregate mathematical formulation is similar to the problem described by Ramasesh and Jayakumar [21], and we make every attempt to use notation consistent with theirs:

$$\min z_{APP} = E_\xi\left[\sum_\tau \frac{1}{(1+\rho_\tau)}\left(\sum_{pr} w_r A_{p,r}^\tau \mathbf{x}_{p,r}^\tau + \sum_p \lambda_p \mathbf{B}_p^\tau\right)\right] \tag{11}$$

subject to

$$\sum_r \mathbf{x}_{p,r}^\tau + \mathbf{B}_p^\tau = \mathbf{D}_p^\tau \qquad \forall\ p, \tau, \tag{12}$$

$$\sum_p A_{p,r}^\tau \mathbf{x}_{p,r}^\tau \le C_r^\tau \qquad \forall\ r, \tau, \tag{13}$$

$$\mathbf{x}_{p,r}^\tau, \mathbf{B}_p^\tau \ge 0 \qquad \forall\ p, r, \tau. \tag{14}$$

If several of the resources (r') are interchangeable, constraint (13) becomes

$$\sum_{p,r'} A_{p,r'}^\tau \mathbf{x}_{p,r'}^\tau \le \sum_{r'} C_{r'}^\tau \qquad \forall\ \tau.$$

Even though the optimization model described by (11)–(14) is smaller than SPP, and separates into one problem for each time block, just as SPP does, solving it directly may not be computationally feasible. However, Ramasesh and Jayakumar [21] develop and test an efficient method for finding approximate solutions. Following the approach of Ramasesh and Jayakumar [21], we assume the demand is known at the beginning of each time block and is different for subsequent time blocks. Again following the approach of Ramasesh and Jayakumar [21], we can estimate the system performance over time by drawing realizations of uncertain parameters from their distributions, and solving the aggregate problem several times. Ramasesh and Jayakumar [21] show that this approach gives solutions very close to optimal solutions to the aggregate problem. However, APP is a relaxation of SPP, and therefore z_{APP} is a lower bound on z_{SPP}.

Now let us analyze APP's estimates for the benefit of a flexible resource. First, let us say that we have the base-line system consisting of a set of resources $\mathcal{R}$, and a new system, consisting of a set of resources $\mathcal{R}'$, where $\mathcal{R}' = \mathcal{R} \cup \{r_{new}\}$. Let $V_{\text{APP}}(\mathcal{R}') = z_{\text{APP}}(\mathcal{R}') - z_{\text{APP}}(\mathcal{R})$ represent the APP estimate of the benefit of the new set of production resources $\{r_{new}\}$, and also let $V_{\text{WS}}(\mathcal{R}') = z_{\text{WS}}(\mathcal{R}') - z_{\text{WS}}(\mathcal{R})$ represent the WS estimate of the benefit of $\{r_{new}\}$. Recall that we postulated that there are three sources of benefit of a resource: (1) lower production cost, (2) capacity, and (3) decision flexibility. The problem APP considers the cost of operating a resource (unlike the Ramasesh and Jayakumar [21] formulation that looks at the time rather than the cost), so the portion of the new resource benefit due to any productivity improvement that results in lower production cost is addressed by APP.

APP only partially accounts for benefit due to capacity. Problem APP has a capacity constraint that preserves the aggregate capacity for the time block. It is possible, however, to observe the aggregate capacity constraint while violating capacity constraints for single periods. For example, if each day has eight hours of capacity, and the time block has 2 days, the aggregate capacity constraint tells us that we cannot exceed the 16 hours capacity in a 2-day period. But a production plan requiring 10 hours on day 1 and 6 hours on day 2 is still aggregate-feasible, although the plan exceeds day 1 capacity and allows an unrealistic shift of available hours. A stronger capacity constraint would force the model to allocate hours properly and highlight the benefits from having the additional capacity on days when it is required. Since APP has a weaker capacity constraint than SPP, the benefit of the new resource due to capacity can be underestimated because the model will not identify benefits on capacitated days.

When decisions are made in APP, all relevant time-block information is known. Benjaafar, Morin, and Talavage [2] show that decision flexibility provides no benefit if no relevant future information is expected. This result implies that APP does not account for any benefit of the new resource due to an increase in decision flexibility, but it does provide an approximation for the benefit from efficiency gain and partial benefit from capacity gain. Similarly, $V_{\mathrm{WS}}(\mathcal{R}') - V_{\mathrm{APP}}(\mathcal{R}')$ provides an estimate for the gains from capacity not captured in $V_{\mathrm{APP}}(\mathcal{R}')$. And most importantly, $V_{\mathrm{SPP}}(\mathcal{R}') - V_{\mathrm{WS}}(\mathcal{R}')$ *provides an estimate for the gains from decision flexibility.*

3.3. The New Method

Both APP and WS make a part of the SPP recourse problem into a deterministic problem and then solve a sequence of deterministic problems with parameters representing realizations of stochastic parameters. The solution to a problem where stochastic parameters are replaced with their realizations is called a *wait-and-see* solution. Birge and Louveaux [3, p. 140] prove that the *wait-and-see* solution is a lower bound on the recourse problem solution (in our terminology, $z_{\mathrm{WS}} \leq z_{\mathrm{SPP}}$). Birge and Louveaux [3, p. 139] also describe the notion of the *expected result of using the expected value solution (EEV).* This is the expected system performance that results if, at the beginning of the planning horizon, we solve a problem replacing all stochastic parameters with their expected values and implement the solution. Birge and Louveaux [3] show that EEV is an upper bound on the recourse problem solution z_{SPP}, because, by construction, EEV is always a feasible solution to the recourse problem. It turns out that typically, EEV is a weak upper bound. Our goal is to develop an algorithm with a stronger upper bound on z_{SPP}.

We begin by looking at the decision-making process under uncertainty. Benjaafar, Morin, and Talavage [2] postulate that there are two general approaches to "flow control decisions in manufacturing." The *planning-based* approach applies when a production plan is determined prior to the beginning of production (at $t = 0$) and is rigidly adhered to. It is similar in flavor to the EEV solution. The *real-time-based,* or *opportunistic* approach, allows decision making to be "contemporaneous with action implementations." Decisions are made based on the state of the system, and no decision is implemented until it has to be. Benjaafar, Morin, and Talavage [2] show that "under conditions of uncertainty, opportunism is superior to planning."

To generate an upper bound on z_{SPP} we apply the opportunistic decision making process and use the rolling horizon strategy (see Bitran and Sarkar [4] and Bitran and Yanasse [5]). Under this strategy, we solve a multiperiod problem each period, but only implement first-period decisions and keep track of the first period performance measures. Algorithm 1 formally outlines the method, and the resulting process for estimating the performance of a production system under uncertainty. We wish to estimate the difference in the total expected cost of operation over the T-period planning horizon of n different systems with different sets of production resources. Let $\mathcal{R}_i$ denote the set of resources in the ith system under consideration. Let $z^t(\mathcal{R}_i)$ be the cost during period t of a system with resources $\mathcal{R}_i$. We estimate the objective function value of the recourse problem z_{SPP} with an estimate $\tilde{z}_{\mathrm{SPP}}$ by repeating the opportunistic decision-making task m times.

