4

Reverse Logistics Network Design

Moritz Fleischmann[1], Jacqueline M. Bloemhof-Ruwaard[1], Patrick Beullens[2], and Rommert Dekker[3]

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

[2] K.U. Leuven, Centre for Industrial Management, Celestijnenlaan 300 A, 3001 Leuven, Heverlee, Belgium, patrick.beullens@cib.kuleuven.ac.be

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

4.1 Introduction

The different chapters of this book highlight manifold examples of reverse logistics programs. While these cases vary substantially with respect to products, actors, and underlying motivations, as discussed in Chapter 1, they share a number of fundamental managerial issues. One of these is the need for an appropriate logistics infrastructure.

Analogous with traditional supply chains, the various transformation processes that turn a returned item into a remarketable good need to be embedded in a corresponding logistics network. In conventional supply chains, logistics network design is commonly recognized as a strategic issue of prime importance (see, e.g. Chopra and Meindl, 2001, Chapter 3). The location of production, storage, and cross-dock facilities, and the selection of transportation links between them, are major determinants of supply chain performance. Analogously, logistics network design has a fundamental impact on the profitability of reverse logistics systems. In order to maximize the value recovered from used products, companies need to set up logistics structures that facilitate the arising goods flows in an optimal way. To this end, one needs to decide where to locate the various processes of the reverse supply chain, as introduced in Chapter 1, and how to link them in terms of storage and transportation. In particular, companies need to choose how to collect recoverable products from their former users, where to inspect collected products in order to separate recoverable resources from worthless scrap, where to re-process collected products to render them remarketable, and how to distribute recovered products to future customers.

In this chapter, we take a detailed look at logistics network design in a reverse logistics context. We start in Section 4.2 by highlighting key business issues and contrasting them with logistics network design for traditional 'forward' supply chains. The core part of the chapter then discusses a number of alternative modeling approaches that support the design of reverse logistics networks and allow for a quantitative analysis of the underlying tradeoffs. Section 4.3 addresses integer-programming-based approaches that build upon traditional facility location models. Section 4.4 considers stochastic programming approaches that focus on incorporating the aspect of uncertainty into the network design decisions, and Section 4.5 presents a stream of research based on continuous approximation techniques. Section 4.6 synthesizes the different modeling approaches by exploring them in a common numerical example that highlights the economics of reverse logistics networks. To conclude, Section 4.7 summarizes the key points of this chapter. Before embarking into a systematic analysis of reverse logistics network design, we illustrate some of the main issues in a real–life business example in the remainder of this section.

4.1.1 Illustrative Case: Reverse Logistics Flows at IBM

The electronics industry has been a key sector in the growth of product recovery management. Ever expanding market volumes on the one hand, and shorter product lifecycles on the other, result in huge amounts of used products being disposed of. In this light, it comes as no surprise that electronic waste has been a prime target of environmental regulation, as reflected in enacted or pending take-back obligations in several countries (see also Chapter 15). At the same time, modular product design and a relatively small extent of mechanical 'wear and tear' sustain the reusability of electronic products and components. Together, both developments result in significant value recovery potential.

Business activities of IBM, as one of the major players in this sector, involve several types of 'reverse' product flows, which together cover most of the categories outlined in Chapter 2. From a business perspective, the most important class concerns product returns from expiring lease contracts. To date, leases account for some 35% of IBM's total hardware sales. Furthermore, IBM has implemented take–back programs in several countries in North America, Europe, and East Asia, which allow business customers to return used products for free or for a small fee. For remarketable products, customers may even receive a positive contribution. In the consumer market, IBM is required to take back end–of–life products in several countries in Europe and East Asia, due to environmental regulation. Besides dealing with used products, IBM, as do most companies, faces a 'reverse' stream of new products, which includes retailer overstocks and canceled orders. This flow very much depends on contractual agreements along the supply chain (see also Chapter 12). Finally, it is worth mentioning returns of rotable spare parts as a fairly traditional type of closed-loop flow related to the service business: defective parts replaced in a customer's machine are sent back for repair and may then be stocked as spare parts again (IBM, 2001; Fleischmann, 2001).

Recognizing the growing importance of reverse logistics flows, IBM set up a dedicated business unit in 1998, which is responsible for managing all product returns worldwide. The main goal of this organization, named Global Asset Recovery Services (GARS), is to manage the dispositioning of returned items to maximize the total value recovered. To this end, GARS operates some 25 facilities all over the globe where returns are collected, inspected, and assigned to an appropriate recovery option (see IBM, 2001). Equipment that is deemed remarketable may be refurbished and then put into the market again. For this purpose, IBM operates nine refurbishment centers worldwide, each dedicated to a specific product range. Internet auctions, both on IBM's own website and on public sites, have become an important sales channel for remanufactured equipment. Equipment that does not yield a sufficient value as a whole is sent to a dismantling center in order to recover valuable components, such as hard–disc assemblies, cards, and boards, which can be fed into IBM's spare parts network or sold on the open market (for a detailed description of this channel, see Fleischmann, 2001 and Fleischmann et al., 2003). The remaining returned equipment is broken down into recyclable material fractions, which are sold to external recyclers. In 2000, IBM reports the processing of 51,000 t of used equipment, of which only a residual of 3.2% was landfilled.

The above processes concern equipment from the business market. Given the much lower market value, consumer returns follow a different road. To work around inefficiencies of individual collection, IBM participates in cooperative, industry-wide solutions for this market sector in several countries. In the Netherlands, for example, IBM supports a system organized by the Dutch association of information and communication technology producers, in response to national product take–back legislation. In this case, used machines from different manufacturers are collected by the municipalities and then shipped to recycling subcontractors. Transportation and recycling costs are shared by the member organizations, proportional to their products' volume contribution (see Nederland ICT, 2002). Yet another system has recently been implemented in the USA. Since November 2000, IBM customers have the option to purchase a recycling service together with any new PC. Once the equipment is no longer needed, the customer sends it by UPS to a dedicated recycling center where it is either prepared for donation to charities or broken down into recyclable materials (IBM, 2000).

4.2 Network Design Issues in Reverse Logistics

4.2.1 Delineation of Reverse Logistics Networks

The above example underscores the need for a logistics infrastructure that accumulates used products and conveys them to recovery facilities and eventually to another user. In general terms, such a structure can be viewed as the logistics link between two market interfaces, which provide a supply of used products and demand for reusable products, respectively. Moreover, this link encompasses the reverse channel functions highlighted in Chapter 1, namely collection, testing and sorting, re–processing, and re–distribution. Figure 4.1 depicts a general scheme of this perspective. It is worth pointing out that the two markets involved may coincide, thereby implying a closed–loop network.

From a logistics perspective, one may characterize the structure illustrated in Figure 4.1 as a many–to–many distribution network. Within this layout, one may distinguish a convergent inbound part corresponding to the collection and acquisition function and a divergent outbound part serving a distribution function. The intermediate part of the network hosts the actual transformation processes. Therefore, its structure very much depends on the type of re–processing involved.

One may argue that it is only the inbound part of the network that actually concerns 'reverse' logistics processes, whereas the remainder very much corresponds with a traditional production–distribution network. However, as discussed in Chapter 1, this segregation may hamper a comprehensive analysis since the different product flows are closely interrelated. In fact, in this light one may wish to extend the scope even further and also include the distribution of the original new products (see Figure 4.1). It should be clear that this broad scope does not mean that the entire network is, or should be, managed by a single company. As in a traditional supply chain context, responsibilities may be allocated to multiple players. However, in line with the supply chain management imperative, one should consider the complete picture in order to understand the economics of reverse logistics networks.

Within the above setup, examples of reverse logistics networks are far from identical. Significant differences concern, for example, the players involved and

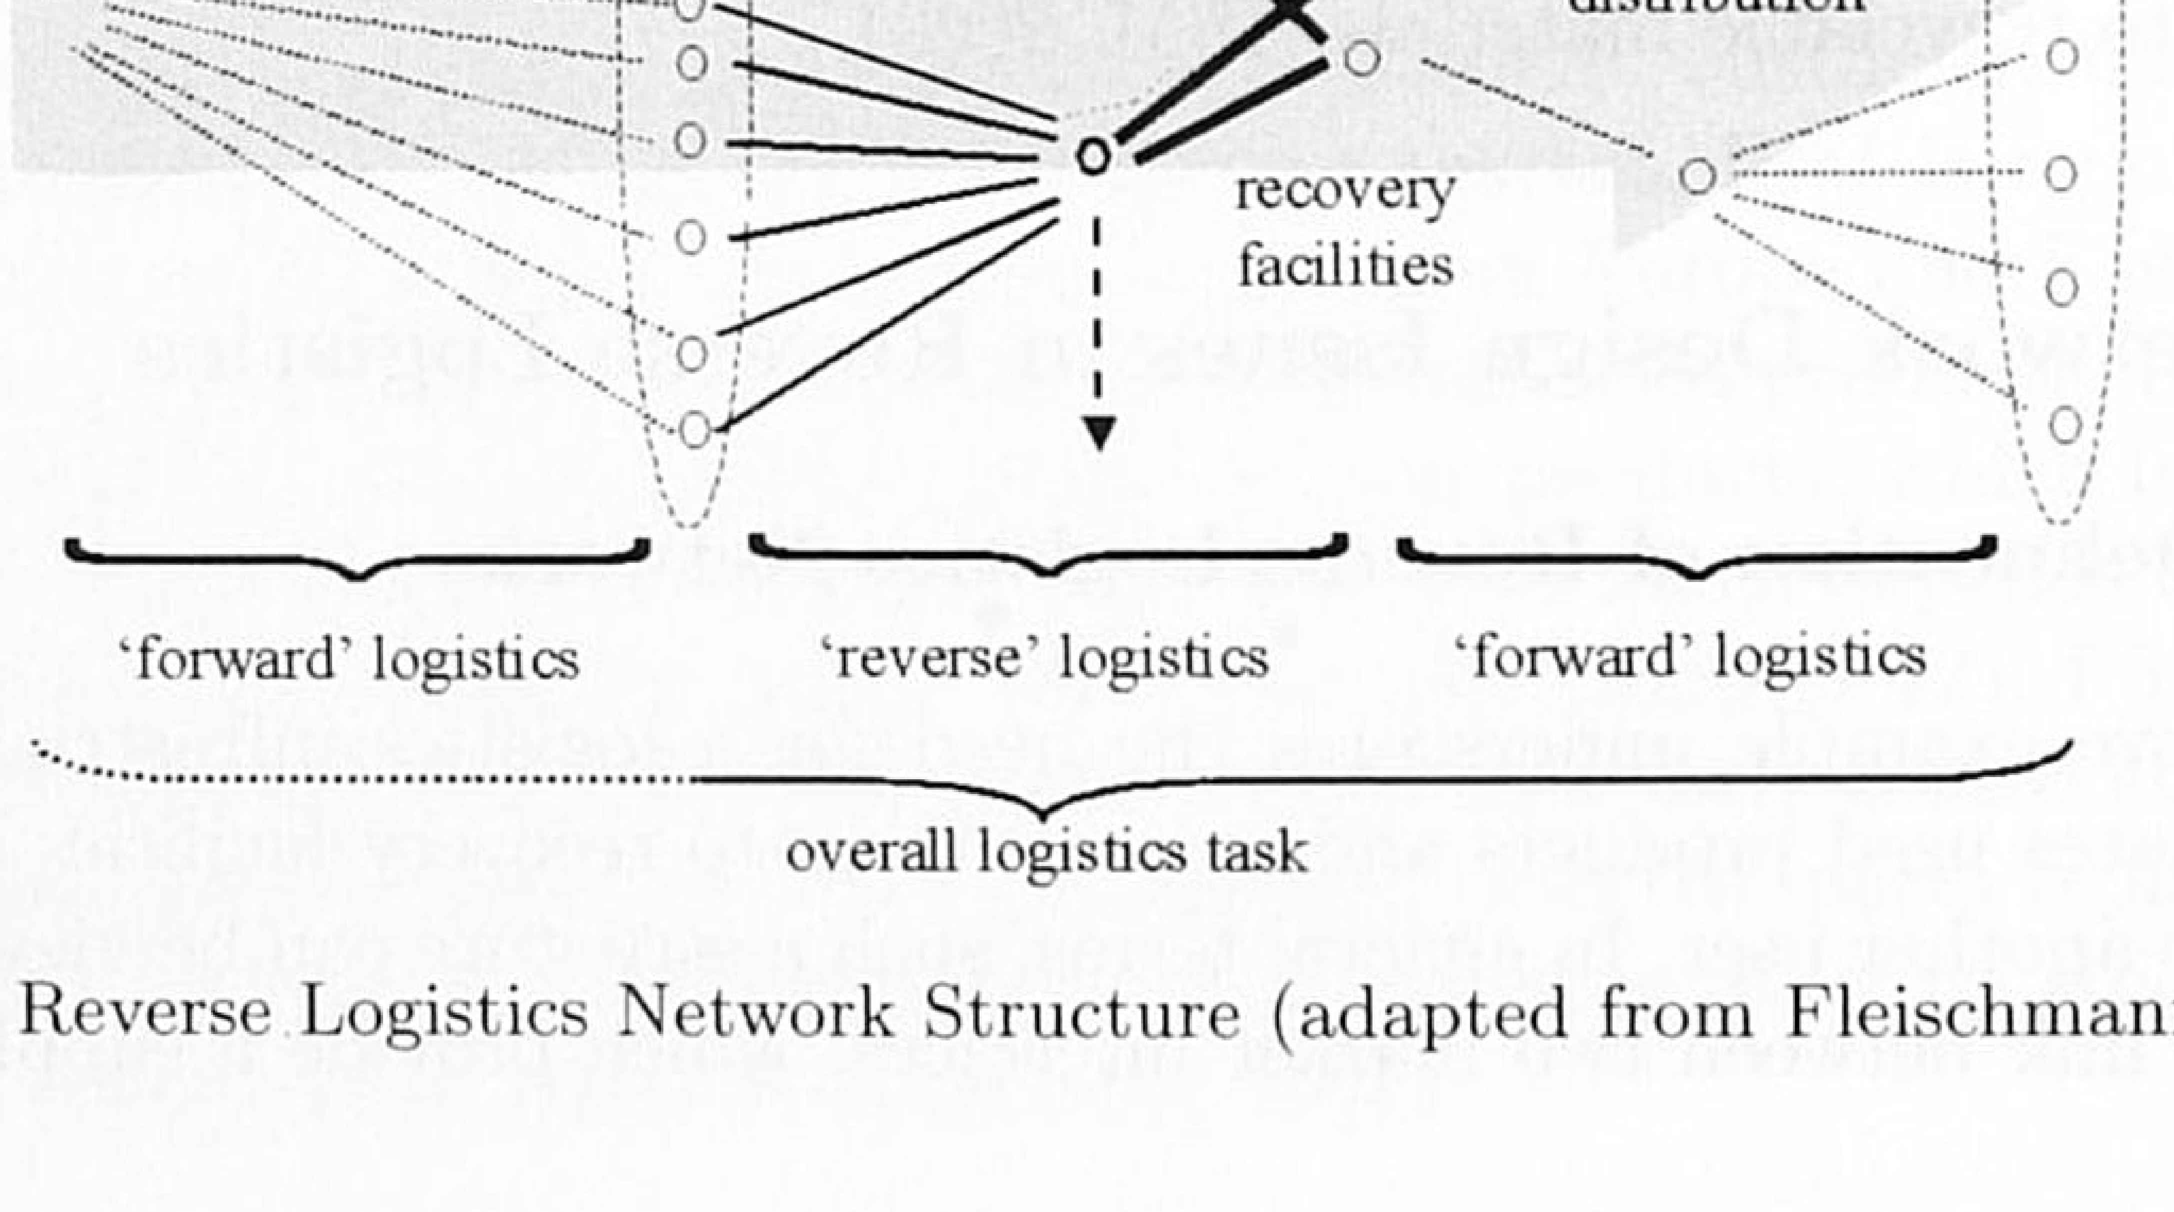

Fig. 4.1. Reverse Logistics Network Structure (adapted from Fleischmann, 2003)

their responsibilities, but also the network structure in terms of centralization and the number of layers. In the literature, several classifications have been proposed for structuring this field.

