APPLICATIONS OF REVENUE MANAGEMENT TO MANUFACTURING

Timothy S. Kniker
Mitchell H. Burman

Analytics Operations Engineering, Inc.
Boston, MA, U.S.A., 02109
(617)-939-0300
mburman@analytics-consulting.com
tkniker@analytics-consulting.com

***Abstract**: Using revenue management, many firms in service industries (airlines, hotels, rental-car firms) have documented huge increases in revenues with relatively small investments in technology. Revenue management can also be applied to manufacturing in a number of areas. In this paper, we consider a production facility with fixed capacity that must decide which incoming orders to accept. Unlike service industry cases, which involve maximizing revenue for discrete events (a flight, a night in a hotel, a specific advertising slot), profit in this scenario is maximized over a time continuum. This opens an untapped area of revenue-management research. To analyze this situation and provide insights, we model the manufacturer's environment as a Markov chain with rewards and policies, which is then solved using dynamic programming. We show that improved rewards (on the order of 10%) can occur by using revenue management.*

***Keywords**: Revenue Management, Manufacturing Order Acceptance, Dynamic Programming, Markov Chains*

1.0 Introduction

1.1 What is revenue management?

Revenue management (also called yield management, inventory control/management) is the business practice of determining which products and/or services to sell to which customers for potentially different prices in order to maximize some utility function (typically profit or revenue) [9]. The main advantage of using revenue management instead of traditional profitability analyses is that revenue management is a systematic method that quantitatively assesses the opportunity costs of all available decisions. The implementation of revenue management typically relies on classifying customers into different segments based on their willingness to pay, then controlling price and available inventory. For the remainder of the paper, we will assume that profit is the utility function being maximized.

To implement revenue management, a firm must perform five steps that are usually interrelated [10]:

- **Segmenting:** The firm must consider how its customers vary and how this affects the price they are willing to pay.
- **Pricing:** The firm must decide how much to charge for its services to each segment.
- **Restricting:** The firm must determine the characteristics of each segment that it can exploit, so customers do not pay a lower price than they are willing.
- **Forecasting:** The firm must estimate the future demand of each segment.
- **Allocation:** The firm must develop an allocation process that determines whether or not it will sell its services to a given customer in real-time.

A number of conditions must be met in order for revenue management to be effective in improving profit [8]:

- **Perishability:** The product or service that is being sold must be perishable. If it is not, then it can be saved and accumulated and matched exactly to demand when it exceeds capacity.
- **Customer segmentation:** There must be an ability to effectively segment customers. If customers are homogeneous, they will all pay the same price; thus, accepting a customer now is always better than reserving the service for a potential future customer.
- **Fixed capacity:** The firm's capacity must be relatively fixed, *i.e.*, it is either time-consuming or prohibitively expensive to add additional capacity. If capacity is easily added, then a simple analysis of comparing marginal revenue and marginal cost can be done, removing the need for revenue management.
- **Uncertain demand:** The demand must be variable and uncertain. If demand is well known ahead of time, the firm can statically determine which customers will be accepted.
- **Capacitated environment:** The resources required to complete all arriving jobs must exceed capacity. If the firm always has ample capacity, then a traditional profitability analysis can be done to determine whether an order should be accepted.

1.2 Order-acceptance problem in manufacturing

Many manufacturing firms face the daily decision of whether or not to accept incoming orders. Orders that come in usually have an associated value and/or a desired due date. It is up to the manufacturing firm either to accept the order as is or, at least, to begin negotiations with the customer by making suggestions for changes in either price or due date.

Figure 1 gives an abstract representation of the problem. Potential jobs (denoted by A, B, C, etc.) arrive at some time (denoted where the top arrows enter the time line). Each job has some reward and capacity requirement associated with it. Also, the due date of the job is specified (denoted where the bottom arrow leaves the time line). At the time of arrival, the firm must decide whether to accept the job. In Figure 1, accepted jobs are represented by black type, while declined jobs are represented by gray type. The middle bar represents the schedule of when accepted jobs were worked.

