8

Stochastic Inventory Control for Product Recovery Management

Erwin A. van der Laan[1], Gudrun Kiesmüller[2], Roelof Kuik[1], Dimitrios Vlachos[3], and Rommert Dekker[4]

[1] Rotterdam School of Management / Faculteit Bedrijfskunde, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands, elaan@fbk.eur.nl, rkuik@fbk.eur.nl

[2] Faculty of Technology Management, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands G.P.Kiesmueller@tm.tue.nl

[3] Department of Mechanical Engineering, Aristotle University of Thessaloniki, P.O. Box 461, 54124 Thessaloniki, Greece, vlachos1@auth.gr

[4] Rotterdam School of Economics, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands, rdekker@few.eur.nl

8.1 Introduction

Essentially, inventory management concerns the process of deciding on 1) how often to review stocks, 2) when to replenish stocks, and 3) how much to replenish. This basic focus of inventory management persists in the presence of item returns that can be recovered and then used for servicing demand. However, the details and complexities with which the three basic decisions manifest themselves can, and usually do, differ greatly due to the presence of recoverable-item flows. This, and the practical relevance of inventory management with recoverables, warrants the development of inventory theory that explicitly includes flows of recoverable items.

A first categorization lists the foundations of the differences between recoverable-item inventory management (RIIM) and traditional inventory management under three headings.

- *Multiple sourcing.* When items can be recovered, requirements of items can be met from multiple sources: the source(s) with newly manufactured or ordered items and the source(s) with recovered items. This extends the three basic decisions with a fourth one: 4) where to replenish from. Added complexities arise since the source of recoverables is capacitated: at any point in time the number of items available from recovery is limited. The sources may not only differ in item-availability, but also in per unit price and supply reliability.

- *Absence of monotonicity.* Return streams are often difficult to control. This means that between regular reviews and replenishments, inventories can go up because of product returns even when inventory levels are sufficiently high to maintain a targeted service level. Therefore (heuristic) analysis based on observation of the inventory just prior to the epochs that a replenishment is due may be highly inaccurate. A further consequence of the loss of monotonicity is that the total inventory in the system is not bounded from above, unless an appropriate disposal policy is implemented.
- *Unreliable sources.* The added complexity of multiple sourcing is frequently compounded by the circumstance that sources of recovered items are unreliable concerning availability, as the item returns may be uncertain both in timing and quality, and thus in suitability for recovery.

This chapter gives an overview of the different approaches that have been presented in the scientific literature to account for the above peculiarities of stochastic inventory control for product recovery. Before discussing the literature in more detail, we first present two examples from practice, each representing a different planning problem.

Case A: Commercial Returns at a Mail-order Company

A large mail-order company in Western Europe (see Mostard and Teunter, 2002) faces a difficult inventory problem for their fashion products. Lead times are long, so well before the selling season starts a replenishment order is placed to accommodate the demand for the whole season. Return rates are usually around 40%, but can be as high as 75%. Therefore, returns really need to be taken into account when determining the order size. To obtain more accurate demand forecasts, preview catalogs are sent out to a selected group of customers. On the basis of customer orders placed in response to the preview mailing, another replenishment order is placed just before the start of the selling season. Additionally, an (expensive) emergency order can be placed some three weeks into the season that makes use of the (limited) information regarding sales and returns during this beginning of the season. At the end of the season, shortages result in lost sales and overages in obsolete products that have to be disposed of. Since fashion products are highly seasonal, sales and return volume forecasts are very crude. Since the last opportunity to order is just three weeks after the season's start, there is little opportunity to adapt to realized demand.

Case B: Product Remanufacturing at Volkswagen

Car parts that have failed during operation on the road, varying from injection pumps to complete engines, are collected by Volkswagen via the car dealers (see Van der Laan, 1997). These parts are subsequently remanufactured by a third party and eventually resold as spare parts for approximately half of the

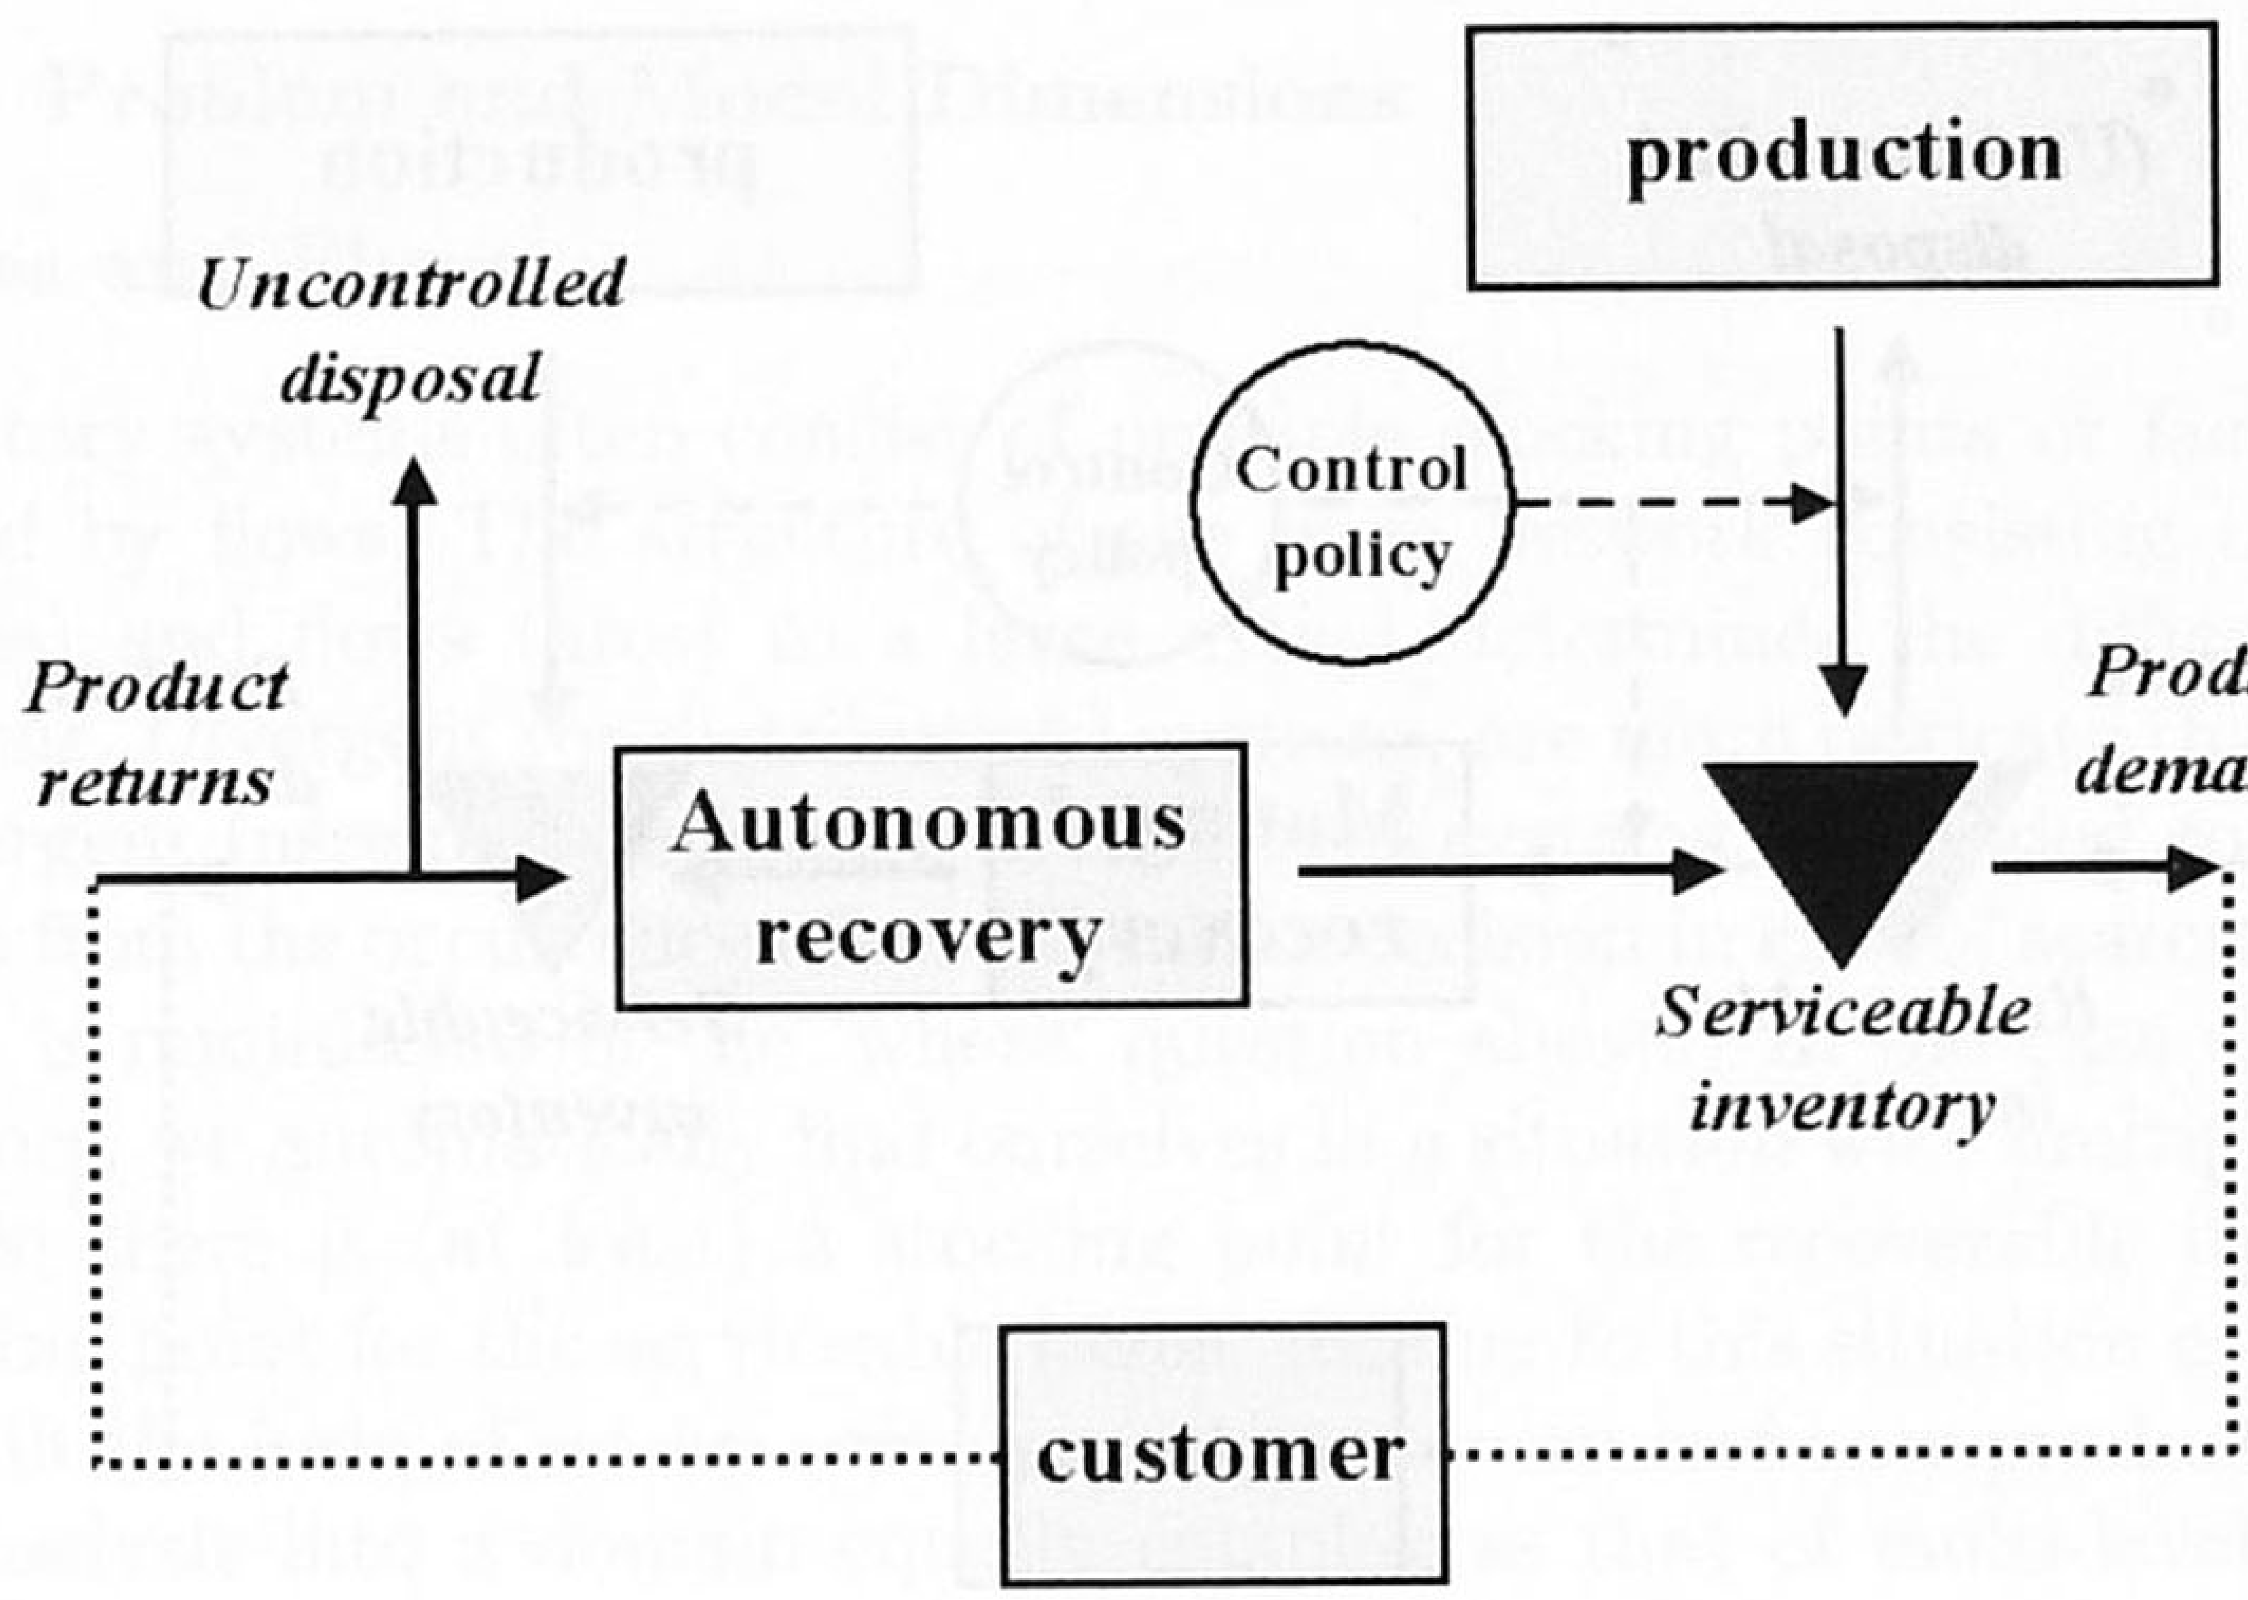

Fig. 8.1. Autonomous recovery

price of newly manufactured parts. Due to normal fluctuations in both demand and supply, and due to a lack of control of the return stream, there may be periods in which the number of collected car parts does not match the number of demands. Then, in times of supply shortage, new parts are ordered from the Volkswagen factories, or, in times of supply overage, recoverable products are disposed of. Lead times for production and remanufacturing may differ. As is the case here, it occurs often that it is this interaction between the supply of new and remanufactured products that makes inventory control more difficult than traditional, single-source inventory control.

Case A admits only a limited number of replenishment decisions, whereas Case B asks for recurrent replenishment decisions. Moreover, in the mail-order company case, returns can be recovered with little management intervention. In the Volkswagen case, the company can pursue an active recovery strategy involving the batching and timing of remanufacturing. This distinction, between cases in which returned items can be made serviceable with little control and cases where management pursues an active recovery policy also seems reflected in the status of the theory of models for these cases. The first type of cases, to which we will refer as the cases of *autonomous recovery*, seems to be more amenable to analysis than the second type of cases, to which we will refer as the cases of *managed recovery*. Note that the cases of autonomous recovery do away with the need for answering the fourth basic question: where to replenish from?

So, the autonomous recovery cases are often those situations in which recoverable products only need to undergo minor operations, such as cleaning and repackaging. These operations are relatively cheap so that it is not really

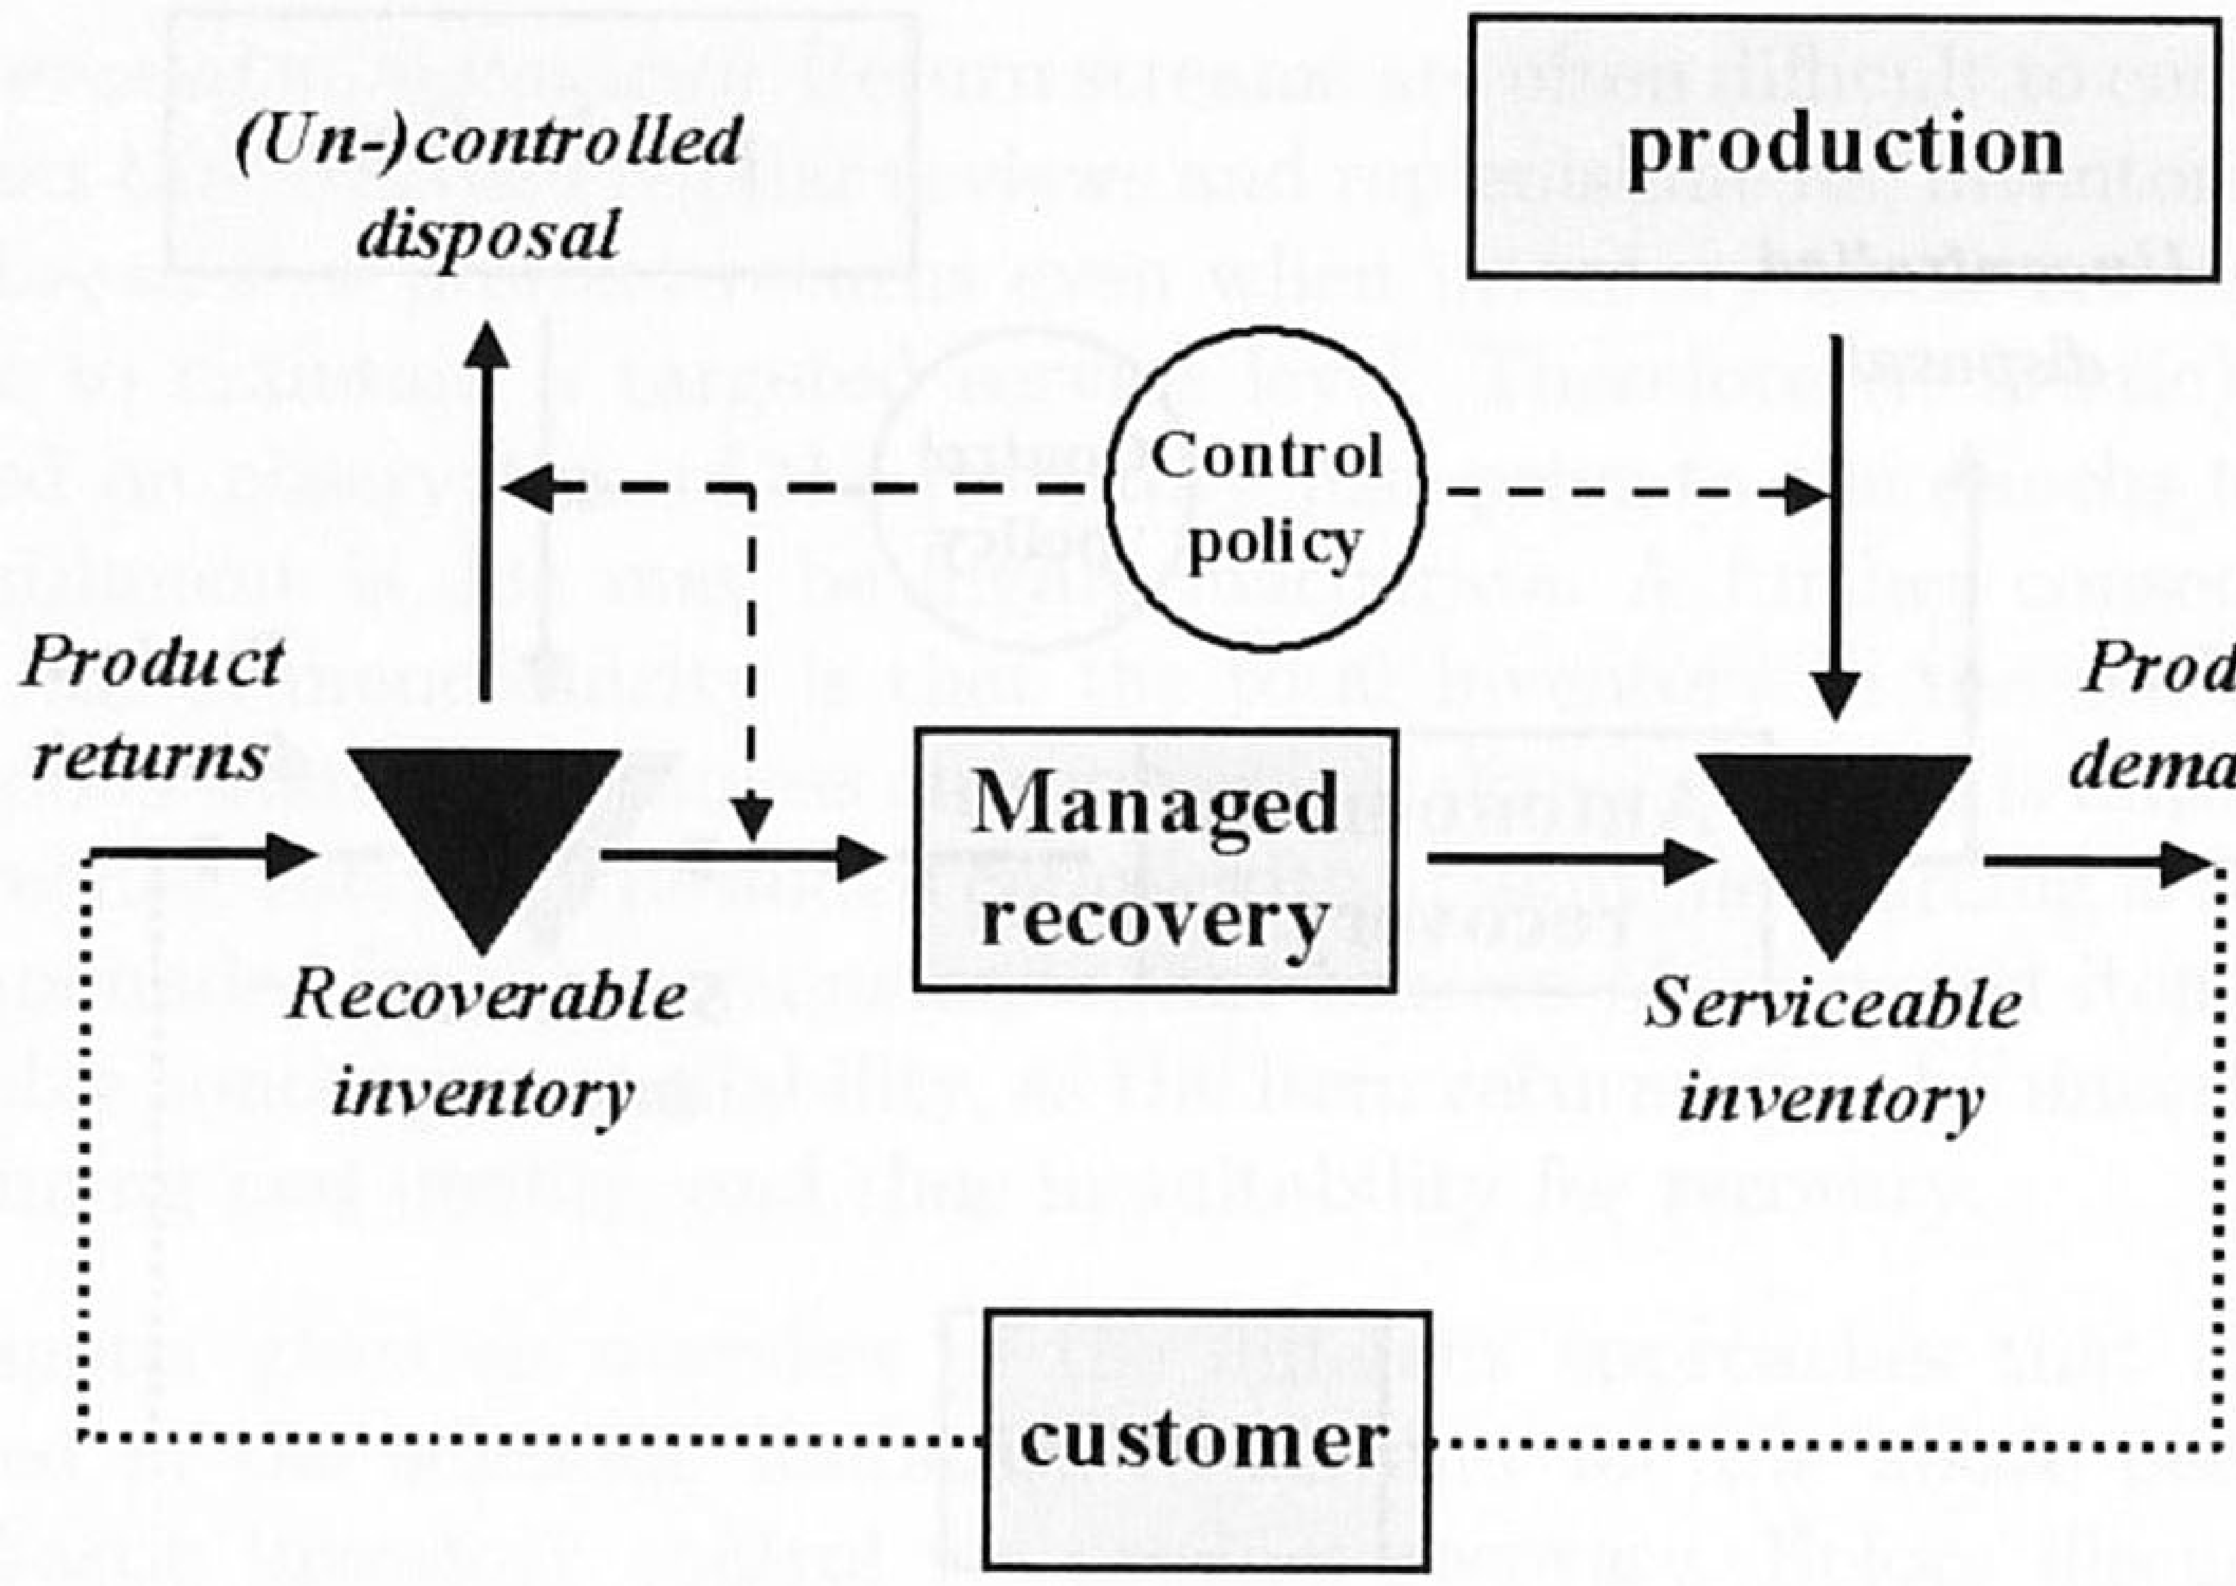

Fig. 8.2. Managed recovery

necessary to delay these operations until actually needed: upon return, items are immediately prepared for reuse and there is no further decision-making involved, hence the provision 'autonomous'. For example, commercial returns are often in very good condition and many of them can be reused directly without much recovery.

