Solving large structured Markov Decision Problems for perishable inventory management and traffic control

Haijema, R.

Publication date
2008

Chapter 2

Perishable inventory theory

In this first part of the thesis we investigate the optimal production and inventory management of a perishable product with a (short) fixed life time. The focus of the study is on blood platelet pools (BPPs), which are the most expensive and most perishable blood products. BPPs are produced at blood banks and transfused to patients in hospitals. In Section 2.1 the platelet production problem (PPP) at blood banks is described in detail. Although some of the problem characteristics are typical for BPPs, one may generalize the problem to other perishable products, such as food as we will discuss in the epilogue in Chapter 9. The focus of the study is on determining (nearly) optimal production volumes. A similar problem arises in an inventory setting, where nearly optimal order quantities need to be set.

Next to the (production) order size, the production-inventory levels are controlled by the so-called issuing policy, which specifies the way in which demand is met. Therefore an overview of standard ordering and issuing policies is presented in Section 2.2. In Section 2.3 the history of modeling (perishable) inventory management is reviewed and many references are given. Finally, in Section 2.6, we formulate research questions to be answered in the next chapters.

2.1 Problem of interest

2.1.1 The platelet production problem (PPP)

The description of the PPP that is given in this section comes from conversations and collaborations with employees of Sanquin, the Dutch association of blood banks. We refer in particular to communications with M. Beun (Physician, [16]), C. Th. Smit Sibinga (Director Sanquin consulting services and member of the WHO, [133]) and B. Hinloopen (logistic manager at Sanquin, division North East, [66]). The Dutch case is to some degree representative for the case at a number of developed countries.

Demand for BPPs

Platelets are of live-saving importance. As small particles in the bloodstream, platelets prevent internal and external bleeding by recognizing and 'repairing' damaged blood vessels. Platelet's quality deteriorates rapidly even inside the blood stream, but most people's production of platelets at the bone marrow is sufficient to retain a safe level. Nevertheless after a major bleeding caused by a trauma or a surgery, patients may temporarily have a lack of platelets. These patients need to be transfused with platelet pools of *any* age up to the maximal shelf life. This demand category is henceforth called the *any-demand*, since there is no strong preference with respect to the age of the pools. The any-demand comprises about a third of the total demand.

The remaining two thirds of the demand is set by patients who suffer from a platelet function disorder. To keep the number of platelets in their blood at a safe level, these patients are transfused frequently (at least once a week) at the Hematology department with *'young'* pools that preferably are at most three days old, counted from the day the pools are added to stock. This demand category is referred to as the *young-demand*. Transfusing older platelet pools is allowed but less desirable.

The distinction between the any-demand and the young-demand is not clearly made at all blood banks. For other blood banks the young-demand may be a smaller portion of the total demand. Although part of the transfusions at the hospitals are scheduled, the demand volume is highly uncertain to blood banks.

According to the eight different blood groups as set by the four ABO-categories (A, B, AB and O) and the presence of the Rhesus-D factor (+ or −), one may distinguish eight different BPP products (A^+, A^-, etc.). Figure 2.1 shows the compatibility of the blood groups and the relative frequencies by which they are present in the Western population.

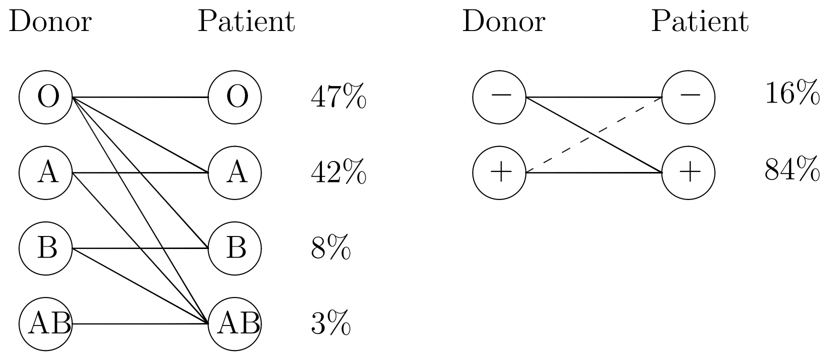

Figure 2.1: Compatibility of blood groups of donor and patient following the ABO-system and Rhesus-D system (+/−)

In Section 3.6 we deliberate on the compatibility of the blood groups. For now, observe that almost 90% of the Western population has blood group O or A and that the O^- donor is the universal donor that can help any patient.

Production process

Blood banks are responsible for the production and distribution of blood products in all sorts. To collect blood, appointments with voluntary donors are scheduled. Most commonly, a donor gives 500 *ml* of whole blood, but a minor part of the donors donates only specific blood components through the much more expensive apheresis technique. In The Netherlands, the fraction of platelet concentrates that comes from apheresis is about 10%. This fraction may be different in other countries, see [96]. We focus on the regular production of products from whole blood donations.

Figure 2.2 shows the product classes and the two processing steps. In a first processing step the red blood cells and plasma are filtered from a whole blood donation. The residual called, the *'buffy coat'*, contains a high concentration of blood platelets (or thrombocytes) and white blood cells. In a second processing step, the blood platelets are extracted from the buffy coats, since white cells might cause complications upon transfusion. The platelets of 5 donors of the same blood group are pooled together in a so-called *blood platelet pool* (BPP). A BPP corresponds to an adult platelet dose or one transfusion unit. Some patients need to be transfused with multiple BPPs. BPPs cannot be shared by patients.

The demand for whole blood donations is, in The Netherlands, dominated by the demand for red blood cell concentrates (RBCs) at hospitals. The residual plasma is used in plasma products and for the production of all kinds of medicines. The number of buffy coats, obtained as a side product of the production of RBCs, usually exceeds the demand for

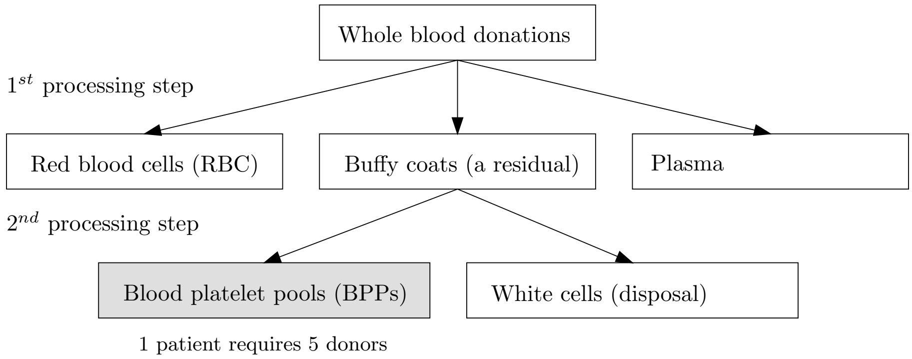

Figure 2.2: Product classes and processing steps at blood banks

platelet pools. As a consequence, on average only a third to a half of the buffy coats is used for the production of BPPs. Furthermore not all buffy coats are used, since production and storage capacity may be limited. Rarely the production of BPPs is hampered by a lack of whole blood donations.

To guarantee safe transfusions, the quality of the whole blood and the processed blood products are tested thoroughly. The whole process of collecting whole blood, processing BPPs and testing their quality takes a full day. The BPPs produced on day 0 are released at the start of day 1, say in the morning, and only then they become available for meeting the uncertain demand at the hospitals. To keep the overall outdating low, BPPs are primarily kept in stock at the blood banks and at a few large hospitals. Delivery of platelet pools happens daily in the morning and during the day upon request of the hospitals. In the Dutch case, each of the four blood banks serves about 20 to 30 hospitals.

In 2004, hospitals paid 458 euros for each BPPs they order at the blood banks, this includes all costs from collection, production, transportation and costs related organizational overhead. This price is set by the association of blood banks, Sanquin, in agreement with the Dutch government. The variable production costs are estimated at 100 to 150 euro. All other blood products are significantly less costly to produce.

Next to being the most expensive blood product, a BPP is also the most perishable blood product. Therefore blood bank managers are interested in a careful production-inventory control of BPPs. In 2004, the maximal shelf life of a platelet pool was only 5 days at most blood banks, counted from the day the pool is released and added to stock (one day after production). Current developments allow raising this maximum to 7 days. All other blood products can be kept in stock for weeks or even months.

Optimization problem

A blood bank's major concern is to produce safe blood products and to distribute these products to hospitals. Both in hospitals and at blood banks efficiency and quality of service get more and more attention. To improve efficiency, one should meet the demand at lowest costs. All costs that we consider in this study are (variable) costs related to the production and inventory management of BPPs.

For keeping a high quality of service, one should have enough pools in stock to meet the demand. Since the quality of BPPs deteriorates in time, the age of the pools upon meeting the demand influences the perceived quality of service. Owing to the demand uncertainty, a blood bank sometimes has to accept low stock levels every now and then. Occasionally a blood bank may even run out of stock. Then demand has to be met from another blood bank's inventory, e.g. by lateral transhipment, or it has to be produced instantaneously (through apheresis). A demanded BPP that cannot be issued from own stock is called a shortage. The variable shortage costs are related mainly to the costs of buying and transporting a BPP from another blood bank and to the loss of goodwill.

When demand is split into two age categories, demand for 'young' BPPs and demand for BPPs of 'any' age, the quality of service depends on the degree to which age preferences are met. Since the quality of a pool is (negatively) correlated with its age, issuing an older BPP than preferred is considered as a mismatch. A mismatch is thus a special type of shortage, namely a shortage of 'young' BPPs. The costs of a mismatch is considerably lower than the shortage costs, but to keep a high quality of service one may account for the risk of losing goodwill.