A first stream of research focuses on the reverse channel structure and the roles of the different players involved. In this vein, Fuller and Allen (1995) distinguish (1) manufacturer–integrated systems, (2) waste–hauler systems, (3) specialized reverse dealer–processor systems, (4) forward retailer–wholesaler systems, and (5) temporary/facilitator systems. The analysis extends earlier suggestions by Guiltinan and Nwokoye (1975) and Pohlen and Farris (1992).

In a different perspective, Bloemhof–Ruwaard and Salomon (1997) and Fleischmann et al. (2000) attribute differences between reverse logistics networks primarily to the form of the recovery process. The authors then distinguish three network types, namely remanufacturing, recycling, and direct reuse networks. Fleischmann (2003) refines this model by including ownership of the recovery process (original equipment manufacturer (OEM) versus third party) and recovery drivers (economic versus legislative) as additional explanatory variables. Based on this analysis, the paper suggests distinguishing the following five network classes: (1) networks for mandated product take–back, (2) OEM networks for value added recovery, (3) dedicated remanufacturing networks, (4) recycling network for material recovery, and (5) networks for refillable containers.

4.2.2 Characteristics of Reverse Logistics Networks

Strategic design decisions related to the type of logistics networks delineated above include the choice of a collection/acquisition method, the location and capacity of the sorting and re–processing operations and corresponding inventory buffers, and the definition of various transportation links in terms of sourcing, modes, and capacities. When comparing these tasks with the design of a conventional production–distribution network, the network structure may seem the most apparent discriminating factor. As pointed out above, reverse logistics implies a many–to–many structure composed of a convergent and a divergent part (see Figure 4.1), whereas production–distribution networks are typically perceived as few–to–many, divergent structures. However, this difference may be a matter of scope rather than an intrinsic element of reverse logistics. Taking a supply chain perspective, and hence taking into account the supplier level in the design of a 'forward' distribution network, one obtains a picture that very much resembles Figure 4.1.

In contrast, the following three factors are specific of reverse logistics networks in a more fundamental sense:

- supply uncertainty,
- degree of centralization of testing and sorting, and
- interrelation between forward and reverse flows.

In traditional supply chains, demand is typically perceived as the main unknown. In a reverse logistics setting, however, it is the supply side that accounts for significant additional uncertainty. Used products are a much less standardized input resource than conventional component supplies or raw materials. As pointed out in Chapter 1, quantity, quality, and timing of product returns are, in general, not known with certainty and may be difficult to influence. Effectively matching demand and supply, therefore, is a major challenge in reverse logistics. Consequently, robustness with respect to variations in flow volumes and composition is an important prerequisite of reverse logistics networks.

The need for testing and sorting operations in reverse logistics is a direct consequence of the above supply uncertainty. The degree of centralization of this stage has a fundamental impact on the transportation needs in a reverse logistics network and is subject to the following tradeoff: testing collected products early in the channel may minimize the total transportation distance since inspected products can be sent directly to the corresponding recovery operation. In particular, this approach helps avoid excessive transportation of worthless scrap. On the other hand, investment costs, for example for advanced test equipment or specially trained labor, may call for centralizing the testing and sorting operations. There appears to be no direct equivalent to this issue in traditional production–distribution networks as product routings are, in principal, known beforehand in this case. Yet, to some extent the underlying trade-off resembles the effect of risk pooling on inventory location decisions.

Another important characteristic of reverse logistics networks concerns potential synergies between different product flows. While traditional distribution networks typically act as one-way streets, closed–loop chains encompass multiple inbound and outbound flows crossing each other's paths. In this setting, it is intuitive to consider integration as a potential means for attaining economies of scale. Opportunities may concern transportation and facilities. For example, integrating the collection of used products with the distribution function may help reduce empty rides. Similarly, integrating operations of the forward and reverse channel in the same facilities possibly reduces overhead costs. At the same time, these opportunities raise a compatibility issue. In many cases, closed-loop supply chains are not designed in a single step but are realized by adding reverse logistics activities to an existing distribution network. It is not clear, however, whether such a sequential approach yields a good solution or whether one should consider an integral redesign of the entire closed–loop network.

In what follows, we review quantitative models that aim at supporting the above network design decisions. Throughout, we pay particular attention to the aforementioned characteristics of reverse logistics networks and discuss how they are captured in the different modeling approaches. In analogy with traditional network design models, we focus on location–allocation decisions and, to a lesser extent, capacity selection. Transportation and the collection strategy are considered on a rather aggregated level here. Chapter 5 zooms in on the transportation operation in more detail. Similarly, Chapter 6 details warehousing and material handling aspects.

4.3 Mixed Integer Location Models for Reverse Logistics Network Design

4.3.1 Literature Review of Reverse Logistics Location Models

The most widespread modeling approach to logistics network design problems in various contexts concerns facility location models based on mixed integer linear programming (MILP). Throughout the decades, an extensive body of literature has been established that ranges from simple uncapacitated plant location models to complex capacitated multi-level, multi-commodity models. At the same time, powerful solution algorithms have been proposed, relying on combinatorial optimization theory. For a detailed overview of models and solution techniques, we refer to Mirchandani and Francis (1989) and Daskin (1995).

Given this extensive body of research, MILP location models appear to be a natural starting point for quantitative approaches to reverse logistics network design. Several authors have followed this route and have presented MILP location models adapted to a reverse logistics context. Table 4.1 provides an overview of the corresponding literature. We distinguish models that encompass the entire network between the two market interfaces sketched in Figure 4.1 and models with a scope restricted to the 'reverse' network part in a strict sense. Moreover, we indicate whether supply of used products is modeled as a push or a pull process, i.e. whether there is a given collection volume that needs to be processed or whether collection primarily responds to demand.

The summary in Table 4.1 indicates that most of the models published to date address the entire network scope and treat supply as a push process. The model of Kroon and Vrijens (1995), which is applied in the context of reusable packaging, essentially is a conventional uncapacitated warehouse

Table 4.1. Reverse Logistics Facility Location Models

	supply push	*supply pull*
integral network	Kroon and Vrijens (1995) Thierry (1997) Spengler et al. (1997) Barros et al. (1998) Marín and Pelegrín (1998) Fleischmann et al. (2001)	Realff et al. (1999) Jayaraman et al. (1999)
reverse network	Berger and Debaillie (1997)	Krikke et al. (1999)

location model with lateral transshipments. Similarly, Marín and Pelegrín (1998) consider a special case of a warehouse location model where each customer's supply equals a fixed fraction of his demand. Jayaraman et al. (1999) analyze a multi–product variant of this model with general supply and demand volumes. Moreover, the supply process is governed by limited core availability rather than by collection obligations.

Thierry (1997) considers a linear programming model that corresponds with the structure outlined in Figure 4.1 with facility locations being fixed. The disposal volume arising at the testing stage is modeled as a fixed fraction of the volume processed. Berger and Debaillie (1997) include location decisions in this model while at the same time limiting its scope to the 'reverse' network part. Moreover, they model the disposal volume as a lower bound rather than a fixed fraction. Krikke et al. (1999) apply a similar model in a case study on copier remanufacturing. In Fleischmann et al. (2001), we analyze a generalization of Thierry's model, including location decisions. This model is discussed in detail in Section 4.3.2 below.

The model presented by Barros et al. (1998) captures a more detailed picture in that it explicitly includes an alternative recovery path rather than an external scrap process for material rejected at the testing stage. Spengler et al. (1997) and Realff et al. (1999) take an even broader perspective by modeling multi–commodity flows in general processing networks. While both cases are motivated by applications in the process industry, they differ in their view of the supply process. The former considers a supply push in a waste recycling context whereas the latter focuses on the recoverable value of the potentially available supply.

The above contributions exhibit much similarity with traditional multi–level location models. From a mathematical perspective, the particular characteristics of reverse logistics identified in the previous section appear to entail only minor modifications. Specific features include additional flow constraints, reflecting supply restrictions. Other variations are due to multiple return flow dispositions and to a possible interaction between forward and reverse channels. As a consequence, most of the models use multi–commodity flow formulations. In the next section, we discuss these aspects in more depth on the basis of a specific model.

4.3.2 A Basic Facility Location Model

To make things specific, let us take a look at a concrete MILP formulation of a reverse logistics network design problem. To this end, we discuss a variant of the model we introduced in Fleischmann et al. (2001). The model picks up the general scheme sketched in Figure 4.1. Specifically, it encompasses the processes between the two market interfaces discussed in Section 4.2. In this setting, the model considers three levels of facilities for a single type of product, namely test centers, factories, and distribution warehouses. Moreover, it