As Case B shows, the remanufacture of end-of-life returns calls for extensive recovery operations. The costs involved with these value-adding operations may be such that it is more economical to stock the returns until they are really needed. Thus, cash outflows are postponed and there is less risk of loss of investment due to obsolescence. In this case, decision-making on recovery operations occurs on a continuous basis, hence the provision 'managed'.

Since in the managed recovery cases one recovers items when deemed necessary, these cases are also referred to as *pull strategies* for inventory. In this vein, the autonomous cases are referred to as *push strategies*. Sometimes recovery management is extended with disposal management. An appropriate disposal policy attempts to bound the total inventory in the system. In practice, this is primarily relevant when the return rate is close to or in excess of the demand rate.

Models for inventory management can be further categorized on several dimensions. Below, we will discuss several of these dimensions and discuss the relationship to the dichotomy 'autonomous-managed'.

8.1.1 Problem and Model Dimensions

Nodes and Flows

Inventory systems often consist of multiple stocking points or facilities connected by flows. The structure of the open network consisting of facilities (nodes) and flows (arcs) to a large extent determines the difficulty of the analysis. Divergent (or distribution) systems, are more intricate than serial or convergent (assembly) systems. In divergent systems, the added complication stems from the occurrence of an allocation decision in case of scarcity of items. (This is reminiscent of the 'where' question above.) In the case of managed recovery, we automatically find ourselves in a situation with multiple stocking points: there is (at least) a stocking point for the recoverable items and a stocking point for the serviceable items. Adding to this situation economies of scale in the form of set-up costs for both recovery and new production brings the analysis into a domain equally complex as that of multi-level inventory management with fixed costs. As, in general, no rigorous results on the optimality of policies in the latter situation are known, it is no surprise that the same holds for the case of managed recovery. Even if fixed recovery costs are absent, the situation already seems to be so complex that one needs to resort to heuristics to find inventory strategies (see Section 8.3).

The case of autonomous recovery seems to be best understood. This is especially true in the case of an inventory system consisting of just one inventory facility from which demand is served (see Section 8.2).

The Model of Time

In the literature, there is a clear distinction between a discrete modelling approach and a continuous modelling approach of time. However, the choice for the time model seems mostly to rest on pragmatic grounds. Use of a continuous time model sometimes presents advantages in computational issues. For example, the continuous time approach seems to be more flexible with respect to cost structure and lead time assumptions. So when it comes to numerical analysis, as will be the case for managed returns, continuous time models are more suitable. Section 8.3.2 reports on a general framework for carrying out numerical analysis in a continuous-time setting.

Costs

As far as costs modelling is concerned, the literature on inventory theory with recoverables recognizes the same costs structures as traditional inventory theory. So, besides linear inventory related costs and variable costs of operation, some models consider fixed production and/or recovery costs and some do not. The discussion on how to set the holding cost parameters in recovery models is less straightforward than in traditional single source models, but this discussion is left to Chapter 11.

Decisions

Also in the domain of decision modelling, the inventory theory with recoverables follows largely traditional inventory theory. So most models incorporate the opportunity to release (re-)manufacturing or replenishment orders. There is one additional decision, though, that is sometimes considered in inventory theory with recoverables (and not in traditional inventory theory). This additional decision is the disposal decision.

The distinction between autonomous and managed recovery (or push and pull strategy) seems to be the most fundamental modelling choice in inventory theory including recoverable returns. To honor this, the following two sections of this chapter follow this dichotomy. Note that the distinction is fundamentally based on a choice related to the way models deal with the issue of multiple sourcing.

8.2 Autonomous Recovery

In the case of commercial returns, often only minor actions like inspection or cleaning suffice to render the returned products suitable for direct reuse. This section deals with the situation that no separate stocking facilities of returned products are necessary. If more complex processing is needed and/or the decision maker has the option to delay these activities, then we are in the case of managed recovery. That is the theme of Section 8.3.

The discussion will be restricted to the situation that the inventory system consists of only one stock facility serving end demand. The evaluation of base stock policies (in the absence of fixed costs) for autonomous product returns in a series or assembly system can be found in the working papers DeCroix and Zipkin (2002a) and DeCroix and Zipkin (2002b), the discussion of which is outside the scope of this chapter.

8.2.1 Naive Netting

The naive way of dealing with product returns is the so-called 'netting on averages'. In this approach, one does not take the returns process into account explicitly. Instead, one assumes that returns are deducted from future demand thus cancelling out part of that demand. The remainder of the *expected* demand is treated with traditional methods for single source inventory control. In principle, this method works and yields fair results in case of small return rates (see Van der Laan et al., 1996a)) or high correlation between demands and returns. However, for high return rates and low correlation between demand and returns, the result can be way off the correct one. The reason for this is that netting on averages blatantly ignores the additional variability that is introduced by the return process. The method correctly reduces the

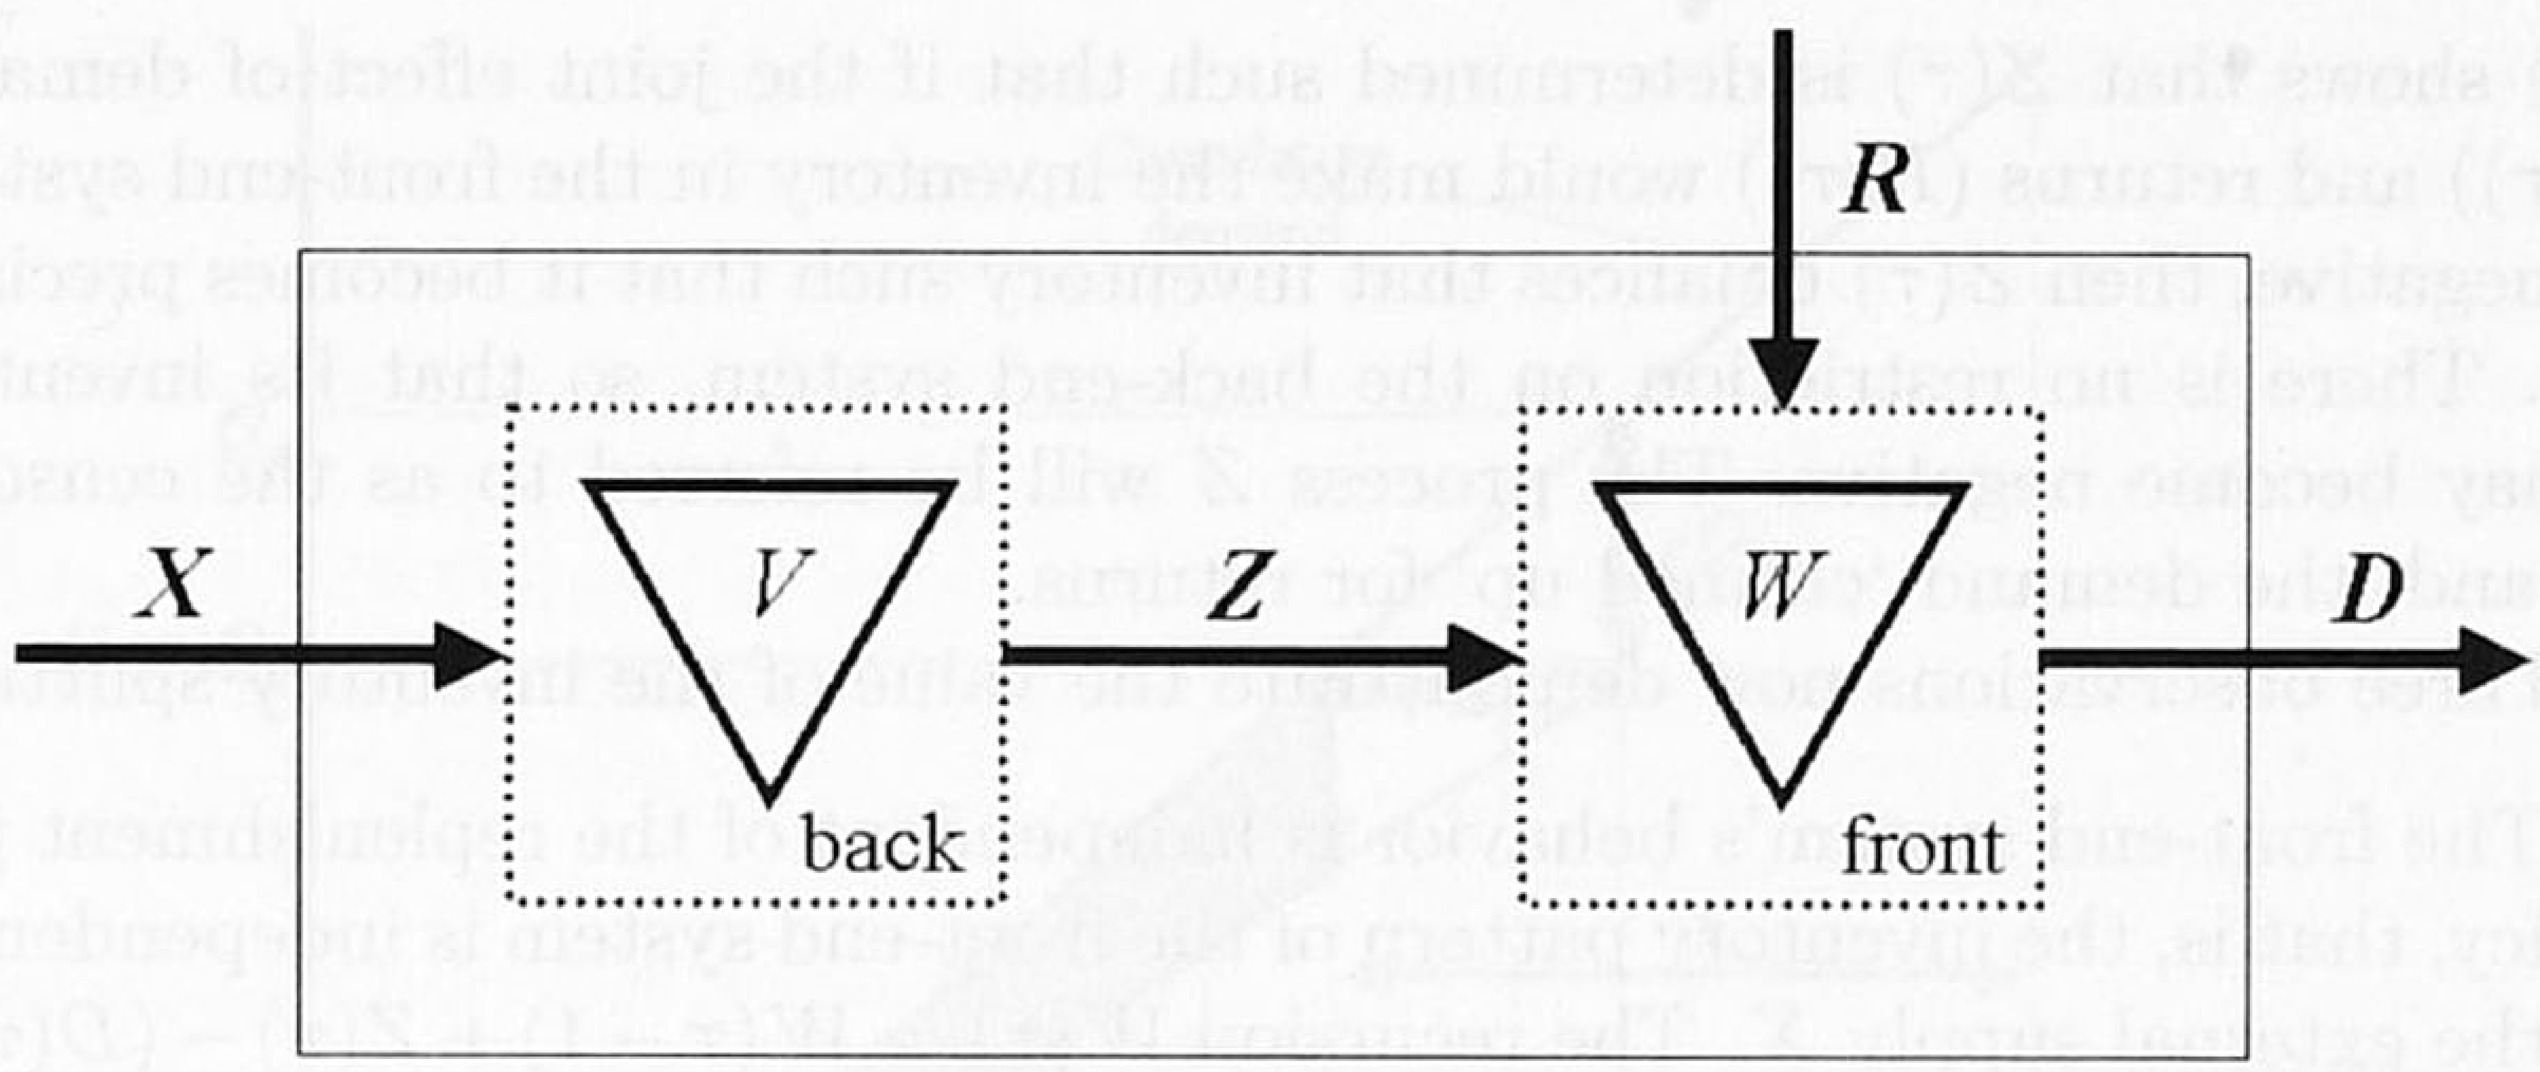

Fig. 8.3. Split inventory model

average demand with the expected returns, but deals incorrectly with the variability of the processes.

8.2.2 Sophisticated Netting: Split Inventory

Naive netting reduces the problem under returns to the problem without returns at the cost of inaccuracy. Perhaps surprisingly, we can do much better by using a netting of the processes and yet cast the problem in one with only non-negative demand. For ease of exposition let us assume that time is discrete. Demand in period τ is written as $D(\tau)$ and returns in period τ as $R(\tau)$. Let $X(\tau)$ be the supply of products from the production facility. Thus, the inventory balance equation for the facility's inventory at the end of period τ, $I(\tau)$ reads

$$I(\tau) = I(\tau - 1) + X(\tau) - (D(\tau) - R(\tau))$$

We now split the inventory $I(\tau)$ into two components, $V(\tau)$ and $W(\tau)$, the sum of which will give back the inventory, through the recursion

$$\begin{cases} V(\tau) \;= V(\tau-1) + X(\tau) - Z(\tau) \\ W(\tau) = W(\tau-1) + Z(\tau) - (D(\tau) - R(\tau)) \end{cases} \tag{8.1}$$

where

$$Z(\tau) \equiv \max\{0, D(\tau) - R(\tau) - W(\tau-1)\} \tag{8.2}$$

and the initial condition $(V_0, W_0) = (I_0, 0)$. The meaning of these recursive equations is best illustrated through Figure 8.3.

The figure shows that the inventory system is divided into two subsystems: a front-end system and a back-end system. It illustrates that (8.1) serves as the balance equation for the front-end system with inventory level W and the balance equation for the back-end system with inventory level V. Expression

(8.2) shows that $Z(\tau)$ is determined such that if the joint effect of demands $(D(\tau))$ and returns $(R(\tau))$ would make the inventory in the front-end system, W, negative, then $Z(\tau)$ balances that inventory such that it becomes precisely zero. There is no restriction on the back-end system, so that its inventory V may become negative. The process Z will be referred to as the censored demand, the demand 'cleaned up' for returns.

Three observations now demonstrate the value of the inventory splitting.

1. The front-end system's behavior is independent of the replenishment policy, that is, the inventory pattern of the front-end system is independent of the external supply X. The recursion $W(\tau) = W(\tau-1) + Z(\tau) - (D(\tau) - R(\tau)) = \max\{0, W(\tau-1) - (D(\tau) - R(\tau))\}$ shows that its dynamics are in fact a random walk on a discrete half-line with random step size $R - D$.
2. The flow Z is independent of production input X. The recursion $Z(\tau) = \max\{0, D(\tau) - R(\tau) - W(\tau-1)\}$ tells us that Z measures the virtual undershoot when the random walk hits the end point of the half-line at 0. In short, the front-end system's behavior, inclusive of its input Z, can be described independently of the replenishment flow X, that is, it can be described and analyzed independently of the replenishment policy. The only system affected by the inventory policy is the back-end system.
3. The back-end system's behavior is that of a traditional inventory system: it faces non-negative demand (represented by Z) and has the opportunity to replenish through the flow X.

Before turning to applications of the inventory splitting framework we make two remarks.

Remark 1. Not only are the behaviors of the processes interesting but so are the costs involved. Suppose the inventory related costs take the form $G(I)$. How do such costs translate into the model with split inventory? In particular, how can we translate these costs into costs for the back-end system, especially the inventory V, which still needs to be optimized? In the case in which the processes are stationary the answer is simple. Introduce the cost function H as

$$H(v) = \mathrm{E}_{(W|v)}(G(v + W)),$$

where $\mathrm{E}_{(W|v)}$ is the expectation with respect to $(W|v)$, that is, with respect to the equilibrium of the process $\boldsymbol{W}$ conditioned on $V = v$.

Remark 2. Clearly, the inventory splitting can be used to transform the deterministic lot-sizing problem with time-varying signed demand (Wagner-Whitin problem with signed demand) into one with only nonnegative demand. Such use of the splitting was made in Kelle and Silver (1989a).

8.2.3 Single-decision Models

In this section, we discuss situations in which a single replenishment decision needs to be taken. This problem occurs in two typical cases. In the case of

