Supply Chain and Operations Decision Making

www.scodm.reapress.com

Sup. Chain Oper. Decis. Mak. Vol. 3, No. 3 (2026) 223–248.

Paper Type: Original Article

When does the Uncertainty Budget Change the Design? A Decision-Response Diagnostic for Robust Closed-Loop Supply Chain Network Design

Mahdi Razaghi[1,*], Mahtab Esmaeili[2], Seyed Esmaeil Najafi[3]

[1] Department of Industrial Engineering, Iran University of Science and Technology, Tehran, Iran; m_razaghi@ind.iust.ac.ir.
[2] Department of Industrial Engineering, Faculty of Engineering, Alzahra University, Tehran, Iran; mahtab80e@gmail.com.
[3] Department of Industrial Engineering, SR. C, Islamic Azad University, Tehran, Iran; najafi1515@yahoo.com.

Citation:

Received: 10 May 2026 *Revised: 16 July 2026* *Accepted: 12 August 2026*	Razaghi, M., Esmaeili, M., & Najafi, S. E. (2026). When does the uncertainty budget change the design? A decision-response diagnostic for robust closed-loop supply chain network design. *Supply Chain and Operations Decision Making*, 3(3), 223-248.

Abstract

Budget-of-uncertainty robust optimization is widely used for demand uncertainty in closed-loop supply chain network design, and the price of robustness is commonly reported as its summary statistic. This paper shows that the statistic conflates responses a decision maker would treat differently, and proposes reporting them separately through a four-layer decomposition: strategic response as Hamming distance between facility-opening vectors, service in aggregate and by region, operational response as normalized L1 flow change per arc family, and penalty alongside the residual change in other cost. The layers can decouple dramatically. In a Bavaria-informed instance whose returns are the observed 2024 collections of all 86 disposal-responsible bodies, cost rises by up to 27.9 percent while the facility configuration and aggregate delivered volume are unchanged across all eighteen exact solves and upstream flows are invariant to the tonne, the operational response confined to terminal arcs. Two secondary results follow. The aggregate-service formulation is a different service policy rather than a repaired counterpart: matched on the protection they buy, the two coincide at the tested endpoints but diverge at every tested interior level, where the aggregate plan delivers less to up to 10 of 12 regions. And perturbation tests attribute the strategic inactivity to an endogenous installed-capacity ceiling 2.1 percent below an immovable physical one. We derive a design-conditional upper service-saturation boundary, confirmed by bisection at 0.9791, which governs the service layer alone and does not determine strategic response. A two-stage adjustable formulation on an aggregated instance reproduces the first-stage invariance at the reference ratio.

Keywords: Closed-loop supply chain, Robust optimization, Budget of uncertainty, Price of robustness, Network design, Waste electrical and electronic equipment.

1 | Introduction

Closed-loop supply chains extend forward logistics with collection, inspection, recycling and redistribution so that end-of-life products re-enter the economy rather than landfill. Designing such networks is a strategic decision: facility locations and capacities are fixed years before the return, and demand quantities they must

Corresponding Author: m_razaghi@ind.iust.ac.ir

https://doi.org/10.48313/scodm.v3i3.67

handle are known. Uncertainty is therefore intrinsic to the problem, and the operations research literature has responded with a large body of stochastic, fuzzy, and robust formulations.

Among these, the budget-of-uncertainty approach of Bertsimas and Sim [1], [2] is widely used. Its appeal is easy to state. A box uncertainty set protects against every parameter deviating simultaneously and yields solutions that practitioners reject as too conservative. The budget parameter, conventionally written Γ, limits how many coefficients in a constraint may deviate at once, and thereby interpolates between the deterministic model and the fully conservative one. Reporting the resulting cost increase against Γ produces the price-of-robustness curve, which is commonly reported as evidence that a robust model is doing useful work.

This paper argues that the curve is weak evidence, because it aggregates responses that a decision maker would treat very differently. A robust model can respond to a larger budget by opening different facilities, by recovering and serving a larger volume, by re-routing flow through the existing network, or by paying shortfall penalties. These have entirely different managerial implications, and only the first commits capital. Yet all four produce the same monotone increasing cost curve, and a study that reports the curve alone leaves the reader unable to tell which occurred.

We therefore propose reporting a four-layer response decomposition, defined in full in Section 3.5. Relative to the zero-budget solution, the strategic response is the Hamming distance between facility-opening vectors. The service response is the change in delivered quantity, in total and by market region. The operational response is the relative L1 change on each arc family, with the number of arcs whose flow changes and the tonnage re-routed. The penalty response is the change in shortfall cost, reported alongside the residual change in all other cost terms. Every quantity comes from solutions that are already computed, so the decomposition costs only bookkeeping and can be applied retrospectively wherever solutions have been retained.

The decomposition matters because the layers can move independently, and in our case study they do.The setting is a Bavaria-informed recovery network for waste electrical and electronic equipment. It is semi-empirical: the return side is observed, while secondary-market demand, site coordinates and the cost parameters are constructed, as Section 4 sets out in full. The Bavarian State Office for the Environment publishes an annual waste balance in which each of the 86 disposal-responsible bodies reports collected quantities by material category, together with the number of civic amenity sites it operates for each appliance class. This yields an instance whose return side is measured rather than assumed: 105,738 tonnes of municipally collected equipment across 13.2 million inhabitants, with collection intensity varying by more than a factor of three between bodies.

The contributions of this study are as follows, in order of weight.

First, and primarily, we propose the four-layer decomposition as a reporting standard for robust network design. It rests on four statistics that any solver already yields: Hamming distance between facility-opening vectors, change in delivered quantity in aggregate and by region, normalized L1 flow change per arc family, and change in shortfall cost set against the residual change in other cost. The claim is not that a robust model must change a decision to be worth having, but that a study which reports only the aggregate cost curve leaves its reader unable to tell which of four quite different things occurred.

Second, and as the primary evidence for the first, we show on measured return data that the layers can decouple dramatically. A textbook-looking price-of-robustness curve, monotone and smooth and rising to 27.9 percent, corresponds to a network whose facility configuration and aggregate service do not change at all and whose upstream material flows are invariant to the tonne, with the entire response carried by terminal allocation and shortfall payment. No single statistic could have revealed this.

Third, as a secondary methodological contribution, we separate the cardinality structure of the budget from the semantics of the service policy. The aggregate-service formulation that removes the cardinality degeneracy identified for closed-loop planning [3], [4] does so by protecting system-wide rather than regional demand, and is therefore a different operational commitment rather than a repaired counterpart; we quantify how differently the two allocate delivery across regions at matched protection.

Fourth, also secondary, we isolate two mechanisms that prevent additional protection from increasing service or inducing strategic redesign. The second is an admissibility ceiling of the kind characterized in general terms by Dehghani Filabadi and Mahmoudzadeh [5], in which the budget is well posed but the protected demand cannot be served. Targeted perturbations of the shortfall penalty, facility fixed cost, echelon capacities, residual fraction, deviation level, and demand allocation separate an endogenous installed-capacity ceiling from an immovable physical one, and show that the two lie close together in the case instance.

Fifth, as validation rather than as a separate finding, a two-stage adjustable robust formulation with openings committed before the realization and flows in recourse, solved by column-and-constraint generation on an aggregated instance, reproduces the first-stage invariance at the reference ratio and so indicates that it is not an artefact of the static timing convention.

Sixth, we identify an upper service-saturation boundary below which additional protected demand can induce service and upstream-flow responses and above which these responses disappear in the baseline instance. We derive the boundary conditional on the installed design and confirm it numerically by bisection. We present it as an explanation of one layer rather than as a headline result: It governs aggregate-service response only, and neither determines nor bounds strategic response, since facility configurations can change through capacity expansion or spatial reconfiguration even when aggregate delivered volume is unchanged.

The remainder of the paper is organized as follows. Section 2 reviews robust closed-loop network design and the known limitations of the budget approach. Section 3 states the model, the two service policies, the decision timing, and the response decomposition. Section 4 describes the Bavarian data. Section 5 presents the computational study, including a two-stage adjustable formulation solved by column-and-constraint generation. Section 6 discusses implications and limitations, and Section 7 concludes.

2 | Literature Review

2.1 | Closed-Loop Network Design under Uncertainty

Closed-loop network design has been studied with every major uncertainty formalism. An early robust counterpart [6] set the template that much of the literature follows. A two-stage stochastic program covering demand, return rates, and returned-product quality [7] found that demand uncertainty had the strongest effect on the value of the stochastic solution, through its influence on facility location and capacity installation, and return-rate uncertainty the weakest. Later work adds a distributionally robust formulation [8], a data-driven robust model for a dual-channel network under carbon cap-and-trade [9] , and a three-stage hybrid combining demand and return scenarios with an uncertainty set for the carbon tax rate [10].

Recent work estimates uncertainty sets from data, with balanced sets whose boundaries sit equidistant from uncovered observations [11]. That addresses calibration of the set, not whether the resulting protection changes any decision.

The SCODM literature shows the same breadth, and is the audience for the diagnostic proposed here: closed-loop spare parts with machine-learning return forecasting [12], opportunistic maintenance in a closed loop [13], robust empty-container positioning [14], bi-objective robust blood supply design [15], food-security chains solved by a genetic algorithm [16], fourth-party logistics in off-site construction [17], [18], and energy-integrated sustainable networks [19]. Treatments of uncertainty in the same venue vary widely: antifragility rather than worst-case cost [20], fuzzy multi-criteria supplier prioritization [21], combined knowledge- and data-driven resilient sourcing [22], and probabilistic arc availability in maximum flow [23]. Each formalism reports differently, and none, including the robust studies above, reports whether the protection purchased changes the decision taken. That is the reporting gap this paper addresses.

2.2 | The Budget of Uncertainty and Its Known Limitations

Bertsimas and Sim [1], [2] introduced the budget of uncertainty to control conservatism. For a constraint indexed by i with uncertain coefficient set Ji, the budget Γi ranges over the interval from zero to the cardinality of Ji, and limits the number of coefficients that may deviate simultaneously. The robust counterpart remains linear, which is why the approach has been adopted so widely.

Two limitations are established. Kim et al. [3] observed that when the cardinality of J subscript i equals one, the robust counterpart with a budget is equivalent to the counterpart with a box uncertainty set, so the budget is ineffective and amounts only to a reduction of the box. They proposed a reformulation in which decision variables express the proportion of available supply used rather than absolute quantities, which places all uncertain parameters in a common expression and restores a meaningful budget. Kim and Lee [4] extended this to collected-material quality uncertainty and reported comparative conditions under which quantity-based and proportion-based formulations are preferable. Notably, in both studies the budget is applied to the reverse-logistics side, and demand uncertainty is handled with a box set.