The occurrence of shortages and mismatches are penalized by introducing (fictitious) costs per BPP mismatched or short. For BPPs that become outdated, we have to accept the loss of labor hours and capital related to the production of the outdated BPPs. These costs are reflected in the so-called outdating costs. Further, minor variable holding costs are accounted for holding BPPs in stock.

From an Operations Research perspective, we are interested in an optimal (stock-age-dependent) ordering strategy, given the variable outdating, holding, shortage and mismatch costs, such that the *long-run average weekly cost* is minimized. Managers of blood banks and hospitals are interested in a practical rule for deriving nearly optimal (production) order sizes of BPPs. These two questions will be answered in this first part of the thesis.

2.1.2 A first modeling step: single-product model

Although a BPP can be of any of the eight blood groups, according to the ABO-system and the RhD- factor, the problem is often studied for a single, say universal, blood group. In the current practice, one produces-to-stock primarily BPPs of blood groups O and A, which fulfill virtually all demand. When BPPs of a specific group are needed, the blood banks commonly receives such a request well before the scheduled transfusion date.

Many perishable inventory studies concentrate on so-called single-product models. Although a great part of the applications concern blood products, only a few studies acknowledge the distinction of blood groups. Given the compatibility scheme, a single-product approach, in which one focusses on BPPs of a single blood group, seems to be a reasonable approximation of the current practice. However, a quantitative justification of such an approximation is not provided in the literature.

2.2 Common issuing and ordering policies

The performance of any perishable inventory system, in terms of outdating and shortages, strongly depends on how the issuing and ordering of products is controlled. The ordering policy prescribes *when* and *how-many* to order at the production department or at an external supplier. The issuing policy prescribes which items are taken from stock upon meeting the demand. To streamline the literature survey in Sections 2.3 and 2.4, we first give an overview of commonly used issuing and ordering policies.

2.2.1 Issuing policies

In the PPP the issuing of BPPs is controlled by the inventory manager. In a different perishable inventory setting, such as at a supermarket, customers may pick themselves the products from the shelf. The order in which products are taken from stock is then called the picking or dispatching order, which may vary between customers: some take an arbitrary product, while others select the freshest product on the shelf. Since our focus in the next chapters is on the PPP, we stick to the term issuing rule.

Inventory managers have some freedom in acknowledging any age-preferences communicated by the doctors in the hospitals. Nevertheless in practice, one tends to issue the oldest products in stock first. Issuing the oldest products in stock first is favored by inventory managers since it keeps outdating at a low level. This issuing policy is referred to as FIFO (First In First Out). Selecting the youngest products in stock first is called a LIFO (Last In First Out) issuing policy. LIFO issuing may contribute to a high quality of service, but it usually results in excessive outdating.

Similar to FIFO is FEFO (First Expired First Out), by which the products that will expire first are issued first. FEFO requires thus a ranking of all products in stock from 'fresh' to 'less-fresh'. Such a ranking is not available for BPPs, but can be approximated by looking at the age, or the residual shelf life of, BPPs. Regarding the 'freshness' of a BPP, the quality of a BPP is measured by the number of active platelets in a pool; this number can be estimated only by laboratory test, using a sample of the platelet concentrate. Since such an evaluation is laborious and expensive, FEFO has not been considered in a blood banking environment. Of course, the quality of a BPP is checked before a BPP is transfused, not only regarding the freshness but also to guarantee a safe transfusion.

Furthermore, demand may be categorized, and each category may prefer a different issuing policy. In the PPP we distinct two demand categories, since for some patients one may more strongly prefer to transfuse them with young platelets than for other patients. In Section 3.3.1, we present a new issuing policy for one of the demand categories, which we call FIFOR and which is a mixture of FIFO and LIFO. Our FIFOR policy is closely related to the upward and downward substitution rule as presented by Deniz, Scheller-Wolf and Karaesmen [35]. More on these and other less commonly studied issuing rules can be read in Section 2.4.1.

2.2.2 Ordering policies

Stock-age-dependent policies versus stock-level-dependent policies

Given a (composite) issuing policy, one needs to specify when replenishment orders are placed and the number of products to order. An optimal ordering strategy takes all relevant information that is available into account. For the optimal ordering of perishables with a fixed shelf life, one should know the number of products of each age category, rather then solely the total number of products in stock. The optimal ordering strategy for perishables is a so-called *stock-age-dependent policy*: it has to take into account the ages of the products in stock.

Since stock-age-dependent policies are more complex to analyze, given the detailed description of the inventory, most studies are on so-called *stock-level-dependent policies*. A stock-level-dependent policy takes as input only the total number of products in stock, irrespective their ages. One may say that the perishable inventory problem is often approximated by the inventory problem of a non-perishable product.

Periodic versus continuous review models

Existing (stock-level-dependent) ordering policies are usually classified as either continuous or periodic review models. In continuous review models orders can be placed at any point in time: typical ordering moments are when the stock level changes, i.e. when demand is met or when outdated products are removed from stock. In periodic review models order may be placed periodically at fixed moments, say at the start of every period. (In some multi-item ordering models an order for item i may be placed every R_i-th period, but for single-product models, R_i is often equal to 1 and therefore not included in the model.)

In the PPP, the orders are set periodically at the start of each working day. In periodic review models one may formulate a dynamic program and study the functional equations. Then often demand is modeled as a continuous distributed variable, since the analysis would become too complicated for discrete distributed variables. Another approach to study the ordering problem, is by considering a continuous review model instead of periodic review.

Some well-studied ordering rules

To streamline the discussion of the relevant literature in the next sections, we introduce at this point some common ordering policies and notations. For a more detailed overview we refer to the standard text book on inventory management, such as [62], [130], and [8]. As the ordering policies in periodic review models and in continuous review models are quite similar, except for the moments of ordering, we discuss them together.

In the models in which the policies below arise, one commonly incurs holding costs, ordering costs, shortage (or backlogging) costs, and outdating costs. These costs are often taken proportional to, respectively, the number of products in stock, the order quantity, the number of products short (or backlogged), and the number of products that become outdated. In addition fixed order costs may apply per order. When the fixed ordering costs are positive, one may decide to order only when the stock levels are 'too low'.

Ordered products may arrive instantaneously, or after a fixed lead time, or after a stochastic lead time. When the lead time is positive, the order volume often depends on the *inventory position* rather than on the number of products that are actually in stock. The inventory position is the actual inventory level plus the stock-on-order minus all demand that is backlogged. When the lead time is zero, the stock position equals the stock level, which we denote by $x = \sum_{r=1}^{m} x_r$. The same holds in periodic review models when the lead time is one period and the products are added to stock before a next order is placed.

Order-up-to S policy – When the fixed ordering costs are zero, the order quantity may be set such that the inventory position is raised to a fixed level S. The order quantity is set to S minus the stock position. In case the products are added to stock instantaneously, the order-up-to level S may be called the *target inventory level*. When the lead time is positive, the target level may not be reached when the order products are added to stock, as demand may have decreased the stock levels. We refer to this policy as an *Order-up-to S policy*. Other commonly used names for the same policy are *Critical-number policy*, or *Total-inventory-to-S (TIS) policy*.

(s, S)**-policy** – When fixed ordering costs apply, a threshold s for ordering is included. Orders are placed only when the stock position is at s or below. Most commonly, the order size is then derived from a fixed order-up-to level S. When the stock position is above s, no order is placed. This policy is commonly called an (s, S) *policy*.

$(S-1, S)$ **policy** – A special case of the (s, S) policy is when $s = S-1$, this policy may be appropriate when the demand is for discrete products. For the ordering of perishable this so-called $(S-1, S)$ policy, is studied by [126] and [112] for cases with positive (stochastic) lead time and demands may be in multiple products at a time.

(s, nQ) **policy** – A last class of ordering policies that we discuss here, is the (s, nQ) policy: every ordering moment the inventory position is checked and when it is at or below s a quantity $n \cdot Q$ is ordered with $n \in \{1, 2, \ldots\}$ such that the new stock position falls in the interval $(s, s + Q]$. Q is the batch size for ordering, which may be set by the supplier: e.g. order per box or per pallet.

2.3 History of production-inventory models

Inventory control is a classical subject of study to Operations Researchers and practitioners. Practitioners are interested in preferably simple policies or tools that help them balancing their quality of service and the efficiency of their operations. Operations Researchers study inventory problems from a mathematical point of view and provide models and simple rules to help the practitioners running their business effectively and efficiently. Before presenting in the next section a detailed literature survey, we give an overview of the history of inventory theory.

The focus is on articles on the ordering of a single product in a single echelon of the supply chain. Endpoint of the discussion is the numerical computation of stock-age-dependent ordering policies. The literature is much more rich, but multi-product and multi-echelon problems are beyond the scope of this thesis.

The timeline in Figure 2.3 shows the evolution of the (scientific) research in inventory management. The timeline starts with the first mathematical models for non-perishable inventory and shows how the problem of interest is shifted and extended over time:

- from deterministic problems in the early stage,
- to stochastic problems in the fifties,
- to perishable inventory theory in the seventies,
- to the approach presented in this thesis.

Deterministic models – EOQ, Camp's formula, and ELS

In 1913, Harris ([59], [60]) introduced one of the first mathematical inventory models: the economic order quantity (EOQ). According to Erlenkotter [43], the EOQ formula became since 1922 also known as Camp's formula [25] and since 1934 as Wilson's economic lot size formula (ELS) [166]. The EOQ formula sets the optimal order quantity for a non-perishable product in a deterministic problem setting, based on a linear cost structure with both fixed and variable cost components. In many text books, e.g. [62], [131], the basic model and a number of extensions are explained in detail.