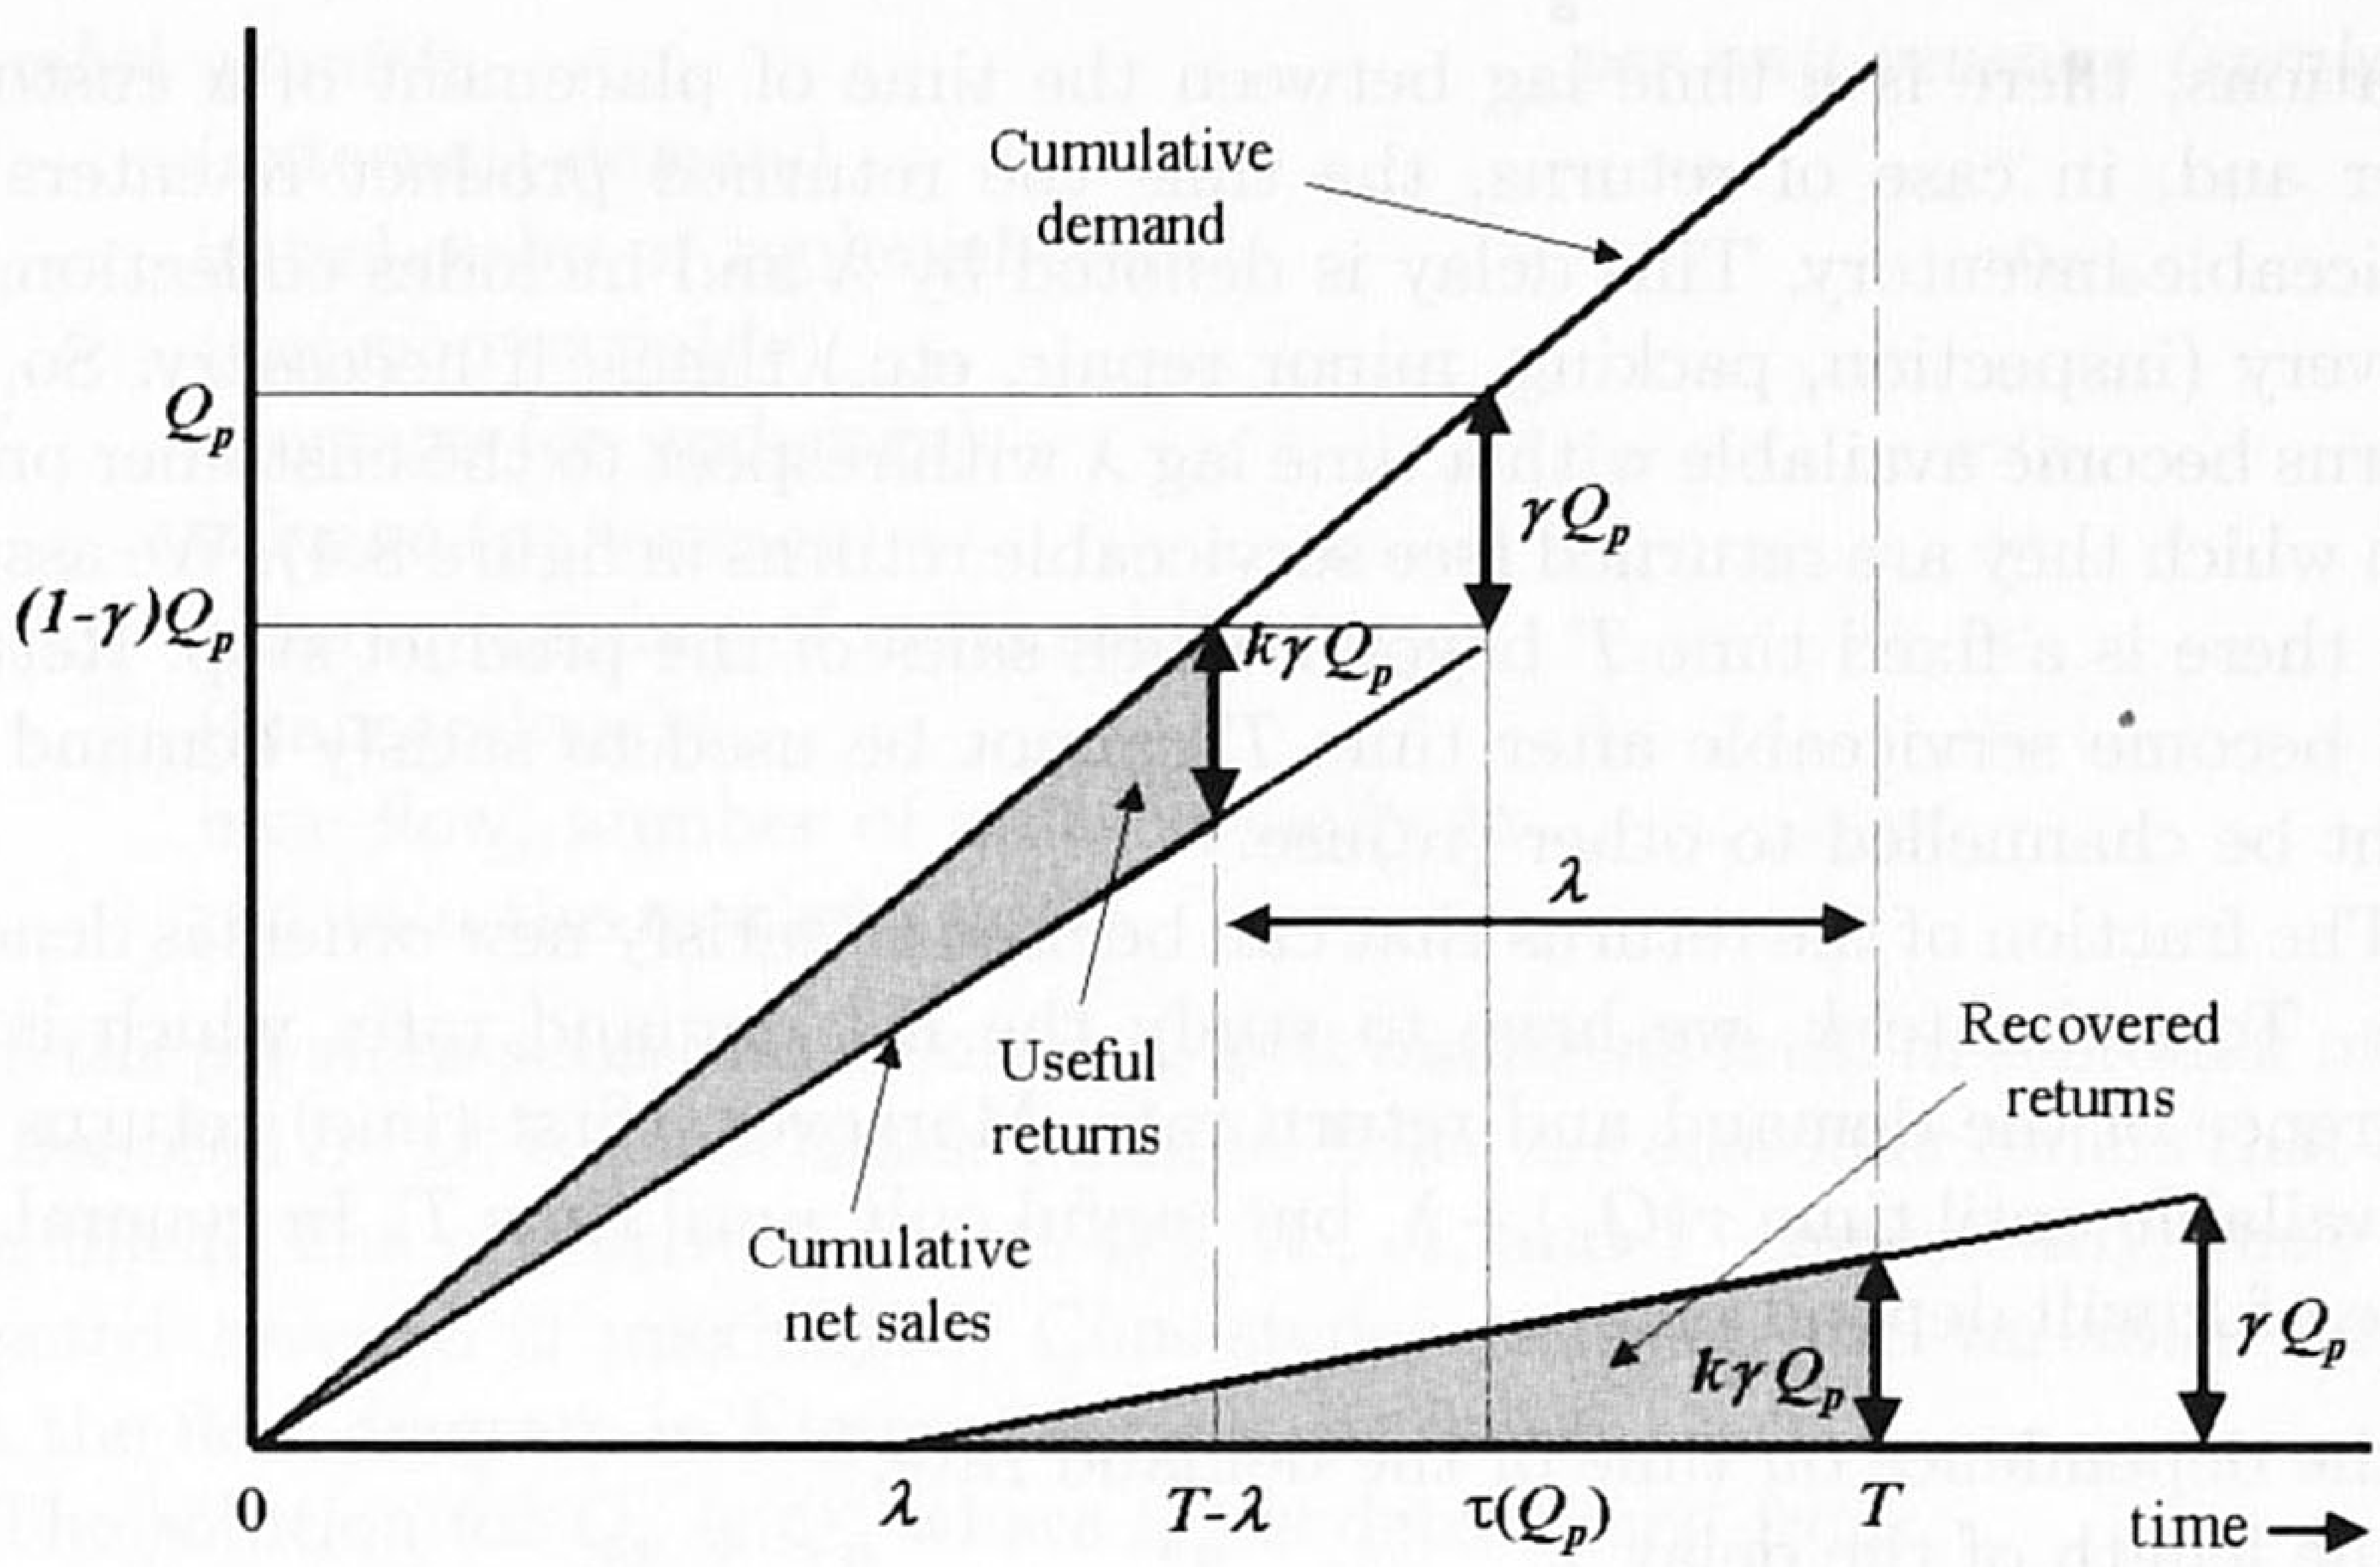

Fig. 8.4. Cumulative orders, net sales, and returns for a constant demand and return rate

fashion products (see Case A in the introduction of this chapter), one needs to determine the initial order quantity before the start of the selling season. Another case is the final order problem, where a manufacturer makes a final production run for his parts/finished products, because he switches over, or already has switched over, to new models. In both cases, the selling period is known in advance. These cases seem similar to the 'news-vendor' problem, but there is a major difference: product returns. Because of the returns, the filled customer orders (transactions) differ from the customer orders filled and not returned by the customer (net sales).

The operation of a single-decision inventory system with returns is explained through Figure 8.4, in which cumulative flows of orders, net sales, and returns are depicted. In Figure 8.4, the solid black line represents the cumulative demand, which attains its maximum value at the end of the selling period, denoted by T.

Suppose one replenishes by ordering (or producing) Q_p units at the beginning of the period over which demand needs to be met. This quantity covers demand until epoch $\tau(Q_p)$. If it were not for recoveries, any demand beyond this epoch would be lost. The total net sales from the initial quantity Q_p are $(1-\gamma)Q_p$ at time $\tau(Q_p)$. Here γ is the return fraction: the fraction of sales that is returned. The surface between cumulative demand and cumulative net sales represents the cumulative returns of the initial order. The maximum value of these returns is uQ_p.

The returns, so it is assumed, may be used to satisfy further demand occurring after $\tau(Q_p)$ and thus can be used to increase (net) sales. In practical situations, there is a time lag between the time of placement of a customer order and, in case of returns, the time the returned product re-enters the serviceable inventory. This delay is denoted by λ and includes collection and recovery (inspection, packing, minor repair, etc.) times, if necessary. So, the returns become available with a time lag λ with respect to the customer orders from which they are returned (see serviceable returns in figure 8.4). We assume that there is a fixed time T beyond which sales of the product stop. Returns that become serviceable after time T cannot be used to satisfy demand but might be channelled to other (re)use.

The fraction of the returns that can be used to satisfy new orders is denoted by k. To estimate k, we have to study the net demand rate, which is the difference of the demand and return rate. Moreover, (first-time) returns will be available until time $\tau(Q_p) + \lambda$, but useful only until time T. In general, the value of k will depend on

- the dependence on time of the demand rate,
- the length of the delay λ,
- the length of the sales period T,
- and the initial order Q_p,

and will be difficult to uncover. However, under the assumptions that 1) the demand rate is constant (as displayed in the figure), 2) the return rate is constant, and 3) products are only returned once, the figure above illustrates that finding k is not difficult. The upshot is that

$$k = \begin{cases} \dfrac{\max\{0, T - \lambda\}}{\min\{Q_p/d, T\}} & \text{if } T \leq \lambda + \frac{Q_p}{d} \\ 1 & \text{if } T > \lambda + \frac{Q_p}{d}, \end{cases}$$

where $d = \frac{D}{T}$ is the demand rate.

Calculation of the Optimal Initial Order Quantity Given k

Now assume that the fraction of returns that can (potentially) be used to service demand is a given value k (independent of the initial order). The subsequent analysis will correspond closely to that of the 'news-vendor' problem. We therefore briefly set forth this analysis first. The following symbols are used.

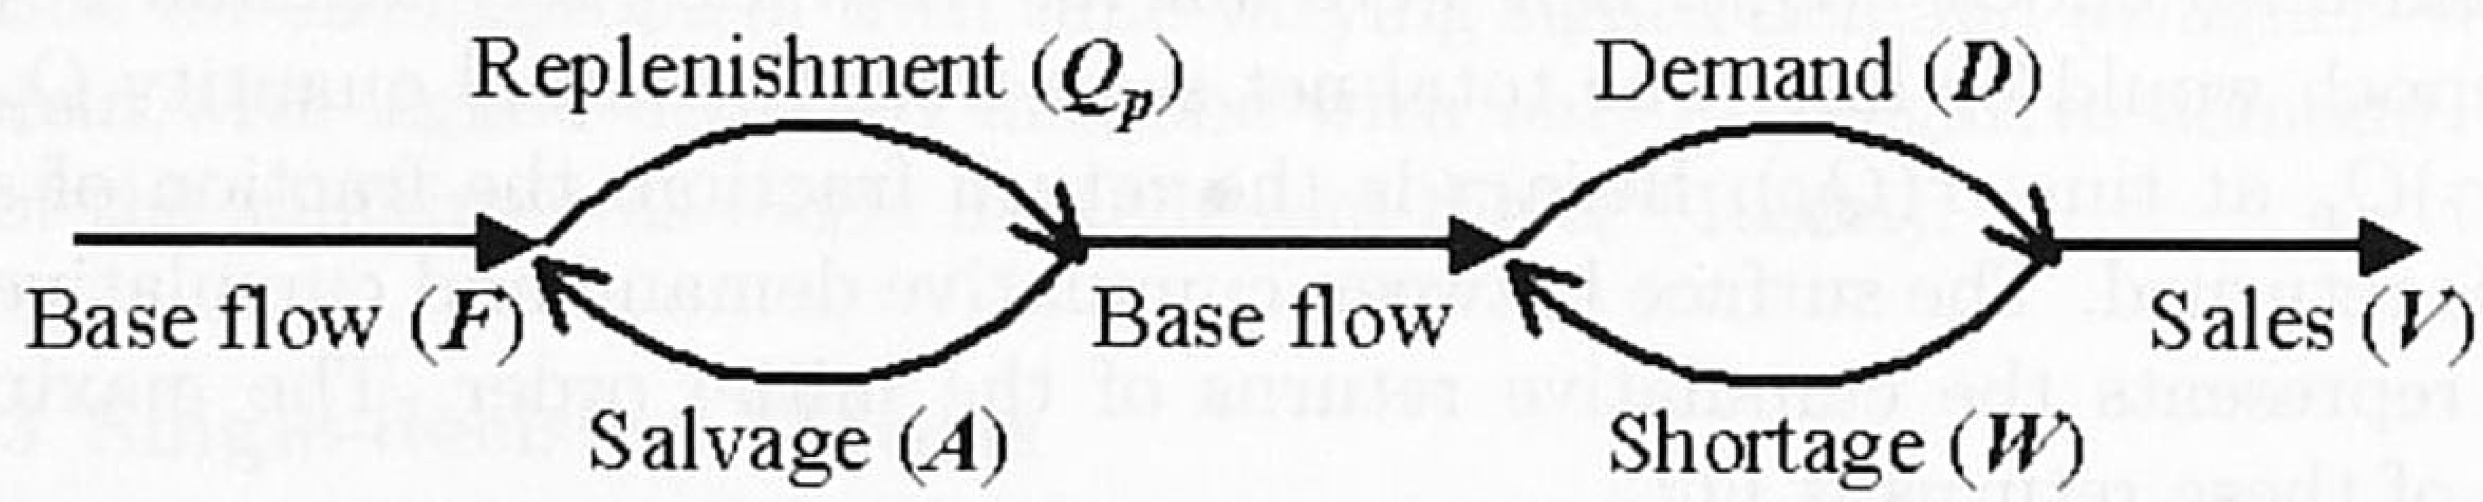

Fig. 8.5. 'News-vendor' problem

symbol	*quantity*	*per unit revenue (symbol)*
D	(external) demand	–
Q_p	initial order or replenishment (decision variable)	$-c_p$
W	shortage (or underage)	$-c_b$
A	salvage (or overage)	ω
V	sales or number of units sold (transactions)	v
F	base flow, number of units actually input into the market (sales)	–

As in the previous subsection, demand, as it manifests itself in customer orders and denoted by D, is a stochastic variable. Sales are customer orders that have been filled. The objective is to set Q_p, W, A, and V consistently, such that expected revenue is maximized. Consistency between the variables implies that the flow diagram in Figure 8.5 must be valid.

The solution for Q_p is Q_p^* where Q_p^* is determined from

$$\text{Prob}(D \leq Q_p^*) = \frac{v - c_p + c_b}{v - \omega + c_b}$$

(see e.g. Silver et al., 1998, p. 387). Of course, the unit costs coefficients appearing in the optimal value for the initial replenishment are related to the flows as given in the figure through the table that introduces them. Now consider the case that includes returns. Returns are a fixed fraction γ of sales. The determination of the optimal initial order quantity, which maximizes the total system profit, depends on the collection and recovery strategy and on the costs involved in all the activities.

Several collection and recovery strategies can be distinguished. The following paragraphs discuss some of these in detail.

Full Recovery and Unlimited Reuse

Suppose that all returns that are suitable for re-use are recovered (full recovery). The following additional symbols are used to model the situation.

symbol	*quantity*	*per unit revenue (symbol)*
U	number of items returned	$-c_r - v$
R	number of returned items used for satisfying demand	–
A'	number of items salvaged after return	ω
Z	censored demand	?
Ω	censored shortage	?

Furthermore, assume that there is no limit[5] on the number of times a unit can be returned and recovered (unlimited reuse). In the analysis, we assume

[5] Also assume that the fraction of returns that is actually used for satisfying demand is k irrespective of the number of times items return.

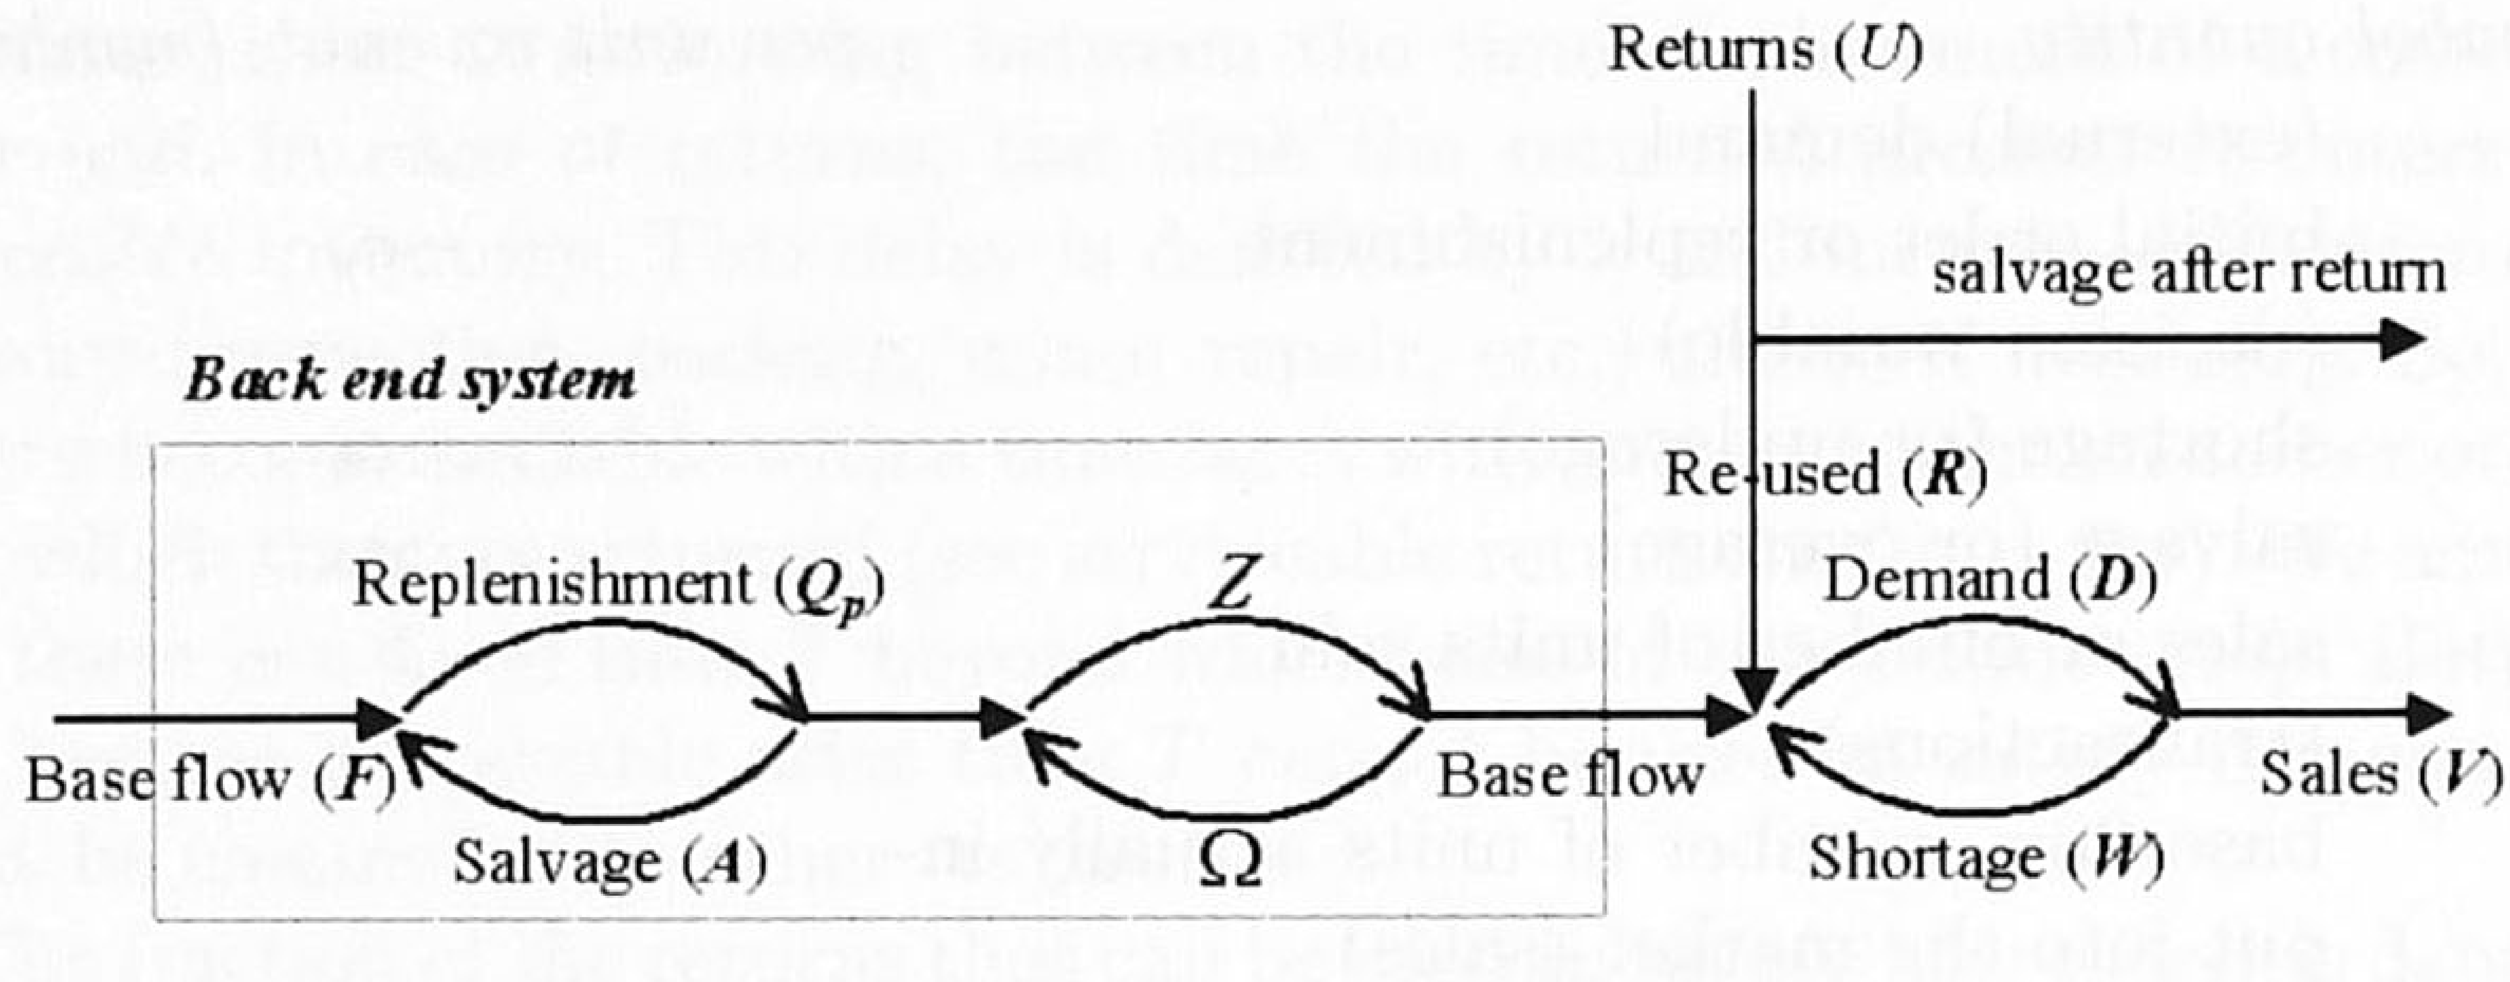

Fig. 8.6. Split model for full recovery and unlimited returns

that demand can be modeled as a continuous variable and can also be filled 'continuously', that is, in non-integer quantities.

We follow the route outlined above, using the split inventory[6] model. However, we have to be a little precise in distinguishing demand from sales, so the splitting requires some extra structure. Also, making returns reuseable requires some processing. Consequently, the part of the model involving returns requires somewhat more care than in the standard case. The ultimate model takes the form as sketched in Figure 8.6.

We first discuss the flows. The auxiliary flows Z and Ω are virtual. We can thus manipulate their values at will as long as consistency with the flow diagram is maintained. Several relationships exist between the recovered flow (R), the censored demand flow (Z), the sales flow (V), the demand flow (D), and the shortage flow (W). For example,

$$\begin{aligned} U &= \gamma V \quad ; F + R + W = D \quad ; \qquad V = D - W; \\ R &= kU \quad ; A' = (1-k)U = (1-k)\gamma V \end{aligned}$$

Now

$$Z - \Omega = F = D - W - R = V - R = (1-k\gamma)V = (1-k\gamma)(D-W) = \frac{1}{\beta}(D-W)$$

where $\beta \equiv \frac{1}{1-k\gamma}$. There is logically no other constraint than $Z - \Omega = F$ that needs to be imposed on the auxiliary flows. However, we choose to put $Z = \frac{D}{\beta}$ and $\Omega = \frac{W}{\beta}$. Second, we turn to costs. A priori the auxiliary flows carried no costs or revenues. We now put a cost of βG per unit on the flow Ω while making the unit costs of W zero at the same time. Then the costs of operating the system remain unaltered. What is the revenue per unit of flow of F that exits the back-end system? As $F = \frac{V}{\beta}$, we can impute a unit costs of β times the revenue of a unit of V while putting the revenue of V to zero without altering the costs. So the question becomes: what is the revenue of

[6] The notion of 'inventory' is somewhat spurious in this context, as there are no factual items held at the end of the period.

