A Stochastic Dynamic Programming Approach to Revenue Management in a Make-to-Stock Production System

Rainer Quante, Moritz Fleischmann, and Herbert Meyr

ERIM REPORT SERIES *RESEARCH IN MANAGEMENT*	
ERIM Report Series reference number	ERS-2009-015-LIS
Publication	March 2009
Number of pages	32
Persistent paper URL	http://hdl.handle.net/1765/15183
Email address corresponding author	mfleischmann@rsm.nl
Address	Erasmus Research Institute of Management (ERIM) RSM Erasmus University / Erasmus School of Economics Erasmus Universiteit Rotterdam P.O.Box 1738 3000 DR Rotterdam, The Netherlands Phone: + 31 10 408 1182 Fax: + 31 10 408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES
RESEARCH IN MANAGEMENT

ABSTRACT AND KEYWORDS	
Abstract	In this paper, we consider a make-to-stock production system with known exogenous replenishments and multiple customer classes. The objective is to maximize profit over the planning horizon by deciding whether to accept or reject a given order, in anticipation of more profitable future orders. What distinguishes this setup from classical airline revenue management problems is the explicit consideration of past and future replenishments and the integration of inventory holding and backlogging costs. If stock is on-hand, orders can be fulfilled immediately, backlogged or rejected. In shortage situations, orders can be either rejected or backlogged to be fulfilled from future arriving supply. The described decision problem occurs in many practical settings, notably in make-to-stock production systems, in which production planning is performed on a mid-term level, based on aggregated demand forecasts. In the short term, acceptance decisions about incoming orders are then made according to stock on-hand and scheduled production quantities. We model this problem as a stochastic dynamic program and characterize its optimal policy. It turns out that the optimal fulfillment policy has a relatively simple structure and is easy to implement. We evaluate this policy numerically and find that it systematically outperforms common current fulfillment policies, such as first-come-first-served and deterministic optimization.
Free Keywords	revenue management, order fulfillment, advanced planning systems, make-to-stock production
Availability	
Classifications	

A Stochastic Dynamic Programming Approach to Revenue Management in a Make-to-Stock Production System

Rainer Quante
Austrian Airlines (TDLC), Vienna Airport, Austria, rainer.quante@austrian.com

Moritz Fleischmann
Rotterdam School of Management, Erasmus University, Rotterdam, The Netherlands, mfleischmann@rsm.nl

Herbert Meyr
Department of Production and Supply Chain Management, Technical University of Darmstadt, Germany, meyr@bwl.tu-darmstadt.de

In this paper, we consider a make-to-stock production system with known exogenous replenishments and multiple customer classes. The objective is to maximize profit over the planning horizon by deciding whether to accept or reject a given order, in anticipation of more profitable future orders. What distinguishes this setup from classical airline revenue management problems is the explicit consideration of past and future replenishments and the integration of inventory holding and backlogging costs. If stock is on-hand, orders can be fulfilled immediately, backlogged or rejected. In shortage situations, orders can be either rejected or backlogged to be fulfilled from future arriving supply.

The described decision problem occurs in many practical settings, notably in make-to-stock production systems, in which production planning is performed on a mid-term level, based on aggregated demand forecasts. In the short term, acceptance decisions about incoming orders are then made according to stock on-hand and scheduled production quantities.

We model this problem as a stochastic dynamic program and characterize its optimal policy. It turns out that the optimal fulfillment policy has a relatively simple structure and is easy to implement. We evaluate this policy numerically and find that it systematically outperforms common current fulfillment policies, such as first-come-first-served and deterministic optimization.

1. Introduction

Successful revenue management (RM) applications abound in the service industries, including the well-known examples of the airline, hotel, and car rental businesses. Recently, practitioners as well as researchers have increasingly been exploring ways to transfer revenue management concepts to manufacturing environments (Harris and Pinder 1995, Swann 1999, Gupta and Wang 2007, Arslan et al. 2007). The core idea is that customer differentiation is beneficial also in manufacturing industries even though they differ markedly from service industries, with respect to, for example, cost structures, the role of inventories, and their replenishment. Even a high service level of, say, 98% still implies that 2% of all customer orders are not fulfilled as requested and it does matter which customers are affected by this. In a manufacturing environment, customers may not only show different willingness to pay (e.g. contract vs. spot market) or different strategic importance (e.g. internal vs. external customers), but also entail different costs (e.g. for production and storage in or distribution to different regions). Therefore, in a manufacturing context, profit maximization seems a more appropriate criterion for customer differentiation than pure revenue maximization (Quante et al. 2009).

Despite the modern trend to establish "lean" processes, many companies cannot avoid significant setup cycles or long replenishment lead times, for cost efficiency reasons. In order to remain competitive and be able to offer short delivery times to their customers, they have to build up finished goods inventory in the mid term. In a make-to-stock (MTS) environment like this, production planning has to be based on forecast demand and cannot be adjusted to short-term demand fluctuations. However, what is still relatively flexible is the allocation of the available inventory to incoming customer orders. This resembles the situation in traditional perishable-asset revenue management, where a fixed amount of capacity is sold to multiple customer classes within a given time interval. However, the situations are not identical. The main difference is that the allocation decision in MTS manufacturing concerns finished goods inventory rather than capacity. Unlike capacity, inventory is storable, at least to some extent, and is replenished at certain times.

In addition, there is also a significant methodological divide between revenue management and manufacturing. Revenue management is strongly rooted in the domain of stochastic optimization and uses probability distributions for assessing opportunity costs. Many manufacturing companies, on the other hand, use so-called advanced planning systems (APS) in their demand and supply management (Stadtler and Kilger 2008). These systems are essentially deterministic, relying on linear programming (LP) technology.

Focusing on short-term demand management, a demand fulfillment module usually handles customer orders in APS. In an MTS situation, this module receives input information from a mid-term master planning and a short-term production planning module, which determine production quantities. These anonymous and usually aggregate quantities are then allocated to the various customer segments by simple allocation rules — for example according to priority rankings, with respect to pre-defined fixed shares, or proportionally to the customers' original forecasts (Kilger and Meyr 2008). Given the resulting "(allocated) available-to-promise"(ATP) quantities, short-term order acceptance decisions then rely on — again simple — ATP consumption rules, such as first-come-first-served, or on input from experienced human planners. For example, Kilger (2008) describes a computer manufacturing setting, in which the company's different sales regions constitute the customer segments, the planned production quantities of the mid–term master plan are allocated to these segments proportionally to the regions' demand forecasts, and two different ATP consumption rules are applied, depending on whether the customer asks for a standard or a customized product.

Our goal in this paper is to develop an RM approach to short-term demand management in MTS production, thereby complementing the current practice reflected in APS. We summarize our problem setting as follows: We consider a make-to-stock production system with exogenously given deterministic replenishments and heterogeneous customers. The objective is to maximize profit over the planning horizon by deciding whether to accept or reject a given order, in anticipation of more profitable future orders. What distinguishes this setup from classical revenue management approaches is the explicit consideration of past and future replenishments and the integration of inventory holding and backlogging costs. If stock is on-hand, orders can be fulfilled immediately, backlogged or rejected. During stockout periods, orders can be either rejected or backlogged to be fulfilled from future arriving supply.

We present a stochastic dynamic programming formulation of this problem and identify structural properties that characterize its optimal policy. It turns out that the optimal policy has an intuitive structure, which makes it easy to implement. In a series of simulation experiments, we compare the performance of this optimal policy with a first-come-first-served rule, as well as with a deterministic allocation policy.

To summarize, the paper makes the following contributions:

• We present a model that reflects important characteristics of order fulfillment in MTS production environments, namely customer heterogeneity, limited short-term supply flexibility, and short-term allocation flexibility. Previous literature has not addressed the interplay between these factors.

• We prove structural properties of this model and derive an optimal demand fulfillment policy. The results integrate the hitherto disconnected streams of revenue management and APS-based order fulfillment in MTS production.

• We show the superiority of a stochastic RM-based fulfillment policy in MTS production over both a first-come-first-served approach and a deterministic allocation approach in numerical examples and illustrate the impact of key parameters on the potential benefits.

The paper is structured as follows. We start with a literature review in §2 and motivate the necessity of our research. In §3, we formulate a stochastic order fulfillment model for MTS production. In §4 we analyze the mathematical structure of this model and derive an optimal fulfillment policy. §5 presents a numerical study that compares the optimal fulfillment policy with two benchmark policies and conducts a sensitivity analysis of key model parameters. §6 summarizes our results and discusses potential extensions and directions for further research.

2. Literature review

The field of revenue management (RM), or demand management more generally, has produced a rich body of literature, encompassing multiple distinct streams. What is common to all of them is the issue of making optimal use of a given resource. Talluri and van Ryzin (2004) structure this field by distinguishing quantity-based RM and price-based RM. Quantity-based RM concerns the allocation of scarce capacity to different customer (fare) classes. Price-based RM addresses optimal price paths over time.

