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

Haijema, R.

Publication date
2008

Chapter 3

A combined Stochastic Dynamic Programming and Simulation approach

In this chapter we present an MDP-based approach for solving the PPP presented in Section 2.1, which may be exemplar for a number of other inventory problems of perishables that have a fixed shelf life. First, the approach is sketched in Section 3.1. Next, the five steps of the approach are presented in Sections 3.2 – 3.6. The MDP formulation is provided and the different steps of the approach are illustrated using data provided by one of the four Dutch blood banks. In Section 3.7, the robustness of the approach and resulting policies is shown by an extensive simulation study.

Besides presenting the methodology, this chapter reports on a case study that may be of interest to both Operations Researchers and Blood bank managers. Although the focus is on the PPP at blood banks, the approach of finding nearly optimal ordering policies may also apply to hospitals and other production-inventory problems.

3.1 Approach

The approach to solve the PPP combines Stochastic Dynamic Programming for solving an MDP and Simulation for obtaining simple nearly optimal ordering strategies for the PPP. The approach can be divided into five steps that are discussed in the next five sections:

Step 1. Provide an MDP formulation of the PPP for a single product,

Step 2. For given data, solve the MDP by SDP, eventually after scaling the data to downsize the MDP,

Step 3. Derive simple ordering rules from an optimal MDP policy by simulation,

Step 4. Check whether the simple rules perform nearly optimal for the (scaled) problem,

Step 5. Verify by a detailed simulation study whether the simple ordering rules perform well and are robust in a more realistic setting, which could not be modeled accurately in the MDP model. If the rule was obtained after scaling the MDP, the parameters of the rule need to be first re-scaled such that they relate to the original problem size.

Aggregation-Disaggregation

In Steps 1–4, aggregation may appear in three different ways:

- The different BPPs are to a high degree substitutes to each other and are thus aggregated into a single, say universal, product,
- Next, if necessary, the BPPs are aggregated into batches such that the scaled MDP is tractable,
- Finally, to develop, a simple rule the stock of different age categories is aggregated.

In Step 5, the production volumes are disaggregated and the aggregation of blood groups into a single, universal group is justified.

3.2 Step 1 – MDP formulation

A first step in our approach is to formulate the PPP as an MDP that can be solved numerically. The state space must be small enough, to allow the computation to be done in a reasonable time window. Therefore blood groups are not included in the model, such that the dimensionality of the problem does not explode. Blood groups are aggregated as if all BPPs are of a universal group. This approach is valid since BPPs of the different blood groups are to a high degree substitutes according to the compatibility scheme in Figure 2.1. The justification of this aggregation is presented in Section 3.6.

The MDP is formulated in discrete time and for the PPP the natural length of a time slot is one day. Further, the MDP is characterized by the state of the production-inventory system, the decisions to take, the transitions from one state to a next state, and the related direct costs structure.

3.2.1 States

In the first place, we are interested in an optimal stock-age-dependent ordering policy, hence the ages of the pools in stock is to be registered. Also to record outdating we need to keep track of the ageing of the products. Suppose that the maximal shelf life is m days. The detailed information about the number of BPPs in stock of each age category is an m-dimensional vector, $\mathbf{x} \in \mathbb{N}^m$, with elements x_r being the number of BPPs with a residual shelf life of r days as inspected early in the morning.

Since production does happen only during weekdays, and demand is weekday dependent, the day of the week, $d \in \{1, 2, \ldots, 7\}$ for Monday to Sunday, has to be included in the state description as well. A complete state description is thus the pair $(d, \mathbf{x})$.

3.2.2 Decision and decision epoch

Every morning a decision is taken about the number of BPPs (a) to produce during the day. In the MDP model, it is assumed that during the day the production volume is not revised, and that enough buffy coats become available to meet this production level. At the end of the day, after 24 hours, just before the start of the next day, the a BPPs are released and added to stock.

Note that the issuing policy I is not part of the optimization problem but given for each demand category. In Section 3.3 a number of issuing policies I is discussed.

3.2.3 Transitions

In the transition from day d to the next day $d+1$ (where from now on $d+1 = 1$ for $d = 7$), a sequence of events happen as depicted in Figure 3.1. After the inspection of the stock state $\mathbf{x}$ the production volume a is set early in the morning and during the day the next events take place:

1. BPPs are taken from stock to meet the young-demand for j 'young' BPPs and the any-demand for k BPPs,
2. mismatches occur when the young-demand is met by 'old' BPPs,
3. shortages might occur when not all demand can be met from stock,
4. at the end of the day all BPPs age 1 day,
5. some BPPs might become outdated,
6. the production is released and added to stock.

In our discrete time model, the events happen in the above order. We thus assume that the production volume is fixed before the demand for the coming day is known. The young-demand for j 'young' BPPs is assumed to be met before the any-demand is met. This order keeps the process simple and is assumed to be more realistic than the opposite order: the any-demand happens a bit more ad hoc during the day, when new patients are brought into the hospital, e.g. due to accidents, and when complications arise during a surgery. The order influences the number of shortages to each demand category. In fact, one may model any other chronological order of meeting the demand, such as spreading the demand evenly over the day (e.g. when $j \approx 2k$ then one may consider the issuing order: 2 young, 1 any, 2 young, 1 any, etc until all demand is met).

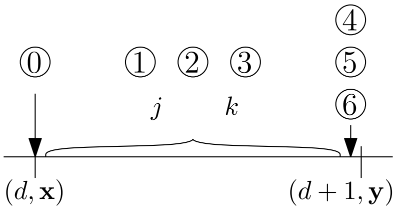

0 = State $(d, \mathbf{x})$ is inspected and production volume a is set,
1 = Meet demand for j 'young' and k BPPs of any age,
2 = Mismatches may happen,
3 = Shortages may arise,
4 = BPPs in stock have aged one day,
5 = Some BPPs may outdate and are removed from stock,
6 = Production is released and added to final stock $\mathbf{y}$.

Figure 3.1: The chronological order of events in the MDP model

Each demand category has a specified issuing policy that prescribes which BPPs to select from stock to meet the demand. The composite issuing policy I is kept simple, such that the number of computations to derive a transition is kept low.

Transition function – The notation of a transition from one state $(d, \mathbf{x})$ to a next state $(d{+}1, \mathbf{y})$ is involved. Therefore we denote the next stock state $\mathbf{y}$ as a function $y(\mathbf{x}, j, k, I, a)$ to emphasizes that it depends on the initial stock $\mathbf{x}$, the composite demand (j, k), the composite issuing policy I and the production volume a. In order to track outdating, shortages and mismatches, the following notations and conventions are introduced:

- i_r^{Y} part of the young-demand that is met by BPPs with r days residual shelf life. i_r^{Y} clearly depends on $(d, \mathbf{x}, j, k, I)$ and the order in which the two types of demands occur.
- i_r^{A} part of the any-demand that is met by BPPs with r days residual shelf life. i_r^{A} clearly depends on $(d, \mathbf{x}, j, k, I)$ and the order in which the two types of demands occur,

and

- i^{Y} number of BPPs that are issued to meet the young-demand $= \sum_r i_r^{\mathrm{Y}}$,
- i^{A} number of BPPs that are issued to meet the any-demand $= \sum_r i_r^{\mathrm{A}}$.

Outdating, shortages and mismatching

With the above notation we get simple expressions for the number of outdated BPPs, for the number of shortages and for the number of BPPs mismatched:

- The number of outdated BPPs, which are removed from stock at the end of the day is $x_m - i_m^{\mathrm{Y}} - i_m^{\mathrm{A}}$,
- The total shortage is $(j - i^{\mathrm{Y}}) + (k - i^{\mathrm{A}})$ BPPs,
- When the young-demand is preferably met by BPPs that have a residual shelf life of l days or more, then the number of mismatches is $\sum_{r=1}^{l-1} i_r^{\mathrm{Y}}$.

Transition probability

The uncertainty in the PPP is solely set by the uncertainty in the demand volumes. For a chosen production volume a the transition probability $P^a_{(d,\mathbf{x}),(d+1,\mathbf{y})}$ from one state $(d, \mathbf{x})$

to a next state $(d+1, \mathbf{y})$ relates thus to the demand distributions of the young-demand and the any-demand. The probability of a demand for j 'young' BPPs and k BPPs of 'any' age on weekday d is denoted by $p_d^{\mathrm{Y}}(j)$ respectively $p_d^{\mathrm{A}}(k)$.

All combinations (j,k) resulting in the same inventory state $\mathbf{y}$, contribute $p_d^{\mathrm{Y}}(j) \cdot p_d^{\mathrm{A}}(k)$ to the transition probability $P^a_{(d,\mathbf{x}),(d+1,\mathbf{y})}$ as in Equation (3.1):

$$P^a_{(d,\mathbf{x}),(d+1,\mathbf{y})} = \prod_{(j,k)\,:\,y(\mathbf{x},j,k,I,a)=\mathbf{y}} p_d^{\mathrm{Y}}(j) \cdot p_d^{\mathrm{A}}(k). \tag{3.1}$$

3.2.4 Costs

The immediate costs include the following cost components:

- c^{H} holding costs per day per BPP in stock at the start of the day,
- c^{O} outdating costs per outdated BPP that is disposed,
- c^{S} shortage costs per BPP not in stock to meet the demand (independent of the demand category),
- c_r^{Y} (age or quality) mismatch costs per issued BPP with residual shelf life r days that violates the age-preference of the young-demand.

The immediate costs when j young BPPs and k BPPs of any age are demanded, is:

$$C(d,\mathbf{x},j,k,I) = \begin{cases} c^{\mathrm{O}} \cdot (x_m - i_m^{\mathrm{Y}} - i_m^{\mathrm{A}}) & \text{outdating costs,} \\ +c^{\mathrm{S}} \cdot (j + k - i^{\mathrm{Y}} - i^{\mathrm{A}}) & \text{shortage costs,} \\ +c^{\mathrm{H}} \cdot x & \text{holding costs,} \\ +\sum_{r=1}^{m} c_r^{\mathrm{Y}} \cdot i_r^{\mathrm{Y}} & \text{(quality) mismatch costs,} \end{cases} \tag{3.2}$$

with the total stock level $x = \sum_{r=1}^{m} x_r$. Note that $C(\cdot)$ depends on d and I via i^{Y} and i^{A}.

Given the issuing policy I, the expected immediate costs to incur in state $(d,\mathbf{x})$ are:

$$\mathbb{E}C^I(d,\mathbf{x}) = \sum_{j,k} p_d^{\mathrm{Y}}(j)\, p_d^{\mathrm{A}}(k)\; C(d,\mathbf{x},j,k,I). \tag{3.3}$$

Remark – The shortage costs may differ between the multiple demand categories. In addition, one may introduce c_r^{A} to capture any soft age-preference of the any-demand. Furthermore, $C(\cdot)$ might depend on the production volume a, when additional cost are made to realize the production of a BPPs, e.g. when the production volume requires working in overtime or when extra effort is needed to have enough donations to realize the targeted production volumes.

Fixed order costs per order can be included, hen it takes time or money to set-up a production run or when one has to pay a fixed amount for a delivery (next to proportional ordering costs). Since these do not apply to the PPP they are left out. Fixed order costs will be considered in the next chapter, in Section 4.5.

3.2.5 Successive approximation

The stochastic dynamic program for the PPP is solved by an SA algorithm, similar to the one introduced in Section 1.2.2. Instead of computing and storing all probabilities $P^{a}_{(d,\mathbf{x}),(d+1,\mathbf{y})}$, we save on RAM memory and only store $p_d^{\mathrm{Y}}(j)$ and $p_d^{\mathrm{A}}(k)$. The transition probabilities are computed using (3.1). SA is implemented as in Section 1.2.2, using the following recursive relation:

$\forall (d,\mathbf{x}) \in \mathcal{X}$:

$$V_n^I(d,\mathbf{x}) = \min_{a\in\mathcal{A}(d,\mathbf{x})} \left(\mathbb{E}C^I(d,\mathbf{x}) \; + \; \sum_{j,\,k} p_d^{\mathrm{Y}}(j)\; p_d^{\mathrm{A}}(k)\; V_{n-1}^I\,(d+1, y(\mathbf{x},j,k,I,a)) \right). \quad (3.4)$$

The SA algorithm start with $n = 1$ and $\mathbf{V}_0^I = \mathbf{0}$ and computes $\mathbf{V}_0^I$ for increasing values of n. The iterative scheme is stopped when $span(\,\mathbf{V}_n^I - \mathbf{V}_{n-7}^I\,)$ is smaller than a small value ϵ. Suppose this happens (at first) at iteration N.

Then the optimal costs level per week is approximated by $\mathbf{V}_N^I - \mathbf{V}_{N-7}^I$. The optimizing actions in the last iteration approximate the optimal policy. If optimal actions under issuing policy I are not stored, then an optimal production strategy π^I follows from Equation (3.5). The optimal production volume on weekday d in stock state $\mathbf{x}$ is:

$$\pi^I(d,\mathbf{x}) = \arg \min_{a\in\mathcal{A}(d,\mathbf{x})} \left(\mathbb{E}C^I(d,\mathbf{x}) \; + \; \sum_{j,\,k} p_d^{\mathrm{Y}}(j)\; p_d^{\mathrm{A}}(k)\; V_N^I\,(d+1, y(\mathbf{x},j,k,I,a)) \right). \quad (3.5)$$

The minimal long-run average costs per week related to this policy is approximated, by any element of $\mathbf{V}_N^I - \mathbf{V}_{N-7}^I$. This approximation is ϵ-accurate. The long-run average costs per day, g, follows from:

$$\left| g - \frac{V_N^I(d, \mathbf{x}) - V_{N-7}^I(d, \mathbf{x})}{7} \right| < \frac{\epsilon}{7}, \tag{3.6}$$

for an arbitrary state $(d, \mathbf{x})$.

3.2.6 Computational complexity

To do the computations in Equation (3.4) for all possible states $(d, \mathbf{x})$, in finite time the state and action space has to be finite. Both the action space and the state space are finite when we imposing an upper bound on the production volumes. This upper bound can be set to the present production capacity, or to an (artificial) finite production capacity for each day of the week. Even when the production capacity is limiting the number of states can be very large. When U is an upper bound on the (artificial) daily production capacity, the number of states is $\mathcal{O}((1+U)^m)$. Hence the number of states grows exponentially fast in the number of age categories m.

We discuss two ways for speeding-up the SA algorithm:

1. The number of states to consider can be reduced by imposing a (fictitious) storage capacity, such that states with an unrealistically-large, sub-optimal number of BPPs in stock are left out.