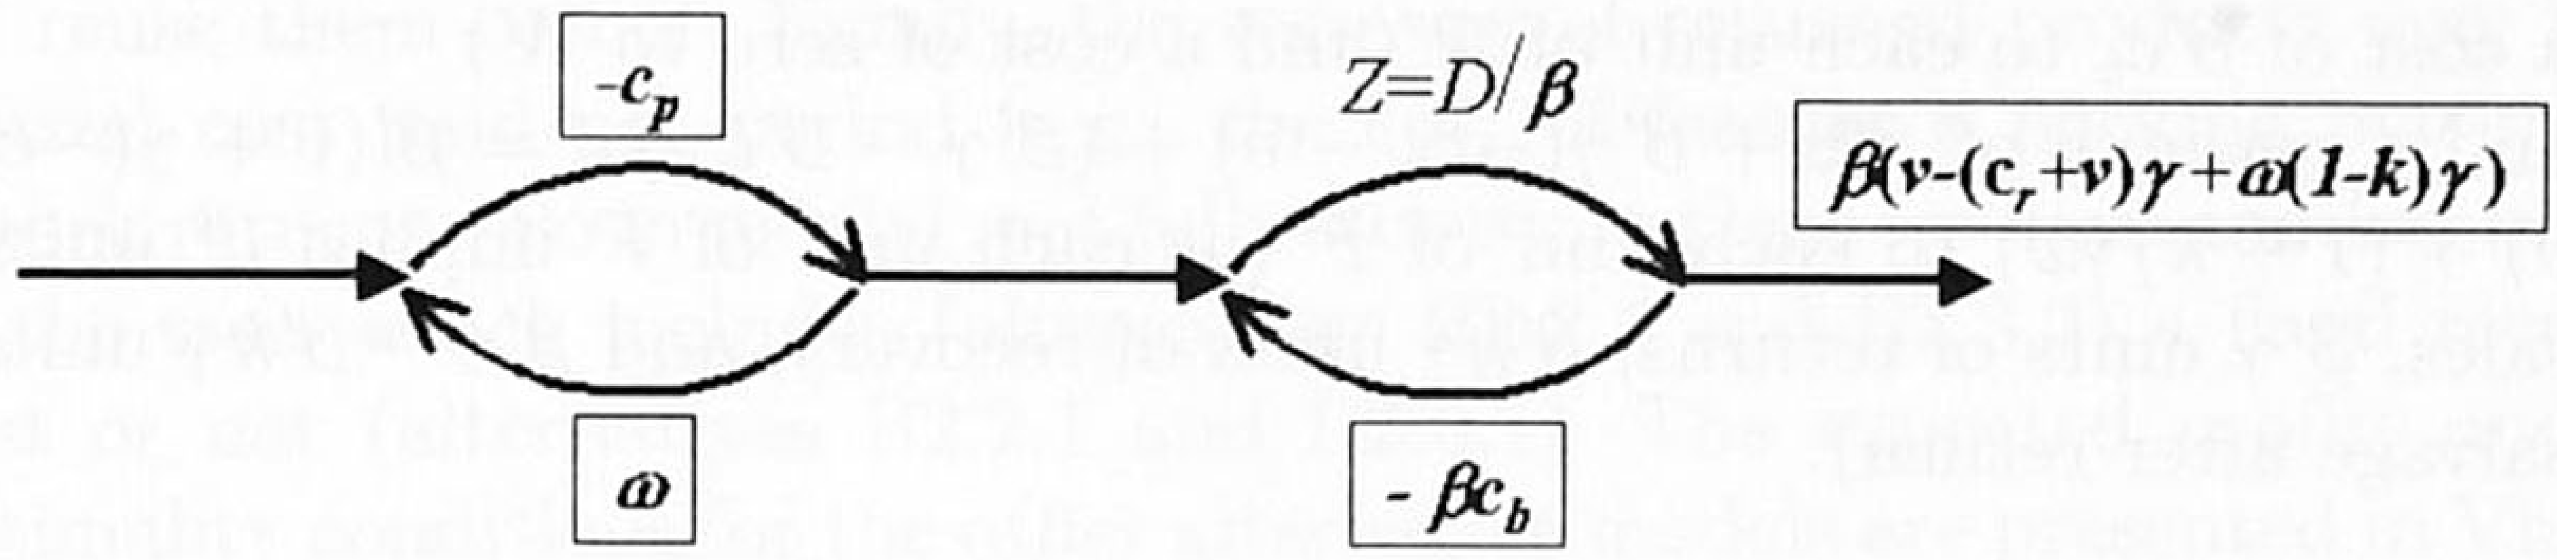

Fig. 8.7. The emulation of the return model: cost and revenue allocation

a unit of V? A unit of flow V implies γ units of returns, $(1-k)\gamma$ units of salvage after returns (A'), and γk units of recovered items. So a unit of V's revenue is $v-(c_r+v)\gamma+\omega(1-k)\gamma$.

Now note that the back-end system (see Figure 8.6) is just a standard 'news-vendor' system with the flow F playing the role of sales in the model without returns. So we in fact need to analyze a standard 'news-vendor' model with demand and costs per unit (boxed values) as given in Figure 8.7. Finding the optimal value for Q_p is now easy. Just apply the unit cost substitution rules $v \leftarrow \beta(v-(c_r+v)\gamma+\omega(1-k))\gamma$, $c_b \leftarrow \beta c_b$ and $c_p \leftarrow c_p$. Accounting for the factor β in the demand, the standard 'news-vendor' solution now gives the rule: set $Q_p = Q_p^*$ where Q_p^* is determined from

$$\begin{aligned}\text{Prob}(\frac{D}{\beta} \leq Q_p^*) &= \frac{\beta(v-(c_r+v)\gamma+\omega(1-k)\gamma)-c_p+\beta c_b}{\beta(v-(c_r+v)\gamma+\omega(1-k)\gamma)-\omega+\beta c_b} \\ &= \frac{(v'+c_b-\omega)\beta+\omega-c_p}{(v'+c_b-\omega)\beta} \equiv \Theta(\beta),\end{aligned}$$

where $v'=(1-\gamma)v+\gamma(\omega-c_r)$.

Full Recovery and One-time Re-use

Under one-time reuse, returns can only be reused once. Once more, we employ the split inventory model (see Figure 8.8).

From one-time recovery, we obtain $U=\gamma F$. From this it follows that $R=kU=k\gamma F$. The flow conditions s(ee Figure 8.8) then further yield

$$Z-\Omega = F = D-W-R = V-R = V-k\gamma F$$

and so $V=(1+k\gamma)F=\beta' F$, where we define β' as $\beta'=1+k\gamma$. We require

$$\beta'(Z-\Omega)=\beta' F = V = D-W.$$

Once more we have some latitude in setting values for Z, the censored demand, and Ω, the censored shortage. We choose to set $Z=\frac{D}{\beta'}$ and $\Omega=\frac{W}{\beta'}$. We now determine costs per unit of Z and Ω which turn the return model into an equivalent standard 'news-vendor' problem. Put

- a cost of $\beta' c_b$ to each unit of Ω (and a cost of zero to W)
- and a revenue of $\beta' v + \beta'\gamma(-c_r - v) + (\beta'\gamma - \beta' k\gamma)\omega = \beta'(v + \gamma(-c_r - v) + (1-k)\gamma\omega)$ to each unit of F (as each unit of F implies β' units of sales, $\beta'\gamma$ units of returns, $\beta' k\gamma$ units of recovery, and $\beta'\gamma - \beta' k\gamma$ units of salvage after return).

The result for determining the optimal value, Q_p^*, of the initial replenishment can now be written down using the result for the standard 'news-vendor' problem as

$$\begin{aligned}\text{Prob}(\frac{D}{\beta'} \leq Q_p^*) &= \frac{\beta'(v + \gamma(-c_r - v) + (1-k)\gamma\omega) - c_p + \beta' c_b}{\beta'(v + \gamma(-c_r - v) + (1-k)\gamma\omega) - \omega + \beta' c_b} \\ &= \frac{(v' + c_b - \omega)\beta' + \omega - c_p}{(v' + c_b - \omega)\beta'} = \Theta(\beta')\end{aligned}$$

with as before $v' \equiv (1-\gamma)v + \gamma(\omega - c_r)$. Note that $1 + k\gamma$ is the expansion to first order in $k\gamma$ of $\frac{1}{1-k\gamma}$. The case where items can be reused precisely n times (and where the system is operated with a full-recovery strategy) then allows that one can apply the foregoing end result for obtaining the optimal initial order by substituting $\beta \leftarrow 1 + k\gamma + (k\gamma)^2 + \cdots + (k\gamma)^n$.

Other Cases

Vlachos and Dekker (2000) have studied many more variants of the single decision problem under returns. Additions are fixed costs for recovery operations and the option to recover only a fraction of the returns suitable for recovery. A tree listing the variants analyzed is given in Figure 8.9.

The first branch (B1) of the tree corresponds to the first model examined, while the second model refers to branch B2.1. Branches B2.2 and B2.3 assume that the recovery cost is significant. The difference between them is that in the first one (partial recovery) this cost is paid for every returned item we reuse, while in the second (full recovery) this cost is paid for all returns whether

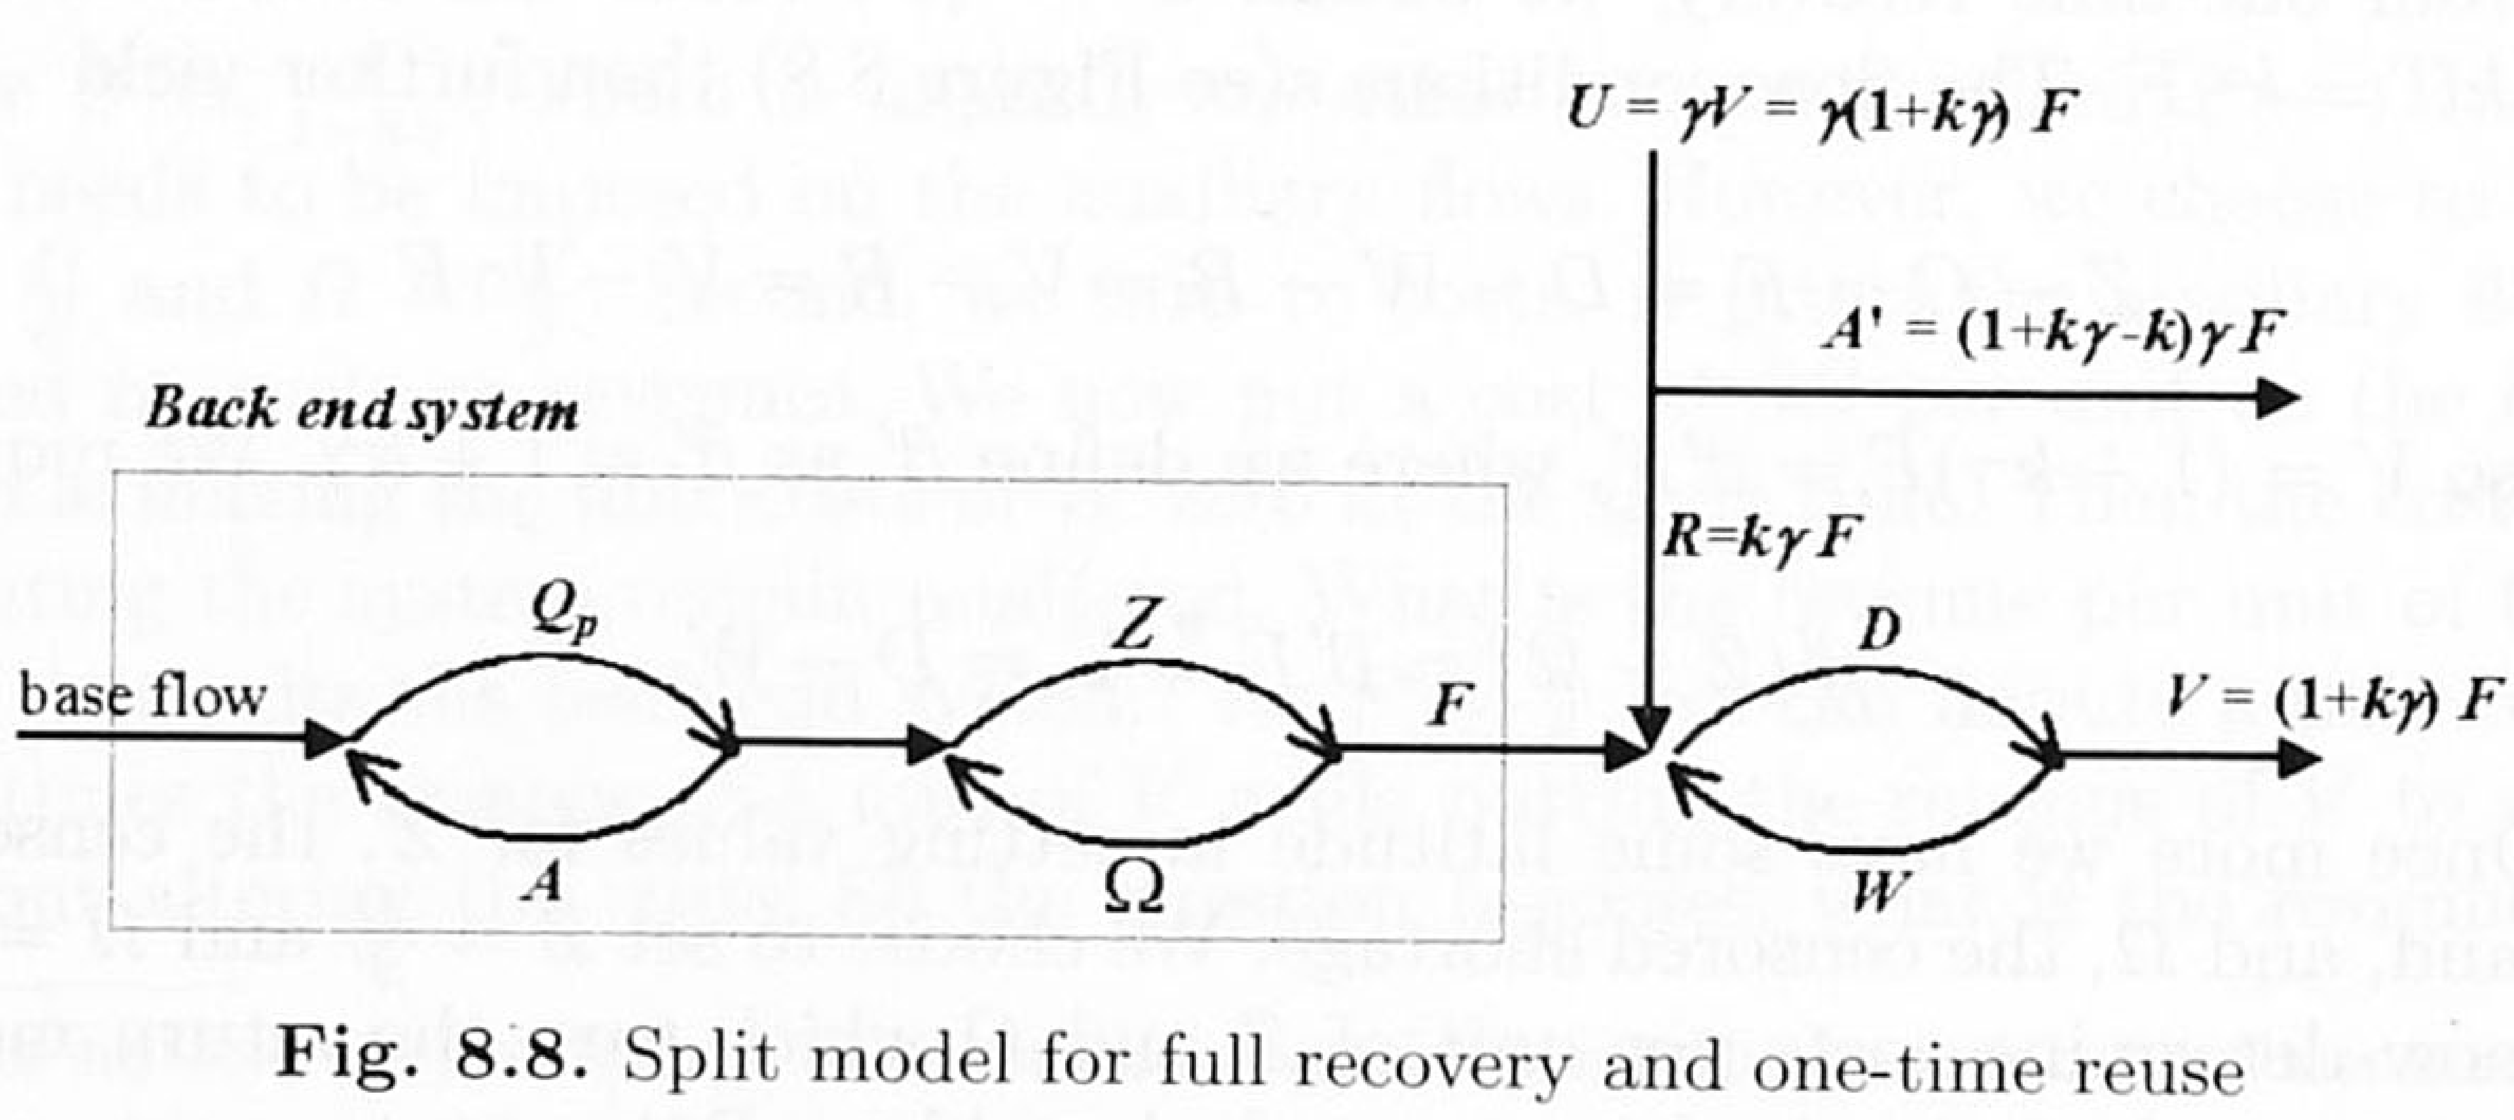

Fig. 8.8. Split model for full recovery and one-time reuse

we reuse them or not. Finally, the recovery of returned products may incur a fixed cost paid per period (e.g., the cost of leasing a packing machine to repack returns). So in partial and full recovery options we have another branch in the tree, which includes (alternatives B2.2.2 and B2.3.2) a fixed recovery cost or not (alternatives B2.2.1 and B2.3.1). The expected profits and the optimality conditions for the other alternative models are presented in Vlachos and Dekker (2000).

A Numerical Example

Figure 8.10 depicts the effect of the return fraction γ and the recovery fraction k to the optimal order quantity and the expected profit for a specific numerical example. The collection and recovery costs are assumed negligible. So, the model allowing re-use only once is used. The other cost parameters are $v = 15$, $c_d = 2$, $c_p = 7$, and $\omega = 5$.

The lines for $\gamma = 0.0$ represent the optimal order quantity and the expected profit of the classical 'news-vendor' problem (without returns). We observe that as the return fraction (γ) or the recovery fraction (k) increase, the effect on Q_p^* and its corresponding expected profit is almost linear. These dependencies also prove the statement that the classical 'news-vendor' optimal quantity is far from optimal when the return rates are high.

Future research directions on single period inventory models include dynamic estimation of expected demand and serviceable returns (quick response) using data from the beginning of the period. This research can be combined with an improved inventory control system for single-period products with returns that includes a second order during the selling period.

In the foregoing models, the ratio between sales (V) and net supply (F) is given by a number β: $V = \beta F$. For the model with unlimited reuse, Mostard and Teunter (2002) go beyond this situation by assuming that each time an

- B1: Sell in the secondary market
- B2: Reuse
  - B2.1: No recovery
  - B2.2: Partial recovery
    - B2.2.1: No fixed cost
    - B2.2.2: Fixed cost
  - B2.3: Full recovery
    - B2.3.1: No fixed cost
    - B2.3.2: Fixed cost

