European Journal of Operational Research 247 (2015) 472–487

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Prioritizing regular demand while reserving capacity for emergency demand

Xinxin Hu[a], Ying Li[b,*], Eunshin Byon[c], F Barry Lawrence[a]

[a] College of Engineering, Texas A&M University, College Station, TX 77843, United States
[b] Mays Business School, Texas A&M University, College Station, TX 77843, United States
[c] College of Engineering, The University of Michigan, Ann Arbor, MI 48109, United States

ARTICLE INFO

Article history:
Received 8 January 2014
Accepted 3 June 2015
Available online 15 June 2015

Keywords:
Capacity allocation
Capacity reservation
Make-to-order
Demand prioritization

ABSTRACT

This research is motivated by the capacity allocation problem at a major provider of customized products to the oil and gas drilling industry. We formulate a finite-horizon, discrete-time, dynamic programming model in which a firm decides how to reserve capacity for emergency demand and how to prioritize two classes of regular demand. While regular demand can be backlogged, emergency demand will be lost if not fulfilled within the period of its arrival. Since backlogging cost accumulates over time, we find it optimal for the firm to adopt a dynamic prioritization policy that evaluates the priorities of different classes of regular demand every period. The optimal prioritization involves metrics that measure backlogging losses from various perspectives. We fully characterize the firm's optimal prioritization and reservation policy. Those characterizations shed light on managerial insights.

1. Introduction

Capacity management is often a challenging task faced by many production managers. This research is motivated by the experience of a major provider of customized products to the oil and gas drilling industry. In recent years, demand for the company's products is fueled by the booming oil market. The rapid growth, however, brings challenges to capacity management. One of the most salient problems is backlogging. The company's records show that customers could wait as long as a year for their orders to be satisfied. As much as the company wants to take full advantage of the current demand increase, the executives of the company are hesitant to invest in capacity expansion for fear of market fluctuations. They are more interested in how to best utilize their existing capacity.

The majority of the company's demand is processed in a first-come-first-serve fashion. Some executives of the company question the validity of this approach because different classes of demand do not have the same gross margins. On the other hand, waiting costs are also a concern because the company hopes to maintain long-term relationships with its clients. Classifying demand based on both revenue and waiting cost is, therefore, necessary. Further complicating the demand prioritization problem is the existence of emergency demand. Unexpected events occurring in the oil and gas field result in emergency orders of specialized equipments. In light of the long-term relationship with critical clients, the company wants to satisfy emergency demand as much as possible, even at the expense of delaying regular demand. Given the frequency of emergency demand, production managers find it necessary to reserve capacity for the uncertain emergency demand. Hence, there are two tasks in the company's capacity management problem. One is reserving capacity for unknown emergency demand. The other is splitting the leftover capacity between different classes of regular demand.

We formulate a discrete-time dynamic-programming model that reflects the two tasks faced by the company and fully characterize the optimal solution. In order to characterize the optimal policy, we develop various metrics that measure the potential losses resulting from backlogging regular demand. Comparisons among these metrics allow us to prioritize regular demand and determine the level of capacity reservation for emergency demand. We believe the optimal solution to our model provides guidance to the company's capacity management.

Although our research was motivated by challenges rising from the oil and gas drilling field, our model and solution approach can be applied to general industrial settings with customized products and capacity constraints. As pointed out by Rahman and Seliger (2013) and Mourtzis, Doukas, Psarommatis, Giannoulis, and Michalos (2014), demand for customized products has been

* Corresponding author. Tel.: +1 979 845 2392; fax: +1 979 845 5653.
E-mail addresses: hu.xinxin@hotmail.com (X. Hu), yli@mays.tamu.edu (Y. Li), ebyon@umich.edu (E. Byon), lawrence@entc.tamu.edu (F. Barry Lawrence).

http://dx.doi.org/10.1016/j.ejor.2015.06.020
0377-2217/

increasing over the years and customization has become a widespread paradigm in manufacturing technology. Alexopoulos, Papakostas, Mourtzis, and Chryssolouris (2011) studies the performance of manufacturing systems with a focus on lifecycle, recognizing the existence of capacity constraints. deKoster, Le-Duc, and Roodbergen (2007) examines order picking systems, which feature both customized orders and capacity constraints.

The rest of the paper unfolds as follows. Section 2 reviews relevant literature. Section 3 explains the problem setting and formulates the model. We present preliminary analysis on demand types and capacity reservation, introduce the concept of dynamic loss, and discuss our solution approach in Section 4. Section 5 describes the optimal policy. Section 6 summarizes our numerical studies, which develop and compare a few heuristics to the optimal policy. Section 7 concludes the main text of the paper. Appendix A presents technical statements that are skipped in the main text for readability. Proofs of technical statements are included in a supplementary document.

2. Relevant literature

Our research is, essentially, about capacity allocation among multiple demand classes. This topic has been studied in a variety of contexts. We review, here, the relevant literature and explain how our work differs.