2. During each iteration it suffices to consider a single value of d: successive values of d may be considered in successive iterations as the PPP is periodic in d.

Imposing an (artificial) storage capacity – Assuming that it is suboptimal to have more than X BPPs in stock, the number of states to consider during each iteration can be reduced by imposing an (artificial) storage capacity of $\overline{x} < mU$ BPPs. All states with $x = \sum_{r=1}^{m} x_r > \overline{x}$ are not included in the state space $\mathcal{X}$. Consequently one should remodel transitions to a state that is outside the (truncated) state space. For example, when the initial total stock level $x = 20$ plus the production $a = 6$ minus the demand $j + k = 3$ exceeds $\overline{x} = 21$, we enforce to stay in the state space by removing the oldest $(x + a - (j + k) - \overline{x})^+ = 2$ BPPs from stock, such that one ends up in a state with $x = \overline{x}$.

We do not want this modeling trick to affect the optimal strategy too much, as the storage capacity is artificially introduced solely to reduce the state space. In a simulation model, we label such actions as *'overproduction'*, as the oldest BPPs would not be removed when the production volume a was set lower. By simulation we check whether overproduction happens very rarely, to justify the height of the chosen artificial capacity. We prefer to set $\overline{x}$ as small as possible, but large enough such that long-run average overproduction is (virtually) zero.

Iterate over one periodic class per iteration – In Equation (3.4) we iterate over all $\mathbf{x} \in \mathcal{X}$ for all days d, while it is more efficient to start with a specific d, say $d = 7$, and iterate only over all $\mathbf{x}$ in the *periodic class* $\mathcal{X}(d = 7)$. In the second iteration we choose $d = 6$ and iterate over all $\mathbf{x} \in \mathcal{X}(d = 6)$, etc. After 7 iterations d is reset to 7. This way the computational burden per iteration is reduced by a factor 7. Thus making use of the periodicity in d results in a more efficient implementation.

3.3 Step 2 – Scaling and Solving

Despite the suggested ways for speeding up an SA algorithm, the number of states to visit during a single iteration can be tremendous for the PPP with realistic data. A straight-forward computation of the optimal production policy is often not possible. In this chapter we use data gathered in 2003-2004 for one of the four Dutch blood banks, Sanquin's Division North-East. The number of BPPs to be kept in stock to meet the demanded is too large to solve the problem by SA. Therefore we suggest to scale the problem. We subsequently discuss, in this section, the data, the scaling, and the results for the scaled problem.

3.3.1 Data

The chosen parameter values are estimates based on data from, and communications with, the Logistics manager and the head Blood Processing of Sanquin's Division North-East (The Netherlands) amongst other [16, 66, 133].

Maximal shelf life

BPPs have a shelf life of at most 5 days, but in the near future the shelf life will be raised to 7 days, depending on improved processing techniques, clinical studies and legislation. The state space, including the weekday, is thus currently 6-dimensional, and is expected to be 8-dimensional in the near future.

Demands

The demand distributions are approximated by fitting the distribution on an estimated mean and standard deviation coefficient of variation. The mean demand is estimated using historical data and expectations regarding the future demand. The standard deviation of the demand is set as if demand is Poisson distributed.

The mean demands on Monday – Sunday are roughly 26, 21, 32, 21, 26, 8, and 10 BPPs. Blood bank managers estimate that (at most) 60-70% of the demand is *young-demand*, which prefers 'young' BPPs of say at most 3 days old (counted from the release date). Young-demand may be mismatched by issuing older pools, although is less favored. In the current practice, no strong distinction is made between young-demand and any-age: both types of demand are met by issuing the oldest BPPs in stock first. Nevertheless,

Table 3.1: Mean daily demand for BPPs at a Dutch blood bank.

Demand	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Weekly
'Young'	20	15	26	15	20	0	0	96
'Any-age'	6	6	6	6	6	8	10	48
Total	26	21	32	21	26	8	10	144

to increase their quality of service, blood bank managers may be interested in different issuing policies for the two types of demand.

In Table 3.1 the average demand of each demand category, as well as the total demand, is reported for each day of the week. Demand is assumed to be stationary but periodic, with period 7 days: e.g. all Mondays are stochastically the same. The average weekly demand is 144 pools and the average daily demand varies over the weekdays. Demand peaks on Wednesdays and continues during weekends although regular production of BPPs stops during weekends. The demand for 'young' BPPs is during weekend virtually zero as no transfusions are scheduled during the weekends. On Wednesday the ratio between young-demand and any-demand is 26 :6, thus more than 80% of the demand prefers 'young' BPPs. On average this ratio is 96 : 48, hence 67% of the demand prefers young BPPs.

Both types of demand are modeled as if demand is approximately Poisson distributed. The coefficients of variation for the total demand range from $1/\sqrt{32} \approx 0.18$ on Wednesday to $1/\sqrt{8} \approx 0.35$ on Saturday. Ledman and Groh [82] report similar values for the coefficients of variation for an empirical case study in region Washington (DC). Apparently the assumption of Poisson distributed demand is quite realistic.

Issuing policy

In the current practice, BPPs are issued in FIFO order, i.e. oldest BPPs are issued first. Since a great part of the demand prefers 'young' BPPs, alternative issuing policies are considered. Next to the common FIFO and LIFO issuing policies discussed in Section 2.2.1, we introduce a new issuing policy that we call FIFOR(l) for FIFO Restricted to age category l. In fact this new issuing policy is a mixture of FIFO and LIFO when mismatching is allowed. Under FIFOR(3) first BPPs with residual shelf life of 3 or higher are issued in the FIFO order, when not all demand can be met this way and if mismatching is allowed older BPPs are issued in a LIFO order. FIFOR(l) is primarily developed for meeting the young-demand.

The FIFOR(3) policy is illustrated in Figure 3.2, where the young-demand for 11 BPPs is met from stock state $\mathbf{x} = (1, 4, 2, 1, 5)$. Since not enough young BPPs are in stock to meet all young-demand, mismatching occurs. Three BPPs are of age category 2, i.e. these BPPs have a residual shelf life of 2 days, consequently one incurs a penalty of $3 \times c_2^Y$.

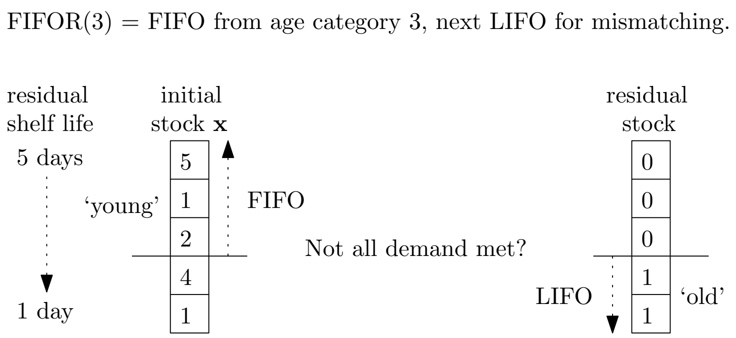

Figure 3.2: An illustration of issuing policy FIFOR(3).

Our FIFOR policy is similar to the upward and downward substitution rule as presented by Deniz, Scheller-Wolf and Karaesmen[35]. When mismatching is not allowed, FIFOR corresponds to the same issuing policy as presented in Pierskalla and Roach [114].

The issuing policy I describes how the two types of demands are met from stock. I is thus a composite issuing policy. In our notation we will first denote the issuing policy for the young-demand followed by the policy with respect to the any-demand. Thus $(\text{LIFO}, \text{FIFO})$ implies issuing the youngest pools to meet the young-demand and issue the oldest pools in stock to meet the any-demand. Combinations that we consider for I are:

- $(\text{FIFO}, \text{FIFO})$: meet all demand by the oldest BPPs in stock,
- $(\text{LIFO}, \text{LIFO})$: meet all demand by the youngest BPPs in stock,
- $(\text{LIFO}, \text{FIFO})$: LIFO for young-demand, FIFO for any-demand, and
- $(\text{FIFOR}(3), \text{FIFO})$: FIFOR(3) for young-demand and age-mismatching is allowed.

Cost structure

The cost parameters may take fictitious values, since only their relative values are of importance for balancing the different criteria rather than their absolute value. The following cost figures are chosen:

- $c^{\mathrm{O}} = 150$ euro per BPP
  The major concern is to balance outdating and shortages. The cost of an outdated BPP is 150 euro, which is the variable production cost that are not paid off.

- $c^{\mathrm{S}} = 750$ per BPP short
  A shortage requires getting a BPP from another blood bank at high transhipment costs, further shortages imply a loss in goodwill. Therefore the variable shortage costs are set to a multiple of say five times the outdating costs.

- $c^{\mathrm{Y}} = (200, 200, 0, 0, 0)$
  Age or quality mismatch costs are incurred when young-demand is met by BPPs that have a residual shelf life of less than $l = 3$ days. In the current practice meeting the age-preferences is getting more-and-more important, although no preference for BPPs of a specific age category is formulated by doctors in the hospital. Instead we make a clear distinction between the two demand types by setting the mismatching cost quite high when BPPs with only 1 or 2 days residual shelf life: $c^{\mathrm{Y}}_1 = c^{\mathrm{Y}}_2 = 200$. The mismatch costs are higher than the outdating costs to reflect that the quality of service is more important than outdating.

- Proportional holding costs are small, say 0.1.

Capacity constraints

Although the show-up of donors, the production capacity, and the storage space is almost never limiting in The Netherlands, we need to set limits to reduce the state space. For the given demand figures we assume that, under an optimal ordering policy, the production stays well below $U = 80$ BPPs per day. On Saturday and Sunday the production capacity is set to zero. According to expression $\mathcal{O}((1 + U)^m)$ obtained in Section 3.2, the number of states is then still in the order of billions when the maximal shelf life is 5 days. A reasonable bound on the storage capacity may be the average demand per week, hence $\overline{x} = 144$ BPPs. The total number of states is still very large. Therefore we suggest to scale data for the PPP.

3.3.2 Scaling

For the data of Sanquin's Division North-East, the MDP has billions of states to consider, whereas the optimal strategy is only tractable for cases with at most a few million states. A way to reduce the computational complexity is by scaling the problem parameters. Therefore we transform the data as if demand happens in batches of B pools. Consequently production volumes, and stock levels x_r are multiples of B pools. The problem can thus be solved at a more coarse or aggregated grid. The number of states reduces roughly by a factor B^m from $\mathcal{O}((1+U)^m)$ to $\mathcal{O}(1+\lceil\frac{U}{B}\rceil^m)$. In the numerical study of the PPP, we set the *'scaling factor'* B to 4.

Scaling the production and storage capacity – The production capacity at each weekday is then set to 20, and the storage capacity is 36 BPPs. Then the number of states in the MDP is 850,414. (In the scaled cases that we consider, the number of states is reduced a bit more by setting a day-dependent production capacity, and reducing the storage capacity even further.) Hence the scaled MDP can be solved in a reasonable amount of time. But first the demand distribution and the costs are to be scaled.

Scaling of demand distributions – The mean demand figures reported in Table 3.1 are divided by $B = 4$. To keep the same degree of uncertainty, the coefficients of variation (cv) are kept the same as those for the unscaled problem. In Table 3.2 the mean demand figures for the scaled MDP are reported, as well as the coefficient of variation.

So far the scaling is simple. A bit more complex is the scaling of the daily demand distributions, which should fit to the scaled means and coefficients of variations (of the original data). Fortunately, Adan, Van Eenige and Resing [2] have developed a procedure for fitting discrete distributions on the first two moments. We do not report their recipe here, but simply report on an example.

Table 3.2: Scaled mean daily demand figures (in batches of 4 pools), and coefficients of variation (cv).

Demand	Statistic	Mon	Tue	Wed	Thu	Fri	Sat	Sun
'Young'	mean	5.00	3.75	6.50	3.75	5.00	0	0
	cv	0.22	0.26	0.20	0.26	0.22	0	0
'Any'	mean	1.5	1.5	1.5	1.5	1.5	2.0	2.5
	cv	0.41	0.41	0.41	0.41	0.41	0.35	0.29

For example, consider the mean any-demand on Monday is $6/4 = 1.5$ batches. The cv is set to the original $cv = 1/\sqrt{6} \approx 0.41$. According to Adan et al., the demand distribution for the any-demand on Monday is approximated by a Binomial distribution: $\widetilde{p}_1^{\mathrm{A}} \sim \mathrm{Bin}(2, 0.75)$. The distribution of the young-demand on Mondays is a mixture of two Binomial distributions, conform Equation (3.7):

$$\widetilde{p}_1^{\mathrm{Y}} \sim 0.59 \cdot \mathrm{Bin}(6, 0.78) + 0.41 \cdot \mathrm{Bin}(7, 0.78) \tag{3.7}$$

One should be warned that for some combinations (μ, cv), no discrete demand distribution exists. We come back to this point in the discussion in Section 3.8.

All other demand distributions are derived in a similar way. To reduce the computational burden per iteration, we do truncate the probabilities at some level $\bar{j}$, such that $\sum_{j=0}^{\bar{j}} \widetilde{p}_d^{\mathrm{Y}}(j) < \theta$ for small θ. The probability distribution p_d^{Y} corresponds to $\widetilde{p}_d^{\mathrm{Y}}$, except for $p_d^{\mathrm{Y}}(\bar{j})$, which equals $\sum_{j=\bar{j}}^{\infty} \widetilde{p}_d^{\mathrm{Y}}(j)$.

Scaling of costs – Finally, the variable costs are multiplied by $B = 4$, since outdating, shortages and mismatching now occurs in batches of 4 BPPs.

3.3.3 Optimal scaled MDP solution for different issuing policies

The scaled MDP can now be solved by SA. The solution gives hopefully a good approximation of the results for the original unscaled problem. Furthermore, the structure of the scaled MDP solution may be visualized by simulation such that one may search for a simple rules that resembles the structure of the optimal policy. This will be checked in the next sections. In this section we investigate the typical results we can expect under an optimal strategy. Since the results strongly depend on the chosen issuing policy, four different issuing polices are considered.

Using the notations introduced in Section 3.2 the parameter values after scaling are as summarized in Table 3.3.

In Table 3.4 we report the annual costs and the related performance measures like the annual outdating and shortages figures and the mismatch of the demand for fresh. Outdating, shortages and mismatches are reported in absolute values (in batches of 4 pools) as well as in percentages of the annual demand. All results are obtained by multiple numerical evaluations of the Markov chain of the optimal strategy with only one costs component included in the cost structure at a time.