Fig. 8.9. Alternative models for single decision cases with returns

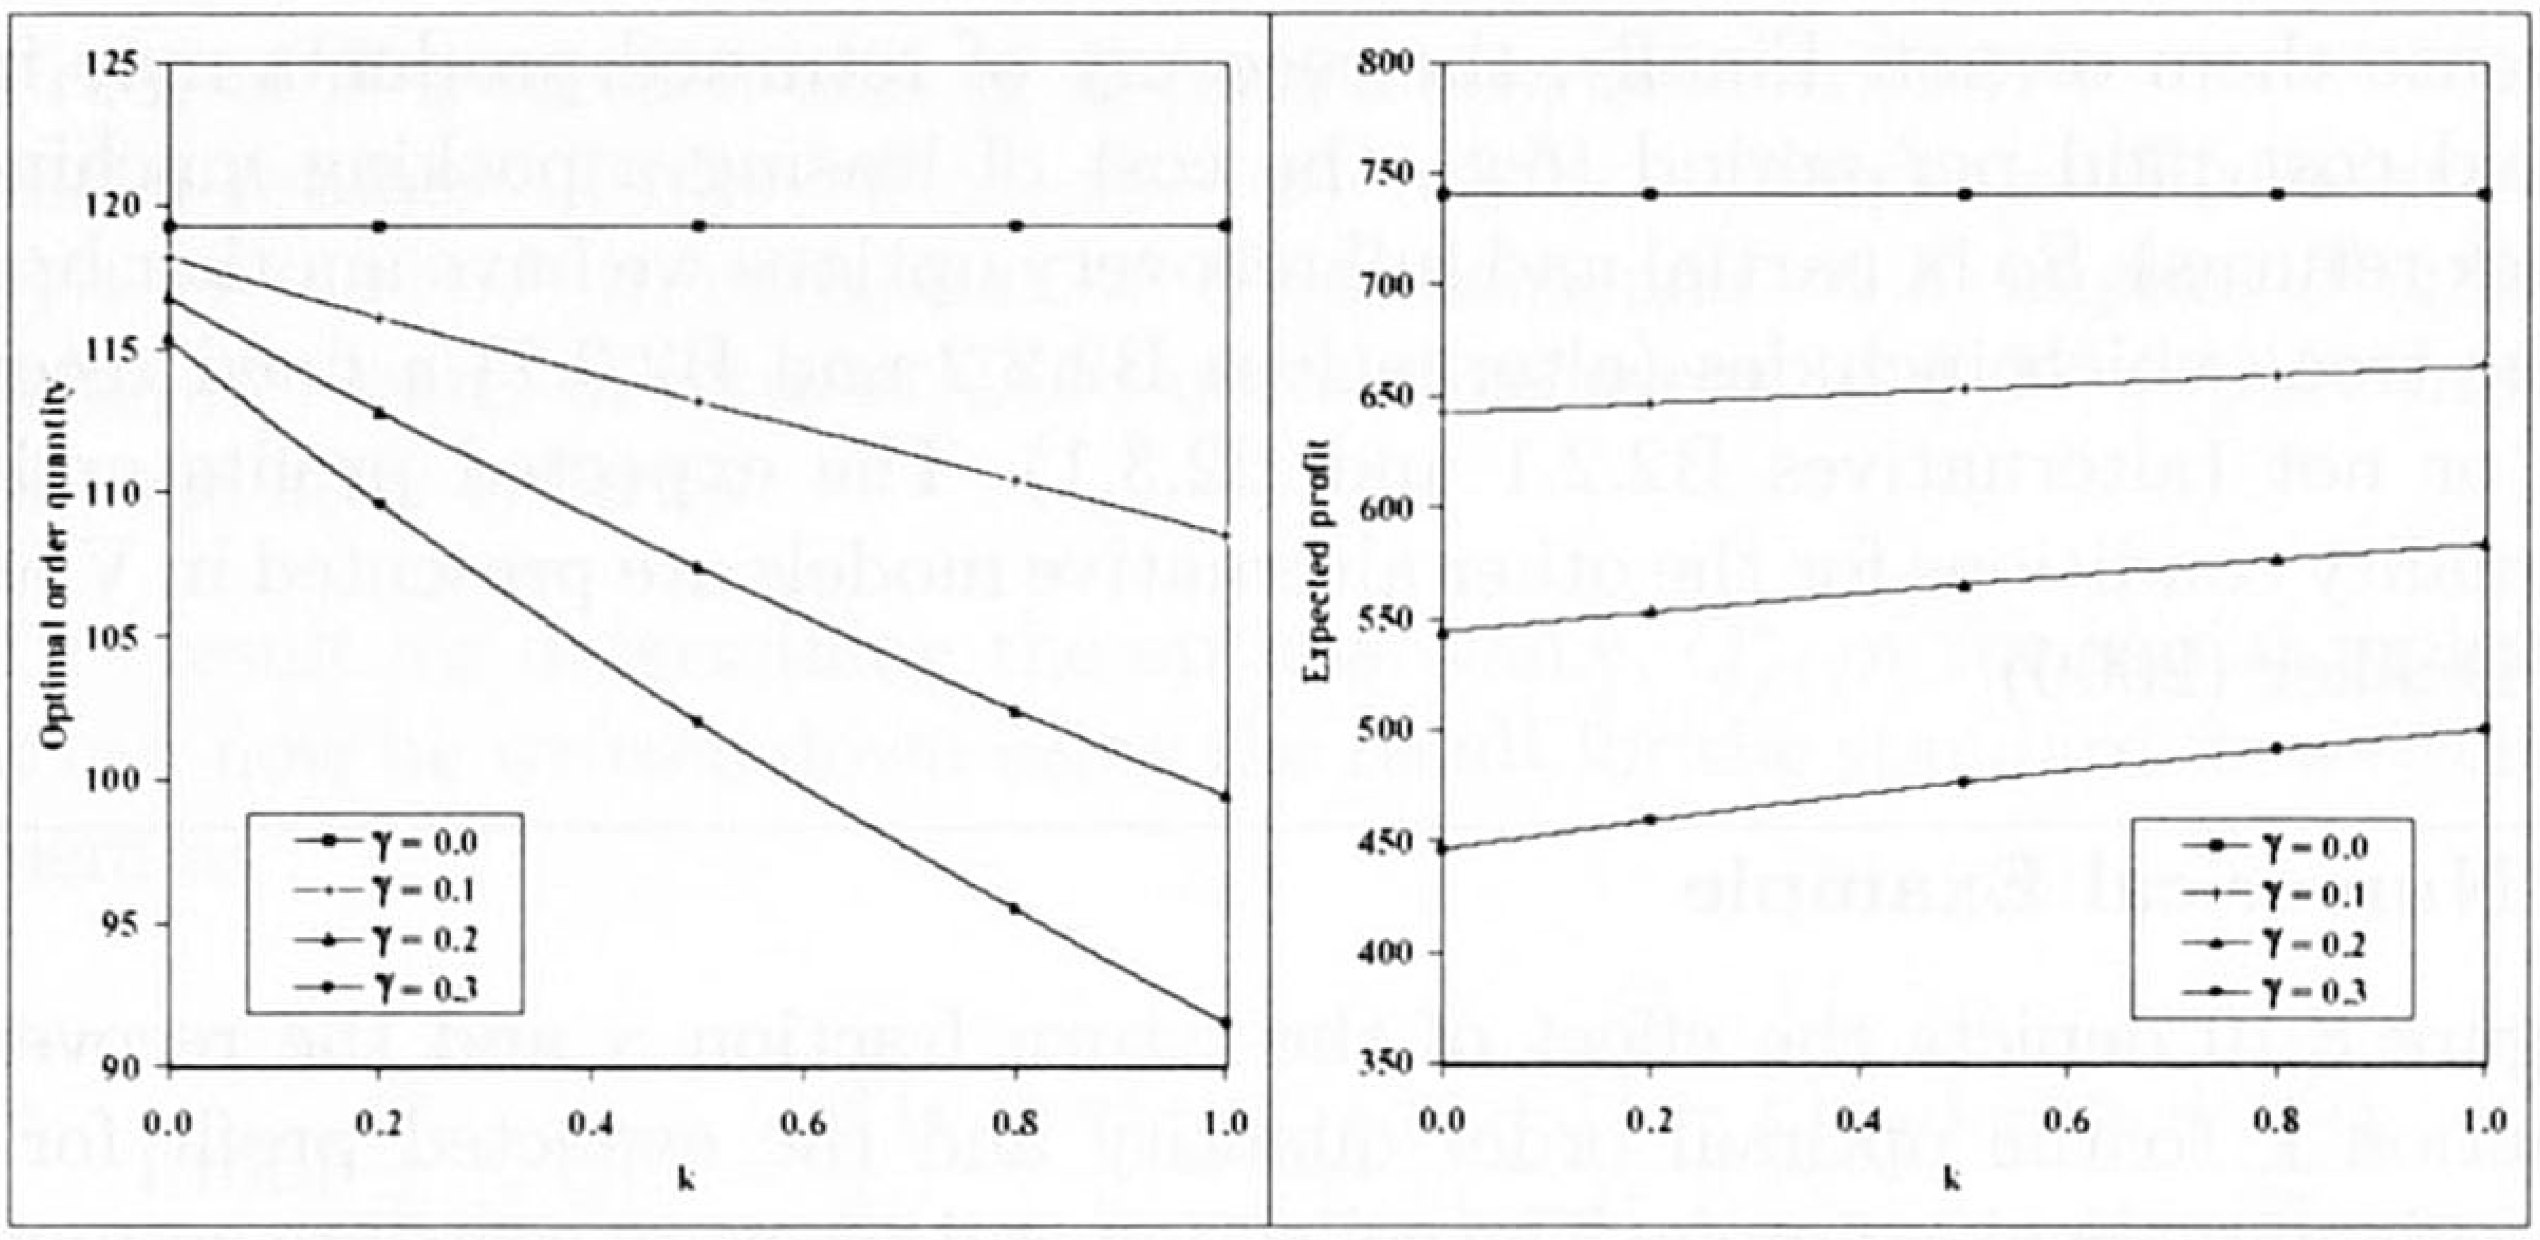

Fig. 8.10. Effect of return fraction γ and recovery fraction k on optimal order quantity Q_p^* and corresponding expected profit

item is sold it has probability of γ of being returned, and that each item returned has a probability of k of being reused. The number of times an item is returned is then geometrically distributed with parameter $1 - k\gamma$: $\text{Prob}(\#\text{ return times} = x) = (1 - k\gamma)(k\gamma)^x$. So the return flow ($R$) can be modeled as $R \cong \sum_{i=1}^{F} \Gamma_i$ where the integer variables $\{\Gamma_i\}$ are independent and each is geometrically distributed with parameter $1 - k\gamma$. Evidently, for any i, $\mathrm{E}(R) = \mathrm{E}(F)\,\mathrm{E}(\Gamma_i) = \mathrm{E}(F)\frac{k\gamma}{1-k\gamma}$ and so from $V = F + R$ one obtains $\mathrm{E}(V) = \frac{\mathrm{E}(F)}{1-k\gamma} = \beta\,\mathrm{E}(F)$ with $\beta \equiv \frac{1}{1-k\gamma}$ as before.

However, knowledge of the value taken by F no longer suffices to determine the values taken by the number of returns (U), the number salvaged after return (A'), the number reused (R), or the number of sales (V). Nonetheless, one readily verifies that each unit of F creates an *expected* revenue of $\beta(v + \gamma(-c_r - v) + (1-k)\gamma\omega)$.

Now define[7] the censored demand (Z) and the censored shortage (Ω) through the equations

$$Z + \sum_{i=1}^{Z} \Gamma_i = D \quad \text{and} \quad \Omega + \sum_{i=1}^{\Omega} \Gamma_i = W.$$

Note that $Z \geq \Omega$ (as $D \geq W$). For model consistency, one requires $Z - \Omega = F = D - W - R$, i.e., $Z - \Omega = D - W - R$. However,

$$Z - \Omega + R + W - D = D - \sum_{i=1}^{Z} \Gamma_i - W + \sum_{i=1}^{\Omega} \Gamma_i + R + W - D$$

[7] The equations may not result in integer values for Z and Ω. In that case, some rounding is necessary. F.e. $Z \equiv \min\{z | z + \sum_{i=1}^{z} \Gamma_i \geq D\}$ which has the advantage that Z is a stopping time for (Γ_i).

$$= R - \sum_{i=\Omega+1}^{Z} \Gamma_i = 0.$$

So Z and Ω are defined consistently. For any i, $\mathrm{E}(\Omega) + \mathrm{E}(\Omega)\mathrm{E}(\Gamma_i) = \mathrm{E}(W)$ and from this it follows $\mathrm{E}(\Omega) = (1-k\gamma)\mathrm{E}(W)$ as $\mathrm{E}(\Gamma_i) = \frac{k\gamma}{1-k\gamma}$. Now attribute to each unit of flow of Ω a cost of $\frac{c_b}{1-k\gamma}$ (and put the unit costs of W to zero). Then, as costs were linear in W, the average costs analysis remains the same.

Under all the above unit costs alterations, for convenience summarized in the following list,

- $v \leftarrow \beta(v + \gamma(-c_r - v) + (1-k)\gamma\omega)$
- $c_b \leftarrow \frac{c_b}{1-k\gamma} = \beta c_b$,

the back-end system now is a model driven by the censored demand Z. Temporarily discounting for the dependence of Z on Q_p, this leads to the following condition that determines the optimal initial replenishment, Q_p^*,

$$\mathrm{Prob}(Z \leq Q_p^*) = \frac{\beta(v + \gamma(-c_r - v) + (1-k)\gamma\omega) - c_p + \beta c_b}{\beta(v + \gamma(-c_r - v) + (1-k)\gamma\omega) - \omega + \beta c_b} = \Theta(\beta),$$

which was previously obtained by Mostard and Teunter (2002).

8.2.4 Multi-period Model

Consider a single inventory facility that carries a single product. Time is segmented into periods. Demand for the product is random with demand in different periods identically and independently distributed (i.i.d.). The facility receives product returns, as well. These product returns in different periods are i.i.d. and independent of demand also. Demand that cannot be met is backlogged. Upon return, products are available for servicing demand after L_r periods. The average of returns per period is smaller than the average demand per period. The difference is furnished by a supplier external to the system considered. To this end, the inventory facility places replenishment orders at the supplier. These orders are delivered with a lead time of L_p periods. Replenishment orders can be placed at the beginning of each period. The decision maker has to determine whether and how much is reordered. Disposal is disallowed. Per unit of product per unit of time, the cost for backlogging demand is c_b and per unit of product per unit of time the inventory carrying cost is h_s. The cost per replenishment order is K_p. As the reader will note, apart from the returns, the model described is the classical inventory model under backlogging.

The analysis of this model lends itself almost immediately as an application of the split inventory model. The one thing that needs some careful consideration is the issue of lead times.

The State Space

In classical stochastic inventory theory, where unmet demand is backlogged, the model just described is best studied on the state space defined by the inventory position. Here, the use of the inventory position avoids the need for dealing explicitly with lead time (lead time is absorbed in the cost function). However, the presence of returns slightly complicates matters, as we now have to deal with two (possibly different) lead times. The state of the system summarizes all the information relevant to the decision at hand. The sequence of events in a period is: 1) arrival of products (returns + ordered) in the serviceable inventory, 2) ordering, and 3) meeting demand and returns. Demand that cannot be met is backlogged. Disposal is disallowed. The state of the system at time period τ is given by a triple $(I_n(\tau), Y(\tau), Z(\tau))$ where

- $I_n(\tau)$ is a scalar representing the net stock at the end of Period τ,
- $Y(\tau)$ is an L_p-tuple $(Y_{L_p-1}(\tau), Y_{L_p-2}(\tau), \cdots, Y_0(\tau))$ where $Y_j(\tau)$ is the amount reordered, in period $\tau + j - L_p$, due to arrive in period $\tau + j$, and
- $Z(\tau)$ is an L_r-tuple $(Z_{L_r-1}(\tau), Z_{L_r-2}(\tau), \cdots, Z_0(\tau))$ where $Z_j(\tau)$ is the amount of returned items becoming available for serving demand in period $\tau + j$.

Introduce the inventory position $I_-(\tau)$ at the beginning of period τ before any events as

$$I_-(\tau) = I_n(\tau - 1) + \sum_{j=0}^{L_p-1} Y_j(\tau) + \sum_{j=0}^{L_r-1} Z_j(\tau) .$$

After product arrivals and reordering, the inventory position is $I_+(\tau) = I_-(\tau) + Y_{L_p}(\tau)$ where $Y_{L_p}(\tau)$ is the replenishment order placed in period τ. The dynamics are given by the equation

$$I_-(\tau + 1) = I_+(\tau) - D_{-1}(\tau + 1) + Z_{L_r}(\tau) ,$$

where $D_j(\tau)$ is the demand during Period $\tau + j$. The net stock I_n can be computed from the inventory position as

$$I_n(\tau + L_p - 1) = I_-(\tau) - \sum_{j=0}^{L_p-1} D_0(\tau + j) + U_{L_r,L_p}(\tau) \qquad (\triangle)$$

with

$$U_{L_r,L_p}(\tau) = \sum_{j=0}^{L_p-1} Z_{L_r}(\tau + j) - \sum_{j=0}^{L_r-1} Z_{L_r}(\tau + L_p - j - 1).$$

Now assume $L_r \le L_p$. Then, clearly, $U_{L_r,L_p}(\tau) = \sum_{j=0}^{L_p-L_r-1} Z_{L_r}(\tau+j)$. Equally important, in terms of distributions, we have for the dynamics

$$I_-(\tau+1) = I_-(\tau) + Y_{L_p}(\tau) - \tilde{D},$$

where $Y_{L_p}(\tau)$ is the amount ordered, and $\tilde{D}$ is distributed as $D - Z$ where demands are distributed as D and returns as Z. That is, I is driven, apart from the quantity reordered, by a stream of independent stochastic variables with distribution $\tilde{D}$, the outcomes of which are unknown at Period τ. Note that $\tilde{D}$ can be interpreted as a signed demand, i.e., as a demand that assumes both positive and negative values. The structure of the dynamics of the inventory position in the basic model, therefore, is the same as that in the standard model except that the demand now is signed. Furthermore, for purpose of performance calculation, we can use ($\triangle$). Indeed, note that the quantity $\sum_{j=0}^{L_p-1} D_0(\tau+j) - U_{L_r,L_p}(\tau)$ is independent of $I_-(\tau)$ (still for $L_r \le L_p$).

The conclusion is that we can consider the system as being a zero-lead-time system subject to signed demand. To this system, we can apply the full strength of the split-inventory technique. When giving explicit results, the following assumes that $L_r = L_p = 0$. However, from the remarks above it will be clear that the analysis can be carried through analogously for any $L_r \le L_p$. Below, results are given for the case of optimizing the expected cost per period.

Analysis and Results

The analysis and results are discussed in detail by Fleischmann and Kuik (2003). One peculiarity is that for the equilibrium process of $(\boldsymbol{V}, \boldsymbol{W})$ the two component processes $\boldsymbol{V}$ and $\boldsymbol{W}$ become independent in the long run: $(W|V = v) = W$ where (V, W) is the equilibrium process. For the stationary case, the cost function for inventory related costs (see Subsection 8.3) takes the simple form

$$H(v) = \mathrm{E}_W(G(v+W))$$

for the inventory in the back-end facility. Somewhat more detailed, the following conclusions hold.

1. An (s, S) inventory policy is optimal for the basic model.
2. The V and W equilibrium processes are independent. The process V can be described through the inventory position process corresponding to an s, S policy for a facility subject to nonnegative demand distributed i.i.d. in periods with distribution

$$(\text{Probability demand} = k) = \begin{cases} \sum_{\ell \ge 0} \pi_\ell \sum_{m \ge 0} \tilde{d}_{\ell-m} & \text{for } k = 0 \\ \sum_{\ell \ge 0} \pi_\ell \tilde{d}_{k+\ell} & \text{for } k \ge 1\,. \end{cases}$$

Here, $\boldsymbol{\pi} = (\pi)_{\ell=0,1,\dots}$ is the invariant distribution for W and $\tilde{d}_i = \text{Prob}(\tilde{D} = i)$. Note that as W is independent of (s, S) so is $\boldsymbol{\pi}$.

3. The costs of the basic model coincide with the costs of the (s, S) process inventory position process described by V with state-dependent period costs given by H. Note that $H(i)$ is the expected cost of a cycle on the half line, incurred by a random walk described by the W-process, starting and ending in the position 0 and with costs in position ℓ given as $G(i+\ell)$.

The three statements together imply that the optimization of the basic model can, at least in principle, be carried out through the processes V and W as an optimization problem for a classic inventory model. This conclusion continues to hold for the case $L_p \geq L_r \geq 0$.

All of the above results immediately carry over to the case of a model in continuous time under continuous review with the obvious modifications (see Fleischmann et al., 2002). The discrete time case is analyzed in more detail in Fleischmann and Kuik (2003).

Remark. The split inventory introduces a front-end facility and a back-end facility. The analysis just reported will carry through in case the back-end system itself is an inventory system consisting of multiple facilities. Thus the analysis, in principle, carries through in case of an assembly system with autonomous returns only at the end-item level (see DeCroix and Zipkin, 2002a).

Heuristics

Several heuristics have been developed for the multi-period model when fixed production costs are disregarded. Simpson (1970) considers a discrete time, periodic review model with a review cycle of m periods. During each period τ there is stochastic demand $D(\tau)$ with mean μ_D and variance σ_D^2 and stochastic recovery output $R(\tau)$ with mean μ_R and variance σ_R^2. Every m periods, a production order is placed with *stochastic* lead time L_p such that the inventory position is raised to S_p. Demand that cannot be fulfilled immediately is backordered. Define the random variable $Y = \sum_{i=\tau+1}^{\tau+m+L_p} D(i) - R(i)$, i.e. Y is the net demand during a review cycle plus a (random) replenishment lead time. The mean and variance of Y are $\mu_Y = (m + \mu_{L_p})(\mu_D - \mu_R)$ and $\sigma_Y^2 = (m+\mu_{L_p})(\sigma_D^2+\sigma_R^2)+(\mu_D-\mu_R)^2\sigma_{L_p}^2$. Under a service objective, i.e. the fraction of demand *not* backordered should be larger than k, it can be shown that the optimal value of S_p is the solution to

$$\int_{S_p}^{0} (y - S_p)g(y)\mathrm{d}(y) = (1 - k)m\mu_D. \tag{8.3}$$

where $g(.)$ is the density of net demand Y. If service is enforced through a penalty cost (c_b per unit backordered) than the optimal S_p is the solution to

$$\int_{S_p}^{\infty} g(y)\mathrm{d}(y) = \frac{mh_s}{c_b}$$

with h_s the serviceable holding cost per unit on stock per period.

A continuous time variant of Simpson's model was recently put forward by Mahadevan et al. (2002). Here, the production lead time, L_p, is a fixed constant and all returned products are initially stocked. As soon as a review epoch occurs, all returns in stock are recovered (fixed lead time L_r) and transferred to serviceable inventory (see Figure 8.11).

One of the heuristics developed for this model approximates the stockout probability $T_p = m + L_p$ time units from the current review epoch, τ *and* the stockout probability at $\tau + T_r$, which is the time at which the last recovery batch (if any) arrives in the interval $[\tau, \tau + T_p)$. The number of recovery batches, N, that arrive during time T_p equals $\lceil L_p/m \rceil$ if recovery batches always arrive *before* manufacturing batches in a review cycle or $\lceil L_p/m \rceil$-1 if recovery batches always arrive *after* manufacturing batches. Hence, $T_r = (N-1)m + L_r$. Assuming that demands and returns follow independent Poisson processes, the net demand during time T_p, Y_p, has mean $\mu_{Y_p} = dT_p - uNm$ and variance $\sigma^2_{Y_p} = dT_p + uNm$. Similarly, the net demand during time T_r, Y_r, has mean $\mu_{Y_r} = dT_r - u(N-1)m$ and variance $\sigma^2_{Y_r} = dT_r + u(N-1)m$. The optimal value of S_p then follows from

$$\int_{S_p}^{\infty} g_p(y) + g_r(y)\mathrm{d}(y) = \frac{mh_s}{c_b}.$$

where $g_p(.)$ and $g_r(.)$ are the densities of Y_p and Y_r respectively. In this particular push policy, the replenishment order is split into a production portion