Algorithm 1 can be viewed as a WS-induced policy, because at every stage t we form a WS problem (with a randomly selected future) and progress towards the end of the planning horizon T, resampling at each stage. Since a new scenario is generated at the beginning of each stage, the solution implemented at each stage is derived without "seeing the future," and hence the value of flexibility can reveal itself. Algorithm 1 estimates the value of a new resource rather than determining an actual policy to be implemented. The approach is somewhat related to the "scenario analysis" work (see Dembo [9]).

By construction, the Algorithm 1 solution is feasible in the m instances of random parameter realizations. As m grows large $P\{\tilde{z}_{\mathrm{SPP}} \geq z_{\mathrm{SPP}}\}$ goes to 1, in other words, as m grows large, it becomes more likely that $\tilde{z}_{\mathrm{SPP}}$ is a valid upper bound to the objective function value of the recourse problem.[3] When it is an upper bound, it is likely to be a stronger upper bound than EEV[4] in expectation, because, again by construction, EEV is a solution where $m = 1$, and all random variables are simply replaced by their expected values, while $\tilde{z}_{\mathrm{SPP}}$ is determined using a larger value of m.

```
j ← 0; z̃_SPP ← 0
foreach ℛ_i (i = 1, . . . , n) do
    while (j ≤ m) do
        foreach t ∈ 1, . . . , T do
            fix decision variables for periods t' < t
            randomly generate realizations of uncertain parameters²
            solve the deterministic flexibility-planning problem
            implement decisions for the current period
            z^t(ℛ_i) ← current period cost
            z^j_SPP ← z^j_SPP + z^t(ℛ_i)
        endfor
    endwhile
    z̃_SPP(ℛ_i) ← (Σ_{j=1}^{m} z^j_SPP) / m
endfor
```

Algorithm 1. Sampling Algorithm for finding a bound on the objective function for a multistage stochastic program.

The major benefit of our modeling technique is that it accounts for the opportunistic decision-making process, explicitly modeling decision flexibility. Therefore, unlike the APP and WS models, our new model will be less likely to underestimate the benefit of a resource as much as APP and WS when that resource provides decision flexibility.

4. JEPPESEN SANDERSON, INC.

The applied portion of this work focuses on Jeppesen's production system for flight manual revision service. For a more complete description of Jeppesen see Katok, Tarantino and

[2] We actually use the same set of scenarios for all the systems we compare.

[3] For example, if we actually solve the problem for every possible scenario, $\tilde{z}_{\mathrm{SPP}}$ becomes a feasible solution to the recourse problem by construction.

[4] Although it is possible to construct examples where it is not a stronger upper bound.

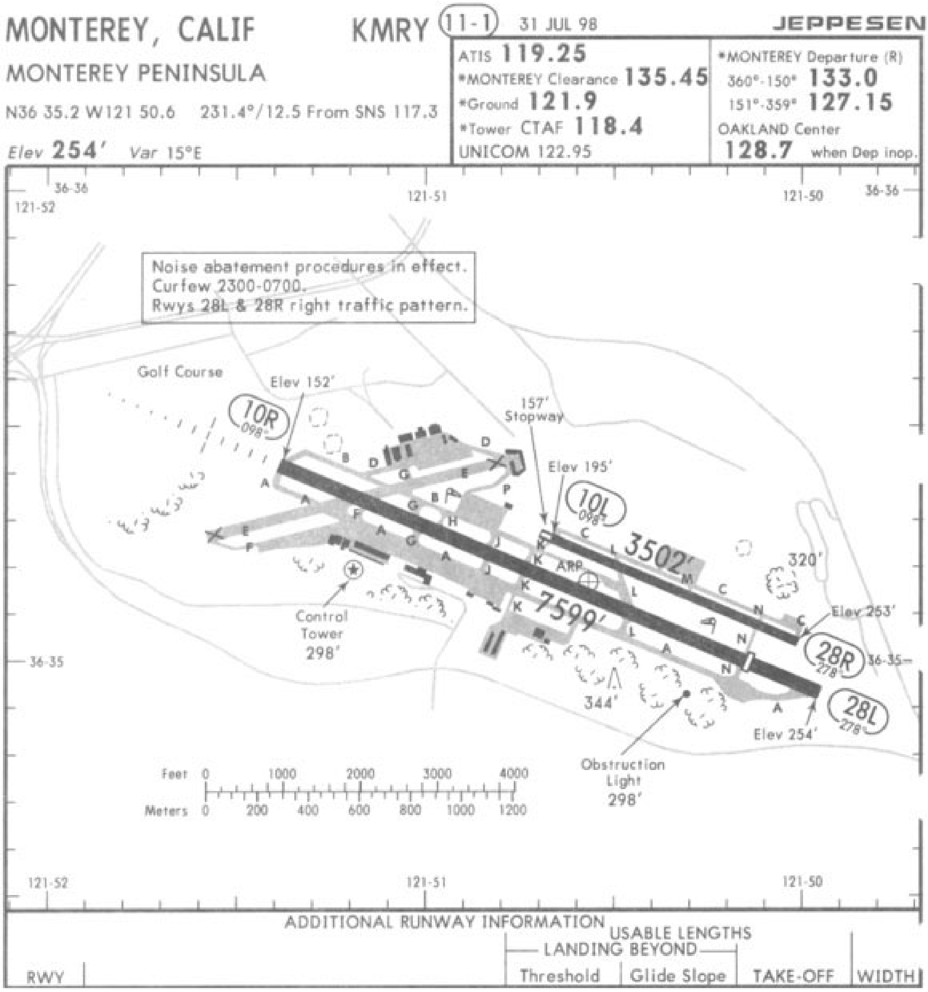

Figure 1. A typical Jeppesen chart.

Tiedeman [19] and Tarantino [28]. Airway safety considerations dictate that all pilots on all flights must have a set of airport maps, enroute charts, and other flight information for the area within a 200-mile radius of the planned route. Flight information changes constantly, so this material must be updated regularly. For example, about 75% of all charts are revised at least once annually, and many charts are amended much more often. Enroute charts that cover large areas change on average four times a year. A typical Jeppesen chart is shown in Figure 1.

Jeppesen usually configures its flight manuals by geographical area. Many pilots subscribe to what Jeppesen refers to as the “Airway service”; however, many of Jeppesen’s large customers, including major airlines such as United, American, and Delta and package delivery services such as FedEx and UPS request special subscription packages. These special packages, called “Air Carrier coverages,” can differ from standard coverages because they contain charts with special information, a customized configuration of pages, or other specific features that a customer might request. Jeppesen maintains over 200 different standard coverages and over 2000 different tailored coverages, made up of over 100,000 distinct images.

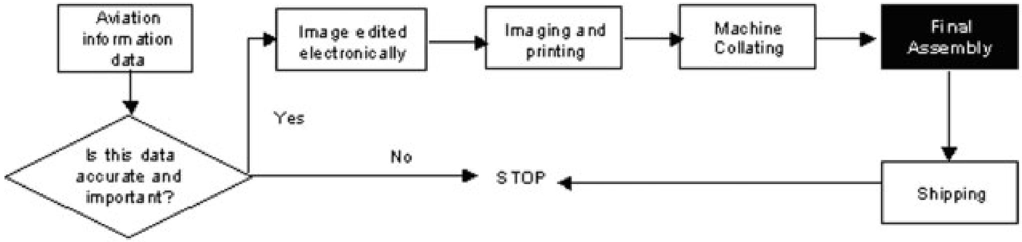

Figure 2. Revision management and production process at Jeppesen.