Table 3.3: MDP parameter values for scaled PPP (Sanquin, NE, 2003/2004).

m	$=$	5 (to 7)
c^{S}	$=$	3,000
c^{O}	$=$	600
c^{H}	$=$	0.4
c^{Y}	$=$	$(800, 800, 0, 0, 0)$
I	$\in$	$\{(\text{FIFO}, \text{FIFO}), (\text{LIFO}, \text{LIFO}), (\text{LIFO}, \text{FIFO}), (\text{FIFOR(3)}, \text{FIFO})\}$
$p^{\text{Y}}_d(j)$	:	fitted on first 2 moments that are reported in Table 3.2
$p^{\text{A}}_d(k)$	:	fitted on first 2 moments that are reported in Table 3.2

In 2004, all demand was met by issuing the oldest BPPs first, $I = (\text{FIFO}, \text{FIFO})$, and outdating was about 10-20% against a shortage of about 1%. Mismatches were not registered. Under the optimal MDP strategy and $I = (\text{FIFO}, \text{FIFO})$ both outdating and shortages are less than 1% and about 6.3% of the young-demand is mismatched. The implied (fictitious) total costs are about $(14 \cdot 600 + 13.9 \cdot 3,000 + 78.9 \cdot 800 \approx)$ 115,727 euro per year.

Issuing the 'youngest' batches in stock results in less shortages and mismatches, but 156 batches outdate under the optimal MDP strategy with $I = (\text{LIFO}, \text{LIFO})$. The annual demand is 1872 batches, hence more than 8% of the production is not transfused because of outdating. The fictitious annual costs are about 125,304 and thus about 10,000 euro higher than under $I = (\text{FIFO}, \text{FIFO})$.

When different issuing policies are related to the two types of demand, the expected annual costs are much lower than under (FIFO, FIFO) or (LIFO, LIFO). The composite

Table 3.4: Mean performance figures for the MDP-optimal ordering policy under different issuing policies I. (Numerical evaluation by value iteration.)

Issuing policy	Outdating [a]		Shortage [a]		Quality mismatch [b]		Annual costs
(FIFO, FIFO)	14	0.76%	13.9	0.75%	78.9	6.30%	115,727
(LIFO, LIFO)	156	8.32%	8.4	0.45%	8.18	0.66%	125,304
(LIFO, FIFO)	37	1.95%	4.9	0.26%	0.09	0.01%	36,605
(FIFOR(3), FIFO)	35	1.89%	4.8	0.26%	0.005	0.00%	35,759

[a] in batches of 4 pools per year; % of total 1872 batches (7488 pools),
[b] in batches of 4 pools per year; % of young-demand of 1248 batches (4992 pools).

issuing policies (LIFO, FIFO) and the more complicated (FIFOR(3), FIFO) both result in very low outdating, shortage and mismatching figures. The cost optimal production policy related to a (LIFO, FIFO) issuing policy results in only 2% outdating, not even 0.3% shortages and virtually no mismatches.

The savings on the annual (fictitious) costs of using (LIFO, FIFO) instead of (FIFO, FIFO) is about 80,000 euro. The optimal production policy under (FIFOR(3), FIFO) performs almost equally well.

Preliminary conclusions – Based on the results for the scaled MDP, we may draw some preliminary conclusions:

- The annual outdating costs are under the MDP and $I = (\text{LIFO}, \text{FIFO})$ about 21,900 euro, whereas in the current practice these costs are $15\% \cdot 7488 \cdot 150 = 168,500$ euro per year. (This cost difference is even greater for $I = (\text{FIFO}, \text{FIFO})$).
- When shortages are 1% in the current practice, then the difference in the total annual shortages and outdating costs are about 187,400 euro: 224,000 ($= 15\% \cdot 7488 \cdot 150 + 1\% \cdot 7488 \cdot 750$) in the current practice, against 36,600 euro under the optimal MDP policy for $I = (\text{LIFO}, \text{FIFO})$.
- The cost difference is slightly bigger for $I = (\text{FIFOR}(3), \text{FIFO})$.

These conclusions are preliminary and require a more solid base

Average age of issued batches

Another relevant performance criterion is the age of the BPPs when they are taken from stock to meet the demand. Through simulation of the optimal policy for 100,000 weeks, the age distribution of the batches is derived. The age is registered for batches that are taken from stock to meet the demand and it is counted from the day of donation. Since the production lead time is one day, the age of a batch is always 1 or higher. Table 3.5 reports the average age of the batches upon meeting the demand. It appears that when all demand is met by FIFO, the mean age under the optimal MDP policy is 2.3 days, while this is only 1.6 days for $I = (\text{LIFO}, \text{LIFO})$.

$I = (\text{LIFO}, \text{LIFO})$ yields the lowest mean age of the batches, but, as we have seen it results in many outdated batches. Mixing an FIFO and LIFO policy gives an average age of 2 days; LIFO for young-demand and FIFO for any-demand seems to be a good

Table 3.5: Mean age of the platelets upon meeting the demand. Results are obtained through simulation of the optimal MDP policy for 100,000 weeks. Age is in days counted from day of donation: platelets are at least 1 day old.

Issuing policy	Mean age Young-demand	Mean age Any-demand	Mean age Total demand
(FIFO, FIFO)	2.3	2.4	2.3
(LIFO, LIFO)	1.5	1.8	1.6
(LIFO, FIFO)	1.4	3.2	2.0
(FIFOR(3), FIFO)	1.8	2.5	2.0

comprise between a low age and low outdating figures. $I = (\text{FIFOR}(3), \text{FIFO})$ yields the same average age of the batches, but the young-demand gets the youngest batches under $I = (\text{LIFO}, \text{FIFO})$.

3.3.4 Conclusions

By aggregating the demand as if it happens in batches of $B = 4$ pools the optimal stock-age-dependent ordering policy for the scaled PPP can be derived. Using realistic data for one of the four Dutch blood banks the optimal MDP policy is computed and evaluated numerically by evaluation of the average costs in Markov chains and by simulation. We may conclude that:

1. a great reduction of outdating from 15-20% in practice to less than 2% seems to be possible, even when two thirds of the demand gets the youngest BPPs in stock.
2. $(\text{LIFO}, \text{FIFO})$ is favored for yielding very low figures for outdating, shortages, and mismatching costs, as well as for providing the young-demand with young BPPs.
3. shortages are low under the optimal MDP policy with $I = (\text{LIFO}, \text{FIFO})$: on average about 0.3% of the demand requires the delivery of BPPs from another blood bank.

The economic cost saving by implementing an optimal MDP policy seems to be big: a rough estimate for the blood bank under consideration indicates that 147,000 euro per year can be saved by the reduction of outdating alone. However, the optimal policy is tractable only after scaling, and the optimal policy may be too complicated for practical use. Furthermore, a number of sensitivities, e.g. regarding the different blood groups, need to be investigated, which we will present in the Sections 3.6 and 3.7.

3.4 Step 3 – Search for simple rules in MDP policy

In the third step of our approach, the structure of the optimal stock-age-dependent ordering policy, as computed for the scaled PPP, is investigated. This may help us in re-scaling (nearly) optimal production volumes fro practical use. We limit the study to ordering policies under the $(\text{LIFO}, \text{FIFO})$ issuing policy. Although the structure depends on the problem data, we keep focussing on the case of Sanquin's Division North-East. Sensitivities regarding the data are investigated in the next sections and in Chapter 4.

In principle an optimal strategy can be very complicated, since optimal production volumes depend on the age of the BPPs in stock. If a simple rule exist that performs nearly optimal, a simple rule is favored for practical use. After a discussion of the optimal stock-age-dependent MDP strategy in Section 3.4.1, we investigate whether stock-*level*-dependent policies can be derived from it. Two stock-*age*-dependent policies are considered in Section 3.4.3.

An open question that we hope to answer is whether a single number gives enough information for setting nearly optimal production decisions for the complicated PPP. This may not be the case as a great part of the demand receives the 'youngest' products in stock and the maximal shelf life is only 5 days.

3.4.1 An optimal stock-age-dependent production policy (MDP)

The MDP solution for the scaled PPP was derived in the previous step of the approach. In Table 3.6 we report the optimal production volumes for a selection of states on a Wednesday morning. Although in all selected states the total inventory level x equals 14 batches, the production volume ranges from 3 to 7 batches. Apparently an order-up-to rule with a fixed order-up-to level does not apply, nor is the production volume fixed. Since the states in the table are only a selection of the so many states at which the system may be on Wednesday morning, we should investigate how likely each state occurs. Therefore we may look at the equilibrium distribution of states under the optimal policy.

Table 3.6: Optimal production volumes on Wednesday under I = (LIFO, FIFO) for a selection of stock states with in total $x = 14$ batches in stock. ($x_2 = x_3 = 0$ due to the production stop in the weekends).

Stock state $\mathbf{x}$					Opt. prod.	Order-up-to
x_1	x_2	x_3	x_4	x_5	volume a	level $S_3 = x + a$
0	0	0	3	11	4	18
1	0	0	5	8	3	17
2	0	0	4	8	4	18
2	0	0	5	7	3	17
3	0	0	1	10	5	19
3	0	0	5	6	5	19
4	0	0	1	9	6	20
4	0	0	4	6	6	20
5	0	0	1	8	7	21
6	0	0	2	6	7	21

Equilibrium distribution and State frequencies

Obtaining the equilibrium distribution $\mathbf{q}^\pi$ of the states under strategy π, can be done iteratively by value iteration. For increasing n, Equation (3.8) is computed starting with $n = 1$ and with initial distribution $q_0^\pi(d, \mathbf{x}) = I(\mathbf{x} = \mathbf{0})$ for all states $(d, \mathbf{x})$. The equilibrium distribution $\mathbf{q}^\pi$ is approximated by $\mathbf{q}_N^\pi$, for N sufficiently large, such that $span(\mathbf{q}_N^\pi - \mathbf{q}_{N-1}^\pi) < \epsilon$. The value of ϵ must be small compared to the probability $q_N^\pi(d, \mathbf{x})$ for the great majority of states $(d, \mathbf{x})$, say $\epsilon = 10^{-15}$.

$$\forall (d, \mathbf{x}) \in \mathcal{X}: \qquad q_n^\pi(d, \mathbf{x}) = \sum_{\mathbf{y}} q_{n-1}^\pi(d, \mathbf{x})\, P^\pi_{(d,\mathbf{x}),(d+1,\mathbf{y})}. \tag{3.8}$$

For details about the convergence, see the paper of Van der Wal and Schweitzer [148].

For studying the structure of the optimal strategy, we do not need an accurate equilibrium distribution as many states occur with a very small probability. Therefore, we suggest to simulate the optimal policy and consider only those states that occurred during the simulation. For each state that occurs we store its frequency, i.e. the number of visits to that state. After a long simulation run the relative frequencies are computed, which approximate the equilibrium distribution. In the next sections we discuss how the thus obtained simulation-based frequency tables help in approximating the optimal ordering strategy for the PPP under consideration.

3.4.2 Stock-level-dependent ordering rules

To investigate whether a simple stock-level-dependent rule fits closely to the optimal MDP policy, we visualize the structure of the optimal strategy through $(state, action)$-frequency tables in Table 3.7. The stock is aggregated into a single number $x = \sum_{r=1}^{m} x_r$. In the top row of the tables one reads the values that x has taken during the simulation; e.g. in Table 3.7(a) the total stock x ranges from 4 to 17 batches on the 100,000 simulated Mondays. The bottom row of each table shows how often each value of x has been observed during the simulation. Most-frequent (24,465 times out of 100,000) a stock state of 9 batches in total is observed on Monday morning.

The optimal actions $\pi(d, \mathbf{x})$ are translated into order-up-to levels S_d, which are simply the sum of x and $\pi(d, \mathbf{x})$. In Table 3.7(a) the first column shows that the S_d on the simulated Mondays ranges from 11 to 22 batches. When $\pi(d, \mathbf{x}) = 0$ the order-up-to level S_d is set to x to emphasize zero production, although any level $S_d \leq x$ relate to zero production. The last column shows how frequent each order-up-to level applies. The most-frequent order-up-to level on Mondays is 16 batches, which fits in 54.5% of the states visited.

When all (state,action)-frequencies on some day d are on a single row, the optimal ordering policy for that day is an order-up-to S strategy with fixed periodic order-up-to levels S_d. When all (state,action)-frequencies are on an upward diagonal, then a fixed order-quantity applies.

In Figure 3.3 we illustrate how three simple ordering rules for Mondays are read from the tables: the simple order-up-to S rule, a bounded order-up-to S rule, and the fixed order quantity rule. For each of the three rules we explain how nearly optimal parameter values are read and how good the rules resemble the optimal strategy. Therefore we define the *goodness-of-fit* of a rule as the (approximate) relative frequency of states in which the rule prescribes exactly the same order quantity as the optimal strategy.

Order-up-to S rule

For the order-up-to S rule one reads (nearly) optimal order-up-to levels S_d from the frequency tables in Table 3.7. Clearly, the order-up-to levels depend on the day of the week. When S_d is the most-frequent order-up-to level read from the table for weekday d then a (nearly) optimal production decision follows from Equation (3.9).

$$a = (S_d - x)^+. \tag{3.9}$$

Table 3.7: Simulation frequency tables of scaled MDP strategy under $I = (\text{LIFO}, \text{FIFO})$.

(a) (State, action)-frequency tables for 100,000 simulated Mondays.

Stock x	4	5	6	7	8	9	10	11	12	13	14	15	16	17	Freq(S_1)
Up-to S_1															
22														1	1
21													10		10
20											2	43			45
19									5	12	314				331
18								26	61	1366					1453
17							76	213	4383						4672
16						24465	19514	10558							54537
15					21002										21002
14				12498											12498
13			4551												4551
12		824													824
11	76														76
10															0
:															:
0															0
Freq(x)	76	824	4551	12498	21002	24465	19590	10797	4449	1378	316	43	10	1	100000

(b) (State, action)-frequency tables for 100,000 simulated Tuesdays.

Stock x	0...5	6	7	8	9	10	11	12	13	14	15	16	17	18	Freq(S_2)
Up-to S_2															
26														1	1
25													2		2
24												9			9
23										1	28				29
22								2	2	170	1				175
21						3	6	15	563	21	5				613
20					9	31	75	1479	209	35					1838
19				31	127	213	2968	913	326						4578
18			17963	20849	25014	19853	7626	1450							92755
17															0
:															:
0															0
Freq(x)	0	0	17963	20880	25150	20100	10675	3859	1100	227	34	9	2	1	100000

(c) (State, action)-frequency tables for 100,000 simulated Wednesdays.