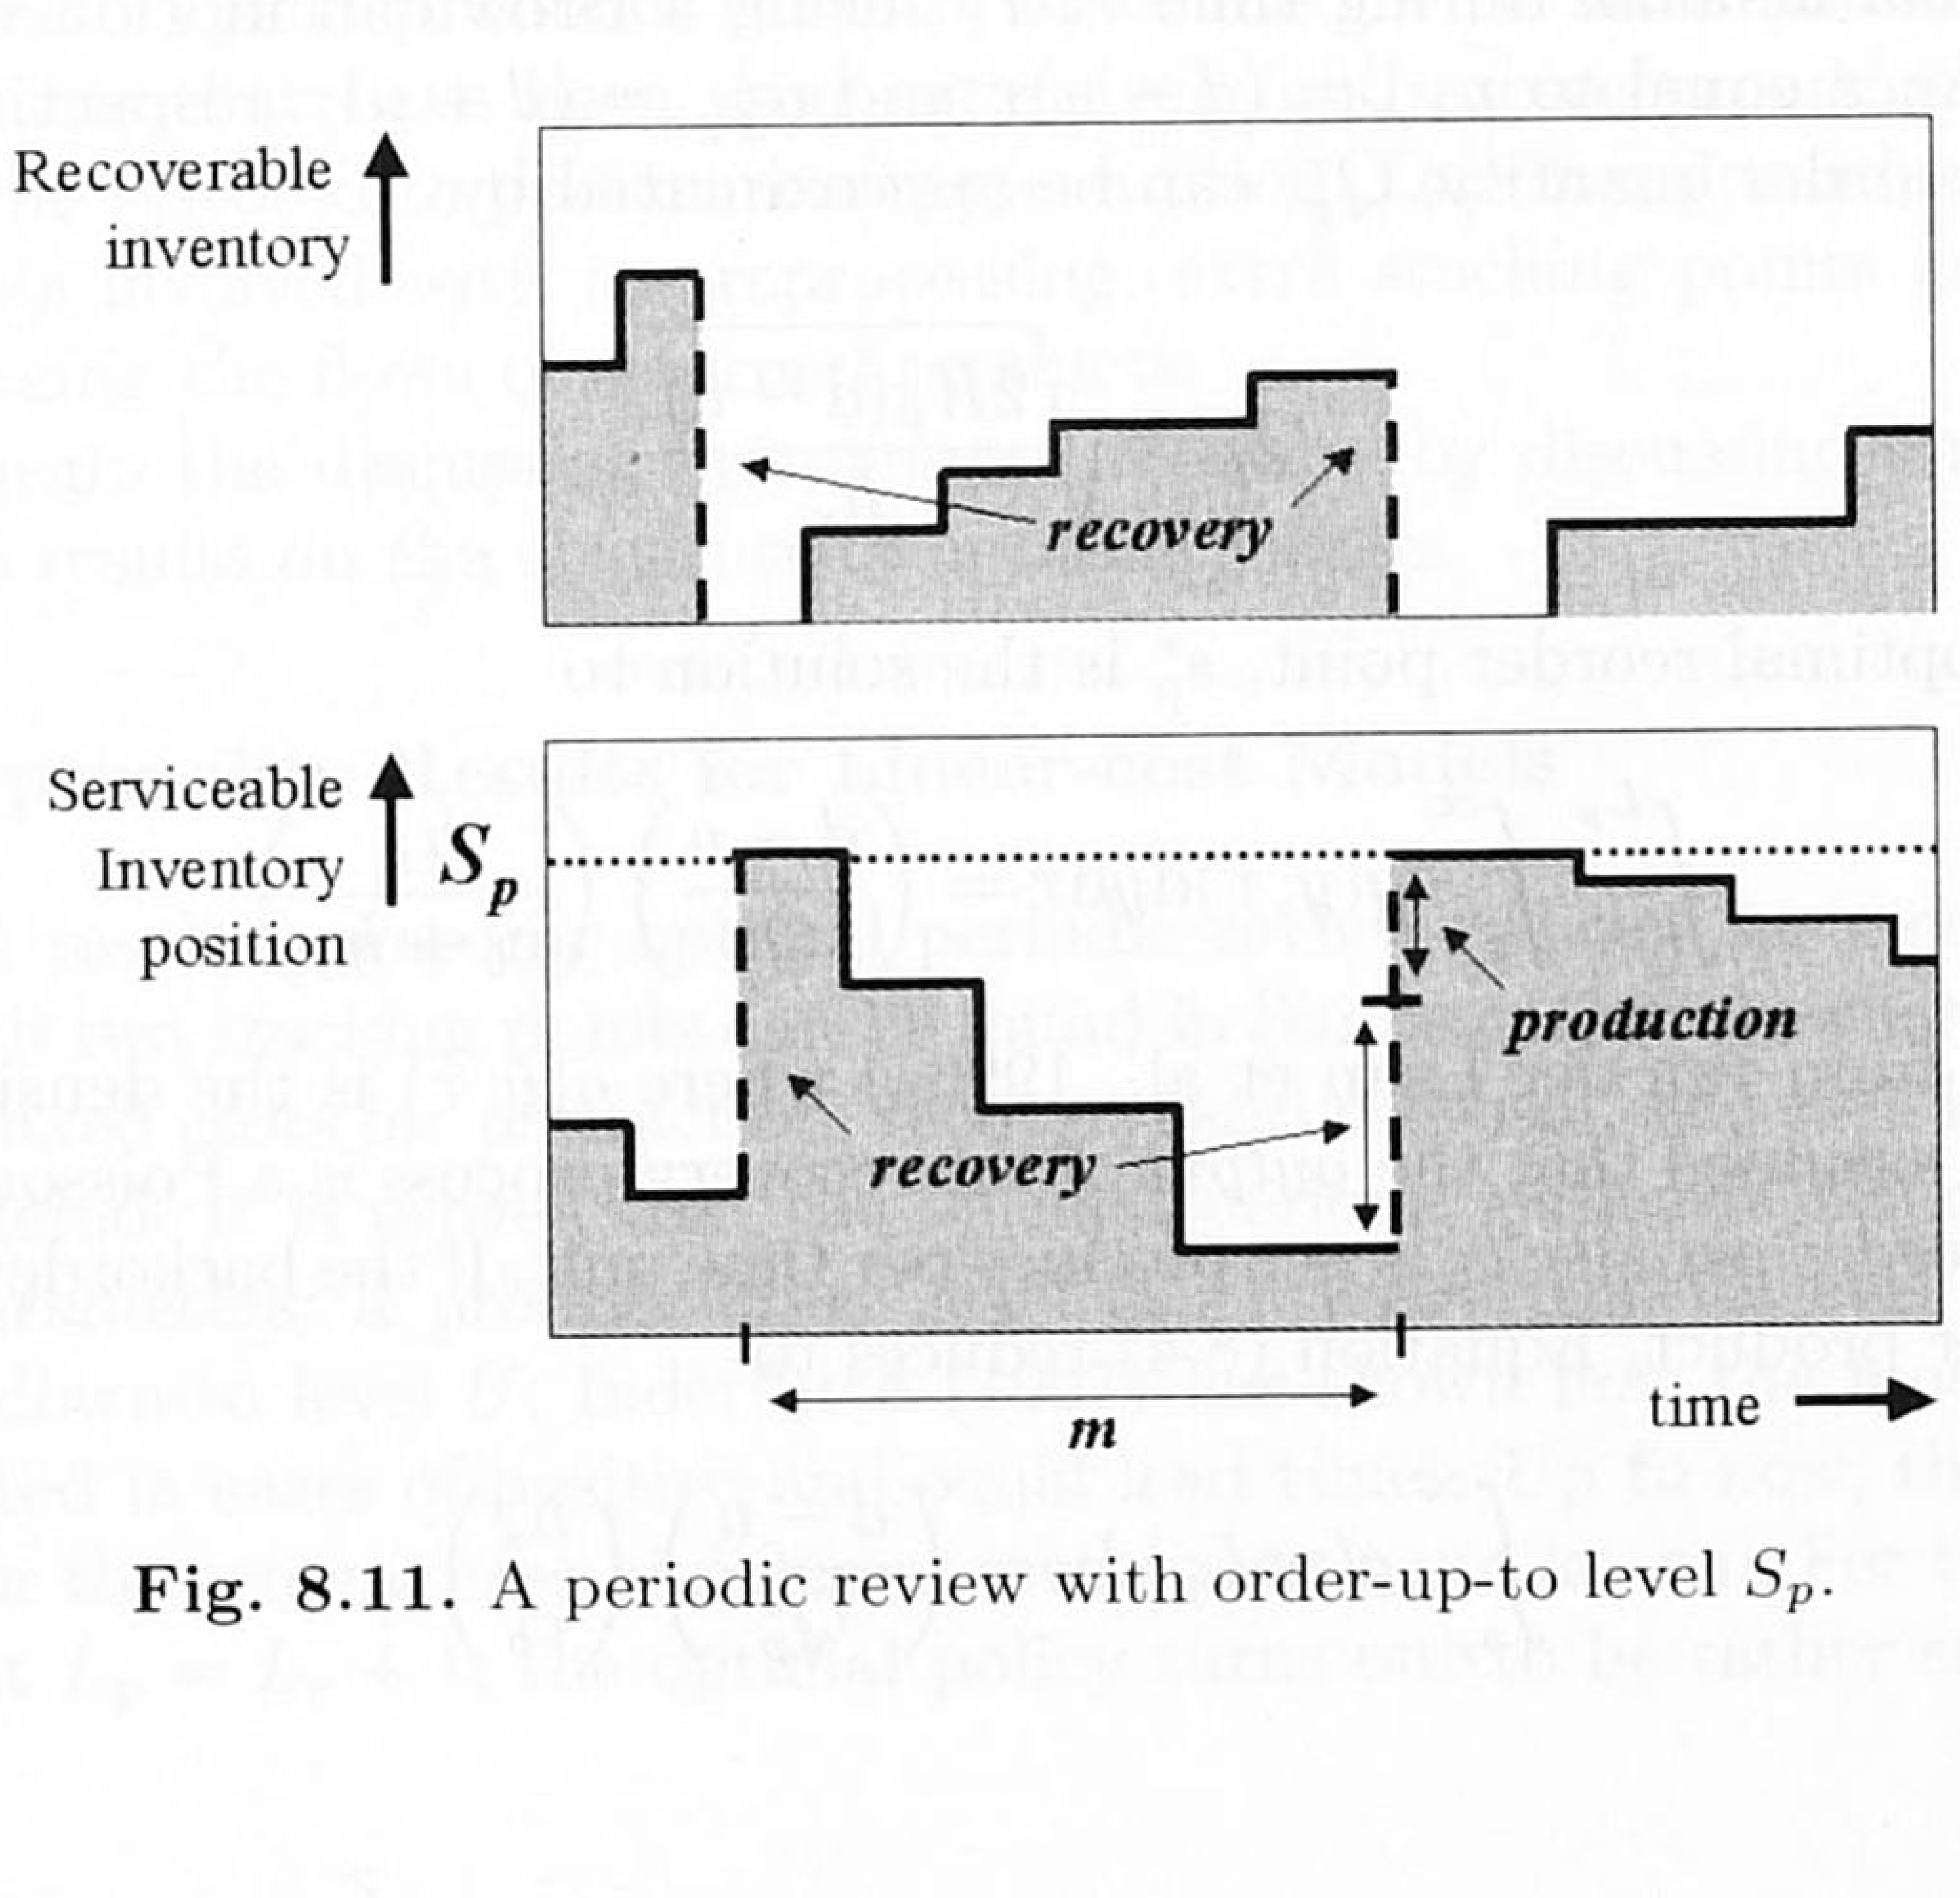

Fig. 8.11. A periodic review with order-up-to level S_p.

and a recovery portion. An interesting side effect is that it pays off to have different lead times for production and recovery so that the incoming serviceable products are more spread out over time. Therefore, increasing one of the lead times may decrease optimal costs.

The only difference between the policies of Simpson ('continuous' recovery) and Mahadevan et al. (periodic recovery) is the way that product returns are handled. Simpson's policy benefits customer service at the price of higher serviceable holding costs, while the policy of Mahadevan et al. delays recovery expenses at the cost of customer service. If recovery is expensive, the latter may be more attractive.

Muckstadt and Isaac (1981) adapted the standard (s, Q) policy to account for a situation with fixed costs for production, K_p. As soon as the inventory position of serviceable items (net serviceable inventory plus everything on order plus all product returns in the recovery facility) drops to s, a production batch of size Q is ordered and will be delivered after a fixed lead time L_p. The recovery facility operates under a one-for-one push policy, and recovery lead times are stochastic. Since the demand and return processes are Poisson streams, the inventory position can be modelled as a Markov chain. The solution procedure is based on the approximation that the net inventory follows a normal distribution and on the assumption that the output of the recovery shop is a Poisson process. The advantage of this procedure is that it results in simple expressions from which one can deduct the optimal values of s and Q very easily. Its disadvantage is that it may not be a very accurate procedure, especially for high return rates (see Van der Laan et al. 1996a). A more accurate procedure was developed by van der Laan et al. (1996a). Instead of approximating the distribution of *net inventory*, the authors approximate the *net demand demand* during time τ, Y_τ, using a Brownian motion with drift and variance equal to $\mu_{Y_\tau} = (d-u)\tau$ and $\sigma^2_{Y_\tau} = (d+u)\tau$ respectively. The optimal reorder quantity, Q_p^*, can be approximated by

$$Q_p^* = \sqrt{\frac{2K_p(d-u)}{h_s}}$$

and the optimal reorder point, s_p^* is the solution to

$$\int_0^{L_p} \int_{s_p}^{\infty} g(y;\tau)\mathrm{d}y\mathrm{d}\tau = \left(\frac{d-u}{Q_p^*}\right)\left(\frac{h_s}{c_b+h_s}\right) \tag{8.4}$$

(adapted from van der Laan et al., 1996a) where $g(y;\tau)$ is the density of Y_τ. Here it is assumed that the *output* of the recovery process is a Poisson stream. The backorder penalty c_b is per product per time unit. If the backorder penalty is just per product, Equation (8.4) reduces to

$$\int_{s_p}^{\infty} g(y;L_p)\mathrm{d}y = \left(\frac{d-u}{Q_p^*}\right)\left(\frac{h_s}{c_b}\right)$$

Under a service objective (at least a fraction k of total demand should be fulfilled immediately), we deduct a similar expression as Equation (8.3) so that s_p^* approximately is the solution to

$$\int_{s_p}^{\infty} (y - s_p) g(y; L_p) \mathrm{d}y = (1 - k) Q_p^* .$$

One of the results of this section is that the standard (s, Q) policy is optimal as long as the recovery process is autonomous and its lead time is shorter than the production lead time. In the original formulation by Muckstadt and Isaac, the recovery facility is modelled as a queuing system with (in)finite capacity and stochastic recovery times. The total lead time, i.e. waiting plus processing time, can be well above L_p especially if the system load is high, implying that the (s, Q) policy combined with autonomous recovery is not necessarily optimal. However, a main drawback of the above heuristic approaches is that they are hard to adapt to managed recovery. The issue of managed recovery is studied in the next section.

8.3 Managed Recovery (or Pull Strategies)

In contrast with situations of direct reuse as discussed in the previous section, Section 8.2, value-added recovery involves more elaborate (re)processing. In this case, the throughput times of returned items can be substantial and varying. Moreover, variation and uncertainty of the quality of returns makes it more difficult to predict reprocessing needs and add to the variability of lead times. Typical products that are associated with value-added recovery are products that have been used extensively and are returned/collected in order to be restored to perfect working condition. Due to value-added and/or fixed costs involved with the reprocessing, extra stocking points are needed for managing the flows of returned products.

We begin the discussion on managed recovery by discussing some of the (limited) results on the structure of optimal policies.

8.3.1 Optimality Results for Linear-cost Models

The first results regarding optimal periodic review policies for recovery systems with two stocking points can be found in Simpson (1978), where a model with no fixed costs for production, recovery, and disposal, and zero lead times is considered. It is proven that the optimal periodic policy is determined by three parameters: a produce-up-to level S_p, a recover-up-to level S_r and a dispose-down-to level U. Inderfurth (1997) has shown how the model has to be adapted in cases of positive and equal lead times. Up to now, the optimal policy for the general case of different lead times is not known. For the special case that $L_p = L_r + 1$, the optimal policy turns out to be rather simple and has been given by Inderfurth (1997). In the following, we will describe the model for the situation $L_r = L_p = L$ in more detail.

It is assumed that the inventories are reviewed periodically and that in each period τ it has to be decided how much to produce ($p(\tau)$), how much to recover ($r(\tau)$) and how much to dispose of ($w(\tau)$). Demands in period τ ($d(\tau)$) as well as returns ($u(\tau)$) are random variables, and the corresponding probability density functions are denoted with $\varphi_{\tau,d}$ for the demands and $\varphi_{\tau,u}$ for the returns.

The recovery system can then be described by two state variables, the physical stock of recoverables I_u and the inventory position of serviceables I_s, which is defined as the stock-on-hand of serviceables plus all outstanding orders minus backorders. The inventory balance equation for the recoverables is given as

$$I_u(\tau+1) = I_u(\tau) - w(\tau) - r(\tau) + u(\tau), \tag{8.5}$$

while the inventory balance equation for the serviceables is given as

$$I_s(\tau+1) = I_s(\tau) + r(\tau) + p(\tau) - d(\tau). \tag{8.6}$$

In Simpson, the optimal policy is determined with respect to the average total relevant cost over a finite planning horizon T. Thereby, production, recovery, and disposal costs are assumed to be proportional to the number of items, and holding and penalty costs are charged to the net inventory at the end of each period. The cost parameters are denoted as follows.

h_s: unit holding cost for serviceable items
h_u: unit holding cost for recoverable items
c_b: unit backorder cost
c_p: unit production cost
c_r: unit recovery cost
c_w: unit disposal cost

Then the average cost for backorders and keeping serviceables in stock in period τ is given as a function of the inventory position Y_s after the reorder decisions ($Y_s(\tau) = I_s(\tau) + p(\tau) + r(\tau)$) as follows:

$$L_\tau(Y_s) := h_s \int_0^{Y_s} (Y_s - z)\varphi_{\tau,L,D}(z)\,dz + c_b \int_{Y_s}^{\infty} (z - Y_s)\varphi_{\tau,L,D}(z)\,dz, \tag{8.7}$$

where $\varphi_{\tau,L,D}$ denotes the density function of the cumulative demands in the periods $\tau - L, \tau - L + 1, \ldots, \tau$. Further, we have to include the average cost for stock keeping of recoverable items in period τ which can be computed as a function of the inventory position Y_u of the recoverable inventory after the reorder decisions ($Y_u(\tau) := I_u(\tau) - w(\tau) - r(\tau)$):

$$C_\tau(Y_u) = h_u \int_{-\infty}^{\infty} (Y_u + z)\varphi_{\tau,u}(z)dz = h_u Y_u + h_u \mathrm{E}[u(\tau)]. \tag{8.8}$$

Since the term $h_u \mathrm{E}[u(\tau)]$ in (8.8) has no influence on the optimization, we can neglect it in the sequel. In order to be able to formulate the problem as a stochastic dynamic programming problem, we introduce the average relevant cost f_n for n remaining periods until the end of the planning horizon. This function depends on two variables: the inventory position of the serviceable inventory I_s and the recoverable inventory I_u.

The functional equation of dynamic programming is obtained as follows

$$f_0(I_s, I_u) \equiv 0 \tag{8.9}$$

and for $n \geq 1$ as

$$\begin{aligned} f_n(I_s, I_u) = \min_{p,r,w \geq 0} \Big\{ & c_p p + c_w w + c_r r + L_{T-n}(I_s + p + r) \\ & + h_u(I_u - w - r) + H_n(I_s + p + r, I_u - w - r) \Big\} \end{aligned} \tag{8.10}$$

where

$$H_n(a, b) := \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f_{n-1}(a - z, b + y)\varphi_{n,d}(z)\, \varphi_{n,u}(y)\, dz\, dy \tag{8.11}$$

with $\varphi_{n,d}(z)$ denoting the density of $d(T-n)$ and $\varphi_{n,u}(y)$ denoting the density of $u(T-n)$. Note that (8.9) and (8.10) also hold for non-identically distributed correlated demands and returns.

An analysis of these equations leads to the structure of the optimal policy (see for details Simpson, 1978). This so-called (S_p, S_r, U) policy (with $S_p \leq S_r \leq U$) is a straightforward extension of the simple (S_p, U) policy which is optimal if stocking of returns is disallowed. The optimal decisions in period τ in the case of two stocking points are determined by the three time-dependent parameters $S_p(\tau)$, $S_r(\tau)$, and $U(\tau)$ as follows.

Items are only disposed of if there are too many in the system, but you can never dispose of items already produced or remanufactured. This leads to:

$$w^*(\tau) = \Big(\min\{I_s(\tau) + I_u(\tau) - U(\tau), I_u(\tau)\} \Big)^+, \tag{8.12}$$

where $(x)^+$ denotes the $\max\{0, x\}$. Production is used if the total number of items in the system is less than $S_p(\tau)$.

$$p^*(\tau) = \Big(S_p(\tau) - (I_s(\tau) + I_u(\tau))\Big)^+ \qquad (8.13)$$

Since there is only a limited number of recoverables available, the serviceable inventory position cannot always be increased to the recover-up-to level $S_r(\tau)$ using recovery, leading to the following recovery decisions:

$$r^*(\tau) = \min\Big\{I_u(\tau), (S_r(\tau) - I_s(\tau))^+\Big\}. \qquad (8.14)$$

For the application of such a control policy in practice, the policy parameters have to be computed.

Omitting the time dependence from the notation, the policy parameters S_p, S_r, and $U \geq 0$ can be determined by solving the following equations.

$$L'(S_p) + \partial_1 H(S_p, 0) = -c_p \qquad (8.15)$$

$$L'(S_r) + \partial_1 H(S_r, U - S_r) = c_w - c_r \qquad (8.16)$$

$$\partial_2 H(S_r, U - S_r) = c_w - h_u \qquad (8.17)$$

where $\partial_1 H$ and $\partial_2 H$ denote the partial derivative of H with respect to the first and second argument respectively. If there exists no solution of (8.17), (8.16) with $U - S_r \geq 0$, one needs to solve (8.16) for S_r with $U = S_r$.

8.3.2 An Exact Modelling Approach

Using the modelling approach of the previous section, it is hard to find, for each period, the optimal policy parameters in the presence of fixed production and recovery costs. This makes it difficult to generalize those results to situations in which batching is necessary. A continuous time, continuous review setting enables one to formulate various inventory control strategies that extend those considered in Section 8.3.1 and that can be optimized and analyzed making use of the theory of Markov Chains. Additionally, this setting facilitates the modelling of stochastic lead times.

Van der Laan et al. (1999b) developed an *exact* procedure that enables one to study a variety of push and pull policies under fairly general conditions, such as stochastic lead times and Markovian return and demand flows. The modelling framework for a system with two stocking points, one for recoverable inventory and one for serviceable inventory, is characterized as follows.

- The *demand and return processes* are stochastic and may be modelled by any Markovian arrival process. The two may even be dependent, but this requires extra state variables to model the number of products in the market (see, for example, Bayindir et al., 2003; Nakashima et al., 2002; Yuan and Cheung, 1998) and/or the time that they have spent there. Although it is common practice to assume simple (compound) Poisson arrivals, an alternative could be to use Coxian-2 arrival processes. These enable one to do a three-moment fit of an arbitrary arrival process, so that a better description of reality can be achieved.

- The *production process* has unlimited capacity. The production costs consist of a variable cost per item and a fixed cost per order. The production lead time L_p is a *discrete* random variable, bounded by some $L_p^{max} < \infty$.
- In principle, the *recovery process* can be any queuing system with Markovian transitions. Alternatively, the recovery process can be modelled as a 'black box' with lead time L_r, a *discrete* random variable bounded by some $L_r^{max} < \infty$. The recovery costs consist of a variable cost per item and a fixed cost per batch.
- If disposals are allowed, the *disposal process* depends on the control policy employed. Next to a variable component, disposal costs may also include a fixed cost per batch.
- The *inventory position* may be defined in various ways (see the discussion in section 8.3). The only restriction is that its transitions are Markovian. For instance, the inventory position may be the net inventory plus all outstanding production orders plus some subset of recoverable products that are currently in the system. As an example, product returns may enter inventory position upon arrival (the arrival process is Markovian) or as soon as some control policy triggers them to be released to the recovery facility. Necessarily, such policies only work on Markovian processes such as the inventory position itself or the stock of recoverables. As we have seen, the optimal policy structure is only known for some very special cases, so in general we have to rely on heuristic policies.
- Although the framework does not pose any restriction on the holding cost parameters, it is reasonable to assume that the recoverable holding cost h_u is smaller than the serviceable holding cost h_s. Moreover, to come to a meaningful performance measure, its numeric values should have a direct relation with the variable costs of production, recovery, and disposal. There is quite some controversy, though, with respect to the correct valuation of holding cost parameters in a reverse logistics setting, but this discussion is left to Chapter 11.
- *Customer service* is modelled in terms of backorder costs, either per product or per product per time unit.
- All system parameters are *stationary*, i.e. do not change over time.

The calculation of the average on-hand serviceable inventory and the average backorder position is difficult since the transitions of the net inventory, $I_n(\tau)$, are usually not Markovian. However, depending on the assumptions with respect to the recovery lead time, we can deduct a handy relation between the net inventory and inventory position $I_s(\tau)$ from which we can calculate the long-run distribution of $I_n(\tau)$. Define

- $W(\tau)$ as the number of recoverables that are included in the inventory position at time τ, but that have not yet been recovered,
- $O(\tau_1 - \tau_2)$ as the output of the recovery process in the interval $(\tau_1, \tau_2]$,
- $D(\tau_1 - \tau_2)$ as the demand in the interval $(\tau_1, \tau_2]$,

- $R(\tau_1, \tau_2)$ as the number of recoverable products that enter the inventory position in the interval $(\tau_1, \tau_2]$ and subsequently enter serviceable inventory at or before time τ_2, and
- $P(\tau_1, \tau_2)$ as the number of products that are ordered at the production facility in the interval $(\tau_1, \tau_2]$ and subsequently enter serviceable inventory at or before time τ_2.

Then we have the following cases.

Case 1: Markovian recovery lead times

$$I_n(\tau) = I_s(\tau - L_p^{max}) - W(\tau) + O(\tau - L_p^{max}, \tau) - D(\tau - L_p^{max}, \tau) \quad (8.18)$$

Case 2a: Discrete recovery lead times bounded by $L_r^{max} \leq L_p^{max}$.

$$I_n(\tau) = I_s(\tau - L_p^{max}) + R(\tau - L_p^{max}, \tau) - D(\tau - L_p^{max}, \tau) \quad (8.19)$$

Case 2b: Discrete recovery lead times bounded by $L_r^{max} \geq L_p^{max}$.

$$I_n(\tau) = I_s(\tau - L_r^{max}) + P(\tau - L_r^{max}, \tau) - D(\tau - L_p^{max}, \tau) \quad (8.20)$$

Relation (8.18) is a generalization of the relation given in Muckstadt and Isaac (1981), whereas Relations (8.19) and (8.20) are taken from van der Laan et al. (1999b). The long-run distribution of $I_n(\tau)$ can be (numerically) found by analyzing long-run and transient behavior of an appropriate Markov chain. For details, we refer to van der Laan (1997). The long-run distribution of net inventory suffices to calculate the long-run expectation of on-hand inventory and the backorder position. All other relevant entities can be obtained from the Markov Chain analysis.

Although the above framework is very general in theory, it suffers from the curse of dimensionality. The state space grows exponentially with the production and recovery lead time and the capacity of the stocking points. Therefore, optimization may be very time consuming and running time grows exponentially in the number of decision variables. The reader should keep in mind, however, that this framework is meant to assess the performance of a wide variety of recovery policies in an *exact* way rather than using approximations. It is not meant as a fast optimization algorithm nor as an efficient numerical recipe. In the case of very large state spaces, simulation, although less accurate, may be an alternative.