Separately, Dehghani Filabadi and Mahmoudzadeh [5] studied problems with resource uncertainty on the right-hand side and showed that the conventional budget may remain over-protective: beyond a threshold, increasing the budget does not change the optimal solution, so the additional budget is ineffective. They proposed a two-stage method that identifies the admissible portion of the uncertainty set and protects only against an effective worst case, demonstrating it on power dispatch and calling for extension to other application domains. Related analyses of right-hand-side uncertainty appear in Minoux [24] and Ouorou [25], and two-stage robust network flow and design under demand uncertainty is treated by Atamturk and Zhang [26].

2.3 | Decision Timing and Two-Stage Robust Network Design

A third literature concern concerns when decisions are taken. Ben-Tal et al. [27] introduced the adjustable robust counterpart, separating variables that must be fixed before the uncertainty is realized from those that may adapt afterwards, and showed that the adjustable counterpart is markedly less conservative than the static one, though generally intractable unless the adaptive variables are restricted to affine functions of the data. Atamturk and Zhang [26] developed two-stage robust formulations for network flow and design under demand uncertainty, and Bertsimas et al. [28] studied robust and adaptive network flows. Zeng and Zhao [29] gave the column-and-constraint generation scheme that makes two-stage robust integer problems solvable in practice, which is the method used in Section 5.4.

The distinction matters for the question posed here, because a static formulation chooses openings and flows together against one protected demand vector, whereas a two-stage formulation commits the openings first and lets the flows adapt. The two can differ: A first stage that must hedge across a whole uncertainty set may install capacity that a static model, optimizing against a single worst case, does not. Recent closed-loop work reflects this awareness, with two-stage distributionally robust designs [8] and multi-stage hybrid formulations combining scenarios for demands and returns with uncertainty sets for policy parameters [10]. Because our central claim concerns whether facility decisions respond to the budget, we report both formulations rather than one.

2.4 | Research Gap

Both limitations reviewed above are established, so the gap must be stated narrowly. Solution invariance under increasing protection is not new: it is the finding of [5], and the intellectual basis for the admissibility mechanism examined here. Nor is the cardinality problem, which [3] and [4] identify and reformulate away.

What remains insufficiently examined is how these two mechanisms interact with the decision structure of closed-loop network design. Research [5] is demonstrated on economic dispatch, an operating problem in which no facility is sited; Researches [3] and [4] are production and sourcing plans in which the discrete decision is collector selection rather than network configuration. Network design differs in that its decisions are heterogeneous and are committed at different times: strategic facility openings, operational flows across six arc families, regional as against aggregate service, and shortfall penalties. Whether a budget that leaves aggregate cost rising but induces no additional service is equally inert in each of those layers, whether addressing the cardinality problem restores a strategic response, and where the boundary between the responsive and unresponsive regimes lies, are open questions. Reviews of the field [30], [31] catalog the uncertainty types modeled but do not report the effectiveness of the protection applied.

This study addresses that gap. It examines both mechanisms in a single network design model on measured data, decomposes the response by layer, distinguishes the service policies that the two mechanisms implicate, and derives the condition separating the regimes, using the numerical experiment as confirmation.

3 | Problem Description and Model

3.1 | Network Structure

The recovery network has six echelons. Collected equipment originates at return sources K, each a disposal-responsible body with a reported annual quantity, and flows to consolidation hubs I, where it splits into a recoverable fraction routed through recycling J and redistribution M to secondary-market regions L, and a residual fraction routed to disposal N. A forward flow from depot P to the sources represents the equipment placed on the market that later returns.

Hubs, recycling facilities, and redistribution centers are siting decisions; disposal sites and market regions are fixed. *Table 1* lists the notation.

Table 1. Model notation.

Symbol	Description
K, I, J, M, N, L	Index sets: return sources, consolidation hubs, recycling facilities, redistribution centers, disposal sites, secondary-market regions
r_k	Annual quantity returned at source k (tonnes)
d_l	Nominal secondary-market demand for recovered material in region l
d_l	Maximum deviation of d_l from its nominal value
s	Technically non-recoverable fraction of collected material
cc_i, cr_j, ce_m, cd_n	Capacities of hubs, recycling, redistribution and disposal
f_i, g_j, h_m	Annual fixed cost of opening a hub, recycling facility, redistribution center
c^{PK}_{pk}	Unit transport cost on the depot-to-return-source arc (p, k)
c^{KI}_{ki}	Unit transport cost on the return-source-to-hub arc (k, i)
c^{IJ}_{ij}	Unit transport cost on the hub-to-recycling arc (i, j)
c^{JM}_{jm}	Unit transport cost on the recycling-to-redistribution arc (j, m)
c^{ML}_{ml}	Unit transport cost on the redistribution-to-market arc (m, l)
c^{IN}_{in}	Unit transport cost on the hub-to-disposal arc (i, n)
C_p, R_p, cs_p	Fixed cost, opening variable and capacity of depot p
π_l	Penalty per tonne of service slack in region l
X_{ki}, U_{ij}, P_{jm}, Q_{ml}, T_{in}, V_{pk}	Flow on the K-I, I-J, J-M, M-L, I-N and P-K arcs (tonnes)
δ_l	Service slack: penalized unmet demand in region l (tonnes)
Y_i, Z_j, W_m	Binary variables for opening a hub, recycling facility, redistribution center
Γ	Budget of uncertainty
ϱ	Recovery-demand ratio: total secondary-market demand divided by recoverable supply
$\varrho^*(Y,Z,W)$	Service-saturation threshold of a given installed design

3.2 | Deterministic Formulation

The model is stated in full so that it can be reconstructed exactly. The objective minimizes annual fixed cost, transport cost on all six arc families, and shortfall penalty:

$$\min Z = \sum_{i\in I} f_i Y_i + \sum_{j\in J} g_j Z_j + \sum_{m\in M} h_m W_m + \sum_{p\in P} C_p R_p + \sum_{p\in P}\sum_{k\in K} c^{PK}_{pk} V_{pk} + \sum_{k\in K}\sum_{i\in I} c^{KI}_{ki} X_{ki} + \sum_{i\in I}\sum_{j\in J} c^{IJ}_{ij} U_{ij} + \sum_{j\in J}\sum_{m\in M} c^{JM}_{jm} P_{jm} + \sum_{m\in M}\sum_{l\in L} c^{ML}_{ml} Q_{ml} + \sum_{i\in I}\sum_{n\in N} c^{IN}_{in} T_{in} + \sum_{l\in L} \pi_l \delta_l. \quad (1)$$

subject to *Eqs. (2)-(15)*. The service requirement is soft. The variable δl is a service slack: realized demand must be covered either through delivered recovered material or through explicitly penalized shortfall, at πl per tonne. The model therefore never becomes infeasible for want of recovered material, and every constraint below should be read as an accounting requirement on demand rather than a guarantee of physical satisfaction. This distinction is central to what follows, because the penalty layer of the decomposition measures exactly the demand that is accounted for but not served:

$$\sum_{m\in M} Q_{ml} + \delta_l \ge d_l, \quad \forall\, l \in L. \quad (2)$$

Delivery is bounded above by what each market can absorb. A region cannot take more than the greatest demand its uncertainty interval admits, so

$$\sum_{m\in M} Q_{ml} \le d_l + \hat{d}\, l, \quad \forall\, l \in L. \quad (3)$$

Without *Eq. (3)*, the service rows are pure lower bounds, and because *Eqs. (8)* and *(9)* are equalities, recovered material that must leave the recycling echelon can be pushed onto markets beyond any demand they could realize. Surplus recovered material is instead routed to disposal at the hub, which the yield ceiling (6) permits and which the *Disposal Capacity (13)* must accommodate; in the case instance, disposal runs at 84,167 tonnes of 100,000 at the lowest ratio tested.

Constraint (3) is a market-absorption bound rather than a protection requirement, and the distinction matters for reading the results. The uncertainty budget determines the demand level against which shortfall is penalized, whereas dl + d̂l represents the largest volume that market l is assumed able to absorb physically. Delivered volume may therefore exceed the protected requirement at low budget levels, when placing recoverable material is cheaper than disposing of it. Still it can never exceed the market's assumed absorption envelope. The case instance shows this directly: at $\varrho = 0.2$ and $\Gamma = 0$, the model delivers 18,572 tonnes against a nominal demand of 17,975 tonnes, because the cheapest route through recycling and redistribution costs 6.23 euros per tonne against 7.20 euros per tonne for the cheapest disposal route, while disposal capacity is not binding at 87,166 tonnes of 100,000. Accordingly, we refer to Q throughout as delivered recovered volume rather than as served demand, and the service layer of the decomposition measures changes in that delivered volume.

The model requires complete processing of the municipal return quantities observed in the case dataset:

$$\sum_{i\in I} X_{ki} = r_k, \quad \forall\, k \in K. \quad (4)$$

Constraint (4) is a modeling choice about the scope of the instance, not a derived legal requirement. Article 3 of the Bavarian Waste Act designates the Landkreise and kreisfreie Städte as public-law waste-disposal authorities and assigns them disposal responsibility, and the reported quantities are those they collected; that regulatory framing motivates treating the observed tonnage as a quantity the system must handle rather than a decision variable. It does not establish that every tonne must traverse the particular echelon structure modeled here, and we do not claim it does.

Material is conserved at each consolidation hub, splitting between the recycling and disposal streams:

$$\sum_{j\in J} U_{ij} + \sum_{n\in N} T_{in} = \sum_{k\in K} X_{ki}, \quad \forall\, i \in I. \tag{5}$$

The recovery yield is bounded by the technically recoverable fraction. At most $(1 - s)$ of the material entering a hub can be routed to recycling, where s is the physically determined non-recoverable share:

$$\sum_{j\in J} U_{ij} \leq (1 - s) \sum_{k\in K} X_{ki}, \quad \forall\, i \in I. \tag{6}$$

Constraints (5) and *(6)* imply disposal of at least s times the collected quantity, and permit more. This is deliberate: material that is technically recoverable but that recycling capacity cannot absorb is landfilled rather than left unaccounted for, whereas fixing disposal at exactly sR would make the model infeasible whenever recycling capacity binds. The relaxation is not vacuous here, and its magnitude should be reported: 17,738 tonnes are disposed against a technical minimum of 15,861, an excess of 1,877 tonnes or 1.78 percent of collections, at 2 of the 12 hubs and reaching 990 tonnes at the largest. These are identical at every budget level under both policies, so the excess reflects a structural shortfall in recycling capacity, not a response to protection.