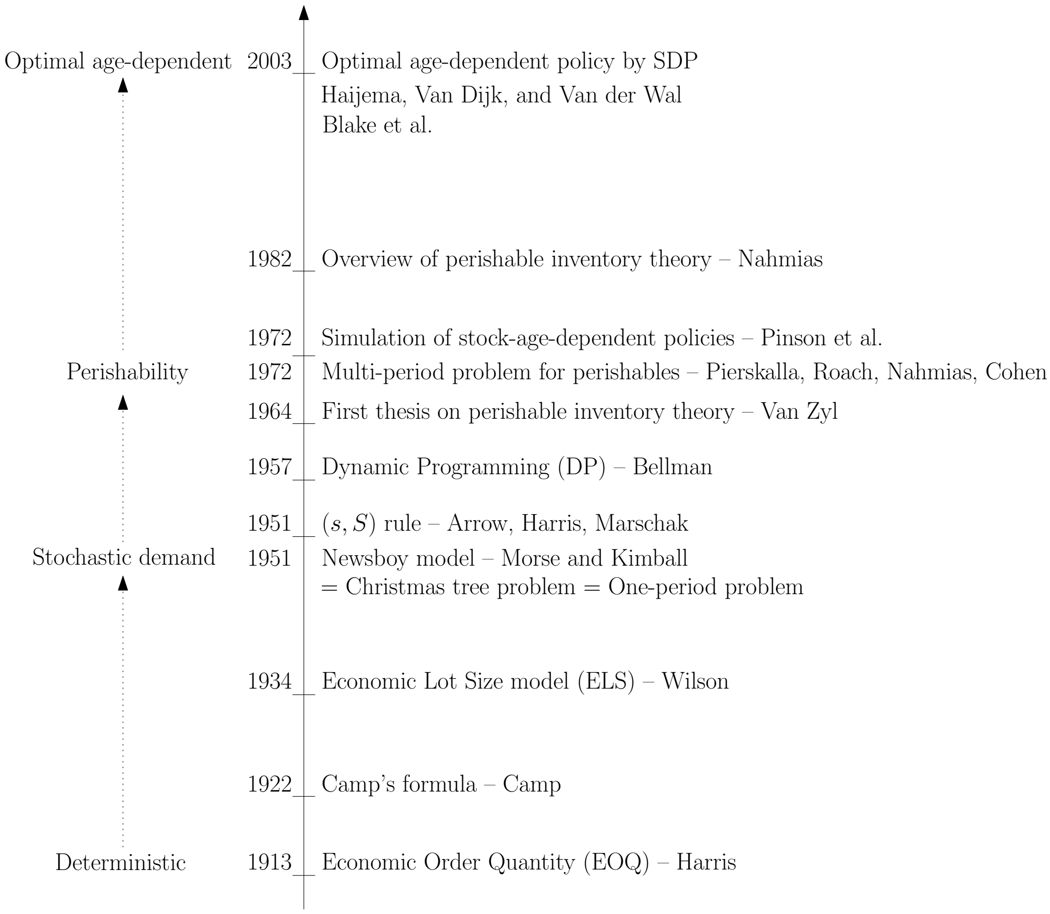

Figure 2.3: The history of inventory management: from simple deterministic models to stochastic perishable inventory models.

Stochastic single-period problems: Newsboy model, Christmas-tree model

One of the first articles on a stochastic inventory model dates back to the early fifties of the last century. In 1951, Morse and Kimball [95] present a one-period problem with stochastic demand. The model is often called the Newsboy model and is solved by the so-called Newsboy equation. Later it became also known as the Christmas tree model.

Kraiselburd [78] describes the very interesting history of the Newsboy problem. In fact, the model comes from classified research for the Navy during World War II. Many of the references to the work of Morse and Kimball recall military applications (mainly inventory problems) rather than the problem of a newsboy deciding on the number of newspapers to buy. Over the years the model is extended and applied to various one-period decision problems. The Newsboy model can be seen as a first stochastic inventory model for perishables with a fixed shelf life of one period.

Stochastic multi-period problems: (s, S) policies and Bellman's principle of optimality

Soon after the introduction of the Newsboy problem, the sequential ordering of a non-perishable products over multiple products was studied. The multi-period problems appeared to be significantly more complicated when fixed order costs apply. In 1951, Arrow, Harris and Marschak presented in Econometrica [7] the so-called (s, S)-type inventory control with stochastic demand. In 1952 and 1953, Dvoretzky, Kiefer and Wolfowitz ([40], [41]) studied the optimal parameter values of the (s, S) strategy. Most of the studies on multi-period problems rely on solving functional equations. The proof that the optimal policy for non-perishables is of the (s, S) type was first given in 1960 by Scarf [125], for a wide class of problems with a linear cost structure.

The study of the optimal control of inventories boosted after the introduction of Bellman's principle of optimality in the fifties. According to Arrow [5], he generalized some previous work on sequential analysis and announced a new technique called Dynamic Programming (DP). The recursive nature of solving a sequential decision problem was already addressed by Arrow, Blackwell and Girshick [6] in 1949. Anyway, DP became a new tool of finding optimal strategies for various decision problems ([14], [68]).

In 1960, Howard introduced a special class of sequential decision problems called Markov decision problems (MDP). In particular, Howard and Bellman deserve the credits for the successful introduction of this field of study. Nevertheless, due to the dimensionality of the state space, the multi-period perishable inventory problem was still considered too complicated to solve.

(Multi-period) Stochastic Perishable inventory models

Virtually all inventory studies before 1960 dealt with non-perishables or with perishables that last for one period only. The inventory management of perishables over multiple periods is considerably more complicated, since products are not only removed from stock to meet the (stochastic) demand but also because they become outdated. The study of inventory management of perishables started already in the sixties with the thesis of Van Zyl [154] and articles Pierskalla and Pierskalla and Roach. In the seventies the research of perishable inventory models flourished and resulted in papers by Cohen, Nahmias and Pierskalla and many others ([113], [114], [101], [29], [30], [97], and [98]). Most studies are on issuing policies and approximations of the outdating and shortages under specific ordering policies, or on finding good parameter values for order-up-to S and for (s, S) policies.

In some numerical studies, in which outdating and shortage figures are low, one claims simple rules to be nearly optimal. The question to what extend these rules are optimal in other cases remained unanswered. A number of features, that may be relevant in real-life perishable inventory problems, were not included in these studies.

Moreover optimal stock-age-dependent policies were hardly studied, in these (early) years. Two exceptions are the studies by Fries [47] and Nahmias [98]: in which some properties of the optimal solution are proved. Numerical approaches to compute optimal strategies for realistic sized problems were doomed to fail at that time. Solving the underlying MDP is seriously hampered by insufficient computer power and the curse of dimensionality in the state space: the number of possible states is too large to solve problems of real size. Only for case with a very short maximal shelf life and zero lead time some results are reported as will be discussed in Section 2.4.2.

Simulation of perishable inventory systems

In the seventies computer simulation rapidly became a more-and-more useful tool to analyze complex dynamic systems. Where the analytical models need to be relatively simple to allow its mathematical analysis, problems could now be studied at a reasonable realistic level by simulation. This different, say more practical, line of research was carried out by quite some researchers and practitioners. Simple order-up-to S rules were simulated to find out whether they yield low shortage and outdating figures under different issuing policies. For simple rules with one or two parameters search algorithms were developed to find optimal parameter values. Analytically obtained insights on the convexity of the cost function proved to be very useful in streamlining the optimization process.

Despite the modeling flexibility of simulation, there is always a tradeoff to make between the degree of realism to add to a model and the required accuracy of the outcomes, given the speed of the computer(s) on which the simulation is executed. This trade-off is inherent to simulation. Since computers were not as fast as nowadays, the level of detail in mimicking a real inventory systems was quite limited; in the early simulation studies the existence of different, limited-compatible blood groups is simply left-out. Most likely, this also explains why no periodic ordering policies are simulated with different order-up-to levels $S_1, \ldots, S_D$ of each of the D periods: i.e. the search for optimal values of all D parameters could have been too time-consuming.

In only a few studies stock-age-dependent ordering rules were developed that acknowledge the perishability of the products. For example, in 1972, Pinson, Pierskalla and Schaefer [115] test by simulation what-they-call *modified* order-up-to S rules. These rules were developed based on common sense arguments to reduce shortages and outdating. Setting new stock-age-dependent rules requires insight into the perishable inventory problem. Simulation is one way to obtain these insights.

Optimal stock-age-dependent ordering policies for perishables

Although it was known that an optimal policy for ordering perishables should be stock-age-dependent, the actual computation of an optimal policy was not considered for perishables with a fixed shelf life of more than 3 days. At the start of the twenty-first century Blake et al. [19], and Haijema, Van der Wal, Van Dijk and Smit Sibinga ([56], [57], [58], [149]) numerically solved realistic (down-sized) Markov Decision Problems of blood platelets pools (BPPs) with a maximal shelf life of up to 7 periods. They apply independently an aggregation-disaggregation approach: individual BPPs are aggregated into batches. Blake et al. focusses on the cost reduction and the tractability at varying batch sizes. Haijema et al. study the structure of the optimal policy (with one and two categories of demand) and derive simple rules with nearly optimal parameter values. Their approach and results are reported in the first part of this thesis.