Stock x	0...8	9	10	11	12	13	14	15	16	17	18	19	20	21	Freq(S_3)
Up-to S_3															
26														1	1
25											4	2			6
24										28	9	1	1		39
23									91	62	21	4			178
22								458	224	79	12	1			774
21							1301	753	243	34	1				2332
20						3129	1880	697	84	17	6				5813
19					4745	3740	2284	679	140	13					11601
18				12888	25368	23761	11559	790	112	20					74498
17							695	3199	598						4492
16								85							85
15															0
:															:
1															0
0									98[a]	81[a]	2[a]				181 [a]
Freq(x)	0	0	0	12888	30113	30630	17719	6661	1590	334	55	8	1	1	100000

(d) (State, action)-frequency tables for 100,000 simulated Thursdays

Stock x	0...3	4	5	6	7	8	9	10	11	12	13	14	15	16	Freq(S_4)
Up-to S_4															
17								399	46						445
16					4126	15584	26550	25562	1256	108	24	3			73213
15				155				314	15703	7305	2269	517			26263
14			3												3
13															0
:															:
0													69[a]	7[a]	76[a]
Freq(x)	0	0	3	155	4126	15584	26550	26275	17005	7413	2293	520	69	7	100000

(e) (State, action)-frequency tables for 100,000 simulated Fridays.

Stock x	0...3	4	5	6	7	8	9	10	11	12	13	14	15	16	Freq(S_5)
Up-to S_5															
23											4	2	2	1	9
22										56	46	34	8	4	148
21								44	345	831	634	944	88		2886
20						3591	17946	30836	27269	13380	3935				96957
19															0
:															:
0															0
Freq(x)	0	0	0	0	0	3591	17946	30880	27614	14267	4619	980	98	5	100000

[a] $S_d = 0$ to emphasize zero production

Stock x	4	5	6	7	8	9	10	11	12	13	14	15	16	17	Freq(S_d)
Up-to S_1															
22														1	1
21													10		10
20											2	43			45
19									5	12	314				331
18								26	61	1366					1453
17							76	213	4383						4672
16						24465	19514	10558							54537
15					21002										21002
14				12498											12498
13			4551												4551
12		824													824
11	76														76
10															0
:															:
0															0
Freq(x)	76	824	4551	12498	21002	24465	19590	10797	4449	1378	316	43	10	1	100000

(a) Reading an order-up-to S rule with a fixed order-up-to level.

Stock x	4	5	6	7	8	9	10	11	12	13	14	15	16	17	Freq(S_d)
Up-to S_1															
22														1	1
21													10		10
20											2	43			45
19									5	12	314				331
18								26	61	1366					1453
17							76	213	4383						4672
16						24465	19514	10558							54537
15					21002										21002
14				12498											12498
13			4551												4551
12		824													824
11	76														76
10															0
:															:
0															0
Freq(x)	76	824	4551	12498	21002	24465	19590	10797	4449	1378	316	43	10	1	100000

Produce up to 16 batches

max. prod. = 7 batches

min. prod. = 5 batches

(b) Reading a bounded order-up-to S rule with bounded order quantities.

Stock x	4	5	6	7	8	9	10	11	12	13	14	15	16	17	Freq(S_d)
Up-to S_1															
22														1	1
21													10		10
20											2	43			45
19									5	12	314				331
18								26	61	1366					1453
17							76	213	4383						4672
16						24465	19514	10558							54537
15					21002										21002
14				12498											12498
13			4551												4551
12		824													824
11	76														76
10															0
:															:
0															0
Freq(x)	76	824	4551	12498	21002	24465	19590	10797	4449	1378	316	43	10	1	100000

(c) Reading a fixed order quantity.

Figure 3.3: Reading simple rules from ($state, action$)-frequency tables of optimal production policy (for Monday) under (LIFO, FIFO) issuing.

From the frequency tables we read most-frequent order-up-to levels $(S_1, S_2, \ldots, S_5) = (16, 18, 18, 16, 20)$. The goodness-of-fit on Tuesday to Friday is high: 92.7%, 74.5%, 73.2%, and 97% respectively. But as illustrated in Figure 3.3(a), the best order-up-to level on Monday fits to only 54.5% of the states visited during the simulation. On Tuesday and Wednesdays an order-up-to S rule with a fixed order-up-to level of 18 batches resembles the optimal production decisions for the most-frequently visited states, but still in about a quarter of the states visited, a higher production volume is optimal. Maybe a more refined order-up-to S rule approximates better the optimal ordering strategy.

Bounded order-up-to rule

Especially on Mondays one observes an upward diagonal that bends at order-up-to level 16, as highlighted in Figure 3.3(b). The fixed order-up-to level 16 applies only when the stock level x is between 9 and 11 batches. When the stock level is 9 or lower a fixed production quantity of 7 batches is optimal. When the stock level is 11 or higher, most often 5 batches are ordered. This suggests that the optimal production volumes on Monday can be approximated by what we call the *'bounded order-up-to S rule'* with a fixed order-up-to level $S_1 = 16$, minimum order quantity $\underline{a}_1 = 5$, and maximum quantity $\overline{a}_1 = 7$ batches.

The bounded order-up-to S rule has thus three parameters for each day of the week:

- S_d the fixed order-up-to level,
- $\underline{a}_d$ the minimum production volume,
- $\overline{a}_d$ the maximum production volume.

For given d and x the production volume a is computed as follows:

$$a = \begin{cases} \underline{a}_d & \text{, if } \; S_d - x < \underline{a}_d, \\ S_d - x & \text{, if } \; \underline{a}_d \leq S_d - x \leq \overline{a}_d, \\ \overline{a}_d & \text{, if } \; \overline{a}_d < S_d - x, \end{cases} \tag{3.10}$$

Or in short: $a = \max\{\, \underline{a}_d \,, \, \min(S_d - x, \overline{a}_d)\}$.

On Mondays, the bounded order-up-to S rule with the given parameter values matches to the optimal production decisions in $(93,488 + 4,383 + 1,366 + 314 + 43 + 10 + 1 =)$ $99,605$ states out of the 100,000. The goodness-of-fit on Mondays is thus 99.6%. For the other weekdays we can read the best upper and lower bounds on the production volume, such that they maximize the fit to the optimal strategy. When the lower bound equals the upper bound a fixed order quantity rule applies.

Fixed-order-quantity rule

A very simple rule one can think of is to produce a fixed quantity every day, which may be week-day dependent. Q_d batches are to be ordered on weekday d, irrespectively of the stock level and the ages of the batches in stock. Especially on Mondays, given the diagonal shape of the elements in the frequency table, the fixed order quantity rule is of interest. In Figure 3.3(c), we illustrate that a fixed order quantity rule is read from the frequency tables by summing the frequencies along the upward diagonals. Each upward diagonal corresponds to a fixed production volume. The sum of the frequencies on the diagonal is the goodness-of-fit.

In the figure only three diagonals with non-zero figures can be drawn. A production volume of 5, 6 or 7 fits to respectively 16.7%, 19.8% and 63.5% of the states visited on the Mondays. Apparently, a fixed production volume of 7 batches is a reasonable approximation of the optimal strategy for Mondays. On the other days of the week, fixed production volumes are less appropriate: the best order quantities on Tuesdays – Fridays are optimal in only 25.3%, 32%, 26.9% respectively 31.2% of the states visited.

3.4.3 (Sub-optimal) Stock-age-dependent ordering policies

From the previous discussion we learn that the optimal ordering policy is stock-age-dependent. Although a bounded order-up-to S rule appears to fit closely to the optimal solution, we consider two stock-age-dependent ordering strategies. Especially the structure of the policy on Tuesdays and Wednesdays appears to be complicated. When many products are in stock, the age of the products appears to play an important role. We therefore propose two stock-age-dependent rules, for which stock is aggregated differently:

1. The *2D-order-up-to rule*, which has two order-up-to levels to acknowledge that a great part of the demand (strongly) prefers 'young' BPPs,

2. The *Final-stock rule* that aggregates the final stock minus any outdated BPPs, when average demand would happen over the day.

2D-order-up-to rule

The what-we-call 2D-order-up-to rule has two order-up-to levels:

- S_d an order-up-to level for the total stock level, and

- S_d^{Y} an order-up-to level for the 'young' stock.

The motivation for two order-up-to levels is that the rule should aim at having enough young BPPs in stock to meet the young-demand, next to the having enough BPPs in stock to meet the total demand. Therefore we propose a 2-Dimensional order-up-to S rule, which explains the name *2D-order-up-to rule* or in short *2D-rule*, by which actions follow from two (fixed) order-up-to levels according to:

$$a = \max\{(S_d - x)^+, S_d^{\mathrm{Y}} - x^{\mathrm{Y}}\}. \tag{3.11}$$

where x^{Y} is the total number of 'young' products in stock. When 'young' means having a residual shelf life of at least 3 days, then $x^{\mathrm{Y}} = \sum_{r=3}^{m} x_r$.

Search for good values $(S_d, S_d^{\mathbf{Y}})$ – The fixed order-up-to levels for the total demand, S_d, are read directly from Table 3.7, as before for the ordinary order-up-to S rule. To obtain good values for S_d^{Y}, the optimal order quantities in states $\mathbf{x}$ are studied.

In Table 3.8, a selection is given of the 183 stock states that have occurred on the 100,000 simulated Wednesdays. The table contains the optimal order quantity $\pi(3, \mathbf{x})$ in each state and the decision prescribed by the ordinary order-up-to S rule with order-up-to level 18 batches. In the last column we report the values of S_3^{Y} under which the 2D-rule prescribes the same order volume as the optimal order quantity. By definition of the 2D-rule the set of values for S_3^{Y} is empty when $S_3 - x$ exceeds the optimal order-quantity. When the optimal production decision corresponds to $S_3 - x$ than multiple values for S_3^{Y} fit. In Table 3.8 it seems that order-up-to level $S_3^{\mathrm{Y}} = 16$ is most frequent, but the state frequencies are not reported in the table and only a selection of states is considered.
In Table 3.9, the state frequencies on Wednesday are reported for the same set of states. For each promising value of S_3^{Y}, the state frequency is copied in the related column. For example in stock state $(3, 0, 0, 1, 8)$ the best choice for S_3^{Y} is 16 and this state occurred 1638 times in the simulation. State $(1, 0, 0, 1, 9)$ was visited 776 times during the simulation. In that state, the 2D-rule fits to the optimal strategy whenever $S_3^{\mathrm{Y}} \leq 17$, but not to $S_3^{\mathrm{Y}} = 18$. Therefore the state frequency 776 is found in all columns except in the last column.

Table 3.8: Illustration of reading 2D-rule for Wednesdays: sets of S_3^{Y} under I = (LIFO, FIFO), given $S_3 = 18$. ($x_2 = x_3 = 0$, since no production in weekends).

Stock $\mathbf{x}$								1D-rule	S_3^{Y} for which 2D-rule is optimal
x_1	x_2	x_3	x_4	x_5	x	x^{Y}	$\pi(3,\mathbf{x})$	$(18-x)^+$	$\{S_3^{\mathrm{Y}} : \max[(S_3-x)^+, S_3^{\mathrm{Y}}-x^{\mathrm{Y}}] = \pi(3,\mathbf{x})\}$
0	0	0	0	11	11	11	7	7	$S_3^{\mathrm{Y}} \leq 18$
0	0	0	4	11	15	15	2	3	∅
1	0	0	1	9	11	10	7	7	≤ 17
2	0	0	6	9	17	15	0	1	∅
3	0	0	1	8	12	9	7	6	16
3	0	0	2	8	13	10	6	5	16
4	0	0	0	8	12	8	7	6	15
4	0	0	3	7	14	10	6	4	16
5	0	0	0	8	13	8	7	5	15
5	0	0	1	8	14	9	7	4	16
5	0	0	4	8	17	12	4	1	16
10	0	0	3	8	21	11	5	0	16
...	...	...	...	...	...	...	...	...	...

Table 3.9: Reading the best S_3^{Y} from the frequency distribution under $I = (\text{LIFO}, \text{FIFO})$, given $S_3 = 18$ is fixed. ($x_2 = x_3 = 0$, since no production in weekends).

Stock **x**					Frequency 2D-rule yields optimal decision				
x_1	x_2	x_3	x_4	x_5	$S_3^{\text{Y}} \leq 14$	$S_3^{\text{Y}} = 15$	$S_3^{\text{Y}} = 16$	$S_3^{\text{Y}} = 17$	$S_3^{\text{Y}} = 18$
0	0	0	0	11	2442	2242	2242	2242	2242
0	0	0	4	11	-	-	-	-	-
1	0	0	1	9	776	776	776	776	-
2	0	0	6	9	-	-	-	-	-
3	0	0	1	8	-	-	1638	-	-
3	0	0	2	8	-	-	1706	-	-
4	0	0	0	8	-	829	-	-	-
4	0	0	3	7	-	-	575	-	-
5	0	0	0	8	-	442	-	-	-
5	0	0	1	8	-	-	788	-	-
5	0	0	4	8	-	-	34	-	-
10	0	0	3	8	-	-	1	-	-
...	...	...	...	...	...	...	...	...	...
				Total	74500	76924	91143	57234	40264
				Goodness-of-fit	75%	77%	91%	57%	40%

By summing all frequencies in a column, one computes the goodness-of-fit is of that value of S_3^{Y}. The table shows only a selection of states, but the next-to-last row reports over all states the frequency that a specific value of S_3^{Y} fits to the optimal MDP policy. On Wednesdays the most-frequent order-up-to level S_3^{Y} equals 16 batches. The goodness-of-fit of the 2D-rule on Wednesdays is 91%.

Similarly we read best order-up-to levels for the other weekdays. The resulting 2D-rule and the goodness-of-fit is reported in Table 3.10. On Thursdays all BPPs in stock are 'young' when $m = 5$, hence S_4^{Y} can be set equal to $S_4 = 16$ or lower.

Table 3.10: The 2D-rule under $I = (\text{LIFO}, \text{FIFO})$ with most-frequent order-up-to levels and their goodness-of-fit as read from the frequency tables.

2D-order-up-to S rule	Mon	Tue	Wed	Thu	Fri
Goodness-of-fit	61%	99%	91%	73%	97%
Total-Order-up-to level (**S**)	16	18	18	16	20
Young-Order-up-to level ($\mathbf{S}^{\text{Y}}$)	5	13	16	-[a]	12

[a] $S_4^{\text{Y}} \leq 16$, as all batches in stock on Thursday are 'young'.