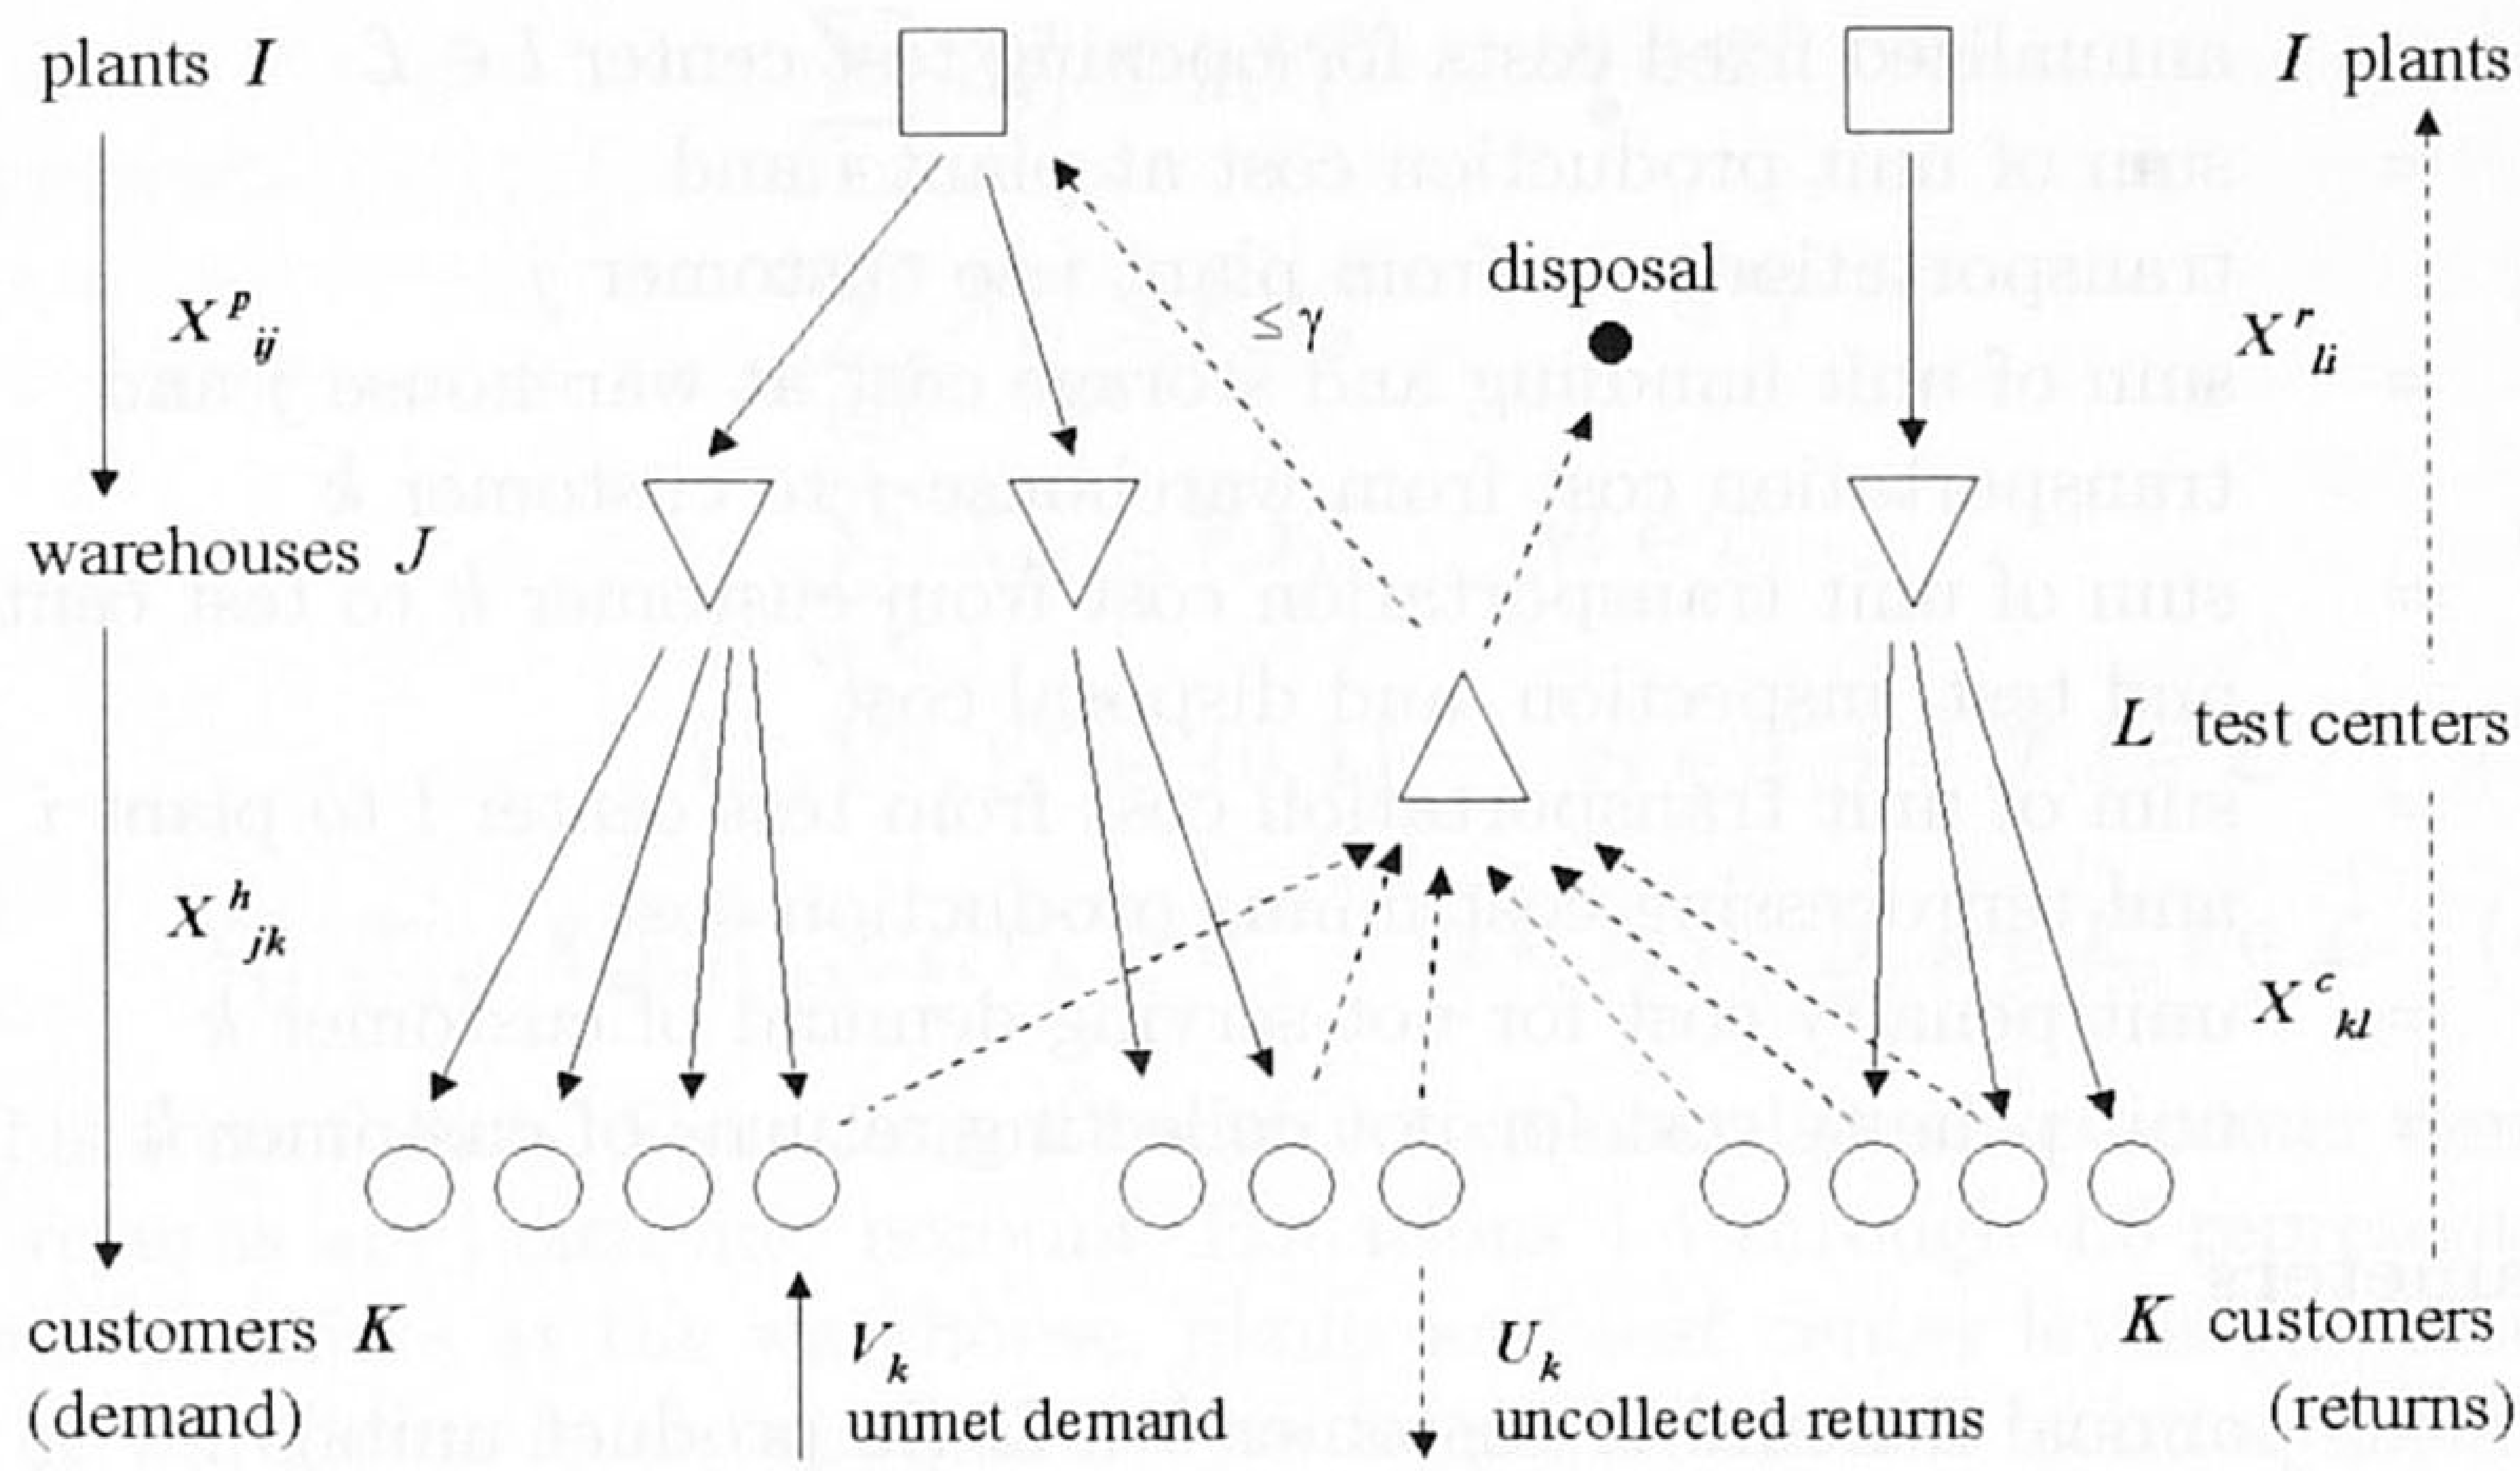

Fig. 4.2. Structure of the Recovery Network Model (adapted from Fleischmann et al., 2001)

includes two generic dispositions for the flow of used products, namely recovery and disposal, where recovery is restricted to a certain maximum yield. Figure 4.2 displays the general structure of this model. The MILP formulation below uses the following notation.

Index sets

$\mathcal{I}$ = set of potential plant locations
$\mathcal{J}$ = set of potential warehouse locations
$\mathcal{K}$ = set of fixed customer locations
$\mathcal{L}$ = set of potential test center locations

Variables

Y^p_i = indicator opening plant $i \in \mathcal{I}$
Y^h_j = indicator opening warehouse $j \in \mathcal{J}$
Y^r_l = indicator opening test center $l \in \mathcal{L}$
X^p_{ij} = product flow from plant i to warehouse j (in product units)
X^h_{jk} = product flow from warehouse j to customer k (in product units)
X^c_{kl} = product flow from customer k to test center l (in product units)
X^r_{li} = product flow from test center l to plant i (in product units)
V_k = unsatisfied demand of customer k (in product units)
U_k = excess supply of used products from customer k (in prod. units)

Costs

f^p_i = annualized fixed costs for opening plant $i \in \mathcal{I}$
f^h_j = annualized fixed costs for opening warehouse $j \in \mathcal{J}$

f_l^r = annualized fixed costs for opening test center $l \in \mathcal{L}$
c_{ij}^p = sum of unit production cost at plant i and transportation cost from plant i to customer j
c_{jk}^h = sum of unit handling and storage cost at warehouse j and transportation cost from warehouse j to customer k
c_{kl}^c = sum of unit transportation cost from customer k to test center l and test, inspection, and disposal cost
c_{li}^r = sum of unit transportation cost from test center l to plant i and reprocessing cost minus production cost
c_k^b = unit penalty cost for not serving demand of customer k
c_k^0 = unit penalty cost for not collecting returns of customer k

Parameters

d_k = annual demand of customer $k \in \mathcal{K}$ (in product units)
u_k = annual returns of used products from customer $k \in \mathcal{K}$ (in product units)
γ = average recovery yield
$\overline{p}_i$ = annual capacity of plant $i \in \mathcal{I}$
$\overline{h}_j$ = annual capacity of warehouse $j \in \mathcal{J}$
$\overline{r}_l$ = annual capacity of test center $l \in \mathcal{L}$

We then formulate the general reverse logistics network design (RLND) model as

$$\begin{aligned} \mathit{min}\ ! \quad & \sum_{i\in\mathcal{I}} f_i^p\, Y_i^p + \sum_{j\in\mathcal{J}} f_j^h\, Y_j^h + \sum_{l\in\mathcal{L}} f_l^r\, Y_l^r \\ & + \sum_{i\in\mathcal{I}} \sum_{j\in\mathcal{J}} c_{ij}^p\, X_{ij}^p + \sum_{k\in\mathcal{K}} \big(c_k^b\, V_k + \sum_{j\in\mathcal{J}} c_{jk}^h\, X_{jk}^h \big) \\ & + \sum_{k\in\mathcal{K}} \big(c_k^o\, U_k + \sum_{l\in\mathcal{L}} c_{kl}^c\, X_{kl}^c \big) + \sum_{l\in\mathcal{L}} \sum_{i\in\mathcal{I}} c_{li}^r\, X_{li}^r \end{aligned} \tag{4.1}$$

subject to

$$\sum_{j\in\mathcal{J}} X_{jk}^h + V_k = d_k \qquad \forall k \in \mathcal{K} \tag{4.2}$$

$$\sum_{l\in\mathcal{L}} X_{kl}^c + U_k = u_k \qquad \forall k \in \mathcal{K} \tag{4.3}$$

$$\sum_{i\in\mathcal{I}} X_{ij}^p = \sum_{k\in\mathcal{K}} X_{jk}^h \qquad \forall j \in \mathcal{J} \tag{4.4}$$

$$\sum_{l\in\mathcal{L}} X_{li}^r \le \sum_{j\in\mathcal{J}} X_{ij}^p \qquad \forall i \in \mathcal{I} \tag{4.5}$$

$$\sum_{i\in\mathcal{I}} X_{li}^r \le \gamma \sum_{k\in\mathcal{K}} X_{kl}^c \qquad \forall l \in \mathcal{L} \tag{4.6}$$

$$\sum_{j\in\mathcal{J}} X^p_{ij} \le \overline{p}_i Y^p_i \qquad \forall i \in \mathcal{I} \tag{4.7}$$

$$\sum_{i\in\mathcal{I}} X^p_{ij} \le \overline{h}_j Y^h_j \qquad \forall j \in \mathcal{J} \tag{4.8}$$

$$\sum_{k\in\mathcal{K}} X^c_{kl} \le \overline{r}_l Y^r_l \qquad \forall l \in \mathcal{L} \tag{4.9}$$

$$Y^p_i, Y^h_j, Y^r_l \in \{0,1\} \qquad \forall i \in \mathcal{I}, j \in \mathcal{J}, l \in \mathcal{L} \tag{4.10}$$

$$X^p_{ij}, X^h_{jk}, X^c_{kl}, X^r_{li}, U_k, V_k \ge 0 \qquad \forall i \in \mathcal{I}, j \in \mathcal{J}, k \in \mathcal{K}, l \in \mathcal{L} \tag{4.11}$$

In this formulation, Equations 4.2 and 4.3 ensure that all customer demand and returns are taken into account. Equations 4.4 through 4.6 represent balance constraints at the warehouse, plant, and test center levels respectively. At the warehouse level, inbound and outbound flows need to be equal. At the plant level, a potential excess outbound volume corresponds with new production. Similarly, the excess inbound volume at the test center level, which is constrained by the recovery yield, corresponds with the disposal volume. Finally, Equations 4.7 through 4.9 are the usual facility opening conditions coupled with capacity constraints.

In order to speed up the solution process, the above formulation can be strengthened by adding the following valid inequalities (compare Bloemhof et al., 1996).

$$X^h_{jk} \le \min(d_k, \overline{h}_j)\ Y^h_j \qquad \forall j \in \mathcal{J}, k \in \mathcal{K} \tag{4.12}$$

$$X^c_{kl} \le \min(u_k, \overline{r}_l)\ Y^r_l \qquad \forall k \in \mathcal{K}, l \in \mathcal{L} \tag{4.13}$$

It should be noted that this model is rather general and can capture a large variety of reverse logistics situations. For example, closed-loop and open-loop structures both can be represented and are reflected in different settings of the parameters d_k and u_k. Specifically, a closed–loop situation is characterized by $d_k \cdot u_k > 0$ for at least some customer k. Similarly, push and pull drivers for used product collection are reflected in different penalty costs c^o_k. Furthermore, it is worth emphasizing that the 'disposal' route may include any form of recovery that is outsourced to a third party, e.g. material recycling.

Mathematically, the above formulation does not differ much from multi–level facility location models in a more traditional production–distribution context. A particular aspect concerns the two sets of exogenous parameters d_k and u_k, which are linked by the different balance conditions. This reflects the need in reverse logistics for matching market conditions on the supply and the demand side. Another element worth pointing out concerns the additional degree of freedom introduced by the yield condition 4.6. By constraining the disposal volume by a lower bound rather than by a fixed fraction, the recovery strategy and the network design are optimized simultaneously. A relevant question concerns the impact of these features on the performance of specific solution methods. To our knowledge, results on this issue are few to date. Verter et al. (2003) have recently presented a Lagrangean decomposition method that exploits the specific problem structure of a combined forward/reverse logistics network design. Initial numerical results seem promising.

4.3.3 Extensions

The (RLND) model introduced above can be extended in manifold ways. Analogous with traditional facility location models, the formulation can be generalized to a dynamic, capacity selection, multi–product setting. We do not elaborate on these features here since they are well known from other contexts. Instead, we indicate a number of additional elements that appear to be specific to a reverse logistics context. For mathematical details, we refer to Fleischmann et al. (2001).

- *Integrating forward and reverse channel facilities*
  As discussed in Section 4.2, integration versus separation of different processes is an important issue in reverse logistics. For example, co–locating a warehouse and a test center may allow for sharing fixed assets and thereby exploiting economies of scale. This effect can be captured in the model by introducing additional indicator variables for combined facilities.

- *Integrating forward and reverse transportation flows*
  Similar synergies may arise from combining transportation routes for forward and reverse goods movements (see also Chapter 5). In the above setup, this can be modeled by means of additional flow variables representing simultaneous flows in both directions between two locations.

- *Distinguishing demand for new and recovered products*
  The above formulation includes only one class of demand, which may be fulfilled through either new production or recovery. Alternatively, one may wish to distinguish between markets for new and recovered products. In essence, this comes down to explicitly including the leftmost part of the scheme in Figure 4.1. Mathematically, this approach results in a multi–commodity network flow formulation.

- *Multiple recovery options*
  The above formulation uses the most basic representation of a recovery strategy in that it distinguishes two return dispositions, namely internal 'recover' versus external 'disposal'. In order to capture a more refined picture, one may wish to distinguish more recovery options. Mathematically, this extension again results in a multi–commodity formulation.

4.4 Stochastic Location Models for Reverse Logistics Network Design

4.4.1 Stochastic Mixed Integer Modeling Approaches

As discussed in Section 4.2, growing uncertainty on the supply side in particular is frequently named as a major characteristic of reverse logistics networks. In the mixed integer network design approaches presented in the previous section, uncertainty is, in general, addressed by means of scenario analyses. Thus a model is solved repeatedly for a set of scenarios and the solution with the best 'overall performance', according to some multi–criteria measure, is retained. In this section, we review modeling approaches that incorporate the aspect of uncertainty more explicitly.

For a general introduction to stochastic programming, we refer to Birge and Louveaux (1997). A stochastic mixed–integer linear program seeks to minimize the expected costs over a given set of scenarios with associated probabilities, subject to linear and integrality constraints. In the model definition, one needs to specify which decision variables need to be fixed before the realization of a scenario is known and which ones can be adjusted afterwards. Let us denote the vectors of both types of decision variables by Y and X, respectively. Moreover, let $\omega \in \Omega$ denote the set of scenarios. Then a stochastic mixed–integer linear program can be written as

$$\min\ c^T Y + E_\omega[c^*(\omega, Y)] \quad \text{s.t.} \quad Y \geq 0,\ Y_\mathcal{I} \in \{0,1\} \quad , \tag{4.14}$$

where c^* is the optimal value of a MILP in decision variables X, which depends on ω and Y, c is a vector of objective coefficients, and $Y_\mathcal{I}$ is some sub–vector of Y. If Ω is finite then (4.14) can be rewritten as an ordinary MILP, though at the expense of an increasing problem size, by introducing scenario-dependent decision variables X_ω. This approach is known as linear programming 'with recourse'.

It is important to note that the optimal solution of (4.14) need not be optimal for any single scenario. In this sense, stochastic programming is more powerful than a simple scenario analysis. This expansion comes at a cost however, since the problem size of the corresponding MILP formulation increases significantly.

In the context of logistics network design, stochastic programming models have been presented to capture the impact of demand uncertainty and price variations (see, e.g. Louveaux, 1986). Typically, these models assume that location decisions are fixed for a longer planning horizon (corresponding to variables Y in our formulation) whereas transportation flows can be adjusted in the short term, according to demand realizations (corresponding to variables X).

Stochastic programming models require a probability to be specified for each scenario. Since in practical applications these probabilities often are hard to define, some authors have argued that other optimality criteria may be more relevant. Instead of expected costs, they suggest considering worst-case criteria, such as minimizing the maximum cost across all scenarios or minimizing the maximum 'regret', i.e. the cost deviation from the corresponding scenario–optimal solution. These approaches do not require any probability specification but seek solutions that provide a good performance guarantee in all cases. For a general introduction to these so–called 'robust' optimization models, we refer to Kouvelis and Yu (1997). It should be noted, however, that despite their name, these approaches may be highly sensitive to the set of scenarios considered since extreme scenarios may strongly dominate the solution.

To our knowledge, two groups of authors have presented robust and/or stochastic extensions to network design models in a reverse logistics context. Realff et al. (2002) report on a case study on the design of a carpet recycling network in the USA. The authors extend a corresponding MILP facility location model to a multi–scenario setting, involving different levels of supply volumes and material prices, and seek to minimize the maximum regret across all scenarios. All binary variables, which represent location choices and capacity levels, are fixed at the beginning of the planning horizon whereas the values of all continuous variables are scenario dependent. In a numerical example, the authors illustrate that the optimal robust solution is not optimal, in general, for any of the individual scenarios considered. Information on the cost deviation between both approaches is not provided, however.