Returned material cannot exceed what was previously shipped to the market, which links the forward depot flow to the reverse network:

$$\sum_{p\in P} V_{pk} \geq \sum_{i\in I} X_{ki}, \quad \forall\, k \in K. \tag{7}$$

Without *Eq. (7)*, the forward variables would be free to take the value zero at no cost, and the depot echelon would be disconnected from the rest of the model.

Material is conserved at recycling facilities and at redistribution centers. Recovery losses are modeled once, upstream at the hub through the *Yield Ceiling (6)*; no second loss or inventory process is introduced downstream, so both balances are equalities and material cannot leave the network at a recycler:

$$\sum_{m\in M} P_{jm} = \sum_{i\in I} U_{ij}, \quad \forall\, j \in J. \tag{8}$$

$$\sum_{j\in J} P_{jm} = \sum_{l\in L} Q_{ml}, \quad \forall\, m \in M. \tag{9}$$

Throughput at each echelon is bounded by installed capacity, which is available only at opened sites:

$$\sum_{k\in K} X_{ki} \leq cc_i Y_i, \quad \forall\, i \in I. \tag{10}$$

$$\sum_{i\in I} U_{ij} \leq cr_j Z_j, \quad \forall\, j \in J. \tag{11}$$

$$\sum_{j\in J} P_{jm} \leq ce_m W_m, \quad \forall\, m \in M. \tag{12}$$

$$\sum_{i\in I} T_{in} \leq cd_n, \quad \forall\, n \in N. \tag{13}$$

$$\sum_{k\in K} V_{pk} \leq cs_p R_p, \quad \forall\, p \in P. \tag{14}$$

Disposal sites in *Eq. (13)* carry no opening variable because they are existing permitted facilities in the case setting. Finally, the variable domains are:

$$X_{ki}, U_{ij}, V_{pk}, P_{jm}, Q_{ml}, T_{in}, \delta_l \geq 0, \quad Y_I, Z_j, W_m, R_p \in \{0, 1\}. \tag{15}$$

In the case study, a single depot with zero fixed cost is used, so R is fixed at one and *Eq. (14)* reduces to a capacity bound on the forward flow; the general form is retained above because it is the form implemented in the code described under Data Availability.

3.3 | Two Service Policies under Demand Uncertainty

Uncertainty enters through secondary-market demand, with the demand in region l lying in the interval of half-width $\hat{d}_l$ around its nominal value d_l. Writing the realized demand as $d_l + \xi_l \hat{d}_l$ with $\xi_l \in [0,1]$, the budgeted uncertainty set is the set of ξ vectors whose entries sum to at most Γ.

We now formulate two models. They are not two counterparts of the same problem. They express different service policies, they carry different feasibility guarantees, and Section 3.3.3 shows that neither is obtained from the other by reformulation. We set them out separately for that reason.

3.3.1 | Regional-service formulation

Under a regional-service policy, the requirement is imposed region by region: For every realization in the set, each region's demand must be accounted for by delivery or by penalized shortfall. This gives the semi-infinite version of *Eq. (2)* and the tractable counterpart:

$$\sum_{M \in M} Q_{ml} + \delta_l \geq d_l + \Gamma_l z_l + p_l, \quad z_l + p_l \geq \hat{d}\,l, \quad \forall\, l \in L. \tag{16}$$

This is the conventional way of applying a budget to a demand-satisfaction constraint, and it inherits a structural limitation. Row l of *Eq. (2)* contains exactly one uncertain coefficient, so the cardinality of the uncertain set for that row is one, Γl is confined to the unit interval, and the protection reduces to min(Γ, 1) multiplied by $\hat{d}_l$. The budget cannot express the intended statement that only some regions spike together, because each row sees only its own region. It is a scalar shrinkage of the box uncertainty set. This is precisely the degeneracy that Kim et al. [3] identify for closed-loop planning, where they show that a budgeted counterpart collapses to a box counterpart when a single uncertain parameter appears in each relevant constraint.

3.3.2 | Aggregate-service formulation

Under an aggregate-service policy, the requirement is imposed on the system as a whole: Total realized demand must be accounted for when at most Γ regions spike together, while every region retains its nominal requirement. The regional rows of *Eq. (2)* are kept at their nominal level, and a robust system-wide row is added:

$$\sum_{l \in L} \sum_{m \in M} Q_{ml} + \sum_{l \in L} \delta_l \geq \sum_{l \in L} d_l + \Gamma z + \sum_{l \in L} p_l, \quad z + p_l \geq \hat{d}\,l, \quad \forall\, l \in L. \tag{17}$$

Here the uncertain coefficients share a single constraint, so Γ ranges over the interval from zero to the number of market regions and carries its intended combinatorial meaning. What the model requires, however, has changed. It requires that aggregate delivered volume plus aggregate shortfall account for the worst-case total demand in the budgeted set. It does not require that any particular region's high realization be accounted for individually, and under neither policy does it require that protected demand be physically served.

This is a substantive modeling commitment, not a technical device, and it is admissible only where recovered material is fungible across destinations, meaning that a tonne placed with one buyer discharges the operator's obligation as well as a tonne placed with another. We adopt fungibility as an explicit assumption for the case instance and do not claim it as a fact. It is a plausible one for this waste stream, whose recovery output is secondary raw material in the form of ferrous and non-ferrous metals, plastic fractions and glass cullet rather than a differentiated finished good, but our data cannot support it: the waste balance records what each body collects and does not describe how recovered fractions are subsequently marketed, to whom, or under what contractual terms. Establishing fungibility would require offtake agreements or transaction records that are not public, and we did not have access to them.

The assumption is falsifiable, and its failure has a clear consequence. Where an operator holds offtake contracts carrying regional service-level penalties, faces a regulatory duty to supply specific territories, or produces a differentiated remanufactured good for which buyers are not substitutable, demand is not fungible, and the aggregate policy does not represent the operator's obligations. In that case, the regional formulation of Section 3.3.1 is the correct model and the modeller must accept its degenerate budget rather than replace it. Section 3.3.3 quantifies what is given up by assuming otherwise, and Section 6.3 records the assumption as a limitation.

3.3.3 | Relationship between the two policies

The two models are related, but neither reduces to the other. Summing the regional requirement over l yields the aggregate requirement, and $d_l + \hat{d}_l \geq d_l$ gives the retained nominal rows, so any solution feasible for the regional formulation at full protection is feasible for the aggregate formulation at full protection. Because both admit the service slack, feasibility is never the issue; the comparison is between the cost each policy incurs and the shortfall each permits. The converse fails. The regional feasible set is therefore contained in the aggregate feasible set, and regional protection is the more conservative policy.

The containment is strict, but where the two policies separate is not obvious, and comparing them at their own maximum budgets is misleading. At full protection they impose the same aggregate upward deviation, 17,975.46 tonnes, because minimizing $\Gamma z + \Sigma p_l$ at Γ equal to the number of regions is attained at $z = 0$; total protected demand is therefore 107,852.76 tonnes under both, against a nominal 89,877.30 tonnes. At that point the two plans coincide exactly, at €13,874,776 each. They also coincide trivially at zero protection. The difference lives in between.

To expose it, we match the policies on the aggregate protection increment they purchase, since their budget parameters are not comparable. The regional policy protects each region by $\Gamma\hat{d}_l$, so at a common $\Gamma \in [0, 1]$ it purchases $\Gamma\Sigma_l\hat{d}_l$. The aggregate policy purchases the protection term at its optimum, min over $z \geq 0$ of $\Gamma z + \Sigma_l \max(\hat{d}_l - z, 0)$, attained at $z = 0$ or at one of the deviations and so found by a scan. For each target increment, we solve the regional model at Γ equal to the increment divided by $\Sigma_l\hat{d}_l$, and the aggregate model at the smallest budget reaching the same increment, located by bisection. *Table 2* reports the comparison at the reference ratio and at $\varrho = 0.8$.

Any study adopting a pooled robust service constraint in place of regional ones should therefore state the fungibility assumption explicitly and report, as above, how the resulting plan reallocates delivery.

Each row solves both policies at the budget that purchases the stated aggregate protection increment, $\Gamma\Sigma_l\hat{d}_l$ under the regional policy and the optimal value of the protection term under the aggregate policy; the two agree to within a tonne at every level. Under-served regions are those the aggregate plan delivers less to than the regional plan does.

Table 2. The two service policies compared at matched protection.

ϱ	Protection (t)	Γ Regional	Γ Aggregate	Cost Regional (€)	Cost Aggregate (€)	Gap (%)	Regions under-Served	Max Deficit (t)
1.0	0	0.000	0.000	10,849,168	10,849,168	0.000	0	0
1.0	2,247	0.125	1.079	11,223,682	11,197,442	+0.234	10	263
1.0	4,494	0.250	2.329	11,598,199	11,552,176	+0.398	9	526
1.0	6,741	0.375	3.677	11,973,831	11,921,106	+0.442	7	1,386
1.0	8,988	0.500	5.087	12,352,841	12,298,156	+0.445	5	1,849
1.0	11,235	0.625	6.529	12,733,211	12,678,578	+0.431	4	2,311
1.0	13,482	0.750	8.089	13,113,733	13,069,481	+0.339	2	2,773
1.0	15,729	0.875	9.919	13,494,254	13,470,328	+0.178	1	1,785
1.0	17,975	1.000	12.000	13,874,776	13,874,776	0.000	0	0
0.8	1,798	0.125	1.079	9,844,162	9,821,798	+0.228	10	210
0.8	7,190	0.500	5.087	10,039,829	9,997,648	+0.422	6	841
0.8	14,380	1.000	12.000	10,432,271	10,432,271	0.000	0	0

3.4 | Decision Timing

The motivation for network design under uncertainty is that facility decisions are committed before the quantities they must handle are observed, so the timing of decisions has to be stated rather than left implicit. We use two formulations, and we are explicit about the role of each.

The static formulation of Sections 3.2 and 3.3 determines openings and flows simultaneously against a single protected demand vector. It is not a claim about how decisions are actually taken. It is common in the applied robust closed-loop literature, and we retain it as the diagnostic benchmark precisely because it is what a reader of that literature will recognize. Claims derived from it are stated as properties of that formulation.

