

2628



Revenue management for operations with urgent orders

Philippe Chevalier, Alejandro Lamas, Liang Lu
and Tanja Mlinar

CORE

Voie du Roman Pays 34, L1.03.01

B-1348 Louvain-la-Neuve, Belgium.

Tel (32 10) 47 43 04

Fax (32 10) 47 43 01

E-mail: corestat-library@uclouvain.be

<http://www.uclouvain.be/en-44509.html>



Decision Support

Revenue management for operations with urgent orders

Philippe Chevalier^{a,*}, Alejandro Lamas^a, Liang Lu^b, Tanja Mlinar^a^a Louvain School of Management – Center for Operations Research and Econometrics, Place des Doyens 1, Université catholique de Louvain, 1348 Louvain-la-Neuve, Belgium^b Logistics Research Centre, School of Management and Languages, Heriot-Watt University, Edinburgh EH14 4AS, UK

ARTICLE INFO

Article history:

Received 6 September 2013

Accepted 9 July 2014

Available online 23 July 2014

Keywords:

Revenue management

Order acceptance

Markov Decision Process

Heuristics

Flexibility

ABSTRACT

This article is motivated by the case of a company manufacturing industrial equipment that faces two types of demand: on the one hand there are the so-called regular orders for installations or refurbishing of existing facilities, these orders have a relatively long lead time; on the other hand there are urgent orders mostly related to spare parts when a facility has a breakdown, the delay in such case is much shorter but higher margins can be obtained. We study the order acceptance problem for a firm that serves two classes of demand over an infinite horizon. The firm has to decide whether to accept a regular order (or equivalently how much capacity to set aside for urgent orders) in order to maximize its profit. We formulate this problem as a multi-dimensional Markovian Decision Process (MDP). We propose a family of approximate formulations to reduce the dimension of the state space via aggregation. We show how our approach can be used to compute bounds on the profit associated with the optimal order acceptance policy. Finally, we show that the value of revenue management is commensurate with the operational flexibility of the firm.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

This article is motivated by the case of a cast iron manufacturer. This company is specialized in the production of cast iron pieces for industrial equipments. Most of its orders are either for the preventive maintenance of installations or for the building of new facilities. Such projects are scheduled with long lead times, but it is very important that the pieces be delivered on time because the plant where the pieces have to be installed will have to be (at least partially) stopped for the maintenance or installation activities to take place and obviously the duration of such stoppage should be minimized. A different type of orders received by the company corresponds to corrective maintenance when a breakdown occurs in a plant, in those cases a new cast iron piece is needed to restart the facility. Given that production at the customer is stopped because of the breakdown, a much faster service is required but the company can charge higher prices for such “emergency” orders. Moreover, the bargaining power of the customer is much weaker in such circumstances. Given its finite production capacity, the company cannot always accept all orders and should sometimes forego a regular order in order to keep some

possibility to accept an urgent order later on. This dilemma is faced by many suppliers confronted with urgent requests that are potentially very profitable but could be very disruptive if not taken into account in their planning. Typical examples from the service sector include heating ventilation air-conditioning companies. The installation of new systems is typically a large project with a relatively long lead time. In contrast, when a system fails it could block the operations of a customer that is then willing to pay a higher fee for speedy action. In a very different context, suppliers of the fashion industry are known to combine orders from large vendors and more profitable orders for high fashion clothes which often require fast delivery (see e.g. deB Harris & Pinder, 1995; Barut & Sridharan, 2005). The question facing the supplier is how much capacity should be set aside for the urgent high margin demand, given the inherent unpredictability of this type of orders.

To address this question, we build a model with a supplier that handles two demand classes, that we will refer to as regular and urgent respectively. The regular orders are typically characterized by longer processing times, longer lead times but lower margins, while the urgent orders have shorter processing times, shorter lead times and higher margins. If the supplier accepts orders without foresight it is likely that at some point, when an urgent order arrives, the supplier will be unable to accept this order as her short term capacity is already entirely committed for regular orders (that were booked earlier with a longer lead time). Given the difference in margin between the two classes, this situation causes some loss

* Corresponding author. Tel.: +32 10 478396; fax: +32 10 474301.

E-mail addresses: philippe.chevalier@uclouvain.be (P. Chevalier), alejandro.lamas@uclouvain.be (A. Lamas), nicowish@gmail.com (L. Lu), tanja.mlinar@uclouvain.be (T. Mlinar).

of revenue. On the other hand, rejecting a regular order in anticipation for potential urgent orders that do not materialize, also causes some revenue loss.

This trade-off has clear similarities with other revenue management problems. The distinct feature is that when an order is accepted the supplier keeps some flexibility. For example, if a regular order necessitates 10 days of work and the lead time is 20 days, in most cases the customer does not care about the days during which the order is effectively produced as long as it is finished on time. In revenue management terms, if we consider that the capacity available during each period is a distinct *product*, an order requiring more than one period of work is in fact reserving several *products*. But the supplier has some flexibility in assigning the *products* to the order and does not need to make a commitment at the time of reservation. Such consideration is reasonable and represent the reality at many firms in which jobs can be interrupted and resumed in later periods. One can draw a parallelism with the network revenue management problem but where the supplier can accept a reservation without committing to specific legs in the network, the only commitment is on the origin and destination points in the network.

Our contributions are summarized as follows:

- We derive a Markov Decision Process (MDP) formulation of the order acceptance problem of two customer classes with lead time constraints. Likewise for the network revenue management problem, the size of this formulation quickly excludes the possibility of solving it exactly for larger size instances. We develop a family of approximate formulations parameterizable to range from a coarse approximation to the original full formulation. This makes it possible to choose between speed of solution and precision of result. What is of particular interest is that for each formulation, we can compute an upper and lower bound on the exact result. This last feature is rather uncommon for revenue management problems and is particularly interesting to make sure the adequate level of approximation is chosen in the proposed family of formulations.
- Through a numerical study we show that the proposed heuristics allow to obtain near-optimal solutions in a tractable time. We also show how the potential benefit of revenue management is commensurate with operational flexibility. In our setting operational flexibility consists in the slack between the promised lead time for an order class and the processing time needed for such order.

The remainder of the paper is structured as follows. In Section 2 we discuss related work. In Section 3 we provide a detailed description of our problem. In Section 4 we introduce an MDP formulation and discuss its resolution to obtain the optimal admission policy. In Section 5 we propose two heuristic formulations of the problem based on different levels of state aggregation and report on a numerical study of our proposed formulations in Section 6. Section 7 investigates the impact of operational flexibility on the benefit of revenue management. Finally, Section 8 summarizes the main conclusions and identifies future research directions.

2. Literature review

Our work belongs to the growing literature of Perishable Asset Revenue Management (PARM) which deals with the problem of allocation of scarce resources to different demand classes. Talluri and van Ryzin (2004) give a comprehensive overview of this topic. The first applications were for the airline industry by Littlewood (1972), and extended by Belobaba (1987), Wollmer (1992), and

Brumelle and McGill (1993). In addition to airlines, typical service applications are in hotel management and car rental (Kimes, 1989; Bertsimas & Popescu, 2003; Talluri & van Ryzin, 2004; Bitran & Mondschein, 1995; Geraghty & Johnson, 1997). Gradually, new applications appeared for very different environments such as: MTO manufacturing (Balakrishnan, Sridharan, & Patterson, 1996; Barut & Sridharan, 2005; Spengler, Rehkoopf, & Thomas, 2007), project management (Herbots, Herroelen, & Leus, 2007; Herbots, Herroelen, & Leus, 2010) and health care (Gupta & Wang, 2008; Dobson, Hasija, & Pinker, 2011).

A stream of literature related to lead time decisions focuses on due-date quotation and scheduling problems in order to allocate the available capacity to incoming orders (see e.g. Kaminsky & Hochbaum, 2004; Keskinocak & Tayur, 2004, for an extensive literature review). Most of these works assign dynamically lead times to incoming orders depending on the state of the system and sequencing policies (see e.g. Duenyas, 1995; Duenyas & Hopp, 1995; Kapuscinski & Tayur, 2007). Kapuscinski and Tayur (2007) propose a dynamic programming approach to address the problem of lead-time quotation for two demand classes when customers are not equally sensitive to waiting. Lead time quotation is used to ensure that the capacity is allocated in such a way that all demands can be delivered on time. So, firms can change the quoted lead time based on the system state. Motivated by the prevalence of static lead time policies (see e.g. Cheng & Gupta, 1989; Hopp & Sturgis, 2000; Keskinocak & Tayur, 2004), we consider a different problem in which lead times are constant and exogenously given. Since the capacity may not be enough to cater all demands, the decision becomes accepting or rejecting orders depending on the system state.

Gupta and Wang (2007) consider an order acceptance problem in which the lead time requirement for regular orders is modeled as a soft operational constraint. It is assumed that tardiness cost is incurred if regular orders are not filled within their lead time window while urgent orders must be filled in the current period once accepted. The authors propose a multi-dimensional MDP whose optimal solution turns out to be a threshold based policy. This solution property is a consequence of the simplicity of their model setting, which leads to a well-structured value function. In contrast, our model is more general, assuming some flexibility in catering urgent orders – the lead time for urgent order does not necessarily need to be one. The state space in our problem is defined in a different way than in Gupta and Wang (2007), because their representation involves tracking the backlogging information for every demand class and becomes particularly inefficient when there are multiple demand classes, which greatly limits its application. Our representation “encodes” in itself the capacity allocation decisions and therefore is more efficient.

The following references focus on acceptance decision problems where the lead times must be strictly respected, as in our case. Germs and Van Foreest (2011) study an order acceptance problem with multiple customer classes with a common lead time, setup times and scheduling constraints. The problem is modeled as a Markov chain controlled by a threshold policy. The authors provide a numerical study for small instances which are computationally tractable. In contrast to their work, we provide efficient alternative methods to treat the state space explosion. Barut and Sridharan (2005) study an order acceptance problem involving multiple demand classes that differ in terms of price, lead time and demand pattern. The authors propose a nested rationing policy which fulfills incoming orders as much as possible while preserving a certain level of capacity for more profitable future orders. The proposed policy is computed using a myopic heuristic method that does not take the evolution of the capacity into account. Consequently, the efficiency of the heuristic is hurt by the simplified estimation of the future available capacity. Our formulation keeps track more

accurately of the capacity evolution for this type of problem, leading to highly efficient policies.