Final-stock rule

A more refined approach to acknowledge the ageing of the stock is by looking one-day ahead. Let $\hat{y}(d, \mathbf{x})$ denote the final stock level, i.e. the next stock level before today's production is added and when average demand is met from. To compute $\hat{y}(d, \mathbf{x})$, one thus subtracts from $\mathbf{x}$ the (rounded) average demand volumes on day d and any batches that would outdate when average demand would happen during day d. The computation of $\hat{y}(d, \mathbf{x})$ thus requires a detailed registration of the outdating and the dispatching of products from stock.

For the *final-stock rule*, the production volume in stock state $\mathbf{x}$ on weekday d are easily computed from a, what-we-call, *target-stock level* F_d for weekday d:

$$a = (F_d - \hat{y}(d, \mathbf{x}))^+. \tag{3.12}$$

Under the final stock rule we thus aim at a fixed target level F_d for the number of products in stock at the start of the next morning. Therefore we may call the Final-stock rule a *target stock rule*. Since it anticipates outdating, it likely performs better than the ordinary order-up-to level, when outdating is an issue. A drawback of the final stock rule is that aggregating stock state $\mathbf{x}$ into a single figure $\hat{y}(d, \mathbf{x})$ easily results in error when done manually.

In Table 3.11, it is illustrated that a Final-stock rule fits better to the optimal MDP policy on Wednesday, than an ordinary order-up-to S rule. The goodness-of-fit of the Final-stock rule with target level $F_3 = 9$ on Wednesday is 91% and is much higher than the good ness-of-fit of 74.5% for the ordinary order-up-to S rule with order-up-to level $S_3 = 18$.

Similarly we generate frequency tables of the other weekdays, resulting in target levels for each working day: $\mathbf{F} = (9, 12, 9, 10, 13)$. On Thursday and Friday a target stock rule does not fit better than an ordinary order-up-to S rule, since outdating does not happen on Thursdays and Friday when the maximal shelf life is 5 days. Outdating happens primarily on Wednesdays: the BPPs produced on a Friday will expire at the end of next Wednesday.

Table 3.11: (State, action)-frequency tables for 100,000 simulated Wednesdays. A final stock rule fits better than an order-up-to rule.

(a) when the present stock **x** is aggregated in x.

Stock x / Up-to S_3	0...8	9	10	11	12	13	14	15	16	17	18	19	20	21	Freq(S_3)
26														1	1
25											4	2			6
24										28	9	1	1		39
23									91	62	21	4			178
22								458	224	79	12	1			774
21							1301	753	243	34	1				2332
20						3129	1880	697	84	17	6				5813
19					4745	3740	2284	679	140	13					11601
18				12888	25368	23761	11559	790	112	20					74498
17							695	3199	598						4492
16								85							85
15															0
:															:
1															0
0									98^a	81^a	2^a				181^a
Freq(x)	0	0	0	12888	30113	30630	17719	6661	1590	334	55	8	1	1	100000

(b) when the final stock as if average demand occurs is aggregated into $\hat{y}$.

Stock $\hat{y}$ / Up-to F_3	0	1	2	3	4	5	6	7	8	Freq(F_3)
9			18970	32038	27087	12373	672	3		91143
8		3674				972	3327	617		8590
7							86			86
6										0
:										:
1										0
0								102^b	79^b	181^b
Freq($\hat{y}$)	0	3674	18970	32038	27087	13345	4085	722	79	100000

[a] $S_3 = 0$ to emphasize zero production
[b] $F_3 = 0$ to emphasize zero production

3.4.4 Selection – simple and good fit

In Table 3.12 we give an overview of the several rules, their (nearly) optimal parameter values and the goodness-of-fit for each day of the week. On Mondays a bounded order-up-to S rule resembles best the optimal strategy, on Tuesday the 2D-rule fits best. On Wednesday the final stock rule fits equally well as the 2D rule. On Thursday and Friday all rules are equally well, except the fixed order quantity which fits very badly. Only on Mondays, the Fixed-order-quantity rule resembles better the optimal strategy than the order-up to rule. The best order-up-to S rule fits to only 54.5% of the simulated Mondays, whereas a fixed production volume of 7 batches fits for 63.5%.

All results presented in Tables 3.7 – 3.12 are under the (LIFO, FIFO) issuing policy. The goodness-of-fit will be different for different issuing policies, as the corresponding optimal MDP policy changes with the issuing policy.

Table 3.12: Goodness-of-fit of some basic production rules under $I = (\text{LIFO}, \text{FIFO})$ as read from the frequency tables.

Rule		Mon	Tue	Wed	Thu	Fri
Fixed order quantity rule						
Goodness-of-fit		64%	25%	32%	27%	31%
Fixed-order-quantity [a]	$\mathbf{Q}$	7	9	6	7	10
Order-up-to S rule						
Goodness-of-fit		55%	93%	75%	73%	97%
Fixed order-up-to level[a]	$\mathbf{S}$	16	18	18	16	20
Bounded order-up-to rule						
Goodness-of-fit		100%	93%	75%	73%	98%
Fixed level[a]	$\mathbf{S}$	16	18	18	16	20
Maximum production[a]	$\overline{\mathbf{a}}$	7	-	-	9	-
Minimum production[a]	$\underline{\mathbf{a}}$	5	6	3	0	7
2D-order-up-to S rule						
Goodness-of-fit		61%	99%	91%	73%	97%
Fixed order-up-to level (Total)[a]	$\mathbf{S}$	16	18	18	16	20
Fixed order-up-to level (Young)[a]	$\mathbf{S}^{\mathrm{Y}}$	5	13	16	-[b]	12
Target level next morning						
Goodness-of-fit		55%	94%	91%	73%	97%
Fixed target level (Total)[a]	$\mathbf{F}$	9	12	9	10	13

[a]in batches of 4 pools.
[b]$S_4^{\mathrm{Y}} \leq 16$, as all batches in stock on Thursday morning are 'young'.

An ordering rule should be practical and easy to use in order to get adopted. Therefore we will not select different types of rule for the different days of the week. The goodness-of-fit of a rule gives an indication how well a policy resembles the optimal MDP policy, but to assess the quality of a rule we are interested the detailed performance statistics.

3.5 Step 4 – Performance of rules versus MDP policy

In the previous step, we have constructed different rules and found good parameter values by investigating the structure of an optimal MDP policy. The quality of these rules is to be checked and to be compared to the optimal MDP strategy. Note that only for the scaled case an exact numerical analysis of both the simple rule and the optimal MDP policy is possible by analyzing the corresponding Markov chains.

In Table 3.13 annual figures for outdating, shortage, and quality mismatches are reported, as well as the expected annual cost under the given cost structure. Next to the annual costs, one reads in the last column how far the costs levels are from the optimal MDP cost level.

All rules, except the Fixed-order-quantity rule, perform nearly optimal with respect to outdating and shortages. The outdating is about 2% of the production and shortages occur less than 0.5%. The difference between the different types of order-up-to rules is visible in the occurrence of mismatching. Although a Fixed-order-quantity rule yields virtually no shortages and mismatches, its implementation is not recommended given the large number or batches that outdate.

Under the simple order-up-to S rule mismatching is 2.17 batches. By specifying bounds on the production volumes, the average annual mismatching is reduced from 2.17 batches to 0.59 batches per year. Under all rules less than 0.2% of young-demand receives batches older than the preferred shelf life. Mismatching happens thus rarely under the (LIFO, FIFO) issuing policy.

Table 3.13: Annual performance statistics of some basic ordering rules under $I = (\text{LIFO}, \text{FIFO})$ (numerical results from solving Markov chains).

Rule	Outdating [a]		Shortage [a]		Mismatch [b]		Annual costs	
MDP-optimal	37	1.95%	4.9	0.26%	0.09	0.01%	36,605	–
Fixed-order-quantity	157	8.41%	1.4	0.07%	0.00	0.00%	98,459	+168.9% [c]
Order-up-to S rule	36	1.95%	5.7	0.30%	2.17	0.17%	40,236	+9.9% [c]
Bounded order-up-to S	36	1.95%	5.6	0.30%	0.59	0.05%	39,077	+6.8% [c]
2D-order-up-to rule	36	1.92%	5.2	0.28%	2.16	0.17%	38,779	+5.9% [c]
Final stock rule	36	1.91%	5.1	0.27%	1.98	0.16%	38,389	+4.9% [c]

[a] in batches of 4 pools per year; % of total (1872 batches = 7488 pools),
[b] in batches of 4 pools per year; % of young-demand (1248 batches = 4992 pools),
[c] percentage above optimal cost level (MDP).

Considering the expected annual cost reported in the last (two) column(s) of Table 3.13, it appears that the simple order-up-to S rule is still 9.9% away from optimal. The bounded order-up-to S rule and the 2D-rule performs slightly better. The target stock rule is closest to the optimal annual cost level of 36,605 euros per year. Given the bad fit of the Fixed-order-quantity rule, it is not surprising that it gives extremely high annual costs.

Conclusions

Under a (LIFO, FIFO) issuing policy, the following conclusions are drawn for different ordering policies:

1. The simple order-up-to S rule, with order-up-to levels S_d that depend on the day of the week, results in an average cost level that is not far from the optimal cost level,
2. Bounding the production volumes improves the order-up-to S rule by reducing the occurrence of mismatching,
3. The 2D-rule reduces mismatching, since it has order-up-to levels for the 'young' stock, next to order-up-to levels for the total stock.
4. Even better results are obtained by Final-stock rule, which aggregates the stock that survives until the next day as if average demand happens during the day.

All rules of the order-up-to type perform almost equally well with outdating of about 2% and shortages less than 0.5%. The difference between the rules may change when another data set of the PPP is used. For practical use we propose to adopt the ordinary order-up-to S rule with only one parameter per day of the week. This rule is subject of further study to justify its use for solving the original unscaled PPP.

3.6 Step 5 – Re-scaling and validation of simple rule

In the last step of our approach, we validate the application of an ordinary order-up-to S rule to problems of realistic size. We return to the original, unscaled PPP, in which both demand and production decisions are in BPPs instead of batches. We consequently discuss:

1. a simulation model by which detailed performance statistics are obtained,
2. the disaggregation of the production volumes by re-scaling the order-up-to levels,
3. the impact of distinguishing different partly compatible blood groups instead of a single universal product, and
4. the robustness of the approach regarding the supply of buffy coats.

In Section 3.7 we report more sensitivity studies with respect to the problem data and we provide answers to several *what-if* questions.

3.6.1 Simulation model and data

Since exact results cannot be obtained for the full size problem, all results reported for the full size problem are obtained through simulation. The simulation model relies on the same assumptions as the MDP model.

Hence the simulation runs in discrete time with fixed time increments of 1 day. The following actions and events happening during the day in the chronological order:

1. the stock is observed at the very start of a day,
2. immediately thereafter the production volume is set and expected costs are incurred,
3. during the day BPPs are selected from stock and distributed to the hospitals to meet the demand: all young-demand is met prior to the any-demand,
4. at the end of the day the ageing of the BPPs left in stock is registered: any outdated pools are removed from stock
5. finally, today's production (if any) is added to stock as BPPs with a residual shelf life of m days.

In a first model, called the *Universal-group simulation model*, the existence of blood groups is not included and the supply by donors is not restrictive. In Section 3.6.3, this model is extended with blood groups and stochastic show-up of donors of different blood groups. We call the latter model the *Multi-group simulation model*.

Data – The unscaled data reported in Section 3.3.1 are used. Most notably, given probability distribution are used or these are fitted on the mean and the standard deviation. We assume, for the moment, that demand is modeled by a Poisson distribution with unscaled means as reported in Table 3.1.

Simulation horizon – To obtain accurate estimates of the several statistics that we are interested in all simulations cover a period of 100 million weeks. According to the following exercise, this is enough to get estimate that are accurate to 3 digits.

Justification of simulation length:

Let X be the number of BPPs that outdate in a week. Presume, X is Poisson distributed with mean μ outdates per week and standard deviation $\sqrt{\mu}$. Through simulation of n independent weeks, one may obtain n independent estimates of μ: $X_1, X_2, \ldots, X_n$. The 95% confidence interval for μ has half-width $2\sqrt{\mu/n}$. The interval has relative half-width $\frac{2\sqrt{\mu/n}}{\mu} = \frac{2}{\sqrt{\mu\, n}}$. The estimate $\mu \approx \frac{1}{n}\sum_{i=1}^{n} X_i$ is with 95% confidence accurate up to 3 digits, when $\frac{2}{\sqrt{\mu\, n}} < 10^{-3}$. Hence n must be at least $4\,10^6/\mu$. When outdating is in the order of 2 per week, then n must be at least $2\,10^6$.

Events that occur less frequent, require a longer simulation horizon, e.g. when μ is in the order of 0.04 per week then a simulation for 100 million is just enough to get estimates that are with 95% confidence accurate up to three digits. Ignoring any dependencies, we assume most of the estimates reported in the next sections for the full-size problem to be accurate up to 3 digits when simulating 100 million weeks.

3.6.2 Re-scaling – disaggregation of batches into BPPs

In the original unscaled or full-size problem the demand distribution is in pools. But in the previous sections the demand and the production happens in batches of four pools. Hence the order-up-to S rule and the corresponding order-up-to levels are also in batches. Below, we discuss two ways to use the frequency tables of the optimal scaled MDP solution in order to estimate the order-up-to levels for the unscaled problem.

Re-scaling the most-frequent order-up-to levels – Re-scaled parameter values of the order-up-to S rule are obtained by multiplying the most-frequent order-up-to levels as read from the frequency tables by the factor $B = 4$. Reasonable order-up-to levels for Monday to Friday for the original PPP of real size are thus: $4 \cdot (16, 18, 18, 16, 20) = (64, 72, 72, 64, 80)$.

Re-scaling average order-up-to levels – Instead of re-scaling the most-frequent order-up-to levels, one could re-scale average order-up-to levels. Average order-up-to levels for the scaled PPP are obtained from the frequency table by multiplying the first and the last columns of the tables in Table 3.7, and divide the result by the simulation length (100,000 weeks). For example the re-scaled average order-up-to level for Thursday morning is computed as follows:

$$S_4 = 4 \cdot \frac{17 \cdot 445 + 16 \cdot 73,213 + 15 \cdot 26,263 + 14 \cdot 3 + 0 \cdot 76}{100,000} \approx 63.$$

After re-scaling and rounding the average order-up-to levels, the following set of order-up-to levels is found: $(62, 72, 73, 63, 80)$. These values are close to the previously found levels.