2.4 Classification of perishable inventory studies

In this section, we classify and review the perishable inventory studies that are most relevant in the light of solving the PPP.

Careful inventory management of perishables is of general interest and has had considerable attention over the past decades. A number of overviews and surveys are published: e.g. in 1984 by Prastacos [118], in 1991 by Raafat, and in 2001 by Goyal and Giri [51]. The most detailed and most cited review is that of Nahmias [100] in 1982. He reviewed almost one hundred articles on the theory of perishable inventory. A great part of the literature in this field is devoted to blood products.

In this section, we present a detailed review of the most relevant papers on inventory management of perishables with a fixed maximal shelf life. Further we limit the discussion to studies in which demand is stochastic. Since our approach – to be presented in the next chapters – is a methodological one, we organize our detailed review by classifying the studies according to the different methods and techniques that are in use. Some articles may fit into several categories.

The partly overlapping categories are:

- **Analytical studies**
  In analytical studies one primarily makes use of analytical techniques (such as calculus, algebra, functional equations, Markov chain analysis and sample path arguments). Closed-form expressions for the number of shortages and outdated products or for the cost function are only obtained for special, trivial cases of perishable inventory problems. In many analytical studies, one relies on numerical techniques in the end, i.e. for the evaluation of complex (multiple) integrals, for root finding and for estimating gradients of complicated analytical expressions. Nevertheless, the analytical studies may provide insights in how to model and approximate perishable inventory problems. Moreover, analytical techniques do not require data: with respect to this point they may be more general, although they less flexible in modeling the problem.

- **Numerical studies**
  In numerical studies numerical or computational procedures are developed and/or used to evaluate or to optimize the inventory control of perishables. By nature, numerical studies are data-driven and suitable for sensitivity analysis. A special class of numerical studies is the class of simulation studies.

- **Simulation studies**
  In simulation studies one mimics the system under some (fixed) policy. Computer simulation is a fast way of doing this and provides insight in the performance of that policy. In combination with a search algorithm, simulation can be used for optimization. It is important to note that simulation itself is a non-optimizing technique.

- **Experimental or empirical studies**
  Experimental or empirical studies report on current practice and on testing or implementing new policies in practice. Computer simulation experiments do not fall in this category. Since experimenting and empirical testing in practice is expensive and time-consuming, most studies do not belong to this class.

In the next subsections we discuss the most relevant papers that compose these four classes of studies.

2.4.1 Analytical studies

Analytical models are parameterized models and thus free of (case-sensitive) data. Often they are used to attain insights into properties of a control rule. For perishable inventory problems closed-form expressions for the number of shortages and outdating can only be derived for simplified models under specific well-structured ordering, such as the (s, S) policy. In most perishable inventory models, one studies the performance of ordering policies that stem from studies of non-perishables. Therefore, we start this review of analytical models with a brief description of the main results derived for non-perishable products. Next, we discuss some studies on perishable inventory management for both single-period and multi-period problems. Since the PPP is a multi-period problem, we pay particular attention to this class of problems.

Ordering policies for non-perishable products

In the early fifties, Arrow, Harris and Marschak [7] and Dvoretzky, Kiefer and Wolfowitz ([40], [41]) studied the (s, S) ordering policy, under a linear cost structure with holding costs, shortage costs and fixed order costs, in the context of stationary continuous (unknown) demand for non-perishable products. Since for the inventory control of non-perishable products, the age of the products is irrelevant, the stock is described by a single variable: the total stock level x.

In 1960, Scarf [125] showed, for the stationary (in)finite horizon problem with stochastic demand, that an (s, S) policy is optimal for ordering non-perishable products when fixed order cost and variable holding and shortage (or backlogging) costs apply. When the fixed order costs are zero an order-up-to S policy is optimal.

Analytical models for perishables

One may say that the study of ordering policies for perishables started with the so-called one-period problem. Soon thereafter the sequential ordering was studied for the so-called multi-period problems.

One-period problem

The one-period problem, often named the *'Newsboy' problem*, was first studied in the forties and fifties by Morse and Kimball [95], as already discussed in Section 2.3. Although the PPP is a multi-period problem, the Newsboy model deserves some attention in this review. In its simplest form the model deals with a newsboy who faces variable order costs c per product and a selling price p per newspaper sold. At the start of a period he decides on the number of newspapers to order, given the uincertainty in the demand and the fact that the unsold newspapers are worthless at the end of the day. Ordering a products costs $c \cdot a$, but yields p times the expected number of sold products.

The number of products sold depends on the demand distribution and is limited by a. Let $P[D = j]$ denote the probability that the demand over a single period is for j newspapers. The optimal order quantity is then the greatest value of a for which holds

$$\sum_{j=0}^{a-1} P[D = j] \leq 1 - \frac{c}{p}.$$

Clearly, the scope of the newsboy model is not limited to the ordering of newspaper, but is also applicable to the ordering of perishable and non-perishable products and services [78].

Multi-period problems for perishables

First, we discuss the properties of an optimal ordering policy for a perishable with fixed maximal shelf life. Next, an overview of commonly adopted, simplifying assumptions is given. Finally we report on most relevant studies on the ordering policy, on the ageing and outdating process, and on the issuing policy.

Structure of an optimal ordering policy for perishables

Multi-period problems for perishables with a fixed shelf life are much more complicated than one-period problems or order problems for non-perishables. Whereas the (s, S) is optimal for a broad class of non-perishable inventory problems, such a policy is in general not optimal for ordering perishables over multiple periods. This was first shown in 1975 by Fries [47] and Nahmias [98]. Both authors analyze independently the functional equations for a dynamic program of the fixed life perishable inventory problem. They show that an order-up-to S policy nor an (s, S) policy is optimal for the general case with maximal shelf life m days. Although the optimal order point might depend on a threshold that is stock-level-dependent, the order quantity should take into account the age distribution of the stock on hand.

In their studies they assume stationary, continuous, independent and identically distributed (i.i.d.) demand, which is met in a FIFO order. The following cost components are included: linear order costs and convex holding, shortage and outdating costs. To limit the dimension of the dynamic program the lead time is assumed to be zero periods. According to Nahmias, the work of Fries is a bit more complete compared to that of Nahmias. Fries obtained results for both finite and infinite horizon models with discrete demand.

Both authors agree upon the computational complexity due to the dimensionality of the dynamic program. Fries states that the direct computation of optimal policies generates further insight into the form of truly optimal policies. Nahmias concludes that future research should be on various simple non-optimal policies, next to the determination of approximate computational methods.

Simplifying assumptions

Given that an optimal ordering policy depends on the ages of the products in stock, its analysis is in general too complicated for an analytical approach. Therefore analytical studies are limited to classes of well-structured strategies, e.g. (s, S) strategies, and additional assumptions are imposed to model the ageing and outdating processes of products in stock. In most studies stock-level-dependent strategies are analyzed, only a few studies consider well-structured stock-age-dependent policies.

Although analytical models are quite general, since they are free of data, the models can be treated only under simplifying assumptions. Commonly adopted assumptions concern:

1. The ordering policy:
   - Most analytical studies concern the evaluation of stock-level-dependent policies that are well structured, e.g. order-up-to S or (s, S) policies, which are proven to be optimal for non-perishables.
   - In a few studies one reports on stock-age-dependent policies that exhibit a special structure, which allows the mathematical analysis under strict assumptions, e.g. Tekin et al. [141].
   - A truly optimal stock-age-dependent policy for a general multi-period problem under stochastic demand is not achieved through analytical methods.
2. The maximal shelf life and the planning horizon:
   - If the (planning) horizon is shorter than or equal to the maximal shelf life, the analysis is simple when the initial stock is zero, since outdating only happens then at or after the horizon, see the work in the early sixties by Hadley and Whitin [52], [53] and Brown et al. [21].
   - A special case is the one-period problem or the Newsboy problem.
   - When the maximal shelf life is at most two (or three) periods, the problem can be formulated as a Markov chain with a low dimensional state space, see Van Zyl [154], Bulinskaya [24], and Nahmias and Pierskalla [101].
3. The ageing and outdating process:
   - The outdating process is in some studies approximated by assuming that outdating is a fixed or unknown fraction of the total stock on hand. When the shelf life is exponentially distributed, the decay of a product is a memoryless process (see Ghare and Shrader [49] and Emmons [42]). This way one avoids a high dimensional state space in a continuous-review model.

- In some other studies, one explicitly assumes that all products in stock have the same expiration date. This assumption may hold when replenishment only arrive when no products are left in stock. Alternatively, products keep a constant quality as long as they are packed in a sealed box, but the deterioration of products start when the box gets unsealed. When boxes are opened only when all old products are removed from stock, a model with only one age category suits ([85], [141]).

4. Other aspects: e.g. lead time and backlogging:

- In many analytical studies one assumes zero lead time, which implies that replenishment happens instantaneously. The state vector describing the number of products in stock of each age category at the start of a period is then of dimension $m - 1$. When the lead time is exactly 1 period, the state vector consists of m dimensions. When the lead time exceeds k periods, one needs to include in the state description the outstanding orders placed in the last for $k - 1$ periods. Assuming zero lead time thus simplifies the analysis. (The PPP is thus a bit more difficult at this point.)

- Introducing positive lead time makes the analysis much more difficult, since next to the age distribution also outstanding orders need to be included in the state description. When demand occurs one-by-one according to a Poisson process, and 1 product is ordered at a time, according to an $(S - 1, S)$ policy, exact results for positive lead-time can be derived, see Schmidt and Nahmias [126] and Perry and Posner [112].