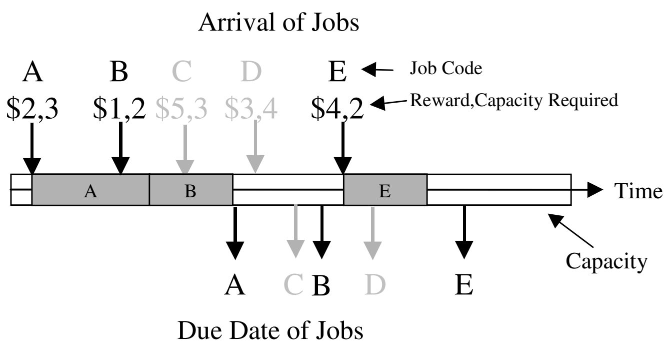

Figure 1. Representation of manufacturing order-acceptance problem.

Typically, in today's environment, a simple analysis (if any at all) is all that is done to determine a job's profitability. These analyses do not consider the profitability of potential future jobs. The authors have not found any systems in manufacturing that use systematic decision-making methods that measure opportunity cost.

There is a significant distinction between using revenue management in a manufacturing environment and using it in the service industries such as airlines, rental-car firms, and hotels. In service industries, revenue management is used for a *discrete event* (*e.g.*, a seat on a specific flight, a day in a car, a night of lodging). In most manufacturing environments, this will not be the case. Manufacturing firms are not selling a pre-specified block of time (*e.g.*, May 5th from 10:00 a.m. to 11:30 a.m.); but a batch of time (*e.g.,* 90 minutes) that the firm is free to schedule as long as the end-item is delivered to the customer by the due date.

2.0 Contributions of this Paper

The research contained in this paper makes the following contributions:

- It models revenue management over a time continuum rather than for a discrete event;
- It presents a dynamic programming formulation of this model to perform analysis; and
- It analyzes the results of the dynamic-programming formulation in order to provide some insights into the potential use of revenue management in a manufacturing environment.

3.0 Literature Review

The literature regarding revenue management is quite rich. The reader is directed to WEATHERFORD and BODILY[9] for a taxonomy of perishable-asset revenue management, BELOBABA[1] for a historical account of revenue management, and MCGILL and VAN

RYZIN[6] for a more recent overview and literature survey. This branch of literature focuses entirely on discrete-event revenue management.

To understand revenue maximization over a time continuum, we gain more insights by looking at the research on admission control to a queuing system. MILLER[7] and LIPPMAN[5] address the situation in which a manager may decide which arriving customers may be admitted to a queuing system. CALDENTEY[3], however, looks at using revenue management in the allocation of units of inventory to incoming orders. The decisions that a manager can make are: accepting an order, declining an order, or idling a machine. Caldentey shows that under heavy traffic conditions the problem can be approximated by a dynamic Brownian motion-control problem.

Depending on the environment in which a firm operates, a possible alternative to revenue management is capacity auctioning. This type of approach is good for industries in which customers do not need immediate responses to their inquiries. GALLIEN[4] presents a survey of the auction literature as it applies to manufacturing.

4.0 Manufacturing Order-Acceptance Problem

4.1 Assumptions

We use the following assumptions in our modeling:

- There is only a single bottleneck in the manufacturing process;
- The arrivals of orders are independent of one another;
- Orders come in for a specific price, volume, and due date;
- The firm must manufacture the order by the due date, *i.e.*, no backlogging; and
- There is no sequence-dependent setup required for the different job classes.

4.2 Basic model

We will model the manufacturing order-acceptance problem as a Markov chain with rewards and policy decisions, and then formulate the problem using dynamic programming. To create a finite-state space, we divide time into discrete periods, each of length Δt. We assume that only one order arrives in each time period. Let us assume that the orders are one of N classes.

4.2.1 Notation

For each order-class $n = 1, \ldots, N$:

R_n: the reward for accepting and manufacturing a class-n job.
U_n: the number of time periods required to manufacture a class-n job.
L_n: the lead-time of a class-n job, measured in time periods.
p_n: the probability that a class-n job arrives in a given time period.