Local search

To check whether re-scaling indeed results in nearly-optimal parameter values for the full size problem, the parameter values are given as input to a simulation-based search procedure. We have constructed an incremental search procedure for this purpose, which is presented Appendix C. The algorithm start with some set of order-up-to levels, e.g. with all $S_d = 0$, and successively determines which order-up-to levels to increment by one. The incremental search stops when improvements stay out for a number of iterations. Alternatively, one may start with very high levels and successively decreases the best levels. As the algorithm may not to all data sets, we simply use the algorithm to verify the re-scaled order-up-to levels.

Since the order-up-to S rule has five parameters there are many local optima at which the incremental search procedure might get stuck. The local search procedure is only used as a verification tool that is not a crucial part of the SDP-Simulation approach. The locally optimal order-up-to levels found by the algorithm are $(65, 74, 80, 64, 82)$ for Monday to Friday.

Evaluation of different sets of order-up-to levels

In Table 3.14 the performance of the order-up-to S rule for different order-up-to levels in the unscaled problem is reported.

- First of all, we observe that the order-up-to S rule is quite robust with respect to the chosen order-up-to levels.
- Secondly, we observe that the relative outdating, shortage and quality mismatch statistics are very close to the scaled case. Not even 2% of the production becomes outdated and less than a half percent of the demand is mismatched or cannot be delivered directly from stock.
- Thirdly, it appears that the annual costs differ from the annual costs for the scaled problem, which were reported in Table 3.12. The discrepancy could be due to the fact that the scaled demand distributions are not fitted on third or higher moments.

Simply re-scaling the most-frequent order-up-to levels suffices to find a nearly optimal order-up-to rule for the full-size PPP. Fine-tuning the order-up-to levels gives only a slight reduction of the cost level, but hardly influences the occurrences of outdating and shortages. As the simulation-based search approach may be time-consuming and could run into a bad local optimum, the SDP-Simulation approach is preferred as it also gives insight in the optimal policy.

Table 3.14: Performance of the order-up-to S rule under different order-up-to levels S. (simulation of 100 million weeks).

Re-scaled Order-up-to levels	Outdating [a]		Shortage [a]		Quality mismatch [b]		Mean age	Annual costs
Most-frequent levels $S = (64, 72, 72, 64, 80)$	142	1.9%	33	0.44%	9	0.18%	2.04	47,726
Average levels $S = (62, 72, 73, 63, 80)$	137	1.8%	33	0.44%	14	0.28%	2.03	47,677
Optimized levels $S = (65, 74, 80, 64, 82)$	191	2.5%	19	0.26%	11	0.21%	2.11	45,205

[a]in pools per year; % of total (7488 pools),
[b]in pools per year; % of young-demand (4992 pools).

3.6.3 Disaggregation of universal BPPs into blood groups

Thus far we ignored the eight different blood groups and the stochastic supply of buffy coats. In this section we show that the aggregation of the BPPs of different blood groups into a single universal product BPP is justified, at least for the Dutch case. Remember from the introduction of the PPP in Section 2.1, that in The Netherlands about a third to a half of the buffy coats is processed to BPPs.

The current practice

The donor management team calls donors by phone or mail to make an appointment for donating whole blood. As we have learned in Section 2.1 one prefers donors with the universal blood group O^-. Therefore one tends to call relatively many O^- donors, but O^- donors are scarce and donors normally donate only twice a year (on a voluntary basis). In Figure 3.4 the compatibility of the blood groups is shown once again.

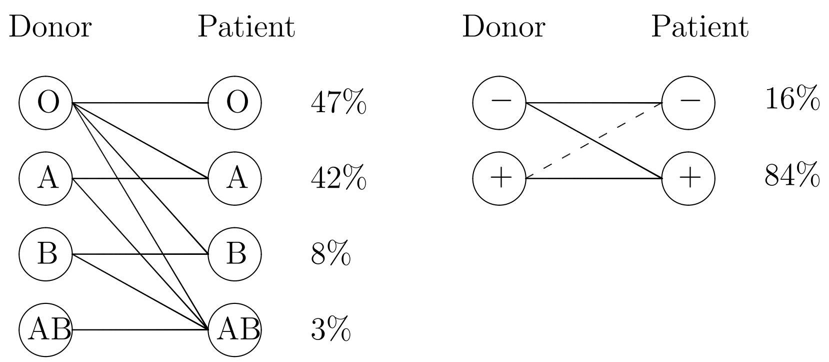

Figure 3.4: Compatibility of blood groups of donor and patient following the ABO-system and Rhesus-D system (+/−)

After a first processing step to separate red blood cells from a whole blood donation, the residual called the buffy coat is processed. To meet the demand for BPPs, on average only a third of the buffy coats needs to be processed. Hence, one may select the most favorite (=most-compatible) buffy coats for processing BPPs. Blood banks primarily produce BPPs from donor groups O and A, since compatible matching is the standard. In Figure 3.4, one reads that about 90% of the Western population has either blood group O or A. Patients of the B and AB type can be helped with O pools. Hence virtually all demand is met by pools of blood group O and A. With respect to the Rhesus-D factor we note that not even 50% of the RhD^- patients requires strict matching. This is only 8% of the patient population.

The Multi-group simulation model and data

In the what-we-call *Multi-group simulation model*, the number of donors of each of the eight different blood groups that show-up, is drawn from eight different Poisson distributions. Similarly eight different Poisson distribution are distinguished for each type of demand.

The model is an extension of the Universal-group simulation model. Below we explain in more detail, how is dealt with the distinction of blood groups.

Donor and patient population – We assume that the (relative) frequencies of the blood groups for the donor and patient population correspond to the Western population. Hence we do not over-represent the more compatible donor groups.

Donations – The number of donors to call is usually set primarily by the demand for RBCs. We assume that the average number of appointments made by the donor management team is the same from Monday to Thursday, but on average on Friday 10% more donors are scheduled to anticipate the production stop during the weekend. One may not always succeed in making enough appointments and some appointments are canceled by the donors on the day of donation. Modeling the number of whole-blood donations by a Poisson distribution may overestimate the uncertainty in the show-up of donors we accept to stay at the safe side.

The average weekly demand for 144 BPPs requires 720 buffy coats. When on average 2160 whole blood donations are taken per week, on average only a third of the buffy coats is used for processing BPPs. The mean number of donations on Monday – Thursday is thus $\frac{2160}{1+1+1+1+1.1} = 423.5$. On Friday the mean is $1.1 \cdot 423.5 \approx 466$ donations. We assume that no blood groups are over-represented compared to in the distribution of the blood groups in most Western countries. Hence, on Friday the number of whole blood donations of blood group O^- is Poisson distributed with mean $47\% \cdot 16\% \cdot 466 = 35$ donations. Thus on average only 7 pools of this blood group can be produced on Friday.

Production – The required aggregated production volume suggested by the order-up-to S rule is reached by producing as many pools as possible of the *most-compatible donations*. With most-compatible donation, we mean donations of a blood group compatible to the greatest part of the patient population. O^- donors can help all patient groups, O^+ is compatible to $84 + 0.5 \cdot 16 = 92\%$ of the patient population, etc. The order of most-compatible to least-compatible donation is O^-, O^+, A^-, A^+, etc.

Order of patient groups – Before distributing the BPPs to the hospitals, the stock is matched to the demand. One has some freedom in setting the order in which the different patient groups are considered. We meet the young-demand prior to meeting the any-demand. Upon meeting the demand for each of these groups, the demand set by the *least-compatible patients* is met first. A least-compatible patient has a blood group that is compatible to only a small fraction of the donations. O^- patients can receive only O^- which is offered by $47\% \cdot 16\% \approx 7.5\%$ of the donors, B^- patients rely on O^- and B^- donors, which is $(47\% + 8\%) \cdot 16\% = 8.8\%$ of the donor population. The order from least-compatible to most-compatible patient is: O^-, B^-, A^-, AB^-, O^+, B^+, A^+, AB^+. One could say the a least-compatible patient is least flexible in receiving BPPs. Since demand is met in this order, AB^+ patients seem to have the highest chance of experiencing a shortage. But a shortage to a AB^+ patients is most easily resolved, since AB^+ patients may receive BPPs of any blood group.

Detailed issuing policy – Now that we have specified an order in which demand is modeled, we need to specify the order in which BPPs are selected. In matching the stock to the demand, one tries to issue the least-compatible BPPs first. Hence in principle stock is issued in the order AB^+, AB^-, B^+, B^-, A^+, A^-, O^+, O^-, etc., provided that the BPPs are compatible with the blood group of the patient. If the issuing policy is (pure) FIFO, then first BPPs of the oldest categories are issued in the order of least-compatible blood group. Next younger pools are issued. Under LIFO the youngest pools in stock that are compatible to the patient's blood group, are issued before older pools are considered. In the PPP we deal with two demand categories; the young-demand receives the youngest compatible pools in stock (LIFO), the any-demand is met by the oldest compatible BPPs in stock (FIFO).

Through the above refined issuing policy, we expect to keep the occurrence of shortages and outdating low. We do assume that all demand is known upon distributing the pools to the hospitals, although in practice a portion of the demand occurs after the distribution early in the morning. Sharing this information with the production department is beneficial, but we consider a more difficult case where the production department has no information about the demand other than probability distributions obtained from historical data as in current practice.

In the remainder, we discuss the results obtained from the Multi-group simulation model. Besides showing how the production, outdating and shortages are spread over the blood groups, we will consider cases where buffy coats are more scarce.

Results – impact of distinction of blood groups

We study the impact of the distinction of blood groups by comparing the results obtained from the Multi-group simulation model against those obtained through the Universal-group simulation model. The results as presented in Table 3.15 are estimates accurate up to 2 or 3 digits. The annual costs in the Multi-group model are only 1% higher than in the Universal-group model. The differences in shortages, outdating and quality mismatches are virtually zero. It thus seems justified to ignore the blood groups in the study of nearly optimal production rules for the Dutch case.

Table 3.15: Annual performance of the order-up-to S rule under $I = (\text{LIFO}, \text{FIFO})$: the impact of blood groups in the Multi-group model (simulation for 100 million weeks).

	Multi-group model		Universal-group model	
Criterion	in pools	relative %	in pools	relative %
Production	7,597		7,597	
Outdating	143	1.9%	142	1.9%
Shortage	33	0.4%	33	0.4%
Quality mismatch	9	0.2%	9	0.2%
Annual costs	48,168		47,726	

Results – production, outdating and shortages per blood group

Although blood groups are not relevant when setting the (aggregated) production volume, it is of interest to study whether certain patient groups face an excessive shortage rate under the re-scaled order-up-to S rule and the (refined) $(\text{LIFO}, \text{FIFO})$ issuing policy. Table 3.16(a) reports on the annual production, outdating and shortages of pools of the eight different blood groups. Table 3.16(b) presents the relative outdating, shortage and quality mismatch figures. We conclude that when on average only a third of the buffy coats is actually used for the production of BPPs, virtually all BPPs are of blood group O. Since O^+ and O^- suits to 92%, respectively 100% of the patients, shortages and mismatches are very rare. The outdating percentage is with 5.5% the highest for the O^- pools, since BPPs of the universal blood group are issued only when no pools of the same age of other blood groups are in stock. Overall outdating is very low: only 1.9% overall blood groups.

Table 3.16: Production, outdating and shortages per blood group under the order-up-to S rule with re-scaled most-frequent order-up-to levels (over 100 simulation runs of 1 million weeks each) when on average a *third* of the buffy coats is processed into BPPs.

(a) Absolute volumes in pools per year

	Total	O^-	O^+	A^-	A^+	B^-	B^+	AB^-	AB^+
Production	7597	1582	5944	62	10	0	0	0	0
Outdating	142.8	86.5	56.3	0	0	0	0	0	0
Shortage	33.2	0.6	7.8	0.8	19.3	0.1	2.2	0.1	2.4
Quality mismatch	9.4	0	0.1	0.2	7.2	0	0.1	0	1.7

(b) Relative volumes

	Total	O^-	O^+	A^-	A^+	B^-	B^+	AB^-	AB^+
Outdating	1.9%	5.5%	0.9%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Shortage	0.4%	0.2%	0.2%	0.3%	0.7%	0.3%	0.4%	0.4%	1.1%
Quality mismatch	0.2%	0.0%	0.0%	0.1%	0.4%	0.0%	0.0%	0.2%	1.2%

Overall shortages and quality mismatches remain well below 0.5%. The AB^+ patients, who constitute less than 3% of the entire population, have the highest probability of experiencing a shortage (1.1%) or a mismatch (1.2%), since the AB^+ demand is met after all other demand is fulfilled.

3.6.4 Sensitivity – scarcity of buffy coats

When on average less donors show up, one uses relatively more of the buffy coats obtained after processing RBCs. As a consequence of the limited availability of O donors, one must use buffy coats of the less compatible blood groups (occasionally one may even not be able to produce the desired number of pools). In the last two rows of Table 3.17 we show that when on average two third of the buffy coats is used instead of only a third, a significant part of the production is of blood group A and on average 108 of the 7597 BPPs are of blood group B or AB.

The scarcity of buffy coats of group O affects the occurrence of shortages and outdating. In Table 3.18 we show that even when on average two third of the buffy coats are used, overall outdating, shortages and mismatches are still very low. Patients from blood group O^- and B^- experience a high shortage probability (2.5% respectively 4.4%), because on

Table 3.17: Production per blood group under the re-scaled order-up-to S rule (over 100 simulation runs of 1 million weeks each) when on average a third, a half or two third of the buffy coats is processed into BPPs.

Scenario	Total	O^-	O^+	A^-	A^+	B^-	B^+	AB^-	AB^+
33% of Bc used									
Annual production	7597	1582	5944	62	10	0	0	0	0
% of total production		20.8%	78.2%	0.8%	0.1%	0%	0%	0%	0%
50% of Bc used									
Annual production	7597	1018	5268	525	786	0	1	0	0
% of total production		13.4%	69.3%	6.9%	10.3%	0%	0%	0%	0
67% of Bc used									
Annual production	7597	739	4201	560	1991	8	86	0	14
% of total production		9.7%	55.3%	7.4%	26.2%	0.1%	1.1%	0.0%	0.2%

average only 9.8% of the produced BPPs is compatible to O^- or B^- patients. Relative outdating figures are high for BPPs of the less favorite blood groups B and AB, but the absolute volume is very low since BPPs of these blood groups are rarely produced.