The two-stage adjustable formulation assigns the facility-opening variables Y, Z, and W to the first stage, taken before the demand realization is known, and all flow and shortfall variables to recourse, taken after it. Writing UΓ for the budgeted uncertainty set, the problem is

$$\min_{Y,Z,W} \{ \textstyle\sum_{i\in I} f_i Y_i + \sum_{j\in J} g_j Z_j + \sum_{m\in M} h_m W_m + \max_{d\in U_\Gamma} \min_{x\in X(Y,Z,W,d)} c^T x \}. \quad (18)$$

Second-stage feasibility has two separate sources. Because the service slack absorbs unserved demand, demand realizations cannot create second-stage infeasibility. Only the *Service Row (2)* sees the realization, becoming $\Sigma_m Q_{ml} + \delta_l \geq d_l + \xi_l \bar{d}_l$; the *Absorption Bound (3)* enters the recourse problem unchanged as $\Sigma_m Q_{ml} \leq d_l + \bar{d}_l$, because it is the physical envelope of what market l can take rather than what it happens to demand. Every remaining constraint is independent of ξ, so recourse feasibility is a property of the first-stage design alone. Feasibility of the fixed return stream is separate, since *Eq. (4)* requires the observed tonnage to clear the network whatever the demand. We impose it on the first stage explicitly. Writing R for total returns, the design must satisfy.

$$\textstyle\sum_{i\in I} cc_i Y_i \geq R, \quad \min\{ \sum_{j\in J} cr_j Z_j , \sum_{m\in M} ce_m W_m \} \geq \max\{0, R - \sum_{n\in N} cd_n\}. \quad (19)$$

The first inequality opens enough collection capacity to receive the returns; the second and third open enough recycling and redistribution capacity to absorb what disposal cannot take, since the *Yield Ceiling (6)* and *Conservation (5)* send the remainder to disposal, which is capacity-limited. With the data condition that total disposal capacity is at least sR, which holds here, these are sufficient for the return stream to clear, and with the first observation above they are sufficient for relatively complete recourse in the network structure considered. We claim no more: A network with additional echelons, return-side uncertainty, or constraints coupling realizations across regions would need its own argument. In the reported runs they never bind, and imposing them changes no result; we state them so the guarantee is constructive rather than incidental.

We solve the two-stage problem by column-and-constraint generation [29], an exact algorithm for two-stage robust programs. The master carries one replicated block of recourse variables per generated scenario and yields a lower bound; the separation problem returns the realization in $U\Gamma$ maximizing recourse cost for the incumbent first-stage decision, and the first-stage cost plus that worst case gives an upper bound. The algorithm stops when the relative gap $(UB - LB) / \max(1, |UB|)$ falls to 10^{-6} or below. Because the master's incumbent supplies the lower bound, the master is solved at a relative MIP tolerance of 10^{-9}, so its incumbent does not sit measurably above its own optimum and produce a spurious negative gap. The recourse value is convex in the right-hand side, so the inner maximum is attained at a vertex of $U\Gamma$; for integral Γ the vertices are the zero-one vectors with at most Γ ones, and we enumerate them exactly, which avoids a dualization step and is tractable at this size.

Section 5.4 reports the two-stage results. Their purpose is to test whether the invariance observed in the static formulation is an artefact of conflating the two stages.

3.5 | A Four-Layer Response Decomposition

Let x_0 denote the optimal solution at $\Gamma = 0$ and $x\Gamma$ the optimal solution at budget level Γ. We characterize the response of the model to the budget along four layers, each of which a decision maker would act on differently.

The strategic response is the Hamming distance between the facility-opening vectors, that is the number of positions in which the vector (Y, Z, W) differs between the two solutions. This is the layer that commits capital.

The service response is the change in recovered volume delivered to the secondary market, reported both in total and as the number of market regions whose delivered quantity changes. Aggregate and regional service can move independently: A fixed total may be re-allocated across regions.

The operational response is the change in the flow variables. For each arc family we report the relative L1 change, the number of arcs whose flow changes by more than a tolerance, and the tonnage re-routed, defined as one half of the L1 change because a tonne that leaves one arc joins another.

The penalty response is the change in shortfall cost, reported together with the residual change in all non-penalty cost terms, so that the two components of the objective change are visible separately.

For the operational layer, we use the normalized statistic:

$$\frac{\|x_\Gamma - x_0\|1}{\max(1,\|x_0\|1)}. \tag{20}$$

Computed per arc family and in aggregate. A budget is termed strategically inactive on an instance when the strategic response is zero across the admissible range of Γ, and upstream-inactive when in addition the flow change on every arc family other than the terminal one is zero. These are weaker and more precise statements than decision invariance, which our results do not support and which we are careful not to claim.

The purpose of the decomposition is to distinguish budgets that induce decision adaptation from budgets whose additional cost is primarily absorbed through shortage penalties. We do not call the latter ineffective. A budget that leaves the design unchanged while quantifying worst-case shortage exposure is doing something a planner may legitimately want, namely pricing insufficient capacity; what it is not doing is changing what gets built or how material moves, and the decomposition exists so that a reader can tell the two apart rather than infer either from a cost curve. It requires no computation beyond recording the solution vectors, which any solver returns, and we recommend it be reported alongside, or in place of, the price-of-robustness curve.

4 | A Bavaria-Informed Case Instance

The Bavarian State Office for the Environment publishes an annual waste balance covering all disposal-responsible bodies in the state. Under Article 3 of the Bavarian Waste Act, disposal responsibility rests with the Landkreise and kreisfreie Städte, several of which act jointly through special-purpose associations. For the 2024 balance, this yields 86 reporting bodies: 59 Landkreise, 21 kreisfreie Städte and 6 associations.

We extracted, for every body, the reported area, population, annual quantity of collected electrical and electronic equipment, and the number of collection sites operated for each of six appliance classes. Extraction succeeded for all 86 bodies on all fields. Aggregate figures provide external validation: the parsed population of 13,202,127 matches the state total, and the parsed 105,738 tonnes correspond to 8.01 kilograms per inhabitant, which is consistent with a municipal collection channel that excludes retail and commercial take-back.

Table 3 summarizes the data, and *Fig. 1* shows the distributions. Collection intensity varies from 4.53 to 15.29 kilograms per inhabitant, with a coefficient of variation of 0.226, and the lowest value belongs to the state capital. We report the variation without attributing it: the balance records quantities collected through municipal channels and does not identify the behavioral or structural reasons they differ, so any explanation would be conjecture. This heterogeneity is the reason for using measured rather than generated return data. Return volumes span more than a factor of twenty-five and collection intensity more than a factor of three, and the two are only weakly related, so the spatial structure is something a synthetic generator would have to be given rather than something it would produce.

Table 3. Bavarian 2024 waste balance, aggregated by reporting body type.

Body Type	Count	Population	WEEE Collected (t)	Mean Kg Per Capita
Landkreis	59	7,930,913	67,189	8.57
Kreisfreie Stadt	21	3,699,883	23,526	8.39
Special-purpose association	6	1,571,331	15,023	9.45
Total	86	13,202,127	105,738	8.01 (aggregate)

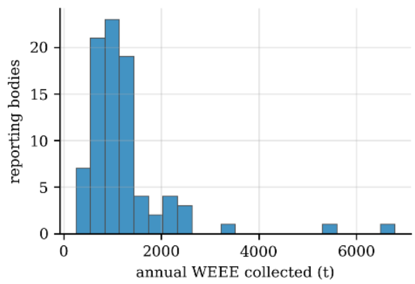

a.

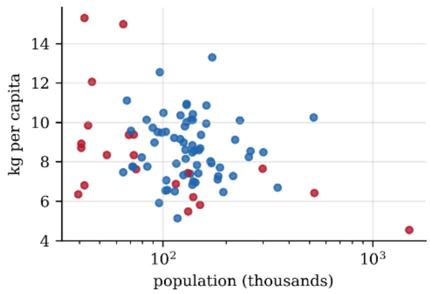

b.

Fig. 1. Distribution of return volumes and collection intensity across the 86 Bavarian reporting bodies; a. Return volumes, and b. Collection intensity.

4.1 | Observed Inputs

Two elements of the instance are taken directly from the waste balance. The return quantity r_k at each of the 86 sources is the reported annual tonnage of collected electrical and electronic equipment for that body, and the count of civic amenity sites it operates is used to select and size candidate consolidation hubs. Nothing else in the instance is observed.

4.2 | Constructed Inputs and the Generation Procedure

The remaining elements are constructed. We state the full procedure and every value, so that the instance can be regenerated exactly; our implementation performs these steps in the order given. All pseudo-random draws use a single NumPy PCG64 generator seeded with 7, and the draws are taken in the sequence listed, since the sequence determines the realization.

Coordinates. The balance reports site counts but not locations, so each body is assigned a coordinate drawn independently and uniformly from the Bavarian bounding box, latitude in [47.3, 50.5] and longitude in [9.0, 13.8]. Distances are great-circle distances multiplied by a road-detour factor of 1.25, and are therefore indicative rather than routed.

Candidate sites. The candidate hubs are the bodies reporting the most civic amenity sites, each inheriting that body's coordinate. Recycling facilities, redistribution centers, disposal sites, and market regions take coordinates drawn without replacement from the 86, in that order.

Market-demand allocation. Total demand is ϱ times recoverable supply, $(1 - s)$ times total returns, distributed across regions by weights drawn uniformly from [0.6, 1.4] and normalized to sum to one, giving a largest-to-smallest ratio of 2.19 in the realization used. Deviations are proportional, $\hat{d}_l = \alpha\, d_l$. One coincidence should be flagged: because demand is ϱ times recoverable supply, at $\varrho = 1$ total nominal demand and recoverable supply are the same 89,877 tonnes, and below that figure denotes recoverable supply unless demand is named. Because the allocation shapes the uncertainty set, Section 5.5 tests three alternatives.

Hub capacities. Hub capacity scales the base figure by the body's site count relative to the mean across candidates, clipped to [0.35, 3.0], giving 168,000 tonnes in total against 105,738 tonnes of returns. All other capacities are uniform across sites.

Table 4 lists every parameter of the instance, marking which values are observed and which are constructed.

Table 4. Complete parameter specification of the case instance. Observed values are marked; all others are constructed as described in Section 4.2.