In our analysis, we assume prices to be exogenous. Therefore, our problem setting falls within the domain of quantity-based RM. In a previous paper, we have refined the structuring of this field through a second dimension, namely the type of replenishment considered (Quante et al. 2009). Underlying models consider either no replenishments, exogenous replenishments, or endogenous replenishments, leading to the three cases displayed in Figure 1. First, *traditional revenue management* (TRM) deals with the rationing of perishable capacity, and replenishments are irrelevant. Second, *allocated available-to-promise* (aATP) models concern the allocation of inventories and take into account exogenously scheduled replenishment orders. Third, *inventory rationing* (IR) models consider allocation and replenishment decisions simultaneously. In what follows, we briefly review literature in each of these three categories, focusing on manufacturing applications. Note that our previously introduced problem setting falls into the second stream, since replenishments are exogenous. We therefore discuss the literature in this stream in more detail and complement it with a briefer discussion of the remaining streams.

Figure 1 Consideration of replenishment in quantity-based demand management models (Quante et al. 2009)

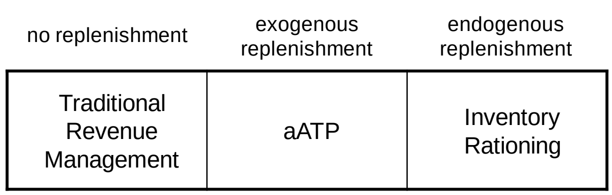

2.1. Traditional revenue management

By far the largest body of literature falls into the category *TRM*. Overviews of applications in different industries, including manufacturing, are provided by Swann (1999), Kimms and Klein (2005) and Chiang et al. (2007). Applications in manufacturing largely focus on make-to-order (MTO) and assemble-to-order (ATO) production environments. This is intuitive due to the correspondence between *production capacities* and *perishable assets*, which are the building blocks of RM in the service industries (Weatherford and Bodily 1992). Harris and Pinder (1995) show an RM application in an ATO setting. For MTO manufacturing, Van Slyke and Young (2000), Kleywegt and Papastavrou (2001) and Spengler et al. (2007) propose knapsack formulations in order to select or reject incoming orders. All of these models assume fixed capacity and do not allow backlogging.

2.2. Allocated available-to-promise

Models of the *aATP* stream, are rooted in the area of advanced planning systems (APS). Available-to-promise (ATP) calculations provide a time-phased view of available supply and serve as a basis for order promising in response to incoming customer requests. Stadtler and Kilger (2008) and Fleischmann and Meyr (2003) discuss ATP and its functions within APS. Kilger and Meyr (2008) focus on demand fulfillment in APS. Other reviews of ATP include Chen et al. (2001) and Ball et al. (2004).

Simple order promising does not differentiate between different customer segments and uses ATP quantities on a first-come-first-served basis. Common approaches for selecting the most profitable orders include *batch promising* or *real-time aATP*. Batch-promising models collect all customer orders of a pre-defined time period (e.g. a day) and promise due dates simultaneously in a "batch" (Chen et al. 2001, 2002). The main drawback of batch promising is the long response time to customer requests. More refined "hybrid" models handle this issue through a two-step approach: First answer the request by an approximate due date and then refine this date in a subsequent batch process. Real-time aATP approaches use another two-step process to answer customer requests immediately. A first step allocates on-hand inventory and future incoming supply to delivery periods and customer segments by means of rather simple allocation rules. In the second step, these (*allocated ATP*) quotas are consumed by incoming customer orders, either first-come-first-served or on the basis of user-defined consumption rules. Fischer (2001) and Pibernik (2006) compare these consumption policies with heuristic and LP-based batch models for practical cases in the lighting industry and the pharmaceutical industry, respectively. While Fischer finds general advantages of the batch approach if customer response times can be neglected, Pibernik recommends switching to batch mode only if shortages are foreseeable. However, both approaches do not focus on the allocation step and therefore only use very simple allocation rules.

APS customer segments typically rely on customer location, e.g. on ZIP codes or countries. Associated revenues or costs are hardly taken into account when building the segments. Meyr (2008) proposes a clustering method that builds customer classes on the basis of profits and additional strategic aspects. Meyr (2009) applies this method to identify better discriminating customer classes, which allow replacing the simple allocation rules of aATP models by a profit-based LP approach. The LP uses sales prices and backlogging and holding costs, as well as point estimates of demand. A numerical simulation study shows this approach to significantly improve profitability if future demand is known. The strength of this model lies in its ability to better differentiate between customers while at the same time allowing for an immediate response to incoming customer requests. Its major drawback is the disregard of demand uncertainty. In the present paper, we extend the model of Meyr (2009) by modeling stochastic demand and applying RM concepts in the allocation decision.

We are aware of only a single paper that considers stochastic demand and exogenous replenishments (Pibernik and Yadav 2009). The authors propose a model that differentiates between two customer classes under a service-level constraint. The model is closely linked to our setting, but with several key distinctions: Their model uses a service-level target whereas ours maximizes expected profit; their model is limited to two customer classes and does not allow backlogging, both of which we do allow.

2.3. Inventory rationing

The third stream of related literature encompasses *IR* models, which include endogenous replenishment decisions. General reviews of this literature can be found in Kleijn and Dekker (1998) and in Teunter and Klein Haneveld (2008).

Taking a manufacturing perspective, Defregger and Kuhn (2007) propose a Markov decision model for an MTO setting, which also includes a finished-goods inventory, used to serve high-margin orders with short lead times. Orders differ with respect to profit and maximum acceptable lead time. Similarly, Gupta and Wang (2007) study a setting with two customer classes, namely transactional and more valuable contractual orders. The manufacturer can choose which transactional orders to fulfill but has to meet the demand of contractual orders within a short lead time. The authors consider both MTS and MTO production of contractual orders.

Further MTS production models are proposed by Ha (1997) and Arslan et al. (2007). Ha (1997) presents a production control and stock rationing problem for two customer classes, which allows backlogging. Limited replenishment capacity is modeled as an M/M/1 queuing system. Decisions are, (1) whether to produce or not, (2) whether to immediately fulfill or backlog a low-priority order, and (3) whether to clear open backorders if stock is available. In this setting, a base stock policy is optimal for production and the optimal rationing policy is decreasing in the number of backorders. Arslan et al. (2007) aim to meet service targets for multiple customer classes with minimal inventory, assuming that replenishment lead times are deterministic and replenishment quantities are unlimited. They develop a heuristic critical level, continuous review policy, which they show to be near optimal.

To conclude, we summarize the positioning of the present paper in the literature. Our problem setting is rooted in the context of advanced planning systems and considers a make-to-stock system. We take inventory replenishments into account and acknowledge that they cannot be changed in the short run. Our analysis therefore falls into the $aATP$ stream of the quantity-based revenue management literature. Unlike previous models in this stream, we consider stochastic demand, a general number of customer classes, and allow backlogging of demand in the case of shortages.

3. Model formulation

We consider a make-to-stock manufacturing system facing stochastic demand from multiple customer classes, which differ in their unit revenues. Scheduled inventory replenishments are known.

Given this information, the manufacturer decides for each order whether to satisfy it from stock, backorder it at a penalty cost, or reject it. The objective is to maximize the expected profit over a finite planning horizon, taking into account sales revenues, inventory holding costs, and backorder penalties. What makes this problem complex is that unlike in traditional revenue management one has to decide not only whether or not to satisfy a given order but also which supply to use, with each supply alternative yielding a different profit. We make the following assumptions to model this situation.

Assumption 1. *Orders from a given customer class follow a compound Poisson process. The order processes of different classes are mutually independent, and they are independent of the available supply.*

The Poisson assumption is common in many RM models, specifically in so-called dynamic demand models (see Talluri and van Ryzin (2004)). As Lautenbacher and Stidham (1999), we allow non-unit order sizes, which appears appropriate in a manufacturing environment. In our analysis, we discretize the planning horizon in such a way that the probability of receiving multiple orders within a single period is negligible. Let T denote the length of the planning horizon and t the period index. Moreover, let $c = 1, \ldots C$ identify the different customer classes.

Assumption 2. *Inventory replenishments are exogenous and known.*

This assumption reflects the APS planning hierarchy. Inventory replenishments are determined by mid-term and short-term production planning and then serve as input for order promising decisions. Let x_i be the available supply arriving at the beginning of period i, $i = 1, \ldots, T$, and let $\bar{x} = (x_1, \ldots, x_T)$ be the vector of all of these replenishments. Note that at time t, x_i corresponds with inventory on-hand if $i \leq t$ and with a future scheduled replenishment otherwise. In APS terminology, $\bar{x}$ denotes the ATP quantities.

Assumption 3. *Order due dates are equal to order arrival times, but orders can be backlogged at a price discount.*

This assumption reflects the MTS context. Customers expect immediate delivery, in principle. Late deliveries are only acceptable at a price discount. Let r_c denote the unit revenue from satisfying an order of class c from stock. Delaying an order gives rise to unit backorder costs b per period, including the customer discount. Analogously, holding costs h are incurred for all units of inventory on hand at the end of a period. Note that unit backorder and holding costs are independent of time and customer class.

Assumption 4. *Partial order fulfillment is allowed.*

This assumption includes splitting an order for partial delivery in different periods. The assumption holds true in many practical situations and is common in literature (e.g. Talluri and van Ryzin (2004, Sect. 2.2) and Fleischmann and Meyr (2004, Sect. 4)). Here we also need this assumption for tractability. Let u_i denote the amount of supply arriving in period i used to satisfy a given customer order, and let $\bar{u} = (u_1, \ldots, u_T)$. Thus $\sum_{i=1}^{T} u_i$ is the total amount used to satisfy that order. Note also that for an order arriving in period t, u_i corresponds with delivery from stock if $i \leq t$ and with backlogging otherwise.

