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

Haijema, R.

Publication date
2008

Chapter 4

Extended SDP-Simulation approach: production breaks and fixed order costs

4.1 Introduction

In this chapter we extend the SDP-Simulation approach, such that it solves non-stationary problems. Although the approach is more general we discuss the extension in the context of the platelet production problem (PPP). In practice the PPP may be virtually stationary, except for a number of periods during which additional production breaks occur because of holidays, e.g. Easter, Christmas and New Year's Day. This kind of short production breaks are known well in advance.

- Setting the right production volumes to anticipate near-future production breaks is in general difficult. Formal support through an extended SDP-Simulation approach is thus needed to support blood bank managers.

The distinction of young-demand from any-demand is not common to all blood banks. Therefore and to simplify the discussion we assume a single demand category. Another reason for doing so is the fact that when the maximal shelf life of BPPs is 5 days only, no 'young' BPPs are available after a production stop of 3 or more days. Mismatch costs are thus excluded from the model. Moreover, we simplify the discussion by meeting all demand with the oldest BPPs in stock first (FIFO), which is common to blood banks.

After applying the extended approach to one of the Dutch blood banks, we consider a case where demand volumes are 10 times as small. For the latter case one may think of a small blood bank or a medium-sized or large hospital. In The Netherlands only a few hospitals keep BPPs in stock. The lead time for regular orders is assumed to be one day.

Further, we provide answers to questions raised in the discussion at the end of the previous chapter:

- At a small blood bank or at a medium-large hospital, the coefficient of variation of the demand usually is higher. We expect that the optimal ordering policy is then more complicated. Therefore we investigate for a few cases whether an order-up-to S rule fits well or that other rules should be considered.
- When operating at a small scale fixed set-up costs may apply for each production run or fixed order costs may apply to hospitals. How do these fixed costs affect the optimal strategy? Is the SDP-Simulation approach also in this case helpful in deriving a nearly-optimal rule and the respective parameter values? In particular we are interested whether an (s, S)-policy fits well and whether nearly optimal thresholds s can be read from the frequency tables.

Answers to these questions are of interest to inventory managers at both blood banks and hospitals as well as to Operations Researchers.

Outline

In Section 4.2 we extend the SDP-Simulation approach such that it can deal with periods where production and demand may be non-stationary. Throughout the whole chapter, we assume that all demand is met by issuing the oldest BPPs in stock first, i.e. all demand is 'FIFO-demand'. Consequently mismatch costs are zero as no distinction is made between young-demand and any-demand. In Section 4.3 we present case study results for a Dutch blood bank that faces short production breaks during Christmas, New Year's Day and Easter.

Next, we consider in Section 4.4 a stationary order problem for a blood bank with much smaller average demand volumes. Alternatively one may think of a hospital that faces an average demand that is only a tenth of the demand at the blood bank. In Section 4.5 we add fixed order costs to the cost structure.

4.2 Extended SDP-Simulation approach

4.2.1 Problem: non-stationary production breaks

A practical question for production-inventory managers, and for blood bank managers in particular, that remained unanswered is: "How should one anticipate irregular production breaks like at Easter and Christmas?" When regular production stops for a few days, while demand continues (as usual) one should anticipate breaks by producing somewhat more some days before the break. Consequently the age distribution of the BPPs in stock is affected and thus the optimal production volumes immediately after a break might be different as well.

In the PPP most weeks are stochastically the same: the supply of whole blood (by voluntary donors) and the demand for BPPs are stationary processes. However, during a short holiday break production might be impossible, since donors do not show up.

As an example we consider, in Section 4.3, two cases of particular interest to blood bank managers:

1. A Christmas period (December 25 and 26) falling on Tuesday and Wednesday, followed by the New Year's Day (NYD) on the next Tuesday.

2. The four-days Easter period from Good Friday to Easter Monday. In The Netherlands the (regular) production is stopped for four consecutive days, while the maximal shelf life of BPPs is only 5 days.

In practice it is difficult to find nearly optimal production volumes on days around holidays. Due to the occasional nature of these events, it is not possible getting experienced. Formal support is thus needed. In the next section we extend the SDP-Simulation approach to include non-stationary production breaks and apply it to the BPP production-inventory management.

4.2.2 Approach

We discuss our approach by considering the breaks at a Christmas period and an Easter weekend. As the time between these breaks is very long compared the maximal shelf life of the products, one may analyze these periods independently of each other. In Figure 4.1 we illustrate our approach for the 4-days Easter weekend. The gray squared blocks on the time bar indicate production stops.

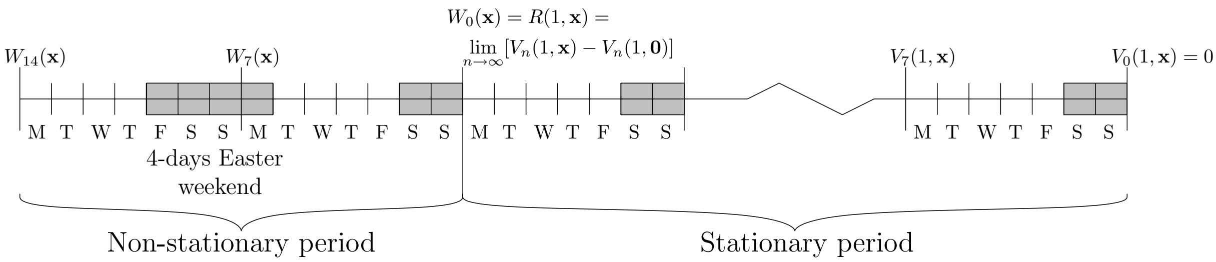

Figure 4.1: The SDP-Simulation approach splits the horizon in two parts, since the system behaves stationary a few days after a break

We model the problem as an infinite horizon problem, that consists of two parts. The first part is the irregular period containing the non-stationary production break(s). The second part relates to the stationary problem. In the example in Figure 4.1, the age distribution of the BPPs in stock on the Tuesday immediately after the break differs from an ordinary Tuesday, the optimal ordering policy may be different than on ordinary Tuesdays. But we may assume that the system returns to the stationary problem a few days after the irregular period has elapsed. We assume that from the next Monday onwards the system behaves indeed stationary.

By stochastic dynamic programming (SDP) the optimal ordering strategy can be computed for both the stationary and the non-stationary problems. We first apply successive approximation to solve the stationary problem and thus obtain the optimal ordering strategy for regular weekdays, Monday to Friday. As a by-product we compute relative values of each stock state that are used as terminal costs to solve in a backward fashion the non-stationary problem. In the next two sections we formalize the approach.

4.2.3 Stationary MDP formulation

The SDP-simulation approach starts with solving the MDP for the stationary problem, since its solution is input to the stochastic dynamic program for the non-stationary problem. The MDP formulation differs only slightly from the one formulated in the previous chapter, as we now deal with only a single demand category, which is met through a FIFO issuing policy.

The stock transition from weekday d to the next day, given that demand is met by a FIFO-issuing policy, follows from the stock transition function $y(\mathbf{x}, k, a)$ and depend on the initial stock $\mathbf{x}$, the demand k, and the production volume a. At the end of the day BPPs that have become outdated are disposed of. The transition probability relate to the demand distribution for weekday d is denoted by $p_d(k)$.

The immediate or direct costs to incur in a slot is now much easier than before, as we consider only a single demand category and no mismatch costs. Let $C(d, \mathbf{x}, k)$ denote the costs to incur in on weekday d, when the demand on that day is k and the initial stock is $\mathbf{x}$. With the total stock level denoted by $x = \sum_{r=1}^{m} x_r$, the direct cost $C(d, \mathbf{x}, k)$ are:

$$C(d, \mathbf{x}, k) = \begin{cases} c^{\mathrm{O}} \cdot (x_m - k)^+ & \text{outdating costs,} \\ +c^{\mathrm{S}} \cdot (k - \sum_{r=1}^{m} x_r)^+ & \text{shortage costs,} \\ +c^{\mathrm{H}} \cdot x & \text{holding costs.} \end{cases} \tag{4.1}$$

Note that in the definition of $C(\cdot)$ we have omitted the parameter I of the issuing policy as in this chapter we consider FIFO issuing only. The expected immediate costs to incur in state $(d, \mathbf{x})$ are simply:

$$\mathbb{E}C(d, \mathbf{x}) = \sum_{k} p_d(k) \cdot C(d, \mathbf{x}, k). \tag{4.2}$$

Using Equation (4.6), the MDP is tractable through a SA algorithm, unless the state space $\mathcal{X}$ is too large.

$$V_n(d, \mathbf{x}) = \min_{a \in \mathcal{A}(d,\mathbf{x})} \left(\mathbb{E}C(d, \mathbf{x}) \; + \; \sum_k p_d(k) \cdot \; V_{n-1} \left(d+1, y(\mathbf{x}, k, a) \right) \right). \tag{4.3}$$

We start the SA algorithm by setting $V_0(1, \mathbf{x}) = 0$ for all states $\mathbf{x} \in \mathcal{X}$. The choice the start at a Monday ($d = 1$) is arbitrary, we could as well choose any other weekday to start with. Next $V_1(7, \mathbf{x})$, $V_2(6, \mathbf{x})$,. . . ,$V_7(1, \mathbf{x})$,etc. are computed for all states $\mathbf{x}$.