- Under some policies the analysis becomes easier when shortages are backlogged. However, in the PPP backlogging is not allowed as meeting the demand is critical as a patient's live could be at risk. In the PPP shortages are resolved by sending BPPs from another blood bank.

The review of analytic multi-period problems is divided into three parts. First, we report on studies with a focus on optimal ordering policies (mostly under a FIFO issuing rule). Next, a few articles on the ageing and outdating process are discussed. Finally, analytical studies on alternative issuing policies are reviewed.

Analytical studies on optimal ordering policies

Most analytical studies are on the (s, S) policy under continuous review. Where the analysis of inventory systems with zero lead time is already complicated, this certainly is the case if the lead time is positive. For a special well-structured strategy, the $(S-1, S)$ policy, and under the assumption of demand to happen by a Poisson process, Schmidt and Nahmias [126] and Perry and Posner [112] have obtained some exact results for the special class of $(S-1, S)$ policies with positive lead time. Under an $(S-1, S)$ policy an order is placed after each occurrence of demand or when products become outdated. The inventory position, which includes any outstanding orders, is thus kept at level S.

Whereas most of the analytical studies for non-perishables could be formulated as continuous review problems, a discrete time model is more natural for the control of perishables with a fixed maximal shelf life. Lian and Liu [85] presented in 1999 a discrete time model for (s, S) policies. All demand that cannot be met directly from stock is backlogged and new stock arrives instantaneously. Under these assumptions they study the (s, S) policy under periodic review with $s \leq -1$: a new order is thus placed only when no items are left in stock is. Consequently all products in stock are of the same age. This property simplifies the analysis significantly, hence a closed-form expression of the cost function can be obtained, which includes fixed order costs, holding costs, shortage costs and outdating costs. Since the properties of the resulting cost function are difficult to study analytically, they conduct a numerical study to investigate the impact on the costs of, what they call, the backlog level $s < 0$ and the order-up to level S.

Also Tekin et al. in [141] adopt the assumption that all products in stock are of the same age category. They compare a (continuous-time) modified lot-size-reorder policy with three parameters (s, Q, T) with an ordinary (s, Q)-policy. Under the (s, Q, T) policy, one orders Q units whenever the total stock level drops below s or when T periods have elapsed since the last time the inventory level was Q. The decay of the batch of Q products starts upon opening the batch for consumption. A batch is opened only when all products of the previous batch are either consumed or removed from stock because they have become outdated. Parameter T makes the rule stock-age-dependent. For the (s, Q, T) policy Tekin et al. derive analytically a cost function by applying renewal theory with the stock level Q as a renewal point. The cost components that they include are fixed order costs and variable holding and outdating cost. Instead of the usual direct costs for lost sales, they impose a service level constraint: the fraction of sales that is lost must be less than some prespecified value. Investigating the properties of the cost function requires an extensive numerical study from which they conclude that the stock-age-dependent (s, Q, T) rule indeed performs better than the (s, Q) policy, especially in cases where outdating is an issue.

In 2002, Zhou and Pierskalla [171] analyze an inventory system with regular orders and emergency orders. Regular orders are less expensive than emergency orders, but have a much longer lead time. At the beginning of a cycle a regular order of Q products is placed. Whenever the inventory drops below s, an emergency order is placed to return the inventory to level s.

Studying the ageing and outdating process

A few perishable inventory studies focus on the ageing and outdating process of perishables.

In 1970, Pegels and Jelmert [111] track a single unit of blood over its maximal shelf life of 20 days. By a Markov chain approach with two absorbing states (expiration and transfusion of the unit), they derive the probability of outdating and the age distribution upon transfusion. The difference between FIFO and LIFO issuing is reflected in the transition probability: at FIFO the probability of transfusing a unit is higher when the unit is older. Through numerical evaluation of the Markov chains with a given transition matrix, the conditional probability of expiration of a unit is computed for each age category as well as the expected days until transfusion. The usefulness of this paper is however limited, as argued by Jennings and Kolesar in [71]. Most critical is Pegels and Jelmert's assumption that the transition probabilities are known. These probabilities should dependent on the demand distribution and the issuing and replenishment policy, but in their paper these components are not modeled explicitly. Even when the transition probabilities are as observed in practice, the contribution of the paper is marginal since the Markov chain analysis provides a seemingly exact way of computing of what is already known from the observations.

Another Markov chain approach that lacks an explicit controlled ordering policy is by Bar-Lev and Perry [9] in 1989. They analyze a discrete time Markov chain of stock levels for a perishable product (having a fixed maximal shelf life). The Markov chain analysis requires the assumption of stochastic supply instead of an explicit policy for ordering products and is therefore beyond the scope of this thesis.

When the replenishment is controlled rather than modeled as a stochastic process, one imposes a well-structured ordering policy and zero lead time to enable an analytical approach. Cohen [27] and Chazan and Gal [26] investigate simultaneously, yet independently, the age distribution of the stock-on-hand under an order-up-to S rule. Whereas Cohen examines the continuous demand case, Chazan and Gal consider the discrete demand case and prove that the expected cost function as well as the expected outdating are convex in S. In addition, Chazan and Gal develop bounds on the expected outdating and generate closed-form solutions for expected outdates for an approximate model with Poisson demands. Usual assumptions in these analytical studies are i.i.d. demand that is met by FIFO issuing and that replenishments arrive instantaneously.

Analytical studies on the issuing policies

The most common issuing and ordering policies are reported in Section 2.2. Often the issuing policy can be kept very simple, since there is only a single category of demand. In case of multiple demand categories, as in the PPP, the issuing policy might be more complicated, e.g. a mixture between FIFO and LIFO.

In 1972, Pierskalla and Roach [114] consider a problem with multiple demand categories that are related to the different age categories. The demand for products of some specific age can be met by products of that age or by younger products, but not by older products. They show that issuing the oldest products available that meet the age preference, i.e. FIFO, is optimal under most common cost structures. For that reason one imposes a FIFO issuing rule in most studies. One of the few exceptions is Cohen and Pekelman [28] who study the less commonly studied LIFO issuing policy.

More recently, in 2004, Deniz, Scheller-Wolf and Karaesmen [35] investigate two types of order-up-to S rules under different issuing policies to meet demand for products of different age categories. The products of different ages are to some degree substitutes to each other, as reflected in the issuing rules, which they call substitution rules. Upward substitution means that the demand can be met in a FIFO way starting with the products of the preferred age category. Downward substitution implies that, when not enough products from the preferred category are available, older products are issued in a LIFO way. Full substitution is the combination of upward and downward substitution: demand for old products is met with products of that age or older, demand for 'young' is met by products of that age or younger. Meeting demand by products of another age category implies substitution costs for each demand category. Two order-up-to S rule are considered: one rule takes the total stock level as input, the other rule looks only at the number of 'new' products in stock. Using sample path arguments they discuss sufficient conditions on substitution costs that ensure dominance relations between the substitution policies.

Another interesting study is by Kopach et al. [76] in 2006. They model the perishable inventory problem as a queueing model with two types of customers: urgent versus non-urgent. Non-urgent customers can be blocked in case of low stock levels to reduce the chance of falling short to the urgent customers. In their queuing model the arrival stream relates to the supply of products, the queue length represents the total stock level, the outdating of products relates to impatience customers leaving the queue, and finally the demand for products is modeled by the server that picks customers from the front of the queue. The first come first served (FCFS) queueing discipline thus relates to the FIFO issuing policy. The mean service time represents the mean inter-demand time. The control problem they face is switching from a high speed server (that accepts both types of demand) to a slow speed server (that accepts only urgent demand) and vice versa, such that one minimizes shortages occurring when the queue is exhausted (out of stock) and outdating (represented by the loss of customers). Since this problem is quite different than ours, we do not report their work here in more detail.

Remark – In a number of studies, not reported in this section, both analytical and numerical techniques are adopted. One often reverts to numerical techniques to evaluate analytically obtained expression for the performance measures and the cost function. In particular, optimal parameter values can only be found through numerical experiments.

Conclusions

1. An optimal policy for perishables with a fixed maximal shelf life is stock-age-dependent.
2. In general, a stock-age-dependent policy is too complicated to study analytically.
3. Simplified analytic models are restricted to well-structured classes of policies, such as order-up-to S and (s, S) policies.
4. Nevertheless, analytic models provide insights into the ageing and outdating process and may approximate the cost function, but closed-form expression are hard to obtain.
5. For insights in the properties of the cost function, e.g. to determine optimal parameter values of an ordering policy, one often reverts to numerical techniques for evaluation.

2.4.2 Numerical studies

Shortly after the introduction of Dynamic Programming by Bellman, Dreyfus stated in 1957 the potential and the limitations of computational procedures for dynamic programming. Dreyfus [38] discussed how to deal with the computational limitations for deriving value functions. As one of the ways to overcome the computational difficulties, he explains the symbioses between sophistical mathematical analysis and digital computation. Mathematical analysis provides approximations and allows doing the computations more efficiently; it is not unusual for analysis to establish the nature of the function without being able to produce the function. In 1962, the book *'Applied Dynamic Programming'* [15] of Bellman and Dreyfus appeared. Other interesting titles are the books by Howard ([68], [69]), in which the focus on DP in the context of Markov models. In these early days, the computational resources were very limited. Nowadays the computational limitations are shifted but still present.