Table 1 summarizes the above notation. We can now formulate our problem as a stochastic dynamic program with state variable $\bar{x}$ and decision variable $\bar{u}$. In principle, one can condense these vectors by removing all components i, for which $x_i = 0$. I.e., the dimension of the state space corresponds with the number of scheduled replenishments. However, for ease of notation we use $\bar{x}$ and $\bar{u}$ as defined above.

The profit $\hat{P}_t(\bar{x}, d, c, \bar{u})$ earned in period t depends on the available supply, order size, customer class, and fulfillment decision as follows

$$\hat{P}_t(\bar{x}, d, c, \bar{u}) = r_c \sum_{i=1}^{T} u_i - b \sum_{i=1}^{T} u_i (i - t)(1 - \delta_{it}) - h \sum_{i=1}^{T} (x_i - u_i)\delta_{it} \tag{1}$$

where δ_{it} is defined as 1 if $i \leq t$ and 0 otherwise, and $\bar{u}$ has to satisfy $u_i \leq x_i$ for all i and $\sum_i u_i \leq d$.

The first term in Equation (1) calculates the revenues received from satisfying the current order of class c. The second term computes backlogging costs that occur when using supply that arrives

Table 1 Notation

Indices:	
$t=1,\ldots,T$	Periods of the planning horizon
$i=1,\ldots,T$	Periods of inventory replenishment
State variables:	
$\bar{x}=(x_1,\ldots,x_T)$	Vector of available supply quantities
Decision variables:	
$\bar{u}=(u_1,\ldots,x_T)$	Vector of supply quantities used to fulfill a given order
Random variables:	
c	Customer class
d	Order quantity
$F(c,d)$	Joint cdf of customer class c and order quantity d
Data:	
r_c	Unit revenue from customer class c
b	Unit backorder costs per period
h	Unit holding cost per period

later than the customer order, i.e. when $\delta_{it}=0$. These costs are computed for the total length of the delay $(i-t)$. The third term represents holding costs that are charged for the on-hand inventory at the end of period t. Note that unlike the backlogging costs, which are charged for the total customer waiting time, holding costs only cover the current period t. To simplify subsequent calculations, we define $P_t(i,c)$ as the incremental profit per unit of supply i used to satisfy one unit of an order of class c in period t. Collecting the terms in Equation (1) that depend on u_i yields

$$P_t(i,c)=r_c-b(i-t)(1-\delta_{it})+h\delta_{it} \tag{2}$$

and

$$\hat{P}_t(\bar{x},d,c,\bar{u})=\sum_{i=1}^{T}P_t(i,c)u_i-h\sum_{i=1}^{T}x_i\delta_{it}. \tag{3}$$

In addition to the current period's profit, we also have to consider the impact of a fulfillment decision $\bar{u}$ on future profits. The state transition is given by $\bar{x}\rightarrow\bar{x}-\bar{u}=(x_1-u_1,\ldots,x_T-u_T)$.

Letting $V_t(\bar{x})$ denote the maximum expected profit-to-go from period t to the end of the planning horizon T for a given supply vector $\bar{x}$ we then obtain the following Bellman recursion

$$V_t(\bar{x}) = E_{d,c}\left[\max_{\substack{\bar{u} \\ 0\leq u_i\leq x_i, \sum_{i=1}^T u_i\leq d}} \left\{\sum_{i=1}^{T}(u_i P_t(i,c) - h x_i \delta_{it}) + V_{t+1}(\bar{x}-\bar{u})\right\}\right] \quad (4)$$

with the boundary condition $V_{T+1}(\bar{x}) = 0$.

4. Structural properties and optimal policy

We now analyze structural properties of the value function of the dynamic program defined in the previous section. This allows us to characterize the optimal fulfillment policy. We start by defining marginal profits:

Definition 1. $\Delta_i V_t(\bar{x}) := V_t(\bar{x}) - V_t(\bar{x} - \bar{e}_i)$ for $x_i \geq 1$,

where $\bar{e}_i$ denotes the i-th unit vector. Definition 1 concerns the expected marginal value of a unit of supply arriving in period i or, equivalently, the opportunity costs of selling this unit. Using this definition, we can rewrite the Bellman recursion of Equation (4) as follows.

$$\begin{aligned} V_t(\bar{x}) &= E_{d,c}[\max_{\substack{\bar{u} \\ 0\leq u_i\leq x_i, \sum_{i=1}^T u_i\leq d}} \{\sum_{i=1}^{T}(u_i P_t(i,c) - h x_i \delta_{it}) + V_{t+1}(\bar{x}-\bar{u})\}] \\ &= E_{d,c}\left[\max_{\substack{\bar{u} \\ 0\leq u_i\leq x_i, \sum_{i=1}^T u_i\leq d}} \left\{\sum_{i=1}^{T}(u_i P_t(i,c) - h x_i \delta_{it})\right.\right. \\ &\quad \left.\left. + V_{t+1}(\bar{x}) - \sum_{i=1}^{T}\sum_{z=1}^{u_i} \Delta_i V_{t+1}\left(\bar{x} - \bar{e}_i(z-1) - \sum_{j=1}^{i-1}(\bar{e}_j u_j)\right)\right\}\right] \\ &= V_{t+1}(\bar{x}) - h\sum_{i=1}^{t} x_i + E_{d,c}\left[\max_{\substack{\bar{u} \\ 0\leq u_i\leq x_i, \sum_{i=1}^T u_i\leq d}} \left\{\sum_{i=1}^{T}\left(\sum_{z=1}^{u_i}(P_t(i,c)\right.\right.\right. \\ &\quad \left.\left.\left. - \Delta_i V_{t+1}(\bar{x} - \bar{e}_i(z-1) - \sum_{j=1}^{i-1}(\bar{e}_j u_j))\right)\right\}\right] \end{aligned} \quad (5)$$

Note that this formulation decomposes $\bar{u}$ into single-unit steps. In this way, the maximization in (5) reflects the trade-off between the profit of selling a unit of supply now and the corresponding opportunity cost. Talluri and van Ryzin (2004, p. 59) give a similar decomposition for the classical single-leg airline yield management problem. What is different in (5) is the summation over i, which introduces an additional dimension to the problem.

We now identify properties of the value function that help us evaluate the above maximization expression. The first step is to compare the marginal values of supplies arriving in different periods.

Proposition 1. *For all $m < n$ and for all $\bar{x}$ with $x_m, x_n \geq 1$ the value function satisfies:*

a) $\Delta_m V_t(\bar{x}) - \Delta_n V_t(\bar{x}) \leq b(\max(n,t) - \max(m,t)) = b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t))$

b) $V_t(\bar{x} + \bar{e}_m) - V_t(\bar{x} + \bar{e}_n) \leq b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t))$

Proof We show part a) by induction. For $t = T+1$ the inequality holds since $V_{T+1} \equiv 0$ and since the right-hand side is non-negative for $n > m$. Now assume that inequality a) holds for $t+1$. We show that it also holds for t. We can rewrite

$$\Delta_m V_t(\bar{x}) - \Delta_n V_t(\bar{x}) \leq b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t))$$
$$\Leftrightarrow V_t(\bar{x} - \bar{e}_n) - V_t(\bar{x} - \bar{e}_m) \leq b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t)).$$

We show that this inequality holds for any given values of c and d, which implies that it also holds in expectation. To this end, let $W_t^{\bar{u}}(\bar{x}, c, d)$ denote the maximum expected profit-to-go when starting in period t with a supply vector $\bar{x}$, receiving demand d from customer class c and taking the fulfillment decision $\bar{u}$. Furthermore, let $\bar{u}_t^*(\bar{x}, d, c)$ be an optimal decision in period t under the same conditions.

Using this notation, we have to show that

$$W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)}(\bar{x} - \bar{e}_n, c, d) - W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_m,d,c)}(\bar{x} - \bar{e}_m, c, d) \leq b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t)).$$

We show that there is a feasible decision $\bar{u}_t(\bar{x} - \bar{e}_m, d, c)$ for which this inequality holds. This suffices since $W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_m,d,c)} \geq W_t^{\bar{u}_t(\bar{x}-\bar{e}_m,d,c)}$.

By definition, we have

$$W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)}(\bar{x} - \bar{e}_n, c, d) = \sum_{i=1}^{T} (u_i^* P_t(i,c) - h x_i \delta_{it}) + h\delta_{nt} + V_{t+1}((\bar{x} - \bar{e}_n) - \bar{u}^*)$$

$$= \dots + u_m^* P_t(m,c) - hx_m\delta_{mt} + \dots + u_n^* P_t(n,c)$$
$$-h(x_n-1)\delta_{nt} + \dots + V_{t+1}(\bar{x} - \bar{e}_n - \bar{u}^*), \quad (6)$$

where we have omitted the arguments of u_i^* for notational convenience. We now construct an appropriate feasible decision for state $\bar{x} - \bar{e}_m$. We distinguish two cases.