Without a loss of generality, we can consider the likelihood of no job arriving in a given time period by constructing a dummy job class, n^*, which has $R_{n^*} = 0$, $L_{n^*} = 0$, $U_{n^*} = 0$, and

$$p_n = 1 - \sum_{m \neq n^*} p_m .$$

4.2.2 *Defining the state space*

Let

$$H = \max_{n=1,\ldots,N} L_n$$

H is the horizon of time that we must consider. We do not need to consider a longer horizon, because we will never receive a job that can be scheduled further out than H time periods.

We now make a key simplification to the problem. We will assume that the manufacturing of orders is done on a first-come, first-served basis, *i.e.*, once an order is accepted, it is put in a queue to be manufactured. Under this assumption, we need only consider the capacity remaining in the horizon, and therefore do not need to store information about specific jobs. This will decrease greatly the size of the state space and make it practical to solve more than trivially sized problems. However, the solution this gives us will yield a smaller reward per stage than if we allowed the manufacturing process to be scheduled.

The model that we have presented is similar to an admission-control queue, with two main differences. First, the research in LIPPMAN[5] looks at an *M/M/n/k* queue (n servers and queue capacity of k) with multiple classes, while this is equivalent to an *M/D/1* queue with infinite capacity. Second, his model considered a cost function that was linear in terms of the length of time a customer waited in queue. In our model, service must be finished by the due date, which is the equivalent of using a step function to model the cost of waiting in queue.

Let S be the state space. Each state, s, in the state space is defined by $\{h,n\}$, where h is the remaining available capacity over the horizon, and an order of class n has arrived during the time period (Figure 2).

4.2.3 *Policies and reward functions*

At each state, there are, at most, two possible decisions: accept the order, or decline the order. Let $Z(\{h,n\})$ be the set of policies available while in state $\{h,n\}$. If $U_n \leq h - (H-L_n)$, then we have enough capacity before the due date to manufacture the order and therefore may accept the order. Also, for convention, we will not allow "accepting" of a null order. Specifically,

$$Z(\{h,n\}) = \begin{cases} \{Accept, Decline\} & \text{if } U_n \leq h - (H - L_n) \text{ and } n \neq n^* \\ \{Decline\} & \text{otherwise} \end{cases}$$

Let us define a gain function, $G(\{h,n\},z)$ when in state $\{h,n\}$ and selecting policy z as:

$$G(\{h,n\}, Accept) = R_n \qquad G(\{h,n\}, Decline) = 0$$

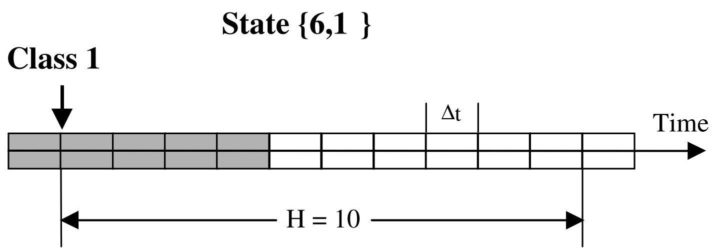

Job Class	Lead Time	Usage	Probability
1	10	3	0.2
2	10	3	0.4
*	0	0	0.4

Figure 2. Representation of a specific state, *{6,1}*.

4.2.4 Transition probabilities

If we are in state *{h,n}* and accept the order, the new capacity will be h - U_n + 1. The addition of one in this expression represents the fact that we gain one more unit of capacity at the end of the horizon as we transition to the new period (Figure 3). Therefore, if we accept the order, our transition probabilities are:

$$P^A[\{h,n\},\{h-U_n+1,m\}] = \begin{cases} p_m & \forall\{h,n\}\in S \quad and \quad m=1,\dots,N \\ 0 & otherwise \end{cases}$$

If we decline the order, then the new capacity will be h + 1, unless h = H, in which case the capacity remains h. Physically, this means we have let a unit of capacity pass by unused. The transition probabilities for declining an order are:

$$P^D[\{h,n\},\{h+1,m\}] = \begin{cases} p_m & \forall\{h,n\}\in \{S: h\neq H\} \quad and \quad m=1,\dots,N \\ 0 & otherwise \end{cases}$$

$$P^D[\{H,n\},\{H,m\}] = \begin{cases} p_m & \forall \quad m,n=1,\dots,N \\ 0 & otherwise \end{cases}$$