Parameter	Value	Source
Return sources K	86	observed: reporting bodies
Returns r_k	251 to 6,787 t; total 105,738 t	observed: 2024 balance
Candidate hubs I	12	top 12 by observed site count
Recycling facilities J	6	constructed
Redistribution centers M	5	constructed
Disposal sites N	4	constructed
Market regions L	12	constructed
Depots P	1, at (48.9, 11.4)	constructed
Residual fraction s	0.15	EU recovery rates
Deviation ratio α	0.20	assumption; no panel available
Recovery-demand ratio ϱ	swept over 0.2 to 1.5; 1.0 at baseline	experimental variable
Demand d_l	$\varrho(1-s)\Sigma r_k$, split by weights ~ U[0.6, 1.4]	constructed
Total demand at $\varrho = 1$	89,877 t	derived

Table 4. Continued.

Parameter	Value	Source
Hub capacity cc_i	9,917 to 31,500 t/yr; total 168,000	14,000 scaled by site count
Recycling capacity cr_j	30,000 t/yr each	constructed
Redistribution capacity ce_m	22,000 t/yr each	constructed
Disposal capacity cd_n	25,000 t/yr each	constructed
Depot capacity cs_p	126,886 t/yr (1.2 × total returns)	constructed
Hub fixed cost f_i	45,000 €/yr	constructed
Recycling fixed cost g_j	850,000 €/yr	constructed
Redistribution fixed cost h_m	250,000 €/yr	constructed
Depot fixed cost C_p	0	constructed
Unit transport cost	0.12 € per tonne-kilometre	German road freight order of magnitude
Arc costs	great-circle distance × 1.25 × 0.12; ranges 0.00 to 59.62 €/t	derived from coordinates
Shortfall penalty π_l	180 €/t	constructed
Random seed	7, single PCG64 stream	fixed

Four of these values carry the analysis, and none is observed: the residual fraction, the deviation ratio, the shortfall penalty and the redistribution fixed cost. Section 5.5 perturbs each of them across a range and reports the effect, so their influence is measured rather than assumed away. The recovery-demand ratio is not treated as a parameter at all but as the experimental variable, precisely because secondary-market demand for recovered material is not a municipal statistic and we decline to fix it at an invented level.

4.3 | Aggregation for the Two-Stage Experiment

Column-and-constraint generation replicates the recourse block once per generated scenario, so the two-stage experiment of Section 5.4 uses a reduced instance. The 86 sources are grouped into 8 collection regions by a longest-processing-time rule, assigning bodies in descending order of return quantity to the region with the smallest accumulated total, which balances volume across regions; the return quantity of a region is the sum of its members and its arc costs are the member means. The candidate sites are reduced to 6 hubs, 4 recycling facilities, 3 redistribution centers, 3 disposal sites and 8 market regions. Each echelon keeps its total capacity and its fixed cost per unit of installed capacity: both are scaled by the ratio of the original to the reduced site count. The reduced instance therefore represents the same physical system over fewer, larger sites rather than a capacity-starved one. Total demand is rescaled to hold ϱ at its intended value.

5 | Computational Results

5.1 | Computational Setup

All models were built with the Gurobi Python interface and solved with Gurobi 13.0.2 under Python 3.12 on Linux (x86-64). Unless stated otherwise, the relative MIP tolerance was 1e-4, the time limit 3600 seconds, and the thread count the solver default; no other parameter was altered, and no tuning, warm start, or callback was used. Instance data are seeded through a fixed FNV-1a hash rather than the interpreter hash, so a given name and seed reproduce identical data across machines.

Every model reached proven optimality within these settings, except the column-and-constraint generation runs, which terminate on the bound criterion of Section 3.4. Instances contain 1,400 to 1,800 variables and solve in well under a second, so no heuristic was required, and none is reported; we draw no wider conclusion, since establishing when metaheuristics are warranted would need a purpose-built benchmark across instance sizes and structures. Every number below comes from scripts that read the extracted waste balance and write the result tables directly.

Every number reported below is produced by scripts that read the extracted waste balance and write the result tables directly; no figure or table is drawn from any other source.

5.2 | Layered Response at the Reference Ratio

We first fix the recovery-demand ratio at unity, so that secondary-market demand equals recoverable supply, and sweep the budget across its full admissible range under each service policy. *Table 5* reports all eighteen.

Table 5. Four-layer response at recovery-demand ratio 1.0.
all 18 solves are shown: Nine budget levels under each service policy, spanning the full admissible range of Γ in each case. Changes are measured against the Γ = 0 solution of the same policy.

Policy	Γ	Cost Δ (%)	Openings Δ	Delivered (t)	Regional Δ	Changed arcs	Rerouted (t)	Penalty Δ (k€)	Other Δ (k€)
Regional	0	0.00	0	88,000	0	0	0	+0	+0
Regional	0.125	3.45	0	88,000	12	15	2,464	+404	-30
Regional	0.25	6.90	0	88,000	12	16	4,929	+809	-60
Regional	0.375	10.37	0	88,000	12	17	7,345	+1213	-89
Regional	0.5	13.86	0	88,000	12	17	9,204	+1618	-114
Regional	0.625	17.37	0	88,000	12	17	12,152	+2022	-138
Regional	0.75	20.87	0	88,000	12	17	15,249	+2427	-162
Regional	0.875	24.38	0	88,000	12	17	18,646	+2831	-186
Regional	1	27.89	0	88,000	12	17	22,043	+3236	-210
Aggregate	0	0.00	0	88,000	0	0	0	+0	+0
Aggregate	1.5	4.31	0	88,000	3	8	4,982	+542	-74
Aggregate	3	8.23	0	88,000	5	12	10,319	+1019	-126
Aggregate	4.5	11.90	0	88,000	7	12	14,210	+1450	-158
Aggregate	6	15.59	0	88,000	9	14	15,381	+1876	-185
Aggregate	7.5	19.14	0	88,000	11	16	18,065	+2279	-203
Aggregate	9	22.33	0	88,000	12	17	22,043	+2633	-210
Aggregate	10.5	25.31	0	88,000	12	17	22,043	+2956	-210
Aggregate	12	27.89	0	88,000	12	17	22,043	+3236	-210

The four layers behave very differently, and the aggregate cost curve hides that fact.

The strategic layer is invariant. Across all eighteen solves, nine per service policy, the Hamming distance between opening vectors is zero: the same eight consolidation hubs, three recycling facilities, and four redistribution centers are selected at every budget level under both policies. No capital decision responds to the budget.

Aggregate service is invariant, at 88,000 tonnes throughout, but regional service is not. Under the regional policy, all twelve market regions receive a different quantity at positive budget; under the aggregate policy, the number of re-served regions grows from two to five as the budget increases. A fixed quantity of recovered material is being re-allocated toward the regions whose protected demand has risen fastest.

The operational layer separates sharply by position in the chain. *Table 6* gives the arc-family detail at the maximum budget. Flows on the collection, recycling, redistribution, disposal and forward arcs are invariant to the tonne, with an L1 change of exactly zero. The entire operational response is carried by the terminal redistribution-to-market arcs, where the relative L1 change reaches 0.50 under the regional policy and 0.81

under the aggregate policy. In other words, the recovery chain up to the point of delivery is frozen, and only the final allocation moves.

Table 6. Operational response by arc family at the maximum budget.

Arc Family	Reg. Arcs	Reg. L1 (T)	Reg. Rel.	Agg. Arcs	Agg. L1 (T)	Agg. Rel.
K-I collection	0	0.0	0.0000	0	0.0	0.0000
I-J to recycling	0	0.0	0.0000	0	0.0	0.0000
J-M to redistribution	0	0.0	0.0000	0	0.0	0.0000
M-L to market	17	44,085.3	0.5010	17	44,085.3	0.5010
I-N to disposal	0	0.0	0.0000	0	0.0	0.0000
P-K forward	0	0.0	0.0000	0	0.0	0.0000

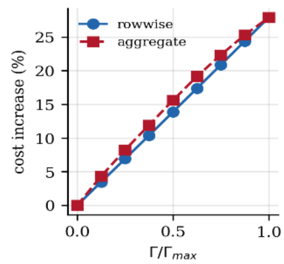

a.

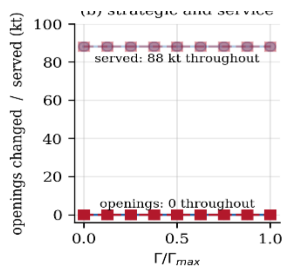

b.

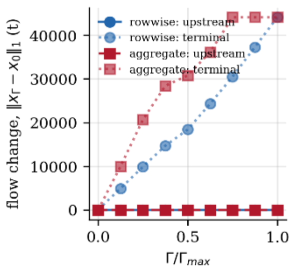

c.

Fig. 2. Layered response at recovery-demand ratio 1.0.Cost rises steadily while facility openings and aggregate delivered volume are constant and upstream flow change is identically zero; the operational response is carried entirely by the terminal arcs; a. Cost, b. Strategic and service, and c. Operational.

Two conclusions follow, and their wording matters. The budget here is strategically inactive and upstream-inactive: it changes no capital decision and no movement in the recovery chain. It is not decision-invariant, since terminal allocation does change, and a study reporting only openings and total delivered volume would have overstated the invariance. Moving from the regional to the aggregate policy alters the shape of the response without activating the strategic layer: at full protection the two coincide exactly, and they differ only at intermediate protection, as Section 3.3.3 shows. Because the two express different guarantees rather than two counterparts of one problem, this is not a repaired budget failing but the ceiling binding under both. The mechanism is the admissibility limitation of [5]: Protected demand cannot be served, because installed capacity binds and the physical return stream lies only 2.1 percent beyond it. Section 5.5 isolates both ceilings and shows that neither responds to the budget.

5.3 | The Upper Service-Saturation Boundary

This section explains the service layer analytically: it accounts for when aggregate delivery stops responding from above, and says nothing about the strategic layer, which Section 5.5 shows is governed differently. The coarse grid brackets a change of regime between $\varrho = 0.9$ and 1.0, and the layers do not switch together. *Table 7* shows delivery and upstream flow responding at all six ratios below the boundary and at none above, while the opening vector changes at five of the six under the static formulation and four of five under the two-stage formulation of Section 5.4. The boundary is sharp for aggregate delivery and is not a boundary for strategic response at all. We treat it as a threshold for the delivery layer alone.

Table 7 reports the sweep over ϱ under the aggregate-service policy, comparing the zero-budget and maximum-budget solutions at each ratio.

Table 7. Layered response across the recovery-demand ratio under the aggregate-service policy, comparing Γ = 0 with Γ = Γmax.