Case (i): $u_m^* > 0$ In this case, the decision $\bar{u}_t^*(\bar{x} - \bar{e}_n, d, c) - \bar{e}_m + \bar{e}_n$ is feasible in state $\bar{x} - \bar{e}_m$ and we get

$$\begin{aligned} W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)-\bar{e}_m+\bar{e}_n}(\bar{x} - \bar{e}_m, c, d) &= \sum_{i=1}^{T} (u_i^* P_t(i,c) - hx_i\delta_{it}) + P_t(n,c) - P_t(m,c) + h\delta_{mt} \\ &\quad + V_{t+1}((\bar{x} - \bar{e}_m) - (\bar{u}^* - \bar{e}_m + \bar{e}_n)) \\ &= \dots + (u_m^* - 1) P_t(m,c) - h(x_m - 1)\delta_{mt} + \dots \\ &\quad + (u_n^* + 1) P_t(n,c) - hx_n\delta_{nt} + \dots + V_{t+1}(\bar{x} - \bar{e}_n - \bar{u}^*). \quad (7) \end{aligned}$$

Taking the difference between Equations (6) and (7), the profits-to-go after period $t+1$ vanish and we are left with the difference in current profits, which equals

$$P_t(m,c) - P_t(n,c) + h\delta_{nt} - h\delta_{mt} = b(n - m + \delta_{nt}(t-n) + \delta_{mt}(m-t)),$$

by definition of P_t.

Case (ii): $u_m^* = 0$ In this case, the decision $\bar{u}_t^*(\bar{x} - \bar{e}_n, d, c)$ is feasible in state $\bar{x} - \bar{e}_m$ and we get

$$\begin{aligned} W_t^{\bar{u}_t^*(\bar{x}-\bar{e}_n,d,c)}(\bar{x} - \bar{e}_m, c, d) &= \sum_{i=1}^{T} (u_i^* P_t(i,c) - hx_i\delta_{it}) + h\delta_{mt} + V_{t+1}((\bar{x} - \bar{e}_m) - \bar{u}^*) \\ &= \dots + u_m^* P_t(m,c) - h(x_m - 1)\delta_{mt} + \dots \\ &\quad + u_n^* P_t(n,c) - hx_n\delta_{nt} + \dots + V_{t+1}(\bar{x} - \bar{e}_m - \bar{u}^*). \quad (8) \end{aligned}$$

Calculating the difference in current profits between Equations (6) and (8) yields the two following terms. First, for the holding costs, we have

$$\begin{aligned} &- hx_m\delta_{mt} - h(x_n - 1)\delta_{nt} - (-h(x_m - 1)\delta_{mt} - hx_n\delta_{nt}) \\ &= h(\delta_{nt} - \delta_{mt}) \le 0, \end{aligned}$$

where the inequality follows from $n > m$. Second, for the difference in expected future profits we obtain

$$\begin{aligned}
& V_{t+1}(\bar{x}-\bar{e}_n-\bar{u}^*) - V_{t+1}(\bar{x}-\bar{e}_m-\bar{u}^*) \\
\leq\ & b(n-m+\delta_{n,t+1}(t+1-n)+\delta_{m,t+1}(m-t-1)) \\
\leq\ & b(n-m+\delta_{nt}(t-n)+\delta_{mt}(m-t)) = \begin{cases} b(n-m) & \text{for } t < m < n \\ b(n-t) & \text{for } m \leq t < n \\ 0 & \text{for } m < n \leq t \end{cases}, \qquad (9)
\end{aligned}$$

where the first inequality follows from the induction assumption and the second inequality follows since (9) is decreasing in t for $m < n$. This completes the proof of Part a).

Part b) follows immediately from Part a) by replacing $\bar{x}$ with $\bar{x}+\bar{e}_m+\bar{e}_n$ and using the definition of $\Delta_i V_t(\bar{x})$.

□

Proposition 1 states that the difference between the marginal value of one unit of a supply arriving in period m and one unit arriving later in period n is bounded by the difference in backordering costs of using each of these supplies in period t. This relationship implies the following important monotonicity property, regarding the alternative fulfillment options.

Proposition 2. *For all $m < n$ and for all $\bar{x}$ with $x_m, x_n \geq 1$ it holds that* $P_t(m,c) - \Delta_m V_{t+1}(\bar{x}) \geq P_t(n,c) - \Delta_n V_{t+1}(\bar{x})$, $\forall c$.

Proof We have the following equivalences

$$\begin{aligned}
& P_t(m,c) - \Delta_m V_{t+1}(\bar{x}) \geq P_t(n,c) - \Delta_n V_{t+1}(\bar{x}) \\
\Leftrightarrow\ & P_t(m,c) - P_t(n,c) \geq \Delta_m V_{t+1}(\bar{x}) - \Delta_n V_{t+1}(\bar{x}) \\
\overset{\text{Equ.(2)}}{\Leftrightarrow}\ & b(n-m+\delta_{nt}(t-n)+\delta_{mt}(m-t)) + h(\delta_{mt}-\delta_{nt}) \geq \Delta_m V_{t+1}(\bar{x}) - \Delta_n V_{t+1}(\bar{x}).
\end{aligned}$$

From Part a) of Proposition 1 we know that $\Delta_m V_{t+1}(\bar{x}) - \Delta_n V_{t+1}(\bar{x}) \leq b(n-m+\delta_{n,t+1}(t+1-n)+\delta_{m,t+1}(m-t-1))$. As in Case (ii) of the proof of Proposition 1, this implies the desired result since for $n > m$ we have $\delta_{n,t+1}(t+1-n)+\delta_{m,t+1}(m-t-1) \leq \delta_{nt}(t-n)+\delta_{mt}(m-t)$ and $\delta_{mt}-\delta_{nt} \geq 0$.

□

The terms on the left-hand-side of the inequality can be interpreted as the net benefit of the current revenues from selling a unit of supply arriving in period m minus the opportunity cost of not having that unit available in the future. Proposition 2 states that this net benefit is decreasing in the arrival time of the supply. Therefore, an order should always be either fulfilled using the earliest available supply or not at all (if the left-hand-side becomes negative). This is a key property, which allows us to drastically reduce the solution space. Instead of the T-dimensional optimization in Equation (5), a simple line-search is sufficient to identify the optimal fulfillment decision for a given order.

The next important property concerns the concavity of the value function along certain axes.

PROPOSITION 3. *Let $\bar{x}$ be such that $\sum_i x_i \geq 2$. Furthermore, let $m = \min\{i|x_i > 0\}$ and let $n = m$ if $x_m > 1$ and $n = \min\{i|i > m, x_i > 0\}$ otherwise. Then*
$P_t(m,c) - \Delta_m V_{t+1}(\bar{x}) \geq P_t(n,c) - \Delta_n V_{t+1}(\bar{x} - \bar{e}_m)$ *,* $\forall c$.

We prove this property jointly with Theorem 1 below. Proposition 3 implies in particular that the value function V_t is concave in the quantity of the earliest available supply. This allows us to characterize the optimal fulfillment policy by a set of critical levels as follows.

THEOREM 1. *Define the following set of critical levels:*
For $i = 1, \ldots, T$ let $\bar{y}_i = \bar{x} - \sum_{k=1}^{i} \bar{e}_k x_k$,
and let $L_t(c,i,\bar{y}_i) = \max\{k|P_t(i,c) < \Delta_i V_{t+1}(\bar{y}_i + k\bar{e}_i)\}$.
Then the following fulfillment decision is optimal in period t, given an order quantity d from customer class c:
Start with $i = 1$;
set $u_i = \max\left(\min\left(x_i - L_t(c,i,\bar{y}_i), d - \sum_{k=1}^{i-1} u_k\right), 0\right)$;
if $u_i < x_i$ set $u_k = 0$ for all $k > i$ and stop, otherwise repeat for i+1.

Proof of Proposition 3 and Theorem 1 We show both properties jointly by induction. For $t = T$, Proposition 3 holds since $V_{T+1}(.) \equiv 0$ and $P_T(m,c) = P_T(n,c)$ for $m,n \leq T$. Now assume that

Proposition 3 holds for Period t. We first show that Theorem 1 then holds for Period t and subsequently that Proposition 3 holds for Period $t-1$.

Equation (5) shows that one can decompose the fulfillment decision into unit steps and that selling a given unit is beneficial if immediate profits outweigh the opportunity cost of losing this unit. Proposition 2 shows that there is an optimal policy $\bar{u}^*$ for which $u_i^* > 0$ implies $x_j - u_j^* = 0$ for all $j < i$. Any optimal policy that does not satisfy this property can be modified by swapping one unit of supply j against one unit of supply i. Proposition 2 implies that this modification does not decrease the objective function value. Therefore, one can obtain an optimal solution through a line search, starting with the earliest available supply, i.e. the smallest i, for which $x_i > 0$. The induction assumption of Proposition 3 shows that the objective function is concave along this search line. Therefore, one can stop the search as soon as immediate profits drop below the opportunity costs. This proves that the procedure defined in Theorem 1 yields an optimal policy.