Listes and Dekker (2001) build upon the work of Barros et al. (1998) concerning a case study on the design of a sand recycling network in the Netherlands (see also Section 4.2). The authors extend the original MILP model to a stochastic model that maximizes expected profit under demand and supply uncertainty. In a first approach, they consider uncertain demand locations and volume. Location decisions for cleaning and storage facilities are assumed to be fixed at the beginning of the planning period, whereas all transportation, processing, and storage decisions may be adjusted to the demand realization. In a second approach, supply volumes are also uncertain. Decisions are now taken in three stages as the scenario realization is revealed successively. In a numerical study the authors document that the optimal stochastic solution need not coincide with the solution for any individual scenario. However, the cost deviation between the stochastic solution and the best solution obtained from a scenario analysis is within a few percentages in each of the cases presented.

4.4.2 A Stochastic Location Model for Reverse Logistics

We now apply the above stochastic modeling approaches to the reverse logistics network design model introduced in Section 4.3. To this end, let Ω denote a finite set of scenarios, and for each scenario $\omega \in \Omega$ let π_ω denote its probability. We assume that scenarios differ in terms of demand and return volumes and recovery yields, which we denote by $d_{k\omega}, u_{k\omega}$, and γ_ω, in analogy with Section 4.3. Then a stochastic version of the model in (4.1)-(4.11) can be formulated by adding a scenario index to the continuous variables X^P, X^h, X^c, X^r, V, and U, taking the expected value of the objective function across all scenarios and imposing restrictions (4.2)-(4.9) per scenario.

It is worth noting that the uncertain volume parameters concern the right-hand side of the MILP formulation, whereas the uncertain recovery yield affects the coefficient matrix. In contrast, all cost parameters are assumed to be fixed. Since all continuous variables depend on ω whereas the binary variables do not, location decisions are taken under uncertainty whereas transportation and processing flows can be adjusted to individual scenario realizations, in line with the above motivation. Furthermore, note that setting $V_{k\omega} = d_{k\omega}$ and $U_{k\omega} = u_{k\omega}$ for all k and ω always provides a feasible solution. Hence, each location decision is feasible for all scenarios.

Comparing this formulation with the original deterministic model in Section 4.3, we observe that the number of continuous variables and the number of constraints has increased with a factor of $|\Omega|$. To improve numerical solution procedures, the MILP formulation can be strengthened by means of valid inequalities analogous with Equations (4.12)-(4.13). We illustrate the relation between solutions of the deterministic and the stochastic model in Section 4.6.

4.4.3 Extensions

The above model can be modified in manifold ways, e.g. to allow for different scenario definitions or information evolution. As an illustration, we take a brief look at alternative optimality criteria and multi–stage decision approaches.

If the minimal costs vary largely across scenarios, then the expected cost criterion used in Section 4.4.2 may result in a biased solution in the sense that it is dominated by a few high-cost scenarios. In this case, minimizing the expected 'regret' may be a relevant alternative. To this end, the term $-\sum_{\omega\in\Omega} \pi_\omega \, c^*_\omega$ should be added to the expected cost function, where c^*_ω denotes the minimal costs for scenario ω in the original deterministic model.

The expected cost criterion may be difficult to apply since estimating the probabilities π_ω may not be straightforward in practical situations. As discussed above, optimizing the worst–case behavior may therefore be a useful alternative. For our model, this so-called 'robust' optimization approach comes down to introducing an additional decision variable Z, which is to be minimized under the additional constraint

$$\begin{aligned}
&\sum_{i\in\mathcal{I}} f^p_i \, Y^p_i + \sum_{j\in\mathcal{J}} f^h_j \, Y^h_j + \sum_{l\in\mathcal{L}} f^r_l \, Y^r_l \\
&\quad + \sum_{k\in\mathcal{K}} \Big(c^b_k \, V_{k\omega} + \sum_{j\in\mathcal{J}} c^h_{jk} \, X^h_{jk\omega} \Big) + \sum_{i\in\mathcal{I}} \sum_{j\in\mathcal{J}} c^p_{ij} \, X^p_{ij\omega} \\
&\quad + \sum_{k\in\mathcal{K}} \Big(c^o_k \, U_{k\omega} + \sum_{l\in\mathcal{L}} c^c_{kl} \, X^c_{kl\omega} \Big) + \sum_{l\in\mathcal{L}} \sum_{i\in\mathcal{I}} c^r_{li} \, X^r_{li\omega} \le Z \quad \forall \omega \in \Omega \,. \qquad (4.15)
\end{aligned}$$

Analogously, one may choose to minimize the maximum regret by combining both of the above approaches. We illustrate the effect of the different cost criteria in Section 4.6.

Finally, it is worthwhile to take another look at how the scenario is revealed and hence at which information is available for which decision. As explained before, the above formulation implicitly assumes that all location decisions are taken before the actual scenario is known, whereas all other decisions are based on its realization. In this sense, the model captures a two-stage decision process. However, as discussed in Section 4.2, the design of a reverse logistics network may involve more stages, in particular if recovery activities are integrated into an existing 'forward' distribution network. One way to capture such a sequential decision process is to separate the scenario space into two independent sets $\Omega = \Xi \times \Psi$ concerning demand–related information (captured by parameters $d_{k\xi}$) and return–related information (captured by $u_{k\psi}$ and γ_ψ), respectively. The degree of information that is available for the different decisions can then be modeled by indexing the decision variables as Y_i^p, Y_j^h; $X_{ij\xi}^p, X_{jk\xi}^h, V_{k\xi}, Y_{l\xi}^r$; $X_{kl\xi\psi}^c, X_{li\xi\psi}^r$, and $U_{k\xi\psi}$ and modifying (4.1) - (4.11) accordingly.

4.5 Continuous Approximation Models for Reverse Logistics Network Design

4.5.1 Approximating Reverse Logistics Costs and Revenues

MILP–based location models as discussed in the preceding sections provide a powerful tool which can be tailored to a variety of different settings. Yet these approaches have some drawbacks when it comes to establishing general insights into the economics of logistics systems. Capabilities for sensitivity analyses in MILP models are limited and, even more importantly, the interrelation between various parameters is not made explicit. Therefore, conclusions on the behavior of a given real-life system often rely on extensive numerical experiments rather than on analytic arguments.

In view of this shortcoming, several authors have considered continuous cost expressions as a basis for alternative approaches to investigating logistics costs and optimizing the design of logistics systems. In particular, Daganzo has promoted this route in what has become known as the 'continuous approximation methodology' (Daganzo, 1999). The key element of this approach is the representation of demand by a continuous density function, as opposed to the discrete demand representations in traditional MILP approaches. If the demand density and other system parameters vary sufficiently slowly across the given service region (which may have spatial and temporal dimensions), logistics costs can be reasonably approximated by appropriately chosen averages, which can be expressed as fairly simple functions in a limited number of parameters. In this way, the cost impact of critical system parameters can be revealed and guidelines for the design of logistics structures can be derived.

In this section, we follow our reasoning in Fleischmann (2003) in applying the 'continuous approximation' approach to the analysis of reverse logistics networks. We consider a setting analogous to the one in Section 4.3. However, for the time being we restrict the modeling scope to the 'reverse' network part in a strict sense, i.e. the logistics structure conveying used products from collection points via inspection and sorting centers to a given recovery facility (compare Figure 4.1). An extension of the model to the entire network, including the redistribution stage, is discussed at the end of this section. Our goal is to approximate the total reverse logistics costs for serving a given area, and eventually to minimize these costs by choosing an appropriate reverse logistics network design. To this end, assume that the return rate of used products per time per unit surface is given by a location–dependent continuous density function, which varies slowly within the service area. The subsequent development is facilitated by considering costs on a per product returned basis. (Note that this criterion differs from total costs just by a scaling factor.) The core idea of the 'continuous approximation' approach then is to express these costs in 'local' problem parameters only and to approximate the overall costs by integrating over the service area.

To assess the unit reverse logistics costs, we distinguish two cases, depending on whether the testing and sorting is carried out at the recovery facility or at a separate location. In what follows, we refer to these cases as 'central' and 'local' testing, respectively. For both cases, one may decompose the total reverse logistics costs into a number of components, namely inbound transportation costs to the test and sort process, outbound transportation costs after sorting, variable sorting and handling costs, and fixed installation costs for the test facility. In what follows, we go through all of these components and discuss the parameters on which they depend. In addition to the symbols introduced earlier, we use the following notation.

A	=	overall service area
$\rho(x)$	=	return rate of used products per time per unit surface at location $x \in A$
$C_R(\ell, \rho)$	=	reverse logistics costs per returned product for a service area with constant return rate ρ at a distance ℓ from the corresponding recovery facility
$C_{RL}(\ell, \rho)$	=	— in the case of local testing
$C_{RC}(\ell, \rho)$	=	— in the case of central testing
c_t	=	low volume vehicle transportation cost per distance
$\tilde{c}_t$	=	high volume vehicle transportation cost per distance
v	=	low volume vehicle capacity
$\tilde{v}$	=	high volume vehicle capacity
c_w	=	disposal costs per product
A_R	=	size of a test facility's service area

A_R^* = optimal size of a test facility's service area
ℓ^* = opt. distance for switching from central to local testing

The first cost component concerns inbound transportation costs to the test and sort process. We assume that used products are collected in milk–runs. The length of a tour can be assessed through a probabilistic analysis of the standard vehicle routing problem. Specifically, it can be approximated by a line-haul distance from and to the test and sort installation plus the sum of the expected distances between two consecutive collection stops (see e.g. Daganzo 1999). Assuming full vehicle loads, we get in the case of central testing

$$\text{unit inbound transp. cost (central)} \approx 2\,\frac{c_t}{v}\,\ell + 0.57\,c_t\,\rho^{-1/2}\,. \tag{4.16}$$

In the case of local testing and sorting, the line-haul distance depends on the size A_R of the area covered by the test facility. Assuming this area to have a circle-like shape with the test facility located at its center, the average line-haul distance approximately equals $2\sqrt{A_R}/3\sqrt{\pi}$ and one obtains

$$\text{unit inbound transp. cost (local)} \approx \frac{4}{3\sqrt{\pi}}\,\frac{c_t}{v}\,\sqrt{A_R} + 0.57\,c_t\,\rho^{-1/2}\,. \tag{4.17}$$

The next chapter presents more detailed versions of these expressions, which also take into account inventory accumulation (see Section 5.4).

The second cost component concerns outbound costs from the test location. In the case of central testing, this term encompasses merely the disposal costs for rejected products, which equal $c_w(1-\gamma)$ per unit. For local testing, one also needs to consider the flow of accepted products to the recovery facility. Recognizing the consolidation function of the test centers, we assume those shipments to be line-hauls rather than multi–stop tours. In the same vein, we assume a larger vehicle capacity than for the collection tours, and a corresponding mileage cost. For full vehicle loads, the outbound costs can then be expressed as

$$\text{unit outbound transp. and disposal cost (local)} \approx 2\,\frac{\tilde{c}_t}{\tilde{v}}\,\ell\,\gamma + c_w\,(1-\gamma)\,. \tag{4.18}$$

As a third cost term we consider the annualized fixed costs for a local test and sort installation. These can be approximated on a per product basis by

$$\text{unit fixed installation cost (local)} \approx \frac{f_r}{\rho A_R}\,. \tag{4.19}$$

Finally, any variable handling and processing costs may be aggregated into a term c_h. Summing up the four cost components yields an expression for the reverse logistics cost per collected product. Specifically, in the case of central testing and sorting, one obtains

$$C_{RC}(\ell,\rho) = 2\,\frac{c_t}{v}\,\ell + 0.57\,c_t\,\rho^{-1/2} + c_w\,(1-\gamma) + c_h\,. \tag{4.20}$$

For the local testing case, the corresponding expression still depends on the size of the collection area A_R. Equations (4.17) and (4.19) characterize the optimal size A_R^* of this area. First order conditions imply

$$A_R^* = \left(\frac{3\sqrt{\pi}\, f_r\, v}{2\, c_t\, \rho}\right)^{2/3} \approx 1.92 \left(\frac{f_r\, v}{c_t\, \rho}\right)^{2/3}. \tag{4.21}$$

Inserting this expression for A_R and summing up leads to the following cost function:

$$C_{RL}(\ell, \rho) = 2\, \frac{\tilde{c}_t}{\tilde{v}}\, \ell\, \gamma + 0.57\, c_t\, \rho^{-1/2} + c_w\, (1-\gamma) + c_h + 1.56 \left(\frac{c_t^2\, f_r}{v^2\, \rho}\right)^{1/3}. \tag{4.22}$$

Comparing $C_{RC}(.)$ and $C_{RL}(.)$ yields an appropriate service area for the central test and sort operation. Specifically, (4.20) and (4.22) define a critical distance ℓ^* from the recovery facility up to which central testing is preferable over local testing. Equating the cost functions yields

$$\ell^* = 0.78 \left(\frac{f_r\, v}{c_t\, \rho}\right)^{1/3} \left(1 - \frac{\tilde{c}_t\, v}{\tilde{v}\, c_t} \gamma\right)^{-1}. \tag{4.23}$$

Putting the above results together, one finally obtains the overall reverse logistics unit cost function $C_R(.)$ as $C_R(\ell, \rho) = \min\{C_{RC}(\ell, \rho), C_{RL}(\ell, \rho)\}$. As discussed above, total reverse logistics costs are then approximated by integrating over the service area $\int_A \rho(x)\, C_R(\ell(x), \rho(x))dx$.

In Section 4.6 we compare the above cost expressions with the results of the previously discussed discrete models and interpret them in the light of the reverse logistics issues identified in Section 4.2. Before doing so, we discuss a number of extensions and refinements to the above approach.

4.5.2 Extensions

The above formulas reflect a very basic cost model which can be extended in manifold ways. In particular, they do not include any inventory considerations and assume all vehicles to operate at full capacity. These assumptions can be relaxed by including decisions on lot sizes and dispatching frequencies. We refer to Chapter 5 for a detailed discussion of this issue. Furthermore, the formulas can be extended to a multi–product setting. However, since these refinements do not appear to exhibit any particular reverse logistics elements, and since they do not change the core of our argumentation, we content ourselves by referring to Daganzo (1999) for a more in-depth discussion of the 'continuous approximation' technique.

In a similar fashion, we can also derive cost expressions for the 'forward' parts of the network (see Figure 4.1). To this end, assume that products are shipped from the factory to the customers via distribution centers. Following the above analysis, one obtains the same formulas, where ρ is replaced by an appropriate demand density δ and γ equals one. In fact, this is the original model discussed by Daganzo (1999). In what follows, we denote these 'forward' logistics costs by $C_F(.)$.

By putting together $C_F(.)$ and $C_R(.)$ one may address the overall network structure. In particular, by considering $C_R(.)$ as inbound and $C_F(.)$ as outbound costs and including investments one may assess the size of a factory's service area. If $\delta(x)$ and $\rho(x)$ are roughly proportional, one can derive expressions similar to (4.21) with ρ replaced by $\delta+\rho\gamma$. However, a critical look seems advisable. On the one hand, the distance approximations may be less accurate since the number of distribution centers and test centers is much smaller than the number of customer locations in the original model. On the other hand, Equation (4.21) assumes the facility to be located close to the center of its service area. While this seems reasonable for a distribution center, it may not be evident for the location of a factory, which depends on additional factors such as tax rates and labor costs.

Finally, note that we have assumed return and disposal rates to be given and therefore have not included any revenues in the analysis. However, the above cost expressions can also be used to assess profitability of a recovery operation. In particular, the tradeoff between reverse logistics costs and production cost savings or additional revenues can be made explicit. To this end, denote by $C_{RN}(.)$ the unit cost for any used product that is not recovered (which may include, for example, lost revenues and/or fees for local recycling). The unit reverse logistics cost function $C_R(.)$ is then obtained by selecting the cheapest among the three options C_{RC}, C_{RL}, and C_{RN} for each value of ℓ and ρ.