In 1999, Smith [134] presented an overview of the development of the relationship between inventory management and dynamic programming (DP). He concluded that DP has one disadvantage: its curse of dimensionality; nevertheless, DP helps developing and evaluating heuristics. He envisioned that numerical techniques can be very useful in solving complex inventory management problems, such as where the products are perishable.

Before presenting the few numerical studies on optimal ordering of perishables, we first discuss the acceptance and application of numerical procedures in a non-perishable inventory setting.

Non-perishable inventory: optimal parameter values

Although inventory management problems have acted as key examples to illustrate the dynamic programming principle, the application of dynamic programming to solve inventory management problems did not really flourish in practice during the first decades. In 1960 Sasieni [123] discussed a stochastic inventory problem of non-perishable products with fixed order costs next to variable holding costs. He thoroughly discusses the dynamic programming principle for minimizing the long-term average costs. Numerically he showed how optimal parameter values for an (s, S)-rule are computed through value iteration.

Johnson [72] presented in 1968 a policy iteration algorithm that exploits the special (s, S) structure; thus speeding up the search for optimal values of the parameters s and S. Most other methods for finding optimal parameter values are essentially based on full enumeration over the two-dimensional grid in the $(s, s+k)$ plane, the interested reader is referred to [156], [10] and [4].

Despite all efforts in developing numerical procedures, these algorithms were hardly used to optimize the parameters. Federgruen and Zipkin presented in [46] a similar algorithm and showed that optimal parameter values could be computed quicker than generally believed.

Further improvement is attained in 1991 by Zheng and Federgruen. In [169] and [170], they present an efficient search algorithm based on an analysis of the convexity of the seemingly ill behaved cost function under an (s, S)-rule. Although the cost function may have several local optima and generally fails to be quasi convex, they pretentiously state that optimizing an (s, S) policy is almost as easy as evaluating a single policy. Although this might be the case for non-perishable products, it may not be the case for perishable products. The analysis and computations are much more difficult when products have a fixed maximal shelf life.

Perishable inventory: optimal stock-age-dependent policy

Truly optimal policies are generally hard to compute, due to the dimensionality of the state space. The state description contains the number of products in stock of each age category. Nevertheless, some efforts were made by Nahmias [99] in 1977, Parlar [110] in 1985, Nandakumar and Morton [102] in 1993, and most recently, in 2003, by Blake et al. [19] and Haijema, Van der Wal and Van Dijk ([56], [57]).

Nahmias [99] and Nandakumar and Morton [102] consider a very simplistic Dynamic Programming (DP) formulation of the perishable inventory problem. The general perishable inventory problem is simplified such that the state space for the truly optimal MDP problem is two-dimensional: the maximal shelf life is set to (at most) three periods ($m = 3$) and products are delivered instantaneously (zero lead time). Several complicating aspects as the periodicity of the demand, production stop during weekends, the distinction of two types of demand and issuing policies other than FIFO are not considered. By successive approximation an optimal production volume is computed. By simulation they compare the truly optimal strategy with an approximately optimal stock-age-dependent ordering policy of lower dimensionality and with an approximate order-up-to S rule. For a more detailed discussion of their work see the next section on simulation studies and the references [99] and [102].

In 1985, Parlar [110] applies Linear Programming (LP) to solve a stationary MDP for setting an optimal ordering policy. He investigate the problems in a very simplistic setting with a maximal shelf life of only 2 periods to allow efficient computation of the truly optimal policy.

From 2003 onwards Blake et al. [19] and Haijema, Van der Wal and Van Dijk [56, 57, 58] show that the problem is nowadays tractable (via successive approximation) for a maximal shelf life of up to 7 periods, when the scale of the demand is not too high. When demand plays at a large scale, the problem is approximated by down-sizing for efficient computation. Blake et al. investigate the tractability and applicability of numerical techniques for solving an MDP formulation of the PPP. The PPP that they investigate under linear cost structure is solved over a finite horizon by assuming that production and demand happens in batches (of multiple pools). They show that reducing the batch size and lengthening the planning horizon results in lower costs but requires more computation time.

In Haijema, Van der Wal, and Van Dijk [56, 57, 58] the structure of numerically computed optimal strategies is investigated for several cases. For simple cases, it is shown that the optimal policy resembles a simple order-up-to S rule with fixed order-up-to levels for each day of the week. When operating at a small scale or when multiple demand groups are distinguished more-advanced replenishment rules fit better to the optimal strategy. Nevertheless they conclude that the outdating and shortage figures of simple order-up-to S rules (with levels read from the optimal policy) are nearly optimal for the realistic platelet production problem that they consider. As will be presented in this thesis the problem can be modeled at a more realistic level than before. From the solution of the dynamic program one learns new ordering rules that closely resemble the structure of the optimal policy.

Conclusions

1. Numerical solution procedures such as successive approximation and linear programming are very suitable for finding optimal stock-age-dependent policies,

2. Nevertheless these techniques are hardly used because of the so-called curse of dimensionality, which prevents the computation of optimal strategies for realistically large problems.

2.4.3 Simulation studies

In simulation studies of perishable inventory systems one mimics the system controlled by some fixed ordering and issuing policy to evaluate the performance in detail. Again blood banking applications and order-up-to S policies get much attention in the literature. A simulation approach allows to model a problem in much more detail than an analytical approach. Nevertheless one needs to carefully avoid over-specification of the simulation model, since obtaining accurate performance estimates from simulation becomes very time consuming when the model gets more and more detailed. A few optimal stock-age-dependent order-up-to S rules are investigated, but since simulation is by nature a non-optimizing procedure the truly optimal stock-age-dependent policy remains intractable.

Simulation is thus mostly used to evaluate a system to gain insights in search for improved policies. One may use a simulation program in combination with a (heuristic) search procedure to find good and hopefully optimal parameter values within a class of policies. Before reviewing the most relevant simulation studies, we briefly give a few references of simulation based search techniques.

Simulation-based optimization concerns the optimization of some objective (or cost) function over different parameter values of some control rule. Since there is in most cases no analytical, closed-form expression available for the cost function, one draws a response curve through points evaluated by simulation. When the curve is proven to be convex, or seems to be so, an efficient search algorithm can be constructed to find (nearly) optimal parameter value(s). An efficient search algorithm for finding an optimal order-up-to level of an order-up-to S rule is Fibonacci search, as described in [115], [29], [30], [97].

Another way for setting parameter values is by simple rules of thumb, such as "Set S equal to the expected demand over some period(s) eventual plus a fixed multiple of the standard deviation of the demand over some time interval".

When order-up-to levels should depend on the day of the week, as in periodic problems, the search for the best parameter values is more complicated. An exhaustive search would be very time-consuming given the number of scenarios to consider. In the next chapters we will present two approaches to deal with this difficulty.

Survey of simulation studies

We start our review with the work of Pinson, Pierskalla and Schaefer in 1972 [115] and that of Cohen and Pierskalla (in 1974-1979, [29], [30], [31]) in which they generalize the class of order-up-to S strategies. By adding weights to the stock level for each age category, they create, what we call, a *weighted* order-up-to S rule: the production volume a is set according to:

$$a = (S - \sum_{r=1}^{m} w_r x_r)^+, \quad \text{with constant weights } w_r \tag{2.1}$$

In case all w_r are equal to 1 then the rule is the order-up-to S rule. Otherwise, the ordering policy in Equation (2.1) is stock-age-dependent.

Pinson et al. and Cohen and Pierskalla assume in their studies increasing weights: 'young' products get higher weight than old products in aggregating the stock-on-hand. The production volume is thus lower when more young products are in stock. When many old products are in stock the production volume is relatively higher to anticipate outdating of old products.

In the simulation experiments of Pinson et al. [115] and of Cohen and Pierskalla ([29], [30], [31]) the focus is on the ordering of whole blood products that have a long maximal shelf life of 21 days ($m = 21$). Typical for the problem they investigate is that (issued) products may return to the inventory after a so-called cross matching. They use data obtained from several blood banks and hospitals in Illinois (USA). Under a FIFO issuing policy and a linear cost structure the impact of the weights (w_r) is studied. For different weights w_r, they draw response curves that visualize the impact of the order-up-to level S on the occurrence of shortages and outdating. From the seemingly convex curves a range of nearly optimal values S is read. Since the curve appears to be rather flat around the optimum, and the accuracy of the simulation is limited, multiple values of S seem to be optimal.

Furthermore, Cohen and Pierskalla study the impact of the issuing policy (LIFO versus FIFO). The products they investigate have a relatively long shelf life and outdating is thus a non-issue when using an order-up-to S rule in combination with FIFO issuing. Weighting the age categories does thus not make much sense.

In 1975 Nahmias [97] reasons that production volumes for a general perishable product should be set lower than in the non-perishable case. He investigates three *modified* order-up-to S rules:

1. production is a fixed fraction of the production volume set by the optimal order-up-to level of the non-perishable case (S^{np}):
   $a = \beta \cdot (S^{\text{np}} - x)^+$

2. production is derived from the order-up-to level S^{np} for the non-perishable case, but is bounded by an upper bound $\overline{a}$ to prevent excessive production:
   $a = \min\{\overline{a}, (S^{\text{np}} - x)^+\}$

3. production is set by a simple order-up-to S rule with order-up-to level S smaller than the optimal ones for the non-perishable case ($S < S^{\text{np}}$):
   $a = (S - x)^+$.