We now show Proposition 3 for Period $t-1$. We use the same notation as in the proof of Proposition 1. Let $\bar{u}^1 := \bar{u}_t^*(\bar{x}, c, d)$, $\bar{u}^2 := \bar{u}_t^*(\bar{x} - \bar{e}_m, c, d)$, and $\bar{u}^3 := \bar{u}_t^*(\bar{x} - \bar{e}_m - \bar{e}_n, c, d)$ denote the optimal decisions in states $\bar{x}$, $\bar{x} - \bar{e}_m$, and $\bar{x} - \bar{e}_m - \bar{e}_n$, respectively, for a given customer class c and demand quantity d. Furthermore, let

$$
\begin{aligned}
A &:= W_t^{\bar{u}^1}(\bar{x}, c, d) - W_t^{\bar{u}^2}(\bar{x} - \bar{e}_m, c, d) \text{ and} \\
B &:= W_t^{\bar{u}^2}(\bar{x} - \bar{e}_m, c, d) - W_t^{\bar{u}^3}(\bar{x} - \bar{e}_m - \bar{e}_n, c, d).
\end{aligned}
$$

We rewrite Proposition 3 for Period $t-1$ as $\Delta_m V_t(\bar{x}) - \Delta_n V_t(\bar{x} - \bar{e}_m) \le P_{t-1}(m, c) - P_{t-1}(n, c)$ and show that

$$
\begin{aligned}
A - B &\le P_t(m, c) - P_t(n, c) - h(\delta_{mt} - \delta_{nt}) \\
&\le P_{t-1}(m, c) - P_{t-1}(n, c),
\end{aligned}
$$

for any values of c and d, which implies that these inequalities also hold in expectation.

The second inequality follows directly from the definition of $P_t(.,.)$ for $m < n$. For the first inequality, we distinguish three cases, based on the value of $\bar{u}^1$.

<u>*Case (i):* $\bar{u}^1 \equiv 0$</u> Theorem 1 implies $\bar{u}^2 = \bar{u}^3 \equiv 0$. From the definition of W_t we get

$$\begin{aligned} A - B &= \Delta_m V_{t+1}(\bar{x}) - h\delta_{mt} - \Delta_n V_{t+1}(\bar{x} - \bar{e}_m) + h\delta_{nt} \\ &\leq P_t(m,c) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}), \end{aligned}$$

where the inequality follows from the induction assumption.

<u>*Case (ii):* $0 < \sum_{i=1}^T u_i^1 < d$</u> Theorem 1 implies $\bar{u}^2 = \bar{u}^1 - \bar{e}_m$. There are two possibilities for $\bar{u}^3$. If $\bar{u}^2 \neq 0$ then $\bar{u}^3 = \bar{u}^2 - \bar{e}_n$, otherwise $\bar{u}^3 \equiv 0$. In the first case, we get

$$A - B = P_t(m,c) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}).$$

In the second case, $B = \Delta_n V_{t+1}(\bar{x} - \bar{e}_m) - h\delta_{nt}$ and Theorem 1 implies that this value is smaller than $P_t(n,c) - h\delta_{nt}$ since it is optimal not to sell another unit of supply n. Thus, $A - B$ satisfies the desired inequality in either case.

<u>*Case (iii):* $\sum_{i=1}^T u_i^1 = d$</u> Theorem 1 implies that either $\bar{u}^2 = \bar{u}^1 - \bar{e}_m$ or $\bar{u}^2 = \bar{u}^1 - \bar{e}_m + \bar{e}_k$ for some $k \geq n$. The first alternative leads to the same calculations as in Case (ii) above. The second alternative leaves two options for $\bar{u}^3$, namely either $\bar{u}^3 = \bar{u}^2 - \bar{e}_n + \bar{e}_l$ for some $l \geq k$ or $\bar{u}^3 = \bar{u}^2 - \bar{e}_n$. In the first case, we get

$$\begin{aligned} A - B &= P_t(m,c) - P_t(k,c) + \Delta_k V_{t+1}(\bar{x} - \bar{u}^1) - P_t(n,c) - P_t(l,c) + \Delta_l V_{t+1}(\bar{x} - \bar{u}^1 - \bar{e}_k) - h(\delta_{mt} - \delta_{nt}) \\ &\leq P_t(m,c) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}), \end{aligned}$$

where the inequality follows from the induction assumption. The other case regarding $\bar{u}^3$ yields $B = P_t(n,c) - h\delta_{nt}$ and therefore

$$\begin{aligned} A - B &= P_t(m,c) - P_t(k,c) + \Delta_k V_{t+1}(\bar{x} - \bar{u}^1) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}) \\ &\leq P_t(m,c) - P_t(n,c) - h(\delta_{mt} - \delta_{nt}), \end{aligned}$$

where the inequality follows from the fact that it is optimal in state $\bar{x} - \bar{u}^1$ to sell an additional unit of supply k.

$\square$

Intuitively, the optimal fulfillment policy successively consumes units of supply, in the order of their arrival, until the immediate marginal profit drops below the opportunity costs. This resembles a booking-limits policy in traditional revenue management (Talluri and van Ryzin (2004)). The values $L_t(c,i,.)$ set nested protection levels that bar some amount of supply i from consumption by classes c and higher. Note that we have separate protection levels for each class and supply arrival. Also note that the protection levels of supply i depend on the available quantities of subsequent arrivals $\bar{y}_i$. However, $L_t(c,i,.)$ is independent of x_i (and of all earlier arrivals x_j for $j<i$) and therefore indeed acts as a protection level. The amount of supply i exceeding $L_t(c,i,.)$ is available for consumption by customer class c at time t. It is worth pointing out that even the most valuable customer class, i.e. $c=1$, may face non-trivial booking limits for future supplies, i.e. for $i>t$: While this class can always consume supply on hand it is not necessarily optimal to backlog demand from this class, due to the incurred backorder penalties. We illustrate the various protection levels $L_t(c,i,.)$ graphically in an example in the next section.

The following proposition shows that protection levels are decreasing in time. This is intuitive since a shorter remaining planning horizon implies less selling opportunities and therefore available supply is of less value. It is worth mentioning however, that this result is only true for stationary demand. Unlike in traditional revenue management models, the holding cost term in our model may destroy the monotonicity of the protection levels if the demand distribution changes across periods.

Proposition 4. *The protection levels $L_t(c,i,\bar{y}_i)$ defined in Theorem 1 are non-increasing in t.*

Proof We show by induction that $\Delta_i V_{t+1}(\bar{x})>0$ implies $\Delta_i V_t(\bar{x}) \geq \Delta_i V_{t+1}(\bar{x})$. Since $P_t(i,c)$ is increasing in t this assures non-increasing protection levels for any i and c.

For $t=T$ the condition is empty since $\Delta_i V_{T+1}(\bar{x})=0$ for all i,c, and $\bar{x}$. In other words, all protection levels vanish at the end of the planning horizon. Assume now that Proposition 4 holds for Period $t+1$. We show that it also holds for Period t. To this end, assume that $\Delta_i V_{t+1}(\bar{x})>0$. Using the Bellman recursion of V_t we have

$$\Delta_i V_t(\bar{x}) = -h\delta_{it} + E_{c,d}\left[\max\left\{P_t(i,c)\delta_{d>0}\,;\,\Delta_i V_{t+1}(\bar{x})\right\}\right], \tag{10}$$

where $\delta_{d>0}$ equals unity if $d>0$ and zero otherwise. For $t<i$ the holding-cost term vanishes and we immediately get $\Delta_i V_t(\bar{x}) \geq \Delta_i V_{t+1}(\bar{x})$. For $t \geq i$ we rewrite (10) for $t+1$

$$\Delta_i V_{t+1}(\bar{x}) = -h\delta_{it+1} + E_{c,d}\left[\max\left\{P_{t+1}(i,c)\delta_{d>0}\,;\,\Delta_i V_{t+2}(\bar{x})\right\}\right], \tag{11}$$

and compare the individual terms. We have $-h\delta_{it} = -h\delta_{it+1}$ and $P_t(i,c) = P_{t+1}(i,c)$ since $t \geq i$. In addition, the induction assumption implies that $\Delta_i V_{t+1}(\bar{x}) \geq \Delta_i V_{t+2}(\bar{x})$ if the maximum in (11) is attained by the last term. Therefore, $\Delta_i V_t(\bar{x}) \geq \Delta_i V_{t+1}(\bar{x})$, which completes the proof.

□

5. Numerical illustration

5.1. Scenario definition

We conclude our analysis with a brief numerical study. Our objective is to illustrate our theoretical results and the performance of our proposed approach relative to other fulfillment strategies. To this end, we compare the fulfillment algorithm defined in Theorem 1, which we denote by RM, to the following benchmarks:

- *First come first served (FCFS)*: This strategy accepts any customer order and fulfills it from stock, up to the amount on hand. Any excess demand is rejected. A comparison with this policy provides an indication of the benefit of customer differentiation in demand fulfillment. We limit the FCFS policy to fulfillment from stock to avoid excessive backordering, which would decrease its performance unreasonably.

- *Single order processing after allocation planning (SOPA)*: This approach was presented by Meyr (2009). It uses an LP to allocate current and future supply to expected demand from different customer classes over the planning horizon (for additional details see Section 2). SOPA differentiates between customer classes but does so in a deterministic fashion. A comparison with RM therefore gives an indication of the benefit of explicitly recognizing demand uncertainty in the fulfillment decision.

- *Global optimum (GOP)*: This policy optimally allocates supply to demand ex-post, once all customer orders are known. Note that this policy cannot be implemented. It serves as a benchmark, indicating the impact of demand uncertainty on expected profits.

In principle, the expected profits of FCFS and SOPA can be evaluated by means of dynamic programming. It suffices to replace the maximization over $\bar{u}$ in (4) by the decisions $\bar{u}^{FCFS}$ and $\bar{u}^{SOPA}$ according to these policies. However, this evaluation of SOPA requires the solution of an LP for each state $\bar{x}$ and each time step t, which is computationally expensive for large scenarios. In addition, GOP cannot be determined through recursion (4). Therefore, we evaluated the different policies by means of simulation.