When critical aviation information changes (such as a runway at an airport is closed or expanded), the change affects multiple Jeppesen charts. Typically, a change affects one airway chart and several customized air carrier charts. When a chart is revised, Jeppesen issues a new manual page to all customers subscribing to coverages containing this page within 1 week of the change. Every week Jeppesen sends out between 3 and 25 million pages to over 300,000 different customers. Some weeks over 1,500 images, affecting over 1,000 different coverages must be altered. Figure 2 shows a diagram of the Jeppesen revision management and production process.

When information regarding a possible change first reaches Jeppesen, a decision is made as to whether this data is important or permanent enough to amend a chart. Some changes do not need to be included on a chart. For example, a runway closing for 20 minutes on a particular day would not require a revision (and will be handled with a "notam"). If a change to a chart is deemed necessary, the first step of the process involves electronically editing the image file. Some alterations are easy to make taking less than 5 minutes, while other changes can require as much as 8 hours of work. After an image file has been edited electronically, a new negative is printed. This negative goes to the first step of the production process, imaging and printing, where it is stripped onto a plate containing 21 negatives, the plate is printed, cut into individual sheets, and specially bound. Sheets then go into the machine collating area, where they are collated into sections. Each section contains up to 25 sheets that will eventually all go into the same coverage. At this point large maps, called folds, are not included in sections, because collating machines cannot handle folded material. Sections and folds go into the final assembly area, where prior to the implementation of our work they were manually assembled into coverages and stuffed into envelopes. Large boxes of envelopes go on to the shipping department. If a coverage completes final assembly on time it is shipped using standard shipping services, but if it is late the service is upgraded to overnight delivery. The bottleneck of the production process forms in final assembly, highlighted in Figure 2. Prior to the implementation of our work, in final assembly sections and folds were arranged and stuffed manually, often by a large number of temporary employees. Figure 3 shows a photograph of a typical Jeppesen assembly process. The use of temporary employees has several disadvantages for Jeppesen. They are often unfamiliar with the work, and tend to be less productive and make more mistakes than full-time employees. Jeppesen customers do not tolerate errors, so all errors are detected and corrected at great expense prior to shipping. The availability of temporary employees can also be unpredictable. Because of these problems, Jeppesen's management wished to evaluate the purchase of a new, automated technology, called a folder collator, for final assembly, and asked us to help them with this decision. The dynamic and complex nature of the Jeppesen operating environment makes properly determining the benefit of the new technology difficult, and hence the application of our method well-warranted.

Figure 3. The manual process

5. EMPIRICAL COMPARISONS

In this section we demonstrate how the three approaches to estimating the benefit of a flexible resource can yield different results. Jeppesen operates on an 8-week revision cycle involving three week-types with differing demand volumes: odd weeks have relatively low volume, even weeks have medium volume, and eighth weeks have the highest volume. Over time, revision characteristics in terms of overall volume (number of customers), the number of different coverages, average volume, and coverage size in terms of both, folds and flats, have been evolving. Figures 4a and 4b show historical trends in weekly revision for relevant dimensions since 1995.

As we mentioned earlier, Jeppesen has two types of customers: The Air Carrier customers, including primarily airlines and package delivery services, subscribe to customized products, while Airway customers subscribe to standard manuals, and include primarily corporate and private pilots. Historically, there are a relatively small number of airway manuals, and each has a large customer base. However, we see from Figure 4a that the number of airway coverages is increasing dramatically, and average quantity per coverage is dropping. The number of air carrier coverages (Fig. 4b) is growing also, but much slower, and the average quantity seems fairly steady. Air carrier coverages, however, are increasing in terms of the number of both, flat and folded charts. When estimating the benefit of new technology for the future, it is important to forecast these various trends into the future as accurately as possible.

The Jeppesen production problem is stochastic because production must begin before the entire weekly demand is known. That is, when the production is scheduled, prior to the first day

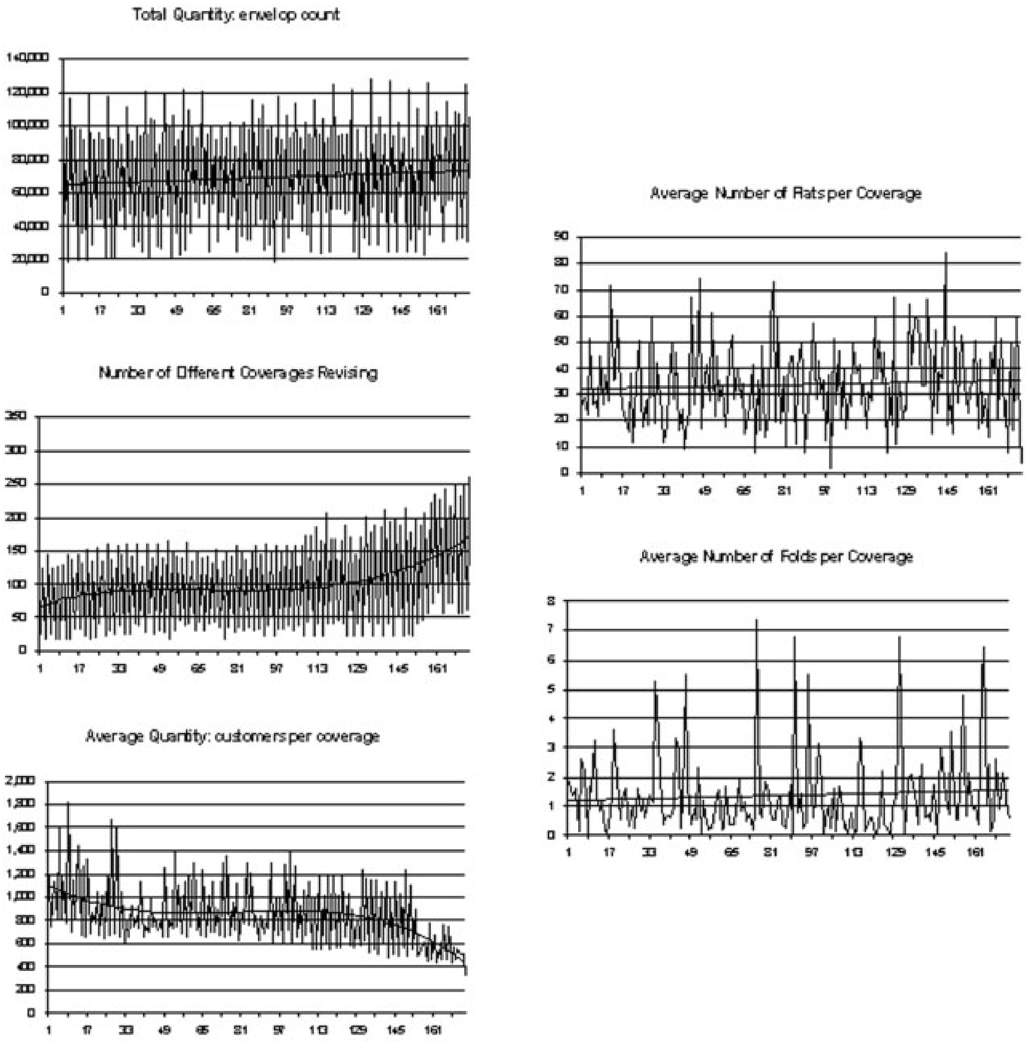

(a)

Figure 4a. *Revision characteristics over time: Airway.* The number of airway coverages revising increases significantly over time, while the average number of Airway customers per coverage declines, highlighting the fact that demand for customized products increases over time.

of the week, the real demand is still a random variable. The precise moment the weekly demand is finalized at Jeppesen is a matter of some debate. Jeppesen assigns official close dates, but they are not always adhered to because Jeppesen goes to great lengths to accommodate its customers. Therefore, for a good part of the week, demand is a moving target.