The span($\mathbf{V}_n - \mathbf{V}_{n-7}$) is checked every 7 iterations, hence at iteration $n = 7$, 14, 21, etc. Suppose $span(\mathbf{V}_n - \mathbf{V}_{n-7})$ is for the first time smaller than a pre-specified small value ϵ at iteration $N - 7$, with N a multiple of 7 days. Optimal actions for each state $\mathbf{x}$ on Monday to Sunday are derived from Equation (4.4) from the last 7 iterations ($n = N - 6$, $N - 5, \ldots, N$). Hence after N iterations an optimal stationary strategy is approximated by:

$$\pi(d, \mathbf{x}) = \arg \min_{a \in \mathcal{A}(d,\mathbf{x})} \sum_k p_d(k) \, V_{N-d}(d+1, y(\mathbf{x}, k, a)). \tag{4.4}$$

In the last 7 iterations the optimal production strategy π is stored. As N is a multiple of 7, the value vector $\mathbf{V}_N$ relates to a Monday. $\mathbf{V}_N$ is a relative value vector that can be used as terminal costs in solving the finite horizon problem of the non-stationary period.

4.2.4 Extension – Including non-stationary periods

In a very similar way the ordering decisions for each day of the non-stationary finite horizon problem are computed. For example, in Figure 4.1, the special period lasts two weeks (14 days) with the 4-days Easter weekend in the middle of the time interval. Note that we have chosen more or less arbitrarily that the finite horizon problem ends a few days after the break on a Monday morning.

For each of the 14 days of the special period the optimal production strategy is determined by Stochastic Dynamic Programming in a backward fashion, using the so-called relative values $R(1, \cdot)$ as terminal costs. The relative values $R(1, \mathbf{x})$ are used to compare stock states $\mathbf{x}$ within the same periodic class, namely the class related to Mondays, the value vector $V_N(1, \mathbf{x}$ can be used for this purpose. Instead of storing the optimal policy for each working day, we now store the policy for all 14 days of the non-stationary period.

The extended SDP-Simulation approach consists of the following five steps:

Step I. Compute relative values of the states for the stationary problem:

- First, the stationary problem is solved (maybe after scaling the problem), using Equations (4.6) and (4.4).
- Next, we choose an arbitrary reference state on Monday, say $(1, \mathbf{0})$, and compute for every possible stock state $\mathbf{x}$ on Monday the difference in expected future costs relative to this reference state:

$$R(1, \mathbf{x}) = \lim_{n \to \infty} [\, V_n(1, \mathbf{x}) - V_n(1, \mathbf{0}) \,] \approx V_N(1, \mathbf{x}) - V_N(1, \mathbf{0}) \tag{4.5}$$

These differences, $R(1, \cdot)$, are feasible relative values, when N is sufficiently large. $V_N(1, \mathbf{x})$ is finite when N is finite. Therefore setting $R(1, \mathbf{x}) = V_N(1, \mathbf{x})$ is also feasible.

Step II. Solving the non-stationary problem by SDP:

Let the irregular period last T days, with the last day being a Sunday. The days of this period are numbered backwards and denoted by t. $t = 1$ thus refers to the last day of the period (a Sunday) and day T is the first day of the finite horizon. Index t thus denotes the number of days to go until the end of the irregular period. In addition to the notation in Equations (4.6) and 4.4 we define:

$p_t^{\text{irr}}(k)$ = the probability of a (composite) demand k on day t. If the demand remains stationary even during a break, then $p_t^{\text{irr}}(k)$ equals $p_d(k)$ for $d = 7 - (t-1) \mod 7$.

$\mathcal{A}'_t(\mathbf{x})$ = action space at day t as bounded by the (artificial) production and storage capacity. Clearly the action space depends on the stock state $\mathbf{x}$. If production is not possible on day t, $\mathcal{A}'_t(\mathbf{x}) = \{0\}$.

$W_t(\mathbf{x})$ = the total expected costs under an optimal strategy from day t onwards when starting in inventory state $\mathbf{x}$ and at the end of the irregular period terminal cost are accounted.

$W_0(\mathbf{x}) \equiv R(1, \mathbf{x})$.

$\pi_t(\mathbf{x})$ = the optimal decision at day t given the inventory state is $\mathbf{x}$.

By stochastic dynamic programming one recursively computes and stores successively for $t = 1, 2, \ldots, T$, for all states $\mathbf{x}$ in the state space $\mathcal{X}'(t)$:

$$W_t(\mathbf{x}) = \min_{a \in \mathcal{A}'_t(\mathbf{x})} \left(\mathbb{E}C_t(\mathbf{x}) \ + \ \sum_k p_d(k) \cdot \ W_{t-1}(y(\mathbf{x}, k, a)) \right). \tag{4.6}$$

with $\mathbb{E}C_t(\mathbf{x}) = \sum_k p_t^{\text{irr}}(k) \cdot C(d, \mathbf{x}, k)$, in which d is the weekday related to t.

The optimal ordering quantity on day t follows from

$$\pi_t(\mathbf{x}) = \arg \min_{a \in \mathcal{A}'_t(\mathbf{x})} \sum_k p_t^{\text{irr}}(k) W_{t-1}(y(\mathbf{x}, k, a)). \tag{4.7}$$

Step III. Read simple rule from simulation-based frequency table
Again the optimal strategy may be fairly complex. Hence simulation is used to investigate the structure of the optimal strategy for each day of the special period. That is: frequency tables for each day t of the irregular period are generated and, if applicable, an order-up-to S rule is read for each day of the week or any other appropriate ordering rule.

Step IV. Finally, by a detailed simulation program the rule is put to the test and its performance, in terms of outdating and shortage figures, is compared to the figures for the optimal stock-age-dependent strategy.

Step V. In case the initial problem is scaled the simple rule is simulated for the full-size problem after re-scaling the parameters of the simple rule.

4.3 Case study – optimal policy around breaks

In this section we apply the extended SDP-Simulation approach using realistic data, as summarized in Section 4.3.1. The results for the problem around Christmas and New Year's Day are presented in Section 4.3.3. In section 4.3.4 we discuss the 4-days Easter weekend. The results are integrated in Section 4.3.5. We will show that a simple rule applies even around breaks, although we do not provide a detailed sensitivity study as we have done in the previous chapter. We limit the illustration to the first four steps of the SDP-simulation approach.

4.3.1 Data

The case study is constructed by modifying the data for one of the four Dutch Blood banks (Sanquin, division North-East). The demand is aggregated into a single category that will receive the oldest BPPs in stock (FIFO issuing). The demand for 144 BPPs a week is spread over Monday to Sunday as in Table 4.1. To make the MDP tractable we scale the problem by a factor 4, resulting in the mean demand figures in the last row of Table 4.1.

Table 4.1: Demand distributions: means and coefficients of variation (cv).

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
original mean	26	21	32	21	26	8	10
cv	0.28	0.31	0.25	0.31	0.28	0.50	0.45
scaled mean	6.5	5.25	8	5.25	6.5	2	2.5

The coefficient of variation (cv) of the demand is set to 1.4 times the cv when demand would be Poisson distributed. In fact, this assumption corresponds to the worst case of Section 3.7.2, where one demand category, comprising two thirds of the average demand, is Poisson distributed, but the cv of the remaining third of the demand is twice as high. The cv over all demand is then $\sqrt{2} \approx 1.4$ times[1] the cv for the Poisson case.

[1]Consider three stochastic variables X, X_1 and X_2. Let $X = X_1 + X_2$, with mean μ. X_1 is Poisson($\mu_1 = \frac{2}{3}\mu$) distributed, X_2 has mean $\mu_2 = \frac{1}{3}\mu$ and coefficient of variation $cv_2 = 2 \cdot \frac{1}{\sqrt{\mu_2}}$. Hence variance of X is $(cv_1 \cdot \mu_1)^2 + (cv_2 \cdot \mu_2)^2 = \mu_1 + 4\mu_2 = \frac{2}{3}\mu + \frac{4}{3}\mu = 2\mu$. Hence the coefficient of variation of X is $\sqrt{2}/\mu = \sqrt{2} \cdot cv(\text{Poisson}(\mu))$.

For each day of the week d we fit a discrete probability distribution $p_d(\cdot)$ on the mean demand and the reported cv. Conform [2] the unscaled demand distributions are mixtures of two negative binomial distributions. For the scaled problem, the demand distributions are fitted using a mixture of two binomial distributions.

The other problem data for the unscaled case remain unchanged:

Annual demand	7,488 BPPs or 1,872 batches
Costs	outdating 150 per outdated BPP, shortage costs 750 per BPP short, all other costs are zero,
Production	Monday – Friday, but no during breaks
Maximal shelf life	$m = 5$ days.