Literature related to development and implementation of different approximate methods to reduce the computation time required to solve large instances of dynamic problems is abundant. Examples of methods developed to solve MDPs include state aggregation techniques (see e.g. Mendelsohn, 1982; Bean, Birge, & Smith, 1987; Aldhaferi & Khalil, 1991; Hu & Wu, 2000; Zhang & Baras, 2001; Van Roy, 2006; Jia, 2011), embedding/time aggregation approaches via value and policy iteration (see e.g. Cao, Ren, Bhatnagar, Fu, & Marcus, 2002; Leizarowitz & Schwartz, 2008; Sun, Zao, & Lux, 2007; Arruda & Fragoso, 2011; Cheng & Zhang, 2012) and convergence acceleration methods for value iteration (Almudevar & Arruda, 2012). In this paper, we apply state aggregation techniques for a problem of order acceptance when customer classes differ in their lead times. There is, however, a difference between the previous works and our paper: they are designed for solving MDP in general forms, while our aggregation method is motivated by the specific characteristics of the problem of interest. In this sense, our work is in the same spirit with Xu, Gao, and Ou (2007). The authors model a ticket queue problem with a Markov chain and used a state aggregation technique to reduce the size of the state space based on the structure of the ticket queue. In this paper, we implement state aggregation techniques to aggregate some information of the states of the system within the lead time window of the incoming orders. Thereby, the infinite horizon problem is solved, and the state space of the problem is reduced.

There are some similarities with the network revenue management problem where the dimension of the MDP models quickly makes it impossible to solve exactly even small size problems. Consequently, research in this area has concentrated on different approximation techniques, notable examples of this stream of investigation include Bitran and Mondschein (1995), Talluri and van Ryzin (1998), Bertsimas and Popescu (2003), Adelman (2007), Kunnumkal and Topaloglu (2008), Zhang and Adelman (2009), and Zhang (2011). The flexibility aspect present in the problem studied here calls for different types of approximations.

3. Model description

We consider the order acceptance problem for a firm serving two customer classes with different profit margins and lead times. The time horizon is infinite and consists of discrete periods. In each period, the firm is subject to a limited processing capacity, which is normalized to 1.

We consider that all uncertainty about the processing time is known when the order is placed. This assumption is motivated by two reasons: on the one hand, in practice most uncertainty is resolved during the ordering process; on the other hand, if the remaining uncertainty is too large it is not possible to promise due dates without either large safety lead times or a low utilization. In the cases that motivated our work, we observed that the remaining processing time uncertainty after the order is placed is dealt with using some type of recourse action such as overtime or renegotiation of the due-date. These recourse actions play only a secondary role in the management of capacity and are beyond the scope of our investigation.

The two classes of demand will be denominated, urgent and regular orders (indexed by $k = \{1, 2\}$, respectively). The demand (in terms of processing time) for class k at time t will be denoted D_{kt} . We suppose the random variables take integer values and are iid between time periods and independent between classes. If $D_{kt} = 0$ there is no demand for class k in period t . The profit margin per unit capacity of a class k customer is r_k and its lead time is L_k . Urgent orders are more lucrative but come with a short lead time, while regular orders are not as profitable, yet have a looser lead




Fig. 1. Partition of lead time window of regular demand.

time. Accordingly, we assume $r_1 > r_2 > 0$ and $L_1 < L_2$. Fig. 1 shows the structure of the problem. We suppose that only a single order of each class can arrive during any period t , this means that the demand D_{kt} cannot be partially accepted, the firm either accepts the order and hence commits to deliver the order before its due date or declines the order and gets no revenue. On-time delivery positively influences customer experience and reinforces the long term image of firms. When confirming orders, firms usually allocate sufficient capacity for processing the orders in a timely manner. Since the processing capacity is deterministic in our model, it is reasonable to plan with no tardiness.

Consequently, the main decision faced by the firm is how many orders to accept in order to maximize its long-run net profit. However, we restrict our attention to the order acceptance decisions of regular orders. We assume that regular orders can be selectively rejected while urgent orders can only be passively rejected as a result of insufficient available capacity. This is because urgent orders often arise in emergency circumstances, and thus are granted with priority over regular orders.¹

In each period, the sequence of events is as follows. A regular order, if any, arrives first, and the firm decides whether to accept it or not. Then, an urgent order, if any, arrives, it will be accepted as long as there is enough available capacity. Finally, the firm uses the capacity in the current period for processing the order with the earliest due date. In fact, the sequence of arrivals does not matter; one can assume that the urgent order arrives before the regular order, or both arrive simultaneously, without complicating the model.

In the remainder of the paper we will assume a somewhat simpler structure for the demand. We assume

$$D_{kt} = \begin{cases} 0, & \text{with probability } 1 - p_k, \\ B_k, & \text{with probability } p_k, \end{cases}$$

and thus the inter-arrival times are geometrically distributed in each demand class. This simplifies the notations in the following sections and our numerical tests indicate that the distribution of the demand has no qualitative impact on our results.

4. MDP formulation for the optimal steady-state policy

The crux of developing the MDP formulation of this problem is to find an efficient representation of the usage of future capacity. Our representation builds on the idea that when an order is accepted, some future capacity will be “reserved”. Reservation refers to *provisionally* allocating the available capacity of the latest periods to process an order without incurring tardiness. This gives the maximum flexibility to accept the future orders. This notion is also used in the model formulation of Kapuscinski and Tayur (2007). Note that, the allocation is provisional because the firm may start processing earlier the order if there is no order to deliver before this one. When an order is (partially) processed before its provisionally allocated time slot, some capacity is freed to process

¹ The modeling framework in our paper is general enough that the option of rejecting urgent orders can be easily incorporated without substantial modification of the model.

future orders. For example, when accepting an order with a lead time of 5 and a size of 2, the capacity of 4th and 5th periods from the current period shall be reserved (provided that they are available for reservation), if the order will be processed as soon as possible in the earliest-due-date (EDD) sequence.

To reserve capacity for new orders, the firm needs to calculate the total available capacity for reservation within the lead time windows of regular and urgent orders. Since the lead time window of urgent orders is contained within that of the regular orders, it suffices to calculate the total available capacity within L_1 periods, and that between $(L_1 + 1)$ th and L_2 th periods, i.e. Interval I and Interval II of Fig. 1, respectively. However, this aggregate information cannot fully characterize the evolution of the system. Note that the $(L_1 + 1)$ th period, (i.e., the first period of Interval II), will be shifted by one period and thus become the L_1 th period in the next period, (i.e., the last period of Interval I). Without the information regarding how the reserved capacity is distributed in Interval II, it is impossible to know whether the shifted capacity is reserved or not in order to update the available capacity in Interval I and II in the next period. Therefore, it is necessary to keep track of the distributional information in Interval II, but only of aggregate information in Interval I.

We now introduce the notation to be used in our formulation.

- $\mathbf{x}$: is the reservation vector, it keeps track of capacity that has been reserved for processing; $\mathbf{x}[0] \in \{0, 1, \dots, L_1\}$ denotes the total reserved capacity until the L_1 th period, (i.e., in Interval I). For $j = 1, 2, \dots, L_2 - L_1$, $\mathbf{x}[j] = 1$ if the capacity of $(L_1 + j)$ th period is reserved, and $\mathbf{x}[j] = 0$ otherwise. Note that in a reservation vector we do not distinguish whether the capacity is reserved for urgent orders or regular orders.
- $\mathbf{y}$: is the cumulative (available capacity) vector² of $\mathbf{x}$; for $j = 0, 1, \dots, L_2 - L_1$, $\mathbf{y}[j]$ denotes the total available capacity until the $(L_1 + j)$ th period, i.e., $\mathbf{y}[j] = L_1 + j - \sum_{i=0}^j \mathbf{x}[i]$. For a given cumulative vector $\mathbf{y}$, its corresponding reservation vector $\mathbf{x}$ can be calculated as follows: $\mathbf{x}[0] = L_1 - \mathbf{y}[0]$, and for $j = 1, 2, \dots, L_2 - L_1$, $\mathbf{x}[j] = \mathbf{y}[j - 1] - \mathbf{y}[j] + 1$.
- a : represents the admission decision for regular orders; $a = 1$ if the firm “admits” the regular order, and $a = 0$ otherwise.
- $\mathbf{D} \equiv (D_1, D_2)$: is the demand vector in each period.
- $R(\mathbf{x}, \mathbf{D}, a)$: denotes the profit generated from $\mathbf{D}$ for a given admission decision a , if the reservation vector at the beginning of the current period is $\mathbf{x}$.

We define the system state as the reservation vector $\mathbf{x}$ at the beginning of a period, before the arrival of regular and urgent orders. It is easy to check that $\mathbf{x}[L_2 - L_1] = 0$ for any system state, because the last element of the system state cannot be reserved by orders that arrived in earlier periods. Thus, the system state space essentially involves $L_2 - L_1$ variables and its size is $(L_1 + 1) \cdot 2^{L_2 - L_1 - 1}$. Let $\tilde{\mathbf{x}}$ be the reservation vector updated from $\mathbf{x}$ after accepting/rejecting D_2 . Additionally, let $\hat{\mathbf{x}}$ be the system state in the next period, this is the reservation vector updated from $\tilde{\mathbf{x}}$ after accepting/rejecting D_1 and processing, if any, is carried out in the current period. Further, let $\tilde{\mathbf{y}}$ and $\hat{\mathbf{y}}$ be the corresponding cumulative vectors of $\tilde{\mathbf{x}}$ and $\hat{\mathbf{x}}$, respectively. Thus, $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ are functions of $\mathbf{x}, \mathbf{D}$ and a : $\hat{\mathbf{x}} = \hat{\mathbf{x}}(\mathbf{x}, \mathbf{D}, a)$, $\hat{\mathbf{y}} = \hat{\mathbf{y}}(\mathbf{x}, \mathbf{D}, a)$. According to the average reward criteria (Ross, 1983; Bertsekas, 1995; Puterman, 2005), we provide the dynamic programming formulation as follows:

$$V(\mathbf{x}) + g = \mathbb{E}_{\mathbf{D}} \left\{ \max_a \{ R(\mathbf{x}, \mathbf{D}, a) + V(\hat{\mathbf{x}}(\mathbf{x}, \mathbf{D}, a)) \} \right\}, \quad \forall \mathbf{x}, \quad (1)$$