In extensive tests we verified that for an increasing number of simulation runs the simulated profit of SOPA converges to the expected profit determined by the dynamic program. In what follows, we report the average simulated profits over 500 simulation runs. In all the cases that we checked, the difference between this simulated profit and the expected profit was much smaller than 1%.

We use the following demand model in our simulations. In line with our analytical model, we assume at most one order arrival per period and denote the probability of no order arrival by p_0. We further assume equal order arrival probabilities for all customer classes, which implies that this probability equals $\frac{1-p_0}{\#classes}$. Given an order arrival, we assume the order size to be strictly positive and discrete, which we model through a negative binomial distribution (NBD). This choice allows us to model large coefficients of variation (CV). In addition, NBDs have a long tradition in the marketing and operations literature, based on strong empirical support (e.g. Ehrenberg 1959, Agrawal and Smith 1998). In order to enforce a strictly positive order size, we shift the NBD to the right by one unit. Thus, letting $NB(\mu, \sigma)$ denote a negative binomially distributed random variable with mean μ and standard deviation σ, we model the order size as $1 + NB(\mu - 1, \sigma)$.

We use the following parameter settings as our base case scenario. The planning horizon is $T = 28$ periods with two receipts of supply of 100 units each, in periods $t = 1$ and $t = 15$, respectively. We assume three customer classes with revenues $r_1 = \$100$, $r_2 = \$90$, and $r_3 = \$80$ and set $h = \$1$ and

$b = \$10$. We further assume exactly one customer arrival per period, i.e. $p_0 = 0$. The order size is negative binomially distributed, as discussed above, with $\mu = 12$ and $\sigma = 8$.

In the next subsection, we evaluate our RM procedure for this base case and compare it to the FCFS, SOPA, and GOP benchmarks. To this end, SOPA takes the expected order size μ as a deterministic and static demand forecast. In subsequent subsections, we then illustrate the impact of several key model parameters on the relative performance of RM. Specifically, we vary demand variability (Section 5.3), customer heterogeneity (Section 5.4), and supply shortage (Section 5.5).

5.2. Base case analysis

We start by applying the RM procedure defined in Theorem 1 to the base case specified in the previous subsection. Figure 2 displays the resulting optimal policy. The left graph shows the customer-class specific protection levels of the first supply x_1 through time, for different levels of the second supply x_{15}. Note that the protection levels are decreasing in the customer class and decreasing over time and that customer class $c = 1$ will always be accepted, as long as x_1 is not depleted. All of these outcomes are intuitive and in line with the results of Section 4.

Figure 2 Protection levels of supply x_1 and x_{15} for different customer classes

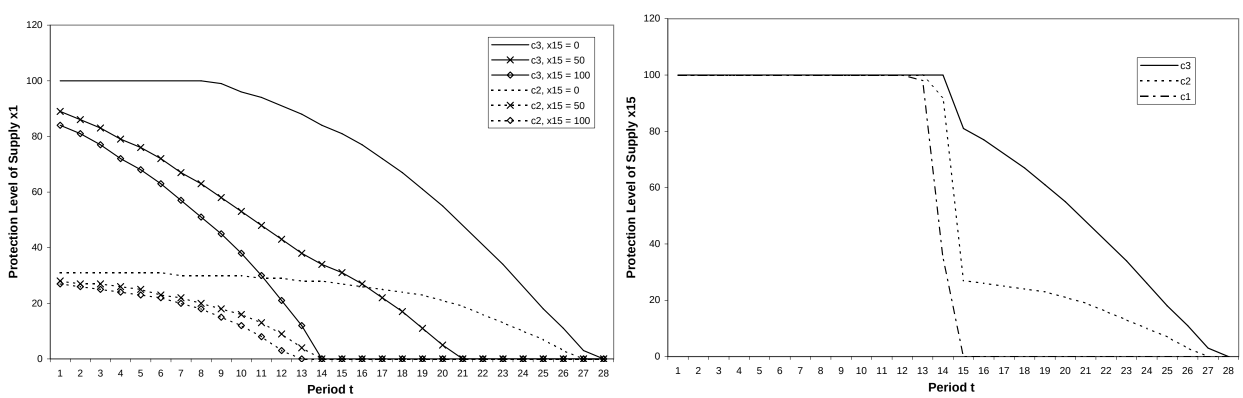

The right graph in Figure 2 shows the protection levels of the second supply x_{15}. Note that until Period 12, the protection levels for all customer classes equal 100 units, implying that the RM policy never backorders any demand for more than two periods. In particular, even customer class

$c = 1$ is facing non-trivial booking-limits in the case of backordering. Note further that from Period 15 onwards the protection levels of supply x_1 (for $x_{15} = 0$) and of x_{15} coincide. This is intuitive since from that period onwards both supplies are available on hand and thus are of the same value to the decision maker.

The base case simulated average profit of the RM policy equals \$ 17,636. Table 2 compares this value to the average profits of the three benchmark policies. RM outperforms both FCFS and SOPA by about 2% and trails GOP by roughly 1% in this example.

Table 2 Base case average profits for different policies

GOP	FCFS	SOPA	RM
17,843	17,247	17,328	17,636

The results in Table 3 help explain the observed performance differences. The table specifies the service levels achieved by each policy. Specifically, it lists for each policy the average fraction of orders lost and backordered of each customer class. Figure 3 displays these values graphically.

Table 3 Customer service levels in the base case

	Lost sales				Backlogging			
Customer class	GOP	FCFS	SOPA	RM	GOP	FCFS	SOPA	RM
1	1.8%	38.7%	7.8%	14.6%	2.62%	0.00%	2.84%	3.10%
2	32.7%	40.4%	29.5%	26.9%	1.51%	0.00%	2.62%	1.45%
3	86.4%	38.5%	83.6%	76.9%	0.03%	0.00%	1.04%	0.00%

Figure 3 Customer service levels in the base case

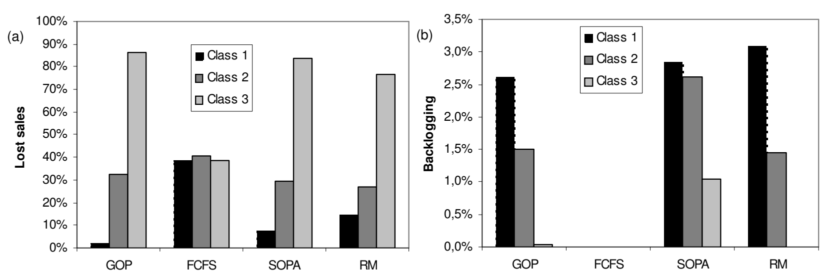

Several differences between the policies stand out. First of all, FCFS does not differentiate between customer classes, which is reflected in roughly uniform service levels across classes. All other policies clearly prioritize high-value orders. While FCFS achieves the lowest total number of lost sales, it loses relatively many high-value orders. Comparing RM and SOPA, we see that SOPA rejects even more orders of Classes 2 and 3. However, the resulting decrease in lost sales of Class 1, relative to RM, is insufficient to compensate the lost revenues in this example. Another difference between RM and SOPA concerns the backordering behavior, where SOPA backorders many more orders from lower customer classes.

5.3. Impact of demand variability

We now expand the analysis of our example by varying a number of key parameters. We start with demand variability and vary the standard deviation of the order size σ from 4 to 16. In addition, we also consider the case of a constant order size of $\mu = 12$ units. This range corresponds with a CV of the order size between 0 and 1.33. The scenario CV=0.67, $\sigma = 8$ corresponds with the base case analyzed in detail in the previous subsection. Note that in addition to the order size, we also have uncertainty in the order arrivals. Therefore, even the scenario with a constant order size is not entirely deterministic.

Figure 4 Average profits for different levels of demand variability

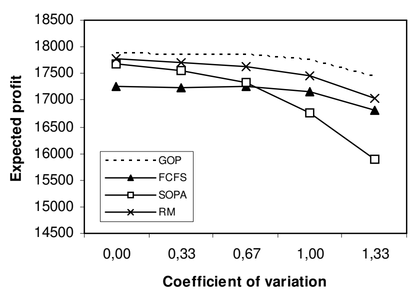

Figure 4 shows the average simulated profits of the four policies for different levels of variability of the order size. Table 4 lists the underlying data. We observe the following: (1) All policies show a decreasing trend in their average profits as demand variability increases. This is intuitive since

Table 4 Average profits for different levels of demand variability

CV / σ	GOP	FCFS	SOPA	RM
0.00 / 0	17,890	17,265	17,682	17,769
0.33 / 4	17,852	17,230	17,554	17,694
0.67 / 8	17,843	17,247	17,328	17,636
1.00 /12	17,749	17,165	16,768	17,449
1.33 /16	17,441	16,804	15,880	17,031

higher variability corresponds with a larger forecasting error, which complicates the planning task. (2) FCFS suffers the least from increasing demand variability. A potential explanation is that FCFS does not use any forecasting information in the fulfillment decision and therefore is not affected by the forecasting error. (3) For low demand variability, SOPA is competitive to RM. For large demand variability, the performance of SOPA deteriorates drastically, even falling behind FCFS. Recall that SOPA is based on expected demand, thus ignoring any forecasting error. Under low uncertainty, this approach can be expected to work well. However, this example shows that if the demand forecast is unreliable, using it for a detailed optimization as in SOPA may become counter-productive. (4) RM yields the best ex-ante decision. Its advantage over FCFS decreases with increasing demand variability. Since RM maximizes the expected profit it necessarily dominates FCFS and SOPA. The main difference with FCFS lies in a more restrictive acceptance of low-value orders. The results show that for higher demand variability, customer differentiation may lose some of its effectiveness since it becomes more risky to reserve demand for potential future orders.