4.6 Quantitative Analysis of Reverse Logistics Network Design Issues

Having reviewed alternative modeling approaches for supporting reverse logistics network design decisions, we now return to the issues highlighted in Section 4.2. In what follows, we exploit the above quantitative tools to analyze these issues and highlight the impact of key parameters on the economics of reverse logistics networks.

We illustrate the analysis in a numerical example adapted from Fleischmann et al. (2001). All computational results are based on an installation of the CPLEX 7.0 standard MILP solver on a Pentium 4, 1495 MHz PC. Consider the situation of an electronic equipment manufacturer operating in the European market (recall the case of IBM from Section 4.1; see also the copier business in Chapter 11 and in Thierry et al., 1995). Assume that used equipment, stemming, for example, from expiring lease contracts, is collected from the customers, remanufactured, and resold. To allow for remanufacturing, used equipment must meet specified quality standards. To this end, all collected equipment is inspected and tested. Rejected equipment is disposed

Table 4.2. Parameter Settings of Network Design Example

Description	Value	Model Parameter	
		discrete	continuous
Fixed cost per factory	5,000,000	f^p	f^p
Fixed cost per warehouse	1,500,000	f^h	f^h
Fixed cost per test center	500,000	f^r	f^r
Transportation costs per km per product			
factory—warehouse	0.0045	c^p	–
warehouse—customer	0.0100	c^h	–
customer—test center	0.0050	c^c	$2\,c^t/v$
test center—plant	0.0030	c^r	$2\,\tilde{c}^t/\tilde{v}$
Penalty cost unsatisfied demand	1,000	c^b	–
Penalty cost uncollected returns	1,000	c^o	–
Capacity factory	500,000	$\overline{p}$	–
Capacity warehouse	150,000	$\overline{h}$	–
Capacity test center	150,000	$\overline{r}$	–
Low volume vehicle capacity	20	–	v
Demand per 1,000 inhabitants	10	$d_k/\#$inh.	$\delta\times$pop.density
Return ratio	[0,0.9]	λ	λ
Recovery yield	0.5	γ	γ
Distance from factory	1,000	–	ℓ

of locally, while the remainder is shipped to the remanufacturing operation, which is co–located with an original manufacturing site.

To implement this example as a MILP model, we assume that customers are located in 50 major European cities (capitals plus cities larger than 500,000 inhabitants) and that demand is proportional to the population size. Moreover, we restrict the potential (re-)manufacturing locations to 20 main metropolitan areas, whereas distribution warehouses and test operations may be located in any of the 50 cities. Table 4.2 summarizes the parameter settings for this example.

We assume that all equipment that passes the test operation has a sufficient contribution margin to be remanufactured rather than disposed of. However, to avoid the cost figures being distorted by large blocks of sunk costs, we do not include variable (re-)manufacturing, handling, and disposal costs. To assess the overall profitability of the remanufacturing operation, these costs as well as sales revenues should be added to the results below.

As a starting point, we compute the optimal 'forward' distribution network for the above example, ignoring any reverse logistics activities. To this end, we solve the conventional two–level facility location model obtained by setting $u_k = 0$ for all k in the MILP model in Section 4.3. The solid lines in Figure

Table 4.3. Results of Network Design Example

Scenario	λ	Test Centers	Min. Cost	Regret in Case of Design		
				Scen. 9	Scen. 3	Robust
			€$\cdot 10^3$	€$\cdot 10^3$	€$\cdot 10^3$	€$\cdot 10^3$
0	0.0	–	0	4,000	1,500	2,000
1	0.1	D	2,600	2,580	603	951
2	0.2	D	4,700	1,660	206	402
3	0.3	GB,D,E	6,610	933	0	44
4	0.4	GB,D,E,I,HU	8,140	592	182	74
5	0.5	GB,D,E,I,HU	9,550	365	477	218
6	0.6	GB,D,E,I,HU	11,000	139	773	361
7	0.7	S,GB,D,F,E,I,HU	12,200	54	1,210	647
8	0.8	S,GB,D_1,D_2,E,I,HU,BG	13,500	0	1,680	964
9	0.9	S,GB,D_1,D_2,E,I,HU,BG	14,600	0	2,200	1,330
stochastic		GB,D,E,I,HU	ø 8,850			
robust		B,D,E,HU				$\leq$ 2,000

4.3(a) illustrate the resulting network structure, which includes one central manufacturing site in Frankfurt and seven regional warehouses. For the sake of clarity, flows from warehouses to customers are omitted. The corresponding annual costs equal €44.8m.

4.6.1 Impact of Supply Uncertainty

As discussed in Section 4.2, reverse logistics network design typically faces significant uncertainty concerning the supply of recoverable resources. In the above MILP model, the supply side is characterized by the parameters u_k and γ. In what follows, we analyze their impact on the optimal solution.

In the model formulation (4.1)-(4.11), the volume parameters u_k occur only on the righthand side. Therefore, standard MILP theory implies that the cost function depends on them piecewise linearly (see e.g. Jenkins, 1982). Moreover, for fixed binary variables, i.e. fixed facility locations, the cost function is convex in u_k for each k. A parametric analysis can be carried out by means of Jenkins's heuristic (Jenkins, 1982).

Let us now assume that $u_k = \lambda d_k$ for all k, i.e. the return ratio λ is identical across locations. Table 4.3 summarizes the results for different values of λ. More specifically, we vary the return ratio in steps of 0.1 in the interval [0, 0.9] and compute for each scenario the optimal reverse logistics network while keeping the forward network fixed to the above layout. The solution time for each scenario is in the order of a few seconds. The dashed lines in Figure 4.3(a) illustrate the solution for $\lambda = 0.4$, which encompasses five regional test centers.

Not surprisingly, the optimal number of test centers and the relevant reverse logistics costs increase with the return volume. However, as discussed