ϱ	PoR (%)	Openings Changed	Delivered Γ = 0 (t)	Delivered Γmax (t)	Upstream L1 (t)	Terminal L1 (t)
0.2	0.46	0	18,572	21,571	8,995	2,998
0.4	2.18	2	35,951	43,141	83,911	14,452
0.6	12.21	1	53,926	64,712	106,287	26,240
0.7	9.55	7	60,000	75,497	298,935	62,404
0.8	6.65	2	71,902	86,282	118,601	41,885
0.9	19.24	7	80,890	88,000	257,350	26,847
1.0	27.89	0	88,000	88,000	0	44,085
1.2	27.02	0	88,000	88,000	0	30,493
1.5	25.69	0	88,000	88,000	0	32,229

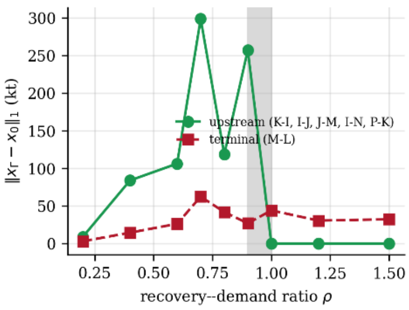

a.

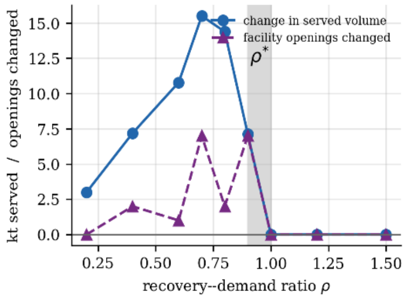

b.

Fig. 3. Layered response across the recovery-demand ratio. Below the upper service-saturation boundary, the service and upstream operational layers respond at all six ratios tested, while the strategic layer responds at five of the six, the exception being ϱ = 0.2. Above the boundary, the service and upstream responses are zero and only terminal allocation moves; a. Operational responseby layer, and b. Strategic and service response.

Let R denote total returns and s the residual fraction. For a given installed design, the maximum deliverable recovered volume is

$$\bar{S}(Y, Z, W) = \min\{(1 - s)\sum_{k\in K} r_k, \sum_{j\in J} cr_j Z_j, \sum_{m\in M} ce_m W_m\}. \quad (21)$$

Total demand is ϱ(1 − s)R by construction. Delivered volume need not equal protected demand everywhere below the ceiling, because the absorption bound (3) leaves room to place material that is cheaper to deliver than to dispose of, as Section 3.2 describes and as the ϱ = 0.2 row of *Table 7* shows. Near the upper boundary, however, once the protected-demand requirement is the active lower bound on delivery, delivered volume increases with protection until the installed service ceiling is reached. The boundary is therefore the ratio at which nominal demand itself reaches that ceiling,

$$P^*(Y, Z, W) = \frac{\bar{S}}{(1 - s)\sum_{k\in K} r_k}. \quad (22)$$

When nominal demand reaches this boundary, additional protected demand cannot increase aggregate service without changing the installed design. The statement is conditional on (Y, Z, W); it concerns the upper boundary only, and it is not a sufficient condition for strategic response.

Two thresholds must be separated, and only one of them is available in advance.

The physical threshold is ϱ*phys = 1. Since the ceiling can never exceed (1 − s)R whatever is installed, no design can raise aggregate delivery once nominal demand reaches recoverable supply. This quantity is genuinely pre-optimization: it needs only the return volumes and the residual fraction.

The installed-design threshold is ϱ*(Y, Z, W), which is conditional on a facility configuration and can be evaluated once that configuration is known. It is not available before optimization in general, because Z and W are decisions; Section 5.5 shows directly that raising the shortfall penalty or halving the redistribution fixed cost changes which capacity is installed and therefore moves the threshold.

We deliberately do not offer a closed-form threshold for strategic response. Whether the budget causes another facility to open depends on discrete economics, the fixed cost of the increment against the shortfall penalty avoided, the transport cost of reaching the additional demand, the size of the capacity increment and the geography of the demand it would serve. Section 5.5 shows each of these changing the outcome at a fixed ϱ. No ratio of volumes can capture that, and presenting one would overstate what the identity supports.

The boundary is confirmed numerically. At baseline, the ceiling is 88,000 tonnes against a recoverable supply of 89,877, so ϱ*(Y, Z, W) = 0.9791. The mechanism is visible in the zero-budget solutions approaching it, where delivery equals nominal demand to the tonne: 80,890 tonnes at ϱ = 0.90 and 87,181 at ϱ = 0.97, each equal to $\varrho(1 - s)R$, then 88,000 from $\varrho = 0.9792$ onward, where nominal demand has met the ceiling. A grid at spacing 0.01 shows aggregate delivery responding at every ratio up to 0.97 and saturated from 0.98 onward, and a six-step bisection brackets the switch in [0.97906, 0.97922], whose midpoint differs from the prediction by 3 x 10^{-5}.

The numerical value belongs to this instance and this configuration. What transfers is the pair of thresholds and their different status: $\varrho^*phys = 1$ as a design-independent bound on aggregate service, and $\varrho^*(Y, Z, W)$ as a conditional statement that becomes checkable once a candidate design is on the table, for instance when assessing whether an existing network will deliver more under demand-side protection. Neither speaks to the strategic layer, as Section 5.5 shows.

The price of robustness is not monotone in ϱ. It rises to 12.21 percent at a ratio of 0.6, falls back to 6.65 percent at 0.8, rises again to 27.89 percent at unity, and then declines slowly to 25.69 percent at 1.5. The non-monotonicity is a further reason to distrust the curve as a summary statistic: its magnitude conflates the cost of adaptation, which occurs below saturation, with the cost of unavoidable shortfall, which dominates above it. The largest values of the curve occur precisely where the model is doing the least.

5.4 | First-Stage Response under Recourse

The results so far were obtained from the static formulation, in which openings and flows are chosen together. If the strategic invariance were an artefact of that conflation, it should disappear once openings are committed before the realization and flows become recourse. We test this directly.

The column-and-constraint generation algorithm converged to a relative upper-lower bound gap below 10^{-6} at every tested budget level, requiring at most two iterations. The largest gap observed was 8 x 10^{-16} across the budget sweep and 4 x 10^{-14} across the ratio sweep, with a single value of minus 2 x 10^{-15} that reflects machine precision rather than a bound crossing. We therefore treat the reported solutions as optimal to numerical tolerance. *Table 8* reports the first-stage response at recovery-demand ratio 1.0. The aggregation used for this experiment is specified in Section 4.3.

Table 8. Two-stage adjustable robust results at recovery-demand ratio 1.0.

Γ	Objective (€)	Cost Δ (%)	Hubs	Recycling	Redist.	Openings Δ	CCG iters
0	12,773,238	0.00	3	2	3	0	1
2	13,724,536	7.45	3	2	3	0	2
4	14,518,883	13.67	3	2	3	0	2
6	15,260,225	19.47	3	2	3	0	2
8	15,822,385	23.87	3	2	3	0	2

The invariance survives. The first-stage design is identical at every budget level: three hubs, two recycling facilities, and three redistribution centers, with a Hamming distance of zero throughout, while the objective rises by 23.87 percent. Committing the openings before the realization and optimizing the flows afterwards does not cause the first stage to hedge in this instance, which is consistent with the fact that at this ratio no configuration can recover more than the return stream yields, so hedging cannot buy additional service. What this establishes is narrower than a general claim: the strategic invariance reported in Section 5.2 is reproduced when the two stages are separated, so it is not attributable to the static timing convention. It remains a finding about this instance at this ratio, on the aggregation of Section 4.3.

The two-stage model does respond at some ratios below saturation, which both complements that finding and bounds it. *Table 9* compares the zero-budget and maximum-budget first-stage designs across ϱ.

Table 9. First-stage response across the recovery-demand ratio, two-stage model.

ϱ	Price of Robustness (%)	Openings at Γ = 0	Openings at Γmax	Hamming
0.2	0.12	(4, 1, 1)	(4, 1, 1)	0
0.4	4.18	(4, 1, 1)	(3, 1, 2)	2
0.6	3.67	(3, 2, 2)	(3, 2, 2)	2
0.8	10.55	(3, 2, 2)	(3, 2, 3)	3
0.9	13.28	(3, 2, 3)	(3, 2, 3)	2
1.0	23.87	(3, 2, 3)	(3, 2, 3)	0
1.2	23.38	(3, 2, 3)	(3, 2, 3)	0

In the two-stage baseline experiment, no first-stage response is observed above service saturation. Below saturation, strategic response remains intermittent: the first stage reconfigures at ratios of 0.4, 0.6, 0.8 and 0.9 but not at 0.2, and the moves differ in kind, opening an additional redistribution center at 0.4 and 0.8 and relocating sites at unchanged counts at 0.6 and 0.9. These results demonstrate association within the tested instance rather than a general relation. Section 5.5 supplies one direct counterexample to necessity, a capacity perturbation at fixed ratio that moves two facilities with aggregate service unchanged. The service-saturation threshold itself governs aggregate-service response, not strategic response, and the two layers decouple in both directions.

5.5 | Isolating the Mechanism

Sections 5.2 and 5.3 established that service freezes, but not why. Several parameters could produce the same appearance, and one observation shows that the physical return stream is not by itself the explanation: The delivered volume settles at 88,000 tonnes while recoverable supply, at a residual fraction of 0.15, is 89,877 tonnes. The freeze occurs 1,877 tonnes short of the physical bound, so something else binds first. We therefore perturb one parameter at a time and record which constraint is tight at the optimum, whether the delivered volume moves, and whether the strategic layer reactivates. *Table 10* summarizes.

Table 10. Mechanism isolation at recovery-demand ratio 1.0.

Perturbation	Supply (t)	Installed Cap. (t)	Delivered Γ = 0 (t)	Delivered Γmax (t)	Delivery Δ (t)	Binding	Openings Δ
Baseline	89,877	88,000	88,000	88,000	0	redistribution capacity	0
s = 0.05	100,451	110,000	88,000	100,451	+12,451	physical supply	4
s = 0.10	95,164	110,000	88,000	95,164	+7,164	physical supply	3
s = 0.15	89,877	88,000	88,000	88,000	0	redistribution capacity	0
s = 0.25	79,304	88,000	79,304	79,304	0	physical supply	0
s = 0.40	63,443	66,000	60,000	60,000	0	recycling capacity	0
Redist. cap. x0.5	89,877	55,000	55,000	55,000	0	redistribution capacity	0
Redist. cap. x0.75	89,877	82,500	82,500	82,500	0	redistribution capacity	0
Redist. cap. x1	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Redist. cap. x1.5	89,877	99,000	89,877	89,877	0	physical supply	0
Redist. cap. x3	89,877	132,000	89,877	89,877	0	physical supply	2
Recyc. cap. x0.5	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Recyc. cap. x0.75	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Recyc. cap. x1	89,877	88,000	88,000	88,000	0	redistribution capacity	0

Table 10. Continued.