² $\mathbf{y}$ is introduced to facilitate the notations related to transitions between system states, which are composed of reservation vectors. There is a one-to-one correspondence between a reservation vector and its cumulative vector.

together with the dynamics described by Eqs. (2)–(7),

$$\tilde{\mathbf{y}}[j] = \begin{cases} \mathbf{y}[j] - (D_2 - (\mathbf{y}[L_2 - L_1] - \mathbf{y}[j]))^+, & \text{if } a = 1 \text{ and } \mathbf{y}[L_2 - L_1] \geq D_2, \\ \mathbf{y}[j], & \text{otherwise,} \end{cases} \quad \text{for } j = 0, 1, \dots, L_2 - L_1, \quad (2)$$

$$\hat{\mathbf{y}}[j] = \begin{cases} \min \{ \tilde{\mathbf{y}}[j + 1] - D_1, L_1 + j \}, & \text{if } \tilde{\mathbf{y}}[0] \geq D_1, \\ \min \{ \tilde{\mathbf{y}}[j + 1], L_1 + j \}, & \text{otherwise,} \end{cases} \quad \text{for } j = 1, 2, \dots, L_2 - L_1 - 1, \quad (3)$$

$$\hat{\mathbf{y}}[L_2 - L_1] = \hat{\mathbf{y}}[L_2 - L_1 - 1] + 1, \quad (4)$$

$$R(\mathbf{x}, \mathbf{D}, a) = R_1(\tilde{\mathbf{x}}, D_1) + R_2(\mathbf{x}, D_2, a), \quad (5)$$

$$R_2(\mathbf{x}, D_2, a) = \begin{cases} r_2 \cdot D_2, & \text{if } a = 1 \text{ and } \mathbf{y}[L_2 - L_1] \geq D_2, \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

$$R_1(\tilde{\mathbf{x}}, D_1) = \begin{cases} r_1 \cdot D_1, & \text{if } \tilde{\mathbf{y}}[0] \geq D_1, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

In the solution of Eq. (1), g represents the steady state expected profit per period. Eqs. (2)–(4) characterize the transition between system states $\mathbf{x}$ and $\hat{\mathbf{x}}$, with the help of their corresponding cumulative vectors $\mathbf{y}$ and $\hat{\mathbf{y}}$. Specifically, Eq. (2) describes how $\mathbf{y}$ is updated to $\tilde{\mathbf{y}}$, and Eqs. (3) and (4) further describe how $\tilde{\mathbf{y}}$ is updated to $\hat{\mathbf{y}}$. When updating the system state, all accepted orders are scheduled in the reservation vector as late as possible within their lead time windows to allow for maximal flexibility to process the orders. Eqs. (5)–(7) calculate the profit generated during transitions. Specifically, $R_1(\tilde{\mathbf{x}}, D_1)$ represents the profit from urgent orders and $R_2(\mathbf{x}, D_2, a)$ represents the profit from regular orders.

Example 1. Consider an instance with the following parameters: $L_1 = 4, L_2 = 8, B_1 = 2, B_2 = 3$. Suppose the reservation vector at the beginning of the current period is $\mathbf{x} = [2, 0, 1, 1, 0]$, meaning there are 2 units of capacity reserved in Interval I and 2 units reserved in Interval II, totaling 4 units of available capacity for fulfilling regular orders. If a regular order arrives ($D_2 = B_2 = 3$), and the decision is to accept it, 1 unit of available capacity in Interval I and 2 units in Interval II will be reserved for processing the order, and the reservation vector will be updated to $\tilde{\mathbf{x}} = [3, 1, 1, 1, 1]$ or $\tilde{\mathbf{y}} = [1, 1, 1, 1, 1]$ (Eq. (2)). However, this leaves with only 1 unit of available capacity within the lead time window of urgent orders, (i.e., $\tilde{\mathbf{y}}[0] = 1$) and therefore, there is not enough room to accommodate any urgent order (since $B_1 = 2$). Finally, the capacity of the current period is used to process one unit of order in Interval I, and one unit of reserved capacity is shifted from Interval II into Interval I. Therefore, the reservation vector observed at the beginning of next period becomes $\hat{\mathbf{x}} = [3, 1, 1, 1, 0]$ or $\hat{\mathbf{y}} = [1, 1, 1, 1, 2]$ (Eqs. (3) and (4)).

5. State reduction heuristics

The formulation described in Section 4 leads to an optimal steady-state policy for accepting/rejecting regular orders. However, the size of the formulation, i.e., the number of states, increases exponentially in $L_2 - L_1$. For example, if $L_1 = 5$ and $L_2 = 25$, the resulting formulation has 3,145,728 states. Thus, we seek to develop more compact heuristic formulations.

The complexity of the formulation is largely related to keeping track of the distributional information in Interval II. To reduce the complexity, one plausible idea is to somehow aggregate the distributional information in Interval II, so that the MDP formulation can be reduced to involve fewer states. The reduced formulation can then be used to generate heuristic policies.

We start with the Full Aggregation Heuristic (FAH) that completely ignores the distributional information in Interval II. Though

being extremely compact, the FAH does not always lead to near-optimal solutions. Consequently, we propose the Partial Aggregation Heuristic (PAH) which keeps the most “important” distributional information intact while aggregating the rest in Interval II. A major advantage of this approach is that one can easily control the tradeoff between computation efforts and optimality.

5.1. Full aggregation heuristic

We propose a new MDP formulation based on *aggregate reservation vectors* as opposed to reservation vectors in the original formulation. For a reservation vector $\mathbf{x}$, we define its corresponding aggregate reservation vector as $\mathbf{x}_f = (\mathbf{x}_f[0], \mathbf{x}_f[1])$, in which $\mathbf{x}_f[0] = \mathbf{x}[0]$ and $\mathbf{x}_f[1] = \sum_{j=1}^{L_2-L_1} \mathbf{x}[j]$, i.e., $\mathbf{x}_f[0]$ corresponds to the total reserved capacity in Interval I, and $\mathbf{x}_f[1]$ corresponds to that in Interval II. Note that there can be multiple reservation vectors mapping to the same aggregate reservation vector. Further, we define $\mathbf{y}_f = (\mathbf{y}_f[0], \mathbf{y}_f[1])$ as the *aggregate cumulative vector*, in which $\mathbf{y}_f[0] = L_1 - \mathbf{x}_f[0]$ and $\mathbf{y}_f[1] = L_2 - \mathbf{x}_f[0] - \mathbf{x}_f[1]$, i.e., $\mathbf{y}_f[0]$ (respectively $\mathbf{y}_f[1]$) corresponds to the total available capacity in the lead time window of urgent (respectively regular) demands. The one-to-one mapping between an aggregate reservation vector and its aggregate cumulative vector still holds.

The new system state is defined as the aggregate reservation vector in the beginning of a period, and therefore only involves two dimensions. One way to think of the new system state is that each one groups multiple system states in the original formulation into a “super state” (Fig. 2), resulting in a significantly shrunk state space. In other words, the size of the aggregated state space is reduced from $(L_1 + 1) \cdot 2^{(L_2-L_1-1)}$ to $(L_1 + 1) \cdot (L_2 - L_1)$.

Then we characterize how one new system state transits to another for a given admission policy a and demand pattern D under the new state space. Note that transitions triggered by accepting/rejecting demands and processing can be easily characterized, but there is one step of transition that cannot be properly defined due to the aggregation, that is, how the aggregated reservation vector at the end of one period evolves to the one at the beginning of the next period. As discussed previously, without distributional information in Interval II, it is impossible to decide whether the shifted unit of capacity is reserved or not. To address this issue, we heuristically account for the transfer of capacity based on assumptions regarding how the reserved capacity is distributed in Interval II. We consider three scenarios as follows:

- **Optimistic scenario.** Assuming that all reserved capacity in Interval II (excluding the last unit) are distributed *as late as possible*, the shifted unit of capacity is reserved only if Interval II

(excluding the last unit) is “full”, i.e., all the capacity in this area has been reserved. “Optimistic” refers to the fact that the assumption results in an overestimation of the available capacity in Interval I.

- **Pessimistic scenario.** Assuming that all reserved capacity in Interval II (excluding the last unit) are distributed *as early as possible*, the shifted unit of capacity is available only if Interval II (excluding the last unit) is “empty”, i.e., none of the capacity in this area has been reserved. “Pessimistic” refers to the fact that the assumption results in an underestimation of the available capacity in Interval I.
- **Realistic scenario.** Assuming that all reserved capacity in Interval II (excluding the last unit) are distributed *uniformly*, the shifted unit of capacity is reserved with a probability, defined as the proportion of reserved capacity in Interval II (excluding the last unit).

Among the three scenarios, we regard the realistic scenario to be closest to the reality. Note that the distribution in Interval II is influenced by two forces: reservation and updating from one period to the next. Reservation fills the capacity from the end to the front, while updating from one period to the next shifts the capacity forward by one period. Hence, we know that the real distribution must be less extreme than the ones in the optimistic and pessimistic scenarios. Nonetheless, the scenario leading to a better policy is yet to be examined.

These assumptions do not deal with the last unit of capacity in Interval II, because whether it is reserved or not can be explicitly characterized without introducing additional dimensions: it is always available at the beginning of one period, and it will become reserved whenever some regular demand is accepted in the current period.

Let $\tilde{\mathbf{x}}_f, \hat{\mathbf{x}}_f, \tilde{\mathbf{y}}_f$ and $\hat{\mathbf{y}}_f$ be the aggregate versions of vectors $\tilde{\mathbf{x}}, \hat{\mathbf{x}}, \tilde{\mathbf{y}}$ and $\hat{\mathbf{y}}$, respectively. In addition, we define $\bar{\mathbf{y}}_f$ as the aggregate cumulative vector after accepting/rejecting urgent orders but before processing is carried out, i.e. $\bar{\mathbf{y}}_f$ serves as an intermediary between $\tilde{\mathbf{y}}_f$ and $\hat{\mathbf{y}}_f$. We characterize the transition between $\mathbf{x}_f$ and $\hat{\mathbf{x}}_f$ as follows.