Table 3.18: Production, outdating and shortages per blood group under the order-up-to S rule with re-scaled most-frequent order-up-to levels (over 100 simulation runs of 1 million weeks each) when on average *two thirds* of the buffy coats (Bc) is processed to BPPs.

(a) Absolute outdating, shortage and mismatching volumes in pools per year

	Total	O^-	O^+	A^-	A^+	B^-	B^+	AB^-	AB^+
Production	7597	739	4201	560	1991	8.0	86	0.0	14
Outdating	175	30.3	52.7	34.4	48.2	0.8	5.0	0.0	3.6
Shortage	64.7	7.1	27.0	1.8	15.4	2.1	9.3	0.2	1.9
Quality mismatch	10.9	1.2	0.4	0.5	5.7	0.6	0.9	0.1	1.4

(b) Relative outdating, shortage and mismatching

	Total	O^-	O^+	A^-	A^+	B^-	B^+	AB^-	AB^+
Outdating	2.3%	4.1%	1.3%	6.1%	2.4%	9.4%	5.8%	37.2%	26.8%
Shortage	0.9%	2.5%	0.8%	0.7%	0.5%	4.4%	1.7%	1.1%	0.9%
Quality mismatch	0.2%	0.6%	0.0%	0.3%	0.3%	1.9%	0.2%	0.7%	1.1%

Remark – We believe that it is not of interest to study under the same issuing and ordering policy a scenario where the fraction of buffy coats used is close to 1. When buffy coats are very scarce, one is not in the luxurious position to issue the youngest BPPs in stock to meet the young-demand. Furthermore, more advanced ordering strategies might perform better in that case. In the extreme case where on average all buffy coats are used to meet the demand, one may expect to process virtually all buffy coats into BPPs to anticipate the uncertain show-up of donors, causing that the desired production volumes are often not realized. We will not study this extreme case.

3.6.5 Conclusions for full-size PPP

The SDP-simulation approach is justified to solve the PPP at a Dutch blood bank. From a simulation study using realistic data we conclude:

1. The shortage, outdating and mismatch figures for the scaled PPP are a good approximation for the full-sized PPP,
2. An ordinary order-up-to S rule performs very well, even when patients and BPPs are of different blood groups,
3. Re-scaling the order-up-to levels yields nearly optimal order-up-to levels,
4. The order-up-to S rule appears to be robust with respect to minor changes in the order-up-to levels,

Consequently, when searching for good order-up-to levels, one can ignore the existence of blood groups and the uncertainty at the supply side. This observation is of practical importance, since it simplifies the search for good order-up-to levels and allows a fast evaluation of different scenarios by the Universal-group simulation model. In the next section we extend the sensitivity study by changing the problem data and we answer some prominent *what-if* questions.

3.7 Sensitivity analysis for the PPP

In the previous sections we have shown that an ordinary order-up-to S rule resembles the optimal ordering strategy for the PPP under consideration. The SDP-simulation approach yields after re-scaling nearly optimal order-up-to levels. In this section we put this rule and our approach for reading the best parameter values to the test. As argued in the previous section the study is done using the 'universal-group' simulation model. We are interested in the robustness and performance under various scenarios, we therefore answer the following questions:

- How robust is the rule when the estimates are far off from the realized mean demand?
- What if (part of) the demand is even more uncertain than Poisson?
- What if shortage costs are set lower or higher than in the base case?
- What if the maximal shelf life is extended to 6 or even 7 days?

The first two questions are interesting scenarios from a practical point of view: no matter how careful the average demand is predicted, the realized demand over a fixed period can be somewhat off. Further the demand may be somewhat more uncertain than Poisson distributed demands.

The latter two questions are more controlled. The cost structure is a matter of choice, that should resemble the preference of the stake holders: the blood bank managers and the doctors at the hospitals. When regulations by law and clinical test allow doing so, the maximal shelf life is in the current practice increased from 5 to 7 days.

We refer back to the previous section for a discussion of the impact of the issuing rule, the distinction of blood groups and the scarcity of donors.

3.7.1 Average demand estimates

Input to the optimization model are estimates for the mean demands, amongst other. These estimates might be a few percent off from what turns out to be the mean over the coming period, of 13 weeks (one quarter). When in the demand turns out to be higher than expected, shortages are more likely. When the average demand is below the expected value, then more BPPs become outdated.

Table 3.19: Performance of the order-up-to S rule with fixed levels $S = (64, 72, 72, 64, 80)$, when the average demand is on all weekdays 5 or 10% above (+) or below (−) the predicted mean demand. (simulation of 100 million weeks).

% off from estimate	Outdating [a]		Shortage [a]		Quality mismatch [b]		Mean age	Annual costs
−10% off	388	5.4%	6.0	0.09%	12.4	0.28%	2.19	65,127
−5% off	247	3.4%	14.1	0.20%	10.5	0.22%	2.12	49,820
0% off	142	1.9%	33	0.44%	9.2	0.18%	2.04	47,726
+5% off	73	0.93%	70	0.89%	8.1	0.16%	1.94	64,856
+10% off	34	0.41%	134	1.6%	7.2	0.13%	1.84	107,052

[a]in pools per year; % of total (7488 pools),
[b]in pools per year; % of young-demand (4992 pools).

In the unscaled case, the mean demand is Poisson distributed with mean 32 BPPs. The expected demand over 13 Wednesdays is thus $13 \cdot 32 = 416$ BPPs. The corresponding standard deviation is $\sqrt{416} \approx 20$ BPPs. A 67%-confidence interval of the demand over any 13 Wednesdays, is $[416 - 20, 416 + 20] = [396, 436]$. A 95%-confidence interval of the demand is $[376, 456]$. When one considers an arbitrary quarter, then with 2.5% chance the average demand is more than 2 standard deviation, or $\frac{40}{416} \approx 10\%$, higher than the true mean 416 BPPs. With probability $\frac{100\%-68\%}{2} = 16\%$, the demand over 13 weeks is $\frac{20}{416} \approx 5\%$ above the expected demand of 416 BPPs.

In Table 3.19, we consider a number of scenarios regarding the difference between the expected demand volumes and the actual demand volumes. Differences may arise by nature: not all periods of 13 weeks are the same, or they may be structural when trends are broken. Suppose the demand on all weekdays is 5% or even 10% below the expected mean demand, then the order-up-to levels are set too high. By simulation we check whether outdating then becomes a big issue. For the case where the mean demand on all weekdays is structurally 5% or 10% higher, we check whether this results in many shortages. Furthermore, we expect that the age of the pools drops when the demand is higher than expected.

The robustness of the order-up-to S rule, is checked for five scenarios: the average demands for each day of the week is consequently −10%, −5%, 0%, +5% and +10% off from the estimated mean demands. The order-up-to levels are fixed to $\mathbf{S} = (64, 72, 72, 64, 80)$. To show the robustness of the approach, we assume all deviations to be in the same direction on all weekdays.

The results show that deviations up to +5% or −5% do not cause many shortages and outdating: shortages are still less than 1% and average outdating ranges from 0.93 to 3.4 percent. In the extreme case where the mean demands deviate +10% or −10%, the outdating and shortages are still moderate compared to the current practice. On the one hand, when the demand is 5% lower than expected, the average number of outdated pools roughly doubles. On the other hand, the average number of shortages roughly doubles when demand on all days is 5% higher than expected. The shortages rate is 1.6%, when all demand is 10% higher than expected. As shortages are more critical than the outdating of BPPs, overestimating the demand is less critical than underestimation.

3.7.2 Demand uncertainty: higher coefficient of variation

In this section we investigate the sensitivity of an order-up-to S rule with respect to the uncertain in the demand. We increase the uncertainty in the any-demand; the young-demand is still Poisson distributed. Under Poisson demand the coefficient of variation, cv_d^{any}, of the any-demand on day d, would be $\frac{1}{\sqrt{\mu_d}}$ for weekday d, where μ_d is the average demand at day d. Now we set $cv_d^{\text{any}} = \frac{2}{\sqrt{\mu_d}}$. We expect to find higher order-up-to levels, such that the occurrence of shortages is kept low at the expense of a few more outdated BPPs.

By simulations of 100 million weeks, we evaluate the ordinary order-up-to S rule for three sets of order-up-to levels:

$\mathbf{S} = (64, 72, 72, 64, 80)$ as obtained by solving the scaled PPP with Poisson demand ($cv_d^{\text{any}} = \frac{1}{\sqrt{\mu_d}}$), and re-scaling the most-frequent order-up-to levels,

$\mathbf{S} = (72, 80, 80, 72, 88)$ as obtained by solving the scaled PPP with $cv_d^{\text{any}} = \frac{2}{\sqrt{\mu_d}}$, and re-scaling the most-frequent order-up-to levels,

$\mathbf{S} = (69, 79, 86, 67, 87)$ the locally optimal order-up-to levels obtained by local search for the unscaled PPP with $cv_d^{\text{any}} = \frac{2}{\sqrt{\mu_d}}$.

The second set of order-up-to levels are 8 pools higher than when demand was less uncertain as in the Poisson demands case. The third set is introduced to show that there is no need to fine-tune the re-scaled levels of the second set, as will become clear from Table 3.20.

In the table one finds the performance of the order-up-to S rule for the given choices of the order-up-to levels. As expected, underestimation of the degree of uncertainty results in

Table 3.20: Performance of the order-up-to S rule under different order-up-to levels S for the case where cv^{any} is doubled and I =(LIFO, FIFO). (simulation of 100 million weeks).

Order-up-to levels	Outdating [a]		Shortage [a]		Quality mismatch [b]		Mean age	Annual costs
Re-scaled, Poisson demand								
$S = (64, 72, 72, 64, 80)$	264	3.5%	129	1.7%	26	0.52%	1.98	141,686
Re-scaled, $cv_d^{\text{any}} = \frac{2}{\sqrt{\mu_d}}$								
$S = (72, 80, 80, 72, 88)$	524	6.6%	50	0.66%	18	0.37%	2.12	119,474
LS, double cv_d^{any}								
$S = (69, 79, 86, 67, 87)$	472	6.0%	53	0.70%	28	0.56%	2.12	116,002

[a]in pools per year; % of total (7488 pools),
[b]in pools per year; % of young-demand (4992 pools).

relatively many shortages. Using accurate estimates of the coefficient of variation results in lower shortages but higher outdating, in the order of 6%-6.6%. The local search reduces the outdating a bit, but this at the expense of more shortages and mismatching. To keep the number of outdated pools low, the order-up-to levels should be kept low, e.g. to those for the Poisson demand case, but than shortages become more prevalent and consequently the average cost level increases.

Return to Step 2 and 3 – investigate structure of optimal strategy

Fine-tuning the order-up-to levels improves the order-up-to S rule only slightly. This might raise the question whether the order-up-to S rule is far from optimal when demand is more uncertain and whether other rules are much better. This question can only be answered by returning to step 2 and 3 of our approach: we investigate the optimal strategy for the scaled PPP with $cv_d^{\text{any}} = \frac{2}{\mu_d}$.

Under the optimal MDP strategy, 7.1% of the produced batches become outdated and shortages and mismatches are on average less than 0.1%. The optimal production volumes depend heavily on the age of the pools in stock. In Table 3.21 we report the goodness-of-fit of some simple rules. These goodness-of-fit which are read from simulation-based frequency tables, which we do not report here.

We conclude that an order-up-to S rule fits in many states visited, but a bounded order-up-to S rule seems to fit much better. Due to the increased uncertainty in the demand, the goodness-of-fit figures are much lower than under Poisson demand. We do not elaborate on setting a better rule for this particular case, but we stress that the SDP-simulation

Table 3.21: Goodness-of-fit of some production rules as read from frequency tables for $I = (\text{LIFO}, \text{FIFO})$ and the any demand (=FIFO demand) is twice as uncertain.

Rule	Mon	Tue	Wed	Thu	Fri
Fixed production rule					
Goodness-of-fit	34%	32%	30%	25%	27%
Fixed production volume [a]	7	9	6	7	10
Order-up-to S rule					
Goodness-of-fit	34%	66%	41%	77%	59%
Fixed order-up-to level[a]	18	20	20	18	22
Bounded order-up-to rule					
Goodness-of-fit	64%	72%	48%	77%	68%
Fixed level[a]	18	20	20	18	22
Maximum production[a]	8	11	9	10	15
Minimum production[a]	6	7	5	0	9

[a]in batches of 4 pools.

approach can be used to find a good ordering rule that is a good compromise between quality and simplicity.

3.7.3 Cost ratios

In order to change the trade-off between outdating and shortages, we show what is the impact of major changes in the shortage costs. The outdating costs are fixed to 150 per outdated pool. Although the mismatch cost is also a kind of shortage cost, we keep it fixed to 200 per pool.

In the base case we have taken shortage costs equal to 5 times the outdating costs: a penalty of 750 euro applies for each pool demanded that cannot be met immediately from stock. In this section we show the impact of lowering and raising the shortage costs on the performance statistics.

In Table 3.22 we report results of a long simulation run of the order-up-to S rule with re-scaled order-up-to levels obtained from the scaled MDP which is solved for shortage costs set to 300, 750, and 1500 euros respectively. When shortage costs are high, one tends to keep many pools in stock. Consequently the available stock is old more BPPs outdate.

Table 3.22: Performance of the order-up-to S rule with re-scaled order-up-to levels for increasing shortage costs. (simulation of 100 million weeks).

Order-up-to levels	Outdating [a]		Shortage [a]		Quality mismatch [b]		Mean age
Shortage cost = 300							
$S = (60, 72, 72, 60, 80)$	127	1.7%	37	0.49%	23	0.46%	2.00
Shortage cost = 750							
$S = (64, 72, 72, 64, 80)$	142	1.9%	33	0.44%	9.2	0.18%	2.04
Shortage cost = 1500							
$S = (64, 72, 72, 64, 84)$	226	2.9%	19	0.26%	26	0.53%	2.05

[a]in pools per year; % of total (7488 pools),
[b]in pools per year; % of young-demand (4992 pools).

Remarkably, compared to the base case, where shortage costs are 750 euro per BPP short, mismatches happen more frequently both when the shortage cost are set higher or lower than 750. When shortage costs are higher, mismatching is higher since the pools in stock are on average older. When shortage costs are low the stock levels are set lower, which makes both shortages and mismatching more likely.

Other changes in the cost-structure are possible. For example, we could study more refined penalties on mismatching. Instead of fixing the mismatch costs to 200 euro per mismatched BPP irrespectively of the age of the BPP. More refined cost incentives for issuing 'young' BPPs can be modeled but are not considered throughout this thesis. Fixed order costs are studied in the next chapter.