Perturbation	Supply (t)	Installed Cap. (t)	Delivered $\Gamma = 0$ (t)	Delivered Γmax (t)	Delivery Δ (t)	Binding	Openings Δ
Recyc. cap. x1.5	89,877	110,000	89,877	89,877	0	physical supply	0
Recyc. cap. x3	89,877	110,000	89,877	89,877	0	physical supply	0
Penalty $\pi = 45$	89,877	44,000	30,000	30,000	0	recycling capacity	0
Penalty $\pi = 90$	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Penalty $\pi = 180$	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Penalty $\pi = 360$	89,877	110,000	89,877	89,877	0	physical supply	0
Penalty $\pi = 900$	89,877	110,000	89,877	89,877	0	physical supply	0
Penalty $\pi = 3000$	89,877	110,000	89,877	89,877	0	physical supply	0
Redist. fixed x0.1	89,877	110,000	89,877	89,877	0	physical supply	0
Redist. fixed x0.5	89,877	110,000	89,877	89,877	0	physical supply	0
Redist. fixed x1	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Redist. fixed x2	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Deviation $\alpha = 0.05$	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Deviation $\alpha = 0.10$	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Deviation $\alpha = 0.20$	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Deviation $\alpha = 0.40$	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Alloc. uniform	89,877	88,000	88,000	88,000	0	redistribution capacity	0
Alloc. skewed 80/20	89,877	110,000	89,877	89,877	0	physical supply	0
Alloc. single dominant	89,877	88,000	88,000	88,000	0	redistribution capacity	0

All 33 perturbations are shown; each varies one parameter with the others held at baseline. Delivery Δ is the change in aggregate delivered volume between $\Gamma = 0$ and Γmax, so a row with a non-zero openings change and a zero delivery Δ is one in which the design moves without delivering more. The tests separate three ceilings, which are nested rather than alternative.

The immediate ceiling at baseline is installed redistribution capacity. The four opened redistribution centers provide exactly 88,000 tonnes of capacity and are fully utilized at every budget level; the fifth candidate remains closed. This ceiling is endogenous: the model declines to open the fifth center because its fixed cost exceeds the value of the 1,877 tonnes of additional service it would permit. Halving the per-site capacity moves the freeze to 55,000 tonnes and halving the fixed cost removes it by making the fifth center worth opening, both of which confirm the attribution.

Behind it lies the physical supply ceiling at $(1 - s)$ times total returns. No perturbation of penalty, fixed cost, capacity, deviation level, or demand allocation ever produces a served volume above 89,877 tonnes. The only parameter that moves it is the residual fraction itself, and when s is lowered, the delivered volume tracks the new bound exactly, reaching 100,451 tonnes at $s = 0.05$ and 95,164 tonnes at $s = 0.10$, in each case equal to $(1 - s)$ times 105,738 to the tonne. This is the ceiling that no design decision can cross.

Between the two lies an economic band. Raising the shortfall penalty to 360 euro per tonne, or halving the redistribution fixed cost, makes the last capacity increment worth installing and lifts service from 88,000 to 89,877 tonnes. Within that band the choice not to serve is economic, and the band is narrow here, 1,877 tonnes or 2.1 percent of recoverable supply.

Two consequences follow, and both qualify claims made earlier in this paper.

First, the attribution in Section 5.2 must be stated as a composite. The budget is strategically inactive at baseline because the binding constraint is installed capacity, and raising protection cannot justify installing more, since the additional volume that could then be served is bounded by physical supply only 2.1 percent away. It is not correct to attribute the freeze to physical supply alone.

Second, and more consequentially, the perturbations separate what governs each layer. A budget can increase aggregate service only when additional protected demand remains physically and operationally servable. Strategic response is governed separately: it may arise through capacity expansion when additional service is economically worthwhile, or through spatial reconfiguration when the geographic pattern of protected demand changes. Consequently, no volume ratio alone determines the strategic layer.

The evidence is thinner than the mechanism. Only three of the thirty-three perturbations change the opening vector: lowering the residual fraction to 0.05 or 0.10, which raises the physical ceiling and moves four and three facilities while delivery rises with it, and tripling redistribution capacity, which moves two facilities while delivery stays at 89,877 tonnes. Only that third case separates the layers, and the two delivery columns of Table 10 make it checkable; we report it as a single demonstration rather than a pattern. Several perturbations change what is installed without changing the opening count: raising the penalty to 360 euros per tonne, halving the redistribution fixed cost, or increasing capacity each lifts installed redistribution throughput from 88,000 to 110,000 or 132,000 tonnes and moves delivery to the physical ceiling.

Two further findings. The deviation level α moves the price of robustness from 6.90 percent at 0.05 to 56.63 percent at 0.40 without changing delivered volume, the binding constraint, or the opened facilities at any level. The curve's height therefore carries no reliable information about whether the strategic or aggregate-service layers respond. Only the terminal operational component moves with it, re-allocation growing from 9,858 to 55,099 tonnes while upstream L1 stays exactly zero. A curve that moves eightfold tracks one layer while three stay fixed, which is what the decomposition exists to expose. Demand allocation behaves differently again: it changes which constraint binds and how much is delivered, a skewed allocation reaching the physical ceiling of 89,877 tonnes rather than 88,000, yet none of the three allocations produces any response to the budget, the opening vector and delivered volume being identical at $\Gamma = 0$ and Γmax. The allocation shapes the plan; the budget does not.

6 | Discussion

6.1 | Implications for Practice

The results answer the question posed in Section 1: Of the four responses a budget can produce, only two moved in this network, and neither was the one that commits capital. We now interpret that for practice and for method.

For a network planner, the finding is a screening rule whose scope depends on purpose. For capital expansion intended to increase aggregate delivery, two conditions must hold together: additional protected demand must remain servable, and the required capacity increment must be worth installing. The saturation boundary of Section 5.3 governs the first and can be checked from volumes; the second turns on the fixed cost of the increment against the penalty it avoids, and cannot. Strategic redesign can also occur without higher aggregate delivery, when protection changes the spatial allocation of demand. Section 5.5 supplies the case: Tripling redistribution capacity moves two facilities while delivery stays at 89,877 tonnes, so a planner reading delivery volumes alone would conclude the design had not moved when it had. Section 3.3.3 supplies a larger one, where matched policies deliver the same total to a different mix of regions.

6.2 | Implications for Research

The response decomposition is inexpensive, and we suggest it become routine. A robust network design study that reports a price-of-robustness curve alone leaves the reader unable to tell which layer moved, and our results show that the layers can move independently: in the same instance, the strategic layer is completely inactive while the terminal operational layer changes by up to 81 percent in relative L1. We also caution against the shortcut we ourselves first took, of characterizing invariance through opened facilities and aggregate delivered volume only. Both were constant here, yet 22,043 tonnes were re-routed. Reporting a flow statistic is necessary; the opening vector alone is not sufficient.

The results also separate two previously independent limitations. The cardinality degeneracy of [3] and the admissibility limitation of [5] are distinct: the aggregate-service policy removes the former while the latter binds under both policies. A modeller who addresses only cardinality sees a changed response profile and may conclude the budget has become meaningful while the strategic layer remains inactive. A second trap sits alongside it: pooling regional constraints into an aggregate one gives the budget something to bind on, and may be the right choice, but it silently replaces a per-region requirement with a system-wide one.

6.3 | Limitations

Seven limitations bound the claims, and two of them concern the formulations themselves rather than the data.

Static formulation as the primary vehicle. The layered results for the full 86-source instance come from the static formulation, in which openings and flows are chosen together. That formulation is not a model of how the decisions are actually taken; we retain it because it is what the applied literature uses, and the two-stage check of Section 5.4 corroborates its central finding on a reduced instance. It does not follow that every static result would survive at full scale under recourse, and the affinely adjustable formulation of [27] and the adaptive network-flow formulations of [28] are not tested here at all. Claims resting on the static model should be read as properties of that model.

Aggregate-service formulation and its fungibility assumption. Results obtained under the aggregate-service policy describe a system-wide requirement and rest on an assumption our data cannot test, namely that recovered material is fungible across buyers. The waste balance covers collection, not downstream marketing, so we adopt fungibility as a modeling assumption for the case instance rather than as an established property of this waste stream. Those results establish nothing about how protection is allocated across regions, and Section 3.3.3 quantifies the difference: at matched aggregate protection, the aggregate-service solution delivers less recovered material to as many as 10 of 12 regions than the regional-service solution does. Where fungibility fails, the aggregate policy is inadmissible and only the regional results apply, so readers should not transfer the aggregate findings to settings with binding regional service obligations.

Mechanism tests. These perturb one parameter at a time from a single baseline, so interactions among them are not mapped; the nested structure is robust across the range tested, but the numerical width of the economic band, 2.1 percent here, is instance-specific.

Two-stage instance size. Column-and-constraint generation replicates the recourse block per scenario, so the two-stage experiment uses an aggregated instance. Total capacity per echelon is preserved, but the spatial detail of the 86 bodies is not, so those results corroborate the static findings rather than standing as a case study in their own right.

Saturation boundary. The value 0.979 belongs to this instance and this installed configuration. The boundary (22) is conditional on (Y, Z, W) and is therefore not in general a pre-optimization quantity; only the physical bound at unity is, and it is the weaker of the two. Both concern the upper boundary only. In this instance, service responds at every feasible ratio below it, but the instance is feasible only above approximately $\varrho = 0.053$, so we have not exercised the low-demand regime in which a standing delivery might already cover protected demand; a network with more disposal capacity could exhibit one. Neither boundary predicts strategic response, and we make no claim that they do.

Unobserved demand side. Secondary-market demand is not measured, and although we sweep it rather than fix it, a case study with observed demand for recovered material would be stronger. Coordinates are likewise not published in the waste balance, so distances are indicative; every layered comparison was made under a fixed distance realization across all budget levels, so the comparisons are internally consistent even though absolute costs are not.

Flow statistics. The decomposition reports flow change through L1 statistics and arc counts, which capture magnitude but not structure; a study concerned with which markets gain and which lose would need a distributional or equity measure in addition.