Accepting/rejecting a regular order:

$$(\tilde{\mathbf{y}}_f[0], \tilde{\mathbf{y}}_f[1]) = \begin{cases} (\mathbf{y}_f[0] - (D_2 - (\mathbf{y}_f[1] - \mathbf{y}_f[0]))^+, \mathbf{y}_f[1] - D_2), & \text{if } a = 1 \text{ and } \mathbf{y}_f[1] \geq D_2, \\ (\mathbf{y}_f[0], \mathbf{y}_f[1]), & \text{otherwise.} \end{cases} \quad (8)$$

Accepting/rejecting an urgent order:

$$(\bar{\mathbf{y}}_f[0], \bar{\mathbf{y}}_f[1]) = \begin{cases} (\tilde{\mathbf{y}}_f[0] - D_1, \tilde{\mathbf{y}}_f[1] - D_1), & \text{if } \tilde{\mathbf{y}}_f[0] \geq D_1, \\ (\tilde{\mathbf{y}}_f[0], \tilde{\mathbf{y}}_f[1]), & \text{otherwise.} \end{cases} \quad (9)$$

Processing and updating to the next period:

$$(\hat{\mathbf{y}}_f[0], \hat{\mathbf{y}}_f[1]) = \begin{cases} (\min\{\bar{\mathbf{y}}_f[0], L_1 - 1\}, \min\{\bar{\mathbf{y}}_f[1] + 1, L_2\}), & \text{with probability } \pi \\ (\min\{\bar{\mathbf{y}}_f[0], L_1 - 1\} + 1, \min\{\bar{\mathbf{y}}_f[1] + 1, L_2\}), & \text{with probability } 1 - \pi. \end{cases} \quad (10)$$

in which the value of π is contingent on whether the last unit of capacity in Interval II is reserved or not. Some additional notation follows: let ψ be the total available capacity in Interval II after processing, i.e., $\psi = \bar{\mathbf{y}}_f[1] - \min\{\bar{\mathbf{y}}_f[0], L_1 - 1\}$. Next, we discuss the value of π in the different cases.

Case 1: $\tilde{\mathbf{y}}_f = \mathbf{y}_f$ (the last unit of capacity in Interval II is available)

- In the optimistic scenario, if $\psi = 1$, then $\pi = 1$; otherwise, $\pi = 0$.
- In the pessimistic scenario, if $\psi \neq L_2 - L_1$, then $\pi = 1$; otherwise, $\pi = 0$.
- In the realistic scenario, $\pi = 1 - \frac{\psi - 1}{L_2 - L_1 - 1}$.




Fig. 2. The new system states defined by “Super States”. This illustrative example shows the states whose first element is 1 for an instance in which $L_2 - L_1 = 4$.

Case 2: $\hat{\mathbf{y}}_f \neq \mathbf{y}_f$ (the last unit of capacity in Interval II is reserved)

- In the optimistic scenario, if $\psi = 0$, then $\pi = 1$; otherwise, $\pi = 0$.
- In the pessimistic scenario, if $\psi \neq L_2 - L_1 - 1$, then $\pi = 1$; otherwise, $\pi = 0$.
- In the realistic scenario, $\pi = 1 - \frac{\psi}{L_2 - L_1 - 1}$.

Fig. 3 shows an example of how an aggregated reservation vector evolves in different scenarios. In this example, $L_1 = 2$ and $L_2 = 7$, the aggregated reservation vector at the end of one period, (i.e., the one after accepting/rejecting demands and processing) is (1, 3). Assuming some regular demand is accepted in this period, the last unit in Interval II is reserved. In the optimistic scenario, the shifted unit of capacity is available because the Interval II (excluding the last unit) is not full, and thus the aggregated reservation vector at the start of next period is $\hat{\mathbf{y}}_f = (1, 3)$; in the pessimistic scenario, the shifted unit is reserved, leading to $\hat{\mathbf{y}}_f = (2, 2)$; in realistic scenario, with probability $\pi = 2/4 = 0.5$, the shifted unit is reserved, leading to $\hat{\mathbf{y}}_f = (2, 2)$, and with probability $1 - \pi = 0.5$, the shifted unit is available, leading to $\hat{\mathbf{y}}_f = (1, 3)$.

The profit generated in each aggregated state can be defined in a similar fashion as in (5)–(7), and a new MDP formulation can be further obtained from optimality Eq. (1) by replacing the system states with their aggregated versions. For simplicity, we omit their representations.

Finally, once an optimal admission policy for the new MDP is obtained, a heuristic policy can be constructed in the following way: for each aggregate state, apply its admission decision to all the corresponding states in the original formulation.

5.2. Partial aggregation heuristic

The lack of accuracy in characterizing the transitions between aggregate states sometimes leads to significant profit gaps in comparison with the optimal profit, as we will show in the numerical experiments. We wonder whether a more precise characterization of the distributional information in Interval II could improve the performance of the heuristic. We propose a generalized heuristic, the PAH, that partially aggregates the distributional information in Interval II. Specifically, we split Interval II into two parts: (i) Sub-Interval II-A, where the distributional information is precisely tracked and (ii) Sub-Interval II-B, where the distributional information is fully aggregated. Fig. 4 shows an example of reservation vectors for different levels of aggregation.

However, we face the same issue as in the full aggregation case, that is, how to account for the transfer of capacity between Sub-Interval II-A and Sub-Interval II-B. Again, we address this issue by heuristic approaches in which the reserved capacity in Sub-Interval II-B is distributed according to the pessimistic, realistic and optimistic scenarios as described in Section 5.1. The MDP formulation is derived in the same way; for simplicity, we omit their presentation.

Let $z \in \{0, \dots, L_2 - L_1 - 1\}$ be the number of periods in Sub-Interval II-A, hereafter referred to as the disaggregation level. This parameter controls the tracking accuracy of the distributional information in Sub-Interval II-A: if $z = 0$, the resulting formulation is actually identical to the full aggregation; if $z \geq L_2 - L_1 - 2$, the resulting formulation coincides with the one without aggregation that provides the optimal steady-state policy. For a given disaggregation level z , the size of the resulting state space is $(L_1 + 1) \cdot (L_2 - L_1 - z) \cdot 2^z$. Obviously more disaggregation leads to larger state space. Moreover, we have the following result comparing the optimal profits for different partially aggregated MDP formulations.

Proposition 1. Let $g_o^*(z)$, $g_p^*(z)$ and $g_r^*(z)$ be the optimal expected profit obtained by solving the partial aggregation model with a level of disaggregation z for the optimistic, pessimistic and realistic scenarios, respectively. We have

$$g_p^*(z) \leq \{g_o^*(z), g_r^*(z)\} \leq g_o^*(z), \forall z,$$

and $g_o^*(z)$ and $g_p^*(z)$ are non-increasing and non-decreasing with z , respectively.

Proof. The proof is based on two facts: (i) it is always feasible to shorten the lead times of the already accepted orders given that there is available capacity in earlier periods and (ii) for any aggregate state, the pessimistic scenario implies shorter lead times for already accepted orders than the realistic scenario, which further implies shorter lead times than the optimistic scenario. Therefore, any transition between two aggregate states of scenarios implying shorter lead times can also be achieved with scenarios implying more relaxed lead times. Consequently, for any sample path for the formulation of scenarios implying shorter lead times, we can obtain an identical sample path that provides the same profit with more relaxed scenarios. Thus, the profit for the formulation of the optimistic scenario is at least as large as the profit for the formulation of the realistic scenario, which is further at least as large as the profit for the formulation of the pessimistic scenario. The same logic can be applied for proving the monotonicity property of $g_o^*(z)$ and $g_p^*(z)$, i.e., more aggregation implies shorter (longer) lead times for the pessimistic (optimistic) scenario. $\square$




Fig. 3. Transition of aggregated reservation vectors between two periods.

	Interval I	Interval II									
Without aggregation $x=[2,1,1,0,1,0,0,1,1,0]$	2	1	1	0	1	0	0	1	1	0	
Full aggregation Optimistic scenario $x=[2,5]$	2	0	0	0	1	1	1	1	1	0	
Partial aggregation Optimistic scenario, $z = 5$ $x=[2,1,1,0,1,0,2]$	2	1	1	0	1	0	0	1	1	0	
		Sub-Interval II-A					Sub-Interval II-B				

Fig. 4. An example of reservation vectors for different levels of aggregation.

Proposition 1 indicates that $g_p^*(z)$ and $g_o^*(z)$ are lower and upper bounds for $g_r^*(z)$ and g^* , but $g_r^*(z)$ and g^* are not directly comparable. It also shows that the proposed bounds get tighter as the level of disaggregation z increases.

6. Numerical results

In this section we investigate the actual performance of the policies presented above, more specifically we analyze:

- the computation time required to obtain an acceptance policy with the different formulations.
- the relative profits obtained with the FAH, the PAH, the optimal steady-state policy and two benchmark policies often cited in the literature: the First-Come-First-Served (FCFS) policy and the protection level based (PLB) policy (see Section 6.3 for a detailed description of this policy).
- the tightness of the bounds on the profit derived for the FAH and the PAH.

In order to answer those questions, we first compute the heuristic acceptance policies and the optimal policy for the instances small enough to do so. We then simulate the different heuristic policies in order to determine their performance in terms of average gain per period. Indeed, although the gains $g_o^*(z)$ and $g_p^*(z)$ computed for the optimistic and pessimistic policies constitute upper and lower bounds on the optimal gain, the actual gains achieved by those policies – as well as the other policies – cannot be directly derived from the solution of the Markov Decision Process with an aggregate state description.

We simulate the long-run net profit for an acceptance policy by generating demand realizations. The simulation consists of an initial warm-up interval of 100,000 periods. Afterwards, the simulation incorporates an additional 100,000 periods for which the accumulated net profit is recorded. We repeat the process until the simulated net profit converges with a precision of 0.001%. We denote $\hat{g}_r$ and $\hat{g}_m$ the simulated long-run net profit achieved by the FCFS and the PLB policies, and $\hat{g}_o(z)$, $\hat{g}_p(z)$ and $\hat{g}_r(z)$ the simulated long-run net profit achieved by the PAH with a disaggregation level z assuming the optimistic, pessimistic and realistic scenarios, respectively. Note that the gain of the optimal steady-state policy is given directly by the solution of the Markov Decision Process, its value is equal to g^* .

Throughout the experiments, we will use linear programming (LP) to solve all MDP formulations (Ross, 1983; Bertsekas, 1995; Puterman, 2005). Among other tools, we have iteration based methods such as value iteration and policy iteration. They are efficient in finding near-optimal solutions in a short time frame.