5.1. Modeling the Jeppesen Problem

In this section we recast the Jeppesen problem as a Stochastic Production-Planning Recourse Problem (SPP). The set of products $\mathcal{P} = \{p|1, \ldots P\}$ include all finished products, as well as

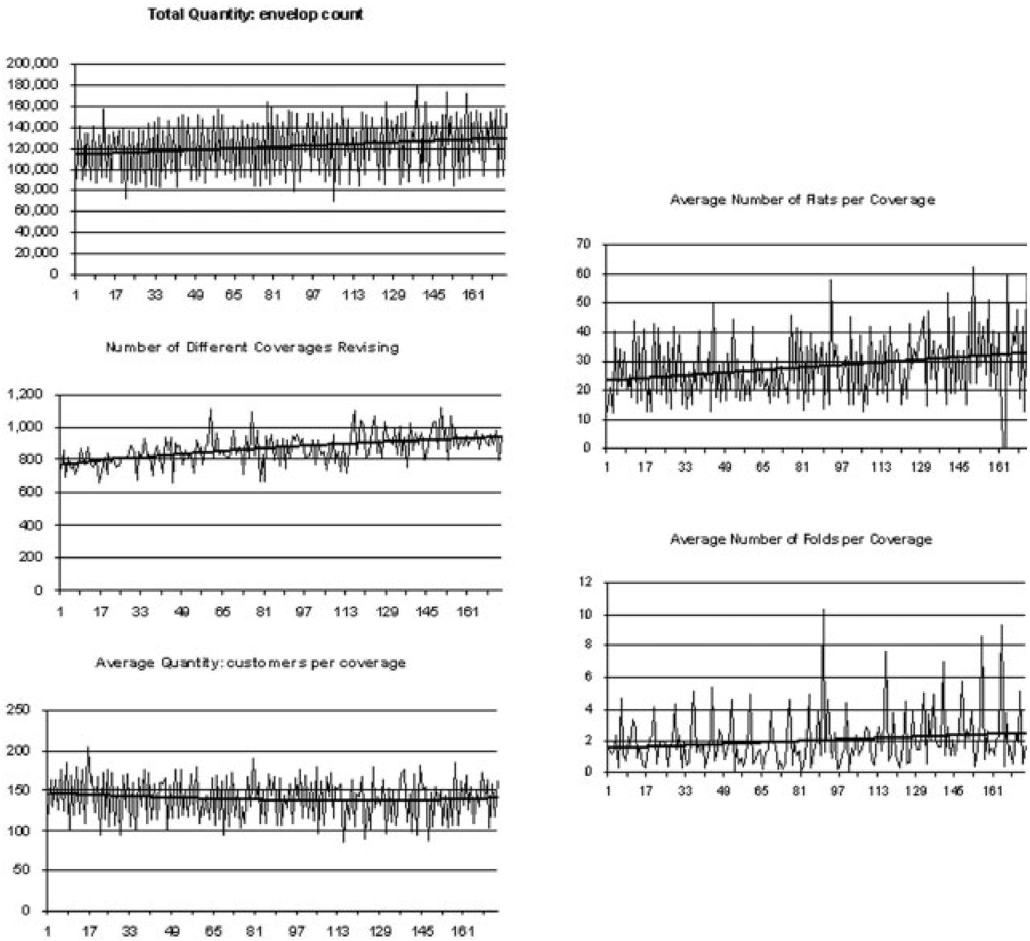

(b)

Figure 4b. *Revision characteristics over time: Air Carrier.* The number of air carrier coverages revising is fairly steady over time, but the average number of flats and folds per coverage is growing. So Air Carrier coverages are becoming larger over time—airlines add information to their customized coverages.

intermediate subassemblies. At Jeppesen, the notion of a "product" changes as the material moves through the production system. In the printing area, and as far as printing vendors are concerned, products are individual charts and folds. In the machine collating area the products are sections, composed of groups of 25 or 36 flat charts. For final assembly, products are coverages. The set $\mathcal{S}_p$ defines the bill of materials (BOM) structure for product p, and Figure 5 shows the BOM for Jeppesen revision products. In general,

$$\mathcal{S}_p = \begin{cases} \emptyset & \text{when } p \text{ is a coverage,} \\ \{coverages\} & \text{when } p \text{ is a section or a fold,} \\ \{sections\} & \text{when } p \text{ is a flat chart.} \end{cases}$$

For Jeppesen $k_{pj} = 1 \ \forall \ p, j$ since coverages never contain multiple copies of charts. The set of production resources $\mathcal{R} = \{r|1, \ldots R\}$ includes four different types of printing presses, a

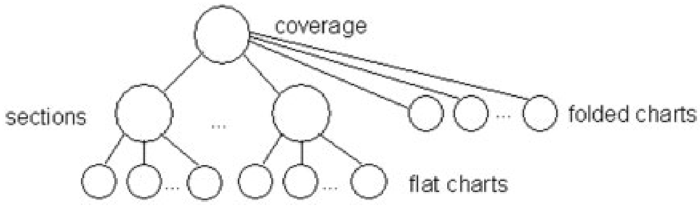

Figure 5. The BOM for Jeppesen revision products.

bindery, several outside printing vendors, two types of collating machines, manual assembly, the new folder collator, and a fold-collating vendor. Capacities of those resources $c_r^{t,\tau}$ are well known at Jeppesen, and are measured in hours a resource is available for operation during a particular day.

Jeppesen revision assembly planning is done on a weekly basis, with no major interactions between weeks. Due to the airway community's 8-week operating cycle, there are a large number of charts scheduled to revise in intervals that are multiples of 8 weeks. So generally, every 8th week Jeppesen faces a very large revision. Even weeks (weeks 2, 4, and 6 of each cycle) are medium-sized, and odd weeks (weeks 1, 3, 5, and 7) are comparatively small. A 1-week problem is a complete planning problem because of the lack of interaction among weeks, so the set of time blocks $\mathcal{J}$ for Jeppesen consists of a single one-week time block, running from Friday afternoon to the following Friday morning. Revision information, however, is only partially known at the beginning of the week, and changes every day, with the main information update occurring each Monday, but minor updates occurring daily. So effectively each weekly time block is broken into eight daily time periods t (where the Friday time periods are actually shorter than one day).