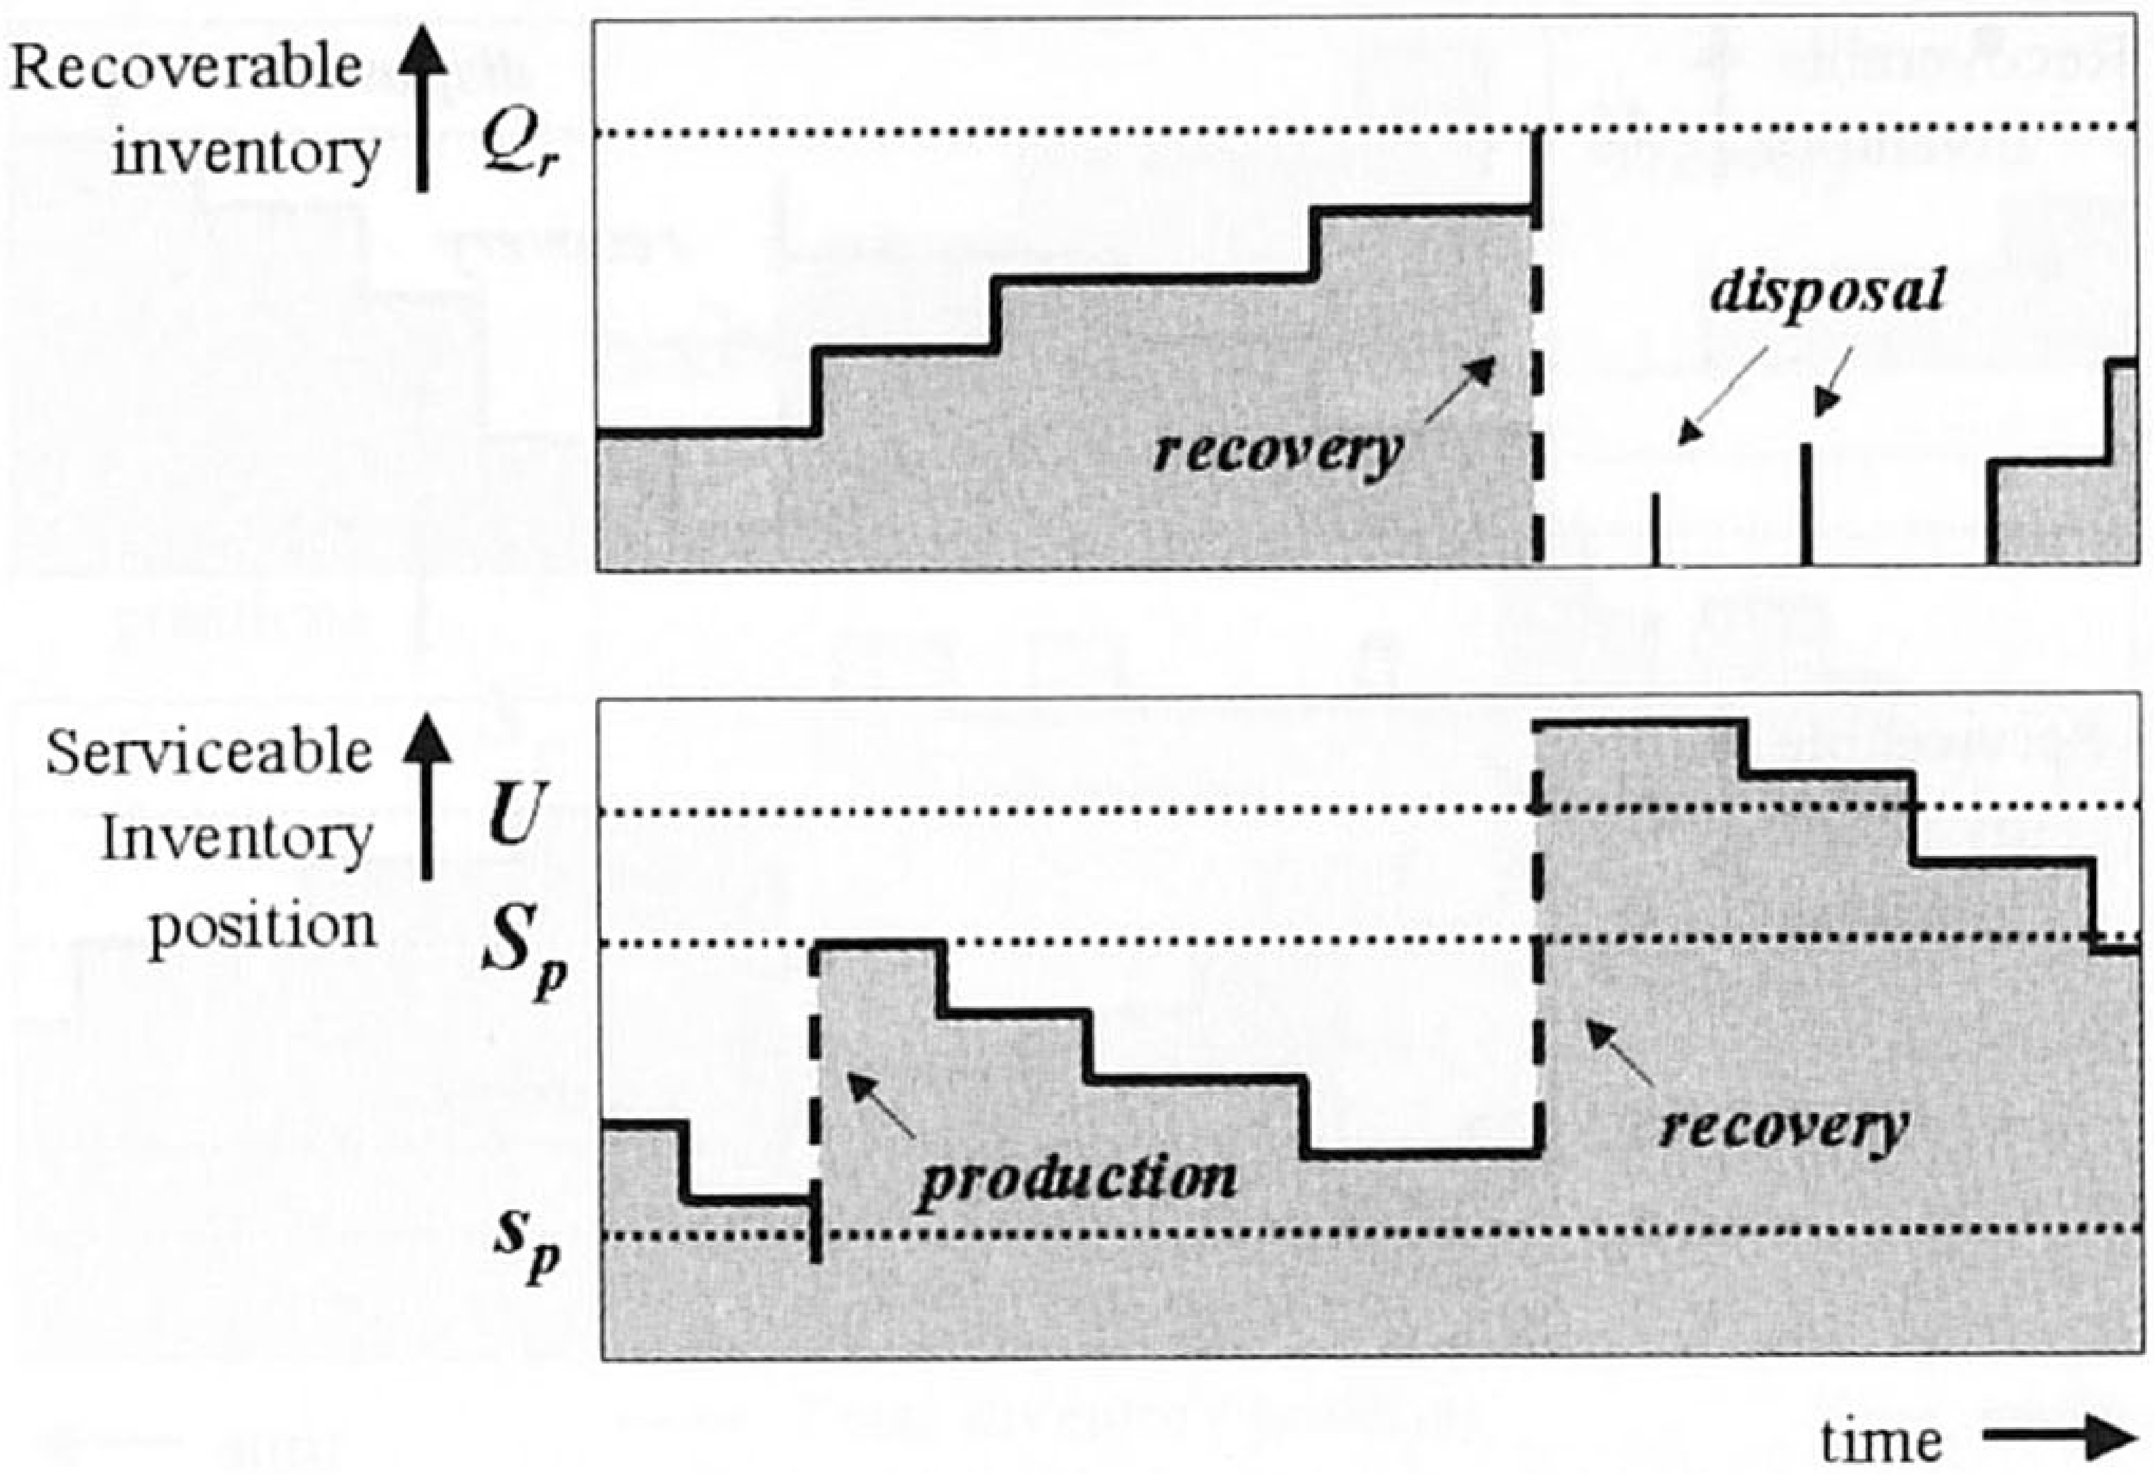

Fig. 8.12. The (s_p, S_p, Q_r, U) PUSH disposal policy.

Push and Pull Policies

The heuristic control policies that have appeared in the literature are all natural extensions of the classical (s, S) policies and may be classified as either push or pull: while production orders are controlled by an (s_p, S_p) policy, recovery batches are either *pushed* through the recovery facility or *pulled* only when they are really needed.

Figures 8.12 and 8.13 give a graphic representation of some of the variants (for details see van der Laan and Salomon, 1997) that can be analyzed with the framework outlined above. With the (s_p, S_p, Q_r, U) PUSH disposal policy, remanufacturing starts whenever Q_r recoverable products are in stock. A production order is placed to increase the serviceable inventory position to S_p as soon as the serviceable inventory position drops to or below the level s_p. Products are disposed of upon arrival as soon as the inventory position exceeds the level U. With the (s_p, S_p, s_r, S_r, U) PULL disposal policy, recovery starts as soon as the serviceable inventory position is at or below s_r, and sufficient recoverable inventory exists to increase the serviceable inventory position to S_r. A production order is placed to increase the serviceable inventory position to S_p as soon as the serviceable inventory position drops to or below the level s_p. Products are disposed of upon arrival as soon as the inventory position exceeds the level U. Note that s_p should never exceed s_r, since otherwise the recovery option would be redundant.

In the above examples, the inventory position is defined as the net serviceable inventory plus all outstanding recovery and production orders. For that case, van der laan et al. (1999a) show that an increase in the recovery lead time or an increase in production lead time variability may lead to lower optimal

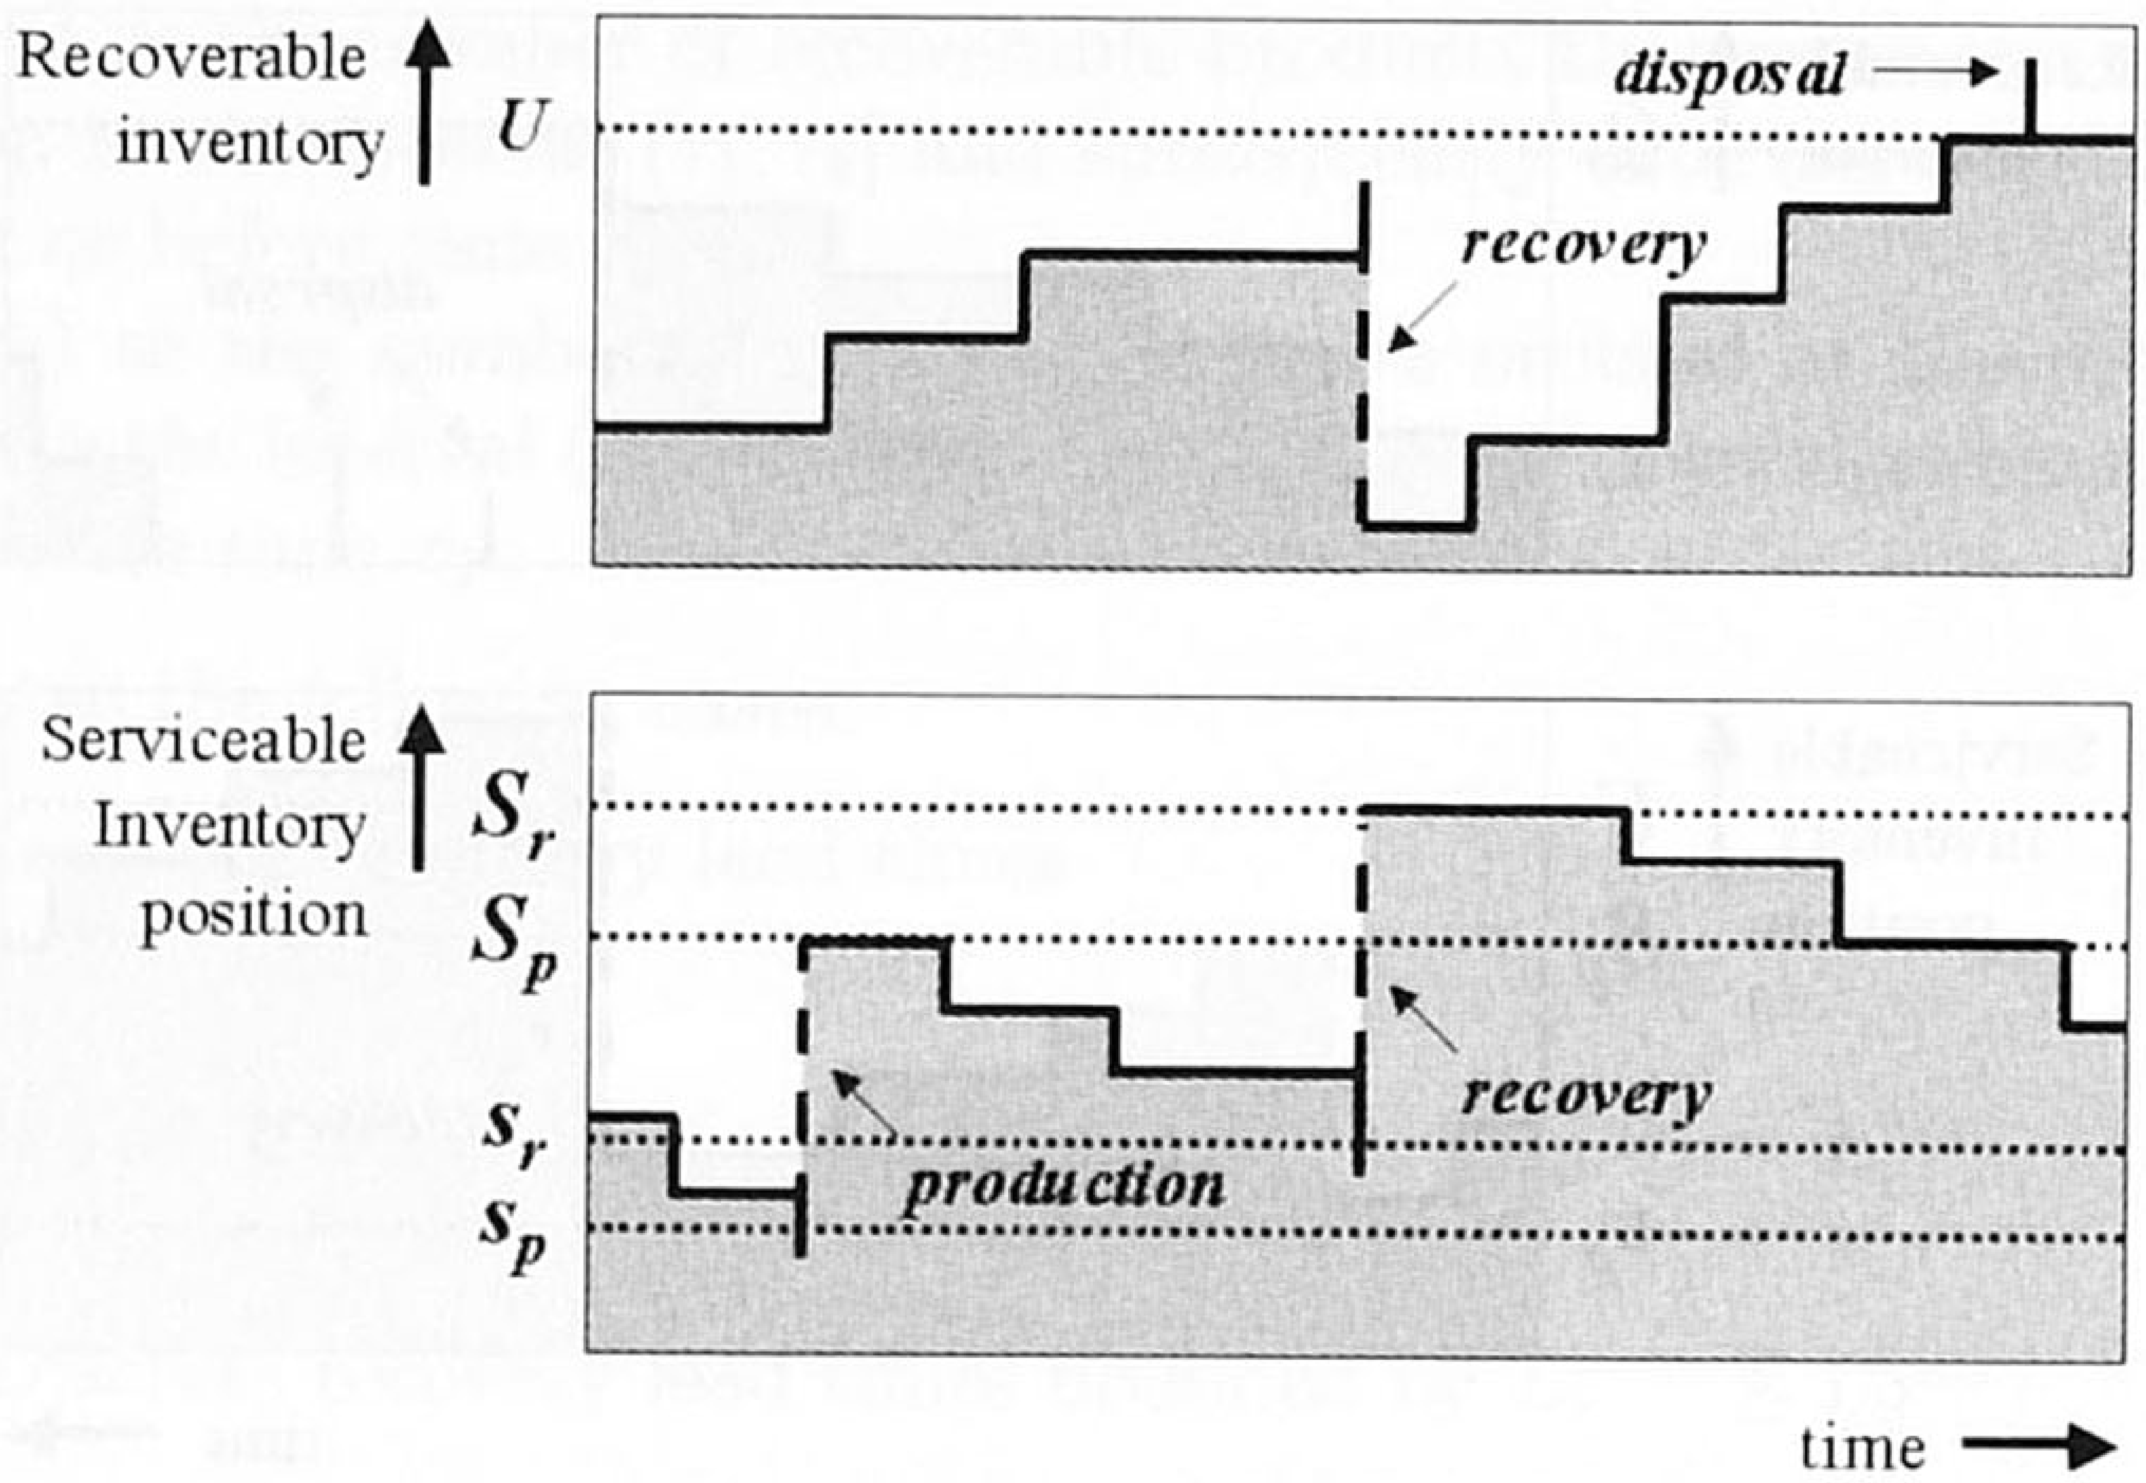

Fig. 8.13. The (s_p, S_p, s_r, S_r, U) PULL disposal policy

costs. It appears that this is due to the sub-optimality of the control policies. Inderfurth and Van der Laan (2001) show that the PUSH-disposal policy is easily improved upon by adjusting the recovery lead time in an appropriate way or by using different information sets for the production, recovery, and disposal decisions. The sub-optimality of the PULL-disposal policy is mainly due to the restriction on the recovery order level ($s_r \geq s_p$). If the recovery lead time is much smaller than the production lead time, s_r will overprotect for recovery lead times. Increasing the recovery lead time then leads to better policy performance. The above effects are further illustrated in Teunter et al. (2002) for push and pull policies without the disposal option.

If the recovery lead time is *smaller* than the production lead time, one could base the production decision on the sum of the serviceable inventory position and the recoverable stock, while the recovery decision could be based on the serviceable inventory position only. If the recovery lead time is longer, one should base both the production and recovery decisions on the serviceable inventory position only. Such policies are easily modelled within the above framework and therefore can be analyzed analytically.

Actually, only orders with a certain remaining service time should be included in the serviceable inventory position (see also the discussion in Section 8.3.3). Such an inventory position, however, is not a Markov process, so our framework cannot be used. In the case that $L_r < L_p$, Teunter et al. (2002) investigate a policy in which the production decision is based on all the inventory in the system, while the recovery decision is based on the serviceable inventory position (serviceables on hand, plus outstanding orders), which includes only those production orders with remaining lead time smaller than L_r (see Figure 8.14). It is shown by simulation that this policy outperforms simple push and pull policies that are based on just one inventory position.