However, when serving as benchmarks one needs to specify parameters such as number of iterations and optimality gap, which is a nontrivial task. Thus, in order to compare different MDP formulations in terms of their computational time on a fair basis, we choose LP instead.

The acceptance policy simulation routine was implemented in the Java language. The computer used was a 6-Core Intel Xeon 2×2.66 GHz with 48 GB of RAM. LP are solved using Gurobi 4.6.1 (<http://www.gurobi.com/>).

In the following subsections we introduce certain characteristics of the demand classes that we used to generate instances of the problem (Section 6.1). Afterwards, we describe the results of three studies: (1) the computation time for the construction of an acceptance policy for different instances (Section 6.2); (2) the profits achieved by the optimal steady-state policy, the FCFS and the PLB policies, and the aggregation heuristics (see Sections 6.3 and 6.4 for the FAH and the PAH, respectively); (3) the profits of the FAH and the PAH for large instances where the optimal steady-state policy cannot be computed (Section 6.5).

6.1. Experiment settings

The instances we generated are characterized by the following attributes:

Profit structure (ρ). It represents the ratio between the net profits of both demand classes, i.e. $\rho = r_2/r_1$. Without loss of generality, the value of r_1 is normalized to 1, so the value of r_2 is obtained directly from ρ . In our experiment we explore high, moderate and low differences in the profit structure, respectively (see Table 1).

Lead time structure L_1 and L_2 . We investigate two aspects of our model that depend on the values of L_1 and L_2 . On the one hand, it is clear from the formulation of our model that the computational complexity is closely related to the difference between L_1 and L_2 . On the other hand we will see that L_1 influences the performance of the different policies.

Order size structure B_1 and B_2 . The sizes of the orders determine the operational flexibility when an order is accepted. We also did some tests with stochastic order sizes, we do not report on this here as the results do not really differ from the deterministic case. The only difference is larger computation times, as a result the computation of the optimal policy is restricted to even smaller instances.

Demand structure (β). It is defined as the ratio between the expected demand rates of both classes, i.e. $\beta = (p_1 \cdot B_1)/(p_2 \cdot B_2)$. We study how this ratio impacts the performance of the proposed formulations. We chose the values of β such that the expected demand of one class is 100% or 50% greater than the other, or the expected demand rates of both classes are equal (see Table 1).

Table 1
Setting values of the parameters.

Common experimental conditions		Computational time experiment		Profit comparison experiment	
Attributes	Values	Attributes	Values	Attributes	Values
ρ	{0.25, 0.50, 0.75}	L_1, B_1	{1, 3, 5, 7}, 1	L_1, B_1	3, {1, 3}
β	{1/2, 2/3, 1, 3/2, 2/1}	L_2, B_2	{13, 15, 17, 19}, 1		7, {1, 3, 5, 7}
τ	{1.2, 1.6, 2.0}			L_2, B_2	{15, 29}, {3, 5, 7, 9}

Global demand rate (τ). It corresponds to the total expected demand for both classes per period i.e. $\tau = p_1 \cdot B_1 + p_2 \cdot B_2$. We focus on scenarios in which $\tau > 1$, inasmuch as in these scenarios the acceptance decision is most meaningful as some demand will have to be refused. If $\tau < 1$, the decision to not accept a demand is not very relevant.

Profit structure, order size structure and global demand rate are common experimental conditions for main numerical experiments we performed. For these parameters, we perform a full factorial experiment based on the values presented in the first part of Table 1. Lead times and order sizes are set depending on the objectives of different numerical experiments performed: analysis of the computational time and the profit comparison of the proposed heuristics with respect to the optimal policy and its bounds. Values of these parameters for each experiment are presented in the second and the third part of Table 1, respectively. Note that, p_1 and p_2 are functions of β, τ, B_1 and B_2 . Their values will be derived from these parameters.

6.2. Computation times

In this section, we compare the average CPU times needed to calculate the optimal steady-state policy and the heuristic policies. We analyze the FAH and the PAH for the realistic scenario, which is the most demanding in terms of computation time among the three scenarios considered. Note that, the PAH is tested for different values for the disaggregation level ($z \in \{2, 4, 6\}$). In addition to the common experimental conditions (see Table 1), in this experiment, we fix the order sizes to 1 because this leads to the longest computation times. Similar insights can be obtained for any other combination of values of B_1 and B_2 . As already explained, the difference between lead times and their combinations have a direct impact on the size of the state space of the system and, thus, the computation time. To explore this dependency we test a wide range of values of L_1 and L_2 shown in Table 1. In total this experiment consists of 342 instances (when discarding the combinations of values of parameters that result in $p_1, p_2 > 1$).

Table 2 shows the average CPU times and the size of the system state space for each combination of L_1 and L_2 . We observe that the time needed to find the optimal steady-state policy increases very quickly when the difference between L_2 and L_1 increases. In fact, for $L_2 - L_1 \geq 15$ we could not determine the optimal order acceptance policy in a reasonable amount of time.

6.3. Efficiency of the FAH

We compare the efficiency of the FAH with the optimal steady-state policy, the FCFS and the PLB policies by measuring their relative profits. We also compare the quality of the FAH under each of the three scenarios. In order to compare the efficiencies we compute the optimality gap of an acceptance policy constructed by the heuristic i as follows: $Gap_i = (g^* - \hat{g}_i) / g^* \times 100\%$. Note that the FAH is represented by $i = \{o(0), p(0), r(0)\}$ and the FCFS and PLB policy by $i = \{f, m\}$, respectively.

The protection level based (PLB) policy applied in this paper is an adaptation of the well-known revenue management approaches

that divide the available capacity into two portions: protected (reserved for the high net-profit class) and unprotected (used for both classes). These approaches are commonly applied for the finite-horizon problems. If they are directly implemented over an infinite horizon, the construction of the policy is computationally as intensive as the construction of the optimal steady-state policy. Therefore, in order to reduce the complexity of the existing approaches while at the same time capturing their essence we implement a myopic method which determines the amount of protected capacity q that maximizes the expected net profit within L_2 . This corresponds to the approximations used in the literature see e.g. Barut and Sridharan (2005). The expected net profit within L_2 is computed with the following expression: $r_1 \cdot E[\min(\bar{D}_1, \max(q, y[L_2] - \bar{D}_2))] + r_2 \cdot E[\min(\bar{D}_2, y[L_2] - q)]$, where the random variable $\bar{D}_i$ represents the amount of demand of class i during L_2 periods. The PLB policy consists in protecting the capacity q that maximizes the previous expression.

In order to evaluate the efficiency of the proposed policies with respect to the optimal steady-state policy, we fix $L_2 = 15$ and $L_1 \in \{3, 7\}$. Under this setting the optimal steady-state policy can be obtained within a tractable time. In order to capture the full essence of different degrees of heterogeneity in demand classes, we consider the different values of order sizes presented in the third part of Table 1. The values of other parameters belong to common experimental conditions. In this experiment, our analysis is based on 984 instances (the instances with $p_1, p_2 \geq 1$ are discarded).

Table 3 reports on the average optimality gaps for the two values of L_1 . The average optimality gap of the FAH is lower than that of the FCFS and the PLB policies. There is also a significant difference between the different implementations of the FAH (realistic, pessimistic and optimistic).

Table 2

Average CPU time and dimension of the state space of the system for the optimal steady-state policy, the FAH and the PAH under the realistic scenario.

L_2	L_1	Average CPU (s)					Size	
		Optimal	FAH	PAH			Optimal	PAH
				$z = 2$	$z = 4$	$z = 6$		
13	1	2.55	0.00	0.01	0.02	0.17	4096	768
13	3	2.99	0.00	0.01	0.06	0.38	2048	1024
13	5	0.26	0.00	0.01	0.06	–	768	768
13	7	0.03	0.00	0.01	0.03	–	256	256
15	1	194	0.00	0.01	0.03	0.30	16,384	1024
15	3	22	0.00	0.01	0.10	1.02	8192	1536
15	5	10	0.00	0.02	0.14	1.20	3072	1536
15	7	0.4	0.01	0.02	0.12	–	1024	1024
17	1	*	0.00	0.01	0.05	0.51	65,536	1280
17	3	*	0.00	0.02	0.16	2.11	32,768	2048
17	5	105	0.01	0.03	0.28	1.72	12,288	2304
17	7	8	0.01	0.03	0.29	3.03	4096	2048
19	1	*	0.00	0.01	0.06	0.73	262,144	1536
19	3	*	0.00	0.02	0.24	2.45	131,072	2560
19	5	*	0.01	0.04	0.44	3.24	49,152	3072
19	7	1015	0.01	0.05	0.54	3.52	16,384	3072

“*” Symbolizes instances for which the optimal steady-state policy cannot be obtained in 1 h and “–” represents the cases where the PAH is equivalent to the optimal policy because $z \geq L_2 - L_1 - 2$.

The lowest average optimality gap of the FAH is achieved with the realistic scenario. The strength of the realistic scenario lies in the balance between the excess of protection of capacity for high profitability orders (pessimistic scenario) and the assumption of maximum flexibility for processing incoming orders (optimistic scenario). The results also reveal that the PLB and FAH tend to perform a bit better when L_1 is larger.

We also study the reliability of the different policies when the values of the parameters vary. Fig. 5 shows the dispersion of the optimality gap of the analyzed methods. The limits of each box represent the first and third quartiles of the measured gaps. The central line corresponds to the mean result. Finally, the bottom and top whiskers represent the fifth and the ninety-fifth percentiles of the optimality gaps. In addition to the advantages achieved in terms of the average optimality gap, the implementation of the FAH with the realistic scenario provides the most reliable performance. Fig. 6 provides more details about the optimality gap of the analyzed instances when utilizing the FAH with the realistic scenario. In particular, the histogram presented in Fig. 6 shows that the first three categories for which $\text{Gap}_r < 1\%$ cover around 90% of the analyzed instances.