For a high accuracy performances statistics will be obtained from 100 detailed simulation runs of 1,000,000 weeks each.

4.3.2 Step I – Stationary problem

After scaling the MDP and solving the scaled problem, we simulate the resulting optimal strategy. Table 4.2 presents the simulation-based frequency tables for Monday to Friday. In the first column we read the order-up-to levels related to the optimal actions. From the last column we read which level is most frequent. On Mondays the most-frequent order-up-to level is 21 and it fits to 59% of the 1 million states visited. Similarly, we find most-frequent order-up-to levels $(21, 21, 21, 19, 24)$ batches (of 4 BPPs) for Monday to Friday.

In Table 4.3 we compare the order-up-to rule with order-up-to levels fixed to $(21, 21, 21, 19, 24)$ against the optimal (age-dependent) strategy. The upper half of the table shows the characteristics of the optimal MDP strategy, the lower half shows the performance of a fixed replenishment rule. As expected from the results and the discussion in the previous chapter, the order-up-to S rule appears to perform nearly optimal. The absolute outdating and shortage figures per week should be related to a weekly demand of 36 batches of 4 pools. The annual results relate to an average annual demand for 1872 batches.

Table 4.2: Simulation frequency tables of scaled MDP strategy with FIFO demand only.

(a) (State, action)-frequency tables for 1,000,000 simulated Mondays.

Stock x	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	Freq(S_1)
Up-to S_1																				
24													4	6	12	153	101	250	27	553
23												4	46	2906	17596	5508	1262			27322
22										5	1156	121403	79791	39802						242157
21								116497	152240	167708	154513									590958
20						39450	74336													113786
19					17209															17209
18				5976																5976
17			1657																	1657
16		332																		332
15	50																			50
14																				0
:																				:
0																				0
Freq(x)	50	332	1657	5976	17209	39450	74336	116497	152240	167713	155669	121407	79841	42714	17608	5661	1363	250	27	1000000

(b) (State, action)-frequency tables for 1,000,000 simulated Tuesdays.

Stock x	0...7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	Freq(S_2)
Up-to S_2																		
26														1				1
25														1	1	3		5
24												2	6	42	33	22	1	106
23											5	277	622	898	247	16		2065
22									40	653	5055	15270	10982	2259	215			34474
21			786	9588	44321	83285	143525	191312	198489	159515	95150	32597	4316	427				963311
20		38																38
19																		0
:																		:
0																		0
Freq(x)	0	38	786	9588	44321	83285	143525	191312	198529	160168	100210	48146	15926	3628	496	41	1	1000000

(c) (State, action)-frequency tables for 1,000,000 simulated Wednesdays.

Stock x	0...9	10	11	12	13	14	15	16	17	18	19	20	21	22	Freq(S_3)
Up-to S_3															
25													3		3
24									2		72	103	87	2	266
23								9	173	1314	4637	1865	592		8590
22						9	988	6708	49510	99883	31271	5334			193703
21				16565	59327	136626	214317	231213	136611						794659
20		167	2612												2779
19															0
:															:
0															0
Freq(x)	0	167	2612	16565	59327	136635	215305	237930	186296	101197	35980	7302	682	2	1000000

(d) (State, action)-frequency tables for 1,000,000 simulated Thursdays

Stock x	0...4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	Freq(S_4)
Up-to S_4																			
22																	3	6	9
21															106	753	204		1063
20											1156	7482	44639	30494	8368	931			93070
19					7114	24181	58816	110748	164500	196650	184235	129167	28890						904301
18			145	1401															1546
17		7																	7
16																			0
:																			:
1																			0
0																		4	4
Freq(x)	0	7	145	1401	7114	24181	58816	110748	164500	196650	185391	136649	73529	30494	8474	1684	207	10	1000000

(e) (State, action)-frequency tables for 1,000,000 simulated Fridays.

Stock x	0...7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	Freq(S_5)
Up-to S_5																
28														1	2	3
27														2		2
26											30	262	223	70		585
25							1	378	2670	17755	20380	9471	1110			51765
24		141	2406	15686	57196	132318	210586	235857	185682	87884	19889					947645
23																0
:																:
0																0
Freq(x)	0	141	2406	15686	57196	132318	210587	236235	188352	105639	40299	9733	1333	73	2	1000000

Table 4.3: Order-up-to S rule vs MDP policy for a regular (stationary) week (statistics in batches of 4 pools obtained by simulation for 100 million weeks).

Weekday	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Weekly	Annual
MDP policy									
# batches outdated	0	0.01	0.06	–[a]	–[a]	0	0.01	0.08	4.30 (0.23%)
# batches short	0.01	0	0	0	0	0	0	0.01	0.71 (0.04%)
Annual costs									4,698 euro
Order-up-to S rule									
Levels S_d	21	21	21	19	24	–	–		
Goodness-of-fit	59%	96%	79%	90%	95%	–	–		
# batches outdated	0	0.01	0.06	–[a]	–[a]	0	0.01	0.08	4.18 (0.22%)
# batches short	0.01	0	0	0	0	0	0	0.01	0.74 (0.04%)
Annual costs									4,724 euro

[a] Outdating must be zero on Thursday and Friday, as $m = 5$ and production stops in the weekends.

From the scaled MDP results we conclude that

- Compared to the current practice, it seems that the annual outdating can be reduced from about 15% to 0.2%, even when demand is more stochastic than Poisson.
- Shortages occur virtually never: only 0.04% of the total annual demand cannot be met immediately from stock.
- The order-up-to S policy as read from the frequency tables closely approximates the structure of the optimal strategy and is only 0.6% off from the optimal cost level.

Remark For the unscaled case, the replenishment levels are re-scaled by multiplication with a factor 4. In the previous chapter we have already shown that the resulting order-up-to levels are nearly optimal. More results for cases where all demand is aggregated into a single category are found in [56], [57], [58], [149], [77].

4.3.3 Steps II to IV – Christmas and New Year's Day

For example, we consider a year in which Christmas falls on Tuesday and Wednesday, and New Year's Day (NYD) on the next Tuesday. On these days (regular) production is stopped because of a lack of donors. As depicted on the timeline in Figure 4.2 there is only a single day between the weekend and the two Christmas holidays. On this Monday one has to produce additional BPPs to anticipate the production stop for the next two days. In this section we consider a worst case where demand for platelet pools continues as usual.

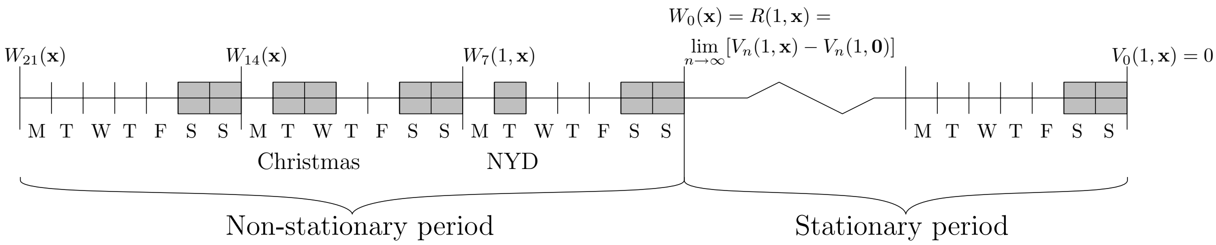

Figure 4.2: The SDP-Simulation approach splits the horizon into two parts, as if the system behaves stationary a few days after a break

What to expect?

When the production capacity on Monday is not restrictive, as in our case, then the production on Monday has to be set high enough to anticipate the production stop on the next two days. When the production capacity on Monday is normally high enough but too restrictive to anticipate the two-day production stop, then the production volume on Friday is raised as well. When no holding costs and quality mismatch costs apply, the capacity restrictions will hardly affect the cost-level. Compared to the stationary case the optimal cost level does increase, since outdating and shortages are likely more prevalent. Shortages mostly occur on the Thursday after Christmas. Excessive production on the Monday before Christmas outdates the next Saturday.

We assume that the production and storage capacity are not restrictive. (Artificial bounds, which are set to make the state space finite, are set high enough to ensure that the optimal policy is not affected.) The length of the non-stationary period can then be reduced from 3 weeks (as in Figure 4.2) to 2 weeks, as the problem in the first week maybe still stationary. When the production capacity on Wednesdays, January 2^{nd}, is not restrictive then the production strategy on Thursday January 3^{rd}, may differ only slightly from that on an ordinary Thursday. Moreover, given $m = 5$, no BPPs will expire on Thursday, as no BPPs in stock on Thursdays are older than three days.

Results for optimal strategy

Table 4.4 shows the impact of the irregular production breaks on the optimal strategy over a ten-days period from Monday (December, 24^{th}) to Wednesday (January, 2^{nd}). The results are presented in batches of 4 pools, since the MDP is scaled to reduce the state space. The average demand over the 10-days period is almost 56 batches (223 pools).

Table 4.4: Impact of production breaks around Christmas and New Year's Day.