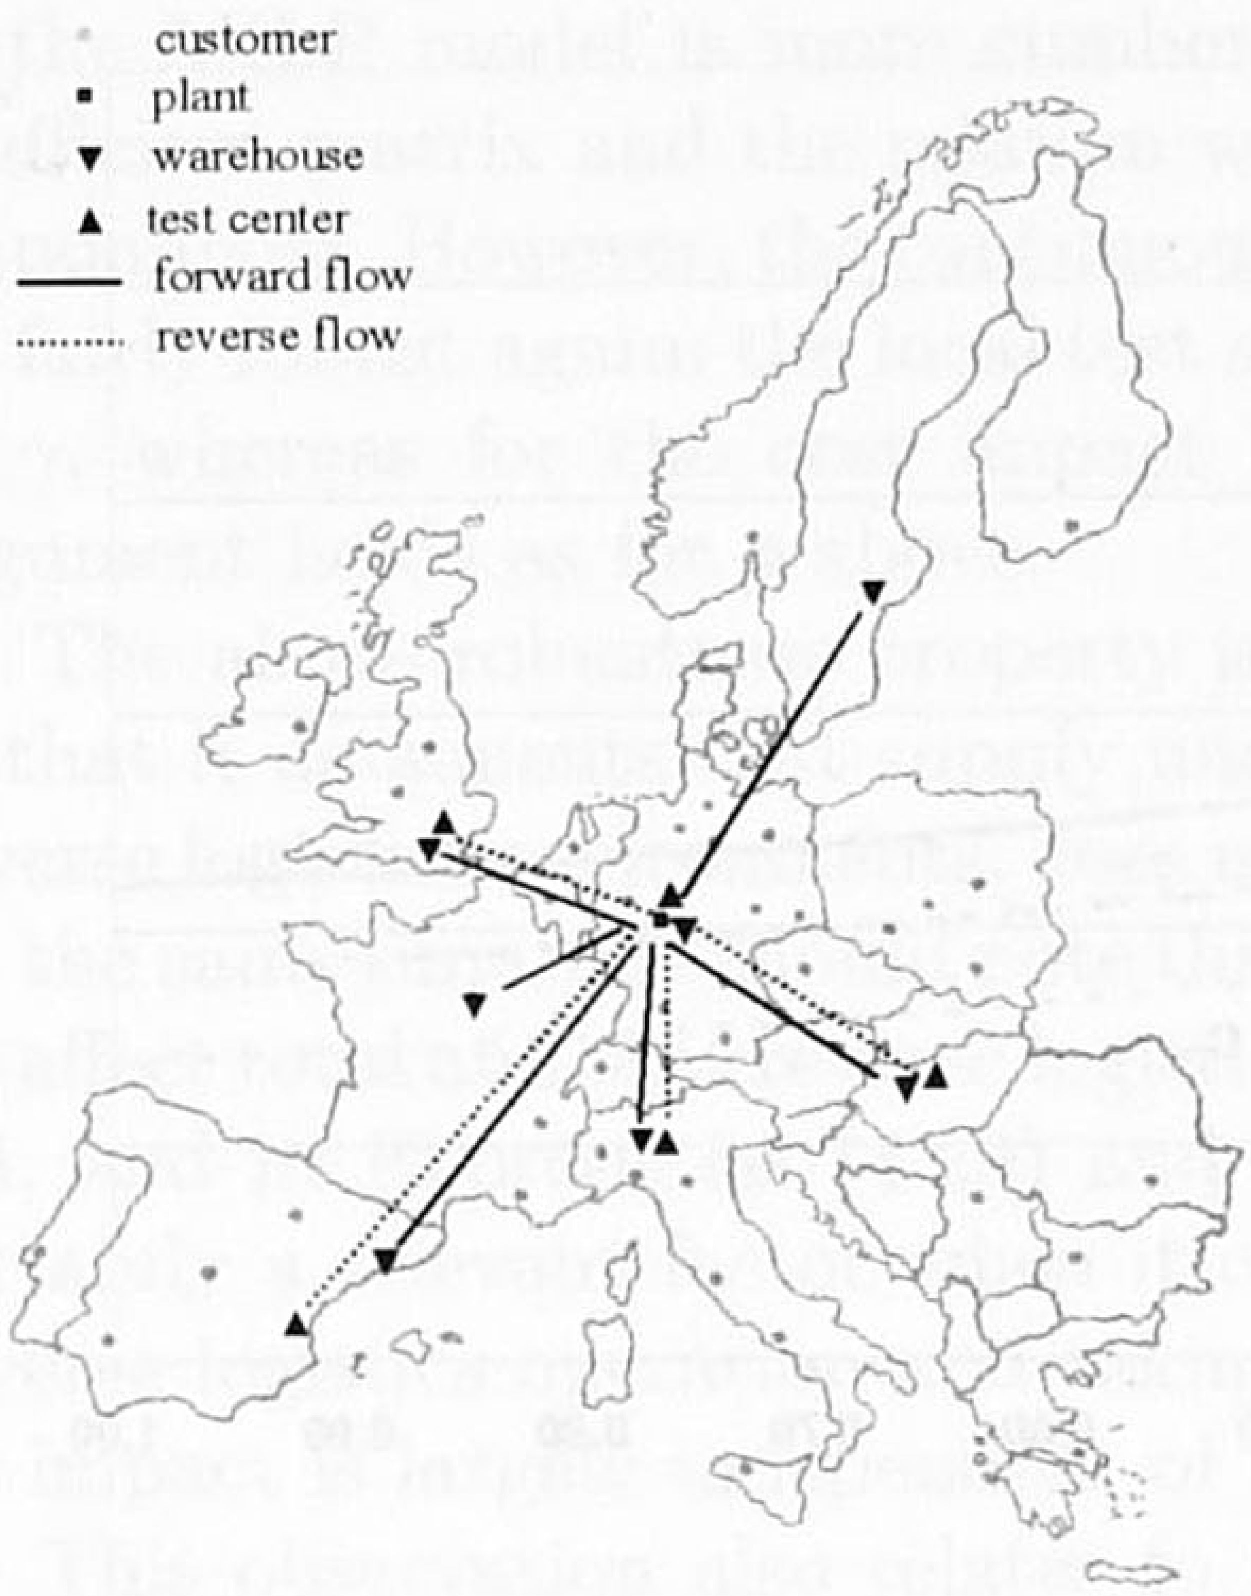

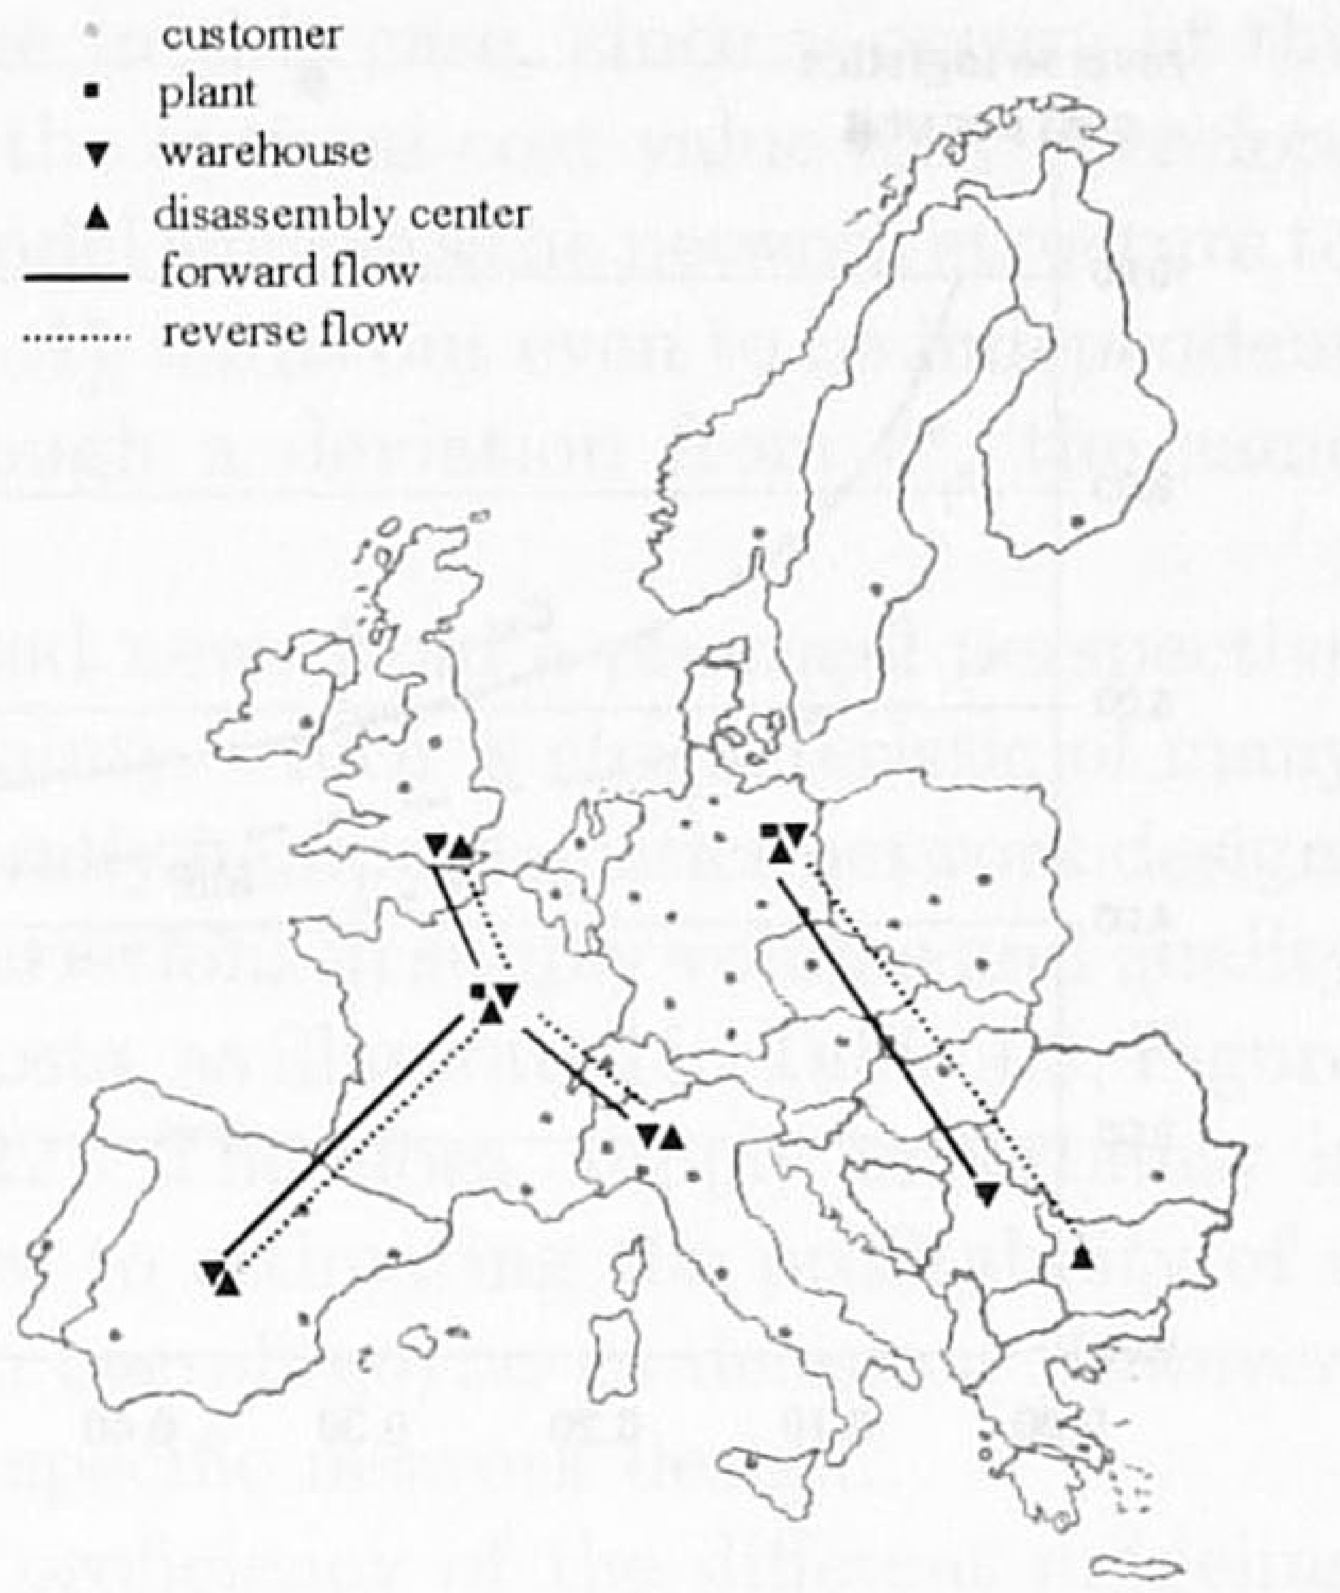

Fig. 4.3. Optimal Forward and Reverse Network Versus Optimal Integral Network

before, the actual return volume is not known, in general, when the location decision is to be taken. In Section 4.4, we have discussed modeling approaches that explicitly take this uncertainty into account. The next-to-last row of Table 4.3 characterizes the network design which minimizes the expected costs for the case of a uniform probability distribution across the above scenarios. The solution turns out to be identical to the optimal design for $\lambda \in [0.4, 0.6]$. As an alternative to this stochastic approach, we also compute an optimal 'robust' solution, which minimizes the maximum cost deviation from the scenario–optimal solution across all scenarios (see Section 4.4). Note that this solution is not optimal for any single scenario.

For conventional facility location models, it is well known that the cost function is fairly 'flat' around its minimum, in the sense that a deviation from the optimal network design entails a rather small cost penalty (see, e.g. Daganzo, 1999). An analysis of the continuous cost model developed in Section 4.5 supports a similar conclusion in a reverse logistics context.

To this end, Figure 4.4 illustrates the relation between the discrete and the continuous model for the above example by depicting the corresponding unit reverse logistics costs per product as a function of the return rate. For the discrete model, this cost curve is obtained by dividing the results in Table 4.3 by the return volume. For the continuous model, the curves display the functions C_{RC} and C_{RL} defined in (4.20) and (4.22). The parameter settings are listed in Table 4.2. Note that to make both modeling approaches compatible one needs to adjust c_t and $\tilde{c}_t$ to account for vehicle capacities and line–haul return trips. The values of ℓ and δ approximate the overall averages. Note

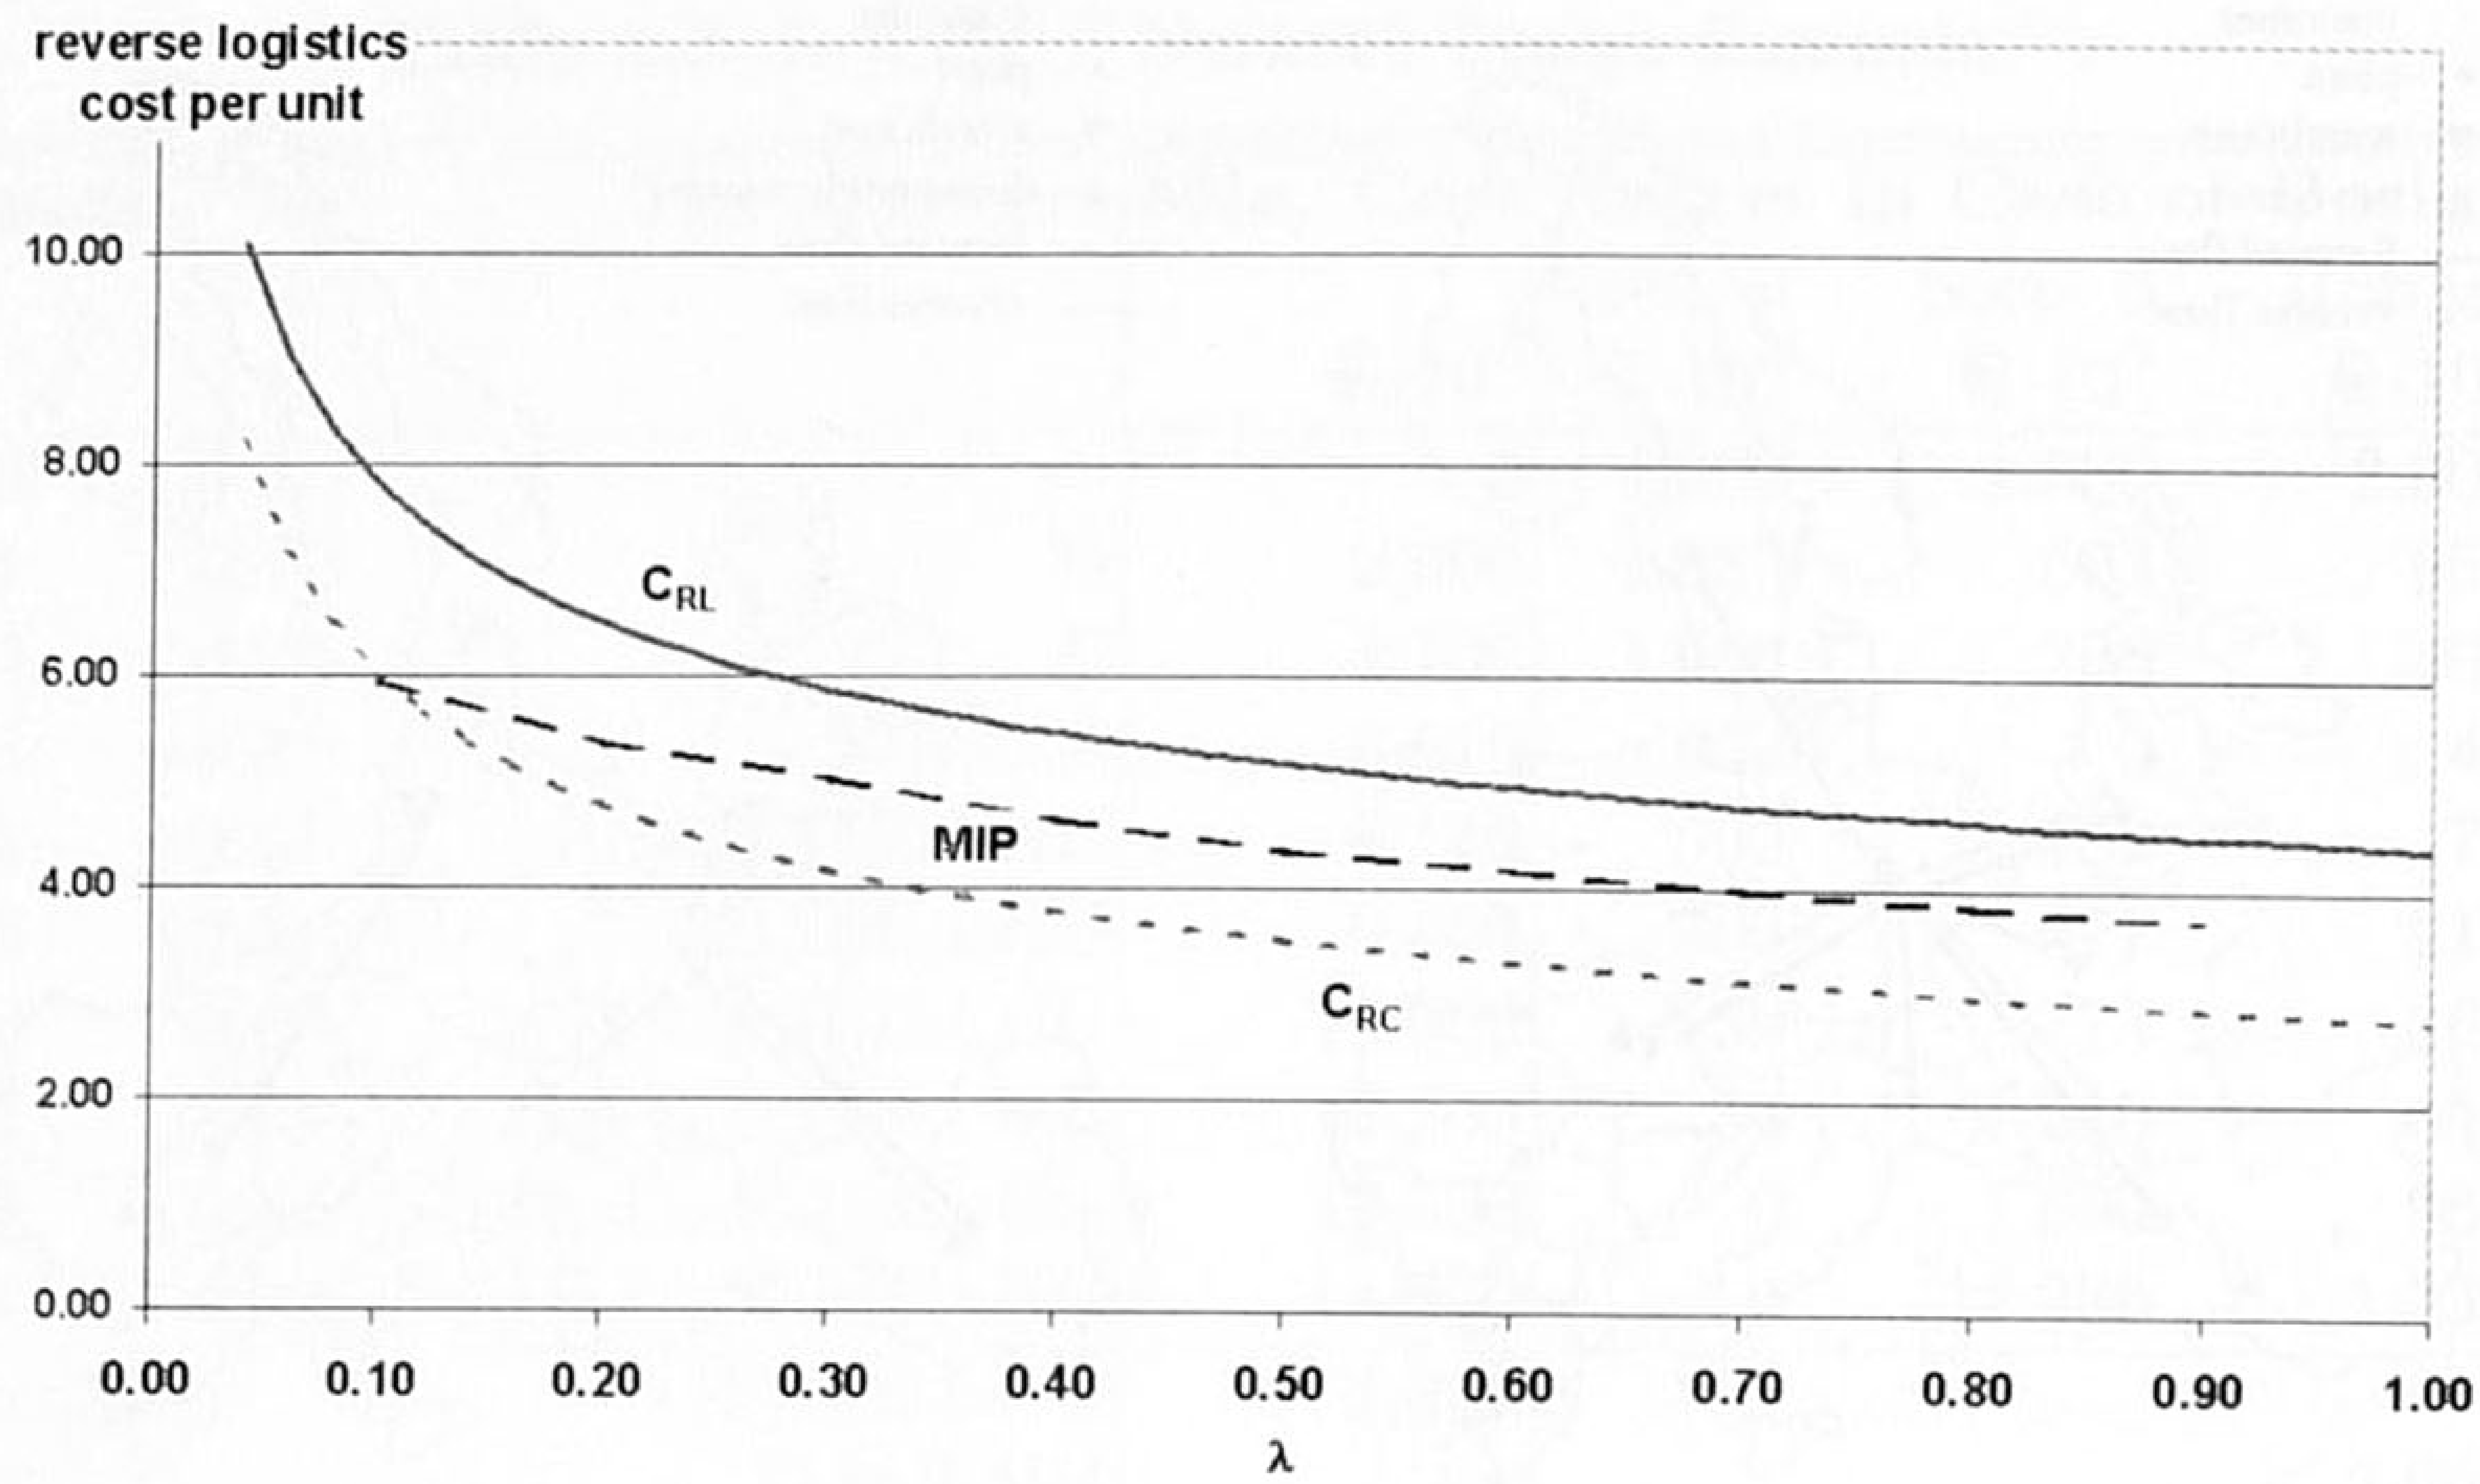

Fig. 4.4. Unit Reverse Logistics Costs

that the eventual unit reverse logistics cost function C_R in the continuous model is a mixture of C_{RC} and C_{RL}. Since the discrete cost function also lies in this interval, Figure 4.4 suggests the results of the discrete model and the continuous model to be compatible.