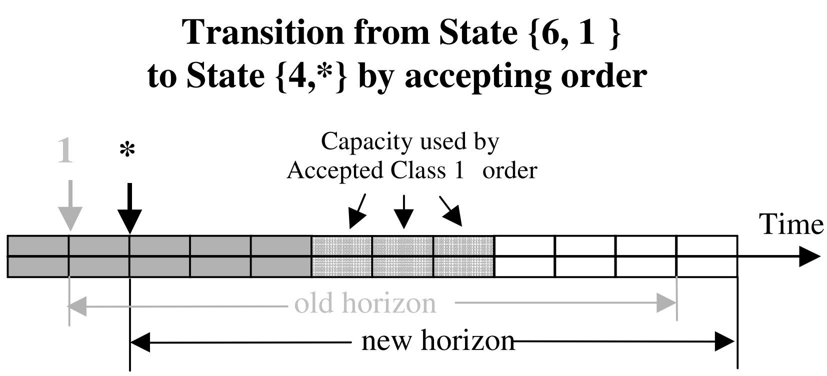

Figure 3. Example of a transition.

4.2.5 Dynamic programming formulation

This problem can be formulated as an average reward per stage, infinite-horizon dynamic program. The optimal average reward Q^* and a vector of state values $d^*(\{h,n\})$ satisfy Bellman's equation:

$$Q^* + d^*(\{h,n\}) = \max_{z \in Z(\{h,n\})} \left[G(\{h,n\}, z) + \sum_{\{h',n'\} \in S} P^z[\{h,n\},\{h'n'\}] d^*(\{h,n\}) \right], \quad \forall \{h,n\} \in S$$

We then use value iteration or policy iteration (BERTSEKAS [2]) to calculate the optimal value vector d^*, the optimal average reward Q^*, and the optimal policy vector z^*.

4.3 Analysis of optimal-reward results

We consider the situation where there are two classes of jobs. The relevant data for each job are as follows:

Job class	Reward	Usage	Lead time	Probability
1	6	30	200	0.02
2	4	30	200	0.04

Table 1. Job characteristics used in numerical analysis.

4.3.1 Comparison to first-come, first-served reward and maximum possible reward

To analyze the magnitude of the optimal reward, we compare it to two values. The first is the First-Come, First-Served (FCFS) reward, Q^{FCFS}. This is the average reward per stage when the policy of First Come, First Served is used to accept orders, *i.e.*, if capacity is available, the order is accepted.

To normalize the data, we will divide by an upper bound, Q^{UB}. This upper bound is obtained by sorting each job class by its reward/usage value. We then assume that all jobs of the highest classes are accepted until the capacity is used up. For example, from the data in Table 1, the maximum average reward per stage would be .17333. We would accept all class-1 orders (which use 60% of capacity) at an average reward of .20/period. Next we would accept 1/3 of all class-2 orders (which use the remaining 40% of capacity) at an average reward of .13333/period.

4.3.2 Effect of lead time on optimal reward

Figure 4 shows the graph of the normalized optimal reward, Q^*/Q^{UB}, (revenue management) versus the normalized FCFS reward, Q^{FCFS}/Q^{UB}, as a function of the lead-time length. When the lead-time is very short (in relation to the capacity requirement of a job), revenue management will not provide much benefit. This is because the revenue-management policy will be similar to the FCFS policy at these levels, mainly because the likelihood that another order will arrive in the short lead-time is small. As the lead-time increases, however, there is more capacity in the horizon that can be optimized. In this small example, revenue management provides a 10% benefit over the FCFS policy.