(a) Impact of breaks on outdating and shortages (in batches) under MDP policy.

10-days period	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed	Total
Stationary case											
# batches outdated	0	0.01	0.06	–[a]	–[a]	0	0.01	0	0.01	0.06	0.15
# batches short	0.01	0	0	0	0	0	0	0.01	0	0	0.03
Irregular breaks		Dec-25	Dec-26						Jan-1		
# batches outdated	0	0.01	0.11	–[a]	–[a]	0.59	0	–[a]	–[a]	0.05	0.77
# batches short	0	0	0	0.06	0	0	0	0.01	0	0.01	0.08

[a]No outdating $m = 5$ days after weekends and holidays, since production is zero.

(b) Impact of breaks on order-up-to levels S_t as read from the optimal MDP policy.

10-days period	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed
Stationary case										
Order-up-to S_t	21	21	21	19	24	–	–	21	21	21
Goodness-of-fit	59%	96%	79%	90%	95%	–	–	59%	96%	79%
Irregular breaks		Dec-25	Dec-26						Jan-1	
Order-up-to S_t	31	–	–	20	24	–	–	27	–	21
Goodness-of-fit	98%	–	–	49%	44%	–	–	93%	–	48%

Table 4.4(a) shows that outdating for this ten-days period has increased from 0.15 batches (0.6 pools) to 0.77 batches (3.1 pools). Expected relative outdating thus is about 1.4%. Since the results relate to the optimal strategy one has to accept this increase, which is primarily due to the outdating of pools produced on the Monday before Christmas. Although not reported in the table, we have observed that only 0.55% of the BPPs (0.08 batch) produced on the Monday before News Year's Day becomes outdated on the next Saturday.

Order-up-to rule vs Optimal strategy

Since in practice one prefers a simple rule, we hope that an order-up-to S rule resembles the structure of the optimal policy. By simulation we generate (state, action)-frequency tables for each day of the special period in a similar way as for the stationary case. In Table 4.4(b) we report the most-frequent order-up-to levels for the ten-days period and compare them against the stationary ones. As expected the order-up-to levels on the Mondays before the two breaks are considerably higher than in the stationary case: e.g. 31 vs 21 before Christmas. Remarkably, an order-up-to S rule fits even better on Mondays just before a break, than on a regular Monday: 98% before Christmas versus 59% on the regular Mondays. On the days after a break the order-up-to S rule fits to almost 50% of the states visited, whereas in the stationary case this figure falls in 79%-95%.

We evaluate the order-up-to S rule with the most-frequent order-up-to levels from the

Table 4.5: Order-up-to S vs MDP policy around Christmas and New Year's Day.

10-days period	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed	Total
Irregular breaks		Dec-25	Dec-26						Jan-1		
MDP policy											
# batches outdated	0	0.01	0.11	–[a]	–[a]	0.59	–[a]	–[a]	0	0.05	0.77
# batches short	0	0	0	0.06	0	0	0	0.01	0	0.01	0.08
Order-up-to S rule											
Levels S_t	31	–	–	20	24	–	–	27	–	21	
# batches outdated	0	0.01	0.11	–[a]	–[a]	0.59	–[a]	–[a]	0.01	0.04	0.75
# batches short	0	0	0	0.06	0	0	0	0.02	0	0.01	0.10

[a]No outdating $m = 5$ days after weekends and holidays, since production is zero.

frequency tables by a long simulation (100 million replications) and compare its performance against the optimal strategy. (The results could be computed in an exact way by solving the underlying Markov chains under modified cost structures.) Although the order-up-to S rule does not fit for 100% at each day, its performance is very close to optimal as reported in Table 4.5. Over the given ten-days period on average only 0.1 batch (out of the 56 batches demanded) cannot be met from stock: the shortage rate is thus less than 0.2%. Outdating is in the order of 1.3%.

Re-scaling After re-scaling the order-up-to levels, one obtains a nearly optimal order-up-to S rule for the real-sized case. Given the robustness of the rule with respect to minor changes in the order-up-to levels as shown in the previous chapter we skip step V of our extended SDP-Simulation approach.

Conclusions – Results breaks during Christmas and New Year's Day

It appears that:

- an order-up-to S policy with increased order-up-to levels for the Monday before Christmas and the Monday before New Year's Day is nearly optimal,
- age plays an more important role after Christmas and New Year's Day, but the absolute impact on outdating and shortages is relatively small, and
- outdating over a ten-days period, including the breaks, is as low as 1.4%, while shortages are less than 0.2%.

4.3.4 Steps II to IV – Four-days Easter weekend

The second example of a non-stationary problem with irregular production breaks is the 4-days Easter weekend, as depicted before in Figure 4.1. The long weekend from Good Friday to Easter Monday is considerably more difficult, since the production is stopped for four consecutive days, while the shelf life is only 5 days and demand remains stationary. We assume that the production and storage capacity are not restrictive.

What to expect?

Again, we expect that primarily the production volumes one day before the break, in this case on the Thursday before Good Friday, are increased dramatically to anticipate the production stops for the next four days. The available stock plus the production on Thursday should be enough to survive until the next Wednesday morning when new stock is released. The order-up-to level on Thursday before the break can be derived by the Newsboy model, since no BPPs will survive until Wednesday.

The marginal return of ordering z batches instead of $z-1$ batches is the savings on shortages. The marginal costs are an increase in the expected outdating costs. Let the stochastic variable Z denote the demand in batches over the six-days period from Thursday to Tuesday, and $P(\cdot)$ is the cumulative distribution of Z. $P(\cdot)$ is obtained from the convolution of the six demand distributions. The optimal order-up-to level on Thursday is the greatest value of z for which the expected marginal savings on shortage costs is still larger than the expected marginal outdating costs, as reflected in Equation (4.8).

$$c^{\mathrm{S}} \cdot P(Z \geq z) \;\geq\; c^{\mathrm{O}} \cdot P(Z \leq z-1) \tag{4.8}$$

Hence the best order-up-to level is the greatest z for which holds:

$$P(Z \leq z-1) \;\leq\; \frac{c^{\mathrm{S}}}{c^{\mathrm{S}}+c^{\mathrm{O}}}. \tag{4.9}$$

For the given demand distributions and the cost figures, the order-up-to level on the Thursday before the break is according to the Newsboy equation 31 batches. The production volume just before the break is thus likely much higher than on a regular Thursday. Consequently, outdating mostly happens five days later on Tuesday, just after the break.

Since production stops from Friday to Monday, shortages primarily happen on Tuesday given that the production lead time is one day. We expect shortages and outdating to be far more prevalent than over the Christmas period.

On Tuesday morning after Easter Monday all products in stock, if any, are of the same age. The optimal production strategy is thus not stock-age-dependent. No BPPs produced before the break will survive until Wednesday morning, hence one may expect that the production volume on Tuesday is fixed to a target inventory level on Wednesday morning. Consequently, the optimal production volume on Wednesday is also fixed, as Tuesdays production becomes available only at the start of Wednesday morning and all products in stock are of the same age. From Thursday onwards the optimal production strategy is again stock-age-dependent.

Results of optimal strategy

After scaling and solving the MDP, we can check our expectations. Through simulation we generate frequency tables from which we read the structure of the optimal strategy. A selection of them is found in Table 4.6. As expected and argued before, according to the upward diagonal in Table 4.6(a), a fixed production volume applies on the Tuesday just after Easter Monday. All batches produced before the break will not survive until the Wednesday morning after Easter.

Since all stock present on Tuesday morning after Easter Monday will perish the same day, the initial stock on Wednesday morning consists only of the 15 batches produced on Tuesday. From Table 4.6(a) we observe that a fixed order-up-to level of 21 batches applies on Wednesday, which implies a fixed production volume of 6 batches. Note that an order-up-to level of 21 batches corresponds to the stationary order-up-to level reported in Table 4.4(b).

The order-up-to level on Thursday after the break equals 19, which corresponds to the stationary order-up-to level. This illustrates that the stationary order-up-to levels apply from Wednesday onwards. Although not reported the outdating and shortages figures from the Thursday after the break onwards do not differ significantly from those on regular days. One positive exception, not visible in the table, is that outdating on the Sunday after the break is significantly lower than usual.

The optimal production policy on Thursday prior to Good Friday resembles for virtually 100% an order-up-to S rule with fixed order-up-to level 32. This is inline with our expectations: the newsboy model suggest an order-up-to level of 31 batches. As the initial stock

Table 4.6: Frequency tables from 1 million simulations of MDP policy around Easter.