Note that some of the variations in the average profits can be explained by changes in the maximum attainable ex-post profit GOP. To eliminate this effect, we consider the relative deviation of FCFS, SOPA, and RM with respect to GOP in the subsequent analyses.

5.4. Impact of customer heterogeneity

The degree of heterogeneity between customer classes is another key parameter that one would expect to impact the effectiveness of different fulfillment policies. Therefore, we complement the scenario of the previous subsection ('low heterogeneity': $r_{1/2/3} = 100/90/80$) with a 'medium heterogeneity' scenario ($r_{1/2/3} = 100/80/60$) and a 'high heterogeneity' scenario ($r_{1/2/3} = 100/70/40$).

The results are displayed in Figure 5. Table 5 lists the underlying data. Note that the results show the average relative profit deviation from GOP. Thus, lower numbers mean higher performance. Also note that the average relative profit deviation differs from the relative deviation of average profits.

Figure 5 Average profit deviation relative to GOP, for different levels of customer heterogeneity: (a) Low: $r_{1/2/3} = 100/90/80$,(b) Medium: $r_{1/2/3} = 100/80/60$, (c) High: $r_{1/2/3} = 100/70/40$

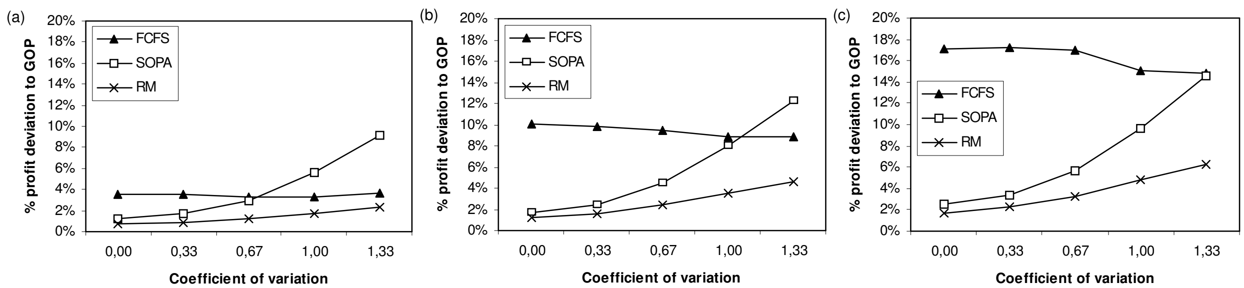

Table 5 Average profit deviation relative to GOP, for different levels of customer heterogeneity

	Low heterogeneity			Medium heterogeneity			High heterogeneity		
CV / σ	FCFS	SOPA	RM	FCFS	SOPA	RM	FCFS	SOPA	RM
0.00 / 0	3.5%	1.2%	0.7%	10.0%	1.7%	1.3%	17.1%	2.5%	1.7%
0.33 / 4	3.5%	1.7%	0.9%	9.8%	2.4%	1.6%	17.3%	3.3%	2.3%
0.67 / 8	3.3%	2.9%	1.2%	9.4%	4.5%	2.5%	17.0%	5.7%	3.3%
1.00 /12	3.3%	5.6%	1.7%	8.8%	8.0%	3.5%	15.1%	9.6%	4.9%
1.33 /16	3.7%	9.1%	2.4%	8.9%	12.2%	4.6%	14.8%	14.6%	6.3%

The results in Figure 5 show the same qualitative patterns as in Section 5.3. However, the degree of customer heterogeneity dramatically affects the performance of FCFS relative to RM and SOPA. For instance, the gap between RM and FCFS, at a CV of 0.67, increases from roughly 2% for low heterogeneity to more than 13% for high heterogeneity, in this example. In contrast, the difference between RM and SOPA is rather insensitive to the degree of customer heterogeneity here.

5.5. Impact of supply shortage

Finally, we briefly consider the impact of the degree of supply scarcity. We define the shortage rate (sr) as one minus the ratio between total supply and total expected demand throughout the planning horizon. Formally,

$$sr = 1 - \frac{\sum_{i=1}^{T} x_i}{(1-p_0) \times E[d] \times T}.$$

Thus, positive values of sr indicate a supply shortage, whereas negative values indicate oversupply. By varying p_0 from 0 to 0.75, we vary sr between 41% and -138%.

Figure 6 Impact of shortage rate on average profit deviation from GOP, for different scenarios: (a) $CV = 0.33$ and low heterogeneity, (b) $CV = 0.67$ and low heterogeneity, (c) $CV = 0.67$ and medium heterogeneity

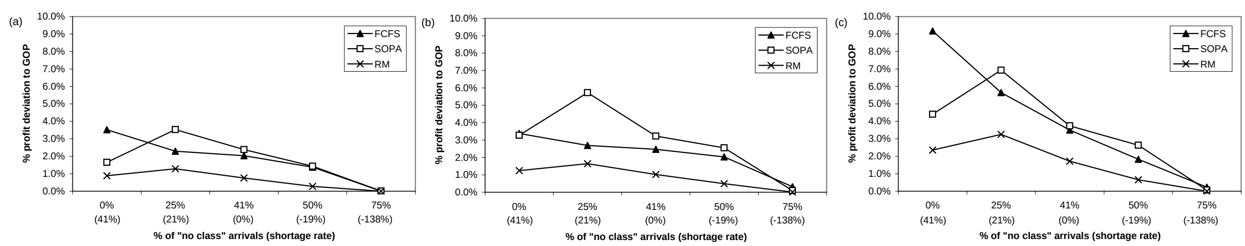

Table 6 Impact of shortage rate on average profit deviation from GOP, for different scenarios

	$CV = 0.33$, low het.			$CV = 0.67$, low het.			$CV = 0.67$, med. het.		
p_0 / sr	FCFS	SOPA	RM	FCFS	SOPA	RM	FCFS	SOPA	RM
0% / 41%	3.5%	1.7%	0.9%	3.4%	3.3%	1.2%	9.2%	4.4%	2.4%
25% / 21%	2.3%	3.5%	1.3%	2.7%	5.7%	1.6%	5.6%	6.9%	3.3%
41% / 0%	2.0%	2.4%	0.8%	2.5%	3.2%	1.0%	3.5%	3.8%	1.7%
50% /-19%	1.4%	1.4%	0.3%	2.0%	2.7%	0.5%	1.8%	2.6%	0.7%
75% /-138%	0.0%	0.0%	0.0%	0.3%	0.1%	0.0%	0.2%	0.1%	0.0%

Figure 6 displays the impact of sr on the relative profit deviations of FCFS, SOPA, and RM from GOP in three different scenarios. The graph in the middle shows the base case scenario of Section 5.2 (CV=0.67, low heterogeneity), the left graph shows a scenario with lower demand variability

(CV=0.33), and the right graph a scenario with higher customer heterogeneity. Table 6 lists the underlying data.

The three graphs show very similar patterns. Note that in each graph the left half of the x-axis corresponds with shortages whereas the right half corresponds with an oversupply. In each scenario, we observe that for significant oversupply the different fulfillment policies coincide — if supply is not a limiting factor, the fulfillment decision is obvious: satisfy all orders. We further observe that the relative performance of FCFS is decreasing for an increasing shortage rate. This is intuitive since the benefit of prioritizing the most profitable orders increases as shortages become more acute. Interestingly, the relative performance of SOPA and RM is not monotonic in sr. In all three scenarios, both policies perform worst for an intermediate shortage rate of 21%. A potential explanation is that for a moderate shortage rate the trade-off between selling a unit of supply for low revenues versus reserving it for potentially higher future revenues is the most difficult. For a high shortage rate, the benefit of reserving supply for high-margin orders becomes more prominent, leaving less room for improving the fulfillment decision ex-post.

6. Conclusion and outlook

In this paper, we present a revenue-management approach to demand fulfillment in make-to-stock manufacturing. Specifically, we consider the situation of a manufacturer who decides on the quantities he is willing to sell to different customer classes. The order-acceptance decision takes into account on-hand inventory as well as planned production quantities, scheduled to arrive in the future. For each order, the manufacturer has to decide whether to accept it, to reject it, or to backlog it against a price discount.

The problem is motivated by the demand-fulfillment task in advanced planning systems. A key characteristic of the problem is that production orders cannot be changed in the short term. This is in line with the hierarchical-planning approach of most advanced planning systems and reflects the reality of many manufacturers. To our knowledge, our paper is the first to apply revenue management in this context.

We model the MTS demand-fulfillment problem as a stochastic dynamic program. What makes this problem complex is that we have multiple fulfillment alternatives for each customer order. However, we are able to show that the optimal policy has a simple, intuitive structure, which can be interpreted as a multi-dimensional booking-limit policy. By explicitly capturing demand uncertainty, our model differs from the rule–based or LP-based deterministic models commonly underlying the demand-fulfillment modules of advanced planning systems.