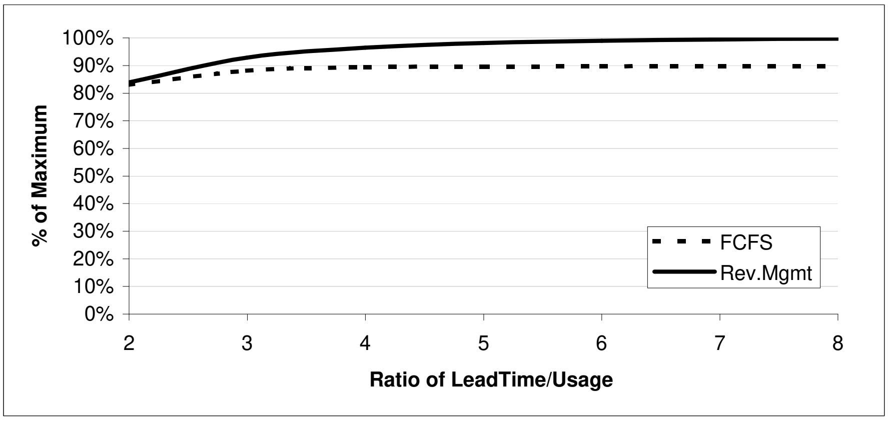

Figure 4. Normailzed rewards as a function of lead-time length.

4.3.3 Effect of reward dispersion on optimal reward

Figure 5 shows the graph of the normalized optimal reward versus the normalized FCFS reward as a function of the reward dispersion. A ratio of 1.0 represents the scenario when the two job classes are the same. As the reward ratio increases, the benefit of revenue management increases. As the ratio of rewards reaches 4, the benefit of revenue management increases to 25%.

.

4.3.4 Effect of load factor on optimal reward

Figure 6 shows the graph of the normalized optimal reward and the normalized FCFS reward as a function of the traffic parameter. The load factor is defined similarly to the traffic parameter in queuing theory. We increased the probability of arrival of both job classes in the same proportion.

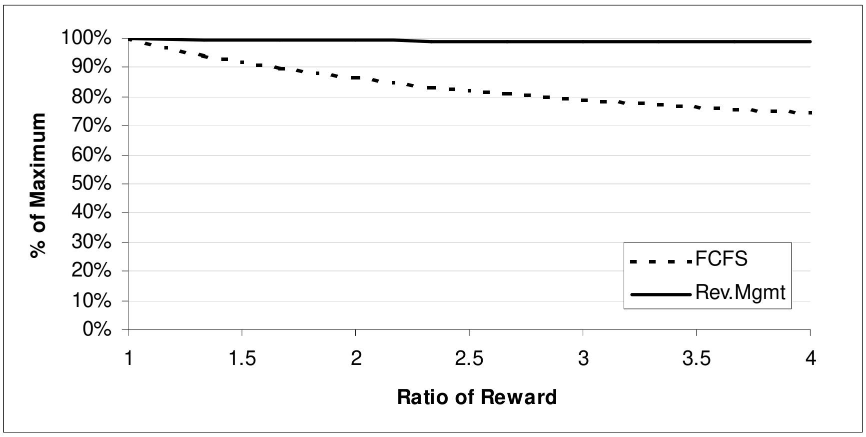

Figure 5. Normalized rewards as a function of reward dispersion.

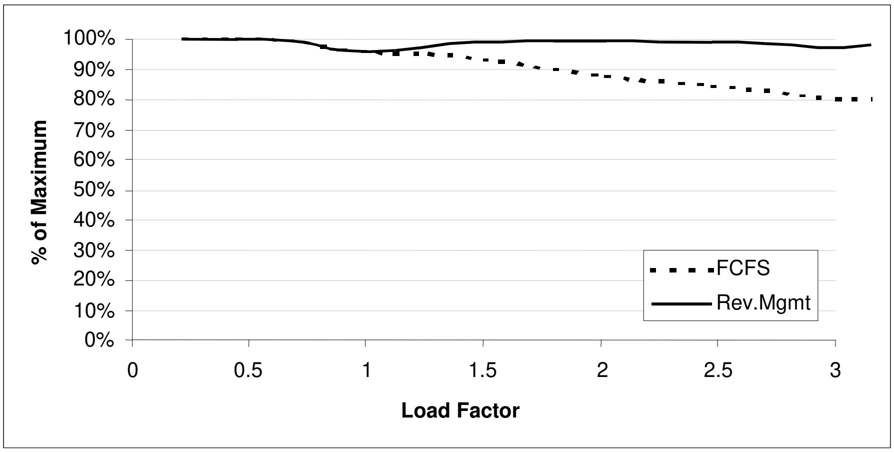

Figure 6. Normalized optimal reward as a function of the traffic parameter.