(a) (State, action)-frequency table for a Tuesday after Easter Monday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	Freq(S)
Up-to S																							
36																						1	1
35																					20		20
34																				176			176
33																			933				933
32																		3578					3578
31																	10375						10375
30																23749							23749
29															43775								43775
28														68301									68301
27													91544										91544
26												110320											110320
25											116201												116201
24										118217													118217
23									110287														110287
22								94938															94938
21							74669																74669
20						53775																	53775
19					35596																		35596
18				21466																			21466
17			11893																				11893
16		5758																					5758
15	4428																						4428
14																							0
:																							:
0																							0
Freq(x)	4428	5758	11893	21466	35596	53775	74669	94938	110287	118217	116201	110320	91544	68301	43775	23749	10375	3578	933	176	20	1	1000000

(b) (State, action)-frequency table for a Wednesday after Easter Monday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	Freq(S)
Up-to S																							
21																1000000							1000000
20																							0
:																							:
0																							0
Freq(x)	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1000000

(c) (State, action)-frequency table for a Thursday after Easter Monday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	Freq(S)
Up-to S																							
21																					245		245
20																		27771	8545	1830			38146
19									8547	27770	66651	122191	174559	196384	174561	122189	66652						959504
18							260	1830															2090
17																							0
:																							:
1																							0
0																						15	15
Freq(x)	0	0	0	0	0	0	260	1830	8547	27770	66651	122191	174559	196384	174561	122189	66652	27771	8545	1830	245	15	1000000

may not survive till the next ordering moment, the best order-up-to level is 1 higher. The order-up-to level 32 is 13 batches higher than on a regular Thursday (see Section 4.3.2).

Just as for the Christmas and New Year's Day period, on the day before a production break an order-up-to S rule (with increased levels) fits even better than in the stationary case. Apparently the age-distribution of the stock is less relevant on a day prior to a production break.

Order-up-to rule vs Optimal strategy

The structure of the optimal policy seems thus, again, to be very well presented by an order-up-to S rule. In Table 4.7 we report over a ten-days period, including the Easter weekend, the outdating and shortage volumes around Easter under both the optimal production policies and the order-up-to S rule. In the table Good Friday and Easter Monday are abbreviated by GF respectively EM.

In the last column of Table 4.7 we observe that the order-up-to S rule performs almost equally well as the optimal MDP policy. As expected, outdating and shortages happen mostly on the Tuesday after Easter Monday. Under both strategies on average approximately 4.3 batches will outdate and stock falls on average 0.28 batches short. The average demand over the ten-days period is on average almost 56 batches and the average production to cover the demand over this period is roughly 60 batches. Relative outdating over the ten-days period is thus $\frac{4.3}{60} = 7.2\%$. The shortage rate over the ten-days period is $\frac{0.28}{56} = 0.5\%$.

Table 4.7: Order-up-to S rule vs MDP policy around Good Friday and Easter Monday.

10-days period Irregular breaks	Mon	Tue	Wed	Thu	Fri GF	Sat	Sun	Mon EM	Tue	Wed	Total
MDP policy											
# batches outdated	0	0.01	0.06	–[a]	–[a]	0	0.03	0.07	4.16	–[a]	4.34
# batches short	0.01	0	0	0	0	0	0	0	0.26	0	0.28
Order-up-to S rule											
Levels S_t	23	22	24	32	–	–	–	–	15[b]	21	
Goodness-of-fit	44%	75%	64%	100%	–	–	–	–	100%[b]	100%	
# batches outdated	0	0.01	0.06	–[a]	–[a]	0	0.03	0.04	4.19	–[a]	4.33
# batches short	0.01	0	0	0	0	0	0	0	0.26	0	0.28

[a]No outdating $5(=m)$ days after weekends and holidays, since production is zero.
[b]A fixed production volume applies since all batches in stock are outdated at the end of the day.

Conclusions – Results four-day Easter weekend

The major conclusions over a ten-days period that includes the four-days Easter weekend are:

- A simple replenishment rule performs nearly optimal and results in:
  - an outdating figure of only 7.2%,
  - shortages to be less than 1%.
- Outdating and shortages happen mostly on the Tuesday after Easter Monday, indicating the difficulty in anticipating a four-days production stop when the maximal shelf life is only 5 days.
- The production levels on the Thursday before and the Tuesday after the weekend are considerably higher than in the stationary case.
  - On the Tuesday after Easter Monday a fixed production quantity applies, since pools produced before Good Friday will not survive until Wednesday.
  - Shortly after the break (from Wednesday onwards) the stationary order-up-to S rule resumes as a very good approximation of the optimal MDP policy.

4.3.5 Conclusions – Extended SDP-Simulation Approach

Extended SDP-Simulation Approach

The SDP-Simulation approach can be applied to solve the order problem of perishables with a (short) fixed shelf life. We have extended the approach such that it can deal with (non-stationary) production breaks. It appears to be a powerful approach for deriving an optimal stock-age-dependent (scaled) policy and for investigating the structure such that a practical rule can be derived from it.

Results over a year including Easter and Christmas

We have tested the approach on a PPP using realistic data for a single category of demand. For the PPP under consideration even around breaks simple order-up-to S rules apply and optimal order-up-to levels are easily read from simulation-based frequency tables.

By combining the results from the previous sections, the following conclusions can be drawn concerning the performance of the SDP-Simulation approach over a year which includes the breaks during Christmas, New Year's Day and the 4-day Easter weekend:

- Average annual shortage = 1.1 batch = 4.2 pools $< 0.1\%$
- Average annual outdating = 9.0 batches = 36 pools $< 1\%$

Compared to the current practice the potential savings are substantial: it seems that the current outdating figure of 15-20% can be reduced to less than 1%, while shortages arise only a few times per year.

4.4 The hospital case with no order costs

The SDP-Simulation approach provides insights into the structure of the optimal strategy. It shows whether an order-up-to S is appropriate, and if so, it suggests nearly optimal order-up-to levels for each weekday. The 'relative uncertainty' in the demand, as measured by the coefficient of variation, depends on the scale at which the problem plays. Therefore we may expect that an order-up-to S rule fits less well to an production-inventory problem that plays at a much smaller scale: i.e. at a much smaller blood bank or at a medium-large hospital.

Then fixed set-up costs per production run or fixed order costs may apply, as we will discuss in the next section. In this section we leave any fixed order costs out of the model, and focus on the impact of the demand uncertainty on the structure of the optimal policy, and the resulting outdating and shortage figures.

4.4.1 Problem and data

Thus far we have studied the BPP inventory problem using data of one of the four Dutch blood banks. Since other (foreign) blood banks might operate at a smaller scale, and surely do most hospitals do, we now consider the order problem at a much smaller scale. We refer to the problem in this section as the *Hospital case*, although we do not change the modeling of the order problem. Again, we assume a fixed lead time of one day: regular replenishment orders are placed early in the morning and these do arrive at the end of the day, say early the next morning.

We investigate the optimal stationary ordering policy for a hospital with an average demand that is 10 times as low as that for the blood bank considered in the previous section. The average weekly demand is thus 14.4 pools. Emergency orders to resolve shortages are delivered from the blood banks more or less instantaneously. The costs of emergency deliveries to resolve any shortages are fixed to 750 per BPP.

Again, we assume that the demand is Poisson distributed. On Mondays the mean demand is 2.6 pools and the coefficient of variation (cv) is thus $1/\sqrt{2.6} \approx 0.62$. Table 4.8 summarizes the daily demand characteristics. Since demands happens occasionally during weekends, the cv on some days can well exceed 1, but is on most days around 0.65. The cv of the demand over an entire week is $1/\sqrt{14.4} \approx 0.26$.

Since demands are Poisson distributed with low means, the coefficients of variation are very high. Therefore we expect an increase in the relative outdating and shortage figures

Table 4.8: Poisson demand distributions: means and coefficients of variation (cv).

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
original mean	2.6	2.1	3.2	2.1	2.6	0.8	1.0
cv	0.62	0.69	0.56	0.69	0.62	1.12	1.00

and the optimal policy to be 'more' stock-age-dependent.

- An open question to answer in the next sections is whether a simple order-up-to S rule is still close to optimal for such a small-scale problem.

Given that BPPs have to be transported from a blood bank to the hospital, daily ordering by the hospital may be inefficient. Instead of daily ordering, one may restrict the number of decision epochs to say Mondays, Wednesdays, and Fridays. An interesting question is thus:

- How does the optimal ordering strategy look like, when orders can be placed only on Mondays, Wednesdays and Fridays.

Hence additional production breaks are introduced on Tuesdays and Thursdays. Alternatively one may incur fixed order costs to reflect the costs involved with the transport of the pools from the blood bank to the hospital.

In Table 4.9, an overview is given of the several studies for the 'Hospital case'. Although, in The Netherlands, these costs are not (directly) billed to the hospital, the effect of fixed order costs is studied in Section 4.5.

Table 4.9: Overview of study of small scale problem.