The backorder structure for the Jeppesen problem is very simple. If there is not enough capacity to meet demand, the product is late. Late products incur a large penalty in the objective function for each day of lateness. This penalty $\lambda_p^{t,\tau}$ represents not only the increased shipping costs (because late products are automatically upgraded to overnight shipping) but also the loss of good will. Although in practice a Jeppesen revision is occasionally late, lateness is generally avoided at all costs, and only happens due to extraordinary circumstances (a machine breakdown at a critical time, or vendor error, for example). When t is the last period of a time block, $\lambda_p^{t,\tau}$ actually represents the cost of meeting the demand with an outside vendor of last resort, so it is very high.

The demand $\mathbf{d}_p^{t,\tau}$ exists only for coverages, and the demand for most coverages occurs on the second Friday of the week ($t = 8$), but some coverages that have long shipping times, such as Australian coverages, are due earlier ($t = 6$, for example).

To create a realistic sample of demand scenarios we used 173 weeks of demand data that started on 6 January 1995 to estimate relevant attributes of the demand. System load depends on: the total quantity demanded, number of different coverages, number of customers per coverage, number of flats per coverage, and number of folds per coverage. Historical trends for those five demand characteristics for Airway and Air Carrier are shown in Figures 4a and 4b, and we forecast all of them to generate realistic demand scenarios. Figures 4a and 4b show that there is a clear cyclical component to revision, and in most cases there is also a trend component. We fit a forecasting model to the historical data, of the form in (15), using Ordinary Least Square (OLS) estimate,

$$\begin{aligned} Lo\hat{a}d_t = {} & Intercept + Trend \times time\ dummy\ variable \\ & + Even \times even\ dummy\ variable + Eighth \times eighth\ dummy\ variable + \epsilon_t, \end{aligned} \quad (15)$$

where *time dummy variable* is a week number starting with week 1 being 6 January 1995, *even dummy variable* is 1 for revision cycle weeks 2, 4, 6, and 8, and *eighth dummy variable* is 1 for week 8. Table 1 shows the regression results for the five relevant demand attributes.

Note that in most cases (15) does a good job of explaining the variability in the data, and generally all variables are significant. One exception is the folds per coverage in airway, where the only significant variable is *eight.* Also, the time trend is not significant in the flats per coverage for airway. The results in Table 1 implies that during an odd week in 1995 around 87,000 envelopes were sent out to air carrier customers and around 32,000 to airway customers. Since then, this quantity has been growing steadily at a weekly rate of about 93 envelopes for air carrier and 42 envelopes for airway (the historical data spans 173 weeks, so during an odd week in 1999 about 103,000 envelopes are sent out to air carrier customers and almost 40,000 to airway customers). During an even week, on average additional 52,000 envelopes are sent out to air carrier customers, and 61,000 to airway customers (bringing 1999's even week total to 155,000 for airway and 101,000 air carriers). During an 8th week (which is also an even week) an additional 52,000 envelopes are mailed out to air carrier customers, and 81,000 to airway on average (bringing an eighth week total in 1999 to over 200,000 air carrier envelopes and over 180,000 airway envelopes). The forecasting model works in a similar way for all five dimensions, so the average number of subscribers per coverage, for example, decreases from week to week. To generate a demand scenario for a particular week we use (15) to estimate the expected values of the five demand attributes and their standard errors, and draw a demand scenario from the resulting distribution.[5]

For the purpose of the tests, we assume inter-stage independence for the vector of random parameters $\xi^{t,\tau}$. Although demand information at Jeppesen is updated daily, new information significantly impacts planning only once, on Monday ($t = 4$) of every week. So a 1-week planning problem is a two-stage stochastic model with recourse, where the initial plan is made on Friday ($\tau = 1$, $t = 1$), production starts and proceeds for 3 days, demand information is revised on Monday ($\tau = 1$, $t = 4$), and the plan is adjusted given the new information.

Parameter w_r represents the labor cost on resource r, and it is generally well known. Unfortunately, accurate processing times for the resources $a_{p,r}^{t,\tau}$ were not as readily available. There were "standard" processing rates, but they did not represent reality. For the purpose of the empirical tests, the processing times are assumed to be deterministic, but the problem of determining accurate estimates of the processing times is interesting, because the time it takes to assemble a coverage, for example, depends on several variables: coverage quantity, the number of sections, the number of folds, and on whether a temporary or a permanent employee performs the work.

Using the manual assembly process as an example, we determined the total processing times by systematically tracking actual processing and setup times for each coverage over a one week period. We then fit the following model:

$$a_{p,assembly} = \alpha_p + \beta_1\, folds_p + \beta_2\, sections_p + \beta_3 d_p + \epsilon_p, \quad (16)$$

[5] Note that we are estimating the load for the week rather than the demand for each individual product. Individual product demands are a function of the week type (even, odd, or 8th), and therefore these demands are correlated.

Table 1. Forecasting assembly load.

	Air carrier Load (units are coverages or charts) (p-value)					**Airway** Load (units are coverages or charts) (p-value)				
	Intercept	Trend	Even	Eight	r^2	Intercept	Trend	Even	Eight	r^2
Total quantity (envelope count)	86,929 (0.0000)	93.23 (0.0000)	52,430 (0.0000)	51,810 (0.0000)	0.88	32,239 (0.0000)	41.84 (0.0000)	61,618 (0.0000)	81,068 (0.0000)	0.93
Number of coverages in revision	744 (0.0000)	0.91 (0.0000)	68.71 (0.0000)	129.57 (0.0000)	0.48	20 (0.0000)	0.35 (0.0000)	127.35 (0.0000)	140.20 (0.0000)	0.93
Customers per coverage	120 (0.0000)	−0.04 (0.0423)	49.78 (0.0000)	38.85 (0.0000)	0.74	1228 (0.0000)	2.00 (0.0000)	436.10 (0.0000)	373.45 (0.0000)	0.67
Flats per coverage	15 (0.0000)	0.06 (0.0000)	14.14 (0.0000)	25.73 (0.0000)	0.66	24 (0.0000)	0.02 (0.3890)	10.74 (0.0000)	30.13 (0.0000)	0.38
Folds per coverage	1 (0.0000)	0.01 (0.0029)	0.46 (0.0251)	3.68 (0.0000)	0.49	1 (0.0000)	0.00 (0.3287)	−0.36 (0.1054)	0.90 (0.0059)	0.08

Table 2. Comparative system performance for 17 weeks.

Date	Week in cycle	Percentage deviation from SPP		Problem size (for a single scenario problem)		Solution time (CPU sec)
		WS	APP	Variables	Constraints	
21-Aug	3	0.00	0.00	52,096	2825	946.04
28-Aug	4	0.09	60.10	49,304	2674	505.81
4-Sep	5	0.00	0.00	65,872	3667	727.1
11-Sep	6	0.00	1.44	32,312	1685	345.6
18-Sep	7	0.00	0.00	32,000	1595	540.01
25-Sep	8	0.19	11.23	63,544	3441	1223.07
2-Oct	1	0.00	0.00	50,136	2729	501.25
9-Oct	2	16.67	22.67	64,304	3566	935.21
16-Oct	3	0.00	0.00	50,040	2801	307.8
23-Oct	4	0.02	62.60	38,712	2077	392.45
30-Oct	5	0.00	1.40	35,928	1846	483.13
6-Nov	6	18.49	26.85	46,696	2513	564.74
13-Nov	7	25.53	25.53	56,864	3172	421.22
20-Nov	8	53.27	76.85	42,864	2234	626.86
27-Nov	1	0.00	0.00	41,784	2224	523.06
4-Dec	2	4.34	12.82	46,304	2450	685.53
11-Dec	3	6.09	21.41	69,480	3854	1334.98

where α_p is the intercept term, $folds_p$ represents the number of folds in coverage p, $sections_p$ represents the number of sections in coverage p, d_p represents the quantity of coverage p demanded, and ϵ_p is an unobservable random error (see Tarantino [28]). Equation (16) gives us an approximation of the total time in assembly. We fit the model using ordinary least squares. All coefficients were significant, and the resulting r^2 was 78.1%. We determined processing times for other resources using the same method.