When the load factor is low, there is little-to-no benefit to using revenue management as we mentioned in Section 1.0. As the load factor increases to the level where the system becomes capacitated (>1.0), revenue management begins to outperform FCFS significantly. The dip that occurs around 1.0 is more a function of the change in the upper-bound value rather than a phenomenon that occurs when the system becomes capacitated.

4.4 Analysis of optimal policy

In Section 4.3, we investigate the optimal reward per stage, which gives us an idea of the quantitative benefit of revenue management. It is even more important to understand the actual process required to achieve these benefits. To do this, we will investigate the optimal policy and remark how it differs from the FCFS policy in future research.

5.0 Conclusion

5.1 Summary

Revenue management is a proven technology in many industries. Some ideas that have been developed for the service industry can be applied to situations in manufacturing. Previous research has examined how inventory stocks can be allocated to customers. In this paper, we examine how manufacturing capacity itself is allocated. We model this situation as a Markov chain with rewards and policy decisions, which we then solve by using dynamic programming.

We show that significant benefit can be gained by using revenue management beyond a simple FCFS approach to manufacturing-order acceptance. Further research should be done to verify this with real-world data. Beyond the quantitative benefits described, numerous benefits are possible, from an organizational and customer-service perspective. In the authors' work in manufacturing consulting, we have interviewed numerous managers who wrestle with order-acceptance problems on a daily basis. Typically, they make their decision by either "gut feel" or a slow, cumbersome organizational structure that pushes decisions to higher levels. Revenue management can eliminate slow service time with an objective decision-making system that is not vulnerable to turnover of a company's key personnel.

5.2 Extensions and future research

Numerous paths of future research stemming from this topic exist. First, we make a key simplification by assuming that each job is manufactured on a first-come, first-served basis. Allowing the scheduling of jobs will greatly increase the complexity of the problem. A dynamic-programming formulation will become intractable due to the amount of state information that needs to be maintained.

Another obvious action is to consider more job classes in our numerical analysis. The number of job classes parallels the number of different fare classes in the airline revenue management problem. Also, we can extend revenue management to a manufacturing process with more than one bottleneck process. This also has a parallel in airline revenue management. Most of the research in the late 1990s has focused on network revenue management, specifically how do you modify "single-leg" revenue management when a passenger wants an itinerary that uses more than one flight leg in the network.

6.0 References

[1] BELOBABA, PETER, "Air Travel Demand and Airline Seat Inventory Control," Ph.D. Thesis, Flight Transportation Laboratory, Massachusetts Institute of Technology, 1987.

[2] BERTSEKAS, DIMITRI P., *Dynamic Programming and Optimal Control, Volume 1*, Athena Scientific, 1995.

[3] CALDENTEY, RENE, "Revenue Management of a Make-to-Stock Queue," Working Notes, Sloan School of Management, Massachusetts Institute of Technology, December 2000.

[4] GALLIEN, JEREMIE, "Design and Analysis of a Smart Market for Industrial Procurement," Working paper submitted to **Management Science**, 2000.

[5] LIPPMAN, S. "Applying a New Device in the Optimization of Exponential Queuing Systems," **Operations Research**, Vol. *23*, 1975.

[6] MCGILL, JEFFREY I. and GARRETT VAN RYZIN, "Revenue Management: Research Overview and Prospects," **Transportation Science**, *33-2*, 1999.

[7] MILLER B., "A Queuing Reward System with Several Customer Classes," **Management Science**, *Vol. 16-3*, 234-245.

[8] PROS Revenue Management web-site, http://www.prosrm.com/.

[9] WEATHERFORD, L.R. AND BODILY, S.E. "A Taxonomy and Research Overview of Perishable-Asset Revenue Management: Yield Management, Overbooking, and Pricing," **Operations Research**, *Vol. 40*, 831-844.

[10] WONG, J.T., F.S. KOPPELMAN, and M.S. DASKIN, "Flexible Assignment Approach to Itinerary Seat Allocation," **Transportation Research**, *27B*, 1993.