The relative advantage of the realistic scenario is born from its accuracy in representing the actual distribution of the reserved capacity in Interval II. Specifically, we expect the assumed distributional scenario leads to a good estimation of the transition of reserved capacity from Interval II to Interval I. With this in mind, we compare the value π for the different distributional scenarios with the observed frequency that such transition occurs. We obtain the observed frequency by simulating the transition between element $x[1]$ and $x[0]$ when implementing the optimal policy. Fig. 7 provides examples of the observed frequency for different values of B_2 and the assumed values of π when the policy is obtained by the FAH with the realistic scenario. When the value of order size of regular orders is small with respect to its lead time (i.e., $B_2 = 2$ in Fig. 7), the values of π under the realistic scenario are the closest to the observations. Nevertheless, when the value of B_2 increases (i.e., $B_2 = 4$) the pessimistic scenario seems to be a better assumption for low reserved capacity in Interval II (we remind the reader that under the pessimistic scenario $\pi = 1$ for $x_f[1] > 0$). However, such states are not likely to occur, as we see in Fig. 7 (probabilities of being in a state for the two values of B_2 are presented by dashed lines with markers). In fact, the realistic scenario estimates π better than the pessimistic scenario for the states where the reserved capacity is equal to 4, which are also the states most likely to occur. Moreover, the sum of the square of the errors of the estimation of the realistic and pessimistic scenario provides more evidence about the advantage of assuming a uniform distribution. In particular, for the case $B_2 = 4$, such sum is 0.019 and 0.013 for the pessimistic and realistic scenarios, respectively.

Our numerical results allow us to get a glimpse of the structure of the heuristic policy obtained through the FAH with the realistic scenario. This structure corresponds to a threshold in the values of $x_f[0]$ and $x_f[1]$. In other words, if we fix the value of $x_f[1]$ incoming regular orders are accepted up to a certain value of the amount of orders in the first interval. Analogously, if we fix the value of $x_f[0]$




Fig. 5. Dispersion of the optimality gap of different heuristic methods and distributional scenarios for the FAH.




Fig. 6. Dispersion of the optimality gap of the realistic scenarios for the FAH.




Fig. 7. Observed values and estimated values (when the FAH under the realistic scenario is implemented) of π . Results for $L_1 = 4$, $L_2 = 11$, $p_1 = 0.1$, $p_2 = 0.5$, $r_2 = 0.5$ and $B_1 = 1$.

incoming regular orders are accepted if the reserved capacity in the second interval is smaller than or equal to some threshold. Fig. 8 shows an example of this type of policy.

We analyse the sensitivity of the threshold value obtained through the FAH with the realistic scenario to changes on some parameters of the problem, as we show in Fig. 9(a)–(d). We can observe that the threshold values for the first and second intervals decrease with the value of B_2 . This is because accepting regular orders with larger order sizes will imply less available capacity for processing the future arrivals of urgent orders. Also, the threshold values have an inverse relation with the values of the probabilities p_1 and p_2 , since larger values of such probabilities lead to a reduction of the probability of having idle capacity in the incoming periods, therefore rejecting regular orders could be convenient in order to make some capacity available for urgent orders. The opposite is true for the threshold values and r_2 , given that the reduction

Table 3

Average optimality gap of the FCFS, the PLB policy and the FAH under the three scenarios for the two values of L_1 .

	FCFS (%)	PLB (%)	FAH		
			Optimistic (%)	Pessimistic (%)	Realistic (%)
$L_1 = 3$	12.92	7.27	5.00	1.90	0.45
$L_1 = 7$	13.52	5.13	4.98	0.57	0.36
Overall	13.33	5.81	4.99	0.99	0.39




Fig. 8. Acceptance/rejection policy of the FAH with realistic scenario for each aggregated state. $L_1 = 4$, $L_2 = 11$, $p_1 = 0.1$, $p_2 = 0.5$, $r_2 = 0.1$, $B_1 = 1$ and $B_2 = 2$.

of the difference between the revenues of the order classes, make regular orders more attractive.

6.4. Efficiency of the PAH

Despite the excellent overall performance of the FAH with the realistic scenario, the optimality gap remains significant for 9.76% of the instances generated ($Gap_r(0) \geq 1\%$), as we observe in Fig. 6. For these instances, we study how the systematic disaggregation of information related to already accepted orders can improve the quality of the acceptance policies found. For this, we calculate $Gap_r(z)$ for $z \in \{1, 2, \dots, 6\}$. The average values of $Gap_r(z)$ are displayed in Fig. 10. The optimality gap of the PAH with the realistic scenario decreases rapidly as the value of z grows. In particular, we note from Fig. 10 that when z increases from 0 to 4, the average gap decreases from 1.54% to 0.73% when $L_1 = 3$ and from 2.05% to 0.17% when $L_1 = 7$. Thus, the optimality gap of the FAH is greatly reduced while the computation time remains small (see Table 2). We remark that for the remaining 90.24% of the instances studied, the average optimality gap is also improved when the value of z increases.

6.5. Efficiency of the PAH and FAH for large lead times

As shown in Table 2 an advantage of the FAH and the PAH is to construct the policies in tractable time even for instances for which the optimal steady-state policies cannot be obtained. In order to





Fig. 10. Average optimality gap of the acceptance policies under the realistic scenario for the Instances with $Gap_r(0) \geq 1\%$ for different values of z and L_1 .

evaluate the efficiency of the studied heuristics for these instances, we fix the value of L_2 to be large i.e. $L_2 = 29$ while the other parameters take the same values as in Section 6.3. As a result our experiments include again 984 instances. We also compute the upper and lower bounds obtained from Proposition 1.


The average net profits of the different policies are plotted as a function of z in Fig. 11(a) and (b) for $L_1 = 3$ and $L_1 = 7$, respectively. We observe the following: first, the FAH and the PAH with the realistic scenario significantly outperform the PLB policy. Second, the effect of increasing z seems much stronger on the quality of the bounds than on the performance of the realistic policy. For a vast majority of instances the disaggregation does not bring any significant improvement. However, in a very similar fashion to Section 6.4, we observe that for the 10% of instances with the worst performance for the FAH, the profit increased by at least 0.5% between the PAH heuristic with $z = 0$ and $z = 6$. The average improvement for these instances is 0.781%. This seems to indicate that the main usefulness of the PAH is to give the possibility of controlling the quality of the proposed solution.

7. Operational flexibility


In the previous section, we illustrated the efficiency of the proposed algorithms for a large set of instances. Here, we try to gain some further insights into the circumstances where revenue




(a) Sensitivity to B_2 . $p_1 = 0.5$, $p_2 = 0.5$, $r_2 = 0.5$.



(b) Sensitivity to r_2 . $p_1 = 0.5$, $p_2 = 0.5$, $B_2 = 2$.



(c) Sensitivity to p_1 . $p_2 = 0.5$, $r_2 = 0.5$, $B_2 = 2$.



(d) Sensitivity to p_2 . $p_1 = 0.5$, $r_2 = 0.5$, $B_2 = 2$.

Fig. 9. Threshold values of the FAH with realistic scenario. $L_1 = 4$, $L_2 = 11$ and $B_1 = 1$.




Fig. 11. Effect of disaggregation of the information on the average long-run net profits achieved with different policies for the large lead time of regular orders.




Fig. 12. Benefits of the proposed revenue management approach (RM) compared to the FCFS policy for different characteristics of the demand streams. Results for $L_1 = 6$, $L_2 = 12$, $\rho = 0.50$ and $\beta = 0.50$. $\Delta_i = L_i - B_i$ where $i = 1, 2$. The Benefits of RM are computed by $(g^* - \hat{g}_f) / \hat{g}_f \times 100\%$.

management would have the most significant impact. We will focus on instances following the pattern of the cases that motivated our work (namely, a class of urgent orders with relatively low demand and a class of regular orders with longer lead times and lower revenues).

As expected, the benefit of revenue management is strongly correlated with the intensity of demand as shown in Fig. 12. We observe that the potential benefit of revenue management compared to the FCFS policy increases with the difference between L_2 and B_2 . This means that the advantage of revenue management is larger when there is more flexibility in processing regular orders. Note that, the effect of the flexibility for urgent orders is more limited.

In order to gain more insights into the impact of revenue management on the operations we compare the acceptance rates of both classes with and without revenue management. Table 4 shows the proportion of accepted orders for each demand class. The results illustrate how the revenue management technique is giving gradually higher priority to the urgent orders when ρ

(the relative margin of regular orders) is decreasing. We also observe again that the operational flexibility once an order is accepted plays an important role: in this experiment when $B_2 = 10$ the impact of revenue management is minimal, and more generally the smaller B_2 (for a fixed value of L_2) the larger the impact. Of particular interest is the case where $B_1 = 5$ and $B_2 = 10$, that is there is hardly any flexibility for both classes. In that case we observe that for $\rho = 0.75$ or 0.5 the optimal policy is essentially FCFS, while for $\rho = 0.25$ the policy is to accept only the high revenue class. In other words, given the small amount of operational flexibility there is an abrupt switch in the acceptance policy from accepting all orders (whenever feasible) to accepting only high margin urgent orders when the difference in margin is high enough.

8. Conclusions and future research

In this paper we studied the order acceptance problem for a firm serving two classes of demand that differ in net profit and lead time over an infinite horizon. We obtained the optimal order acceptance policy by formulating the problem as a multi-dimensional Markov Decision Process. However the resolution for the optimal policy can require high computational requirements. To overcome this difficulty, we proposed an efficient heuristic consisting in a parametric aggregation of the state space. The parameter makes it possible to find the best trade-off between the computation time and the quality of the solution. We propose several variants based on different assumptions about the dynamics of the aggregate state. The variant which assumes that all reserved capacity is distributed uniformly in the aggregation interval gives significantly better solutions than the other approaches studied in the extant literature. The other variants give lower and upper bounds that make it possible to obtain a guarantee about the quality of the solution and the possible gap with respect to the optimal solution. Finally, the computation time remains very short even for instances with large order lead times.

Table 4

Acceptance rate under the FCFS and the optimal policy and the benefits of RM for different order sizes and revenues.