<table>
<tr><td colspan="4">Hospital case or Small blood bank
'FIFO-demand' only : 14.4 BPPs per week</td></tr>
<tr><td colspan="2">No fixed order costs ($c^F = 0$)</td><td colspan="2">With fixed order costs c^F</td></tr>
<tr><td>Order daily
(Mon to Fri)</td><td>Order only on
Mon, Wed, Fri</td><td>low
$c^F = 75$</td><td>high
$c^F = 150$ or 500</td></tr>
<tr><td>Section 4.4.2</td><td>Section 4.4.3</td><td>Section 4.5.2</td><td>Section 4.5.3</td></tr>
</table>

4.4.2 Daily ordering

The formulation of the MDP is similar to that in Section 4.2.3. Since this problem plays at a small scale direct computation of an optimal strategy is possible without scaling the problem. Still we investigate the structure of the resulting optimal strategy in search for a practical rule. From simulation-based frequency tables (see Table 4.10) we read an ordinary order-up-to S rule with order-up-to levels $(11, 10, 11, 9, 11)$ for Monday to Friday. The order-up-to S rule fits to 47-66% of the states visited during the simulation. The performance of the rule and the optimal MDP strategy is evaluated by a long simulation for 100 million weeks. The main results and conclusions are:

1. Even under the optimal MDP policy, the relative outdating and shortage figures for the hospital case are much higher than those reported for the blood bank (in Table 4.3), due to the high coefficient of variation of the demand and the small scale at which the hospital problem plays:
   - outdating = 8.5%,
   - shortages = 1%.
2. An order-up-to S rule rule appears to perform nearly optimal with slightly more shortages:
   - outdating = 8.5%,
   - shortages = 1.3%,
   - The order-up-to S rule shows a 10% increase in costs compared to the cost-optimal MDP policy:
     - weekly shortage and outdating costs under MDP: 303,
     - weekly costs of order-up-to S rule: 337

     .

For now, we conclude that, even for the small scale problem under consideration, the order-up-to S rule performs quite well.

Table 4.10: Frequency tables from 1 million simulations of MDP policy with daily ordering independent of the fixed order costs.

(a) (State, action)-frequency table for Monday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_1)
Up-to S_1															
14												55	121		176
13										194	6543	4242			10979
12								226	44582	104321	22168				171297
11							31536	252641	185848						470025
10					59645	101872	134222								295739
9			13301	30744											44045
8		5276													5276
7	2463														2463
6															0
:															:
0															0
Freq(x)	2463	5276	13301	30744	59645	101872	165758	252867	230430	104515	28711	4297	121	0	1000000

(b) (State, action)-frequency table for Tuesday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_2)
Up-to S_2															
15														11	11
14												113	413	10	536
13											260	2744	1256		4260
12										204	10805	17370			28379
11								25	7719	184917	108229				300890
10					13789	39903	103890	167811	217563	119774					662730
9			8	1131											1139
8															0
:															:
1															0
0												1764	291		2055
Freq(x)	0	0	8	1131	13789	39903	103890	167836	225282	304895	119294	21991	1960	21	1000000

(c) (State, action)-frequency table for Wednesday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_3)
Up-to S_3															
14												4	15		19
13											599	3001			3600
12										21291	45764	829			67884
11								249	127232	323870	45411				496762
10						843	79943	156663	148673						386122
9				2819	10810	31400									45029
8			538												538
7		46													46
6															0
:															:
0															0
Freq(x)	0	46	538	2819	10810	32243	79943	156912	275905	345161	91774	3834	15	0	1000000

(d) (State, action)-frequency table for Thursday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_4)
Up-to S_4															
12												379			379
11										3378	18643				22021
10								8315	263097	77619					349031
9				17646	48379	86116	150772	319554							622467
8			4613												4613
7		76													76
6															0
:															:
1															0
0												1413			1413
Freq(x)	0	76	4613	17646	48379	86116	150772	327869	263097	80997	18643	1792	0	0	1000000

(e) (State, action)-frequency table for Friday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_5)
Up-to S_5															
16													48		48
15												584			584
14											2305	2197			4502
13									283	20966	43832	265			65346
12							67	17357	246413	150981	2899				417717
11					30246	78344	158916	220214	12319						500039
10			2382	9345											11727
9		37													37
8															0
:															:
0															0
Freq(x)	0	37	2382	9345	30246	78344	158983	237571	259015	171947	49036	3046	48	0	1000000

4.4.3 Order only on Mondays, Wednesdays and Fridays

When orders can be placed only on Monday, Wednesday and Friday, one can expect to order more at a time to anticipate the additional stationary ordering breaks on Tuesday and Thursday. Consequently one expect to observe an increase in both outdating and shortages. The frequency tables in Table 4.11 show the structure of the optimal strategy. Again a simple order-up-to S rule seems to fit reasonably well to the optimal MDP policy, with goodness-of-fit percentages close to 50%.

Both the optimal MDP strategy and the order-up-to S rules are simulated for 100 million weeks. In Table 4.12(a) and Table 4.12(b) we compare daily ordering on Monday to Friday against ordering on Monday, Wednesday and Thursday only.

We observe:

- Under the optimal MDP policy with ordering on Monday, Wednesday and Friday only the outdating and shortages rates are respectively 11.7% and 1.2%. Compared to daily ordering the additional breaks on Tuesday and Thursday thus have significant negative effect on the outdating of pools.

- Under the order-up-to S rule, with order-up-to levels $(13, 0, 12, 0, 12)$ for Monday to Friday, outdating and shortages are respectively 11.7% and 2.0%. An order-up-to S rule yields thus only 0.8% more shortages than the optimal stock-age-dependent strategy, but at the same amount of outdated BPPs.

- The resulting cost level of the order-up-to S rule is about $\frac{500-416}{416} = 20\%$ above the cost level of the optimal MDP strategy with ordering on Mondays, Wednesdays, and Fridays only.

The main conclusion is that the costs under daily ordering are much lower than when ordering on Tuesday and Thursday is prohibited: when applying an order-up-to S rule the gap is about 163 euros per week. When ordering would costs 75 euros per order, then daily ordering is $163 - 2 \cdot 75 = 13$ euros per week cheaper. But when the fixed order costs are twice as high, 150 euro per order, then not ordering on Tuesdays and Thursdays implies a cost saving of $163 - 2 \cdot 150 = 137$ compared to daily ordering.

Table 4.11: Frequency tables from 1 million simulations of MDP policy with ordering every Monday, Wednesday and Friday independent of the fixed order costs.

(a) (State, action)-frequency table for Monday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_1)
Up-to S_1															
18												44	78		122
17											1624	1984	258		3866
16										4860	8519	2065	426		15870
15									5845	31922	31769	5329			74865
14								4125	162516	114021	5688				286350
13						16628	153762	224439	80156						474985
12				18923	43547	71030									133500
11		2348	7184												9532
10	910														910
9															0
:															:
0															0
Freq(x)	910	2348	7184	18923	43547	87658	153762	228564	248517	150803	47600	9422	762	0	1000000

(b) (State, action)-frequency table for Tuesday: zero production.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_2)
Up-to S_2															
0						625	6171	29367	67457	129870	374060	309674	72958	9818	1000000
Freq(x)	0	0	0	0	0	625	6171	29367	67457	129870	374060	309674	72958	9818	1000000

(c) (State, action)-frequency table for Wednesday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_3)
Up-to S_3															
17														306	306
16													2150	279	2429
15												15339	8300	644	24283
14										2729	39526	43900	1103		87258
13									50893	208093	107255				366241
12						55810	101847	160175	158537						476369
11				10953	26371										37324
10			3987												3987
9		1276													1276
8	527														527
7															0
:															:
0															0
Freq(x)	527	1276	3987	10953	26371	55810	101847	160175	209430	210822	146781	59239	11553	1229	1000000

(d) (State, action)-frequency table for Thursday: zero production.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_4)
Up-to S_4															
0			1	216	4424	16363	41376	95387	153179	330485	268262	78406	11637	264	1000000
Freq(x)	0	0	1	216	4424	16363	41376	95387	153179	330485	268262	78406	11637	264	1000000

(e) (State, action)-frequency table for Friday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	Freq(S_5)
Up-to S_5															
17													119	32	151
16												872	446		1318
15											6817	4741	966		12524
14									379	22776	37499	7176			67830
13								427	120672	109750	12028				242877
12						141	163203	198773	70900						433017
11				35466	67646	112238									215350
10			16588												16588
9		6850													6850
8	3495														3495
7															0
:															:
0															0
Freq(x)	3495	6850	16588	35466	67646	112379	163203	199200	191951	132526	56344	12789	1531	32	1000000

Table 4.12: Impact of ordering on Monday, Wednesday and Friday only, compared to daily ordering on Monday to Friday.

(a) Order-up-to S rule.

Order-up-to S rule	S_1	S_2	S_3	S_4	S_5	Outdating	Shortage	Weekly costs
Daily	11	10	11	9	11	8.5%	1.3%	337
Only on Mon-Wed-Fri	13	0	12	0	12	11.7%	2.0%	500

(b) Optimal stock-age-dependent policy (MDP).

MDP	Outdating	Shortage	Weekly costs
Daily	8.5%	1.0%	303
Only on Mon-Wed-Fri	11.7%	1.2%	416

4.5 The hospital case with fixed order costs

In this section we will investigate the impact of making the order costs explicit in the MDP model. The modification of the cost structure in the MDP model is discussed in Section 4.5.1. We investigate the structure of the optimal ordering policy for two scenarios: one with fixed order costs of 75 euro per order, another with fixed order costs of 150 euro per order. Numerical results are presented in Sections 4.5.2 respectively Section 4.5.3.