Estimating processing times with the new collator was a more difficult task because we did not have the opportunity to observe the collator's performance in the Jeppesen production environment. Instead, a team of Jeppesen managers observed the collator's performance at the vendor's site. They collected the production data that we ultimately used to estimate collator processing times.

5.2. Comparative Results

To begin our empirical comparison of the three flexibility evaluation approaches, we picked 17 actual consecutive weeks (two complete 8-week cycles, and one additional week following the second cycle): 21 August 1998 through 11 December 1998. The date 21 August 1998 is the Friday of week 3, 28 August is the Friday of week 4, 4 September is the Friday of week 5, and so forth. We had the actual data that went into the revision, and the sequence in which this data was becoming available to the planning group.

We ran the three models on the 17 weeks of data, which involved solving 17 separate two-stage stochastic problems. We modeled the current state of the system and the hypothetical system configuration with the new folder-collator. To determine the benefit of the collator each week, we take the difference of revision cost with and without the collator. Table 2 compares the SPP estimates of collator benefit with those of APP and WS, and reports the size of the one-scenario problem. All numbers are presented as percentage difference with SPP. We estimate $V_{SPP}(\mathcal{R}')$ by running 30 replications of Algorithm 1 on a 2-stage problem, and the

Table 3. Summary of solution results for estimating collator benefit.

	SPP	WS		APP	
Week	Average weekly benefit ($) (standard error)	Average weekly benefit (standard error)	Deviation from SPP	Average weekly benefit ($) (standard error)	Deviation from SPP
Odd	1652.71	1220.67	26.1%	1164.69	29.5%
	393.36	263.29		258.20	
Eight	190,086.45	111,280.79	41.5%	88,343.54	53.5%
	73,680.73	39,722.19		46,271.29	
Even	7357.41	5450.43	25.9%	4735.58	35.6%
	1542.56	884.53		1344.10	

solution time we report is for $m = 30$. We learn several things from Table 2. First, we see that in every case

$$V_{\text{APP}}(\mathcal{R}') \leq V_{\text{WS}}(\mathcal{R}') \leq V_{\text{SPP}}(\mathcal{R}')$$

We also observe that in many of the weeks APP and WS models underestimate the benefit of the collator relative to SPP. All three models give the same solution in several of the weeks. Those are all small odd weeks, with low load.

Most of the savings from the collator are due to the two 8th weeks, since the 8th weeks are the only weeks where internal capacity is insufficient to fulfill demand and an outside vendor is used for fold assembly. The outside vendor is much more expensive than internal fold assembly, even if overtime and temporary employees are used. With the new collator, the use of the outside vendor can be avoided.

5.3. Estimating the Total Collator Impact

We now compare how the three models estimate the benefit of the new collator over the 3-year planning horizon. The previous section showed that the benefit of this resource increases with system load.

We estimate the benefit of the collator by simulating the three-year Jeppesen production environment, based on (15). In other words, we generated 156 weeks of demand consistent with demand characteristics as presented in Table 1. Each week is a separate two-stage stochastic model, and the SPP estimates were obtained for each week separately by running 30 replications of Algorithm 1. Table 3 summarizes average weekly benefit estimates for all three models, along with their standard errors. Figure 6 presents our analysis graphically.

In the odd week problems, WS and APP results are generally quite close because odd weeks have low volume, on average, so capacity virtually never becomes an issue. Results look very different in even and 8th weeks. We clearly see that WS and APP underestimate the benefit of the collator relative to SPP. In even weeks, APP undervalues the collator by about 36% and WS by 26%. In 8th weeks, APP undervalues the collator by 54% and WS by 41%. In odd weeks, both APP and WS undervalue the collator by about 26%. All of these differences are highly statistically significant, using a two sample t-test assuming unequal variances.

Our analysis showed that the annual discounted savings a new collator will generate are around $1.4 million. Jeppesen accepted our analysis and recommendation, and purchased the collator in July 1998.