A body-level multi-year panel would allow the deviation ratio to be estimated from observed year-to-year variation rather than assumed. The per-body files published by the state office cover a single year, and we chose not to construct a panel we could not verify against the primary source.

7 | Conclusion

A budget of uncertainty can raise cost by more than a quarter without changing anything a planner would recognize as a decision. That is what the Bavaria-informed instance shows: under both service policies the facility configuration was invariant, aggregate delivered volume fixed, and flows on every arc family except the terminal one invariant to the tonne. The price-of-robustness curve reports none of this, because it aggregates four responses that move independently.

The contribution is the decomposition that separates them into strategic, service, operational, and penalty responses computed from quantities any solver already returns. Reporting it costs only bookkeeping and lets a reader tell adaptation from accounting.

Three results delimit the finding. The aggregate-service formulation is a different service policy rather than a repaired counterpart, and at matched protection it reallocates delivery across regions rather than protecting each of them. Perturbation tests attribute the inactivity to an installed-capacity ceiling just beneath an immovable physical one, with penalty and fixed-cost structure rather than volumes deciding whether the last increment is built. A two-stage formulation on an aggregated instance reproduces the first-stage invariance, so static timing is not the cause.

For practice, this yields a screening rule. Where protected demand approaches what the return stream can yield, demand-side protection will not change what is built, and collection coverage, recovery yield, and processing capacity are the levers that will.

Three directions follow: adapting the effective-budget construction of [5] to network design, estimating deviation ratios from a multi-year panel of body-level collections, and applying the decomposition retrospectively to published robust recovery-network models.

Acknowledgments

The authors thank the Bavarian State Office for the Environment for publishing the waste balance data in a form that permits independent analysis.

Author Contribution

Conceptualization, M. R. and S. E. N.; methodology, M. R.; software, M. R.; validation, M. E. and S. E. N.; formal analysis, M. R.; investigation, M. E.; resources, M. E.; data curation, M. R.; writing of the original draft, M. E. and M. E.; review and editing, S. E. N.; visualization, M. R.; supervision, S. E. N.; All authors have read and agreed to the published version of the manuscript.

Funding

This research received no external funding.

Data Availability

The primary data are published by the Bavarian State Office for the Environment and are openly available through its waste balance portal. The extracted table covering all 86 reporting bodies, together with the instance generator and the scripts that produce every table and figure in this paper, are available from the corresponding author on reasonable request.

Conflicts of Interest

The authors declare no conflict of interest. Funders played no role in the design of the study, in the collection, analysis, or interpretation of the data, in the writing of the manuscript, or in the decision to publish the results.

References

[1] Bertsimas, D., & Sim, M. (2004). The price of robustness. *Operations Research*, *52*(1), 35–53. https://doi.org/10.1287/opre.1030.0065

[2] Bertsimas, D., & Sim, M. (2003). Robust discrete optimization and network flows. *Mathematical programming*, *98*(1), 49–71. https://doi.org/10.1007/s10107-003-0396-4%0A%0A

[3] Kim, J., Do Chung, B., Kang, Y., & Jeong, B. (2018). Robust optimization model for closed-loop supply chain planning under reverse logistics flow and demand uncertainty. *Journal of Cleaner Production*, *196*, 1314–1328. https://doi.org/10.1016/j.jclepro.2018.06.157

[4] Kim, J., & Lee, S. (2025). Robust optimization model for closed-loop supply chain planning with collected material quality uncertainty. *Operations Research Perspectives*, *15*, 100368. https://doi.org/10.1016/j.orp.2025.100368

[5] Dehghani Filabadi, M., & Mahmoudzadeh, H. (2022). Effective budget of uncertainty for classes of robust optimization. *INFORMS Journal on Optimization*, *4*(3), 249–277. https://doi.org/10.1287/ijoo.2021.0069

[6] Pishvaee, M. S., Rabbani, M., & Torabi, S. A. (2011). A robust optimization approach to closed-loop supply chain network design under uncertainty. *Applied Mathematical Modelling*, *35*(2), 637–649. https://doi.org/10.1016/j.apm.2010.07.013

[7] Biçe, K., & Batun, S. (2021). Closed-loop supply chain network design under demand, return and quality uncertainty. *Computers Industrial Engineering*, *155*, 107081. https://doi.org/10.1016/j.cie.2020.107081

[8] Ge, C., Zhang, L., & Yuan, Z. (2022). Distributionally robust optimization for the closed-loop supply chain design under uncertainty. *AIChE Journal*, *68*(12), e17909. https://doi.org/10.1002/aic.17909

[9] Gao, Y., Lu, S., Cheng, H., & Liu, X. (2024). Data-driven robust optimization of dual-channel closed-loop supply chain network design considering uncertain demand and carbon cap-and-trade policy. *Computers Industrial Engineering*, *187*, 109811. https://doi.org/10.1016/j.cie.2023.109811

[10] Haddadsisakht, A., & Ryan, S. M. (2018). Closed-loop supply chain network design with multiple transportation modes under stochastic demand and uncertain carbon tax. *International Journal of Production Economics*, *195*, 118–131. https://doi.org/10.1016/j.ijpe.2017.09.009

[11] Zhao, Y., Ghiasvand, M. R., & Tosarkani, B. M. (2025). Balanced uncertainty sets for closed-loop supply chain design: A data-driven robust optimization framework with fairness considerations. *Expert Systems with Applications*, *276*, 127170. https://doi.org/10.1016/j.eswa.2025.127170

[12] Foroozanfar, A., Amou Jafari, A., & Mehrabi, M. (2026). Sustainable design of a closed-loop automotive spare parts supply chain using machine learning–based product return rate forecasting. *Supply Chain and Operations Decision Making, 3*(1), 26–45. https://doi.org/10.48313/scodm.v3i1.50

[13] Najafi, M., & Faroughi, H. (2025). Opportunistic maintenance modeling in a closed-loop supply chain: The impact of spare parts material, technician skill, and environmental conditions. *Supply Chain and Operations Decision Making, 2*(4), 243–260. https://doi.org/10.48313/scodm.v2i4.47

[14] Tavakoli, P., Sheikhmohammady, M., & Nikbakhsh, E. (2026). Positioning empty containers among ports considering leasing and purchasing options using robust optimization. *Supply Chain and Operations Decision Making, 3*(1), 10–25. https://doi.org/10.48313/scodm.v3i1.49

[15] Zarei, M., & Varmazyar, M. (2026). Designing a blood supply chain network: A bi-objective robust optimization approach. *Supply Chain and Operations Decision Making, 3*(1), 63–75. https://doi.org/10.48313/scodm.v3i1.52

[16] Farokhiani, M., Tabashir, S., Hassan Pour, H. A., & Ghaffari Touran, H. (2026). Designing a food security supply chain for sustainable oil production using a non-dominated sorting genetic algorithm (NSGA-II) approach. *Supply Chain and Operations Decision Making, 3*(1), 46–62. https://doi.org/10.48313/scodm.v3i1.51

[17] Karimi, B., & Omran, M. M. (2026). A blockchain-based optimization model for smart 4PL network design in off-site construction. *Supply Chain and Operations Decision Making*, *3*(1), 76–86. https://doi.org/10.48313/scodm.v3i1.54

[18] Karimi, B., & Mohammadpour Omran, M. (2025). Fourth-party logistics in off-site construction: A synergy-based service strategy. *Supply Chain and Operations Decision Making, 2*(4), 229–242. https://doi.org/10.48313/scodm.v2i4.43

[19] Ghashghaei, M., & Golpîra, H. (2025). Advanced optimization of energy production distribution system integrated with sustainable supply chain networks. *Supply Chain and Operations Decision Making, 2*(2), 70–79. https://doi.org/10.48313/scodm.v2i2.31

[20] Fasihia, M., & Ebrahimi, S. A. (2025). Designing antifragile and agile food supply chains. *Supply Chain and Operations Decision Making, 2*(1), 39–49. https://doi.org/10.48313/scodm.v2i1.28

[21] Tanhaie, F., & Izanloo, A. (2025). A fuzzy multi-criteria decision-making approach for green supply chain management: Supplier prioritization in the automotive industry. *Supply Chain and Operations Decision Making, 2*(3), 105–114. https://doi.org/10.48313/scodm.v2i3.36

[22] Azarmina, S. S., Fallahnejad, M. S., & Khademi Zare, H. (2025). A combined knowledge-driven and data-driven approach to selecting resilient suppliers based on heterogeneous information (Case study: Steel industry). *Supply Chain and Operations Decision Making, 2*(4), 201–219. https://doi.org/10.48313/scodm.v2i4.44

[23] Dolatyari, M., Gholami Koldehi, A., & Zegordi, S. H. (2026). Calculating the access probability of network arcs to determine the maximum flow in paths. *Supply Chain and Operations Decision Making, 3*(1), 1–9. https://doi.org/10.48313/scodm.v2i4.46

[24] Minoux, M. (2012). Two-stage robust LP with ellipsoidal right-hand side uncertainty is NP-hard. *Optimization Letters*, *6*(7), 1463–1475. https://doi.org/10.1007/s11590-011-0341-z%0A%0A

[25] Ouorou, A. (2016). Robust models for linear programming with uncertain right hand side. *Networks*, *68*(3), 200–211. https://doi.org/10.1002/net.21693

[26] Atamtürk, A., & Zhang, M. (2007). Two-stage robust network flow and design under demand uncertainty. *Operations Research*, *55*(4), 662–673. https://doi.org/10.1287/opre.1070.0428

[27] Ben-Tal, A., Goryashko, A., Guslitzer, E., & Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, *99*(2), 351–376. https://doi.org/10.1007/s10107-003-0454-y%0A%0A

[28] Bertsimas, D., Nasrabadi, E., & Stiller, S. (2013). Robust and adaptive network flows. *Operations research*, *61*(5), 1218–1242. https://doi.org/10.1287/opre.2013.1200

[29] Zeng, B., & Zhao, L. (2013). Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, *41*(5), 457–461. https://doi.org/10.1016/j.orl.2013.05.003

[30] Hu, Y. (2023). *Perspectives in closed-loop supply chains network design considering risk and uncertainty factors*. https://doi.org/10.48550/arXiv.2306.04819

[31] Jahani, H., Abbasi, B., Sheu, J. B., & Klibi, W. (2024). Supply chain network design with financial considerations: A comprehensive review. *European Journal of Operational Research*, *312*(3), 799–839. https://doi.org/10.1016/j.ejor.2023.02.033