We tested our model numerically and compared it against first-come-first-served order fulfillment and against a deterministic optimization approach. Our results show that explicitly accounting for demand uncertainty can significantly improve the performance of demand fulfillment, unless forecasts are highly accurate. The results also show that customer differentiation can yield a substantial profit increase, in particular if differences in profitability are large across orders and if supply is scarce. In conclusion, our results highlight substantial opportunities for improving the current practice of order fulfillment in make-to-stock manufacturing.

Our model makes a first step in this direction. However, many challenges remain. Possibly the biggest limitation of our model is its limited scalability. For large problem instances, the full evaluation of the value function is computationally demanding. Developing efficient heuristics is therefore an intuitive next step. Techniques such as randomized linear programming that have been successfully applied in traditional revenue management may provide a good starting point. It will be interesting to see how such approximations compare to a deterministic approach. In order to gain further insight into the relative performance of different methods, they should also be compared based on empirical data, in addition to theoretical demand distributions. Another relevant extension to our model is to include different customer due-dates. It is not immediately clear which effect this will have on the structure of the optimal fulfillment policy.

Acknowledgments

The authors are grateful to the Vienna Science and Technology Fund (WWTF) for funding this research.

References

Agrawal, Narendra, Stephen A. Smith. 1998. Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Research Logistics* **43**(6) 839–861.

Arslan, Hasan, Stephen C. Graves, Thomas Roemer. 2007. A single-product inventory model for multiple demand classes. *Management Science* **53**(9) 1486–1500.

Ball, Michael O., Chien-Yu Chen, Zhen-Ying Zhao. 2004. Available-to-promise. David Simchi-Levi, S. David Wu, Zuo-Jun Shen, eds., *Handbook of Quantitative Supply Chain Analysis – Modeling in the E-Business Era*, chap. 11. Kluwer Academic Publishers, 447–483.

Chen, Chien-Yu, Zhen-Ying Zhao, Michael O. Ball. 2001. Quantity and due date quoting available to promise. *Information Systems Frontiers* **3**(4) 477–488.

Chen, Chien-Yu, Zhenying Zhao, Michael O. Ball. 2002. A model for batch advanced available-to-promise. *Production and Operations Management* **11**(4) 424–440.

Chiang, Wen-Chyuan, Jason C. H. Chen, Xiaojing Xu. 2007. An overview of research on revenue management: current issues and future research. *International Journal of Revenue Management* **1**(1) 97–128.

Defregger, Florian, Heinrich Kuhn. 2007. Revenue management for a make-to-order company with limited inventory capacity. *OR Spectrum* **29**(1) 137–156.

Ehrenberg, Andrew S. C. 1959. The pattern of consumer purchases. *Applied Statistics* **8**(1) 26–41.

Fischer, Markus E. 2001. *"Available-to-Promise": Aufgaben und Verfahren im Rahmen des Supply Chain Management*. Theorie und Forschung Band 704, Wirtschaftswissenschaften Band 63, Roderer, Regensburg.

Fleischmann, Bernhard, Herbert Meyr. 2003. Planning hierarchy, modeling and advanced planning systems. A.G. de Kok, S.C. Graves, eds., *Supply Chain Management: Design, Coordination and Operation*, *Handbooks in Operations Research and Management Science*, vol. 11, chap. 9. Elsevier, Amsterdam, 457–523.

Fleischmann, Bernhard, Herbert Meyr. 2004. Customer orientation in advanced planning systems. Harald Dyckhoff, Richard Lackes, Joachim Reese, eds., *Supply Chain Management and Reverse Logistics*, chap. 3. Springer, 297–321.

Gupta, Diwakar, Lei Wang. 2007. Capacity management for contract manufacturing. *Operations Research* **55**(2) 367–377.

Ha, Albert Y. 1997. Stock-rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Research Logistics* **44**(5) 457–472.

Harris, Frederick H., Jonathan P. Pinder. 1995. A revenue management approach to demand management and order booking in assemble-to-order manufacturing. *Journal of Operations Management* **13**(4) 299–309.

Kilger, Christoph. 2008. Computer assembly. Hartmut Stadtler, Christoph Kilger, eds., *Supply Chain Management and Advanced Planning*, 4th ed. Springer, Berlin, Heidelberg, 381–398.

Kilger, Christoph, Herbert Meyr. 2008. Demand fulfilment and ATP. Hartmut Stadtler, Christoph Kilger, eds., *Supply Chain Management and Advanced Planning*, 4th ed. Springer, Berlin, Heidelberg, 181–198.

Kimms, Alf, Robert Klein. 2005. Revenue Management im Branchenvergleich. *Zeitschrift für Betriebswirtschaft* **1** 1–30.

Kleijn, Marcel J., Rommert Dekker. 1998. An overview of inventory systems with several demand classes. Econometric institute report 9838/a, RSM Erasmus University.

Kleywegt, Anton J., Jason D. Papastavrou. 2001. The dynamic and stochastic knapsack problem with random sized items. *Operations Research* **49**(1) 26–41.

Lautenbacher, Conrad J., Shaler Jr. Stidham. 1999. The underlying markov decision process in the single-leg airline yield-management problem. *Transportation Science* **33**(2) 136–146.

Meyr, Herbert. 2008. Clustering methods for rationing limited resources. Lars Mönch, Giselher Pankratz, eds., *Intelligente Systeme zur Entscheidungsunterstützung*. Multikonferenz Wirtschaftsinformatik, München, 26.02.2008-28.02.2008, SCS Publishing House e.V, San Diego et al., 19–31.

Meyr, Herbert. 2009. Customer segmentation, allocation planning and order promising in make-to-stock production. *OR Spectrum* **31**(1) 229–256.

Pibernik, Richard. 2006. Managing stock-outs effectively with order fulfilment systems. *Journal of Manufacturing Technology Management* **17**(6) 721–736.

Pibernik, Richard, Prashant Yadav. 2009. Inventory reservation and real-time order promising in a make-to-stock system. *OR Spectrum* **31**(1) 281–307.

Quante, Rainer, Herbert Meyr, Moritz Fleischmann. 2009. Revenue management and demand fulfillment: Matching applications, models, and software. *OR Spectrum* **31**(1) 31–62.

Spengler, Thomas, Stefan Rehkopf, Thomas Volling. 2007. Revenue management in make-to-order manufacturing – an application to the iron and steel industry. *OR Spectrum* **29**(1) 157–171.

Stadtler, Hartmut, Christoph Kilger, eds. 2008. *Supply Chain Management and Advanced Planning*. 4th ed. Springer, Berlin, Heidelberg.

Swann, Julie L. 1999. Flexible pricing policies: Introduction and a survey of implementation in various industries. Contract report, General Motors Corporation.

Talluri, Kalyan T., Garrett J. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers.

Teunter, Ruud H., Willem K. Klein Haneveld. 2008. Dynamic inventory rationing strategies for inventory systems with two demand classes, poisson demand and backordering. *European Journal of Operational Research* **190**(1) 156–178.

Van Slyke, Richard, Yi Young. 2000. Finite horizon stochastic knapsacks with applications to yield management. *Operations Research* **48**(1) 155–172.

Weatherford, Lawrence R., Samuel E. Bodily. 1992. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Operations Research* **40**(5) 831–844.

Publications in the Report Series Research[*] in Management

ERIM Research Program: “Business Processes, Logistics and Information Systems”

2009

How to Normalize Co-Occurrence Data? An Analysis of Some Well-Known Similarity Measures
Nees Jan van Eck and Ludo Waltman
ERS-2009-001-LIS
http://hdl.handle.net/1765/14528

Spare Parts Logistics and Installed Base Information
Muhammad N. Jalil, Rob A. Zuidwijk, Moritz Fleischmann, and Jo A.E.E. van Nunen
ERS-2009-002-LIS
http://hdl.handle.net/1765/14529

Open Location Management in Automated Warehousing Systems
Yugang YU and René B.M. de Koster
ERS-2009-004-LIS
http://hdl.handle.net/1765/14615

VOSviewer: A Computer Program for Bibliometric Mapping
Nees Jan van Eck and Ludo Waltman
ERS-2009-005-LIS
http://hdl.handle.net/1765/14841

Nash Game Model for Optimizing Market Strategies, Configuration of Platform Products in a Vendor Managed Inventory (VMI) Supply Chain for a Product Family
Yugang Yu and George Q. Huang
ERS-2009-009-LIS
http://hdl.handle.net/1765/15029

A Mathematical Analysis of the Long-run Behavior of Genetic Algorithms for Social Modeling
Ludo Waltman and Nees Jan van Eck
ERS-2009-011-LIS
http://hdl.handle.net/1765/15181

A Taxonomy of Bibliometric Performance Indicators Based on the Property of Consistency
Ludo Waltman and Nees Jan van Eck
ERS-2009-014-LIS
http://hdl.handle.net/1765/15182

A Stochastic Dynamic Programming Approach to Revenue Management in a Make-to-Stock Production System
Rainer Quante, Moritz Fleischmann, and Herbert Meyr
ERS-2009-015-LIS
http://hdl.handle.net/1765/15183

Some Comments on Egghe’s Derivation of the Impact Factor Distribution
Ludo Waltman and Nees Jan van Eck
ERS-2009-016-LIS
http://hdl.handle.net/1765/15184

---

[*] A complete overview of the ERIM Report Series Research in Management:
https://ep.eur.nl/handle/1765/1

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship