

Order acceptance in food processing systems with random raw material requirements

Onur A. Kilic · Dirk Pieter van Donk ·
Jacob Wijngaard · S. Armagan Tarim

Published online: 11 June 2010

© The Author(s) 2010. This article is published with open access at Springerlink.com

Abstract This study considers a food production system that processes a single perishable raw material into several products having stochastic demands. In order to process an order, the amount of raw material delivery from storage needs to meet the raw material requirement of the order. However, the amount of raw material required to process an order is not exactly known beforehand as it becomes evident during processing. The problem is to determine the admission decisions for incoming orders so as to maximize the expected total revenue. It is demonstrated that the problem can be modeled as a single resource capacity control problem. The optimal policy is shown to be too complex for practical use. A heuristic approach is proposed which follows rather simple decision rules while providing good results. By means of a numerical study, the cases where it is critical to employ optimal policies are highlighted, the effectiveness of the heuristic approach is investigated, and the effects of the random resource requirements of orders are analyzed.

S. Armagan Tarim was supported by Hacettepe University (BAB) and the Scientific and Technological Research Council of Turkey (TUBITAK) under Grant No. SOBAG-108K027.

O. A. Kilic (✉) · D. P. van Donk · J. Wijngaard
Department of Operations, University of Groningen, P.O. Box 800,
9700 AV Groningen, The Netherlands
e-mail: o.a.kilic@rug.nl

D. P. van Donk
e-mail: d.p.van.donk@rug.nl

J. Wijngaard
e-mail: j.wijngaard@rug.nl

S. A. Tarim
Department of Management, Hacettepe University, Beytepe, 06800 Ankara, Turkey
e-mail: armagan.tarim@hacettepe.edu.tr

Keywords Processing system · Order acceptance · Raw material

1 Introduction

The food processing industry is characterized by divergent product structures where a small number of (agricultural) raw materials are used to produce a large variety of customer-specific end products (e.g. [Akkerman and van Donk 2009](#)). Due to the large variety of end products, it is often not possible or at least inefficient to produce all end products in make-to-stock fashion. Hence, make-to-order (MTO) is the typical production setting in the food processing industry.

Food processing industry involves highly perishable raw materials which are usually replenished periodically during their relatively short harvest seasons. The raw material procurement costs are relatively high compared to operational costs. Therefore, firms often face the issue of covering demands with limited amounts of raw materials being significantly of less value at the end of the season.

Another important characteristic of the food processing industry is the variability in production yield. This issue basically derives from two sources. First, the food processing industry involves raw materials whose qualities are often variable ([Fransoo and Rutten 1994](#)). Most quality parameters, such as protein, fat, and sugar content are usually hard to measure reliably. Some others, such as texture, smell, and taste can only be measured in a subjective way. Hence, it is hard to know the exact amount of raw material that is needed to process a given amount of end product ([Somsen and Capelle 2002](#)). Second, the production process itself involves variability ([Fransoo and Rutten 1994](#); [Flapper et al. 2002](#)). The yield variability due to the production process is often associated with the type of the operation and production quantities involved ([Murthy and Ma 1996](#)). The yield of a production run is affected by the inconsistencies in processing operations (involving chemical reactions), disturbances (i.e. starting up, changeovers, finishing), and packaging operations. Henceforth, either a part of the batch or all of it may fail to fulfill certain quality specifications and may need to be disposed of as waste or by-product. In such cases, additional production runs, and hence, additional raw materials are required. This issue is particularly important in MTO environments where demands are rigid and shortages are not allowed ([Grosfeld-Nir and Gerchak 2004](#)).

In the production environments discussed above, an important planning problem is how to allocate available raw material to incoming orders over time according to their relative importance in order to achieve better operational performance ([Fransoo and Rutten 1994](#); [Van Donk 2000](#)). This study is motivated by this practical and pervasive issue encountered in food processes.

A typical real-life example of the aforementioned problem can be found in the potato starch industry. The basic and most obvious process is the conversion of potatoes into starch during the harvest season. Starch is used in many different applications such as food, textile, paper, adhesives, and detergents. Given the size of the industry and technologies involved, products are made and marketed in different business units. Since the main aim of a potato starch company is to sell all starch during the year in order to get the highest value, it is common practice to allocate a certain amount of raw material to each business unit with the explicit demand to transform it into

marketable products at the highest possible price. Some allocations depend upon the specific characteristics of the starch, for example, starch from modified potatoes. It should be obvious that there is a considerable drawback in still having inventory at the end of the year, if new potato starch is then available. Some products made out of potato starch are technologically advanced and highly customized, and they show a rather erratic almost lumpy demand pattern. In these cases product orders require specific types of processing operations usually involving chemical reactions, which lead to order-specific revenues and variable raw material requirements, often because the production process is insufficiently under control. Throughout the year, orders for various products arrive, and they are subsequently either accepted or rejected, as a business unit aims at maximizing its returns. The admission decisions are given according to revenues and raw material requirements associated with incoming orders. The revenue gained by accepting a particular order is known at the time of making the acceptance decision. However, the raw material requirement of an order is not known with certainty due to the variability in yield. As already noted, given the nature of the company and its policies, and due to the relatively perishable nature of the raw material, the remaining inventory (if there is any) has to be disposed of at a low value, for example, as waste or a by-product. The key problem here is to establish decision rules that coordinate the admission decisions for incoming orders. Similar types of decision situations can be found in milk processing (where a certain amount has to be processed into products) and other food industries.

The problem we consider in this study falls into the category of capacity control problems in revenue management which have attracted great interest from both practitioners and researchers. Revenue management is used in situations where a finite amount of products/services have to be allocated to several classes of customer. The reader is referred to [Talluri and Van Ryzin \(2004\)](#) for an overview of this field. Revenue management literature offers a large variety of studies concentrating on establishing optimal policies for capacity control problems, especially for airline management practices. [Lee and Hersh \(1993\)](#) is one of the first studies to characterize the structure of the optimal policy for the basic capacity control problem. They showed that the optimal policy is threshold-type relying on the remaining time and the remaining resource level. In other words, for a given resource level, a given order is accepted only if the remaining time is less than an order-specific time threshold; and similarly, for a given remaining time, a given order is accepted only if the remaining inventory is larger than an order-specific resource level threshold. However, these easily implementable threshold-type policies are only optimal when the resource requirements of orders are unit-sized as is the case for airline seat allocation problems (see e.g. [Lee and Hersh 1993](#); [Papastavrou et al. 1996](#); [Van Slyke and Young 2000](#); [Kleywegt and Papastavrou 2001](#); [Brumelle and Walczak 2003](#)). From a modeling point of view, order-specific resource requirements do not pose much difficulty. However, they have a profound impact on the structure of the optimal policy since the optimal expected revenue function no longer preserves some of the basic monotonicity properties ([Talluri and Van Ryzin 2004](#)). In the case of non-unit resource requirements, the behavior of the optimal admission decision is rather complex, since the optimal policy is not threshold-type. As a result, practical use and implementation is limited, since a very careful and precise examination of the resource level throughout time is required.

The aforementioned literature provides a strong background for the problem we address. However, there are some specific characteristics of the food processing industry, such as the random resource requirements of orders and disposal costs, which have not yet been addressed. In this study, we stylize and streamline the raw material allocation problem in the food processing industry by addressing the aforementioned characteristics. We build on the well-established revenue management models. We do not aim to characterize the optimal policy since it is known that it does not possess a simple structure. Rather, we are rather interested in (i) pointing out the cases where it is critical and necessary to employ optimal/near-optimal admission policies; (ii) developing a heuristic approach possessing a rather simple structure while providing satisfactory performance; and (iii) analyzing the effects of the stochasticity of resource requirements of orders.

The remainder of this paper is organized as follows: In Sect. 2, we provide the formal problem definition. In Sect. 3, we present a DP to compute the optimal policy. In Sect. 4, we discuss the structural properties of the optimal policy. In Sect. 5, we propose two simple and easily implementable heuristics for the problem. In Sect. 6, we conduct a numerical study and investigate the effects of different problem settings on the performance of the optimal policy and heuristics. Finally, in Sect. 7, we draw our conclusions and propose some extensions of the study.

2 Problem definition

Consider a food processing system where a key perishable resource (raw material) is used to process a set of order types indexed by $i = 1, \dots, m$. The planning horizon is composed of t discrete time periods indexed by $n = 0, \dots, t - 1$. The resource inventory at the beginning of period 0 involves s units of material. The remaining resource at the end of period $t - 1$ (or the fictitious period t) is disposed of as waste or by-product with a unit disposal cost of c . Customer orders arrive throughout the planning horizon. In each time period at most a single order may arrive. In time period n there is a probability p_{in} of a type- i order arrival and $p_{0n} = 1 - \sum_{i=1}^m p_{in} \geq 0$ of no order arrival. Orders are either accepted or rejected as a whole (i.e. complete admission). Upon the arrival of an order, its type and associated revenue become known. The revenue gained by accepting a type- i order is denoted as r_i . However, the resource requirement of a type- i order, denoted as w_i , is random as it emerges during processing the order. Hence, the decision maker accepts or rejects an incoming order without knowing the exact resource requirement. The resource requirement of a type- i order follows a known probability mass τ_i which depends on the process technology and product recipe used in processing type- i orders. When the resource inventory is insufficient to fulfill an accepted order, the shortage can be covered from an external source at a unit shortage (penalty) cost of z . We assume that z is large enough that it is not profitable to accept an order when there is no resource on hand. Otherwise, it would be optimal to accept all incoming orders. The resource inventory is reviewed throughout the planning horizon, and the decision maker knows the resource level when an order arrives.

The admission decisions depend on (i) arrival processes, profitabilities, and resource requirements of incoming orders; (ii) the current resource level and the time remaining until the end of the planning horizon; and (iii) associated cost parameters. The basic intuition is that, when the resource level is low and the time remaining until the end of the planning horizon is long, it would be reasonable to reject less profitable orders in order to preserve resources so as to be able to accept more profitable future orders. Also, for high levels of raw material and a short planning horizon, accepting every order seems reasonable. We analyze the admission policies, characterized by decision rules for given resource levels, and time periods for accepting/rejecting orders, maximizing expected revenue accumulated throughout the planning horizon.

3 Dynamic program

The problem we define in Sect. 2 can be modeled as a dynamic program. Let $g_n(\cdot)$ be the optimal revenue function at period n , that is, $g_n(x)$ represents the expected revenue if the initial resource level is x units and the optimal admission decisions are made throughout the rest of the planning horizon. Then, we can write

$$g_n(x) = \sum_{i=1}^m p_{in} v_{in}(x) + p_{0n} g_{n+1}(x) \quad (1)$$

where

$$v_{in}(x) = \begin{cases} r_i + E[g_{n+1}(x - w_i)] & \text{if } r_i \geq g_{n+1}(x) - E[g_{n+1}(x - w_i)] \\ g_{n+1}(x) & \text{otherwise} \end{cases} \quad (2)$$

with the terminal revenues

$$g_t(x) = -c(x)^+ - z(x)^- \quad (3)$$

incorporating both disposal and penalty costs where $(x)^+ := \max\{0, x\}$ and $(x)^- := \max\{0, -x\}$. As stated in the problem definition, the penalty cost of shortage is independent of time and z is large enough to prevent any acceptance decision when there is no resource available. These enable us to express the penalty cost in the terminal revenue function and set $g_n(x) = -zx$ for all $x \leq 0$.

It is clear that an optimal policy for the dynamic program accepts a type- i order when the resource level is x in period n only if

$$r_i \geq g_{n+1}(x) - E[g_{n+1}(x - w_i)]. \quad (4)$$

The left-hand side of (4) represents the immediate incremental revenue, whereas the right-hand side is the expected loss in future revenue by accepting a type- i order.

In order to solve the DP, one needs to compute $g_n(x)$ for $x = 1, \dots, s$ and $n = 0, \dots, t - 1$ by backward recursion. Note that there is no need to evaluate $g_n(x)$ for $x \leq 0$ explicitly since they are all equal to $-zx$.

4 Structural properties

In this section, we explore some monotonicity properties of the DP given in Sect. 3 which may result in acceptance policies that are easy to implement. We begin our discussion by considering a simple version of the problem. Let us assume that each order is characterized by the same constant unit-sized resource requirement, that is, $w_i = 1$ for all $i = 1, \dots, m$. Furthermore, let us assume that the disposal cost is zero and there is no shortage option. In this case, the problem is reduced to the basic seat allocation problem in airline revenue management with multiple demand classes. It can be shown that the optimal revenue function possesses some important structural properties (see e.g. Lee and Hersh 1993). For this problem, Eq. (4) can be re-written as

$$r_i \geq g_{n+1}(x) - g_{n+1}(x - 1). \quad (5)$$

Notice that now the expected loss in future revenue by accepting an order is independent of the order type. Based on this observation, Lee and Hersh (1993) state the main properties of $g_n(x)$ by means of the following theorem:

Theorem 1 (Lee and Hersh 1993) *When all orders have unit-sized resource requirements,*

1. $g_n(x) - g_n(x - 1)$ is non-increasing in x for any given n ,
2. $g_n(x) - g_n(x - 1)$ is non-increasing in n for any given x .

Theorem 1 shows that the expected loss in future revenue by accepting an order (or the marginal value of an additional unit of resource) is higher when the resource level is relatively low and/or the remaining time until the end of the planning horizon is relatively long. The monotonicity of $g_n(x)$ leads to the following implications:

1. For each order type i and any given period n there exists a critical resource level x^* satisfying $r_i \geq g_{n+1}(x) - g_{n+1}(x - 1)$ for all $x \geq x^*$ such that a type- i order is rejected whenever $x < x^*$ and accepted otherwise.
2. For each order type i and any given resource level x there exists a critical time period n^* satisfying $r_i \geq g_{n+1}(x) - g_{n+1}(x - 1)$ for all $n \geq n^*$ such that a type- i order is rejected whenever $n < n^*$ and accepted otherwise.

Let us illustrate these results by means of a simple numerical example.

Example 1 Consider a five-period problem with two order types (Type-1 and Type-2) both having unit-sized resource requirements. The arrival probabilities of order types are stationary over time and they are both equal to 0.5, that is, $p_{1n} = p_{2n} = 0.5$ for all $n = 0, \dots, 4$. The respective rewards of orders are $r_1 = 1$ and $r_2 = 2$. Since there are only two order types to be considered, each with unit-sized resource requirements, it is clear that the more profitable order type, that is, Type-2, would be accepted whenever there are sufficient resources available (i.e. $x \geq 1$). However, Type-1 orders can be rejected in order to allocate the available resources for Type-2 order arrivals in the later periods.

Table 1 The optimal expected revenues and admission decisions for Example 1

x/n	0	1	2	3	4
5	7.500 (1, 1)	6.000 (1, 1)	4.500 (1, 1)	3.000 (1, 1)	1.500 (1, 1)
4	6.469 (0, 1)	6.000 (1, 1)	4.500 (1, 1)	3.000 (1, 1)	1.500 (1, 1)
3	5.281 (0, 1)	4.938 (0, 1)	4.500 (1, 1)	3.000 (1, 1)	1.500 (1, 1)
2	3.781 (0, 1)	3.625 (0, 1)	3.375 (0, 1)	3.000 (1, 1)	1.500 (1, 1)
1	1.969 (0, 1)	1.938 (0, 1)	1.875 (0, 1)	1.750 (0, 1)	1.500 (1, 1)
0	0.000 (0, 0)	0.000 (0, 0)	0.000 (0, 0)	0.000 (0, 0)	0.000 (0, 0)

Rows and columns stand for the respective resource level x and the decision period n and each entry presents the expected revenue and the admission decisions (1 = acceptance and 0 = rejection) for Type-1 and Type-2 orders, respectively

We evaluate initial resource levels $[0, 5]$. Since there are only five periods in each of which at most one order can arrive, the initial resource level of five units is the maximum amount of resources that could possibly be used. Table 1 presents the optimal expected rewards and optimal admission decisions corresponding to each decision period and resource level.

Let us consider Type-1 orders which are less preferable as compared to Type-2 orders. The critical resource levels of Type-1 orders are 5, 4, 3, 2, 1 in periods 0, 1, 2, 3, 4, respectively. For the given periods, Type-1 orders are accepted only when the resource level is higher than the critical level. The critical time periods for Type-1 orders are 4, 3, 2, 1, 0 for resource levels 1, 2, 3, 4, 5, respectively. For the given resource levels, Type-1 orders are accepted only when the time period is later than the critical time period. The relationship between the acceptance decisions and the optimal average rewards can also be observed in Table 1. For instance, consider the last period, that is, Period 4. Here, any incoming order would be accepted as long as $x \geq 1$. Hence, the optimal expected reward equals $0.5 \times 1 + 0.5 \times 2 = 1.5$ for all $x \geq 1$ and 0 for $x = 0$. Then, we can say that the marginal value of an additional resource at $n = 4$ is 0 for all $x \geq 1$ and 1.5 for $x = 0$. Consequently, since the revenue gained by accepting a Type-1 order is equal to 1, a Type-1 order at period 3 is accepted for all $x \geq 2$ and rejected for all $x < 2$.

The above example illustrates the threshold-type behavior of the optimal policy. However, the problem addressed in this paper possesses a number of additional features which may affect this behavior. In the following, we investigate whether the simple threshold policies apply to the problem we address.

First, we consider the non-zero disposal costs and shortage penalty costs. Introducing non-zero disposal costs and shortage penalty costs leads to a terminal reward $g_t(x) = -c(x)^+ - z(x)^-$ which is linearly decreasing on $x \geq 0$ with rate c and linearly increasing on $x \leq 0$ with rate z . Henceforth, in this case, $g_t(x)$ is a concave function of x which preserves the property that $g_t(x) - g_t(x - 1)$ is non-increasing in x . This shows that non-zero disposal costs and shortage penalty costs do not interfere with the critical resource levels and decision periods. Hence, the optimal policies are still threshold-type.

Table 2 The optimal expected revenues and admission decisions for Example 2

x/n	0	1	2	3	4
10	12.500 (1, 1)	10.000 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
9	12.375 (1, 1)	10.000 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
8	11.906 (1, 1)	10.000 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
7	10.938 (1, 1)	9.750 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
6	9.969 (0, 1)	9.063 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
5	8.313 (1, 1)	7.813 (1, 1)	7.000 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
4	7.344 (0, 1)	6.875 (0, 1)	6.125 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
3	4.875 (1, 1)	4.750 (1, 1)	4.500 (1, 1)	4.000 (1, 1)	2.500 (1, 1)
2	3.906 (0, 1)	3.813 (0, 1)	3.625 (0, 1)	3.250 (0, 1)	2.500 (1, 1)
1	0.969 (1, 0)	0.938 (1, 0)	0.875 (1, 0)	0.750 (1, 0)	0.500 (1, 0)
0	0.000 (0, 0)	0.000 (0, 0)	0.000 (0, 0)	0.000 (0, 0)	0.000 (0, 0)

Rows and columns stand for the respective resource level x and the decision period n and each entry presents the expected revenue and the admission decisions (1 = acceptance and 0 = rejection) for Type-1 and Type-2 orders, respectively

Second, we consider non-unit resource requirements. When requirements are not unit sized, the behavior of the optimal policy is rather complex. In general, the monotonicity properties of $g_n(x)$ discussed so far do not hold in case of non-unit resource requirements (see e.g. [Lee and Hersh 1993](#); [Brumelle and Walczak 2003](#)). This is due to the combinatorial behavior of the problem which derives from the large variety of options to match available resources with the resource requirements of orders of different types. We illustrate the effect of non-unit resource requirements with a simple example.

Example 2 Consider again the problem explained in Example 1. Here we make a simple change in the parameter values of order Type-2. Let the resource requirement of Type-2 orders w_2 be 2 units rather than 1, and let the revenue gained by fulfilling a Type-2 order be 4 rather than 2. We leave the rest of the parameters unchanged.

We evaluate initial resource levels $[0, 10]$. Since $w_2 = 2$, the initial resource level of ten units is the maximum amount of resource that could possibly be used in this example. Table 2 presents the optimal expected rewards and optimal admission decisions corresponding to each decision period and resource level.

It is easy to observe that the critical resource levels are non-existent in this case. Let us consider the admission decisions regarding Type-1 orders in period 0. The optimal admission decision here is to accept Type-1 orders at resource levels 1, 3, 5, 7, 8, 9, 10 and to reject them at resource levels 2, 4, 6. Hence, there is no critical resource level for Type-1 orders. It is easy to interpret this result. When the resource level is an odd number, after allocating the available resource to more profitable orders, that is, Type-2 orders which have a resource requirement of 2 units, the remaining one unit of slack resource can only be allocated to Type-1 orders. This is a simple illustration of matching available resources to the resource requirements of different order types. It is not profitable to preserve, e.g., the first, the third, or the fifth unit of resource for possible future orders of Type-2. Hence they should be allocated to orders of Type-1.

It is also possible to show that the optimal policy possesses an irregular behavior with respect to the remaining decision periods. The interested reader is referred to [Brumelle and Walczak \(2003\)](#) for further examples illustrating this type of irregularity.

Example 2 clearly shows that the optimal admission policy presents an irregular behavior with respect to the available resource level in case of non-unit resource requirements. In other words, the optimal admission decision regarding an order type may switch from acceptance to rejection and then from rejection to acceptance a number of times on the resource level axis given a decision period. Implementing such a policy in practice is very difficult, since it would require a careful examination of the resource level upon arrival of an order over time. There has been some work on characterizing the special cases where optimal admission policies are still of threshold-type (see e.g. [Papastavrou et al. 1996](#); [Brumelle and Walczak 2003](#)). For example, if splitting orders is allowed (i.e. partial admission), then the optimal policy is still threshold-type. However, these special cases are rather restrictive and do not hold for the problem addressed in this study.

Finally, we will discuss the stochasticity of resource requirements. So far we have not explicitly considered this specific characteristic of the problem we address in this study. Nonetheless, the discussion provided in this section can be generalized to the problem with stochastic requirements. The problem with deterministic resource requirements is a special case of the problem with random resource requirements. Consequently, we know that the optimal revenue function of the stochastic problem shows an irregular behavior as in the deterministic case. Thus, the non-optimality of the simple threshold-type policies also applies to the problem with stochastic resource requirements.

Taken all together, these observations show that it is fairly easy to model and solve the resource allocation problem in food processes via standard approaches from the literature. However, the resulting policies are rather complex and difficult to implement in practice.

5 Heuristic approaches

We have shown that the optimal admission policy of the problem under consideration does not have an easily implementable structure. In this section, we propose two heuristic approaches which follow simple decision rules and, therefore, can easily be implemented in practice. In the following subsections we provide the details of these approaches which we refer to as two-band heuristic and first-come-first-served heuristic.

5.1 The two-band heuristic

The two-band (TB) heuristic limits the irregular behavior of the optimal policy and provides simple decision rules regarding resource levels. The underlying intuition of the TB heuristic is based on two simple arguments:

1. It must be profitable to accept an order when the resource level is “sufficiently high” that it is not necessary to preserve resources for future orders with higher rewards.
2. It must be profitable to accept an order when the resource level is “sufficiently low” that it is not possible to accept future orders with higher rewards because of their larger resource requirements.

Henceforth, one can think of two bands on the resource level axis for each order type such that an incoming order is accepted whenever the resource level lies within one of these bands. We refer to those bands as the higher and the lower acceptance bands. Each band can be characterized by the critical resource levels setting its upper and lower bounds. In other words, the higher and lower acceptance bands of order type- i in period n involves the respective resource levels within $[\underline{x}_{in}^{\text{high}}, \bar{x}_{in}^{\text{high}}]$ and $[\underline{x}_{in}^{\text{low}}, \bar{x}_{in}^{\text{low}}]$. Since there are only two acceptance bands along the resource levels axis, the resulting admission policy under the TB heuristic is very simple. The admission decision regarding a given order type only switches at the boundaries of the two acceptance bands and remains the same for all other resource levels.

When those bounds characterizing the higher and the lower acceptance bands are known, the revenue function of the TB heuristic $g_n^{\text{TB}}(x)$ can be written as

$$g_n^{\text{TB}}(x) = \sum_{i=1}^m p_{in} v_{in}^{\text{TB}}(x) + p_{0n} g_{n+1}^{\text{TB}}(x) \quad (6)$$

where

$$v_{in}^{\text{TB}}(x) = \begin{cases} r_i + E[g_{n+1}^{\text{TB}}(x - w_i)] & \text{if } \underline{x}_{in}^{\text{low}} \leq x \leq \bar{x}_{in}^{\text{low}} \text{ or } \underline{x}_{in}^{\text{high}} \leq x \leq \bar{x}_{in}^{\text{high}} \\ g_{n+1}^{\text{TB}}(x) & \text{otherwise} \end{cases} \quad (7)$$

with the terminal revenue function given in Eq. (3).

We design the TB heuristic as a simplified version of the optimal policy, which ignores most of the irregularities in the optimal revenue function. Hence, the critical resource levels can be obtained by a simple search procedure within the backward recursion used for the optimal DP. Since the higher acceptance band corresponds to sufficiently high resource levels, we can assume that $\bar{x}_{in}^{\text{high}} = \infty$. In other words, type- i orders are accepted in period n whenever the resource level is higher than $\underline{x}_{in}^{\text{high}}$. The remaining bounds $\underline{x}_{in}^{\text{high}}$, $\underline{x}_{in}^{\text{low}}$, and $\bar{x}_{in}^{\text{low}}$ are then obtained by means of the following equations:

$$\underline{x}_{in}^{\text{high}} = \sup \left\{ x + 1 : r_i < E[g_{n+1}^{\text{TB}}(x - w_i)] - g_{n+1}^{\text{TB}}(x) \right\} \quad (8)$$

$$\underline{x}_{in}^{\text{low}} = \inf \left\{ x : r_i \geq E[g_{n+1}^{\text{TB}}(x - w_i)] - g_{n+1}^{\text{TB}}(x) \right\} \quad (9)$$

$$\bar{x}_{in}^{\text{low}} = \inf \left\{ x - 1 : r_i < E[g_{n+1}^{\text{TB}}(x - w_i)] - g_{n+1}^{\text{TB}}(x), x > \underline{x}_{in}^{\text{low}} \right\} \quad (10)$$

Table 3 The expected revenues and admission decisions for Example 3

x/n	0	1	2	3	4
10	12.500 (1, 1)	10.000 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
9	12.375 (1, 1)	10.000 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
8	11.906 (1, 1)	10.000 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
7	10.938 (1, 1)	9.750 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
6	9.969 (0, 1)	9.063 (1, 1)	7.500 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
5	8.250 (0, 1)	7.813 (1, 1)	7.000 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
4	7.344 (0, 1)	6.875 (0, 1)	6.125 (1, 1)	5.000 (1, 1)	2.500 (1, 1)
3	4.813 (0, 1)	4.688 (0, 1)	4.500 (1, 1)	4.000 (1, 1)	2.500 (1, 1)
2	3.906 (0, 1)	3.813 (0, 1)	3.625 (0, 1)	3.250 (0, 1)	2.500 (1, 1)
1	0.969 (1, 0)	0.938 (1, 0)	0.875 (1, 0)	0.750 (1, 0)	0.500 (1, 0)
0	0.000 (0, 0)	0.000 (0, 0)	0.000 (0, 0)	0.000 (0, 0)	0.000 (0, 0)

Rows and columns stand for the respective resource level x and the decision period n and each entry presents the expected revenue and the admission decisions (1 = acceptance and 0 = rejection) for Type-1 and Type-2 orders, respectively

There are three possible cases regarding the existence of the acceptance bands for a given order type i and decision period n : (i) both acceptance bands exist, that is, there are at least two non-consecutive resource levels where an order will be accepted ($\underline{x}_{in}^{\text{high}} > \bar{x}_{in}^{\text{low}}$); (ii) only a single acceptance band exists, that is, there is at least one block of resource levels where an order will be accepted ($\underline{x}_{in}^{\text{high}} = \bar{x}_{in}^{\text{low}}$); and (iii) neither the upper nor the lower band exists, that is, there is no resource level where an order will be accepted ($\{x : r_i < E[g_{n+1}^{\text{TB}}(x - w_i)] - g_{n+1}^{\text{TB}}(x)\} = \emptyset$).

For all periods, the search procedure systematically evaluates resource levels in terms of the condition $r_i \geq E[g_{n+1}^{\text{TB}}(x - w_i)] - g_{n+1}^{\text{TB}}(x)$ which specifies whether accepting a type- i order leads to a non-negative increment in the expected revenue. The procedure first checks the existence of acceptance bands. Then, it sets the upper and lower bounds of the higher and lower bands from the initial resource level downwards and from 0 upwards, respectively. At each iteration the procedure also computes the revenue function of the TB heuristic following Eqs. (6) and (7).

We illustrate the basic principles of the TB heuristic by means of a simple example.

Example 3 Let us consider again the problem sketched in Example 2 and use the TB heuristic rather than the optimal policy. Table 3 presents the expected rewards and admission decisions corresponding to each decision period and each resource level.

We can observe the acceptance bands in all decision periods characterizing the TB heuristic. For example, in period 0, for Type-1 orders, the higher acceptance band involves resource levels $[7, \infty)$, and the lower acceptance band involves the resource level of 1 unit. When compared to the results regarding the optimal policy (see Table 2) we can see that the admission policy of the TB heuristic is more stable. However, there is also a loss in the expected revenue due to not making the optimal admission decisions. Consider period 0 with the initial resource level of 5 units. If the optimal admission policy is implemented then the expected revenue will be 8.313, whereas,

if the heuristic admission policy is implemented then the expected revenue will be 8.250.

5.2 The first-come-first-served heuristic

The first-come-first-served (FCFS) heuristic is an approach where incoming orders are attended to in the sequence they arrive. It thus addresses the case where no action is taken to ration the resource inventory. Henceforth, it is a logical benchmark to gauge the effectiveness of any allocation policy. In our numerical analysis, we use the FCFS heuristic to assess the optimal policy and the TB heuristic.

The optimal admission policy depends on the revenues gained as immediate results of the acceptance/rejection actions and the expected revenues associated with ensuing resource levels. The FCFS heuristic considers only the former and accepts any incoming order which will lead to an immediate nonnegative increment in the expected revenue. That is, a type- i order is accepted in period n if $r_i \geq z E[(x - w_i)^-]$. Note that the decision rule is independent of the possible revenues associated with the subsequent periods. Furthermore, it does not rely on the period in which the decision is made. Hence, it can be translated into a static acceptance threshold for each order type. That is, the FCFS policy accepts a type- i order only if the resource level is higher than an order specific threshold. We denote this threshold by \underline{x}_i . We can then write

$$\underline{x}_i = \inf \{x : r_i \geq z E[(x - w_i)^-]\}. \quad (11)$$

Consequently, the revenue function of the FCFS heuristic $g_n^{\text{FCFS}}(x)$ can be written as

$$g_n^{\text{FCFS}}(x) = \sum_{i=1}^m p_{in} v_{in}^{\text{FCFS}}(x) + p_{0n} g_{n+1}^{\text{FCFS}}(x), \quad (12)$$

where

$$v_{in}^{\text{FCFS}}(x) = \begin{cases} r_i + E[g_{n+1}^{\text{FCFS}}(x - w_i)] & \text{if } x \geq \underline{x}_i \\ g_{n+1}^{\text{FCFS}}(x) & \text{otherwise} \end{cases} \quad (13)$$

with the terminal revenue function given in Eq. (3).

6 Numerical study

We conduct numerical studies in order to analyze

1. The performances of the FCFS and the TB heuristics.
2. the effects of considering the stochasticity of resource requirements, and
3. the effects of penalty and disposal costs.

6.1 Experiment settings

As we discussed, the complex structure of the optimal admission policy is due to the combinatorial behavior of the problem. This behavior derives from the large variety of options to match available resources with the resource requirements of orders of different types. Hence, in our numerical study our aim is to investigate cases with a variety of order types and parameters characterizing different system configurations.

An arbitrary problem instance is characterized by a set of parameters: rewards r_i , probability mass functions τ_i , and arrival probabilities p_{in} associated with each order type i ; the number of periods t , the initial inventory level s , the penalty cost z , and the disposal cost c . We assume that the resource requirements of orders follow a discretized truncated normal distribution. We characterize the stochasticity of the resource requirements by a coefficient of variation ρ common for all order types. Notice that this assumption enables us to uniquely characterize the probability mass of each order type τ_i for a given mean \bar{w}_i and a coefficient of variation ρ .

We consider three main classes of random instances involving $m \in \{2, 5, 10\}$ types of orders. For each of these classes, we generate four sub-classes by imposing the coefficient of variation levels $\rho \in \{0.00, 0.05, 0.15, 0.25\}$. For each sub-class, we randomly generate 10^3 instances with various rewards r_i , average resource requirements \bar{w}_i , and arrival probabilities p_{in} . The rewards and average resource requirements of each order type are selected from the set $\{10, 20, \dots, 100\}^2$ with uniform probability. The arrival probabilities are assumed to be stationary over time, that is, $p_{in} = p_i$. The probability of no order arrival p_0 is set to 0.2. The arrival probability of each order is selected from $(0, 1)$ with uniform probability. They are then normalized such that the arrival probabilities of orders and the no arrival probability sum up to 1. The number of periods t is fixed at 20.

For all instances we set the number of periods and compute the expected revenues of both the optimal policy and the heuristic approaches g , g^{FCFS} , and g^{TB} (we omit indices for simplicity's sake). In order to characterize the respective performances of the FCFS and the TB heuristics we define $\Delta^{\text{TB}} = (1 - g^{\text{TB}}/g) \times 100$ and $\Delta^{\text{FCFS}} = (1 - g^{\text{FCFS}}/g) \times 100$.

We consider the initial inventory levels $x \in [0, 2\xi]$ where ξ is the expected total resource requirement, i.e., $\xi = \sum_{i=1}^m \sum_{n=0}^{t-1} p_{in} \bar{w}_i$. It is hard, however, to reflect all x values within the given range individually. Hence, rather than reporting the expected revenue for each x , we report the average expected revenues for a set of x values. In order to do so, we divide the whole domain $[0, 2\xi]$ into 20 intervals with equal lengths each covering 10% of the domain. For each interval $k \in \{1, \dots, 20\}$ we report the average g , g^{FCFS} , and g^{TB} for $(k-1)\xi/20 < x \leq k\xi/20$. The number of random instances sums up to 4×10^3 for each sub-class and to 12×10^3 in total. We believe that this broad collection of instances should allow us to address some practical cases found in real-life applications.

Having generated a collection of random instances, we can now investigate the performance of any admission policy given a penalty cost z and a disposal cost c .

6.2 Numerical results and insights

In what follows, we discuss our findings in detail regarding the points raised at the beginning of this section.

6.2.1 The performances of the FCFS and the TB heuristics

In this sub-section we analyze the performances of the FCFS and the TB heuristics with respect to the optimal policy. We conduct a set of experiments considering all random instances with the respective penalty and disposal costs $z = 10$ and $c = 0.5$.

For all resource level intervals $k \in \{1, \dots, 20\}$ we report on g , Δ^{TB} , and Δ^{FCFS} . The results can be found in Table 4, 5 and 6. These results show that g , Δ^{FCFS} , and Δ^{TB} are severely affected by the resource level, the coefficient of variation, and the number of order types.

To start with, it is interesting to examine the behavior of the revenue function g . From our discussion we know that g is not necessarily concave on the resource level. Nevertheless, we can observe that g first tends to increase and then to decrease with

Table 4 The optimal expected reward and relative errors of heuristic approaches averaged over the class of random instances involving two order types with $z = 10$ and $c = 0.5$

k	$\rho = 0$			$\rho = 0.05$			$\rho = 0.15$			$\rho = 0.25$		
	g	Δ^{FCFS}	Δ^{TB}	g	Δ^{FCFS}	Δ^{TB}	g	Δ^{FCFS}	Δ^{TB}	g	Δ^{FCFS}	Δ^{TB}
20	473.57	0.00	0.00	473.57	0.00	0.00	473.57	0.00	0.00	473.55	0.00	0.00
19	517.27	0.00	0.00	517.27	0.00	0.00	517.27	0.00	0.00	517.19	0.00	0.00
18	560.96	0.00	0.00	560.96	0.00	0.00	560.95	0.00	0.00	560.81	0.00	0.00
17	604.62	0.00	0.00	604.62	0.00	0.00	604.59	0.00	0.00	604.38	0.00	0.00
16	648.20	0.00	0.00	648.19	0.00	0.00	648.13	0.00	0.00	647.81	0.00	0.00
15	691.55	0.01	0.00	691.54	0.01	0.00	691.40	0.01	0.00	690.85	0.01	0.00
14	734.25	0.02	0.00	734.21	0.02	0.00	733.85	0.02	0.00	732.75	0.02	0.00
13	774.67	0.06	0.00	774.53	0.06	0.00	773.42	0.06	0.00	770.83	0.07	0.00
12	806.21	0.18	0.01	805.84	0.18	0.01	803.18	0.19	0.00	797.94	0.21	0.00
11	816.16	0.58	0.05	815.30	0.58	0.02	810.84	0.58	0.01	803.36	0.60	0.00
10	795.15	1.62	0.12	794.37	1.58	0.06	789.52	1.52	0.01	781.54	1.51	0.01
9	751.99	3.50	0.23	750.57	3.44	0.12	745.18	3.32	0.02	737.68	3.22	0.01
8	694.37	6.33	0.31	693.23	6.24	0.16	688.21	6.01	0.03	681.13	5.78	0.01
7	630.52	9.73	0.37	629.21	9.65	0.19	624.20	9.36	0.04	617.39	9.03	0.01
6	560.80	13.52	0.42	559.44	13.45	0.22	554.21	13.13	0.04	547.52	12.69	0.02
5	482.39	17.60	0.48	481.39	17.50	0.26	476.79	17.04	0.05	470.28	16.50	0.02
4	396.37	21.29	0.56	395.20	21.25	0.32	390.08	20.85	0.06	383.73	20.17	0.02
3	296.91	24.95	0.71	296.05	24.85	0.41	291.77	24.08	0.08	285.69	23.12	0.03
2	183.56	26.13	0.96	182.63	26.09	0.59	178.23	25.25	0.15	172.65	23.72	0.06
1	49.62	14.57	0.45	49.14	14.51	0.36	46.38	13.26	0.24	42.44	11.33	0.23

Table 5 The optimal expected reward and relative errors of heuristic approaches averaged over the class of random instances involving five order types with $z = 10$ and $c = 0.5$

k	$\rho = 0$			$\rho = 0.05$			$\rho = 0.15$			$\rho = 0.25$		
	g	Δ^{FCFS}	Δ^{TB}	g	Δ^{FCFS}	Δ^{TB}	g	Δ^{FCFS}	Δ^{TB}	g	Δ^{FCFS}	Δ^{TB}
20	475.86	0.00	0.00	475.86	0.00	0.00	475.86	0.00	0.00	475.85	0.00	0.00
19	519.85	0.00	0.00	519.85	0.00	0.00	519.85	0.00	0.00	519.78	0.00	0.00
18	563.83	0.00	0.00	563.83	0.00	0.00	563.82	0.00	0.00	563.69	0.00	0.00
17	607.78	0.00	0.00	607.77	0.00	0.00	607.75	0.00	0.00	607.53	0.00	0.00
16	651.62	0.00	0.00	651.61	0.00	0.00	651.54	0.00	0.00	651.17	0.00	0.00
15	695.13	0.01	0.00	695.11	0.01	0.00	694.90	0.01	0.00	694.18	0.01	0.00
14	737.51	0.03	0.00	737.44	0.03	0.00	736.86	0.03	0.00	735.35	0.04	0.00
13	776.28	0.09	0.00	776.08	0.09	0.00	774.66	0.09	0.00	771.59	0.11	0.00
12	805.36	0.29	0.01	804.95	0.29	0.01	802.19	0.30	0.00	796.97	0.34	0.00
11	816.16	0.89	0.02	815.51	0.89	0.01	811.51	0.90	0.01	804.56	0.93	0.00
10	804.62	2.35	0.05	803.80	2.34	0.03	799.16	2.30	0.01	791.58	2.27	0.01
9	775.14	5.04	0.08	774.20	5.02	0.05	769.36	4.89	0.01	761.80	4.72	0.01
8	733.92	8.91	0.13	732.89	8.88	0.07	727.95	8.63	0.02	720.54	8.28	0.01
7	683.45	13.47	0.18	682.35	13.44	0.10	677.32	13.07	0.03	670.05	12.54	0.01
6	623.85	18.35	0.26	622.66	18.30	0.15	617.56	17.83	0.04	610.47	17.13	0.02
5	554.33	23.35	0.37	553.08	23.30	0.21	547.96	22.72	0.05	541.06	21.84	0.02
4	473.25	28.36	0.51	471.97	28.30	0.29	466.90	27.61	0.07	460.30	26.54	0.03
3	377.07	33.18	0.67	375.84	33.12	0.41	370.97	32.27	0.10	364.74	30.91	0.04
2	258.51	36.28	0.88	257.44	36.25	0.57	253.05	35.07	0.16	247.50	33.05	0.07
1	93.34	26.20	0.60	92.83	26.25	0.43	90.27	24.90	0.21	86.76	22.22	0.15

increasing resource levels in general. It is obvious that the optimal policy is more selective in accepting orders for resource levels where g tends to increase. Thus employing admission policies is mainly critical when resource levels are low.

The behavior of g is also reflected in the performance of the heuristic approaches. Both Δ^{FCFS} and Δ^{TB} tend to decrease with increasing resource levels. That is, the importance of making the optimal admission decisions decreases with increasing resource levels. One exception is the case with extremely low resource levels. Then Δ^{FCFS} and Δ^{TB} may increase moving from the resource level interval $k = 1$ to $k = 2$ (see e.g. Table 4). This result is also intuitive since the number of order types for which sufficient resources can be provided is very limited for those resource levels. As a result, the optimal policy in this case cannot be very selective in accepting orders. It is important to note that the TB heuristic is very competitive for all resource levels with a maximum Δ^{TB} of 0.96%.

Obviously the expected revenue decreases with the degree of stochasticity of the resource requirements of orders, regardless of the policy employed. Since the optimal policy is the one best suited to handle stochasticity, one may expect that it will perform relatively better than the other heuristics for high ρ values. However, the numerical

Table 6 The optimal expected reward and relative errors of heuristic approaches averaged over the class of random instances involving ten order types with $z = 10$ and $c = 0.5$

k	$\rho = 0$			$\rho = 0.05$			$\rho = 0.15$			$\rho = 0.25$		
	g	Δ^{FCFS}	Δ^{TB}	g	Δ^{FCFS}	Δ^{TB}	g	Δ^{FCFS}	Δ^{TB}	g	Δ^{FCFS}	Δ^{TB}
20	459.91	0.00	0.00	459.91	0.00	0.00	459.91	0.00	0.00	459.91	0.00	0.00
19	503.70	0.00	0.00	503.70	0.00	0.00	503.70	0.00	0.00	503.64	0.00	0.00
18	547.47	0.00	0.00	547.47	0.00	0.00	547.46	0.00	0.00	547.35	0.00	0.00
17	591.22	0.00	0.00	591.21	0.00	0.00	591.19	0.00	0.00	590.99	0.00	0.00
16	634.87	0.00	0.00	634.86	0.00	0.00	634.78	0.00	0.00	634.41	0.00	0.00
15	678.15	0.01	0.00	678.12	0.01	0.00	677.88	0.01	0.00	677.10	0.01	0.00
14	720.07	0.02	0.00	719.98	0.02	0.00	719.31	0.03	0.00	717.65	0.04	0.00
13	757.84	0.09	0.00	757.62	0.09	0.00	756.08	0.11	0.00	752.86	0.13	0.00
12	785.78	0.33	0.00	785.37	0.33	0.00	782.62	0.36	0.00	777.46	0.40	0.00
11	797.07	1.02	0.01	796.48	1.02	0.01	792.66	1.04	0.00	785.96	1.08	0.00
10	789.08	2.62	0.02	788.37	2.63	0.01	783.97	2.60	0.01	776.60	2.57	0.00
9	765.44	5.55	0.02	764.66	5.55	0.02	760.04	5.44	0.01	752.56	5.26	0.01
8	730.17	9.65	0.03	729.34	9.65	0.02	724.64	9.41	0.01	717.21	9.03	0.01
7	684.61	14.37	0.05	683.74	14.36	0.03	679.02	14.01	0.01	671.72	13.43	0.01
6	629.02	19.32	0.07	628.13	19.32	0.04	623.43	18.86	0.02	616.32	18.09	0.01
5	563.08	24.36	0.10	562.17	24.36	0.06	557.54	23.81	0.02	550.69	22.85	0.01
4	485.68	29.45	0.15	484.75	29.45	0.09	480.26	28.79	0.03	473.75	27.62	0.02
3	393.79	34.38	0.23	392.89	34.40	0.14	388.65	33.60	0.05	382.64	32.14	0.03
2	279.78	38.14	0.34	278.98	38.19	0.23	275.26	37.11	0.09	270.05	35.03	0.05
1	114.17	28.99	0.40	113.75	29.12	0.32	111.57	27.83	0.19	108.42	24.98	0.15

results show that both Δ^{FCFS} and Δ^{TB} decrease with increasing ρ , especially for low resource levels. This result shows that the optimal policy is not robust with respect to the degree of stochasticity, whereas both the FCFS and the TB heuristics are. This is a rather interesting result in the sense that simple control rules perform relatively better when a complicating factor such as stochasticity is higher.

It is clear that the optimal policy is more selective when the number of order types is large, since this leads to a variety of options to preserve resources for more profitable orders. This can be observed by considering the performance of the FCFS policy which does not preserve resources for future orders. Regardless of the resource level or the degree of the stochasticity, Δ^{FCFS} increases with the number of order types. In contrast to the FCFS heuristic, the performance of the TB heuristic improves as the number of order types increases. This can be clarified by considering the structure of the TB heuristic. The gap between the optimal policy and the TB heuristic stems from the irregular behavior of the revenue function which is mostly neglected by the TB heuristic. This irregular behavior arises because of the dissimilarity of the order types in terms of revenues and resource requirements. Note that the order types are picked randomly from a bounded set. As a consequence the similarity between them

Table 7 The expected rewards of the optimal policy and heuristic approaches: The comparison of policies computed for the stochastic and deterministic resource requirements cases both applied in the stochastic resource requirements case. Results are averaged over the class of random instances involving five order types with $z = 10$ and $c = 0.5$

		$\rho = 0.05$			$\rho = 0.15$			$\rho = 0.25$		
		g	g^{FCFS}	g^{TB}	g	g^{FCFS}	g^{TB}	g	g^{FCFS}	g^{TB}
5	STOCH	553.08	424.22	551.93	547.96	423.46	547.70	541.06	422.87	540.95
	DET	552.86	423.96	549.57	546.16	418.03	541.90	536.61	407.65	529.44
4	STOCH	471.97	338.39	470.59	466.90	337.99	466.59	460.30	338.13	460.17
	DET	471.73	338.11	467.54	464.92	332.11	458.92	455.52	321.73	445.53
3	STOCH	375.84	251.34	374.29	370.97	251.25	370.60	364.74	251.98	364.60
	DET	375.61	251.06	370.15	368.91	245.11	360.18	359.84	234.72	345.45
2	STOCH	257.44	164.11	255.97	253.05	164.30	252.66	247.50	165.69	247.33
	DET	257.27	163.81	250.43	251.23	157.78	239.29	242.98	147.31	223.23
1	STOCH	92.83	68.47	92.43	90.27	67.79	90.08	86.76	67.48	86.63
	DET	92.79	68.36	89.84	89.63	65.10	84.55	84.86	59.82	76.67

increases with the number of order types. Thus, larger number of order types results in a more lenient optimal policy and thus positively affects the performance of the TB heuristic.

6.2.2 The effects of considering the stochasticity of resource requirements

We analyze here what happens if we ignore the stochasticity of resource requirements and follow the admission decisions tailored to the set of instances with deterministic resource requirements (i.e. $\rho = 0.00$) for instances characterized by stochastic resource requirements (i.e. $\rho \in \{0.05, 0.15, 0.25\}$). Here, we only consider the class of random instances involving five order types for simplicity's sake, since the results are analogous with the other sub-classes. We use the respective penalty and disposal costs $z = 10$ and $c = 0.5$. We only consider the initial inventory levels corresponding to $k = \{1, \dots, 5\}$, since the importance of stochasticity becomes negligible for higher resource levels. The results related to this set of experiments are given in Table 7.

One obvious observation is that as ρ increases, the gap between the stochastic and deterministic approaches gradually increases. In addition to this, for all policies, the gap between the deterministic and stochastic approach is relatively higher when the resource level is lower. This is due to the fact that the variation with respect to the total resource requirements of all prospective orders is lower than the sum of the variations of each prospective order. This is usually referred to as the pooling effect.

We can also observe that ignoring the uncertainty results in relatively larger losses for the heuristic approaches as compared to the optimal policy. Thus, especially when a heuristic approach is being used one should be certain that the stochasticity in material requirements is correctly accounted for.

6.2.3 The effects of penalty and disposal costs

Finally, we analyze the effects of penalty and disposal costs. Here, we only consider the sub-class of random instances involving five order types for simplicity's sake. Nevertheless, we would like to note that the results are very similar for the other sub-classes. We first fix the number of periods at 20 and consider the initial inventory levels corresponding to $k = \{1, \dots, 5\}$. In order to analyze the effect of penalty cost, we fix the disposal cost at 0.5, and consider the penalty costs $\{10, 12, 14\}$. Similarly, in order to analyze the effect of disposal costs, we fix the penalty cost at 10, and consider the disposal costs $\{0.5, 1.0, 2.0\}$. We compute and report on the expected revenues of all proposed policies, that is, g , g^{FCFS} , and g^{TB} . The results can be seen in Tables 8 and 9.

The effects of penalty and disposal costs on the proposed policies are rather straightforward. As can be observed, increasing penalty and disposal costs negatively effects the expected rewards of all proposed policies. This effect is stronger when the resource level is rather low. Furthermore, the effects of those cost parameters are more severe when the stochasticity of resource requirements is larger.

7 Conclusions and extensions

We addressed the problem of determining the order acceptance/rejection decisions in a food processing system where a single raw material is processed into a set of different orders. We considered some specific characteristics of the food processing industry, such as random raw material requirements of orders, shortage penalty costs and disposal costs which have not yet been addressed in the literature. Our contribution is threefold. First, we showed that the problem can be modeled and solved as a dynamic program. Second, since the optimal admission policy does not follow simple decision rules, we provided a heuristic approach, which we referred to as the TB heuristic, based on intuitive decision rules which can obtain good results. Third, with an extensive numerical study we examined the effects of various parameters on admission policies and pointed out those cases where it is critical to employ admission policies.