3.7.4 Shelf life m

Blood banks are allowed to raise the maximal shelf life of the platelet pools, when clinical studies shows that the quality of a BPP upon transfusion is guaranteed to be high enough[1]. In some countries the shelf life is effectively only 4 days, whereas in most other countries the shelf life is 5 days and current efforts are focused on raising the shelf life to 6 or even 7 days. In this section we will study the importance of increasing the shelf life from an inventory management point of view.

[1]The quality of a pool (of 5 donor units) is sufficient when it contains more than 250 10^9 active platelets. Upon production immediately after donation this number is normally around 350 10^9, but decreases over time.

Table 3.23: The most-frequent order-up-to levels for the scaled case at varying maximal shelf life and young-demand prefers a residual shelf life of at least 3 days.

Maximal shelf life	Order-up-to levels as read from scaled MDP	Re-scaled order-up-to levels
$m = 4$	$(15, 17, 17, 16, 20)$	$(60, 68, 68, 64, 80)$
$m = 5$	$(16, 18, 18, 16, 20)$	$(64, 72, 72, 64, 80)$
$m = 6$	$(16, 18, 18, 16, 21)$	$(64, 72, 72, 64, 84)$
$m = 7$	$(16, 18, 18, 16, 22)$	$(64, 72, 72, 64, 88)$

When the maximal shelf life m is increased, one should redefine the 'young' demand. With $m = 5$ all young-demand strongly prefers pools with a residual shelf life of at least $l = 3$ days. The young-demand is thus mismatched when BPPs with a residual shelf life of 2 days or less are issued. We stick to this definition: independent of the maximal shelf life m, the age preference for the young-demand is a residual shelf life of $l = 3$ days. Alternative definitions are considered but not reported here. Mismatch costs are kept to 200 euro per mismatched pool (or 800 per batch of 4 pools).

Under this definition all young-demand on Monday will be mismatched when the maximal shelf life is only 4 days.

By executing the first four steps of our approach, for maximal shelf life varying from 4 to 7 days, we read simple order-up-to S rules with most-frequent levels for the scaled case. Since we will focus on the quality of the simple order-up-to S rule, we do not consider alternative production rules that might better fit to the optimal strategy. Table 3.23 display the several order-up-to levels for the scaled cases.

Apparently the maximal shelf life hardly affects the order-up-to levels. Most prominent are the higher order-up-to level for Friday when the maximal shelf life increases from 5 to 7 days. The higher order-up-to level on Friday reduces shortages on Monday, while the longer shelf life makes outdating less of a problem. These statements are verified by looking at the details from the simulation of the re-scaled order-up-to S rule for the full-size PPP.

Again the order-up-to S rules are simulated for 100 million weeks. The results are tabulated in Table 3.24. Raising the maximum shelf life from 4 to 5 days is most beneficial considering the reduction of outdating and shortages. As expected mismatches are very high when $m = 4$, since all any-demand on Monday must be mismatched under the given definition of the 'young' demand. Outdating and shortages roughly halves when the shelf life is increased by one day. Consequently also the cost figures halves when the shelf life

Table 3.24: Performance of the order-up-to S rule with re-scaled order-up-to levels for varying maximal shelf life and young-demand prefers BPPs with a residual shelf life of at least 3 days. (Simulation results over 100 million weeks).

	Outdating [a]		Shortage [a]		Quality mismatch [b]		Mean age	Annual costs
$m = 4$ days								
$S = (60, 68, 68, 64, 80)$	311	4.0%	62	0.82%	1056	21.15%	1.86	304,015
$m = 5$ days								
$S = (64, 72, 72, 64, 80)$	142	1.9%	33	0.44%	9.2	0.18%	2.04	47,726
$m = 6$ days								
$S = (64, 72, 72, 64, 84)$	88	1.2%	15	0.21%	6.7	0.13%	2.15	26,068
$m = 7$ days								
$S = (64, 72, 72, 64, 88)$	42	0.56%	8.1	0.11%	2.1	0.04%	2.25	12,760

[a] in pools per year; % of total (7488 pools),
[b] in pools per year; % of young-demand (4992 pools).

is lengthened by 1 day, except for $m = 4$: the costs when $m = 4$ are excessively high due to the mismatches on Monday. At maximal shelf life 7, both outdating and shortages are very low, hence lengthening the shelf life even more gives only a slight improvement.

A drawback of an increased shelf life is that the average age of the BPPs upon meeting the demand is somewhat higher. By definition the age of a BPP is at least one day, at $m = 7$ the mean age of the issued BPPs is 2.25 days, while this is only 1.86 days when $m = 4$. The mean age of a BPP is of secondary importance in our main study. In Kortbeek, Van der Wal, Van Dijk, Haijema, and De Kort [77] a number of ways is discussed to use the SDP-Simulation approach for finding good order-up-to levels when transfusing young pools is even more important.

3.8 Discussion and Conclusions

This section is split into two parts. First the SDP-Simulation approach is discussed, next its application to the PPP and the results are examined.

3.8.1 SDP-Simulation approach – Discussion and conclusions

Although the platelet production-inventory problem (PPP) is a high-dimensional problem, the problem exhibits enough structure to formulate an MDP problem and solve it. Solving the MDP requires the aggregation of states such that an SA algorithm can be executed in a reasonable amount of time. Aggregation is needed at several levels. First of all, the blood groups are to be aggregated as if blood platelet pools (BPP) are of a single universal blood group. Secondly, BPPs are aggregated in batches when the PPP plays at a scale that is too large. Finally, to come up with a simple rule, the stock consisting of batches of different age categories are aggregated into a single number. By simulation the structure of the optimal strategy is investigated and approximated by simple rules.

For the PPP with realistic data from a Dutch blood bank, aggregation at all three levels is possible and needed. For the given data, it appears that an order-up-to S rule roughly fits to the optimal ordering decisions. When other data, e.g. from a smaller blood bank or with more uncertainty in the demand, is used, more advanced stock-age-dependent rules may be needed.

The combined SDP-Simulation approach appears to be a very powerful approach to obtain insight in the structure of the optimal ordering strategy for perishable inventory problems, such as the PPP. It is a fast approach for:

- learning about simple rules that resemble the optimal strategy and
- obtaining their respective nearly optimal parameter values; e.g. nearly optimal order-up-to levels (and bounds) are easily read from frequency tables generated by simulation of the optimal strategy.

In this discussion section we elaborate on solutions to overcome some computational difficulties that one may experience when applying the approach to other perishable inventory problems. Next to the computational difficulties related to the available computation time, one may have difficulties in fitting a scaled demand distribution. In rare cases fitting a distribution on scaled demand figures is infeasible. We discuss these difficulties and ways of resolving them. First, we discuss the computation times of the solution procedures to give an idea of the computational complexity and the necessity of scaling the problem. Finally we raise the question to which extend the approach can be extended to solve a finite horizon problem, such that it can deal with non-stationary problems.

Computation times

Core of the SDP-Simulation approach is solving the (scaled) MDP and simulating the thus obtained optimal strategy. The MDP is solved by an SA algorithm, an iterative scheme that is relatively easily coded in any programming language, i.e. in our case in Delphi-Pascal. The computational complexity of an SA algorithm is already discussed in Section 3.2.6. To give an indication of the computation time of the algorithm, we report in Table 3.25 the time it takes to solve the PPP for varying values of the maximal shelf life m by SA on a PC with a Pentium-D 2.8GHz (dual) processor with 1Gb RAM. The absolute running times serve as an indication (and can be reduced by streamlining the code and cutting out any flexibilities of the current program).

The total number of states of the scaled MDPs are significantly reduced by the production and storage capacity, such that states that are very unlikely to happen under an optimal strategy are not considered. The optimal MDP strategy can thus be computed in a reasonable amount of time. For example if the storage capacity was not imposed, the total number of states for $m = 7$ would be about 19 million, the computation time would be much higher and the storage of the value vectors would require about 300Mb. Clearly, the number of states and the running times grow more than linearly in the maximal shelf life.

Table 3.25: The computational complexity of solving the scaled PPP for varying maximal shelf live and production capacity limited to $(20, 20, 16, 16, 20)$ batches for Monday to Friday and the storage capacity set to 35 batches.

Maximal shelf life	# of states scaled PPP	Running time SA (in s.)
4	135,066	38
5	701,832	72
6	1,427,694	172
7	3,890,824	1860

Solving the MDP takes thus less than a minute for $m = 4$ to half an hour for $m = 7$.

The next step in the SDP-Simulation approach is to simulate the optimal strategy to generate frequency tables for identifying a simple rule. Since we are not interested in a very accurate estimation of the state-frequencies, a short simulation run of say 100,000 weeks suffices. Such a short simulation run takes only a couple of seconds.

The more detailed simulations for obtaining accurate estimates of the performance statistics should last much longer, say 100 million weeks. The simple Universal-group simulation model executes in 20-30 minutes, which is relatively fast given the number of statistics that are tracked. For the scaled problem one could also estimate the statistics by solving related Markov chains under different cost structures (as argued in Section 3.2). The evaluation of the average costs of the Markov chains is done in 5 seconds for $m = 4$ up to 216 seconds for $m = 7$. These figures can be lower when the required accuracy is smaller.

The Multi-group simulation model executes much slower than the Universal-group model, since much more stochastic processes have to be simulated and the issuing policy is more evolved. Given the eight different blood groups, random numbers are generated for the demand by 16 patient groups and for the supply by eight donor groups. For the full-size problem the Multi-group simulation of 100 million weeks last about 5 hours. The running time of both simulation programs grow only linearly when the maximal shelf life m is increased, whereas solving the MDP takes an in-m-exponentially growing running time.

A simulation-based search for optimal order-up-to levels using the Multi-group simulation program seems to be very time-consuming. Therefore aggregation of the blood groups and scaling the problem is favored for simulation as well as for solving the MDP.

MDP intractable for high values of m?

As argued above, and in Section 3.2.6, the running time of the SA algorithm strongly depends on the maximal shelf life m, since the state space is $(m + 1)$-dimensional. When considering the approach for solving other perishable inventory problems of products with a longer maximal shelf life, one needs to limit the time it takes to solve the MDP. We have shown that by scaling the PPP the optimal solution to the MDP can be solved (approximately) by SA for m up to 7 days. However when we consider an inventory problem for a perishable product with a higher fixed shelf life, of say 14 periods or even more, scaling the demand figures as if demand happens in batches may not resolve the computational problems (unless the demanded volumes are very low).

A simple solution is to approximate the MDP by an MDP with a shorter shelf life $m' < m$

for which the MDP is tractable (after scaling). Thus the dimensionality of the MDP is reduced from m to m', but this approximation is accurate only when in reality rarely products with residual shelf life $(m' + 1)$ days or older are in stock. This approach was first considered by Nahmias [99]. In [77], the approach is adopted as an alternative to aim at enough young BPPs in stock to meet the demand for young under an order-up-to S rule.

Another solution to reduce the dimension is to model the problem on a more coarse time scale by aggregating time, e.g. two periods are aggregated into one. When the problem is modeled as if production decisions are scheduled every two days instead of every single day, then the age of the products in stock can be expressed in multiples of two days. This modeling assumption slightly affects the accuracy in tracking outdating. The quality of such approximations are not considered in this thesis.

Infeasible fit of demand distribution?

After downsizing the average demand figures by a factor B, it might be impossible to fit a demand distribution on the first two moments. Not for all combinations of the mean demand and the coefficient of variation of the demand a fit of a discrete probability distribution is feasible, as shown in [2]. In particular when after scaling the mean demand μ is low, say less than 1, then a fit is only possible for 'high' coefficients of variation cv. As cv is copied from the full-size problem, the problem may arise at the scaled problem but not at the full-size problem.

A way to overcome such a fitting problem is choosing a smaller scaling factor B. Thus the computation time for solving the MDP will be somewhat higher due to an increase of the number of states (amongst other). Alternatively, one may modify the demand data slightly: one fits the distribution to a higher value of either cv or μ. One thus needs to accept such an modeling error. When the coefficient of variation, cv, of a particular demand category on a specific day is too low to get a fit on the first two moments, one may model the problem as if the cv is higher. Thus the uncertainty in the demand is increased, resulting in a (slight) overestimation of shortages and an underestimation of the number of outdated products.

When one chooses to increase μ for a specific patient group and day such that a fit becomes possible, then one may compensate by decreasing the mean demand of another patient group or the mean demand at some other day(s). For example consider a weekend day where only a small portion of the total demand requires 'young' BPPs. Instead of discriminating between the two demand types all demand is modeled as demand for pools of any age. When only a small fraction of the total demand is modeled incorrectly, the impact on the overall mismatching is negligible. In the simulation of the full-sized problem the real mismatching figures can be computed to validate the results.

Non-stationary demand or production breaks

When modeling the PPP as an MDP we explicitly have assumed that the problem is periodic but stationary: every Monday is stochastically the same and production breaks happen only in weekends. In practice, blood bank managers need to anticipate irregular production breaks. For example on the two Christmas days production is hampered by a lack of voluntary donors, while demand for BPPs at hospitals continue (maybe at a somewhat lower level). The SDP-Simulation can be extended to deal with such periods, as will be discussed in the next chapter.

3.8.2 Application – Discussion and conclusions

Results for PPP

The SDP-Simulation approach is applied to the PPP introduced in Section 2.1 using data for one of the four Dutch blood banks.

- In 2003 outdating in The Netherlands was 15-20%, and shortages happened at division North-East roughly once every one or two weeks.
- We showed that outdating can be 2% (or even less) while shortages hardly happen, even when about two thirds of the demand gets the 'youngest' pools in stock.
- Just as in the current practice BPPs are produced primarily of the blood groups O and A, since compatible matching is the standard.
- A detailed sensitivity study is executed including practical issues, such as an increase of the maximal shelf life, 'biased' demand estimates, and different costs to trade-off between outdating and shortages.

Generalization

We believe that the results for the PPP at a Dutch blood bank gives an indication of what can be achieved at other blood banks of comparable size. The SDP-Simulation approach can also be applied to other blood banks and to hospitals. However, one has to be careful as the processes and the scale of the operations may differ between blood banks. Some differences one may observe in other regions and countries, and at hospitals are that:

- In some regions and countries BPPs are (also) stored in hospitals,
- The order problem at other blood banks and hospitals may play at a much smaller (or a much larger) scale,
- Delivery of fresh BPPs at hospitals may not happen daily: e.g. because fixed order costs are accounted.

We address these issues in the next chapter.