To quantify the impact of supply uncertainty on the network design, let $C_{RL}(A_R)$ denote the unit reverse logistics costs as a function of the test service area A_R in the case of local testing. From Equations (4.19) and (4.17) one gets that $C_{RL}(A_R)$ can be written as $a + b\sqrt{A_R} + c/A_R$, with some positive constants a, b, and c. Similar to the well-known case of the EOQ formula this function is very flat around its minimum. Specifically, for $\varepsilon > 0$ one gets

$$[C_{RL}((1+\varepsilon)A_R^*) - C_{RL}(A_R^*)]\,/C_{RL}(A_R^*) \quad \leq \quad \varepsilon^2/3(1+\varepsilon)\,. \qquad (4.24)$$

Furthermore, Equation (4.21) implies, for example, that a relative error of ε in ρ causes a relative error of at most $0.67\,\varepsilon$ in A_R^* and therefore by (4.24) a relative cost penalty of at most $0.22\varepsilon^2/(1.5+\varepsilon)$. This implies that a forecasting error of 50% in the return rate results in an eventual cost penalty of less than 3% in the network design decision.

For the design parameter ℓ^*, which characterizes the domain of central testing, one observes a similar robust behavior. Equation (4.23) shows that the impact on ℓ^* of an error in ρ is limited. Moreover, moving to a critical distance ℓ' different from ℓ^* only affects the costs for customers located at a distance between ℓ' and ℓ^*, which again has a dampening effect on the overall cost deviation.

So far, we have restricted our attention to variations in the return *volume.* To round off our analysis, let us take a brief look at the impact of the return quality, characterized by the yield parameter γ. An exact sensitivity analysis in the MILP model is more cumbersome in this case, since γ occurs in the coefficient matrix and the relation with the optimal cost value may therefore be nonlinear. However, the continuous model suggests the network structure to be fairly robust again: the local test area A_R^* turns out even to be independent of γ, whereas for the cost impact through a deviation from ℓ^*, the same argument holds as for ρ above.

The above robustness property is good news from a practical perspective in that it documents that supply uncertainty, which is characteristic of many reverse logistics environments, does not really hamper logistics network design. At the same time, one should note that variations in supply volume and quality do affect total and unit reverse logistics costs, as illustrated in Table 4.3, Figure 4.4, and in Expressions (4.20) and (4.22). Therefore, supply uncertainty is certainly a relevant factor when it comes to estimating the profitability of a reverse logistics operation and taking an overall go/no-go decision. However, its impact is largely independent of the specific network design.

This observation also relates to the proficiency of the different modeling approaches. Specifically, it explains why in this context performance differences tend to be small between the solution of a scenario analysis and those of theoretically more powerful yet computationally more demanding methods, such as stochastic or robust models (see also Table 4.3). It is worth underlining the impact of the scenario selection, however. A scenario analysis may require a much finer gradation of scenarios than a stochastic or a robust model. For an illustration, consider the last three columns of Table 4.3. For a given network structure, the 'maximum regret' will, in general, be assumed for one of the extreme scenarios $\lambda = 0$ or $\lambda = 0.9$. For the robust model, a scenario space consisting of these two cases is therefore sufficient. For a conventional scenario analysis, however, this is not true. Choosing only between the optimal design for $\lambda = 0$ and $\lambda = 0.9$ respectively yields a 100% increase in the maximum regret. However, there does exist an intermediate scenario ($\lambda = 0.33$), whose corresponding optimal solution performs equally well across all scenarios as the optimal robust solution.

4.6.2 Compliance with Forward Networks

The robustness property analyzed in the previous subsection also plays an important role when it comes to the compliance of reverse logistics networks with 'forward' logistics infrastructure already in place. As discussed in Section 4.2, this is an important issue since companies, in many cases, do not set up reverse logistics networks from scratch but on top of an existing 'forward' network.

In this vein, the forward and reverse network parts have been optimized sequentially in the above examples. To assess the consequences of such a two-stage approach, let us compare its outcome with an integral design, which optimizes both network parts simultaneously. Figure 4.3(b) illustrates the optimal solution in this case for $\lambda = 0.4$. All parameters are kept equal to the values in Table 4.2. Comparing Figures 4.3(a) and (b), one observes that an integral design approach indeed leads to a significantly different network structure. However, the costs of both solutions are almost identical, namely €52.9m in the sequential approach versus €52.7m in the integral approach. This result generalizes to other values of λ. Specifically, the cost penalty for adding the reverse logistics network on top of a previously designed forward network rather than optimizing both parts together increases from 0% for $\lambda = 0$ to not more than 1.6% for $\lambda = 0.9$. We have observed similar results for many other parameter settings, including the case that demand and return volumes are not proportional (see Fleischmann et al., 2001).

One can explain this observation as follows. First, forward flows outweigh reverse goods flows, in general, in terms of volumes and costs. Therefore, the overall optimal solution can be expected to be 'close' to the optimal forward network. A deviation from this structure must allow for substantial savings per unit in the reverse channel in order to set off against the resulting increase in distribution costs. Second, the flat cost structure highlighted in the previous subsection results in a very limited cost penalty for deviating from the optimal reverse network structure due to constraints imposed by existing infrastructure. This is the more true if demand and returns have a similar geographical distribution.

This observation is again good news from a business perspective since it suggests that setting up an efficient reverse logistics network in many cases does not require a fundamental redesign of a company's existing logistics networks. In addition to limiting investment costs, this conclusion simplifies the

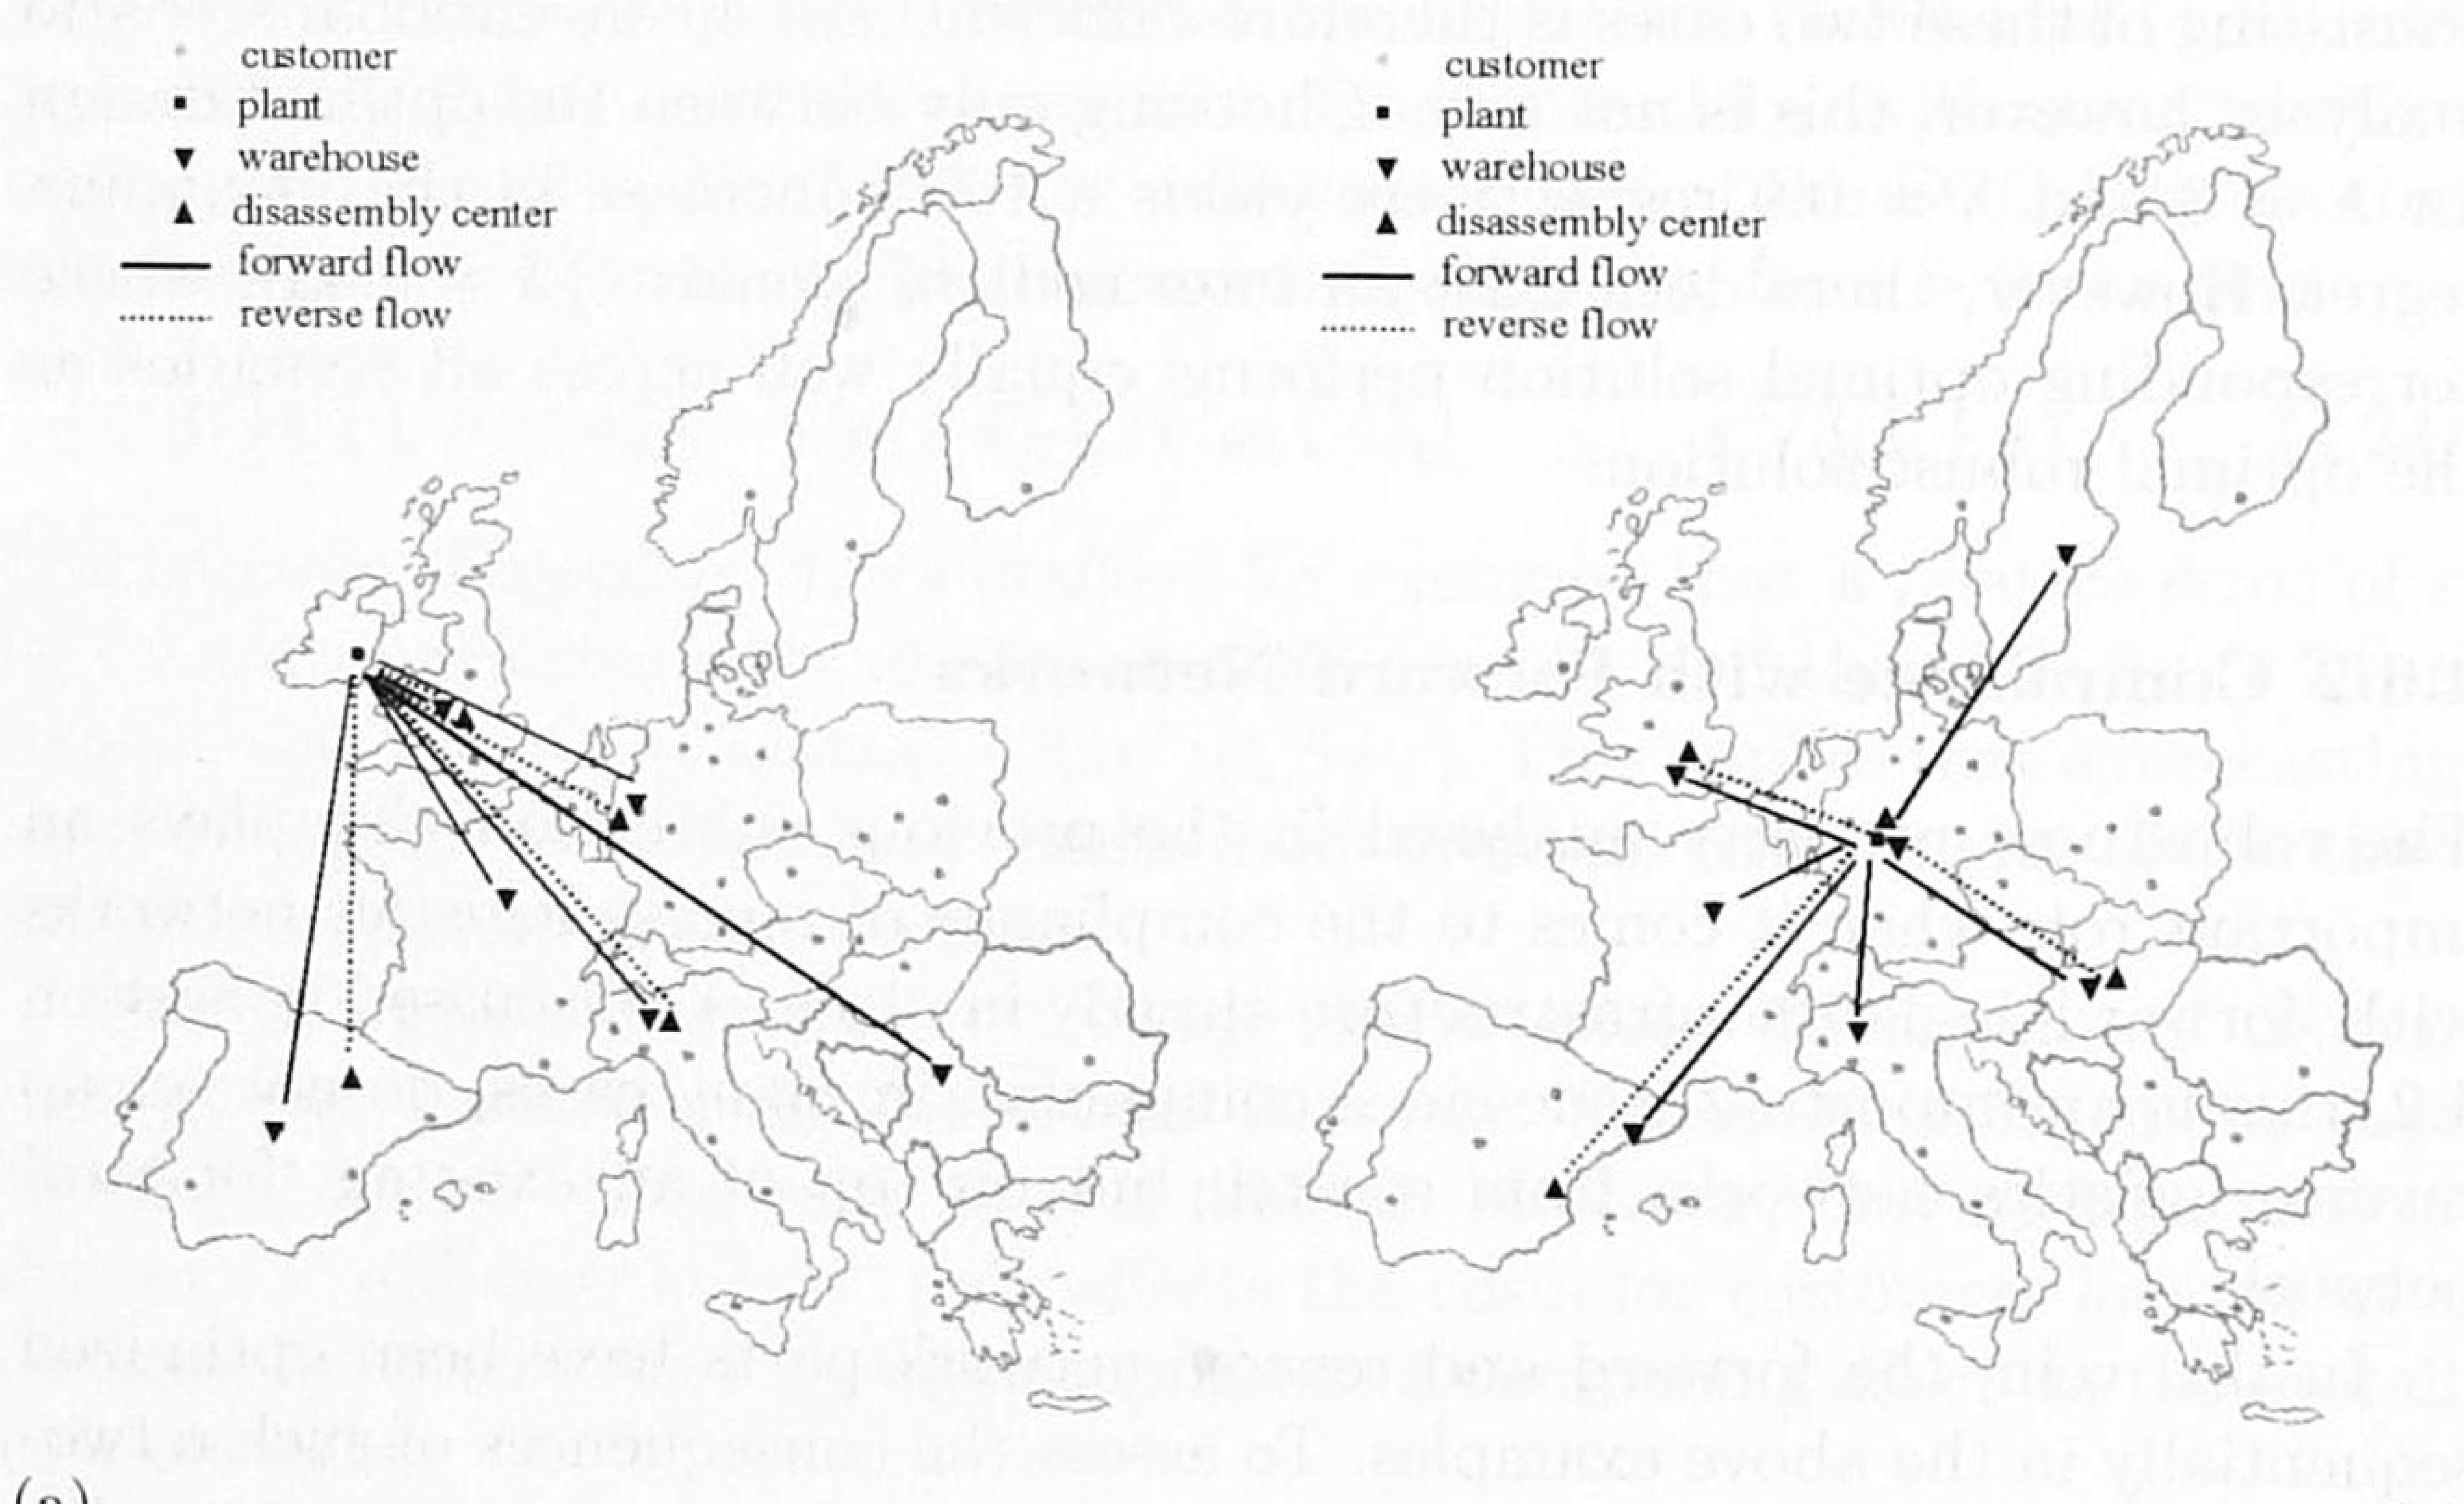

Fig. 4.5. Sequential Versus Integral Network Design With Regional Cost Differences

organizational implementation of reverse logistics initiatives. From a modeling perspective, this observation results in a significant reduction of complexity by optimizing forward and reverse network structures separately.