The main conclusions of our numerical study can be summarized as follows. We compared the optimal policy with the FCFS heuristic in order to see how critical it is to employ the optimal policy. We observed that employing the optimal admission policy is essentially important in case of limited resource levels. Obviously, when the initial inventory level can be set freely, there is hardly any need to use an admission policy since it will be optimal to accept most of the orders. We also saw that the penalty of not using the optimal policy is higher when there is a larger number of order types with a lower degree of stochasticity in their resource requirements. We observed that the overall performance of the TB approach is very good. The relative gap between the optimal policy and the TB heuristic narrows down quickly as the resource level increases. Also, the TB heuristic performs relatively better in cases characterized by a large number of order types and a high degree of stochasticity of the resource requirements of orders. We saw that considering the stochasticity of the resource requirements is very critical, especially when heuristic approaches are being

Table 8 The expected rewards of the optimal policy and heuristic approaches: The comparison of policies with respect to different levels of penalty costs $z \in \{10, 12, 14\}$. Results are averaged over the class of random instances involving five order types with $c = 0.5$

k	z	$\rho = 0$			$\rho = 0.05$			$\rho = 0.15$			$\rho = 0.25$		
		g	g^{FCFS}	g^{TB}	g	g^{FCFS}	g^{TB}	g	g^{FCFS}	g^{TB}	g	g^{FCFS}	g^{TB}
5	10	554.33	424.87	552.40	553.08	424.22	552.02	547.96	423.46	547.77	541.06	422.87	541.01
	12	553.60	424.00	551.44	552.19	423.55	551.08	546.58	423.11	546.38	541.06	422.45	539.13
	14	553.08	423.59	551.03	551.52	423.17	550.38	545.49	422.83	545.28	537.68	421.90	537.62
5	10	473.25	339.04	470.94	471.97	338.39	470.67	466.90	337.99	466.67	460.30	338.13	460.24
	12	472.45	338.15	469.82	471.00	337.73	469.64	465.44	337.75	465.19	460.30	337.89	458.26
	14	471.88	337.76	469.42	470.28	337.39	468.87	464.29	337.58	464.03	456.74	337.47	456.67
5	10	377.07	251.97	374.62	375.84	251.34	374.38	370.97	251.25	370.68	364.74	251.98	364.67
	12	376.16	251.06	373.26	374.76	250.68	373.21	369.38	251.11	369.08	364.74	251.93	362.56
	14	375.50	250.68	372.90	373.95	250.37	372.37	368.14	251.01	367.83	360.96	251.62	360.87
5	10	258.51	164.73	256.30	257.44	164.11	256.06	253.05	164.30	252.75	247.50	165.69	247.42
	12	257.37	163.81	254.60	256.14	163.46	254.68	251.25	164.25	250.93	247.50	165.80	245.07
	14	256.54	163.44	254.23	255.18	163.18	253.69	249.85	164.23	249.52	243.30	165.61	243.21
5	10	93.34	68.89	92.89	92.83	68.47	92.55	90.27	67.79	90.20	86.76	67.48	86.74
	12	91.65	67.76	90.64	91.05	67.40	90.75	88.11	66.89	88.04	86.76	66.37	84.16
	14	90.43	67.09	90.02	89.75	66.78	89.45	86.48	66.23	86.41	82.19	65.28	82.17

Table 9 The expected rewards of the optimal policy and heuristic approaches: The comparison of policies with respect to different levels of disposal costs $c \in \{0.5, 1.0, 2.0\}$. Results are averaged over the class of random instances involving five order types with $z = 10$

k	c	$\rho = 0$			$\rho = 0.05$			$\rho = 0.15$			$\rho = 0.25$		
		g	FCFS _g	g ^{TB}	g	FCFS _g	g ^{TB}	g	FCFS _g	g ^{TB}	g	FCFS _g	g ^{TB}
5	0.5	554.33	424.87	552.40	553.08	424.22	552.02	547.96	423.46	547.77	541.06	422.87	541.01
	1.0	544.16	419.42	541.01	542.61	418.76	540.90	536.45	417.56	536.17	528.32	416.10	528.26
	2.0	527.87	408.52	521.29	525.77	407.83	522.35	517.90	405.75	517.39	507.96	402.54	507.87
4	0.5	473.25	339.04	470.94	471.97	338.39	470.67	466.90	337.99	466.67	460.30	338.13	460.24
	1.0	464.52	334.04	460.84	462.94	333.37	460.88	456.88	332.56	456.53	449.17	331.94	449.09
	2.0	450.43	324.03	442.86	448.28	323.33	444.23	440.51	321.72	439.90	431.14	319.55	431.03
3	0.5	377.07	251.97	374.62	375.84	251.34	374.38	370.97	251.25	370.68	364.74	251.98	364.67
	1.0	369.77	247.22	365.87	368.26	246.59	366.00	362.47	246.16	362.07	355.30	246.21	355.21
	2.0	357.88	237.73	349.78	355.83	237.08	351.46	348.44	235.99	347.73	339.78	234.66	339.66
2	0.5	258.51	164.73	256.30	257.44	164.11	256.06	253.05	164.30	252.75	247.50	165.69	247.42
	1.0	252.72	160.21	249.33	251.43	159.59	249.36	246.27	159.45	245.85	239.93	160.22	239.83
	2.0	243.18	151.18	235.94	241.45	150.54	237.51	234.87	149.76	234.17	227.29	149.28	227.15
1	0.5	93.34	68.89	92.89	92.83	68.47	92.55	90.27	67.79	90.20	86.76	67.48	86.74
	1.0	89.23	64.60	88.47	88.65	64.18	88.23	85.74	63.25	85.65	81.80	62.44	81.77
	2.0	82.06	56.04	79.61	81.34	55.60	80.09	77.82	54.16	77.66	73.17	52.34	73.13

used. We also observed that the effects of disposal and penalty costs are larger when the degree of stochasticity of the resource requirements of orders is higher.

There are two directions for further research worth exploring. First, the production capacities and lead-times could be considered. Our model neglects the production side of the system. As a result, our results do not readily apply to cases where production capacities are limited and/or lead-times are not negligible. It would be specifically interesting to consider a case where during processing an order (with unknown material consumption), other orders might arrive that have to be accepted or rejected without exactly knowing the resource level. Second, the model can be extended for systems involving multiple raw materials. Referring to the analogy with revenue management problems, this case corresponds to capacity control problems with multi-leg flights.

Acknowledgments The authors acknowledge the valuable comments and suggestions provided by two anonymous referees.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Akkerman R, van Donk DP (2009) Analyzing scheduling in the food-processing industry: structure and tasks. *Cogn Technol Work* 11(3):215–226
- Brumelle S, Walczak D (2003) Dynamic airline revenue management with multiple semi-Markov demand. *Oper Res* 51:137–148
- Flapper SDP, Fransoo JC, Broekmeulen R, Inderfurth K (2002) Planning and control of rework in the process industries: a review. *Prod Plan Control* 13(1):26–34
- Fransoo J, Rutten W (1994) A typology of production control situations in process industries. *Int J Oper Prod Manag* 14:47–57
- Grosfeld-Nir A, Gerchak Y (2004) Multiple lotsizing in production to order with random yields: review of recent advances. *Ann Oper Res* 126(1):43–69
- Kleywegt AJ, Papastavrou JD (2001) The dynamic and stochastic knapsack problem with random sized items. *Oper Res* 49:26–41
- Lee TC, Hersh M (1993) A model for dynamic airline seat inventory control with multiple seat bookings. *Transport Sci* 27:252–265
- Murthy DNP, Ma L (1996) Material planning with uncertain product quality. *Prod Plan Control* 7(6):566–576
- Papastavrou JD, Rajagopalan S, Kleywegt AJ (1996) The dynamic and stochastic knapsack problem with deadlines. *Manage Sci* 42:1706–1718
- Somsen D, Capelle A (2002) Introduction to production yield analysis—a new tool for improvement of raw material yield. *Trends Food Sci Technol* 13:136–145
- Talluri KT, Van Ryzin GJ (2004) The theory and practice of revenue management. Kluwer, Boston
- Van Donk DP (2000) Customer-driven manufacturing in the food processing industry. *Br Food J* 102:739–747
- Van Slyke R, Young Y (2000) Finite horizon stochastic knapsacks with applications to yield management. *Oper Res* 48:155–172