4.5.1 Optimal control under fixed order costs

An optimal ordering policy can be computed in a similar way as described in Section 4.2.3. The SA algorithm is unaltered except that $C(d, \mathbf{x}, k)$ is to be replaced by $C(d, \mathbf{x}, a, k)$:

$$C(d,\mathbf{x},k) = \begin{cases} c^{\mathrm{F}} \cdot \mathrm{I}(a > 0) & \text{fixed order costs,} \\ c^{\mathrm{O}} \cdot (x_m - k)^+ & \text{outdating costs,} \\ +c^{\mathrm{S}} \cdot (k - \sum_{r=1}^{m} x_r)^+ & \text{shortage costs,} \\ +c^{\mathrm{H}} \cdot x & \text{holding costs.} \end{cases} \tag{4.10}$$

After a sufficiently large number of iterations, say N, an the optimal strategy with fixed order costs is known, or, if optimal actions are not stored, is to be computed from:

$$\pi(d,\mathbf{x}) = \arg \min_{a \in \mathcal{A}(d,\mathbf{x})} \left(c^{\mathrm{F}} \cdot \mathrm{I}(a > 0) + \sum_k p_d(k)\, V_{N-d}(d+1, y(\mathbf{x}, k, a)) \right), \tag{4.11}$$

with $V_N(d, \mathbf{x})$, the expected costs over an horizon of $N - d$ days when the the horizon starts in state $(d, \mathbf{x})$.

After simulation of the resulting optimal policy, one aims in the SDP-Simulation approach at finding a more simple nearly optimal rule. From the literature review, we may expect that an (s, S) policy could be nearly optimal. This is to be checked, and further one need to find for each working day d nearly optimal parameter values s_d and S_d. Hopefully these can be read from the simulation-based frequency tables.

4.5.2 Structure optimal policy when $c^{\mathbf{F}} = 75$

For numerical results we use the following costs figures:

- c^{F} = fixed order costs: 75 per order,
- c^{O} = proportional outdating costs: 150 per outdated pool,
- c^{S} = proportional shortage costs: 750 per pool short.

The fixed order costs of $c^{\mathrm{F}} = 75$ euro is a reasonable estimate of the costs involved in the (scheduled) transport of one or more pools from a blood bank to the hospital. Since multiple deliveries at hospitals can be combined in a route by the same vehicle, the fixed order costs do not relate to a single trip from a blood bank to a hospital. When the hospitals are geographically more dispersed the transportation costs should be set higher. Higher fixed order costs are considered in the next section.

An optimal stock-age-dependent strategy is computed and simulated to investigate its structure. In a simulation of 1 million weeks we count for each observed total stock level x, how often the MDP policy implies a certain order-up-to level. The results are tabulated in the frequency tables in Table 4.13. Optimal order size 0 is made explicit by translating it into order-up-to level 0, as reported in the next-to-last row of each table.

At first sight the tables look similar to those in Table 4.10, but the great difference is read in the next-to-last row of each table from which we read how often the optimal strategy implies zero ordering. Especially on Tuesday and Thursday we observe that in more than 75% respectively 51% of the states visited the optimal order size is zero.

Reading an (s, S)-rule:

An (s, S)-policy seems to apply, since zero ordering is prevalent when the stock level x exceeds some threshold value s. The thresholds are visualized in the Table 4.13(a) to Table 4.13(e) by the additional vertical line. For example, the threshold on Thursday is $s_4 = 7$: when $x > 7$ no orders are placed. Positive order quantities on Thursdays are found only when $x \leq 7$; the optimal order quantity can be approximated by an order-up-to level $S_4 = 9$, which is the most-frequent positive order-up-to level. Apparently for the given cost structure the minimum order-quantity on Thursday is $S_4 - s_4 = 2$ BPPs.

Similarly one reads thresholds s_d and order-up-to levels S_d for the other weekdays. On Friday no threshold is found. On Monday, Tuesday and Wednesday, the optimal strategy

Table 4.13: Frequency tables from 1 million simulations of MDP policy with $c^{\mathrm{F}} = 75$.

(a) (State, action)-frequency table for Monday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	Freq(S_1)
Up-to S_1																
18									1			14	32			47
17								10			34	90				134
16							8			59	368	18				453
15						2			161	5251	375					5789
14					1			912	32311	57378						90602
13				1			133	139106	195198							334438
12				24140	51679	97652	176127	110056								459654
11		3541	9682													13223
10	1539															1539
:																0
0										52174	35584	6018	345			94121
Freq(x)	1539	3541	9682	24141	51680	97654	176268	250084	227671	114862	36361	6140	377	0	0	1000000

(b) (State, action)-frequency table for Tuesday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	Freq(S_2)
Up-to S_2																
16											88	7				95
15										830	229					1059
14									3706	3012						6718
13								8861	8680							17541
12							27055	20073								47128
11						12065	985		67639							80689
10				2131	5189	3030	16968	44963	27564							99845
9			773													773
8	82	228														310
:																0
0										167138	306259	204034	59062	9326	23	745842
Freq(x)	82	228	773	2131	5189	15095	45008	73897	107589	170980	306576	204041	59062	9326	23	1000000

(c) (State, action)-frequency table for Wednesday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	Freq(S_3)
Up-to S_3																
16									19							19
15								82								82
14							19									19
13						6				85540						85546
12					1				84241	109089						193331
11				4494			71595	156108	148289							380486
10			1128		14066	37276	14488									66958
9		279														279
8	72															72
:																0
0										42389	180333	44803	5503	180		273208
Freq(x)	72	279	1128	4494	14067	37282	86102	156190	232549	237018	180333	44803	5503	180	0	1000000

(d) (State, action)-frequency table for Thursday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	Freq(S_4)
Up-to S_4																
10							93309	29467								122776
9				11265	32605	133097	71791	109159								357917
8			3407													3407
7		1281														1281
6	561															561
:																0
0									167515	287990	49481	8763	309			514058
Freq(x)	561	1281	3407	11265	32605	133097	165100	138626	167515	287990	49481	8763	309	0	0	1000000

(e) (State, action)-frequency table for Friday.

Stock x	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	Freq(S_5)
Up-to S_5																
17											4		42			46
16										157	23	456				636
15									927	778	2684	343				4732
14								1519	3833	14525	12191	344				32412
13							1530	5962	33791	61866	8473					111622
12					518	4741	104397	196390	196975	47995						551016
11				14453	47594	99474	83962	48177								293660
10			4409													4409
9		1092														1092
8	352															352
:																0
0												23				23
Freq(x)	352	1092	4409	14453	48112	104215	189889	252048	235526	125321	23375	1166	42	0	0	1000000

Table 4.14: Goodness-of-fit of (s, S)-policy at order costs $c^{\mathrm{F}} = 75$.

(s, S)-policy	Mon	Tue	Wed	Thu	Fri
Threshold s_d	9	8	9	7	S_5
Order-up-to levels S_d	12	10	11	9	12
Goodness-of-fit	55%	85%	65%	87%	55%
Freq. $x \leq s_d$ (do order)	96%	25%	77%	49%	100%

is more stock-age-dependent such that a threshold might not hold uniformly. In some states the order quantity is positive even when the total stock level exceeds the threshold. The threshold s_d on day d is chosen such that at column $x = s_d + 1$ in the majority of states the production is zero and at column $x = s_d$ production is positive at the majority of states. For determining the most-frequent order-up-to level one should restrict to the states where $x \leq s_d$. (Alternatively one may compute the average order-up-to level over all positive (non-zero) order-up-to levels.)

The goodness-of-fit of the (s, S)-ordering strategy is the frequency of states in which the (s, S)-policy prescribes the same order quantity as the optimal MDP policy. The (s, S)-policy on Thursday fits to $357,917 + 514,058 = 871,975$ out of the 1 million states. The goodness-of-fit of the (s, S)-rule on Thursday with $(s_4 = 7, S_4 = 9)$ is thus 87%.

In Table 4.14 we summarize the best parameter values as read from the frequency tables, when the fixed order costs are 75 euro. On Friday there seems to apply no threshold s_5, to emphasize this we report $s_5 = S_5$. Considering the goodness-of-fit values of the (s, S)-policy we conclude that it resembles for a great part the optimal MDP strategy. The last line in the table shows how frequent orders are placed. Virtually every Monday and Friday an order is placed. Only on 25% of the Tuesday BPPs are ordered. On 77% of the Wednesdays and 49% of the Thursdays an order is placed.

In Table 4.15 we compare three policies with respect to their shortage and outdating rates and the average number of orders placed per week. The cost optimal MDP policy gives the best trade-off between these three criteria for the given cost structure. The (s, S)-policy and order-up-to S rule performs nearly-optimal with respect to the shortage and outdating rate. The (s, S)-policy results on average in 1 order more over a period of $\frac{1}{0.2} = 5$ weeks. Ignoring the thresholds s_d results in much more frequent ordering: 4.3 orders per week against 3.4 orders under the optimal strategy.