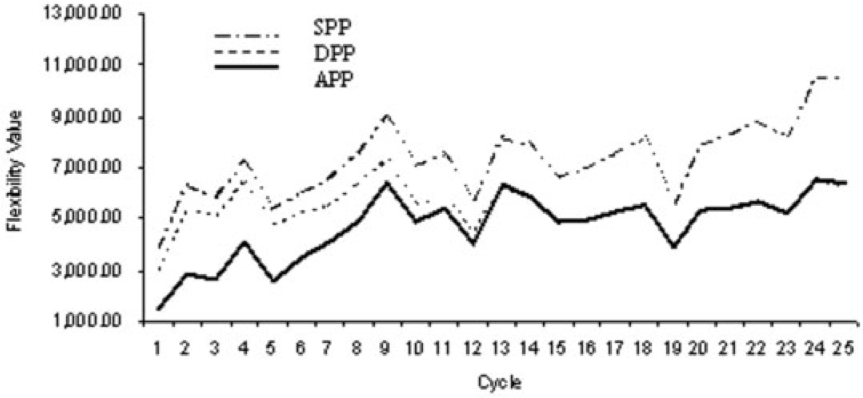

(a) Typical even week

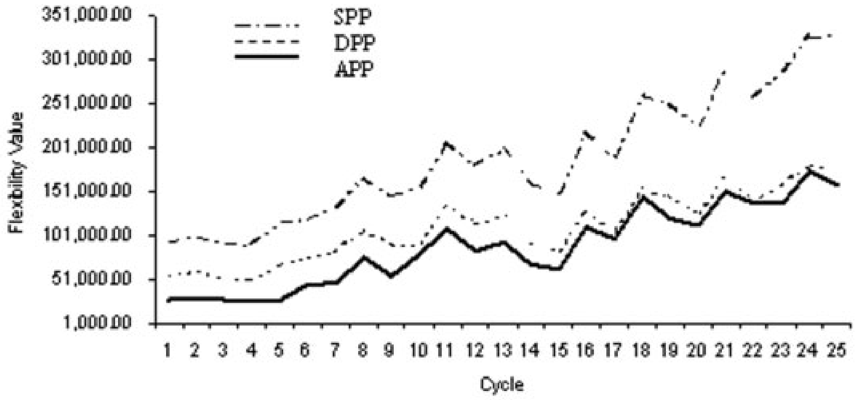

(b) Typical eighth week

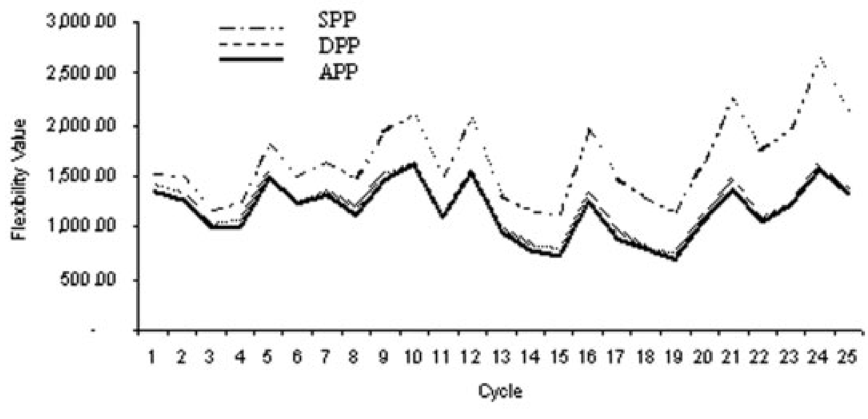

(c) Typical odd week

Figure 6. Graphical representation of results.

Figure 7. Collator in use at Jeppesen

6. COLLATOR IMPACT

The collator, referred to as “Longford” at Jeppesen, after its manufacturer, was custom-built and delivered in late December 1998. Figure 7 shows the collator in action at Jeppesen. After an initial training period, Jeppesen’s assembly area started using the collator on 6 January 1999. The assembly operations manager kept track of all work performed on the collator per our request. Table 4 summarizes this data for the period of 8 January 1999 through 21 May 1999, and compares the actual savings with savings forecasted by the three alternate models.

We determine actual savings each week by considering that week’s entire revision and determining which coverages should be assembled using the collator. We then compare the actual cost of assemblying these coverages on the collator, and what it would have cost to manually assemble them. The difference is in the column labeled “internal” in Table 4. On 8th weeks internal capacity without the collator is not sufficient to meet the load, so a large number of folds would have been assembled by an outside vendor. We use the vendor’s actual price schedule to determine the outsourcing cost that would have been incurred if the collator was not available. This figure is shown in the column labeled “external” in Table 4.

The internal savings are due to lower cost, in terms of man-hours, for using the collator instead of the manual assembly method, the increased capacity the collator provides, and increased decision flexibility the collator offers. The vast majority of the savings, however, are “external,” meaning that the collator allowed Jeppesen to bring much of the work in-house that was previously subcontracted out to a vendor. These “external” savings illustrate how the

Table 4. Savings and forecasts over the test period.

Week of	Cycle	Actual savings			Forecasted savings		
		Internal	External	Total	APP	WS	SPP
08-Jan-99	6	477		477	2657	5145	5865
15-Jan-99	7	3322		3322	1045	1085	1251
22-Jan-99	8	2959	134,728	137,687	28,537	56,365	92,815
29-Jan-99	1	22		22	1478	1527	1801
05-Feb-99	2	402		402	4076	6378	7175
12-Feb-99	3	34		34	1178	1261	1505
19-Feb-99	4	661		661	2572	4825	5411
26-Feb-99	5	234		234	1300	1365	1650
05-Mar-99	6	4178		4178	3537	5316	6101
12-Mar-99	7	2583		2583	1183	1200	1478
19-Mar-99	8	7425	89,679	97,104	28,888	60,521	100,345
26-Mar-99	1	0		0	1465	1538	1928
02-Apr-99	2	2021		2021	4086	5559	6558
09-Apr-99	3	0		0	1550	1630	2104
16-Apr-99	4	2647		2647	4912	6334	7593
23-Apr-99	5	529		529	1040	1127	1495
30-Apr-99	6	3854		3854	6406	7314	8973
07-May-99	7	3625		3625	1471	1531	2043
14-May-99	8	10,317	82,808	93,125	27,737	53,555	93,172
21-May-99	1	733		733	933	1010	1316
28-May-99	2	841		841	4915	5640	7133
04-Jun-99	3	19,141		19,141	751	836	1173
11-Jun-99	4	3130		3130	5409	5815	7522
18-Jun-99	5	7846		7846	725	815	1142
25-Jun-99	6	1582		1582	4005	4541	5732
02-Jul-99	7	7936		7936	1299	1331	1943
09-Jul-99	8	6994	68,529	75,523	27,542	51,905	90,722
16-Jul-99	1	14,961		14,961	938	986	1478
23-Jul-99	2	6758		6758	6310	6294	8197
30-Jul-99	3	8825		8825	778	809	1295
06-Aug-99	4	0		0	5901	5889	7946
13-Aug-99	5	7600		7600	718	755	1151
20-Aug-99	6	340		340	4891	4843	6647
27-Aug-99	7	1567		1567	1013	1108	1635
03-Sep-99	8	8615	72,560	81,175	27,394	69,216	115,584
10-Sep-99	1	19,706		19,706	1381	1464	2239
17-Sep-99	2	3909		3909	4968	4937	6989
24-Sep-99	3	18,482		18,482	1063	1110	1747
01-Oct-99	4	3415		3415	5317	5286	7584
08-Oct-99	5	10,481		10,481	1194	1230	1966
15-Oct-99	6	4874		4874	5552	5517	8164
22-Oct-99	7	5955		5955	1579	1601	2640
29-Oct-99	8	13,510	72,820	86,330	46,190	75,388	120,834
05-Nov-99	1	24,136		24,136	1344	1365	2130
12-Nov-99	2	24,180		24,180	3915	3842	5604
19-Nov-99	3	4117		4117	1160	1221	1939
26-Nov-99	4	3223		3223	5350	5321	7897
03-Dec-99	5	2439		2439	1160	1222	1961
10-Dec-99	6	361		361	5413	5372	8322
17-Dec-99	7	8513		8513	1160	1223	1984
24-Dec-99	8	13,542	78,761	92,303	46,984	84,510	132,955
31-Dec-99	1	13,716		13,716	1160	1224	2006
07-Jan-00	2	7623		7623	5681	5667	8836
Total savings to date:		323,862	599,885	**923,748**	356,556	587,723	**929,812**

collator increased Jeppesen's volume flexibility. In an internal memo dated 14 May 1999, Paul Vaughn, the assembly operations manager wrote: "The bottom line is that the Longford continues to meet expectations and we are saving dollars!"

The 53 weeks of data presented in Table 4 demonstrate that our new method for determining equipment benefit (SPP) is much more accurate in forecasting actual savings than the other two common methods. SPP estimates are determined by running 30 replications of Algorithm 1 for each week's problem, while APP and WS estimates are determined by solving corresponding deterministic problems. Using a matched pair t-test, we cannot reject with the null hypothesis that the SPP forecast and the actual data are the same (p-value = 0.4530). We can reject this null hypothesis at 1% level for both, APP and WS forecasts (p-value = 0.0006 for APP and 0.0016 for WS). The mean square error (MES) and bias measures tell the same story. MSE is 613,644,680 for APP, 250,796,163 for WS, and 172,665,224 for SPP. So clearly SPP provides the highest overall quality forecast. The average bias is −10,661 for APP, −6,252 for WS, and 216 for SPP, so both APP and WS significantly undervalue the collator, and SPP does not.