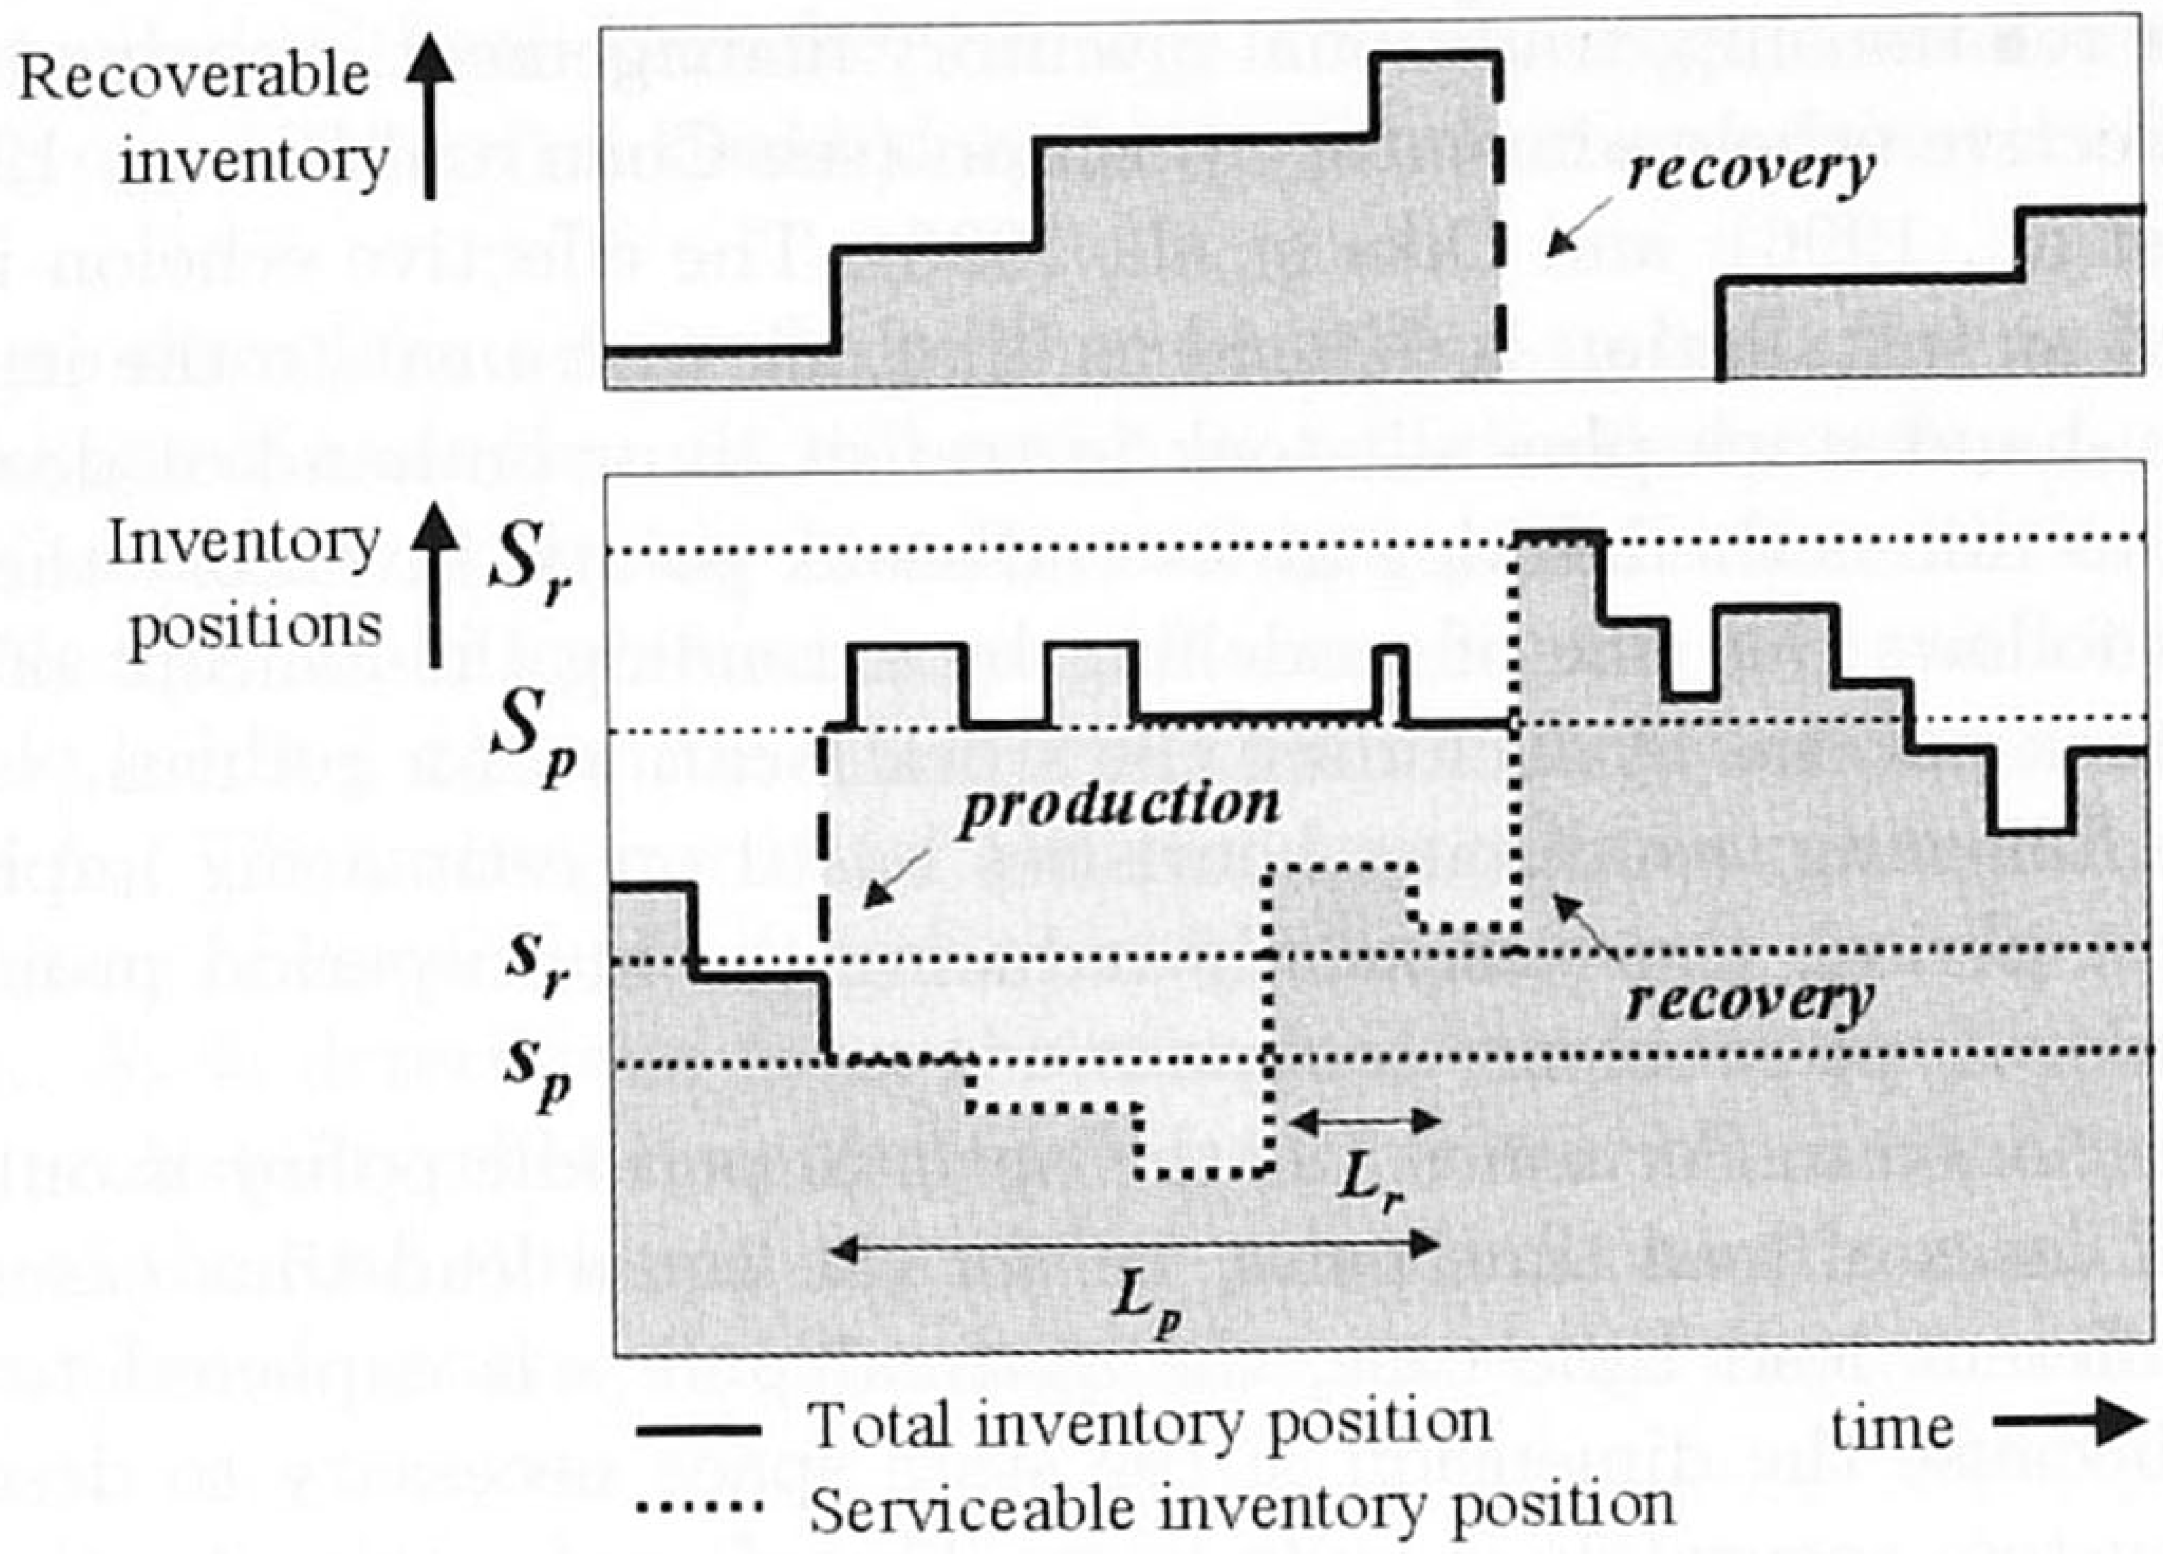

Fig. 8.14. A schematic representation of the double-PULL policy

The drawback of this policy is that it cannot be modelled as a Markov Chain and is therefore difficult to analyze exactly.

8.3.3 Heuristics

In the case of models with managed recovery, the difference with traditional inventory theory goes beyond signed demand. Inventory control now explicitly needs to deal with recovery flows. The source of returned items is unreliable and limited. This makes its control particularly difficult. In traditional inventory theory, one also has to deal with unreliability since supplying inventory facilities might be out of stock. Also in that case, unless special structures in costs and the inventory network exist, one has to resort to rules of thumb or heuristics to find reasonably good solutions. Frequently one finds heuristics and approximations based on the notion of effective echelon inventory position.

Let us first introduce the relevant definitions and concepts. We will do this with an assembly network structure in mind (or special cases thereof, such as series systems or a single location). The installation net stock of a facility is defined as its on-hand minus its backlog. The installation inventory position of a facility is its net stock plus the number of products on order (with its supplier). The nominal echelon stock of a facility is its installation echelon stock plus the installation echelon stock position of all its downstream facilities. In multi-echelon systems, the nominal echelon stock position has no simple relation to the net stock (shifted by a lead time and up to demand during a period of lead-time length), yet in other models this is one of the primary reasons for introducing the concept of inventory position. To restore the simple relationship, traditional inventory management introduces the concept of effective echelon inventory position (see Chen and Zheng, 1994); Van Houtum et al., 1996); and Diks et al., 1996). The effective echelon inventory position of an installation is defined as all stock *in transit* to the installation plus its on-hand stock plus all stock in transit to or on-hand at downstream installations minus the backlog at its end-stock points. Inventory theory with recoveries follows this line of modelling by extending the concept of effective echelon stock policies by including the stock locations for returns.

In the following, we discuss heuristics based on estimating implied costs after a first period, thus essentially reducing the multi-period problem to a 'news-vendor' type problem.

As mentioned in Section 8.3.1, the optimal periodic policy is only known for special cases of lead time pairs, i.e. for the equal lead time case. For the general different lead time case, the optimal policy is expected to be very complex because the dimension of the state space necessary to describe the recovery system accurately is quite large. Therefore, for this situation several heuristics are developed. But also for the equal lead time case it is reasonable to use heuristics since the exact computation of the optimal policy parameters using Equations (8.15), (8.16), and (8.17) can be very time consuming.

Equal Lead Times

Kiesmüller and Scherer (2002) provide two computational approximation schemes to determine nearly optimal policy parameters S_p, S_r, and U. One is based on an approximation of the value-function in the dynamic programming problem (8.10) and leads to excellent results. They obtain in their numerical study an average relative deviation of the average cost of 0.1% and a maximum relative deviation of less than 2%. The other approximation uses a common decomposition technique which is based on a deterministic model and safety stocks. This approach leads to the shortest computation time but also to less accurate policy parameters. The same numerical examples lead to an average relative deviation of the average cost of about 3% and a maximum relative deviation of 10%. Especially in the case of large return rates and large standard deviations for the demands and returns, the first approximation outperforms the latter substantially.

Kiesmüller and Scherer also illustrate that the computation of the policy parameter U can be numerically ill-conditioned when U has hardly any influence on the costs. This is always the case when disposal of items does not play a significant role in the optimal policy. A detailed simulation study by Teunter and Vlachos (2002) shows that in stationary models this is almost always the case. Therefore, an (S_p, S_r) policy, where S_p denotes the produce-up-to-level and S_r the recover-up-to-level is a near-optimal policy for a stationary situation.

For such an (S_p, S_r) policy, simple 'news-vendor' type formulas for the computation of the parameters can be found in Kiesmüller and Minner (2002).

In order to derive these formulas, overage and underage costs are estimated depending on whether it is decided upon the production or the recovery quantity.

First we show how a formula for the recover-up-to-level S_r can be derived. Any underage of a single unit will result in a backorder penalty c_b, assuming that the backorder will only last for a single period. In addition, another recoverable unit could have been remanufactured in the past without affecting serviceable holding costs. Therefore, the cost for each unit short is given by $c_s = c_b + h_u$. The overage cost for the recovery of one unit too many equals the difference of serviceable and recoverable holding cost $c_o = h_s - h_u$. Using c_s and c_o, S_r is determined from the marginal 'news-vendor' approach such that the probability that cumulative demand within the next $L+1$ periods (the sum of the lead time and the review period) does not exceed S_r is equal to the fraction of underage and overage plus underage cost per unit and unit of time

$$F_{L+1}(S_r) = \frac{c_b + h_u}{c_b + h_s} \qquad (8.21)$$

with F_{L+1} denoting the cumulative distribution function of the demands in $L+1$ periods.

For the produce-up-to-level, the underage cost simply equals the unit backorder cost per unit of time ($c_s = c_b$). Producing one unit too much leads to serviceable holding costs h_s and might also influence future recovery decisions. Then, the recovery of a unit has to be postponed for some time. This time i can be approximated with a random variable which is distributed according to a geometric distribution with probability $F_{d-u}(0)$. Thereby, $F_{d-u}(0)$ denotes the probability that the net demands $d-u$ are smaller than zero. Therefore, the overage costs can be estimated as follows:

$$c_o = h_s + h_u \sum_{i=1}^{\infty} i \cdot F_{d-u}(0)^i \cdot (1 - F_{d-u}(0)) = h_s + h_u \frac{F_{d-u}(0)}{1 - F_{d-u}(0)}. \qquad (8.22)$$

This leads to the following equation for the determination of the produce-up-to-level:

$$F_{L+1}(S_p) = \frac{c_b}{c_b + h_s + h_u \frac{F_{d-u}(0)}{1 - F_{d-u}(0)}}. \qquad (8.23)$$

The accuracy of these formulas has been tested in a detailed simulation study (see Kiesmüller and Minner, 2002) and reveals excellent results.

Different Lead Times

Now we drop the assumption of equal lead times. Although many authors do not distinguish between the situation with recovery lead time being longer than the manufacturing lead time and the reverse situation, we believe this to be necessary because the control problems are different. In the following, we will provide an approach for the control of a hybrid stochastic production/recovery system with different lead times which can also be found in Kiesmüller (2002). Since we assume a linear cost model and stationary demands and returns, it seems to be reasonable to extend the (S_p, S_r) policy mentioned above to the situation with different lead times. In the following, we assume that the processes for demands and returns are identically distributed in each period.

Long Recovery Lead Time

In case of a long recovery lead time, the system is quite easy to control. For moderate return rates, it is reasonable to push all returns into the recovery process while the faster production supply mode takes care of the items that remain to be produced. Only in the case of large return rates and large lead time differences may some problems occur, if there are low demands in subsequent periods. Then it may happen that more items than required are in the recovery process and there is no more chance for adaptation using the faster supply mode.

In order to use an (S_p, S_r) policy to control the system, we have to know which information should be used for the decisions. Some authors (for example, Gotzel and Inderfurth, 2002, or Inderfurth and van der Laan, 2001) suggest using one inventory position for both the recovery and the production decision. In the context of dual supplier models, this approach would be called a single index policy. Another possibility is to aggregate information in two different variables, which we also call inventory positions (dual index policies). In Kiesmüller (2002), it is shown that the dual index policies outperform the single index policies for recovery systems, especially for large lead time differences. The approach is described in the following.

For the recovery decision, the inventory position $I_m(\tau)$ in period τ includes the current serviceable net-stock plus all outstanding production and recovery orders, including the production order placed at time τ

$$I_m(\tau) := I_n(\tau) + \sum_{i=0}^{L_p} p(\tau - i) + \sum_{i=1}^{L_r} r(\tau - i) . \tag{8.24}$$

The information which is used for the production decision is aggregated in the second inventory position I_s. Since the production decision in period τ influences the stock-on-hand in period $\tau + L_p$, we assume that it is only necessary to consider the outstanding production and recovery orders which arrive in the periods $\tau, \tau + 1, \ldots, \tau + L_p$. This leads to the following definition of I_s:

$$I_s(\tau) := I_n(\tau) + \sum_{i=1}^{L_p} p(\tau - i) + \sum_{i=0}^{L_p} r(\tau - (L_r - L_p + i)) . \quad (8.25)$$

Based on these inventory positions the following decisions rules are obtained:

$$p(\tau) = (S_p - I_s(\tau))^+ \quad (8.26)$$

and

$$r(\tau) = \min\{I_u(\tau), (S_r - I_m(\tau))^+\}. \quad (8.27)$$

In Kiesmüller (2002) it is shown by simulation that such an (S_p, S_r) policy outperforms a similar policy where the decisions are only based on one inventory position because less safety stock is needed. With increasing lead time differences, the cost improvements are increasing and they are much larger than can be obtained with the policy improvement procedure proposed in Inderfurth and van der Laan (2001).

For the policy given by (8.24), (8.25), (8.26), and (8.27), simple formulas exist for the computation of nearly optimal policy parameters (see Kiesmüller and Minner, 2002). Estimating overage and underage costs and using 'newsvendor' type formulas leads to the following two equations for the determination of near optimal policy parameters:

$$F_{L_r+1}(S_r) = \frac{c_b + h_u}{c_b + h_s} \quad (8.28)$$

and

$$F_{L_p+1}(S_p) = \frac{c_b}{c_b + h_s + h_u \frac{F_{d-u}(0)}{1 - F_{d-u}(0)}} . \quad (8.29)$$

Long Production Lead Time

In the case of a long production lead time compared to the recovery lead time, the control situation is much more difficult. Due to the longer production lead time, we have to include information with respect to future incoming returns in the production decision. Thus, in this case, the control problem is more complicated.

As an heuristic, we again suggest using an (S_p, S_r) policy based on two inventory positions. For the definition, we use the same principle as above: for the decision with the longer lead time include all outstanding orders in the inventory position and for the decision with the shorter lead time include only the orders which will arrive until the new released order comes in. Therefore, for the production decision the following inventory position is used:

$$I_s(\tau) := I_n(\tau) + I_u(\tau) + \sum_{i=1}^{L_p} p(\tau - i) + \sum_{i=1}^{L_r} r(\tau - i) . \quad (8.30)$$

For the recovery decision, we use

$$I_m(\tau) := I_n(\tau) + \sum_{i=0}^{L_r} p(\tau - (L_p - L_r + i)) + \sum_{i=1}^{L_r} r(\tau - i) \,. \tag{8.31}$$

Using (8.30) and (8.31) for the decisions (8.26) and (8.27) leads again to a much better cost performance compared to a policy with one inventory position, although both policies lead to nearly the same system-wide stock-on-hand. The reason for the cost reduction is the partition of the system-wide stock-on-hand in the two stocking points. The policy presented above keeps returned items in the recoverable inventory as long as possible while in the other case most of the items are pushed in the serviceable inventory (see for details Kiesmüller, 2002). Further, the lead-time paradox, decreasing average cost with increasing recovery lead time, cannot be observed for the policy defined by (8.26), (8.27), (8.30), and (8.31).

A nearly optimal recover-up-to-level can be obtained by

$$F_{L_r+1}(S_r) = \frac{c_b + h_u}{c_b + h_s} \tag{8.32}$$

and a-produce-up-to-level by

$$F_\Delta(S_p) = \frac{c_b}{c_b + h_u \frac{1}{1-F_{d-u}(0)} + h_s F_{d-u}(0)} \,. \tag{8.33}$$

Thereby, F_Δ denotes the cumulative distribution function of $\sum_{i=0}^{L_p} d(\tau + i) - \sum_{i=0}^{L_p-L_r-1} u(\tau + i)$.

Extensions

As illustrated above, information plays an important role when deciding about production and recovery. Many problems arise when there is not much information available about returns. On the other hand, additional information can be quite valuable. For the situation when product returns are dependent on the demands, which holds, for example, in the case of rented or leased products, a discussion can be found in Kiesmüller and van der Laan (2001). There it is assumed that the number of returns in a period τ depends on the number of demands in a previous period $\tau - \ell$, where ℓ is some fixed number of periods. If the probability that an item can be recovered is assumed to be known, then the number of returns can be estimated using this probability and the information about the known demands. Using this estimation for system control, it is shown that costs can be reduced compared to the situation where the dependency is ignored. Further it is illustrated that the variance of the inventory processes is reduced. Toktay et al. (2000) come to a similar conclusion. They take a queueing model approach and explicitly model the dependence relation between the demands and returns to investigate the value of return information (see Chapter 3 in this book for further details).

In the models discussed up to now, it is assumed that there is only one option to recover the returned products. But in many situations an old product can be reused in different ways, each yielding different costs and profits. A model with multiple recovery options, one disposal option, but no additional production facility is investigated in Inderfurth et al. (2001). Here the problem is to allocate the limited amount of recoverable products to the different recovery options. The structure of the optimal policy is extremely complicated, due to the inherent allocation problem in the case of scarce recoverables. But under a linear allocation rule, a fairly simple near-optimal policy exists which is characterized by a single dispose-down-to-level and a specific recover-up-to-level for each reuse option.

The research presented in this section is dealing with remanufactured products which are assumed to be as good as new. In Inderfurth (2002), a model is introduced which assumes that remanufactured products differ significantly from new ones and that higher-value new products are offered to the customer if there is a stock-out of remanufactured products. For a single-product, single-period problem with stochastic demands and returns and positive lead times, the optimal policy is determined. It is given by produce-up-to and recover-up-to order functions (of the on-hand serviceable inventories) for manufacturing and remanufacturing.

Other Approaches

A special case of an (S_p, S_r) policy is examined in Tagaras and Vlachos (2001) and Vlachos and Tagaras (2001). These publications refer to an inventory system with two replenishment modes (corresponding to production and recovery), one of which may act as an emergency supply channel which can deliver on short notice. An approximate cost model is provided which can be easily optimized with respect to the decision parameters. This model is used as the basis for an heuristic algorithm, which leads to solutions that are very close to the exact optimal solutions determined through simulation.

8.4 Discussion and Outlook

Inventory control for product recovery is very much different from traditional inventory control due to the highly variable and uncertain nature of the extra resource: product returns. The development of specialized inventory models is essential to analyze and understand the complicated dynamics of stochastic inventory control for product recovery. So far this chapter mainly dealt with *modeling* aspects of stochastic inventory control for product recovery management. Below we list some of the managerial implications of the modeling efforts.

Autonomous versus managed recovery

Specialized techniques should simultaneously determine trigger levels and quantities for production, recovery, and disposal operations given the characteristics of its specific environment, such as procurement lead time, the form of the product life cycle, and seasonal influences. In this context, we distinguished between autonomous recovery and managed recovery. Autonomous recovery (or push recovery) mainly relates to those situations in which minor operations against limited costs suffice for successful recovery and reuse. Managed recovery (or pull recovery) is appropriate if recovery operations are more costly and/or setup costs are present so that it is better to stock recoverables until they are really needed.

Split inventory analysis

In the case of autonomous recovery with only one stocking point (for serviceables) present, the inventory process can be cleverly split into two components: the back system and the front system. The front system depends on the demand and the return process, but is independent of the production orders. The front end can be interpreted as a modified demand process that acts on the back system. In this way, the inventory system can be viewed as a standard single-source system with demand process that takes on positive as well as negative values. In Subsections 8.2.3 and 8.2.4, we showed how this modelling framework applies to both single-period and multi-period decision problems. The managerial implication for the mail-order company of Case A is that in principle it could use standard techniques to calculate the initial order sizes. Finding the correct cost parameters, however, requires a thorough understanding of the underlying model, so this still could be quite problematic.

Naive netting versus sophisticated netting

The 'sophisticated netting' of the demand process that is used in the split inventory technique is very much unlike 'naive netting' through which a part of the expected demand rate is cancelled with the expected return rate. The resulting demand process is then input for traditional inventory control models. This is a very simple approach that unfortunately renders a very poor performance, unless return rates are very low. The variance that is introduced by the return process is completely ignored, while the variance of the demand process is moved away from its real value. It is therefore not recommended to use naive netting in practice unless return rates are very small.

Lead-time effects

In the case that managed recovery is more appropriate, things are more complicated, since we have to take into account non-zero recovery lead times and/or fixed setup costs for the recovery process. Only for very specific situations (no setup costs, fixed production lead time equals fixed recovery lead time) do we know the optimal structure of the inventory policy. In other cases, we have to rely on heuristics. Without fixed setup costs and neglecting disposals one can, as an approximation, reduce the multi-period decision problem to a single-period decision problem that is basically the standard 'news-vendor' problem. Crucial for a good policy is that the recovery decision and the production decision be based on different inventory positions that include different information sets.

Inventory systems that do include fixed recovery costs have mainly been modelled in continuous time rather than discrete time. An exact modeling framework, using the theory of Markov chains, for analyzing a wide variety of (heuristic) push and pull inventory policies was presented in Section 8.3.2. It can be shown that the cost performance of pull policies dominate the performance of push policies for most relevant values of the cost parameters. In practice, a push policy could still be preferred though, since it is easier to implement and its performance could be reasonable in the case of long and/or highly variable recovery lead times. Disposal policies are of importance only when return rates are close to or exceeding demand rates or products are very slow-moving.

Value of information

A crucial assumption of most available models to ensure tractability is that product returns are independent of product demands. In reality, though, there is always some dependency relation between product demands and returns. Knowledge about this dependency relation enables more sophisticated forecasts with respect to timing and quantity of product returns. If good forecasts are available these could be incorporated in the inventory policy and well improve system performance. The issue of forecasting and its impact on inventory management is studied in Chapter 3 of this book.

From a modelling perspective, it is sometimes convenient to assume that the time in market is negative exponentially distributed. This way it is sufficient to keep track of the *total* number of products in the market, rather than all individual products. In de Brito and Dekker (2003), it is shown by company data analysis that this assumption is not always according to real behavior. More research is needed to assess the impact of wrongly assuming exponential lags.

From the massive growth of the literature on inventory control for product recovery we may conclude that at least the scientific community believes that this field is worth studying. Due to the complexity of these systems, however, attention has been limited to single-product, single-component models, while a remanufacturing company like Volkswagen (dis)assembles and recovers thousands of components. At the same time, there seems to be a lack of communication between academics and practitioners, considering the very limited amount of empirical studies that have been conducted up to now.

More case studies will undoubtedly help in bridging the gap between theory and practice.