One stream of research considers demand that is always backlogged if unsatisfied. The same product is made and sold to multiple classes of customers. At any given time, one needs to decide whether the facility should produce or not. Once a product is made, one needs to decide whether it should be allocated to on-hand inventory, or any class of customers. An optimal solution should strike a balance between holding cost and various levels of backlogging costs. Ha (1997) started this line of research in a setting with two demand classes. de Vericourt, Karaesmen, and Dallery (2002), Huang and Iravani (2008), and Gayon, de Vericourt, and Karaesmen (2009) look at extensions or variations of Ha's (1997) work. Iravani, Liu, and Simchi-Levi (2012) look at a manufacturer who accepts all demand for a generic product with make-to-stock system and can accept or reject demand for customized products with make-to-order system. Both types of demand can be backlogged after being accepted in the setting of Iravani et al. (2012). One of the key differences between this line of research and ours is our consideration of emergency demand, which cannot be backlogged. The uncertain arrival of emergency demand during a period in our setting makes capacity reservation necessary in addition to allocating demand between two regular demand classes.

Another stream of research analyzes production and allocation decisions when a firm faces two classes of demand: contractual and transactional. The contractual demand must be met immediately while the firm chooses to accept or reject transactional demand. Carr and Duenyas (2000) formulate a continuous-time, make-to-stock model in which transactional orders are either accepted or rejected upon arrival. They show that the firm's optimal acceptance and production decisions are determined by an acceptance threshold curve and a production threshold curve. Frank, Zhang, and Duenyas (2003) adopt a discrete-time approach. Their contractual orders are deterministic and their transactional orders are stochastic. Unsatisfied transactional orders are lost. Frank et al. (2003) find that their optimal policy can be characterized by a modified and much more complicated (s, S) policy. Gupta and Wang (2007) formulate a discrete-time model and allow transactional demand to be satisfied within a fixed number of periods without penalty. When both contractual and transactional demand are make-to-order, Gupta and Wang (2007) show that the optimal acceptance decision is a threshold policy with the threshold being determined by the firm's capacity and the current period's contractual demand. The contractual demand in this research stream differs from the regular demand in our research because contractual demand is fulfilled immediately or within a fixed number of periods without penalty. Our regular demand, in contrast, can be backlogged throughout the whole planning horizon and with accumulating penalties.

Wang, Liang, Sethi, and Yan (2014) consider two groups of customers: short lead-time ones who demand the product immediately and long lead-time ones whose orders can be fulfilled either immediately or in the next cycle. Backlogging costs accumulate for unsatisfied demand. The supplier needs to determine how to allocate on-hand inventory between the two groups of customers and how to allocate replenishment inventory among backlogged orders. The paper focuses on three priority rules and derive inventory commitment policies based on those rules.

Gao, Xu, and Ball (2012) consider pseudo orders that become actual orders after being confirmed and use a Markov chain model as a short-term forecast for pseudo orders. The manufacturer decides whether to accept or reject a confirmed order. Accepted orders must be fulfilled during a fixed time window and productions are constrained by both capacity and component inventory availability. Orders from different classes bring in different revenues.

Capacity allocation research in medical settings often deals with demand uncertainty (Gerchak, Gupta, & Henig, 1996; Patrick, Puterman, & Queyranne, 2008; Green, Savin, & Wang, 2006; Patrick & Puterman, 2007; Gupta & Wang, 2008; Min & Yih, 2010; Dobson, Hasija, & Pinker, 2011; Geng & Xie, 2012 etc.). Green (2012) discusses how operations research can improve healthcare delivery. Chen and Robinson (2010) and May, Spangler, Strum, and Vargas (2011) examine capacity allocations when scheduling appointments or surgeries, respectively. Ramirez-Nafarrate, Hafizoglu, Gel, and Fowler (2014) study ambulance diversion policies in order to minimize patients' average waiting times. Gerchak et al. (1996) is the closest to our research. They study how to share operating room capacity between elective surgeries and emergency surgeries. Before learning the amount of emergency demand arriving on any given day, the scheduler must decide how many elective surgeries are admitted for operation on the current day. Delayed elective surgeries result in postponed revenue collection and waiting costs. The elective surgeries are analogous to our regular demand but have just a single class. With a single class, the prioritization problem goes away.

Revenue management is an area that naturally deals with capacity allocation among multiple demand classes. It is often studied in contexts including airlines, hotels, and car rentals. McGill and van Ryzin (1999) and Talluri and van Ryzin (2004) provide detailed reviews. A key difference between our work and revenue management literature is the lost sales of unsatisfied demand in revenue management settings. Backlogged regular demand is more appropriate in our model because the company that motivates this research rarely turns away any demand.

A key element of our model is the existence of customized products. Given its popularity, customization has been the focus of a series of research in manufacturing settings. Mourtzis, Doukas, and Psarommatis (2013) proposes a genetic algorithm for mass customization. Mourtzis et al. (2014) examines a web-based platform that facilitates customization. Adaptability or flexibility is crucial to the manufacturers of customized products. Papakostas and Mourtzis (2007), Alexopoulos, Mourtzis, Papakostas, and Chryssolouris (2007), and Alexopoulos et al. (2011) examine the evaluations of system flexibility or adaptability. Makris and Chryssolouris (2010) incorporates demand uncertainty in manufacturing planning and applies their Bayesian model to the automotive industry. This series of research studies the design and evaluation of manufacturing systems, which often consist of multiple stages. Our model analyzes the tension between demand and supply, assuming away the details regarding manufacturing stages and layouts.