None of these rules is stock-age-dependent. Nahmias is primarily interested in the comparison of the performance of these rules. Therefore he considers a clean simple problem setting, leaving out periodicity of demand, production stops and any other problem specific complicating aspects. It is assumed that any shortages are backlogged. By simulation and Fibonacci search techniques he evaluates and optimizes the rules under FIFO issuing, for varying choices of the variable ordering, outdating, and shortage costs and for different demand distributions (exponential, Erlang-2 and Erlang-5). From the execution of simulations he concludes that all rules perform almost equally well since outdating is not a big issue under these rules. A simple order-up-to S rule is preferred for its simplicity in practice.

In 1977 Nahmias [99] is the first to report numerical results for the truly-optimal policy as obtained through successive approximation. Since the computation of a truly optimal policy is seriously hampered by the dimensionality of the state space, he suggest an approximate policy based on a reduction of the dimensionality by leaving out the oldest age categories. He argues that such an approximation is justified under FIFO issuing, since the number of old products is normally relatively small compared to the number of young products.

To allow 'efficient' computation the problem is simplified such that the state space for the truly optimal MDP problem is two-dimensional: the maximal shelf life is set to three periods ($m = 3$) and he assumes zero lead time. Thus the number of products with residual shelf life three days, x_3, equals today's production volume. The number of products left behind from yesterdays replenishment is x_2. The x_1 oldest products do perish at the end of the day if not taken from stock. In an approximate MDP model Nahmias leaves out the oldest products as if no products will be in stock for more than 2 days. The optimal order quantity in this approximate model thus solely depends on stock level x_2. The approximate strategy is compared to the optimal stock-age-dependent MDP policy, which depends on both x_1 and x_2. Nahmias shows by simulation that the approximate MDP policy performs nearly optimal under the given linear cost structure (which includes variable holding, ordering, shortage and outdating costs). Note that in the (late) seventies solving the two-dimensional MDP was already very time-consuming. The approximation strategy was obtained in a few minutes, whereas computing the optimal MDP policy took more than an hour.

Nahmias thus derives numerical bounds on outdating and shortages and thus provides a benchmark to simple rules. Nahmias shows that the approximate MDP policy is better than the approximate order-up-to S rule formalized in Equation (2.2):

$$a = (S - \sum_{r=i}^{m} x_r)^+. \tag{2.2}$$

The approximate order-up-to S rule is a simple stock-age-dependent rule based on the available 'young' stock with a residual shelf life of at least i periods (with $1 \leq i < m$). Note that for $m = 3$ and $i = 2$ this rule can be seen as the weighted order-up-to S rule with $w_1 = 0$ and $w_2 = w_3 = 1$ of Equation (2.1). In 1993, Nandakumar and Morton [102] revisit Nahmias' work for testing some new approximations based on upper and lower bounds on the outdating and shortages figures using the 'Newsboy theory' and other results for non-perishable inventory.

When production does not happen daily or when demand strongly depends on the day of the week, an order-up-to level with (different) order-up-to levels for each day of the week is more appropriate. This clearly complicates both the analysis as well as the search for nearly optimal parameter values. Therefore one is interested in a simple rule for setting good parameter values. In 1983, Katz et al. [75] study the $(\mu + T \cdot \sigma)$-rule for setting order-up-to levels. In the formula, μ is the mean demand for the particular production day and σ is the standard deviation of the demand.

In 1991, Sirelson and Brodheim [132] report in more detail on a case study, in which they compare by simulation several ways to set order-up-to levels S for the order-up-to S rule. One approach is to set $S = (\mu + 3 \cdot \sigma)$, where μ is the average demand until the next order point, and σ is the standard deviation. This approach results in the $S = (\mu + 3 \cdot \sigma)$-rule.

For practical use, they suggest order-up-to levels that do not depend on the second moment of the demand (σ), but solely on μ. Therefore they propose to let S be $k \cdot \mu$, for any real value k. For $\mu = 1$ they report on, what-they-call, normalized stock levels S_μ. The value of S_μ follows from simulation experiments and typically falls in the range of 1.5 to 2.5 depending on the desired balance between shortages and outdating. Through quadratic regression Sirelson and Brodheim estimate outdating and shortages curves in terms of S_μ.

For other $\mu \neq 1$ order up-to levels S are derived from the normalized stock levels S_μ according to

$$S = \mu \cdot S_\mu. \tag{2.3}$$

In their study, Sirelson and Brodheim do distinct A, B and O platelet inventories at 14 hospitals of the Greater New York Blood Program. They do assume that products of different blood groups are no substitutes to each other.

In 1993, Goh, Greenberg and Matsuo [50] present a simulation study, in which two demand categories are distinguished. The demand of one category strongly prefers 'young' products, while the other demand category can be met by somewhat older items. Both demands are satisfied by issuing first the oldest pools in stock that satisfy the respective age restriction. The 'young' pools may or may not be available to meet the demand for old pools. In their simulation experiments the demand for 'young' is 10% up to 50% of the total demand. Analytically obtained upper and lower bounds on shortages and outdating are validated by simulation. The simulation study concerns perishables with a maximal shelf life of 42 days (young products are products up to 10 days old), and the daily demand is Poisson distributed. The study the problem at three scales: the average total demand is respectively 1, 5 or 10 products per day. Outdating appears to be in most cases still 4-6% (and even higher) as some products are return after being issued for a number of days to cross match the RBC product with the blood of the patient. At most a few percent of the demand is lost due to shortages.

In 1997, Hesse, Coullard, Daskin and Hurter [64] report on a case study for a blood bank that serves 35 hospitals in the Chicago area. Over 1995-1996, the overall outdating of platelet concentrates was more than 20%. Hesse et al. classify the hospitals in groups and set 'optimal' periodic review (s, S)-policies for each group. Although they do optimize they do not report on the best parameter values for s and S. In their study they assume a shelf life of 5 days and that the products are issued in a FIFO order. By deterministic simulation using a data set over 1995 – 1996, they show that outdating at individual hospitals can be reduced from more than 20% to 3.5% – 12%.

Recently, a simulation study of the UK blood supply chain is executed by Katsaliaki and Brailsford [74]. Their study covers many different blood products. The simulation model includes the transhipment of blood products from storage centers with excessive stock to centers with deficient stock, as well as emergency and ad-hoc deliveries. Further they acknowledge the limited compatibility of blood from different blood groups. Although they report in [74] significant savings by improved inventory management of RBCs, they do not report on platelet concentrates.

Conclusions

From the above simulation studies we learn that

1. issuing the oldest available products first greatly reduces outdating, but is inappropriate when demand is age-specific,
2. outdating and shortages happen frequently when a great part of the demand requires 'young' products,
3. great savings on outdating, shortages, and operating costs seem to be possible by applying simple order-up-to S rules.

2.4.4 Experimental and Empirical studies

Experimental and empirical studies report on current practice and on testing or implementing new policies in practice. The results reported are thus utmost realistic, but the conclusions hold primarily for the case under consideration. Since the experiments are performed real time and are not replicable, sensitivity analysis is not possible. Thus it is hard to draw general conclusions other than that the outcome of the experiments does or does not support conclusions from more general numerical and simulation studies. For the case of inventory management of blood products, such as platelets pools, the experimental studies, which we discuss below, indeed show that outdating can be reduced significantly.

In 1984, Ledman and Groh [82] from the American Red Cross Blood Services showed that the outdating of platelet pools can be reduced from about 20% to 3%. Their empirical results come from an experiment over a period of 6 months in the Washington region with one regional blood center supplying 60 hospitals with about 60,000 donor platelet concentrates (= 12,000 platelet pools) per annum. At their study platelets have a shelf life of 5 days, but during the first two months of the experiment a fraction of the storage bags can keep platelets for only 3 days. Ledman and Groh did not report on a formal model; instead a team of experts cooperated to carefully set the production levels based on the expected supply of whole blood and on predictions of the demand for BPPs and other blood products.

For RBCs a similar approach is used in part of The Netherlands, as presented in 2003 by Verhoeven and Smid [159] of Sanquin blood bank (region North West, The Netherlands). Through cooperation and sharing knowledge and information concerning the demand for RBCs at hospitals, they manage to smoothen the demand faced by the regional blood bank. In close cooperation with the donor management team, the call for voluntary donors is managed effectively. In [159] and in [157], they explain the ordering policy that inventory managers at hospitals should use. The ordering policy is set by three stock level parameters: a target level S (to determine the size of an order), a threshold $s^{(r)}$ (that triggers regular orders) and a critical level $s^{(e)}$ (for emergency orders). The target level S is set to the average demand per week, $s^{(r)}$ is set to $\frac{3}{7}S$, and $s^{(e)}$ is $\frac{1}{7}S$. This way one strives for a reduction of outdating at the hospitals and great savings in the transhipment of products.

In 2006, Verhoeven reports on experimenting with Vendor Managed Inventory [158]. Sanquin Blood bank North West and three hospitals experiment in letting the blood bank manage the inventories at the hospitals using the rules settled in cooperation with the hospitals. Hospital managers are positive and they admit hospitals are not well educated for the administration of inventories. The applicability of the approach to platelet pools, which are stored centrally at the regional blood bank because of their very short shelf life, is however not discussed neither are figures reported on how many hospitals do adopt the suggested ordering policy and what the savings in outdating and distribution cost actually are.

Mercer and Tao [92] study the allocation and distribution of inventory to regional depots for a food manufacturer. When not all depots are visited daily the transhipment costs needs to be taken into account in deciding when to visit which depots. The work on Vendor Managed Inventory (VMI) is beyond the scope of this thesis.