B_1	B_2	FCFS		Optimal, $\rho = 0.75$			Optimal, $\rho = 0.50$			Optimal, $\rho = 0.25$		
		μ_1 (%)	μ_2 (%)	μ_1 (%)	μ_2 (%)	b (%)	μ_1 (%)	μ_2 (%)	b (%)	μ_1 (%)	μ_2 (%)	b (%)
1	2	55.6	91.8	96.6	70.8	4.9	99.3	68.9	14.1	99.9	67.7	31.8
1	6	77.1	70.7	85.1	65.8	0.3	93.4	59.8	3.6	97.6	55.2	11.4
1	10	85.0	52.4	85.1	52.3	0.0	85.1	52.3	0.0	90.2	45.6	1.6
3	2	44.8	95.4	80.4	75.2	2.8	87.3	70.3	12.5	92.4	62.3	33.6
3	6	57.7	74.2	57.8	74.2	0.0	80.5	57.4	4.4	88.5	45.6	17.5
3	10	60.4	56.5	60.4	56.5	0.0	60.4	56.5	0.1	97.1	0.0	9.6
5	2	36.1	97.1	66.0	79.3	1.8	66.0	79.3	9.1	71.9	70.1	26.4
5	6	40.4	77.1	40.4	77.1	0.0	53.4	66.1	1.8	71.0	48.4	20.6
5	10	44.4	56.8	44.4	56.8	0.0	44.4	56.8	0.0	79.6	0.0	9.3

Results for $L_1 = 6$, $L_2 = 12$, $\tau = 1.25$ and $\beta = 0.50$. Note that, μ_i is the service level of demand stream i and b is the benefit of RM compared to the FCFS.

We showed that the amount of operational flexibility (in our case this means how much slack there is between the promised lead time and the effective processing time needed) has a large impact on the performance of revenue management. The more flexibility there is the larger the potential benefit of implementing a revenue management based order acceptance policy. The study of the potential benefits of revenue management associated with flexibility in other contexts is an interesting avenue for further research. In the airline industry for example, the hub and spoke organization is widespread. The major airlines have several hubs and for a journey that is not starting or ending in a hub, it might be interesting to keep some flexibility in terms of the legs traveled by a passenger (i.e. through which hub) as long as a time-slot is respected for the departure and arrival times at the origin and destination respectively. It would be interesting to investigate how the state aggregation policy presented in this article could be extended for the more general network structure of airline operations.

Another direction for future work is the exploration of some properties of the value functions in the MDP formulation. Based on the concept of L^1 -convexity, some recent works (see e.g. Zipkin, 2008; Pang, Chen, & Feng, 2012) partially characterize the optimal policies for inventory problems with lead time issues. We conjecture that similar properties exist in the context of our problem. If so, it might be possible to develop more efficient computational methods.

Acknowledgements

This research was partially supported by the ARC project “Managing Shared Resources in Supply Chains” (Project No. ARC 08-13-008). The authors are thankful to the anonymous referees for providing useful suggestions that helped to improve this paper, and to Jean-Sébastien Tancrez for his insightful comments on an earlier version of the manuscript. Liang Lu also would like to thank the support of CORE Fellowship from Université catholique de Louvain during 2011–2013.

References

- Adelman, D. (2007). Dynamic bid-prices in revenue management. *Operations Research*, 55(4), 647–661.
- Aldhahebi, R. W., & Khalil, H. K. (1991). Aggregation of the policy iteration method for nearly completely decomposable Markov chains. *IEEE Transactions on Automatic Control*, 36(2), 178–187.
- Almudevar, A., & Arruda, E. F. (2012). Optimal approximation schedules for a class of iterative algorithms, with an application to multigrid value iteration. *IEEE Transactions on Automatic Control*, 57, 3132–3146.
- Arruda, E. F., & Frago, M. D. (2011). Time aggregated Markov decision processes via standard dynamic programming. *Operations Research Letters*, 39(3), 193–197.
- Balakrishnan, N., Sridharan, V., & Patterson, J. (1996). Rationing capacity between two product classes. *Decision Sciences*, 27(2), 185–214.
- Barut, M., & Sridharan, V. (2005). Revenue management in order-driven production systems. *Decision Sciences*, 36(2), 287–316.
- Bean, J. C., Birge, J. R., & Smith, R. L. (1987). Aggregation in dynamic programming. *Operations Research*, 35(2), 215–220.
- Belobaba, P., 1987. *Air travel demand and airline seat inventory management*. Ph.D. Thesis, Cambridge, Massachusetts: Massachusetts Institute of Technology.
- Bertsekas, D. P. (1995). *Dynamic programming and optimal control*. MA: Athena Scientific Belmont.
- Bertsimas, D., & Popescu, I. (2003). Revenue management in a dynamic network environment. *Transportation Science*, 37(3), 257–277.
- Bitran, G. R., & Mondschein, S. V. (1995). An application of yield management to the hotel industry considering multiple day stays. *Operations Research*, 43(3), 427–443.
- Brumelle, S. L., & McGill, J. I. (1993). Airline seats allocation with multiple nested fare classes. *Operations Research*, 41, 127–137.
- Cao, X. R., Ren, Z., Bhatnagar, S., Fu, M., & Marcus, S. (2002). A time aggregation approach to Markov decision processes. *Automatica*, 38(6), 929–943.
- Cheng, T. C. E., & Gupta, M. C. (1989). Survey of scheduling research involving due date determination decisions. *European Journal of Operational Research*, 38(2), 156–166.
- Cheng, K., & Zhang, K. J. (2012). Potentials based optimization with embedded Markov chain for stochastic constrained system. *Nonlinear Dynamics*, 67(2), 1067–1073.
- deB Harris, F. H., & Pinder, J. P. (1995). A revenue management approach to demand management and order booking in assemble-to-order manufacturing. *Journal of Operations Management*, 13(4), 299–309.
- Dobson, D., Hasija, S., & Pinker, E. J. (2011). Reserving capacity for urgent patients in primary care. *Production and Operations Management*, 20(3), 456–473.
- Duenyas, I. (1995). Single facility due date setting with multiple customer classes. *Management Science*, 41(4), 608–619.
- Duenyas, I., & Hopp, W. J. (1995). Quoting customer lead times. *Management Science*, 41(1), 43–57.
- Geraghty, M. K., & Johnson, E. (1997). Revenue management saves national car rental. *Interfaces*, 27(1), 107–127.
- Gems, R., & Van Foreest, N. D. (2011). Admission policies for the customized stochastic lot scheduling problem with strict due-dates. *European Journal of Operational Research*, 213(2), 375–383.
- Gupta, D., & Wang, L. (2007). Capacity management for contract manufacturing. *Operations Research*, 55(2), 367–377.
- Gupta, D., & Wang, L. (2008). Revenue management for a primary-care clinic in the presence of patient choice. *Operations Research*, 56(3), 576–592.
- Herbots, J., Herroelen, W., & Leus, R. (2007). Dynamic order acceptance and capacity planning on a single bottleneck resource. *Naval Research Logistics*, 54(8), 874–889.
- Herbots, J., Herroelen, W., & Leus, R. (2010). Single-pass and approximate dynamic-programming algorithms for order acceptance and capacity planning. *Journal of Heuristics*, 16, 189–209.
- Hopp, W. J., & Sturgis, M. L. R. (2000). Quoting manufacturing due dates subject to a service level constraint. *IEE Transactions*, 32, 771–784.
- Hu, G., & Wu, C. (2000). Relative value iteration algorithm with soft state aggregation. *Control Theory and Applications*, 17(3), 415–418.
- Jia, Q. S. (2011). On state aggregation to approximate complex value functions in large-scale Markov decision processes. *IEEE Transactions on Automatic Control*, 56(2), 333–344.
- Kaminsky, P., & Hochbaum, D. (2004). Due date quotation models and algorithms. In J. Y. T. Leung (Ed.), *Handbook of scheduling: Algorithms, models, and performance analysis* (pp. 20–1–20–22). Boca Raton, FL: CRC Press.
- Kapuscinski, R., & Tayur, S. (2007). Reliable due-date setting in a capacitated MTO system with two customer classes. *Operations Research*, 55(1), 56–74.
- Keskinocak, P., & Tayur, S. (2004). Due-date management policies. In D. Simchi-Levi, S. Wu, & Z. J. Shen (Eds.), *Handbook of quantitative supply chain analysis: Modeling in the e-business era. International series in operations research and management science* (pp. 485–553). Norwell, MA: Kluwer Academic Publishers.
- Kimes, S. E. (1989). Yield management: A tool for capacity-constrained service firms. *Journal of Operations Management*, 8(4), 348–363.
- Kunnumkal, S., & Topaloglu, H. (2008). A refined deterministic linear program for the network revenue management problem with customer choice behavior. *Naval Research Logistics*, 55(6), 563–580.
- Leizarowitz, A., & Schwartz, A. (2008). Exact finite approximations of average-cost countable Markov decision processes. *Automatica*, 44(6), 1480–1487.
- Littlewood, K. (1972). Forecasting and control of passengers bookings. In *AGIFORS 12th annual symposium proceedings* (pp. 95–128).
- Mendelsohn, R. (1982). An iterative aggregation procedure for Markov decision processes. *Operations Research*, 30(1), 62–73.
- Pang, Z., Chen, F. Y., & Feng, Y. (2012). A note on the structure of joint inventory-pricing control with lead times. *Operations Research*, 60(3), 581–587.
- Puterman, M. L. (2005). *Markov decision processes: Discrete stochastic dynamic programming*. New York: John Wiley & Sons.
- Ross, M. S. (1983). *Introduction to stochastic dynamic programming*. New York: Academic Press Inc.
- Spengler, T., Rehkopf, S., & Thomas, V. (2007). Revenue management in make-to-order manufacturing an application to the iron and steel industry. *OR Spectrum*, 29, 157–171.
- Sun, T., Zao, Q., & Lux, P. B. (2007). Incremental value iteration for time aggregated Markov decision processes. *IEEE Transactions on Automatic Control*, 52(11), 2117–2182.
- Talluri, K., & van Ryzin, G. (1998). An analysis of bid-price controls for network revenue management. *Management Science*, 44(11), 1577–1593.
- Talluri, K., & van Ryzin, G. (2004). *The theory and practice of revenue management*. Boston, Dordrecht, London: Kluwer Academic Publishers.
- Van Roy, B. (2006). Performance loss bounds for approximate value iteration with state aggregation. *Mathematics of Operations Research*, 31(2), 234–244.
- Wollmer, R. D. (1992). An airline seat management model for a single leg route when lower fare classes book first. *Operations Research*, 40, 26–37.
- Xu, S. H., Gao, L., & Ou, J. (2007). Service performance analysis and improvement for a ticket queue with balking customers. *Management Science*, 53(6), 971–990.
- Zhang, D. (2011). An improved dynamic programming decomposition approach for network revenue management. *Manufacturing and Service Operations Management*, 13(1), 35–52.
- Zhang, D., & Adelman, D. (2009). An approximate dynamic programming approach to network revenue management with customer choice. *Transportation Science*, 43(3), 381–394.
- Zhang, C., & Baras, J. S. (2001). *A new adaptive aggregation algorithm for infinite horizon dynamic programming*. Tech. Rep. TR 2001-5, CSHCN.
- Zipkin, P. (2008). On the structure of lost-sales inventory models. *Operations Research*, 56(4), 937–944.