Compared to the results for ordering on Mondays, Wednesdays, and Fridays only, as reported in Table 4.12, the savings on outdating and shortage is significant. Compared to daily ordering and ignoring any fixed order costs, outdating and shortages have increased, but the number of orders placed per week is much lower: for the optimal stock-age-dependent policy the average number of orders drops form 5 to 3.4 times per week.

Table 4.15: Performance of different policies when $c^{\mathrm{F}} = 75$

$c^{\mathrm{F}} = 75$	Outdating	Shortage	# orders per week	Total weekly costs
Optimal MDP policy	10%	1.2%	3.4	601
(s, S)-policy:				
$s = (9, 8, 8, 7, S_5)$				
$S = (12, 10, 11, 9, 12)$	10%	1.2%	3.6	633
Order-up-to S				
$S = (12, 10, 11, 9, 12)$	10%	1.1%	4.3	660

Conclusions

- The SDP-Simulation approach with fixed order costs successfully solves the MDP and finds nearly optimal (s, S)- policies.
- Order frequencies on Tuesday and Thursday are (much) lower than on the other weekdays.
- On Fridays an ordinary order-up-to S rule applies.

In the next section we investigate the impact of higher fixed order costs on the structure of the optimal strategy.

4.5.3 Higher fixed order costs

We consider two case in which the fixed order costs are increased to $c^{\mathrm{F}} = 150$, respectively $c^{\mathrm{F}} = 500$ euro per order.

Doubled fixed order costs: $c^{\mathbf{F}} = 150$

When the fixed order costs are doubled, one tends to order less frequently and thus more at a time. Consequently one may expect the BPPs in stock to be somewhat older and that outdating thus becomes more prevalent. The optimal strategy likely is thus more stock-age-dependent than when the fixed order costs are low.

In terms of the (s, S)-policy, the thresholds s are expected to be lower and the order-up-to levels S are expected to be higher than for the case with $c^{\mathrm{F}} = 75$. The minimum order quantity $S_d - s_d$ is thus higher when the order costs are increased.

The above expectations are supported by numerical results obtained by the SDP-Simulation approach with the fixed order costs equal to 150 euros. The nearly optimal parameter values that one reads for each working day are summarized in Table 4.16. Compared to Table 4.14 where the fixed order costs were 75, the thresholds s of the (s, S)-policy are indeed somewhat lower and the order-up-to levels S are slightly raised, but not dramatically. Although the optimal strategy is a bit more stock-age-dependent, given that in relatively more states where $x > s_d$ the order-size is positive, the (s, S)-policy fits even better to the optimal strategy when $c^{\mathrm{F}} = 150$ than for $c^{\mathrm{F}} = 75$.

In Table 4.17 we compare the performance of the (s, S)-rule against that of the optimal stock-age-dependent policy. The (s, S)-rule performs nearly optimal, besides an additional 0.3% of the demand (or on average 2 BPPs per year) cannot be met from stock. The average number of orders per week is only 2.9.

Table 4.16: Goodness-of-fit of (s, S)-policy at order costs $c^{\mathrm{F}} = 150$.

(s, S)-policy	Mon	Tue	Wed	Thu	Fri
Threshold s_d	8	7	8	6	S_5
Order-up-to levels S_d	13	12	11	9	12
Goodness-of-fit	85%	94%	78%	81%	52%
Freq. $x <= s_d$ (do order)	87%	15%	50%	41%	99.9%

Table 4.17: Performance of different policies when $c^{\mathrm{F}} = 150$.

$c^{\mathrm{F}} = 150$	Outdating	Shortage	# orders per week	Total weekly costs
Optimal MDP policy	11%	1.2%	2.9	836
(s, S)-policy:				
$s = (8, 7, 8, 6, S_5)$				
$S = (13, 12, 11, 9, 12)$	11%	1.5%	2.9	865

Very high fixed order costs: $c^{\mathrm{F}} = 500$

When the fixed order costs are extremely high, say 500 euro, we observe that the unnatural (s, Q) policy then performs equally well as the more natural (s, S)-policy. In Table 4.18 we report on the performance when the order costs are extremely high (500 euro per order). In general, shortages increase when fixed order costs are raised. The optimal strategy limit shortages since its order quantities depend on the ages of the BPPs in stock. When ages are ignored as in the stock-level-dependent policies (e.g. the (s, Q) and (s, S)- policies), then shortages are much more prevalent. The resulting average weekly costs of the stock-level-dependent policies are more than 20% above the optimal cost level. If this is unacceptable one should search for better rules by developing simple rules that are stock-age-dependent, similarly as we did in Section 3.4.

Table 4.18: Performance of different policies when $c^{\mathrm{F}} = 500$

$c^{\mathrm{F}} = 500$	Outdating	Shortage	# orders per week	Total weekly costs
Optimal MDP policy	17%	1.8%	2.0	1,600
(s, S)-policy:				
$s = (7, 6, 6, 6, 8)$				
$S = (14, 14, 14, 13, 13)$	18%	4.8%	1.9	1,943
(s, Q)-policy:				
$s = (7, 6, 6, 6, 8)$				
$Q = (10, 8, 11, 9, 8)$	19%	4.9%	1.9	1,952

4.6 Summary and conclusions

The contribution of this chapter is two-fold:

1. The SDP-Simulation approach is extended such that it can deal with non-stationary periods (e.g. additional production breaks during holidays), in a further stationary horizon.

2. It is shown how the SDP-Simulation approach derives efficient ordering rules when average demand volumes are much smaller than in the previous chapter. In particular, the problem is studied with fixed order costs.

We have studied the PPP in a setting where no distinction is made between young-demand and any-demand: any age-preference regarding the issued BPPs is ignored. All demand is met by issuing the oldest BPPs first (FIFO), and no mismatch costs apply. In the next two subsections, we summarize the main conclusions for a number of cases.

Production breaks at blood banks

We have investigated the impact on the ordering strategy of some additional production breaks that occur during a year: i.e. on Good Friday, Easter Monday, the two Christmas days and New Years Day. Therefore, the mean demand figures from one of the Dutch blood banks are used. In fact, a bit more difficult case is studied as we assume the demand uncertainty to be somewhat higher than under Poisson distributed demands. We draw the following conclusions:

- Additional outdating and shortages on the day(s) after a production break are to be accepted even under the truly-optimal MDP policy.

- An order-up-to S rule resembles the optimal policy even better on the day prior to an additional production break than on ordinary weekdays.

- When the production break last $m - 1$ days, the order problem on the day before the break is a single period problems. The best order-up-to level S just before the break is then a bit higher than the one suggested by a Newsboy equation, as the initial stock will not survive until the next order moment.

- Simulation results over a year, including the production breaks during Christmas, New Year's Day and the 4-days Easter weekend, indicate that compared to the current practice overall outdating and shortage figures can be reduced significantly:
  - outdating from 15-20% to less than 1%,
  - shortages from about 1% to less than 0.1%).

Fixed order costs for hospitals

For Poisson distributed demands the coefficient of variation (cv) of the demand is much higher when the mean demand figures are much lower. Consequently, the trade-off between shortages and outdating is then much more difficult to make. We have executed a numerical study for a case where the demand is only 14.4 BPPs per week, or on average only 2 per day. The cv is then about 0.7. Typically, one may think of a hospital that keeps BPPs in stock, or a small blood bank.

When any fixed ordering costs are neglected, the SDP-Simulation shows that the optimal ordering policy is still well structured: an ordinary order-up-to S rule performs nearly optimal. Nevertheless, for the given data the outdating is substantial: 8.5% under both the optimal policy and the order-up-to S rule. When operating at this scale, the outdating cannot be reduced more through the ordering policy, unless one accepts more shortages.

As daily ordering may be not efficient, regarding the efforts to set-up an order and to arrange the transportation of BPPs to an hospital, ordering on Monday, Wednesday and Friday only will reduce the number of orders to place. Alternatively, we can add any set-up and transportation costs in the form of a fixed order costs per order. The fixed order costs can be high when the hospital is quite isolated from other hospitals and blood banks.

The SDP-Simulation approach successfully solves the PPP with fixed order costs. From the simulation-based frequency tables (s, S) policies can be read. On day d a threshold value s_d applies for ordering: if the stock level is above s_d, no orders are placed. When the total stock level x is on or below the threshold value $S_d - x$ BPPs are ordered.

1. Incorporating the fixed order costs in the MDP model results in cost optimal balancing of shortage, outdating and order costs,
2. An (s, S)-ordering strategy appears to be nearly optimal even for the given relatively short fixed shelf life and the high uncertainty in the demand,

3. Optimal parameter values for the seven thresholds s_d and seven order-up-to levels S_d are easily generated by the SDP-Simulation approach,

4. When the fixed order costs are increased the optimal ordering strategy becomes more stock-age-dependent: the stock-level-dependent (s, S)-policy results in many shortages and is far from optimal when the order costs are extremely high.