In Fleischmann et al. (2001), we have indicated the limits of the above observation by means of an example of a paper recycling network. In that case, the recycling operation does have a fundamental impact on the entire logistics network by reducing the impact of virgin raw material sources. While the structure of the original forward network is strongly dominated by pulp wood production close to the Scandinavian forests, recycling 'pulls' the business activities closer to the main markets in Western Europe. We illustrate the underlying economics in the context of our previous example.

To this end, consider the potential differences in labor and investment costs across Europe. To make things specific, assume that salaries and tax effects result in a manufacturing cost advantage of €2.- per unit in Ireland and of €1.50 in Eastern Europe, compared to the remaining countries. Moreover, assume that these cost differences are less prominent in the recovery channel due to a lower level of labor skills. For the sake of argument, let us assume that effective differences in remanufacturing costs across countries are negligible. In what follows, we set $c^p =$ €7.0m and $\gamma = 0.8$, while all other parameters remain unchanged with respect to the previous examples.

Figures 4.5 (a) and (b) depict the optimal network structures according to a sequential and an integral design, respectively ($\lambda = 0.4$). Moreover, Figure 4.6 shows the corresponding cost functions. Apparently, the cost advantage of an integral design approach is much more significant in this case than in the previous examples. This can be explained as follows. The structure of the

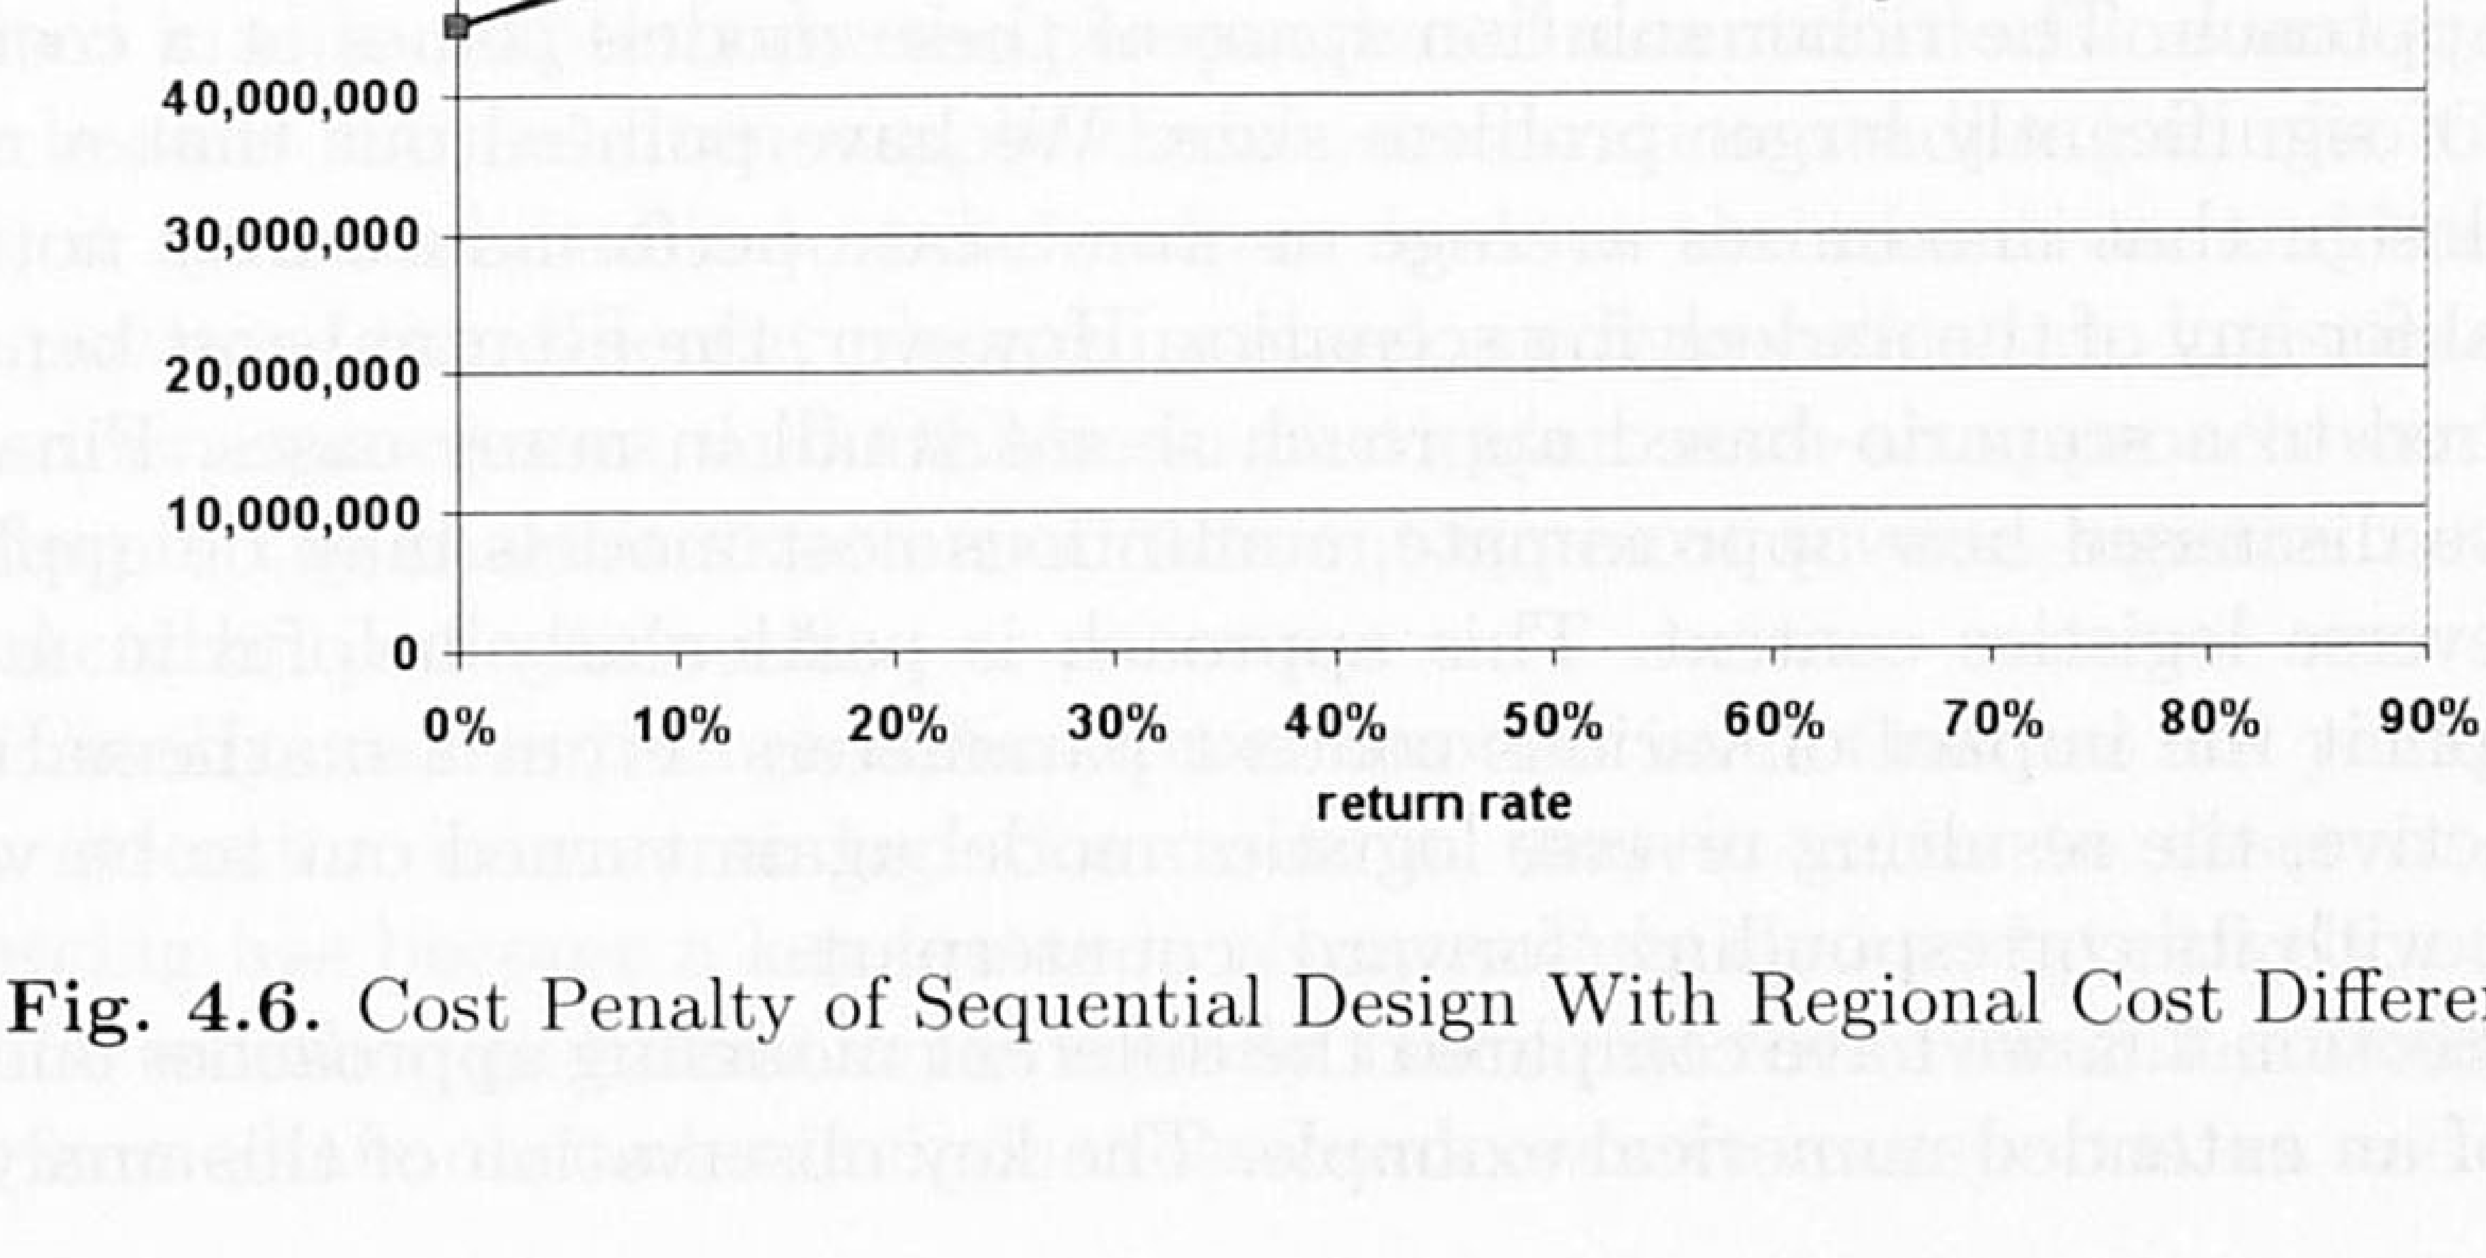

Fig. 4.6. Cost Penalty of Sequential Design With Regional Cost Differences

original forward network is dominated by the labor cost differences, which cause the production facility to be located in Ireland. However, when moving towards remanufacturing this factor is losing importance and the network structure is determined by transportation distances primarily. In other words, it is the *substitution* of virgin products by recovered ones that have structurally different cost drivers, which is key to the above result. It is under these conditions that one can expect reverse logistics to have a more fundamental impact on a company's overall logistics network.

4.7 Conclusions and Outlook

In this chapter, we have considered the design of appropriate infrastructure for companies engaged in reverse logistics programs. In Sections 4.1 and 4.2, we have argued that logistics network design is a key determinant of the overall profitability of closed–loop supply chains. By comparing this setting with more traditional production–distribution networks, we have distilled three important issues that appear to be characteristic of reverse logistics networks. First, the supply side is subject to significant uncertainty. Second, the need for testing and sorting used products before assigning them to an appropriate recovery option leads to a particular centralization–decentralization tradeoff. Third, reverse logistics requires the coordination and integration of different inbound and outbound flows.

The core part of this chapter, encompassing Sections 4.3 through 4.5, reviews quantitative models for supporting reverse logistics network design. Analogous with traditional network design problems, many of the approaches rely on MILP formulations. We pointed out that these models, in general, very much resemble conventional multi–level facility location models. Specific features that can be attributed to reverse logistics are few. They include a set of balance equations that link exogenous conditions on the supply and the demand side and an additional degree of freedom in optimizing the logistics network structure and the recovery policy simultaneously. A few models explicitly incorporate the aspect of uncertainty through a stochastic programming approach. The richer solution space of these models comes at a cost in terms of significantly larger problem sizes. We have pointed out that a network design that maximizes average or worst–case performance may not be optimal for any of the underlying scenarios. However, the eventual cost benefit compared to a scenario–based approach seems small in many cases. Finally, we have discussed how approximate, continuous cost models may be applied in a reverse logistics context. This approach is particularly helpful in making explicit the impact of various context parameters. From a mathematical perspective, the resulting reverse logistics model again turned out to be very similar with its corresponding 'forward' counterpart.

In Section 4.6, we have compared the different modeling approaches on the basis of an extended numerical example. The key observation of this analysis concerns the robustness of reverse logistics costs with respect to moderate changes in the network structure. This result, which concurs with what is known about conventional production–distribution networks, has important practical implications. On the one hand, it limits the impact of the aforementioned supply uncertainty when it comes to choosing an appropriate network structure. On the other hand, it indicates that in many cases reverse logistics networks are flexible enough to comply with existing network structures.

Given the short history of reverse logistics research, it goes without saying that many issues are yet to be explored. We conclude this chapter by highlighting a number of issues that we believe to be important for furthering the understanding of reverse logistics networks.

From a methodological perspective, the research focus to date has been on model formulations and on output analyses. In contrast, attention to algorithmic aspects has been limited so far. As consensus about the modeling foundations is growing, looking for efficient solution algorithms is gaining relevance. In particular, the question arises whether solution methods from traditional location theory are also adequate in a reverse logistics context, and which features may require modified approaches. Although the results in this chapter hint at a close similarity between reverse logistics network design and traditional location models, a thorough understanding of algorithmic implications is still lacking. Recent work by Verter et al. (2003) provides a promising starting point for future research in this direction.

Regarding the modeling choices, an important aspect that has, to date, been left aside in reverse logistics network design is the role of inventories. Note that none of the models reviewed throughout this chapter accounts for inventory effects, except possibly as part of some constant unit handling cost term. Yet, inventory is well known to be an important parameter in distribution network design. Risk-pooling and postponement are major factors in the centralization/decentralization trade-off. We see a clear need for corresponding analyses concerning the effects of inventory considerations on reverse logistics networks.

Another aspect that deserves extra attention in future research concerns the multi-agent character of reverse logistics networks. All of the models presented in this chapter take the perspective of a central decision-maker. Given the lessons learned from supply chain coordination studies (see also Chapter 12), it seems advisable to take a closer look at the incentives and the channel power of the different players involved, such as collectors, logistics service providers, processors, and OEMs. In particular, such an approach would help underpin the characterization of different types of reverse logistics networks, as sketched in Section 4.2.

Finally, we mention globalization as another issue that has not yet received the attention in reverse logistics research that it seems to deserve. Global sourcing has become a key factor in many supply chains. Intuitively, one may doubt whether globalization is equally beneficial for reverse chains. Some of the immediate obstacles include cross-border waste transportation and tax requirements. However, there may be even more fundamental arguments against global reverse logistics flows, which relate to the contribution of each player within the chain. A thorough analysis of these issues, which seems instrumental for a good understanding of the differences between forward and reverse logistics networks, again calls for a broadening of the modeling approaches available to date.