CORE Reprints -- Recent titles

2582. Augustin LEFEVRE, François GLINEUR and Pierre-Antoine ABSIL. A convex formulation for informed source separation in the single channel setting. *Neurocomputing*, 141, 26-36, 2014.
2583. Dunia LÓPEZ-PINTADO and Juan D. MORENO-TERNERO. On discrimination in the optimal management of teams. *Journal of Mathematical Economics*, 51, 154-162, 2014.
2584. Agostinho AGRA, Henrik ANDERSSON, Marielle CHRISTIANSEN and Laurence WOLSEY. A maritime inventory routing problem: discrete time formulations and valid inequalities. *Networks*, 62, 297-313, 2013.
2585. Pierre GIOT, Ulrich HEGE and Armin SCHWIENBACHER. Are novice private equity funds risk-takers? Evidence from a comparison with established funds. *Journal of Corporate Finance*, 27, 55-71, 2014.
2586. David de la CROIX and Clara DELAVALLADE. Why corrupt governments may receive more foreign aid. *Oxford Economic Papers*, 66(1), 61-66, 2014.
2587. Yurii NESTEROV and Vladimir Y. PROTASOV. Optimizing the spectral radius. *SIAM Journal on Matrix Analysis and Applications*, 34(3), 999-1013, 2013.
2588. David CRAINICH, Louis ECKHOUDT and Olivier LE COURTOIS. Decreasing downside risk aversion and background risk. *Journal of Mathematical Economics*, 53, 59-63, 2014.
2589. Marc FLEURBAEY, Marie-Louise LEROUX and Gregory PONTIERE. Compensating the dead. *Journal of Mathematical Economics*, 51, 28-41, 2014.
2590. Jacques H. DREZE and Alain DURRE. Fiscal integration and growth stimulation in Europe (with an appendix by Jacques Drèze and Jean-François Carpentier). *Recherches Economiques de Louvain*, 80, 5-45, 2014.
2591. Kris BOUDT and Mikael PETITJEAN. Intraday liquidity dynamics and news releases around price jumps: evidence from the DJIA stocks. *Journal of Financial Markets*, 17, 121-149, 2014.
2592. Florian MAYNERIS and Loriane PY. The efficiency of enterprise zone programs. Some conflicting results? *Région et Développement*, 38, 209-224, 2013.
2593. Yurii NESTEROV. Subgradient methods for huge-scale optimization problems. *Mathematical Programming Series A*, 146(1-2), 275-297, 2014.
2594. Olivier DEVOLDER, François GLINEUR and Yurii NESTEROV. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming Series A*, 146(1-2), 37-75, 2014.
2595. Mathieu LEFEVRE, Pierre PESTIEAU and Gregory PONTIERE. Measuring poverty without the mortality paradox. *Social Choice and Welfare*, 40(1), 285-316, 2013.
2596. Pierre PESTIEAU and Gregory PONTIERE. Optimal fertility along the life cycle. *Economic Theory*, 55(1), 185-224, 2014.
2597. Claude d'ASPREMONT and Rodolphe DOS SANTOS FERREIRA. Household behavior and individual autonomy: an extended Lindahl mechanism. *Economic Theory*, 55(3), 643-664, 2014.
2598. Johanna M.M. GOERTZ. Inefficient committees: small elections with three alternatives. *Social Choice and Welfare*, 43(2), 357-375, 2014.
2599. Nguyen Thang DAO and Julio DAVILA. Implementing steady state efficiency in overlapping generations economies with environmental externalities. *Journal of Public Economic Theory*, 16(4), 620-649, 2014.
2600. Kristian BEHRENS, Susana PERALTA and Pierre M. PICARD. Transfer pricing rules, OECD guidelines, and market distortions. *Journal of Public Economic Theory*, 16(4), 650-680, 2014.
2601. David A. WUTTKE, Constantin BLOME and Michael HENKE. Focusing the financial flow of supply chains: an empirical investigation of financial supply chain management. *International Journal of Production Economics*, 145(3), 773-789, 2013.
2602. Constantin BLOME, Tobias SCHOENHERR and Matthias KAESSER. Ambidextrous governance in supply chains: the impact on innovation and cost performance. *Journal of Supply Chain Management*, 49(4), 59-80, 2013.
2603. Constantin BLOME, Tobias SCHOENHERR and Dominik ECKSTEIN. The impact of knowledge transfer and complexity on supply chain flexibility: knowledge-based view. *International Journal of Production Economics*, 147(2), 307-316, 2014.
2604. Francesco DI COMITE, Jacques-François THISSE and Hylke VANDENBUSSCHE. Verti-zontal differentiation in export markets. *Journal of International Economics*, 93(1), 50-66, 2014.
2605. Koen DECANQ, Luc VAN OOTEGEM and Elsy VERHOFSTADT. What if we voted on the weights of a multidimensional well-being index? An illustration with Flemish data. *Fiscal Studies*, 34(3), 315-332, 2013.

2606. Koen DECANQ. Copula-based measurement of dependence between dimensions of well-being. *Oxford Economic Papers*, 66(3), 681-701, 2014.
2607. Kristof BOSMANS, Koen DECANQ and André DECOSTER. The relativity of decreasing inequality between countries. *Economica*, 81(322), 276-292, 2014.
2608. Olivier BARGAIN, Mathias DOLLS, Clemens FUEST, Dirk NEUMANN, Andreas PEICHL, Nico PESTEL and Sebastian SIEGLOCH. Fiscal union in Europe? Redistributive and stabilizing effects of a European tax-benefit system and fiscal equalization mechanism. *Economic Policy*, 28(75), 375-422, 2013.
2609. Dominik DORSCH, Hubertus Th. JONGEN, Jan-J. RUCKMANN and Vladimir SHIKHMAN. On the local representation of piecewise smooth equations as a Lipschitz manifold. *Journal of Mathematical Analysis and Applications*, 411(2), 916-930, 2014.
2610. Dirk VAN DE GAER, Joost VANDENBOSSCHE and José Luis FIGUEROA. Children's health opportunities and project evaluation: Mexico's oportunidades program. *The World Bank Economic Review*, 28(2), 282-310, 2014.
2611. Mathieu VAN VYVE, Laurence A. WOLSEY and Hande YAMAN. Relaxations for two-level multi-item lot-sizing problems. *Mathematical Programming Series A*, 146(1-2), 495-523, 2014.
2612. Monalisa SEN and Anil K. BERA. The improbable nature of the implied correlation matrix from spatial regression models. *Regional Statistics*, 14(1), 3-15, 2014.
2613. Johanna M.M. GOERTZ and François MANIQUET. Condorcet jury theorem: an example in which informative voting is rational but leads to inefficient information aggregation. *Economics Letters*, 125(1), 25-28, 2014.
2614. Thierry BRECHET, Carmen CAMACHO and Vladimir M. VELIOV. Model predictive control, the economy, and the issue of global warming. *Annals of Operations Research*, 220(1), 25-48, 2014.
2615. Giuseppe DE FEO and Jean HINDRIKS. Harmful competition in insurance markets. *Journal of Economic Behavior & Organization*, 106, 213-226, 2014.
2616. Jean HINDRIKS, Susana PERALTA and Shlomo WEBER. Local taxation of global corporation: a simple solution. *Annals of Economics and Statistics (Special Issue on the Economics of Taxation)*, 113-114, 37-65, 2014.
2617. Paul BELLEFLAMME, Thomas LAMBERT and Armin SCHWIENBACHER. Individual crowdfunding practices. *Venture Capital: An International Journal of Entrepreneurial Finance*, 15(4), 313-333, 2013.
2618. Feng QIU, Shabbir AHMED, Santanu S. DEY and Laurence A. WOLSEY. Covering linear programming with violations. *INFORMS Journal on Computing*, 26(3), 531-546, 2014.
2619. Gautier M. KRINGS, Jean-François CARPANTIER and Jean-Charles DELVENNE. Trade integration and trade imbalances in the European Union: a network perspective. *PLoS One*, 9(1), e83448, 2014.
2620. Jean-Charles DELVENNE and Henrik SANDBERG. Finite-time thermodynamics of port-Hamiltonian systems. *Physica D*, 267, 123-132, 2014.
2621. Volker BRITZ, P. Jean-Jacques HERINGS and Arkadi PREDTETCHINSKI. A bargaining theory of the firm. *Economic Theory*, 54(1), 45-75, 2013.
2622. Volker BRITZ, P. Jean-Jacques HERINGS and Arkadi PREDTETCHINSKI. On the convergence to the Nash bargaining solution for action-dependent bargaining protocols. *Games and Economic Behavior*, 86, 178-183, 2014.
2623. Mathieu VAN VYVE. Fixed-charge transportation on a path: optimization, LP formulations and separation. *Mathematical Programming Series A*, 142(1-2), 371-395, 2013.
2624. Philippe MICHEL and Pierre PESTIEAU. Social security and early retirement in an overlapping-generations growth model. *Annals of Economics and Finance*, 14-2(B) 705-719, 2013.
2625. Jacques-François THISSE. The new science of cities by Michael Batty: the opinion of an economist. *Journal of Economic Literature*, 52(3), 805-819, 2014.
2626. Bas de GEUS, Bart DEGRAEUWE, Grégory VANDENBULCKE, Luc INT PANIS, Isabelle THOMAS, Joris AERTSENS, Yves DE WEERDT, Rudi TORFS and Romain MEEUSEN. Utilitarian cycling in Belgium: a cross-sectional study in a sample of regular cyclists. *Journal of Physical Activity and Health*, 11, 884-894, 2014.
2627. Helmuth CREMER and Pierre PESTIEAU. Social long-term care insurance and redistribution. *International Tax and Public Finance*, 21(6), 955-974, 2014.
2628. Philippe CHEVALIER, Alejandro LAMAS, Liang LU and Tanja MLINAR. Revenue management for operations with urgent orders. *European Journal of Operational Research*, 240(2), 476-487, 2015.