Hesse et al. [64] investigate the platelet inventory management in the Chicago area. In contrast to the situation in The Netherlands the 35 hospitals in the Chicago area do stock platelets. Hospitals do not pay for excessive ordering: hospitals only pay for the platelet pools that are actually transfused. This policy results too often in 'over-ordering', but reduces the involved transportation cost carried by the blood bank. As a consequence, outdating is significant: at the individual hospitals outdating is 5% – 46%. By simulation they show that periodic review (s, S) policies in combination with a FIFO issuing policy reduces outdating to 3.5% – 12%.

Van Donselaar et al. [150] argue that for a better control of perishables in supermarkets, the order-up-to levels or what they call maximum inventory levels should depend on the (maximal) shelf life of the products. Furthermore, existing ordering systems should make better use of the impact of substitution on outdating: outdating of one product can be reduced by order or display less fresh items of competing products.

A different line of research of perishable inventory control is on pricing and discounting strategies. Alternatively one might prioritize demand and refuse low-priority demand when stock is low. Since we focus on ordering policies we will not consider other policies to control outdating.

For more case studies we refer to a few studies based on simulation ([30], [89], [75], [132], [44], and [74]) or on other numerical techniques [19] using realistic data, which are already discussed in the previous sections.

Conclusions

From experimental studies we learn that

1. on average 5 to 20% of BPPs at blood banks and hospitals become non-transfusable due to outdating,

2. at small hospitals outdating can be even close to 50%,

3. outdating can be reduced significantly through careful production and inventory management.

2.5 Blood platelet pool studies

Many of the above mentioned studies on blood bank inventory management deal with RBCs. As RBCs have a long maximal shelf life of five to six weeks, the inventory management seems to be more complicated for BPPs, which have a very short shelf life of only 5 to 7 days. Moreover the demand for BPPs arises at a much smaller scale than the demand for RBCs.

Besides being the most perishable blood product, a BPP is also the most expensive blood product, due to the high processing cost. Disposing outdated pools is thus a waste of money and undesirable for ethical reasons. Nevertheless, a number of studies report high outdating figures of 10-20% or even higher at European and American blood centers (see [155], [64], [132], [82], and [75], [163]).

A major complication of the PPP is that a great part of the demand strongly prefers 'young' BPPs, of say at most 3 days old, whereas the remaining part of the demand may get older BPPs. Nevertheless in most studies this distinction is not made. Current developments (in The Netherlands) of allowing pools to be stocked for a longer time (up to seven days) make the distinction between demand for 'young' pools and the demand for pools of any age more important.

In the remainder of this section, we summarize the results from studies on the inventory management of blood platelets. In chronological order, the most relevant studies on platelet concentrates are reviewed below. Most of the citations stem from the previous sections, in which the focus was on the different methodologies and techniques in use.

Two of the first platelet studies, that we have found, are the PhD thesis of Deuermeyer [36] in 1976 and the paper of McCullough [89] in 1978. Based on a data set of realized demand figures McCullough examines several ways to improve platelet availability. Deuermeyer uses simulation techniques to investigate the relationship between the (total) stock level and the occurrence of outdating and shortages.

In 1983, Katz et al. [75] showed in a simulation study that an order-up-to S strategy may perform very well, when all demand is met by issuing the oldest BPPs in stock first. The different order-up-to levels S are considered according to $S = \mu + T \cdot \sigma$. The mean demands (μ) as well as the standard deviations (σ) for each day of the week, are estimated based on historical data for two years. Production happens 6 days a week. For given values of T they report almost no outdating and shortages. When the shelf life is only 3 days, outdating is about 2% when T is set to 2.

Ledman and Groh [82] showed empirically that by carefully choosing the production volumes, the outdating of pools can be reduced from about 20% to 2.6% in the Washington region. Instead of using a model, a team of experts met every working day to carefully set the production volume based on the expected needs at hospitals.

In 1991, Sirelson and Brodheim [132] investigate, just as Katz et al., how the replenishment level S of an order-up-to S rule affects the occurrence of outdating and shortages. An appropriate replenishment level S can be read from curves that show the relation between the shortage rate, the outdating rate and the what they call normalized base stock levels S_μ. The curves are obtained by quadratic regression of simulation results. After choosing the normalized stock level S_μ that reflects an acceptable trade-off between shortages and outdating, the order-up-to level S is set to $\mu \cdot S_\mu$, with μ being the mean demand till the next ordering point. Since the replenishment levels are normalized their rule of thumb allows periodic order-up-to S rules and non-stationary demand.

Hesse et al. [64] investigate the platelet inventory management in the Chicago area. In contrast to the situation in The Netherlands, many of the 35 hospitals in the Chicago area keep inventory of platelets. In 1995 – 1996, about 5% – 46% of the BPPs outdate at the hospitals, since hospitals do not pay for excessive ordering. In their study they assume a shelf life of 5 days and FIFO issuing at the hospitals. By setting 'optimal' (s, S)-policies for each stock-keeping hospital outdating is reduced to 3.5% – 12%. A practical question left for future research is on the limited substitutability of platelets units of compatible blood groups.

More recently, Blake et al. [19] modeled the PPP as a finite horizon MDP. Since the number of states may be prohibitively large for realistic problems, Blake et al. round inventory levels and production volumes to what-they-call 'buckets': production and demand happens in multiples of 5 pools. Thus the number of states to consider is reduced significantly. In [19] is shown that the MDP is tractable for varying values of the downsizing factor and how the downsizing affects the cost optimal minimum as well as the computation time.

In 2006, Veihola et al.[155] report on a European study on the efficiency of blood platelet production in 10 European countries. The study includes 17 Red Cross, national and hospital blood centers. It appears that on average about 14% of the platelet units produced is discarded, primarily due to outdating. This figure actually ranges from 7 to 25% when looking at the individual institutes, which vary in annual production scale from 3,345 to 103,643 donor units (or 669 to 20,728 BPPs per year). Veihola et al. state that blood centers may benefit from a closer look at their platelet supply chain. The supply of whole blood donation is not restrictive for platelet production: about 41% of the platelets obtained by whole blood donations is actually used for the preparation of platelet pools. Although much effort is put into extending the shelf life of platelets, great savings can be expected by careful setting the production volumes and sharing information about stock levels and expected demand.

Conclusions

From the studies on the production and inventory management of platelet concentrates we learn that

1. the inventory management of platelets is complicated because of the very short life time,
2. in most of the (theoretical) studies preferences concerning the age of the platelets are not included, and the distinction of blood groups is commonly not addressed,
3. on average 14% of the production becomes outdated according to a recent European study at 17 European institutes spread over 10 countries.

2.6 Research questions

The PPP is a high-dimensional MDP. Although the number of dimensions can be greatly reduced by aggregating the different blood groups into a single blood group, one needs to distinct between products of different ages for deriving an optimal stock-age-dependent ordering policy. No formal quantitative justification is provided in the literature for approximating the PPP by a single product model, while virtually all analytical, numerical and simulation studies concern single product models.

Despite all the research efforts, anno 2002 the existing simple inventory policies are hardly adopted by hospitals, which still make decisions based on experience [171]. In 2003, the importance of effective strategies and cooperation in managing the blood supply chain is once again stressed in the scientific journal *Transfusion* [73] by Robert Jones of the New York Blood Center. The focus was on improving the quality of the BPPs and extending their maximal shelf life, rather than on the logistical issues to improve both the quality and the efficiency. Why did the work done so far not break through? Were blood bank managers not convinced of the value of the models too simplistic? And if so, why not: were the suggested models and solutions considered too 'mathematical', too simplistic, too specific? Or were they afraid of high costs in money and time to get these models validated and implemented?

Anyway, many characteristics of the PPP, which can also be found in other problem settings, are not included in the single product models found in the literature. Of growing importance to blood bank managers and doctors, two types of demand should be distinguished: one category of patients needs 'young' platelets while other patients may receive BPPs of any age (up to the maximal shelf life). Consequently, it seems that the need for age-depending policies becomes more and more urgent, unless one shows and justifies that stock-level-dependent policies suffice. To what extent stock-level-dependent policies are nearly optimal for the PPP is an open question, especially when a great part of the demand gets the 'youngest' products in stock.

Research questions

We formulate the following research questions to be answered in the next two chapters:

1. Is it justified to approximate the PPP by a single-product model?
2. What is the structure of an optimal stock-age-dependent ordering policy?
   (a) What simple rules closely resemble the optimal stock-age-dependent ordering policy?
   (b) How well do stock-level-dependent ordering policies fit to an optimal stock-age-dependent policy for perishables?
   (c) How close to or far from optimal are such simple rules with respect to outdating and shortages?
   (d) How does the structure of the optimal strategy change when operating at a small scale and when fixed order costs apply?
3. How robust are the approach, the optimal policy and the simple rules?
   (a) What is the impact on outdating and the age distribution when a significant part, of say 30–70%, of the demand is met by issuing the youngest BPPs in stock, and the remaining part of the demand receives the oldest products in stock?
   (b) Do changes in the values of problem parameters dramatically affect the results?
   (c) How do the results change when the supply of buffy coats for the production of BPPs is uncertain and occasionally limiting?
   (d) What if production is not possible during holidays?

By combining the strength of successive approximation for computing optimal policies and the power of simulation to read a simple rule and to justify the resulting strategies, we will answer the above questions. Effective production-inventory strategies are to be derived that eventually can be used by blood bank inventory managers.

Outline of next chapters of Part I

In Chapter 3, the combined SDP-Simulation approach for the PPP is presented. In Chapter 4, first we extend the approach to include non-stationary periods (e.g. irregular production breaks), and next we discuss the order problem when operating at a much smaller scale, such as at hospitals, and fixed order costs apply.