The APP model undervalues the Longford, particularly in high-volume 8th weeks in large part because of aggregation. Since APP does not contain inventory variables to provide links across periods, it is limited in its ability to properly model systems with limited capacity. But disaggregating the APP model does not entirely address the problem, because even the disaggregated model (WS) does not properly model decision flexibility. By increasing volume flexibility at Jeppesen, the Longford gave management the ability to postpone finalizing production plan for a few days, until demand becomes known with certainty. The reason our sampling-based optimization algorithm captures such benefits, and the other two methods do not, is that our method models how a production system responds to uncertainty.

7. CONCLUSIONS

We have demonstrated how some commonly used techniques for evaluating investment decisions in new production resources can severely underestimate the benefit of the resources when the resources provide capacity and decision flexibility. Simple analysis techniques, however, such as the one developed by Ramasesh and Jayakumar [21] provide a useful benchmark for analyzing the problem. We see that a simple aggregate method provides a conservative estimate and can systematically underestimate the benefit of flexible equipment, sometimes quite substantially. If a new technology can be justified using a simple conservative method, no further analysis is required. If, however, a conservative method cannot justify an investment in new flexible equipment, it may be worthwhile to consider our method to determine the benefit of flexible equipment more accurately. Naturally, our framework does not guarantee a perfect estimate for the benefit of flexibility either. No model can systematically account for human error, for example, so the SPP estimate is likely still a lower bound on the true collator benefit. It is, however, a more accurate estimate than the two simpler methods, as our results demonstrate.

We applied our method to a real investment problem faced by Jeppesen Sanderson, Inc., the major aviation information provider in the world. In July 1998, Jeppesen accepted our recommendation to invest in a new piece of equipment, the Longford folder collator, for its final assembly department. The Longford was built to Jeppesen specifications, and delivered in December 1998. Between 6 January 1999 and 7 January 2000, the Longford has been used consistently, including seven "8th weeks," generating savings in excess of $900,000. Prior to our work, the Jeppesen finance department had rejected the Longford proposal, estimating that it would take 6 years to pay for itself. As a result of our work they subsequently reversed their decision. As our model predicted, the Longford has paid for itself in less than 6 months. The management was so pleased with this outcome that subsequently, in July 2000 they purchased a second Longford collator.

ACKNOWLEDGMENT

The authors wish to thank the Associate Editor for his detailed and thoughtful comments that greatly helped us to improve this paper.

REFERENCES

[1] G. Azzone and V. Bertele, Measuring the economic effectiveness of flexible automation: a new approach, Int J Prod Res 27(5) (1989), 735–746.
[2] S. Benjaafar, T.L. Morin, and J. Talavage, The strategic value of flexibility in sequential decision making, Eur J Oper Res 82(3) (1995), 438–457.
[3] J.R. Birge and F. Louveaux, Introduction to stochastic programming, Springer-Verlag, New York, 1997.
[4] G.R. Bitran and D. Sarkar, On upper bounds of sequential stochastic production planning problems, Eur J Oper Res 34(1) (1988), 191–207.
[5] G.R. Bitran and H.H. Yanasse, Deterministic approximations to stochastic production problems, Oper Res 32(2) (1984), 999–1018.
[6] D.R. Carino, T. Kent, D.H. Myers, C. Stacy, M. Sylvanus, A.L. Turner, K. Watanabe, and W.T. Ziemba, The Russell-Yasuda Kasai model: An asset/liability model for a Japanese insurance company using multistage stochastic programming, Interfaces 24(1) (1994), 29–49.
[7] S. Das and P. Nagendra, Investigations into the impact of flexibility on manufacturing performance, Int J Prod Res 31(11) (1993), 2337–2354.
[8] X. de Groote, The flexibility of production processes: A general framework, Manage Sci 40(7) (1994), 933–945.
[9] R. Dembo, Scenario optimization, Ann Oper Res 30(1) (1991), 63–80.
[10] G.D. Eppen, R.K. Martin, and L. Schrage, A scenario approach to capacity planning, Oper Res 37(4) (1989), 517–527.
[11] C. Fine and R. Freund, Optimal investment in product-flexible manufacturing capacity, Manage Sci 36(4) (1990), 449–466.
[12] D. Gerwin, Manufacturing flexibility: A strategic perspective, Manage Sci 39(4) (1993), 395–410.
[13] E.M. Goldratt and J. Cox, The goal, North River Press, Great Barrington, MA, 1992.
[14] D. Gupta, Y. Gerchak, and J. Buzacott, The optimal mix of flexible and dedicated manufacturing capacities: Hedging against demand uncertainty, Int J Prod Econ 28(3) (1992), 309–319.
[15] Y.P. Gupta and S. Goyal, Flexibility of manufacturing systems: Concepts and measurements, Eur J Oper Res 43 (1989), 119–135.
[16] G. Infanger, Planning under uncertainty: Solving large-scale stochastic linear programs, The Scientific Press, Danvers, MA, 1993.
[17] W.C. Jordan and S.C. Graves, Principles on the benefits of manufacturing process flexibility, Manage Sci 41(4) (1995), 577–594.
[18] E. Katok, Planning manufacturing flexibility in an uncertain production environment: general framework and implementation, unpublished Ph.D. dissertation, The Pennsylvania State University, University Park, 1996.
[19] E. Katok, W. Tarantino, and R. Tiedeman, Flexibility planning and technology management at Jeppesen Sanderson, Inc., Interfaces 31(1) (2001), 7–29.
[20] J.M. Mulvey, G. Gould, and C. Morgan, An asset and liability management system for Towers Perrin-Tillinghast, Interfaces 30(1) (2000), 96–114.
[21] R.V. Ramasesh and M.D. Jayakumar, Inclusion of flexibility benefits in discounted cash flow analyses for investment evaluation: A simulation/optimization model, Eur J Oper Res 102 (1997), 124–141.
[22] K.S. Sethi and S.P. Sethi, Flexibility in manufacturing: A survey, Int J Flexible Manuf Syst 2 (1990), 289–328.

[23] D. Sinha and J. Wei, Stochastic analysis of flexible process choices, Eur J Oper Res 60(2) (1992), 183–199.
[24] J.E. Smith and K.F. McCardle, Valuing oil properties: Integrating options pricing and decision analysis approaches, Oper Res 46(2) (1998), 198–217.
[25] J.E. Smith and R.F. Nau, Valuing risky projects: Option pricing theory and decision analysis, Manage Sci 41(5) (1995), 795–816.
[26] F.F. Suarez, M.A. Cusumano, and C.H. Fine, An empirical study of manufacturing flexibility in printed circuit board assembly, Oper Res 44(1) (1996), 223–240.
[27] N.C. Suresh, Towards an integrated evaluation of flexible automation investments, Int J Prod Res 28(9) (1990), 309–319.
[28] W. Tarantino, Flexibility planning: A general framework and application for investment under uncertainty, unpublished Ph.D. dissertation, The Colorado School of Mines, Golden, 1999.
[29] D.M. Upton, The management of manufacturing flexibility, Cal Manage Rev (Winter 1994), 72–89.
[30] J.A. van Mieghem, Investment strategies for flexible resources, Manage Sci 44(8) (1998), 1071–1078.